A Superconducting Software-Defined Radio Frontend with Application to the Square-Kilometre Array

Mark Hans Volkmann

Dissertation presented for the degree Doctor of Philosophy in the Faculty of Engineering at Stellenbosch University

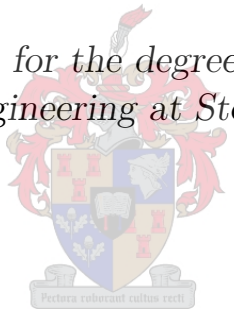

Promoters:

Prof. C. J. Fourie Prof. D. B. Davidson Prof. W. J. Perold

December 2013

Declaration

By submitting this dissertation electronically, I declare that the entirety of the work contained therein is my own, original work, that I am the sole author thereof (save to the extent explicitly otherwise stated), that reproduction and publication thereof by Stellenbosch University will not infringe any third party rights and that I have not previously in its entirety or in part submitted it for obtaining any qualification.

Signature: M. H. Volkmann

Date: September 2013

Acknowledgements

Writing this section comprehensively and accurately was a challenging task. I have received more help, guidance and support than I can mention in this restricted format. Nonetheless, I would like to express my sincere gratitude to the people mentioned in this section. I apologise to those that I forget to mention here.

Most grateful I am to my supervisor, Professor Coenrad J. Fourie, who has guided me expertly through the last three years. I could not have wished for a better supervisor, teacher and mentor. Specifically, I thank him for introducing me to the field of applied superconductivity, for teaching me about the fundamental concepts, for supervising and guiding my work, for doing and supervising important work parallel and essential to mine, and for accompanying me to and supporting me at various international conferences.

Next I thank Dr. Oleg Mukhanov, who has supported and mentored me through both my internships at Hypres and taught me far more than I could have hoped to learn in a comparatively short time. I also thank him for financing my second trip to Hypres and his flexibility in all regards. Finally I thank him for his unending support in demanding administrative matters.

Special thanks go to Dr. Thomas Ortlepp, whose frankness and openness have eased many of my endeavours. I thank him specifically for teaching me fundamentals in superconducting circuit design and experimental procedure, for hosting me in Erfurt, for putting me up at his home for a substantial period, for accompanying me to a conference and for reintroducing me to German culture.

I am grateful to Professor D. B. Davidson for co-supervising my thesis, much general and specific advice and his support for my SKA bursary application, as well as other applications.

I thank my co-supervisor Professor W. J. Perold for much advice and the arrangement of my first internship at Hypres.

To Professor H. C. Reader I am grateful for advice and instruction regarding topics ancillary to my research. I also thank him for imparting on me the importance of electromagnetic compatibility and for taking the black magic out of this topic. Finally I thank him for supporting various applications.

I am also grateful to Dr. Dirk de Villiers and Professor M. Botha for advice on various topics. Dr. De Villiers deserves special thanks for accompanying

me to a conference.

I am grateful to Anubhav Sahu and Saad Sarwana for their training and patient endurance of my missteps. Special thanks go to Anubhav Sahu and Andrey Dotsenko for conducting crucial experiments in my absence. Thanks also go to Dr. Igor Vernik, Andrei Talalaevski, and Vladimir Dotsenko for much advice in experimental procedure.

To Dr. Steve Kaplan I am grateful for many fruitful discussions about my research at Hypres and many fascinating stories, as well as an entertaining evening in Cambridge.

I am very grateful to Dr. Deep Gupta for many discussions on all topics imaginable. I am particularly grateful for his inspiring idealism.

I am grateful to Retief Gerber and NioCAD Pty (Ltd), for providing me with a temporary license for NioPulse and introducing me to many members of the superconducting electronics community.

To Dr. Elie Track I am grateful for much encouragement, for support of several applications as well as generally easing my attempts at networking in the community.

I am grateful to Professor Alan Kadin for expending much time in helping me understand the most mystifying physical (and political) phenomena.

To Amol Inamdar, Alex Kirichenko, Dr. Dimitri Kirichenko and Dr. Timur Filippov I am grateful for enduring my constant intrusions and their patient willingness to satisfy my curiosity.

Gail Reuter and Karen Vaillancourt deserve special thanks for making me feel at home in a foreign place.

I thank Daniel Yohannes for his advice and his substantial efforts at fabricating my chips.

Thanks go to the fabrication teams at Hypres and the Institute for Photonic Technology for fabricating my experimental chips.

To Dr. Olaf Wetzstein I am grateful for hosting me in Jena and all his advice on experimental procedure. I also thank him for conducting experiments for me in my absence.

To Dr. Ronny Stolz, Dr. Jürgen Kunert and Professor H.-G. Meyer I am grateful for hosting me in Jena repeatedly and for their role in fabricating my chips.

Special thanks go to Dr. Scott Holmes for much advice and extraordinarily clear explanations of some superconducting phenomena, as well as an entertaining and informative day in the Oregon countryside.

I am grateful to SKA South Africa for financing my post-graduate studies.

Then I thank my family, specifically my father Hermann Volkmann, for advice and support.

Last, but by no means least, I thank Michelle, without whose boundless love, support and understanding I could not have achieved what I have.

Abstract

A Superconducting Software-Defined Radio Frontend with Application to the Square-Kilometre Array

M. H. Volkmann

Dissertation: PhD

2013

Superconducting electronics can make the Square Kilometre Array (SKA) a better instrument. The largest radio telescope in the world will consist of several arrays, the largest of which, consisting of more than 3000 dishes, will be situated primarily in South Africa. The ambitions of the SKA are grand and their realisation requires technology that does not exist today.

Current plans see signals in the band of interest amplified, channelised, mixed down and then digitised. An all-digital frontend could simplify receiver structure and improve its performance. Semiconductor (analog-to-digital converters) ADCs continue to make great progress and will likely find applications in the SKA, but superconductor ADCs benefit from higher clock speeds and quantum accurate quantisation. We propose a superconducting software-defined radio frontend.

The key component of such a frontend is a superconducting flash ADC. We show that employing such an ADC, even a small- to moderately-sized one, will significantly improve the instantaneous bandwidth observable by the SKA, yet retain adequate signal-to-noise ratio so as to achieve a net improvement in sensitivity. This improvement could approach factor 2 when compared to conventional technologies (at least for continuum observations). We analyse key components of such an ADC analytically, numerically and experimentally and conclude that fabrication of such an ADC for SKA purposes is certainly possible and useful.

Simultaneously, we address the power requirements of high-performance computing (HPC). HPC on a hitherto unprecedented scale is a necessity for processing the vast raw data output of the SKA. Utilising the ultra-low-energy switching events of superconducting switches (certain Josephson junctions), we develop first demonstrators of the promising eSFQ logic family, achieving experimentally verified shift-registers and deserialisers with sub-aJ/bit energy

requirements. We also propose and show by simulation how to expand the applicability of the eSFQ design concept to arbitrary (unclocked) gates.

Uittreksel

'n Supergeleidende Sagteware-Gedefineerde Radio Kopstuk met Toepassing tot die "Square-Kilometre Array"

M. H. Volkmann

Proefskrif: PhD

2013

Supergeleier-elektronika kan 'n beter instrument maak van die "Square Kilometre Array" (SKA). Die wêreld se grootse radioteleskoop sal bestaan uit etlike skikkings, waarvan die grootste — met meer as 3 000 skottels — hoofsaaklik in Suid-Afrika gesetel sal wees. Die SKA is ambisieus en vereis tegnologie wat nog nie vandag bestaan nie.

Volgens huidige planne sal seine in die band van belang versterk, gekanalisieer, afgemeng en dan versyfer word. 'n Heel-digitale kopstuk sal die ontvangerstruktuur kan vereenvoudig en sy prestasie kan verbeter. Halfgeleier analoog-na-digital omsetters (ADOs) verbeter voortdurend en sal waarskynlik toepassings in die SKA vind, maar supergeleier ADOs trek voordeel uit hoër klok spoed en kwantumakkurate kwantisering. Ons stel 'n supergeleier sagteware-gedefinieerde radio kopstuk voor.

Die sleutelkomponent van so 'n kopstuk is 'n supergeleier "flash" ADO. Ons toon hoe die gebruik van so 'n ADO, selfs een van klein tot matige bisgrootte, die oombliklike bandwydte waarneembaar deur die SKA aansienlik sal verbeter en 'n voldoende sein-tot-ruis verhouding sal behou, en gevolglik 'n netto verbetering in sensitiwiteit sal bereik. Hierdie verbetering kan, vergeleke met konvensionele tegnologie, 'n faktor van 2 nader (ten minste vir kontinuum waarnemings). Ons analiseer belangrike komponente van so 'n ADO analities, numeries and eksperimenteel en lei af dat die vervaardiging van so 'n ADO vir SKA doeleindes beide moontlik en nuttig is.

Terselfdertyd spreek ons die drywingsverkwisting van Hoë-verrigting rekenaars aan. Sulke rekenaars van 'n tot dusver ongekende skaal is 'n noodsaaklikheid vir die verwerking van die enorme rou data uitset van die SKA. Deur die gebruik van die ultra-lae-energie skakels van supergeleier skakelaars (sekere

Josephson-vlakke), ontwikkel ons die eerste demonstratiewe hekke van die veelbelowende eSFQ logiese familie, en toon eksperimenteel bevestigte skuifregisters en deserieëliseerders met sub-aJ/bis energievereistes. Ons stel verder voor en wys met simulasies hoe om die toepaslikheid van die eSFQ ontwerpkonsep na arbitrêre (ongeklokte) hekke uit te brei.

Contents

Declaration i

Acknowledgements ii

Abstract iv

Uittreksel vi

Contents viii

List of Figures x

List of Tables xx

1 Introduction 1
- 1.1 Motivation . . . 1
- 1.2 Superconductivity at a glance . . . 3
- 1.3 Objectives of dissertation . . . 8

2 A software-defined radio for the SKA 11
- 2.1 Introduction . . . 11
- 2.2 The Square-Kilometre Array (SKA) . . . 12
- 2.3 Introduction to SFQ electronics . . . 16
- 2.4 Digitisation of the astronomical signal . . . 28
- 2.5 Chapter summary and conclusion . . . 42
- 2.6 Chapter research output . . . 43

3 Addressing the Software Situation 46
- 3.1 Introduction . . . 46
- 3.2 Design process . . . 47
- 3.3 Capabilities and limitations of available software . . . 58
- 3.4 Addressing the limitations . . . 65
- 3.5 Chapter summary and conclusion . . . 81
- 3.6 Chapter research output . . . 82

4 Superconducting Flash ADC Architecture **83**
4.1 Introduction . . . 83
4.2 A multi-threshold comparator . . . 83
4.3 The superconducting flash ADC . . . 92
4.4 Chapter summary and conclusion . . . 124
4.5 Chapter research output . . . 125

5 ADC Design **126**
5.1 Introduction . . . 126
5.2 Designing a QOS comparator . . . 127
5.3 Benchmarking comparators . . . 147
5.4 Experimental verification of comparator performance . . . 153
5.5 Redesigned Comparators . . . 173
5.6 Applying the signals . . . 180
5.7 Chapter summary and conclusion . . . 182
5.8 Chapter research output . . . 184

6 Ultra-Low-Power Data Transport **187**
6.1 Introduction . . . 187
6.2 Motivation . . . 188
6.3 RSFQ power requirements . . . 189
6.4 Lowering RSFQ power consumption . . . 194
6.5 Designing eSFQ data transport cells . . . 197
6.6 Experimental verification . . . 206
6.7 Towards broader application of eSFQ concepts . . . 236
6.8 Chapter summary and conclusion . . . 244
6.9 Chapter research output . . . 245

7 Conclusions **249**
7.1 Summary of results . . . 249
7.2 Conclusions drawn . . . 252
7.3 How to proceed . . . 254
7.4 Last word . . . 256

Bibliography **257**

A Circuits and Layouts **271**
A.1 Comparators . . . 271
A.2 Test-beds . . . 291
A.3 Comparator chip 1 . . . 292
A.4 Comparator chip 2 . . . 293
A.5 eSFQ cells . . . 294
A.6 eSFQ chips — functional tests . . . 305
A.7 eSFQ chips — high-speed tests . . . 312

B Selected Code **325**
B.1 Selected gEDA modifications (Scheme) 325
B.2 Selected signal-processing code (Python) 326

List of Figures

1.1 Heike Kamerlingh-Onnes, Dutch physicist and Nobel laureate credited with the discovery of superconductivity (a), and his historic measurement of the resistance of a mercury wire (b) [1]. 4
1.2 History of discovered superconductors. Blue entries are high-temperature superconductors (HTS), red entries low-temperature superconductors (LTS). Note the position of niobium (Nb) above the liquid helium line [2]. 4
1.3 Levitation of a magnet over a superconductor [3]. 5
1.4 Comparison between voltage-state logic encoding, common in semiconductor logic families (a) and pulse-based logic encoding, common in superconducting logic families (b). In voltage-state logic, the logical state is measured on the rising edge of the clock. In pulse-based logic, the logical state is encoded in the presence or absence of a pulse in the clock period. 7

2.1 Atmospheric electromagnetic opacity across the spectrum [4]. . . 13
2.2 Projected dish locations for the SKA. A quasi-self-similar pattern with random perturbations [5] (a) ensures wide spatial-frequency coverage. The dishes will be placed across much of Africa (b). . . 15
2.3 Artists impression of the eventual SKA full dish array (a). Possible projected offset Gregorian antenna design for MeerKAT (b). Cryocooled feed of KAT-7 antenna (c). The helium compressor referred to in (b) is a part of the cryocooler [6, 7]. 17
2.4 A Josephson junction and its circuit symbol. 20
2.5 Simplified IV-curve of the Josephson junction. R refers to the resistance of the (parallel) resistive channel, and may be dependent on I. 21
2.6 RCSJ model of the Josephson junction. 22
2.7 Quantities of a damped pendulum. 23
2.8 Switching behaviour of overdamped (a), critically damped (b) and underdamped (c) junction in the presence of a resistive channel. . 24
2.9 Hysteretic IV-curves of the Josephson junction [8]. 25

2.10 Schematic of Josephson transmission line (JTL). 26
2.11 Generating multiple SFQ pulses using the $\sqrt{2}$-rule, as suggested by Likharev et al [9]. 27
2.12 Schematic of D-cell [9]. 28
2.13 Converting $x(t)$ to $x[k]$, *sampling* and *quantisation*. 30
2.14 Advantages of superconducting ADCs. The low-loss characteristics of superconducting circuits are set off against the finite size of the magnetic flux quantum, which makes achievable ADC sensitivity an open question. 32
2.15 Sigma-delta architecture: a popular oversampling ADC. 34
2.16 Superconducting sigma-delta modulator employing flux quantisation. 34
2.17 Possible conventional receiver (a) and proposed superconducting receiver (b) of an SKA antenna. 38
2.18 Angular resolution may obviate need for high dynamic range. . . 41

3.1 Typical design process for superconducting circuits 48
3.2 Ideal high-level design deliverables. 49
3.3 Circuit verification as part of circuit design. 51
3.4 DFF circuit topology and parameters. 52
3.5 Results of a typical margin analysis. The critical margin is the lowest margin of the investigated circuit, and is often a good indication of where work on the circuit could be most rewarding. 53
3.6 Parametric modelling support in gEDA: when values are specified, they are fixed in the netlist (a), whereas their omission sees them replaced by variables (b). 67
3.7 The conventional FastHenry process (without preconditioning). . 69
3.8 Our parallel FastHenry process (without preconditioning). 70
3.9 FastHenry execution time measurement for the conventional (a) and parallel (b) output. M_{prep} and τ_{prep} are approximately 500 MiB and 80 s respectively. M_{sol} is approximately 750 MiB. Of importance to the designer is the `real` entry, which best describes elapsed wall-clock time. 71
3.10 Our tool `verify` operating on a circuit (a), for the successful (b) and an unsuccessful (c) case. The specification is from Listing 3.3. 76
3.11 DFF in test environment to show usage of `analyse`. 77

4.1 Parallel combination of Josephson junction and inductor. 84
4.2 Qualitative response of parallel inductor (a) and Josephson junction (b). 85
4.3 A quantitative $I\phi$-curve of the circuit from Figure 4.1. The dashed line is $I_L(\phi)$, included for clarity. 86
4.4 State response of the parallel inductor-junction combination. . . . 88
4.5 Static response of I_J as a function of the independent variable I in the injective case. 90

4.6 Abstract MTC, a multi-threshold comparator with binary output (a). A first-order input signal $x(t)$ (b) is quantised by the QOS circuit, which produces $\tilde{x}(t)$ (c). The quantised signal is fed into a comparator with threshold at zero, yielding coded output signal $\hat{y}(t)$ (d). 91
4.7 A *decision-making pair* made from two Josephson junctions. . . . 92
4.8 Schematic of a QOS multi-threshold comparator. 93
4.9 Bins of interrogation potential into which I_{sig} is divided by a single, ideal MTC. 94
4.10 Schematic of ADC4, a 4-bit superconducting flash ADC (a), and the interrogation potential of the four comparators as I_{sig} varies (b). 95
4.11 Quantisation by truncation, a natural scheme for unipolar inputs and the flash ADC architecture. Here shown for $n = 3$. 98
4.12 Uniform distribution of the continuous random variable e_q (a) and the discrete random variable x_q (b) for quantisation by truncation. 99
4.13 Quantisation by rounding, with modification for bipolar input signals. Here shown for $n = 3$. 101
4.14 Uniform characteristics of the continuous random variable e_q (a) and the discrete random variable x_q (b) for quantisation by rounding. 102
4.15 An ideal DMP with threshold I_{th} (a) and the associated switching probability as a function of I (b). 104
4.16 A noisy DMP with threshold I_{th} and random noise source I_N. . . 104
4.17 Switching probability of a noisy DMP in response to a clock event. The grey patch indicates the grey-zone. 106
4.18 A grey DMP imposes grey-zones on every threshold of the QOS-based multi-threshold comparator. 107
4.19 Single sweep of $N = 100\,000$ samples across the full range of a 4-bit flash ADC based on grey QOS comparators. 109
4.20 Distribution of total error e_t for the results from Figure 4.19. . . . 110
4.21 Distributions of the components of the error signal confirm that e_q is still uniformly distributed (a) and e_d is discrete (b). This is for $N = 100\,000$ samples and $n = 4$ bits. 111
4.22 Components of the error signal for $N = 100\,000$ samples and $n = 7$ bits. 112
4.23 Program `adcn_sample`, an ADC sample generator for statistical evaluation of ADC performance. The input paramters (n, N, I_c, I_g, Δ) are the number of comparators, number of samples, critical current of the quantisation loop, absolute grey-zone and LSB respectively. Output signals $(y[k], e_t[k])$ are discrete signals of length N representing sets of random samples of the ADC output and total digitisation error respectively. URNG and NRNG refer to uniform and normal random number generators respectively. Note that every comparator $\hat{q}'(\cdot)$ has an independent NRNG. 113

4.24 Distribution of $\widehat{\texttt{SNR}}$ before (a) and after its transformation to the dB scale (b). 115
4.25 Output of `snr_estimator`, compared against (4.11), for $I_g = 0$, $\alpha = 0.001$, $\beta = 0.02$ (and other arbitrary practical parameters). The error bars indicate the interval (maximum 2% away from the plotted point) within which the true achievable SNR lies with 99.9% confidence. 117
4.26 Simulation of relationship between I_g and conservatively predicted SNR for several values of n. 118
4.27 The DMP grey-zone which depresses by 3 dB the achievable SNR of a QOS-based n bit superconducting flash ADC. 119
4.28 Conventional (a) and improved (b) methods of biasing the comparators in a superconducting flash ADC (demonstrated here for $n = 3$). The improved method exhibits a lower grey-zone across the MSB transition as the conventional method. 120
4.29 Results of the repeated experiment for the improved architecture, compared to the conventional method. 121
4.30 Effective grey-zones of comparator transitions varying with β_L. . . 123
4.31 The achievable SNR as ν varies. 124

5.1 The single-threshold comparator. 128
5.2 An example single-threshold comparator layout. 129
5.3 Two options for decision-making pair layout: minimisation of Loop 1 or minimisation of Loop 2. Illustration is neither to scale, nor is the relief of the layers indicated correctly. 130
5.4 A QOS-quantiser with parasitic load L_p. 131
5.5 Response of QOS-quantiser with parasitic output 133
5.6 The response of an inductively-loaded QOS-quantiser with low L_q. 137
5.7 Layout of a QOS-based MTC with minimal L_p. The input current enters from the left. 139
5.8 Schematic of a QOS-based MTC with minimal L_p. 140
5.9 Response of the quantiser in the unclocked, preliminary QOS-based MTC. 141
5.10 Response of the quantiser in the unclocked, preliminary QOS-based MTC after reducing I_{cq}. 142
5.11 Simulated response of clocked QOS-based MTC. The solid rectangular features represent SFQ pulses in quick succession. 143
5.12 The Complementary QOS multi-threshold comparator [10, 11]. . . 145
5.13 Response of equivalently parameterised QOS and CQOS quantisers. Schematics were depicted in Figures 4.8 and 5.12. The traces are simulations for $I_c = 125\,\mu\text{A}$ and $\beta_L = 0.1$ (a), $\beta_L = 0.95$ (b), and $\beta_L = 2.0$ (c). 146
5.14 Excerpts of circuit schematic (a) and layout (b) for MC2 (for full details, see Section A.1.2). 148

5.15 Demonstration of how some `Clock` events produce an `Out` event and others do not for `MC1` at I_{in} near threshold. 151
5.16 An illustration of the fitting procedure employed to determine threshold and grey-zone of an MTC transition. The example investigates a rising (a) and a falling (b) threshold of `MC1`. 151
5.17 Co-damping to maximise speed/grey-zone figure of merit. 155
5.18 Test structure to evaluate comparator `MC3`. This particular structure is located in the South-West of `mvv2-gz-inv`. 156
5.19 Microphotograph of test structure to evaluate comparator `MC1` (also see Figure 5.18). 156
5.20 Experimental test setup to measure the grey-zone of various multi-threshold comparators. 158
5.21 Correct operation of a fabricated comparator (here `MC4`), with a comparatively low (left) and high (right) frequency clock (both in the $\sim 10\,\text{kHz}$ range). Every rising edge of the `Clock` signal signifies one clock event and every rising and falling edge of the `Out` signal signifies one output event. (Note that the experiments were not recorded at the same time or necessarily the same operating point, which explains the differences in signal level.) 161
5.22 Measured switching probabilities and fit curves for the first (a) and fourth (b) measured transition of `MC3`. Both the threshold and the grey-zone can be determined in this way. Visually the fits look good. 162
5.23 Additional measured switching probabilities and fit curves, here for comparator `MC5`, for the third (a) and fourth (b) transitions. Visually the fits do not look as good as in the case of `MC3` (Figure 5.22). 162
5.24 Measured switching probabilities and improved fit curves, here for comparator `MC5`, for the third (a) and fourth (b) transitions. . . . 164
5.25 Estimated achievable SNR for the comparators in our test set. . . 167
5.26 Comparator response of `MC4` over multiple I_m. While $[-2I_m, 2I_m]$ seems correctly represented, outside of approximately $[-3I_m, 3I_m]$ the response is erroneous. 169
5.27 Input current flow for large I_{in} through `MC5`. The current flows underneath the quantising junctions, which probably accounts for the observation depicted in Figure 5.26. 170
5.28 Simulated switching probabilities and correlation, here for comparator `MC3`, at 10 GHz (a) and 22 GHz (b). 171
5.29 A simulation of `MC3` adjusted for $J_c = 4.5\,\text{kA/cm}^2$ at $f_c = 21\,\text{GHz}$. The observed $f_{\max} = 20.2\,\text{GHz}$. 173
5.30 `MX` cell layout, new experimental QOS comparator, currently in fabrication. Details, such as extracted parameters, can be found in Section A.1.6. 176

5.31 `MCX` cell layout, new experimental CQOS comparator, currently in fabrication. Details, such as extracted parameters, can be found in Section A.1.7. 177
5.32 Testbed for `MC` and `MCX`, currently in fabrication. Both comparators share a high-speed testbed, due to limited available chip pads. . . 178
5.33 Traces of possible experimental signals for the testbed from Figure 5.32. Blue signals are input signals, green signals are (unobservable) internal signals and red signals are observable output signals. . . . 178
5.34 Applied signal for parallel comparators. 183

6.1 Inability of conventional technology to acquire ADC output. . . . 188
6.2 Deposited flux quanta reroute bias currents, leading to different biased junctions in a non-storing (a) and storing (b) DFF. 190
6.3 Conventional, resistive RSFQ biasing. 191
6.4 Schematic of SFQ-DC converter proposed by Likharev and Semenov [9]. 193
6.5 The current-limiting properties of the Josephson junction (a) can be exploited to achieve the desired bias current distribution (b). . 195
6.6 eSFQ biasing principle: Exploiting the current-limiting properties of the Josephson junction without the need for a large L_b is possible when the phase gradients between different bias terminals remain constant ($\dot{\phi}_i(t) \approx \dot{\phi}_j(t)\ \forall i, j$). This can be achieved by ensuring that all bias terminals are connected to the clock net, as each point on the clock net experiences the same phase shift ($+2\pi$) during each clock period. Most RSFQ gates contain a decision-making pair (DMP) which is interrogated by the clock signal, making for a suitable bias terminal. 196
6.7 A 1-to-8 deserialiser reduces the data rate by factor 8. 198
6.8 Function of an RSFQ deserialiser. The serial input is distributed across several outputs. 198
6.9 Concept of making ADC output available to conventional technologies. 199
6.10 Conversion of D-cell from RSFQ to eSFQ. 200
6.11 `eSR` - eSFQ shift register cell: schematic (parasitics omitted) (a), simulation (b) and in a typical configuration (c). 202
6.12 Comparison of DMPs in RSFQ (ERSFQ) and eSFQ. Injecting the bias current in the conventional DFF leaves J_1 unbiased (a). Moving the injection point to the DMP forces the phase in both DMP junctions in the same direction during ramp-up (b). 203
6.13 `MeSR` - eSFQ shift register cell with corrective flux bias: schematic (parasitics omitted) (a), simulation (b) and in a typical configuration (c). 204

6.14 Results of simulation with a 20 GHz clock for 1.0 aJ/bit `eSR`- (a) and 0.8 aJ/bit `MeSR`-based (b) 16-bit eSFQ shift registers at bias current corresponding to the centre of the operational region. Bias currents distribute correctly, with accetable distortion through switching events. For the eSR-based shift register, 16 output bits are immediately observed after starting the clock, before any input is applied. In both cases, the input pattern is reproduced at the output. Lower bias current requirements are evident for the magnetic flux-biased shift register (at higher margins). 205
6.15 `eDES` - eSFQ deserialiser cell: schematic (parasitics omitted) (a), simulation (b), in a typical configuration (c). The flux bias is ramped up from 0.1 ns to 0.2 ns in the simulation from (b), as evident from the bottom trace. Initially, eDES is non-storing, but when a pulse arrives at `DIn`, it is stored, readable by both a `Shift` or a `Read` pulse. 207
6.16 Simulation of a 4-bit eDES-based deserialiser at 20 GHz. Deserialisation is clearly evident, the serial input signal (a 4-bit count from 0 to 15) is parallelised, resulting in 4 output streams. 208
6.17 Microphotographs of two of the five $5 \times 5\,\mathrm{mm}^2$ ICs fabricated with Hypres's $4.5\,\mathrm{kA/cm}^2$ process, one with three shift registers (Section A.6.2) (a), and one with a 16-bit deserialiser (Section A.6.5) (b). Each bias line is fed from multiple contact pads, which enables experimental investigation of different bias fan-in configurations. . 209
6.18 Photos of eSFQ shift register cells: eSR with overdamped J_1 (left) and MeSR with critically damped J_1 and a flux bias line inductively coupled to cell storage inductor (right). Cell sizes (indicated by red boundary): eSR: $80 \times 110\ \mu\mathrm{m}^2$, MeSR: $80 \times 105\ \mu\mathrm{m}^2$. These dimensions do not include the bias line. 210
6.19 Photographs of deserialiser test structures: the deserialiser base cell `eDES` (a), and a corresponding 4-bit eSFQ deserialiser (b). Bias lines and some key devices are indicated. A set of DC-SFQ and SFQ-DC converters enable the interface with room temperature electronics. 211
6.20 Photographs of $5 \times 5\,\mathrm{mm}^2$ shift register (a) and $10 \times 10\,\mathrm{mm}^2$ deserialiser (b) chips suitable for high-speed testing. 212
6.21 Microphotograph of 184-bit `eSR`-based shift register test chip. . . 213
6.22 An average voltage pad is connected to SFQ lines, providing the means to measure the average (dc) voltage on the line, where $V_{\mathrm{avg}} = \Phi_0 f$, with f the frequency of pulse incidence on the line. 214
6.23 A crude mixer that enables synthesis of a wideband data signal (of experimental value) from two narrowband signals (easily transmitted) by means of an on-chip voltage-controlled switch. 215

6.24 Schematic of a voltage-controlled switch as employed in the mixing concept introduced in Figure 6.23. $V_{sw} = 0$ leaves J_2 unbiased, so any SFQ pulse arriving at `In` promptly escapes through J_1. When a sufficient voltage is applied at V_{sw}, J_2 is biased sufficiently and pulses arriving at `In` are propagated to `Out`. 215
6.25 Octopux test setup employed for functional tests. 216
6.26 High-speed test setup for shift registers in Suite II. 217
6.27 High-speed test setup for shift registers with average voltage pin and on-chip modulator. 219
6.28 High-speed test setup for deserialisers. 221
6.29 Measured correct functionality of 16-bit eSFQ shift registers: simple pattern (a), randomly generated data and clock pattern (b). All 9 shift registers were fully operational. 221
6.30 Observed bias margins for magnetically biased shift registers as a function of magnetic bias current. 223
6.31 Amplified output of a 32-bit shift-register. The output pattern matches the input pattern, and the (correct) spacings of τ_{off} are shown for different clock frequencies. 224
6.32 Observed operational range for eSFQ shift-registers without (eSR) (a) and with (MeSR) (b) a magnetic corrective bias. Upper and lower lines of the same style form boundaries of the region of operation. 225
6.33 Calculated operational range for the second high-speed test setup, for eSFQ shift-registers without (`eSR`) (a) and with (`MeSR`) (b) a magnetic corrective bias. Upper and lower solid lines of the same style form boundaries of the region of operation, whereas dashed lines indicate the corresponding centers of the region. 227
6.34 Oscilloscope traces for a functional 184-bit shift register test. The complex test pattern is reproduced with the correct delay. A 3.2 MHz clock and data signal are applied. 228
6.35 Oscilloscope traces for high-speed 184-bit shift register test. The upper set of traces corresponds to $f_c = 1$ GHz, the lower to $f_c = 24.9$ GHz clock. In each case, the upper trace is the input data signal and the lower trace is the output data signal. 229
6.36 High-speed test results for 184-bit shift register. The ordinate refers to the delay difference between the investigated clock frequency and the reference frequency (1 GHz). 230
6.37 Measured correct functionality of 16-bit deserialiser for different input patterns, (a) and (b). This circuit divides the input serial stream into 16 output parallel streams. Random data and shift clock pulses are used for testing. Note that measured voltages are scaled automatically, making some constant-level bits appear noisy. 232
6.38 Observed bias margins for (magnetically biased) deserialisers as a function of magnetic bias current. 233

6.39 Output of free-standing deserialiser (a) and deserialiser behind modulator (b). . . . 235
6.40 Schematic of the RSFQ TFF. . . . 237
6.41 A decision-making tuple (DMT), interrogatable by a clock signal. 239
6.42 Schematic of cell `eTFF`, an eSFQ compatible version of the ubiquitous TFF cell. . . . 240
6.43 Schematic of TFF core of designed `eTFF` implementation. Details can be found in Section A.5.4. . . . 241
6.44 A 20 GHz simulation of the designed `eTFF` cell within an 8-bit shift register. . . . 242
6.45 Realised physical implementation of `eTFF` cell in Hypres's 4.5 kA cm^{-1} process. Details, including extracted paramters, may be found in Section A.5.4. . . . 243
6.46 Realised physical implementation of `eTFF` test in Hypres's 4.5 kA cm^{-1} process. . . . 243

A.1 Final circuit diagram of QOS-based multi-threshold comparator MC1 (Parameter values in Table A.1). . . . 272
A.2 Final layout of QOS-based multi-threshold comparator (MC1). . . 273
A.3 Final circuit diagram of CQOS-based multi-threshold comparator MC2 (Parameter values in Table A.2). . . . 274
A.4 Final layout of CQOS-based multi-threshold comparator MC2. . . 275
A.5 Final circuit diagram of CQOS-based multi-threshold comparator MC3 (Parameter values in Table A.3). . . . 277
A.6 Final layout of CQOS-based multi-threshold comparator MC3. . . 278
A.7 Final circuit diagram of CQOS-based multi-threshold comparator MC4 (Parameter values in Table A.4). . . . 279
A.8 Final layout of CQOS-based multi-threshold comparator MC4. . . 280
A.9 Final circuit diagram of CQOS-based multi-threshold comparator MC5 (Parameter values in Table A.5). . . . 283
A.10 Final layout of CQOS-based multi-threshold comparator MC5. . . 284
A.11 Final circuit diagram of QOS-based multi-threshold comparator `MX` (Parameter values in Table A.6). . . . 286
A.12 Final layout of QOS-based multi-threshold comparator `MX` (Parameter values in Table A.6). . . . 287
A.13 Final circuit diagram of CQOS-based multi-threshold comparator `MCX` (Parameter values in Table A.7). . . . 288
A.14 Final layout of CQOS-based multi-threshold comparator `MCX` (Parameter values in Table A.7). . . . 289
A.15 Circuit diagram of grey comparator test-bed. . . . 291
A.16 Chip `mvv2-gz-inv`, to experimentally verify grey-zones of (complementary) QOS comparators. . . . 292
A.17 Chip `mvw4-gz-c1-inv`, to experimentally verify grey-zones and correlation of (complementary) QOS comparators. . . . 293

A.18 Final circuit diagram of eSR shift register cell (Parameter values in Table A.8). 294
A.19 Final layout of eSR shift register cell (Parameter values in Table A.8). 295
A.20 Final circuit diagram of MeSR shift register cell (Parameter values in Table A.9). 296
A.21 Final layout of MeSR shift register cell (Parameter values in Table A.9). 297
A.22 Final circuit diagram of eDES deserialiser cell (Parameter values in Table A.10). 299
A.23 Final layout of eDES deserialiser cell (Parameter values in Table A.10). 300
A.24 Final circuit diagram of eTFF toggle flip-flop cell (Parameter values in Table A.11). 302
A.25 Final layout of eTFF toggle flip-flop cell (Parameter values in Table A.11). 303
A.26 Chip `mv348-esr-lqvar`, to experimentally verify eSR shift registers and investigate effect of L_b. 306
A.27 Chip `mv348-esr-lbvar`, to experimentally verify eSR shift registers and investigate effect of L_b. 307
A.28 Chip `mv348-esr-mag`, to experimentally verify MeSR shift registers and investigate effect of L_b. 308
A.29 Chip `mv348-edes-mag`, to experimentally verify eDES deserialiser. 309
A.30 Chip `mv348-edes8a4-mag`, to experimentally verify shorter eDES deserialisers. 310
A.31 Chip `mv348-etff-8b`, to experimentally verify functional operation of the `eTFF` cell, embedded in various configurations in 8-bit shift registers. 311
A.32 Chip `mv350_esr_lqvar_hf`, to experimentally verify high-speed operation of eSR shift registers and investigate effect of L_b. 313
A.33 Chip `mv350_esr_lbvar_hf`, to experimentally verify high-speed operation of eSR shift registers and investigate effect of L_b. 314
A.34 Chip `mv350_mag_hf`, to experimentally verify high-speed operation of MeSR shift registers with corrective flux bias. 315
A.35 Chip `mv350_pmd_edes16_hf`, to experimentally verify high-speed operation of eDES deserialisers. 316
A.36 Chip `mv352_magcomp_hf`, to experimentally confirm power dissipation of eSR and MeSR cells. 318
A.37 Chip `mv352_magcomp_nodiv_hf`, to experimentally confirm power dissipation of eSR and MeSR cells, without a clock divider. 319
A.38 Chip `mv352_esrlong_hf`, to experimentally confirm operation and power dissipation of longer shift register. 320

A.39 Chip mv352_esrlong_nodiv_hf, to experimentally confirm operation and power dissipation of longer shift register, without a clock divider. 321
A.40 Chip mv352_esrlong_dsw_nodiv_hf, to experimentally verify correct pattern reproduction of longer shift register at high speeds, without a clock divider. 322
A.41 Chip mv352_pmd_edes16_b_hf, to experimentally verify high-speed operation of eDES deserialisers, with correction. 323
A.42 Chip mv352_pmd_edes16_it2_hf, to experimentally verify high-speed operation of eDES deserialisers, with newer output drivers and better biasing. 324

List of Tables

2.1 Preliminary technical specifications for the SKA. 14

3.1 Differences between superconducting and conventional microelectronics and their effects on software toolsets. 47
3.2 Resources required for FastHenry execution. 70
3.3 Parameter attributes for margin analyses. 77
3.4 Software used for project. 82

4.1 Important parameters of the QOS multi-threshold comparator . . 94
4.2 4-bit Gray-to-decimal conversion table 96
4.3 Parameter choices for I_g experiment. 117

5.1 Extracted inductances and other circuit parameters for the single-threshold comparator. 130
5.2 Critical currents per unit line width of typical thin-film Nb superconductors [12] 135
5.3 Niobium layer thicknesses of IPHT process [13] 136
5.4 Parameters of QOS-based multiple threshold comparator with minimal L_p. 138
5.5 Grey-zones as estimated from SPICE simulations. 152
5.6 Structures developed to investigate various comparators 157
5.7 Measured performance metrics of the first batch of fabricated CQOS comparators. 165
5.8 Possible output sequence entry pairs and corresponding probabilities. 171

6.1 Experimentally determined functional bias margins for eSFQ shift registers across three chips (a comprehensive set of devices in different bias configurations) measured to establish functional correctness of the devices and attempt to identify desirable traits of the bias line layout. . . . 222
6.2 Bit energy E_b of 32-bit `eSR`- and `MeSR`-based shift registers for different bias configurations. . . . 226
6.3 Measured bias margins and associated power dissipation for the 184-bit shift register for several bias configurations. . . . 230
6.4 Experimentally determined bias margins for eSFQ deserialisers across two chips (a comprehensive set of devices in different bias configurations) measured to establish functional correctness of the devices and attempt to identify desirable traits of the bias line layout. . . . 233
6.5 States of `eTFF` cell and response to interrogation. . . . 240

A.1 Extracted parameters of MC1. . . . 276
A.2 Extracted parameters of MC2. . . . 276
A.3 Extracted parameters of MC3. . . . 281
A.4 Extracted parameters of MC4. . . . 282
A.5 Extracted parameters of MC5. . . . 285
A.6 Extracted and nominal parameters of `MX`. . . . 290
A.7 Extracted and nominal parameters of `MCX`. . . . 290
A.8 Extracted (inductances) and nominal parameters of eSR. . . . 298
A.9 Extracted (inductances and coupling factor) and nominal parameters of MeSR. . . . 298
A.10 Extracted (inductances and coupling factor) and nominal parameters of eDES. . . . 301
A.11 Extracted (inductances and coupling factor) and nominal parameters of `eTFF`. . . . 304
A.12 Structures developed to functionally investigate various eSFQ shift registers and deserialisers (Suite I). . . . 305
A.13 Structures developed to investigate high-speed operation of various eSFQ shift registers and deserialisers (Suite II). . . . 312
A.14 Structures developed to investigate high-speed operation of various eSFQ shift registers and deserialisers (Suite III). . . . 317

Chapter 1

Introduction

1.1 Motivation

On 25 May 2012 it was announced that the Square Kilometre Array (SKA) will be built primarily in South Africa [14]. The SKA is an ambitious project with a budget of $\sim 2 \cdot 10^9$ € and envisaged performance metrics that exceed those of all comparable contemporary instruments. Throughout the bidding process, optimism pervaded the South African team, emboldened by the perceived strength of its bid, available talent and as yet untapped potential. There is little doubt that the SKA will change South Africa for the better, with an increased appreciation for science and engineering already budding, in no small way thanks to an effective marketing campaign of the South African SKA team. The influx of money, international talent and attention should prove a catalyst for the development of science, technology and mathematics (STEM) fields all over South Africa and hopefully beyond.

But now, we have to build it. While the successful bid was naturally followed by elation, the monumental task of designing, optimising, building and commissioning the SKA is taking shape and, consequently, elation is replaced by zest, and optimism by realism. The scope of the SKA combined with competing constraints, most of which boil down to the "budget versus performance" kind, channel designer thinking towards postulated, not well-understood and experimental technologies that allow excursions in thinking outside the well-known limits and trade-offs. This dissertation is such an excursion.

The Square Kilometre Array will be the world's largest radio telescope, comprising $\sim 3\,000$ dishes and many other antennas. The bulk of the SKA equipment (antennas) will be located in Southern Africa, with the heart of the instrument situated near Carnarvon in the Karoo, South Africa, but some dishes located as far away as Ghana and Mauritius. A significant part of the SKA will also be built in Australia and neighbouring countries.

The central site is within the Northern Cape province of South Africa, re-

moved from large cities, extensive power generation and distribution networks, powerful radio towers (such as those used for mobile phone coverage), and frequent meteorological upheaval. This makes the site ideal for radio-astronomy, which relies on low-interference reception of faint signals.

Currently the site boasts KAT-7, an array that serves as an engineering pathfinder and demonstrator for multi-receiver radio-astronomy in South Africa. Although the objectives of KAT-7 were mainly engineering-related, it was recently commissioned for actual science, underlining the quality of the instrument. MeerKAT, a second pathfinder telescope and an official precursor to the SKA, will be completed in 2016.

The SKA itself will be built in two phases. The first phase, consisting of about 10% of the proposed equipment, will be completed in 2020. The second phase is currently slated for completion only in 2024.

The proposed sensitivity of the SKA exceeds that of all known radio telescopes. This is required for the scientific objectives of the SKA, which include probing hitherto unobservable celestial events (due to the faintness of their emissions). Other specifications are no milder, such as a raw data output of unprecedented volume and bandwidth. The technology required to build the SKA does not exist at this time, and reliance on its development does, of course, invite risk. The timeline associated with SKA construction is deemed sufficient to develop these technologies.

Superconductors have many desirable properties. Until several decades ago, superconductors were a somewhat esoteric subject with few engineering applications. Today, more than a century after their discovery, superconductors are found in medical equipment (such as MRI machines), ultra-sensitive meters (such as those employed for aerial prospecting), high-energy physics apparatus (such as particle accelerators) as well as electricity transmission networks. Although many aspects of superconductivity are not yet fully understood, their properties are so desirable, predictable and accessible that they have been employed in the definition of the international metrological standard of the volt.

Superconducting electronics are, in many ways, superior to conventional electronics. Integrated circuits that rely on superconductors rather than semiconductors run at greater speeds, greater sensitivity and greater power efficiency, and that with comparably inferior contemporary fabrication methods.

Superconducting analog-to-digital converters are a great example of the applicability of these features. Relatively simple circuits, they are capable of digitising signals of greater bandwidth when compared to conventional analog-to-digital converters. It is this capability that inspired our research: Digitisation of the entire band of interest of the SKA — which will span approximately 10 GHz or more — is currently proposed to be done with a (conventional) superheterodyne-inspired receiver assembly. We propose instead to harness the superior properties of superconductors to digitise the entire band of interest instantaneously, obviating the need for a (noisy, complicated)

superheterodyne-like receiver assembly and gaining the advantages of continuum sampling.

Digitising the entire signal as soon as possible (after the antenna) and doing all processing (such as demodulation of spectral "channels") in the digital domain constitutes a *software-defined radio* system. Such systems have become popular of late as digital technology has increased sufficiently in capability to handle such a demanding task.

Naturally, this places a burden on the digital backend of the SKA, which has to cope with greater instantaneous data bandwidths. While we do not address this issue comprehensively, we clearly demonstrate how superconductors can certainly play a role here too.

1.2 Superconductivity at a glance

1.2.1 Discovery

Superconductivity was discovered accidentally in 1911 by Heike Kamerlingh-Onnes (Figure 1.1a), to whom the Nobel Prize in Physics was awarded in 1913. Also the first person to liquefy helium, he was investigating the resistances of pure metals at increasingly low temperatures, discovering that the resistance of mercury fell sharply to "essentially zero" [15] at approximately 4.2 K (Figure 1.1b).

Today, many more superconductors are known. A superconductor is characterised by its *critical temperature* T_c, above which it loses its superconductivity (some become a *normal* conductor, others do not). At present, the highest known critical temperature of any superconductor is approximately 135 K at ambient pressure (critical temperatures at up to 164 K have been observed at extremely high pressure). This still makes superconductivity a phenomenon observable exclusively at cryogenic temperatures. A timeline of the discovery of superconducting materials is depicted in Figure 1.2.

1.2.2 Large-scale and analog applications

Superconductors exhibit many properties that are exceedingly desirable for electrical and electronic systems. Famously, they exhibit no resistance to dc currents, opening up the possibility of lossless energy transmission. Whereas the requirement of cryogenic cooling is certainly inconvenient, ingenious ways have been devised that make superconducting transmission lines of great scale possible, such as the LIPA Long Island Power Cable [16].

Other so-called large-scale applications include the use of superconducting magnets, for example in magnetic resonance imaging (MRI) or nuclear magnetic resonance (NMR) spectrometry. The response of a superconductor to magnetic flux is curious and still mystifies theoreticians and experimentalists

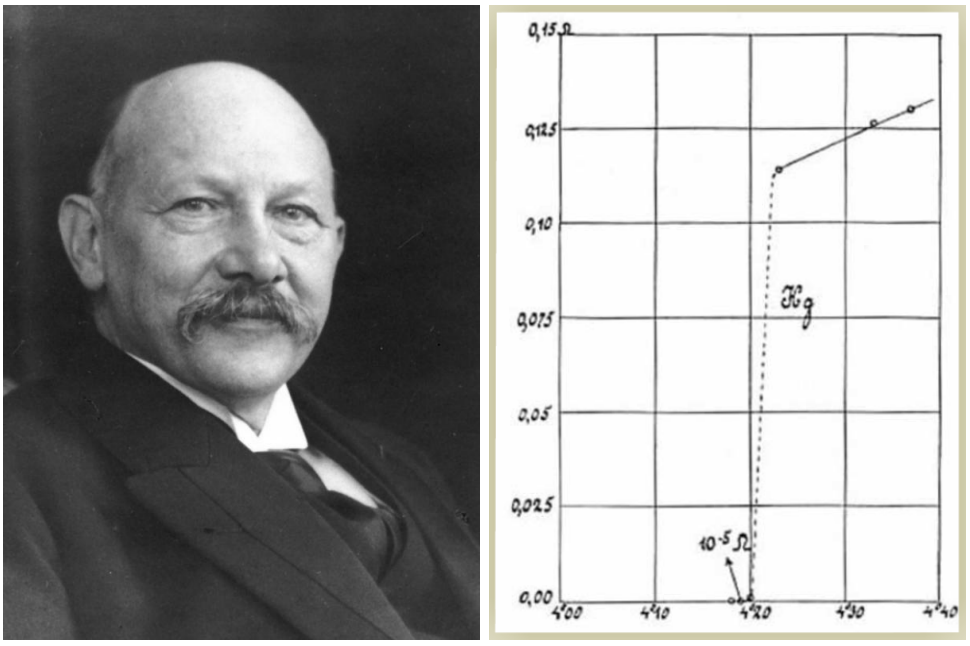

(a) (b)

Figure 1.1: Heike Kamerlingh-Onnes, Dutch physicist and Nobel laureate credited with the discovery of superconductivity (a), and his historic measurement of the resistance of a mercury wire (b) [1].

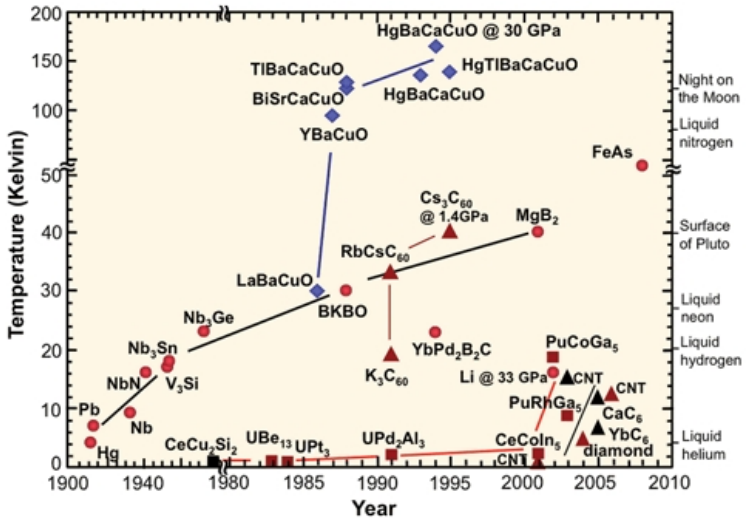

Figure 1.2: History of discovered superconductors. Blue entries are high-temperature superconductors (HTS), red entries low-temperature superconductors (LTS). Note the position of niobium (Nb) above the liquid helium line [2].

alike. Figure 1.3 shows a permanent magnet levitating over a superconducting surface, a result of the property of *flux exclusion*.

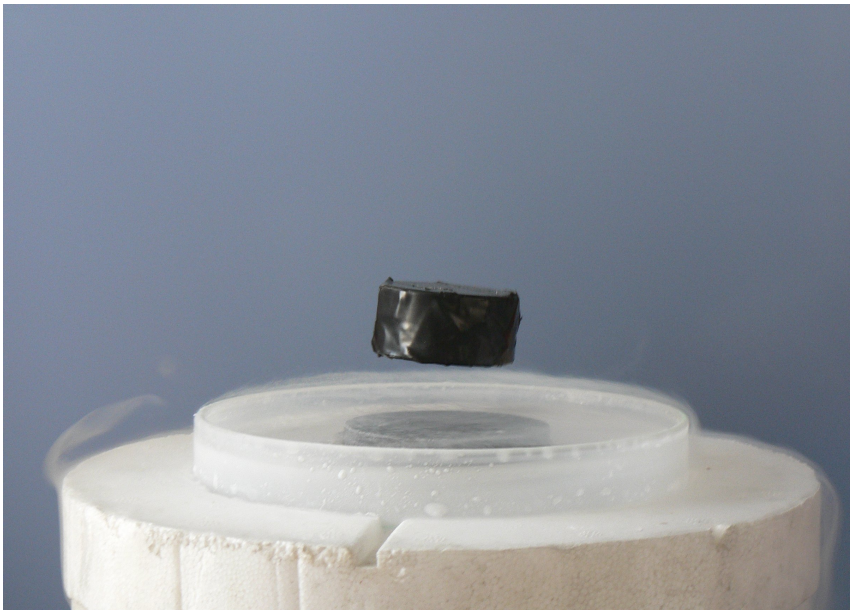

Figure 1.3: Levitation of a magnet over a superconductor [3].

We are more interested in the myriad small-scale applications of superconductors. The discovery of the Josephson junction, a superconducting two-terminal device that exhibits macroscopic quantum-mechanical properties, propelled superconducting electronics into a new era. Josephson junctions are two-terminal devices that are superconducting in certain conditions and non-superconducting (resistive) in others.

This well-known response can be exploited in several ways. For example, the Josephson junction can be exposed to an unknown environment and its response measured. From those observations, deductions can be made about the nature of the enclosing environment. Perhaps the most famous example is the *superconducting quantum interference device* (SQUID), which can be employed as a magnetometer of unparalleled sensitivity [17].

Similar methods are employed for non-invasive monitoring of bioelectromagnetic activity, for example in the practice of magnetoencephalography.

1.2.3 Digital applications

Due to their low-loss characteristics (among other properties), superconducting devices are usually very sensitive and dissipate very small amounts of power (if any). Furthermore and for the same reasons, characteristic times of superconducting devices are usually very low, making them ideal for high-speed (and high-bandwidth) applications. The transition between the superconducting

and resistive regime in the Josephson junction, for example, can be designed to occur very quickly. In this way, the Josephson junction can be employed as a "switch", paving the way for the development of logic circuits.

Controlled switching is the cornerstone of all digital logic. Digital data are encoded in analog signals. Examples of such encodings include voltage levels in *complementary metal-oxide-semiconductor (CMOS)* technology (and many other semiconductor logic families) and *single flux quantum* (SFQ) pulses in superconducting logic families such as the *rapid single flux quantum* (RSFQ) family. Consider Figure 1.4 for an illustration of the differences between these to logic families. In a voltage-state logic family, the logic bits are encoded as a voltage level (high for `1`, low for `0`). In the pulse-based logic family RSFQ, the logic bits are encoded as the presence of absence of a pulse (a pulse in a clock period means `1`, the absence of such a pulse means `0`).

Controlled generation of the voltage pulses from Figure 1.4b in superconducting circuits can occur at very high frequencies, a consequence of the low characteristic times associated with the Josephson junction. Whereas voltage-state logic families have also been successfully demonstrated in superconducting circuits (such as the COS-Logic family by Perold [18]), the greatest speed and power advantage so far has been demonstrated by pulse-based logic gates. An RSFQ TFF (toggle-flip-flop) running at 770 GHz was reported [19]. Superconducting logic circuits are in the form of integrated circuits, fabricated in custom foundries in a process akin to that employed for semiconductor integrated circuits (in fact, the fabrication process for superconducting circuits is arguably somewhat simpler, as no *doping* is required and clock frequencies of tens of GHz can be achieved in a µm technology, whereas sub-µm technology is required to achieve comparable performance with known semiconductors).

The advantages of logic circuits operable at high clock frequencies are manifold and seem most immediately applicable in high-performance computing (HPC). Whereas HPC is very much a goal of superconducting electronics development (explored further later), the limitations here lie in the level of *integration* achievable with superconducting electronics. Integration is a rough measure of the size and complexity of the circuits that can be fabricated and operated reliably. Semiconducting integrated circuits achieve extremely high levels of integration, up to several billions of switching elements (transistors) at present [**?**]. By contrast, superconducting integrated circuits are hampered by outdated fabrication methods, achieving only several tens of thousands of switching elements (Josephson junctions) at present [20]. Since the fabrication process of superconducting circuits is not much different from that of semiconducting circuits, it is reasonable to assume that levels of integration for superconducting circuits will continue to rise as fabrication methods improve. We are unaware of an immediate or impending encumbrance by fundamental limitations that may place an upper bound on superconducting circuit integration.

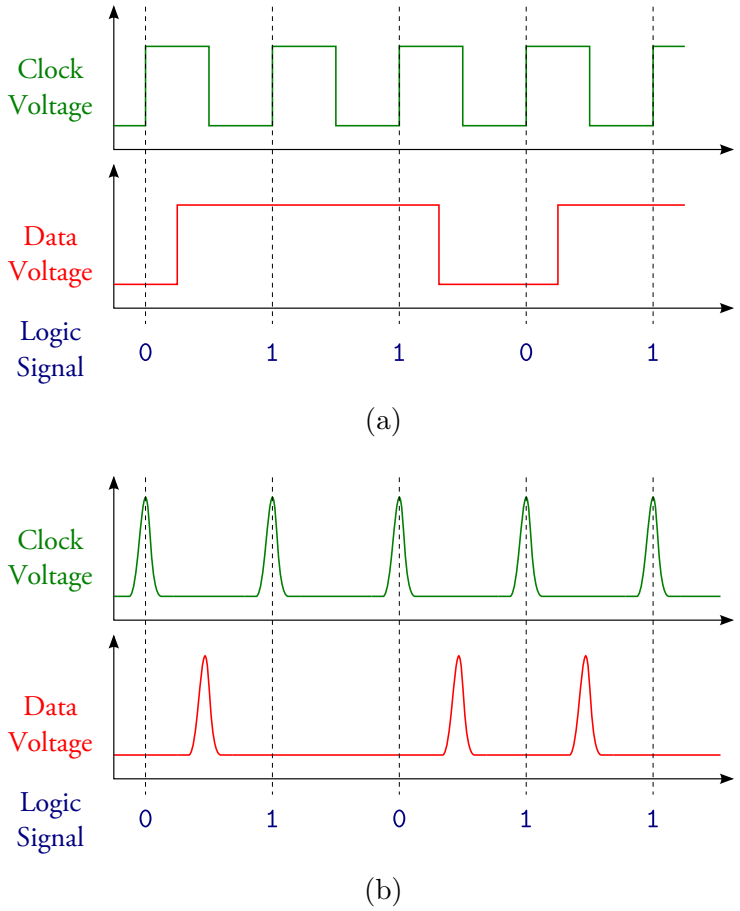

Figure 1.4: Comparison between voltage-state logic encoding, common in semiconductor logic families (a) and pulse-based logic encoding, common in superconducting logic families (b). In voltage-state logic, the logical state is measured on the rising edge of the clock. In pulse-based logic, the logical state is encoded in the presence or absence of a pulse in the clock period.

1.2.4 Our application

High speeds of operation are desirable for applications of digital logic. The SKA will produce immense quantities of raw data that need to be transported and processed. The bandwidth of the datapaths involved is vast, and total processing requirements are projected to approach the exascale (10^{18} operations per second) barrier.

Employing superconducting logic as processing elements is tempting and would undoubtedly bring many advantages to the SKA. A very real concern of building any exascale backend that is suitable for processing SKA data output is the power requirements of such a system, projected to approach 30 MW [5] or greater [21]. Distributing such quantities of power requires considerable engineering effort, magnified by the projected sensitivity of the SKA and asso-

ciated electromagnetic compatibility (EMC) concerns. Superconducting logic circuits dissipate far less power than their semiconductor counterparts. When extrapolating from power requirements of current circuits, a superconducting supercomputer of exascale proportions would have power requirements two to three orders of magnitude below even optimistic estimates for a semiconductor system of similar proportions [22]. A good summary of the state of the art from this perspective was recently published by Ortmann [23] (This discussion is continued in Chapters 6 and 7.)

But even at currently achievable levels of integration, superconducting electronics present a promising platform for developing another essential component of the SKA, the frontends of the individual receivers. Here the astronomical signal has just been converted to an electrical signal by the receiving antenna and needs to be digitised before processing. The device performing this conversion (from analog electrical signal to digital signal) is called an analog-to-digital converter (ADC).

ADCs are relatively small (simple) devices. A moderate level of integration is sufficient to achieve useful ADCs. Superconductors, characterised by their sensitivity and switching speed, constitute an ideal platform for the development of high performance analog-to-digital converters. An immediate advantage that superconducting ADCs have over semiconducting ADCs is the higher clock speed with which superconducting ADCs can digitise (and output) signal samples. The frequency with which a signal can be sampled determines the bandwidth of the analog signal by the Nyquist-Shannon Sampling Theorem [24].

Other capabilities of superconductors may also be exploited. For example, a SQUID may be used to provide quantised feedback into a signal, ensuring accurate quantisation. Furthermore, straightforward implementation of desirable ADC building blocks, such as near-perfect integrators and sensitive comparators, make superconductors eminently suitable for signal digitisation.

1.3 Objectives of dissertation

In this work, we propose components of a superconducting frontend for the Square Kilometre Array, with the intention of digitising the signal as early as possible in the receiver chain. This eliminates a large portion of the receiver chain, consisting of microwave components such as filters and mixers. The elimination of these components does not only make the design simpler and, correspondingly, reduces associated maintenance effort, but also reduces the number of links in the chain where noise (and other uncertainties) can disturb the signal.

Furthermore, we explore the merits of a newly invented ultra-low-power superconducting logic family, eSFQ. Other ultra-low-power superconducting logic families exist, but have several disadvantages that could hamper integra-

tion efforts [25]. As integration will be of critical importance for systems on the scale of the SKA, we feel that developing eSFQ is an essential step on the way to superconducting high-performance computing for radio-astronomical purposes.

More formally, the objectives of this dissertation are:

1. To introduce the reader to the relevant facets of the SKA, including the scientific objectives of this ambitious instrument, the engineering challenges they provoke and how we propose to meet them,

2. to introduce the reader to superconducting electronics — a somewhat lesser-known subject even among electronic engineers — as well as their properties and how these pertain to the SKA,

3. to describe the design process associated with superconducting circuits, as well as the tools we developed to facilitate it,

4. to describe our development of a high-level mathematical model of a superconducting ADC and investigate the parameters of said model as well as the performance achievable,

5. to relate our efforts at designing a superconducting ADC at the circuit level and fabricating key components,

6. to summarise our exploration of ultra-low-power superconducting logic families, as well as our efforts at designing and fabricating test circuits in a new ultra-low-power logic family that serve as useful demonstrators and would also benefit the SKA,

7. to describe experimental test setups for superconducting circuits, as these require specialised equipment and a controlled cryogenic environment,

8. to report the experimental results of our efforts, and link these to our expectations, and

9. to develop the conclusions we draw from our work and explore what these mean for the SKA.

In the next chapter, Chapter 2, we introduce the Square Kilometre Array as well as superconducting electronics. We introduce the Josephson junction and describe its place in superconducting integrated circuits. We develop where the superior switching speeds and other desirable properties come from and attempt to quantify the advantages for our purposes. Finally, we offer a system-level description of a proposed front-end for SKA receivers.

Next, in Chapter 3, we relate the design methods employed for superconducting circuits. This is a fertile area of discussion, as the state of superconducting design software (such as CAD tools) is dire and needs to be addressed.

A considerable amount of time went into the development of suitable software tools. We describe the tools we developed and modified to design and analyse our circuits.

After setting the stage, we continue to Chapter 4, where we develop a mathematical model of our chosen ADC architecture and analyse its performance and response to parameter variations.

With our newly-gained insight, we continue in Chapter 5 to design the key components of our ADC at the circuit level, describing fabrication limitations and how we overcome them. Next, we analyse different iterations of key components and design an experiment to obtain measurements of real, fabricated circuits. We report the results of those experiments, and how we propose to proceed.

Next, in Chapter 6, we address data transport. For this purpose we choose the new and undeveloped eSFQ logic family and develop first demonstrators. We choose these in such a way to be useful as readout circuitry for the ADC we propose and fabricate key versions of our designs for experimental verification. We also describe the experiments and results obtained.

Finally, Chapter 7 summarises our conclusions. We summarise the insights we developed and the results of our efforts before exploring what these mean to the SKA. We describe what must be done next to continue this fruitful research.

Chapter 2

A software-defined radio for the SKA

2.1 Introduction

The Square Kilometre Array (SKA) as well as superconducting electronics are introduced in this chapter. This is to set the stage for the remainder of our work and this text. Whereas general radio astronomy practices are only introduced superficially, the concepts of superconducting electronics warrant deeper focus here. This is because it is the fundamental properties of superconductors that make them so attractive not only for radio astronomy, but also for other high-end applications.

Two specific areas of superconducting electronics will be illuminated. Firstly, superconducting analog-to-digital converters will be introduced, as it is probably at this crucial point in the receiver chain where superconductors can bring great advantage to the SKA. Secondly, ultra-low-power superconducting digital electronics will be mentioned, as data readout and perhaps even post-processing will require such electronics.

A formal literature study is not presented here. There are three reasons for this:

1. The SKA is an ongoing topic with few hard references. Much speculation about eventual bands of interest, receiver structure, back-end requirements, etc. render a comprehensive technical summary about the current state of the project not useful to us.

2. An excellent, concise summary of superconducting analog-to-digital converters has been published very recently by Mukhanov, perhaps the most notable authority of this specific field [26], and should be consulted by the reader if desired.

3. Ultra-low-power superconducting electronics are a fairly new field, with little published literature available.

Nonetheless, the important facets of the influences guiding our work are introduced and discussed here (with appropriate citations). In the remainder of this work, discussions on literature surrounding certain aspects of our research will be presented contextually.

2.2 The Square-Kilometre Array (SKA)

The Square Kilometre Array is a planned radio telescope to be built in Southern Africa and Australia. Upon its completion in the early 2020s, it will be the world's largest and most sensitive radio telescope [5, 27].

2.2.1 Objectives of the SKA

The SKA "will address a wide range of fundamental questions in physics, astrophysics, cosmology and astrobiology. It will be able to investigate previously unexplored parts of the distant universe" [27].

The "science case" for the SKA includes the following objectives [28]:

1. **Probing the Dark Ages and the Epoch of Reionization**: By measuring the *redshift* of celestial bodies, we can determine their age (redshift is discussed in Section 2.4.3). We can detect *cosmic microwave background* (CMB) radiation at redshifts of 1000 and greater, but the most-redshifted celestial bodies we can observe are at redshifts of about 6 and below. Somewhere between redshifts of 6 and 1000, *neutral hydrogen*, which made up all *baryonic matter* at some point, became ionised. This was the first step towards the formation of structures in the universe. The SKA will be employed to probe this *Epoch of Reionisation*.

2. **Galaxy evolution, cosmology and dark energy**: The assembly of galaxies, how gas is converted to stars, will be examined by the SKA. In addition, *baryon acoustic oscillations (BAOs)* will be tracked as a function of redshift to provide information about the early evolution of galaxies.

3. **Origin and Evolution of Cosmic Magnetism**: While gravity has been present since the earliest times of the universe, magnetic fields may not have been. They could have first been generated by dynamo activity of galaxies or galaxy clusters. By measuring *Faraday rotation* of various sources over large fractions of cosmic time, the evolution of magnetic fields can be studied.

4. **Strong Field Tests of Gravity Using Pulsars and Black Holes**: The SKA will find *millisecond pulsars* and measure their timing precisely. A *Pulsar Timing Array* can then be constructed and employed for the

detection of nanohertz gravitational waves. By probing the space-time environment around extreme locations such as black holes, the *equation of state* of nuclear matter can be analysed.

5. **The Cradle of Life**: High sensitivity and resolution of the SKA will enable centimetre-wavelength thermal radiation from dust in *proto-planetary discs*. Changes while planets form can be monitored, providing clues towards the understanding of the planetary formation process. The SKA can also probe molecular clouds for complex *prebiotic molecules*. Finally, the SKA will be able to detect, for the first time, "leakage" or unintentional emissions from other civilisations.

All of these scientific objectives are ambitious, requiring data that cannot be obtained reliably with current telescopes. This explains the scientific justification for the SKA.

2.2.2 Why Radio-astronomy?

Radio astronomy refers to the observation of celestial bodies in the radio frequency spectrum. Radio frequencies roughly refer to the range from 3 kHz to 300 GHz, but there is a particular range suitable for radio astronomy. Consider the atmospheric opacity with respect to frequency, as depicted in Figure 2.1.

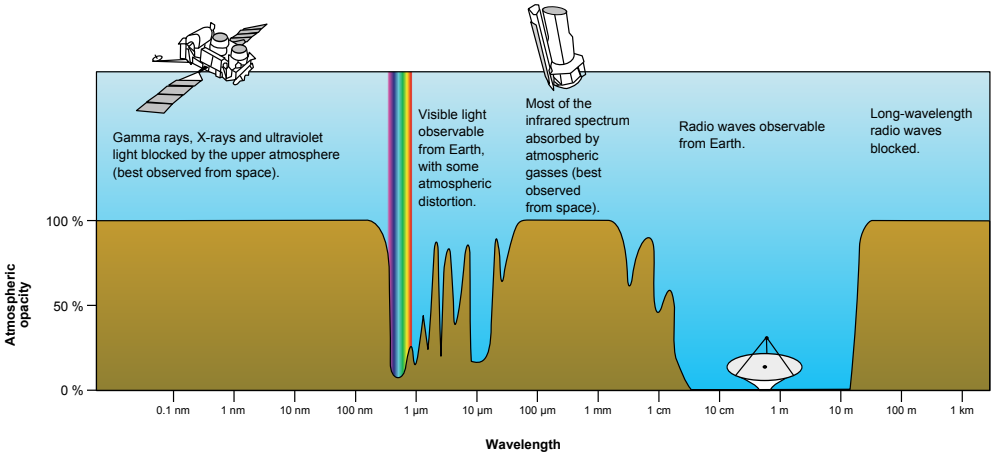

Figure 2.1: Atmospheric electromagnetic opacity across the spectrum [4].

The atmospheric electromagnetic opacity (here equivalent to atmospheric attenuation) is a measure of the ability of an earthbound observer to receive electromagnetic radiation of extra-terrestrial origin. The Earth's atmosphere reflects or absorbs radiation in large swathes of the spectrum. For example, the electromagnetic resonances of water vapour, oxygen and ozone are broadened by atmospheric pressure, ensuring absorption over large bands.

Fortunately, there is a region of excellent atmospheric transmission (low opacity) from wavelengths of several centimetres to several 10s of meters. This corresponds roughly to the frequency range 10 MHz – 10 GHz. It is here that much radio-astronomy takes place.

2.2.3 Technical capabilities

Technical specifications of the SKA stem from both scientific objectives and engineering principles. Finding an optimal agreement between the two is the subject of much current discussion and speculation. No specification for the SKA can be regarded as final, but preliminary specifications have been made available [29]. These are listed in Table 2.1.

Table 2.1: Preliminary technical specifications for the SKA.

Parameter	**Specification**
Effective collector area	$> 1\,000\,000$ m^2
Frequency range	70 MHz – 10 GHz
Sensitivity (area/system temperature)	5000 m^2/K (400 µJy in 1 minute) between 70 and 300 MHz
Survey figure-of-merit	$4 \cdot 10^7 - 2 \cdot 10^{10}$ m^4 deg^2/K^2 depending on sensor technology and frequency
Field-of-view	200 square degrees between 70 and 300 MHz, 1 – 200 square degrees between 0.3 and 1 GHz, 1 square degree maximum between 1 and 10 GHz
Angular resolution	< 0.1 arcsecond
Instantaneous bandwidth	Band centre ±50%
Spectral (frequency) channels	16 384 per band per baseline
Calibrated polarisation purity	10 000:1
Synthesised image dynamic range	$> 1\,000\,000$
Imaging processor computation	10^{18} operations/second
Final processed data output	10 GB/second

2.2.4 Projected topology of the SKA

The SKA will consist of three components: the low-frequency aperture array (0.07–0.45 GHz), the mid-frequency aperture array (0.4–1.4 GHz) and the full dish array (0.45-10 GHz) [5]. The low-frequency aperture array is to be built in Australia, whereas the high-frequency aperture array and the full dish array will be built in Southern Africa.

The SKA will be an *interferometer*. Interferometry is a widespread practice in modern radio astronomy, employing strategically placed dishes to obtain an image comparable to that from a much larger telescope. To cover the spatial frequencies of interest, both long and short baselines should be present in the interferometer. A baseline refers to the distance between two dishes in the interferometer. An interferometer with n_d dishes will thus have $\sim n_d^2$ baselines.

The projected topology of the full dish array is depicted in Figure 2.2.

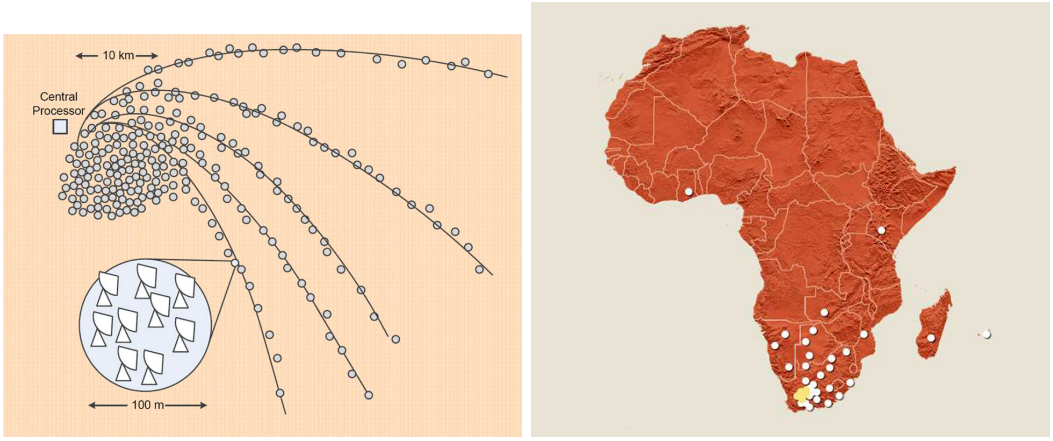

Figure 2.2: Projected dish locations for the SKA. A quasi-self-similar pattern with random perturbations [5] (a) ensures wide spatial-frequency coverage. The dishes will be placed across much of Africa (b).

The SKA will be built in two phases. The first phase, to be completed by 2020, consists of about 10% of the antennas. Science operations will be performed with Phase I. Phase II is to be completed by 2024 and will contain the remainder of the antennas.

Worldwide efforts towards the SKA have included and will include *precursor telescopes*. Some notable examples are South Africa's MeerKAT, to be completed by 2016 [30], Australia's ASKAP [31], and LOFAR in the Netherlands [32] (which is currently the largest, most sensitive radio telescope).

MeerKAT will consist of 64 dishes, with construction to be started soon. It will probably become a part of (and be superseded by) Phase I of the SKA, which will see an additional 190 dishes built. The "KAT" in MeerKAT stands for "Karoo Array Telescope".

Note also that KAT-7 [33], an engineering prototype for MeerKAT and the SKA consisting of 7 dishes, has already been built in the Karoo and has recently become operational for science activities.

2.2.5 Antenna structure

An artists impression of the SKA full dish array is reproduced in Figure 2.3a. The probable antenna design for MeerKAT is displayed in Figure 2.3b. This is a good indication of an eventual SKA antenna design.

Several *receivers* may adorn one antenna, each for a different frequency range. The receivers consist of a *horn* (or other antenna, such as the log-periodic array employed in the ATA [34]) which is followed by the *feed assembly*, such as the feed assembly of the KAT-7 telescope antennas, as depicted in Figure 2.3c.

At the horn (or other antenna), the transition between free-space signal (the astronomical signal) and guided signal occurs. In the feed assembly, the guided signal is conditioned to make it more suitable for transport and further processing. This includes amplification and, perhaps, synchronisation or conversion to a different form (such as *RF-over-fibre*).

We believe that the SKA would benefit from superconducting electronics in the feed. The feeds for the SKA will be cryocooled to lower noise contamination, making it a suitable location for superconductors.

2.3 Introduction to SFQ electronics

In recent decades, superconductivity has progressed from a phenomenon of mainly theoretical interest to an engineering tool of considerable practical value. The low-loss, low-dispersion characteristics of superconductors make them ideally suitable for low-power, high-sensitivity applications.

Superconducting filters exhibit unparalleled performance in terms of return loss, transfer loss and sharpness of roll-off [35].

Superconducting magnets are employed where large magnetic fields are required. The superb current-carrying capacity of superconductors enables fields of $\sim 10\,\mathrm{T}$, enabling a large number of ambitious engineering feats, such as particle acceleration [36] and magnetic resonance imaging [37].

Superconducting mixers rely on the Josephson junction, examined later. Frequencies in the sub-THz range are the hallmark of such mixers [38].

SQUIDs, or superconducting quantum interference devices, are ranked as the most sensitive magnetometers known [17]. They have diverse applications ranging from prospecting to measuring biological activity in the brain [39].

Digital superconducting electronics have not been as popular until about two decades ago. The invention of RSFQ logic [9], promising high clock frequencies and low power dissipation, has catalysed the field of digital super-

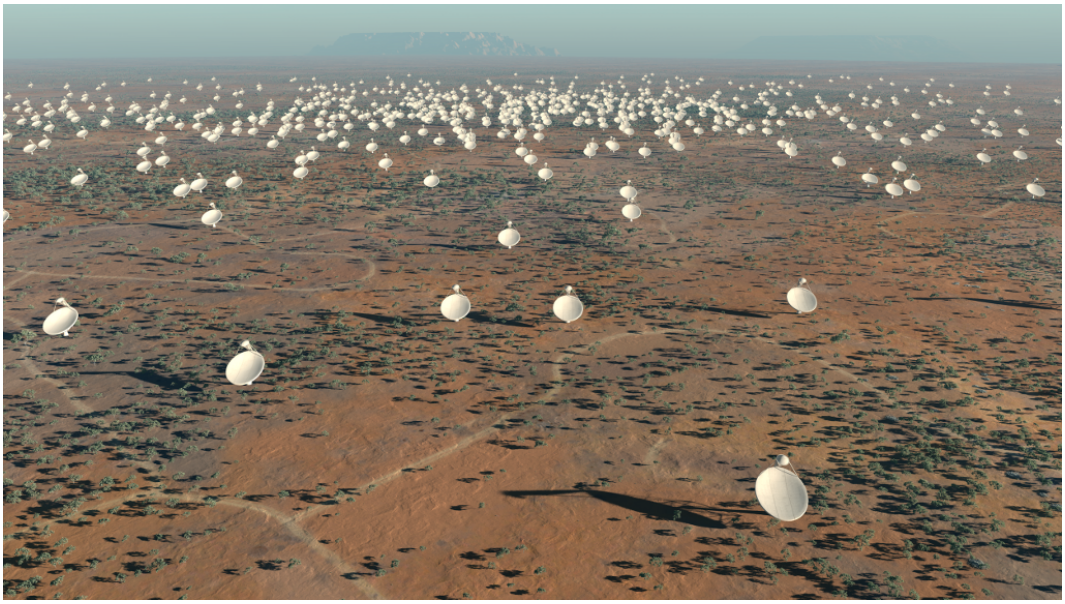

(a)

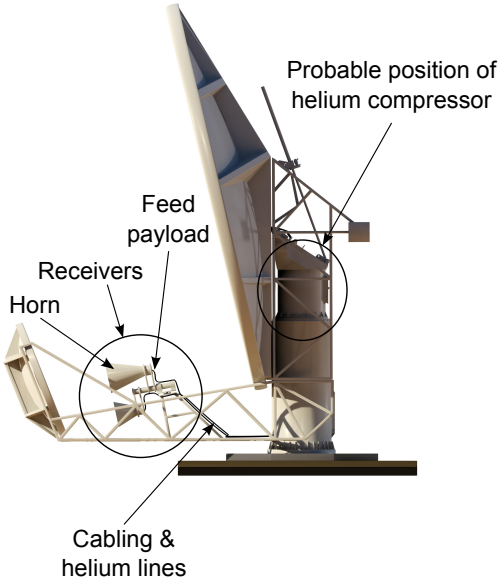

(b)

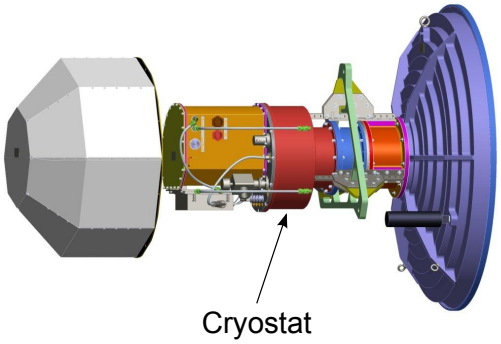

(c)

Figure 2.3: Artists impression of the eventual SKA full dish array (a). Possible projected offset Gregorian antenna design for MeerKAT (b). Cryocooled feed of KAT-7 antenna (c). The helium compressor referred to in (b) is a part of the cryocooler [6, 7].

conducting electronics and spawned much research activity across the globe, although commercial applications remain elusive.

We explore the properties of superconductors, elementary superconductor devices and basics of superconducting circuits in the remainder of this section. Superconductors were a surprise discovery and continue to puzzle us today. The present author must apologise that the coverage here will not progress far beyond the superficial in this restricted format. An excellent summary of the history and applications of superconductivity has recently been published [40].

2.3.1 The applicability of superconductors

Superconductors are well-known for their property of no dc resistance. In other words, a superconductor acts as an experimentally verifiable perfect electrical conductor [41], for dc currents. Many applications of this property come to mind, such as low-loss transmission of power and the conservation (storage) of energy [42].

Superconductors are characterised by a critical temperature, T_c and a critical magnetic field, H_c. At temperatures above T_c or magnetic fields above H_c, a (conventional) superconductor reverts to normal metallic, generally Ohmic, behaviour.

The earliest superconductors discovered had critical temperatures of $< 5\,\mathrm{K}$, succeeded shortly by superconductors with $T_c \sim 25\,\mathrm{K}$. To this day, superconductors with $T_c < 25\,\mathrm{K}$ are known as conventional superconductors.

In 1986, *high-temperature superconductivity* (HTS) was discovered. Cuprates such as $YBa_2Cu_3O_7$ (known as YBCO) were found to become superconducting at $T_c \approx 93\,\mathrm{K}$. Later, compounds with $T_c \approx 130\,\mathrm{K}$ (and up to 160 K at high pressure) were discovered. While these compounds still require cryogenic cooling, YBCO already has a T_c comfortably above 77 K, the boiling point of nitrogen. As liquid nitrogen is far cheaper and easier to handle than, say, liquid helium, much hope was placed into the development of HTS materials as a technology suitable for electronics.

The discovery of HTS does not supersede low-temperature superconductors (LTS). Until now, the material properties of HTS compounds, as well as certain fundamental properties, have rendered them far less suitable for practical electronics than LTS materials [42, 43].

Perfect electrical conductors do not allow magnetic flux to enter, which is a direct consequence of perfect conductivity [41]. Superconductors, however, do not only not allow flux to enter (with certain exceptions in Type II superconductivity), but will expel flux present in the conductor when it reaches superconductivity. This property is usually referred to as the Meissner effect [44].

A curious situation occurs when examining a superconducting ring. If superconductors do not allow flux to enter, can flux ever thread a supercon-

ducting ring? The answer is yes, as long as the flux entered the ring when it was not (entirely) superconducting.

If such a superconducting ring is heated above T_c, flux can penetrate the metal (screening currents will die down because of its non-zero resistance) and thread the ring. If the temperature now falls below T_c, the ring becomes superconducting. Any flux threading the ring is "frozen" in place, it cannot exit the ring (as to do so would require it to enter a superconductor).

This phenomenon of flux-freezing follows directly from the property of perfect conductivity. Another property of superconductivity, however, cannot be explained by perfect conductivity alone: the property of *flux quantisation*.

Flux threading a superconducting ring is not only "trapped", it is also quantised:

$$\Phi = n\Phi_0, \tag{2.1}$$

where Φ_0 is a quantity known as the *magnetic flux quantum*. It is a constant, defined as

$$\Phi_0 = \frac{h}{2e} \approx 2.068 \cdot 10^{15}\text{Wb}. \tag{2.2}$$

The property of flux quantisation in a superconducting ring follows directly from the quantum mechanical underpinnings thought to explain superconductivity. Flux quantisation can be experimentally observed, it is a *macroscopic quantum effect*.

The absolute value of the flux quantum is fairly small. Other useful representations are

$$\Phi_0 \approx 2.07\,\text{mV ps}, \text{ and } \Phi_0 \approx 2.07\,\text{pH mA}.$$

Flux quantisation has many tantalising practical applications. SQUIDs are essentially superconducting loops with a hatch (or several hatches), which allows flux to enter or leave, one flux quantum at a time. By exposing the SQUID to a magnetic field, then counting the flux quanta stored in the SQUID, a precise measurement of the magnetic field can be made. The resolution of such a *magnetometer* can be improved by adding a pickup transformer. Practical SQUIDs are the most sensitive magnetometers known [17, 45].

Flux quantisation provides a means to precisely discriminate between different energy states. This property is of immense value to superconducting electronics.

2.3.2 The Josephson junction

The Josephson effect

Brian Josephson's prediction of the the Josephson effect [46, 47] earned him a (jointly awarded) Nobel Prize in 1973. His prediction predated its experimental

observation. The *stationary Josephson effect* predicts that, if two superconductor "banks" are separated by a thin insulating barrier, a supercurrent I can flow across the barrier:

$$I = I_c \sin \phi_J, \tag{2.3}$$

where ϕ_J is the difference in superconducting phase between the two banks and I_c is a quantity termed the *critical current* of the junction. This structure is illustrated in Figure 2.4.

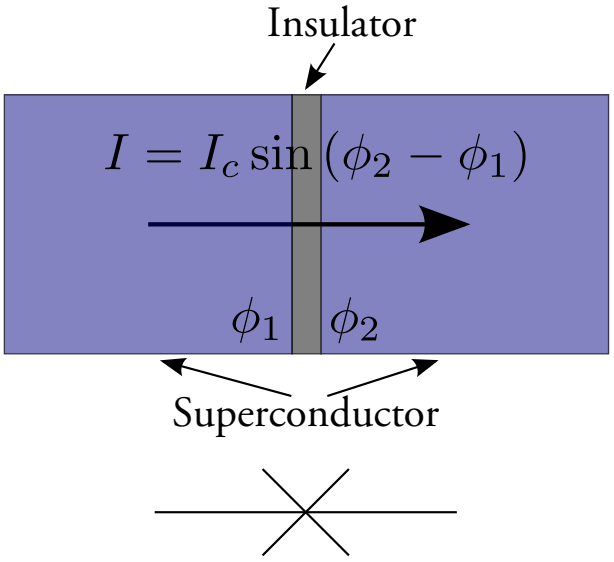

Figure 2.4: A Josephson junction and its circuit symbol.

The *non-stationary Josephson effect* describes the evolution of superconducting phase across the barrier with voltage v_J:

$$\frac{d\phi_J}{dt} = \frac{2ev_J(t)}{\hbar} = \frac{2\pi}{\Phi_0} v_J(t), \tag{2.4}$$

where Φ_0 is the magnetic flux quantum, as introduced above.

After making peace with the concept of currents flowing through conductors without any voltage drop, we are now confronted with the remarkable idea of such currents flowing across an insulator too. Indeed the Josephson junction is superconducting, up to its critical current $\pm I_c$ it behaves the same way in the presence of dc currents as other superconductors do.

When the critical current of the junction is exceeded, it behaves as a normal Ohmic conductor (save the behaviour at the transition, which to some extent depends on the way the junction is manufactured). This partitions the IV-curve of the junction, depicted in Figure 2.5, into a superconducting and a non-superconducting region. As $I \gg I_c$, the junction resembles a resistor with

resistance R_n. R_n is termed the *normal resistance* of the junction. The normal resistance results from current flowing across the insulator in a conventional manner.

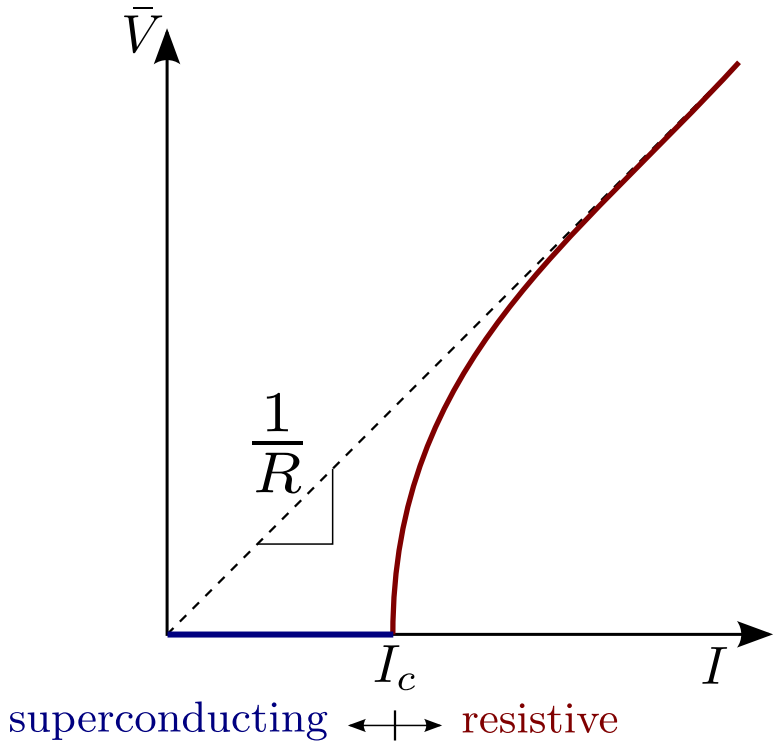

Figure 2.5: Simplified IV-curve of the Josephson junction. R refers to the resistance of the (parallel) resistive channel, and may be dependent on I.

Non-ideal junctions

Josephson junctions for our applications are generally realised as a "sandwich" feature on superconducting integrated circuits. Two superconducting banks are stacked vertically in a column, sandwiching an insulator of area A between them. Here A is proportional to the I_c realised. This is discussed in greater detail in Section 3.2.4.

This structure resembles a parallel-plate capacitor. Since the insulator should be thin to enable tunnelling [41], the distance between the parallel plates d is small. As the capacitance of such a structure is inversely proportional to d (see (3.1) later), its effect is significant.

So far, besides the ideal Josephson junction characterised by (2.3) and (2.4), we have identified a parallel resistance and capacitance that accompanies every real junction. We model this real junction using the RCSJ model (resistive and capacitively-shunted junction), shown in Figure 2.6.

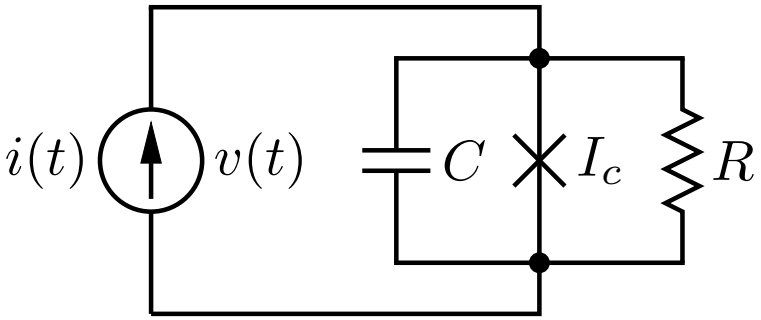

Figure 2.6: RCSJ model of the Josephson junction.

Considering this model and the Josephson effect, we conclude the relationship

$$i(t) = C\frac{\Phi_0}{2\pi}\frac{d^2}{dt^2}\phi(t) + \frac{\Phi_0}{2\pi R}\frac{d}{dt}\phi(t) + I_c \sin\phi(t). \tag{2.5}$$

We may rewrite this as a dimensionless quantity, obtaining

$$\frac{i(t)}{I_c} = \beta_c\frac{d^2\phi(t)}{d\tau'^2} + \frac{d\phi(t)}{d\tau'} + \sin\phi(t), \tag{2.6}$$

where

$$\tau' = \frac{t}{\tau_J}, \tag{2.7}$$

$$\beta_c = \frac{\tau_{RC}}{\tau_J}, \tag{2.8}$$

and

$$\tau_J = \frac{\Phi_0}{2\pi I_c R}, \tag{2.9}$$

$$\tau_{RC} = RC. \tag{2.10}$$

The quantity β_c from (2.8) is known as the Stewart-McCumber parameter [?, ?].

The pendulum analogy

We now introduce an interesting analogy. Consider the damped pendulum from Figure 2.7. Quantity τ refers to the torque applied to the pendulum, whereas θ represents the angle the pendulum makes with the vertical axis intersecting its fulcrum. Furthermore, ℓ is the length of the pendulum and m its mass. The gravitational constant is g.

The dynamics of the pendulum are governed by

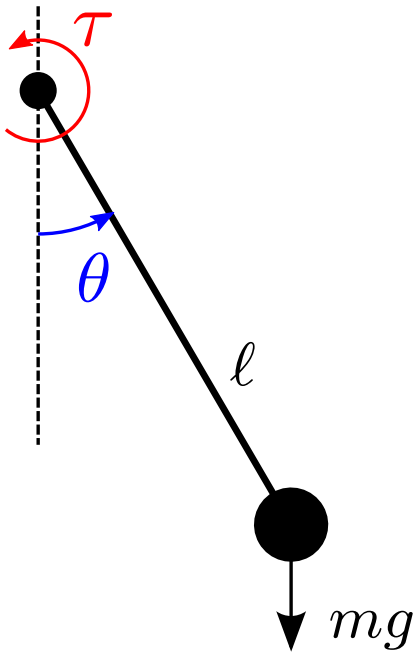

Figure 2.7: Quantities of a damped pendulum.

$$\tau(t) = m\ell^2 \frac{d^2}{dt^2}\theta(t) + D\frac{d}{dt}\theta(t) + mg\ell \sin\theta(t), \tag{2.11}$$

where D is the damping factor.

Comparing (2.5) and (2.11), we observe that they have the exact same form, with torque τ corresponding to current i, angle θ corresponding to phase ϕ and various constant coefficients.

Hence, the behaviour of the Josephson junction can be described by the dynamics of a damped pendulum.

If to a damped pendulum is applied a short *pulse* of torque, one of two things may happen:

1. If the applied pulse has a high enough amplitude, the pendulum will make at least one revolution, returning to rest after some time. The number of revolutions depends on the damping factor D. Each revolution is associated with a 2π increase in θ.

2. If the amplitude is too low, the pendulum angle θ will not reach $\frac{\pi}{2}$ and the pendulum will return to rest (with some possible zero-crossings of θ) without making any revolution.

If a constant torque is applied to the pendulum before the torque pulse, giving it an initial angle of $\theta \lesssim \frac{\pi}{2}$, the amplitude required of the pulse to

precipitate a revolution is smaller. It must be just high enough to apply "critical torque", which corresponds to making θ cross the $\frac{\pi}{2}$ mark.

Similarly, a Josephson junction biased by a current I_b so that $\phi \lesssim \frac{\pi}{2}$ will experience at least one 2π "phase jump" when receiving a current pulse of sufficient amplitude. We say that the junction "switches". The closer ϕ is to $\frac{\pi}{2}$, the smaller the amplitude required of the pulse. Note that $\phi = \frac{\pi}{2}$ when $i(t) = I_c$.

Junction switching

The Stewart-McCumber parameter β_c introduced in (2.8) describes the "damping" behaviour of the shunted junction. When $\beta_c \gg 1$, the junction is "underdamped", whereas $\beta_c \ll 1$ describes an "overdamped" junction. "Critical damping" is achieved when $\beta_c = 1$. The behaviour of the voltage response of the junction during switching is determined by the damping, as illustrated in Figure 2.8.

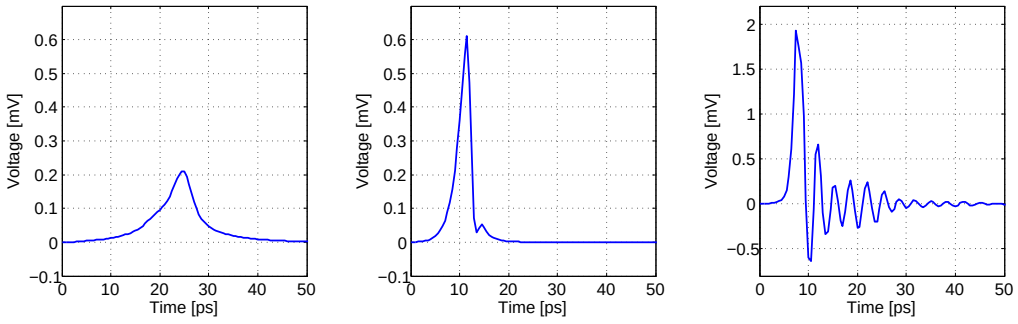

Figure 2.8: Switching behaviour of overdamped (a), critically damped (b) and underdamped (c) junction in the presence of a resistive channel.

Similarly, the IV-curve of the Josephson junction is ascribed hysteresis in the RCSJ model, as depicted in Figure 2.9. Once a junction has traversed out of the superconducting region, the current must fall to a value below I_c (generally called $I_{c_{\min}}$, which depends on β_c) to return the junction into the superconducting region. For the underdamped case $I_{c_{\min}}$ approaches 0, whereas for the overdamped case it approaches I_c.

From (2.8) we know that only when $\beta_c = 1$ are τ_J and τ_{RC} equal. This allows the shortest settling time of the junction, as depicted in Figure 2.8b. When $\beta_c > 1$ ($\beta_c < 1$), then $\tau_J < \tau_{RC}$ ($\tau_{RC} < \tau_J$), which means that while one of the characteristic times is shorter than for $\beta_c = 1$, the other is longer, drawing out the switching process. This is clearly evident in Figure 2.8, where

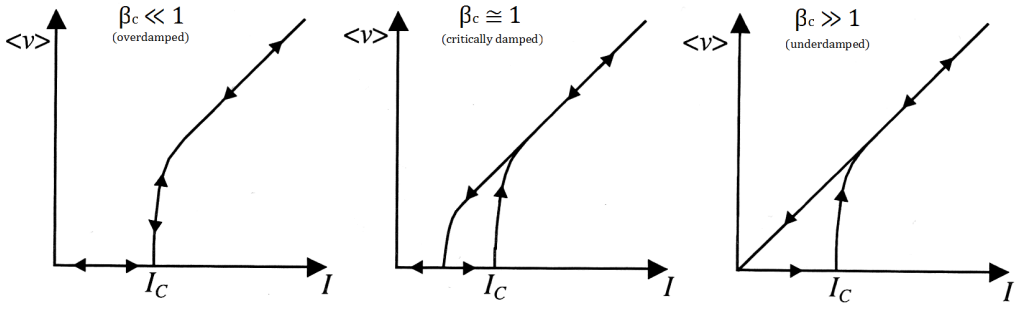

Figure 2.9: Hysteretic IV-curves of the Josephson junction [8].

the switching process takes longer (albeit for different reasons) for both $\beta_c \gg 1$ and $\beta_c \ll 1$ (when compared to $\beta_c = 1$).

2.3.3 RSFQ Electronics

The voltage responses depicted in Figure 2.8 differ substantially from each other, yet they have one thing in common: their areas are equal. While this is not obvious from the images, it follows directly from the non-stationary Josephson effect. Applying 2.4, we conclude that

$$\text{Area of pulse} = \int_{t_1}^{t_2} v(t)dt = \int_{t_1}^{t_2} \frac{\Phi_0}{2\pi}\dot{\phi}(t) = \frac{\Phi_0}{2\pi}[\phi_1 - \phi_2] = \Phi_0, \qquad (2.12)$$

where the pulse occurs between t_1 and t_2 and we know that the difference in phase, $\phi_2 - \phi_1$, is 2π. Thus, junction switching is accompanied by a quantised voltage pulse. This pulse is known as a *single flux quantum* (SFQ) pulse.

Since approximately 1985, a logic family has developed around this phenomenon, known as the *rapid single flux quantum* (RSFQ) logic. Research pioneered at Moscow State University by Mukhanov and others has resulted in a landmark paper by Likharev and Semenov [9], the gist of which we briefly introduce in the remainder of this section.

The RSFQ logic family relies on three basic elements treating SFQ pulses:

1. A *transfer element*, which transfers SFQ pulses from one physical location to another (usually adjacent) location,
2. a *storage element*, which stores SFQ pulses, and
3. a *decision element* that decides whether or not to transmit an SFQ pulse.

We examine two staple gates of the RSFQ logic family, the Josephson transmission line, made from transfer elements, and the D-cell (also known as the DFF, D-Flip-Flop or DRO, Destructive ReadOut gate), made from a storage element and a decision element.

The Josephson transmission line

The discrete Josephson transmission line, or JTL, is schematically depicted in Figure 2.10. The bias current I_b is usually set at about $0.7I_c$, which ensures that ϕ_J of the junctions is sufficiently close to $\frac{\pi}{2}$ to enable even a small "kick" to precipitate a 2π phase jump.

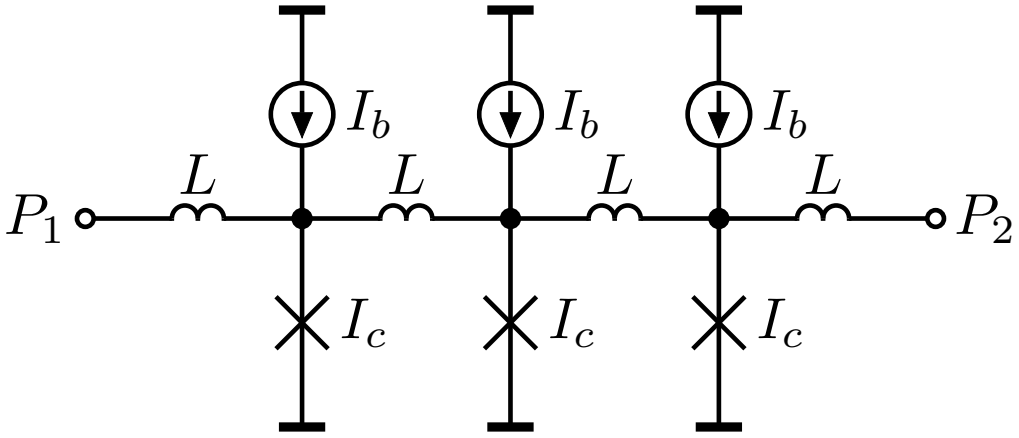

Figure 2.10: Schematic of Josephson transmission line (JTL).

Generally it is best to employ critical damping for junctions that generate SFQ pulses. Minimal oscillations ensure that flux quanta propagate in a predictable manner, which maximises operating margins (the concept of which is explained in Chapter 3).

Junctions fabricated in the processes commercially available to us ([48, 49]) are naturally underdamped, with β_c of the order 100. To obtain critical damping, a shunt resistor R_s is placed in parallel with the junction. Note that this simply lowers the R already present in the RCSJ model. The magnitude of R_s follows from solving (2.8) for $\beta_c = 1$.

It is briefly worth noting here that the quantity Φ_0, the magnetic flux quantum, has a relatively small magnitude. Hence, the energy associated with an SFQ pulse is small, but non-zero, it is dissipated in the resistor R from Figure 2.6. This explains the need for the bias current I_b, it replenishes the energy lost to the generation of SFQ pulses.

It may not be immediately clear, but a good way to supply the required kick to the junction is by applying an SFQ pulse to it. Such a voltage pulse applied at P_1 temporarily increases the current through the left-most inductor L, which, in turn, increases the current through the left-most junction. If the increase is high enough, the junction switches, regenerating the pulse. The regenerated pulse propagates along the JTL in an analogous manner.

By stringing transfer elements in series like this, SFQ pulses can be transported across essentially arbitrary distances, regenerated at each junction along the way. By carefully choosing critical currents and biases, an SFQ pulse

can be employed to switch two parallel subsequent junctions, as illustrated in Figure 2.11.

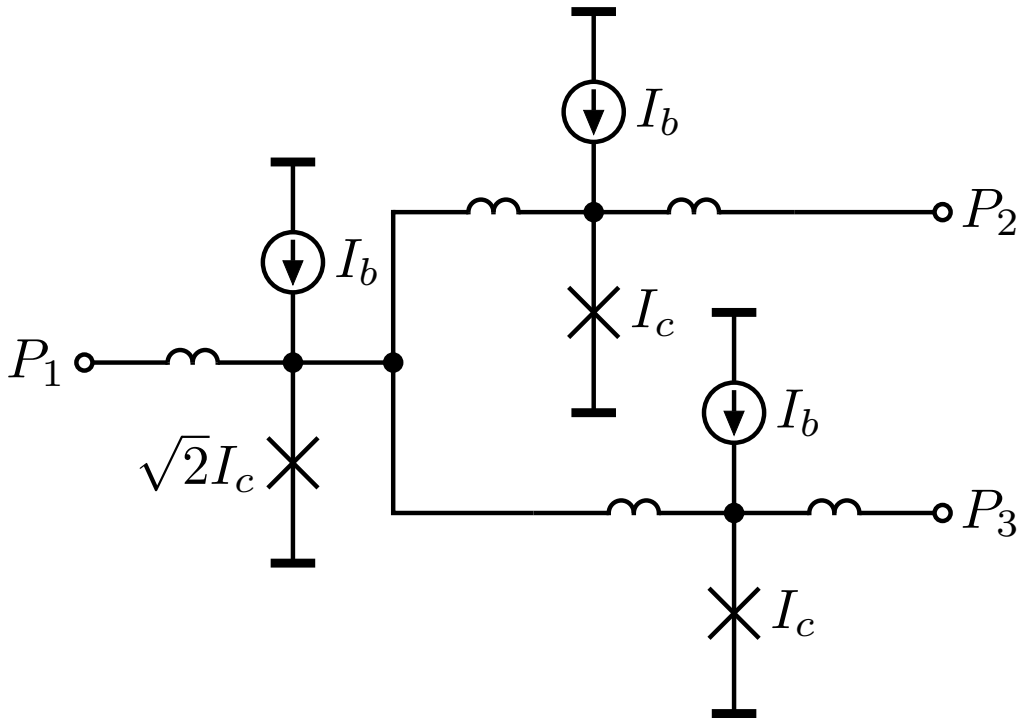

Figure 2.11: Generating multiple SFQ pulses using the $\sqrt{2}$-rule, as suggested by Likharev et al [9].

Pulse-splitting is a ubiquitous requirement, featuring in, for example, distribution of the clock network.

The D-cell (DFF)

Creating a D-cell in RSFQ logic requires two things: the ability to store an SFQ pulse, and the ability to detect whether such a pulse is stored. A popular DFF schematic is depicted in Figure 2.12.

Initially, the DFF is non-storing. Junction J_d is biased by I_b. An SFQ pulse arriving at `Set` would switch J_d, depositing a flux quantum into the loop made up of J_d, L, J_2. This loop represents the storage element. As J_2 is unbiased it does not switch, but becomes biased by the current now circulating in the superconducting loop. We say that the DFF is storing.

Junctions J_1 and J_2 make up a *decision-making pair* (DMP) that can be interrogated by a clock pulse to determine whether or not the DFF is storing. The critical current of J_1 is designed slightly smaller than that of J_2. If a pulse arrives at `Clock` in the non-storing state, J_1 switches as J_2 is unbiased. If a pulse arrives in the storing state however, J_2 switches instead.

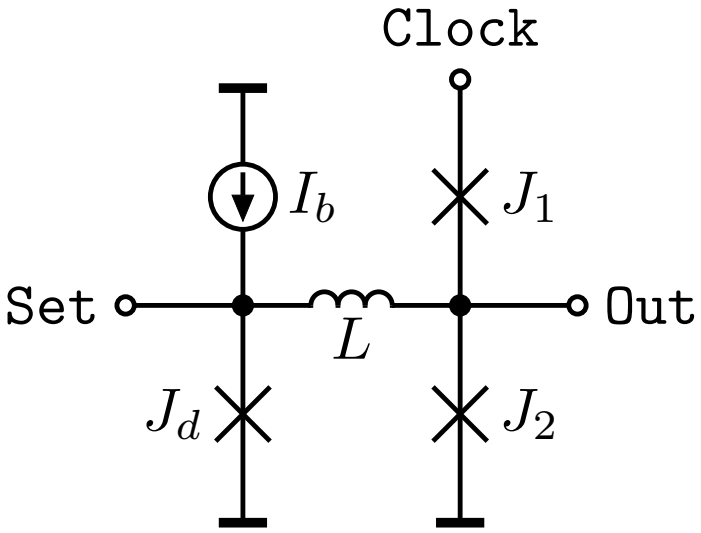

Figure 2.12: Schematic of D-cell [9].

When J_2 switches, an SFQ pulse appears at `Out`, signifying that the DFF was storing. The decision-making process is a destructive one, whether or not the DFF was storing before the `Clock` pulse arrived, it will be non-storing after interrogation. This destructive read-out is what the "D" in DFF refers to.

The DFF illustrates the way binary information is represented in RSFQ logic. It is tightly coupled to the concept of a clock. A `1` is represented by the presence of a pulse in a clock period, whereas a `0` is represented by the absence of such a pulse.

The sequential nature, clock requirements and low-power aspects of RSFQ logic are explored in greater detail in Chapter 6.

2.4 Digitisation of the astronomical signal

2.4.1 Analog-to-digital convertors (ADCs)

Role of analog-to-digital conversion

An *analog-to-digital converter* (ADC) is a mixed-signal device that transports a signal from the analog to the digital domain. The converse device is a digital-to-analog converter (DAC). ADCs normally have a way to apply an analog input signal and output a number of digital bits that indicate the level of the input signal at a certain point in time.

Many modern systems rely on ADCs and DACs to switch between analog and digital domains wherever convenient. Analog signals are required when physical distances must be bridged, whether these are in the form of a cable or free space (wireless transmission). Digital signals are required in modern

signal processing techniques that take place on DSPs or personal or embedded computers. Particularly communication systems rely on both domains to be effective.

In radio astronomy, signals are received by an antenna (or many antennas, as in our case). As described in Section 2.2, interferometry requires correlation of data received by different antennas that may be very far away from each other. To do this (entirely) in the analog domain would be very impractical, hence the signals received should be digitised and *timestamped* to allow for realtime or later processing.

While the signal is in the analog domain (such as in space, or beyond the antenna in an electrical network), it is subject to noise contamination. Heat from the environment contributes stochastically to the fundamental quantities that determine the signal level (such as voltage, fields, heat), distracting from the true level of interest. In radio astronomy received signals may be very faint (the faintest signals resolvable are a key performance indicator of radio telescopes), exacerbating the problem of noise contamination.

Once a signal is in the digital domain, it can be transmitted and manipulated without any loss of information. When physical distances are to be crossed the signal must be transported in the analog domain, but as the signal is in the form of digital bits this can usually be assured with negligible error probability. Thus, noise contamination is no longer a concern once the signal is in the digital domain.

To limit the exposure of the signal to contaminants, it thus seems prudent to transport the received signal to the digital domain as soon as possible in the receiver chain.

Components of analog-to-digital conversion

Analog-to-digital conversion is formally defined as the combination of three processes: *sampling*, *quantisation* and *coding* [50]. The input to an ADC is an analog signal $x(t)$ and the output is the discrete signal $x[k]$. *Sampling* refers to "discretisation in time" whereas *quantisation* refers to "discretisation in amplitude" (see Figure 2.13). *Coding* refers to the way in which the output bits of the ADC are to be interpreted as the level of $x[k]$.

Sampling is usually achieved by *triggering* the ADC at the desired points in time, often by a clock signal. The points in time may be equally spaced, as in Figure 2.13, but this is not necessarily so.

A popular means of representing the levels of $x[k]$ involves, of course, binary code. The output bits supplied by the ADC may be interpreted as an unsigned integer that is in turn interpreted as a proportion of the maximum output. This proportion represents the value of the input signal as a proportion of the maximum input signal (allowed by the ADC). For example, consider an 8-bit ADC that has a specified *full-scale input* of 1 V.

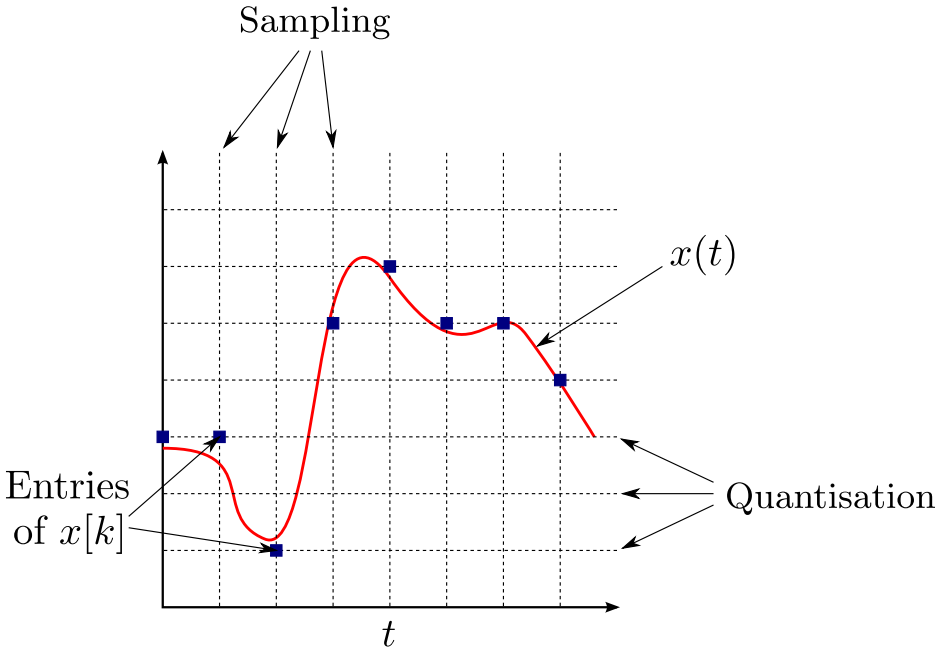

Figure 2.13: Converting $x(t)$ to $x[k]$, *sampling* and *quantisation*.

In this coding scheme, if this ADC produces an output of $x[1] = 10010111_2$ at sample time t_1, this corresponds to $x[1] = 151$. As the maximum output of this ADC is $x_{\max} = 256$, our test output corresponds to an input of

$$x(t_0) = \frac{x[1]}{x_{\max}} \cdot 1\,\mathrm{V} \approx 0.590\,\mathrm{V}.$$

Key performance indicators

As ADCs are mixed-signal circuits, with a substantial part of the functionality achieved in the analog domain, noise contamination does occur. As predictable from Figure 2.13, the process of quantisation results in a loss of fidelity. Whereas $x(t)$ can vary on a continuous scale, $x[n]$ is discrete in amplitude, meaning that a closest match to $x(t)$ at a particular time can be chosen, but will include some error. This error is known as *quantisation noise* and can be thought of as an additional noise source in the ADC. Quantisation noise is discussed in detail in Chapter 4.

The total noise in the output signal can be compared to the actual value of the output signal, resulting in a key performance indicator known as the *signal-to-noise ratio (SNR)*. This dimensionless quantity is generally reported in logarithmic units (dB).

Sampling the signal in time may result in a further loss of fidelity. The Nyquist-Shannon sampling theorem [51] states that as long as the sampling

frequency f_s is at least twice the highest frequency component of the sampled signal, no information is lost during sampling.

Hence, an ADC sampled at f_s has a *bandwidth* associated with it. We consider *bandwidth* another key performance indicator of ADCs. For *Nyquist-rate ADCs*, this bandwidth is

$$B = \frac{1}{2} f_s. \tag{2.13}$$

Undersampling refers to the process of sampling a bandlimited signal spreading from kB to $(k+1)B$, by sampling at $f_s = 2B$, and is sometimes employed in radio astronomy. For any dish antenna, the "Ruze-factor", which is related to surface roughness of the antenna, determines the highest frequency receivable [**?**], and as the specification for the initial SKA antennas calls for a maximum receivable frequency of 20 GHz [**?**], this is indeed a possibility. However, also in this case, (2.13) remains valid.

ADC *sensitivity* refers to the size of the digitisable signals. It is often specified as the minimum difference in signal levels that can be resolved on the ADC, also known as the *least significant bit (LSB)*. In this text, sensitivity and the LSB are closely related. The 8-bit ADC we considered above has a full-scale input of 1 V and can distinguish between $2^8 = 256$ different levels. Consequently, it has an LSB of

$$\text{LSB} = \frac{1\,\text{V}}{256} \approx 3.91\,\text{mV}.$$

In our opinion, this summarises the three most important performance indicators of ADCs for our application: *signal-to-noise ratio*, *bandwidth* and *sensitivity*. There are many other key performance indicators of ADCs, such as *linearity* and *effective number of bits (ENOB)*. These will be treated contextually where important.

Note that we consider the signal-to-noise ratio as equivalent to the *dynamic range* of the ADC for our purposes.

2.4.2 Superconducting ADCs

Many properties of superconductors suggest that they are extraordinarily well-suited to building analog-to-digital converters:

1. The macroscopic quantum-mechanical effect of flux quantisation, discussed above in Section 2.3.1, enables us to quantise signals at a fundamental, quantum-accurate level. This is particularly evident in oversampling ADCs which rely on quantum-accurate feedback in the form of SFQ pulses (see Section 2.4.2).

2. Low characteristic times of Josephson junctions promise high bandwidths. Clock signals may be distributed at high speeds in RSFQ electronics.

3. Cryogenic cooling required by superconductors ensures a low noise temperature of the environment. The noise temperature is nonetheless of great importance, as becomes evident in Chapters 4 and 5.

4. Low-loss systems ensure good to excellent sensitivity.

The current state of superconducting ADCs was succinctly summarised recently [26]. We distinguish between *Nyquist-rate ADCs* and *Oversampling ADCs*. The advantages of superconducting ADCs are summarised in Figure 2.14. Both Nyquist-rate and Oversampling ADCs are discussed in greater detail next.

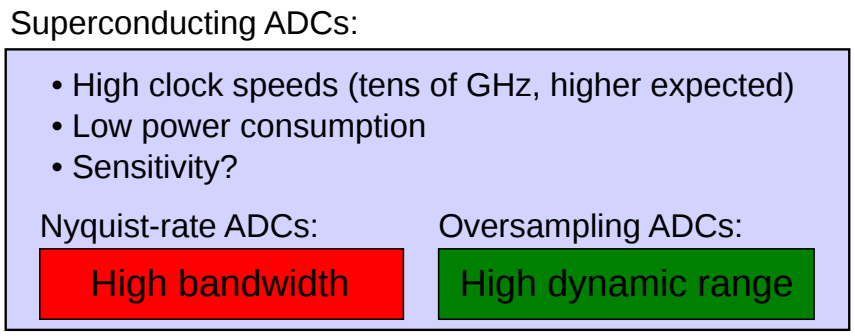

Figure 2.14: Advantages of superconducting ADCs. The low-loss characteristics of superconducting circuits are set off against the finite size of the magnetic flux quantum, which makes achievable ADC sensitivity an open question.

Nyquist-rate ADCs

In a Nyquist-rate ADC, the sampling frequency f_s is also the frequency with which the output is produced. One output bitvector (entry in $x[k]$) is produced each time the signal is sampled, ensuring the applicability of (2.13). Usually Nyquist-rate ADCs consist of parallel comparators to which different proportions of the signal are applied. This is achieved with particular elegance in the case of the superconducting flash ADC.

The superconducting flash ADC was first proposed by Zappe of IBM in 1975 [52]. Relying on the Josephson junction in a SQUID configurations, the proposed ADC (later demonstrated [53, 54]) relied on the periodic response of SQUIDs to enable synthesis of n bits from only n comparators, whereas conventional technologies required $2^n - 1$ comparators instead. This remains the foundation of modern superconducting flash ADCs, extensively investigated in Chapter 4.

Because flash ADCs are governed by (2.13), they have high bandwidths. This is of particular value for us, as the SKA will target wide frequency bands. If we continue to assume that the SKA will investigate frequencies up to $f_{\max} \sim$

10 GHz, then conventional technologies (based on semiconductor ADCs) will probably not be able to digitise the entire band-of-interest instantaneously, at least not without undersirable constraints (see Section 7.2.3). While this may be acceptable, an instantaneous bandwidth covering the band-of-interest would certainly be desirable.

As covered before, superconducting logic circuits can operate at clock frequencies f_{clk} of tens of GHz, with 100s of GHz demonstrated [19]. (These 100s of GHz frequencies are projected to become consistently achievable with future lithographical improvements [55, 56, 57].)

The parallel comparators employed in superconducting flash ADCs must be matched to operate efficiently. Fabrication tolerances make this problematic in any integrated circuit technology. For this reason, the number of reliable bits achieved by flash ADCs is normally fairly low. Flash ADCs with 4-6 bits were demonstrated [58, 26, 59, 60, 11, 61], but were not investigated extensively (presumably due to the great experimental effort required for this) and seemed to operate at the edge of their capabilities. Nonetheless, flash ADCs have proven fairly successful with many applications limited not by the ADC itself but the periphery instead.

The signal-to-noise ratio of flash ADCs can roughly be tied to the number of bits, where each bit contributes about 6 dB to the SNR. Hence, a 4-bit ADC should have an SNR (and dynamic range) of approximately 24 dB.

Oversampling ADCs

As opposed to Nyquist-rate ADCs, oversampling ADCs do not produce an output each time the signal is sampled. Instead, two clocks are employed, the sampling clock f_s and the read clock f_r. The *oversampling ratio* is given by

$$R_o = \frac{f_s}{f_r}.$$

Often, only one comparator is employed in an oversampling ADC. The 1-bit digitised version of the signal, say $x_1[k]$, is then post-processed by a method called *noise shaping*. The principle of noise shaping is as follows: sampling at f_s enables recovery of the signal bandwidth up to $\frac{1}{2}f_s$, but, if the band of interest B is smaller than $\frac{1}{2}f_s$, then one can make sure that the majority of the quantisation noise is pushed into those parts of the spectrum $\bar{B}$ that are not of interest.

A popular such ADC is the delta-sigma ADC, depicted schematically in Figure 2.15. The digitised signal $x[k]$ is subtracted from the input signal, with the quantiser essentially digitising the difference rather than the full signal level. In this way, together with the integrator, noise-shaping is achieved (more detailed explanations are abundant [50, 62]). Compared to Nyquist-rate ADCs, this is achieved at the expense of bandwidth.

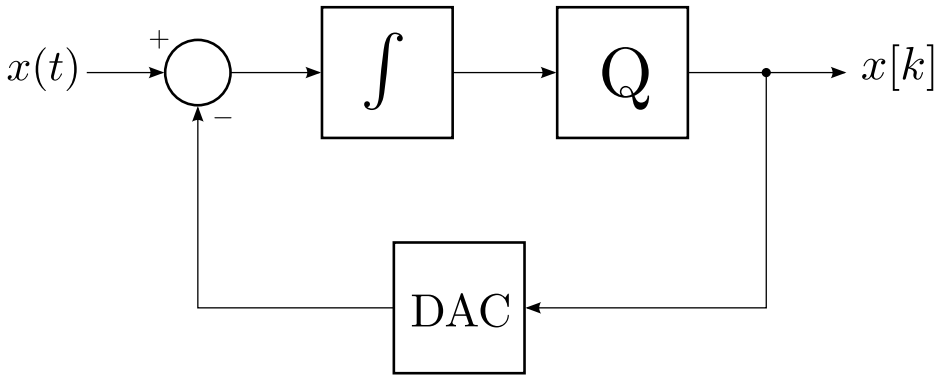

Figure 2.15: Sigma-delta architecture: a popular oversampling ADC.

The superconducting delta-sigma modulator was invented by Przybysz in 1992 [63]. Such a modulator is depicted in Figure 2.16. The integrator is achieved by an (omitted) LR circuit. Quantisation is performed by the decision-making pair (DMP), which, when clocked, outputs a 1-bit quantised version of the current applied.

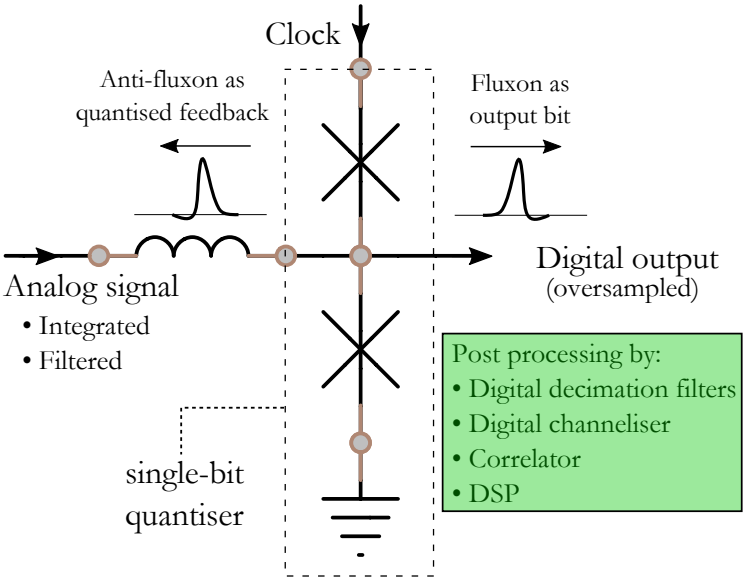

Figure 2.16: Superconducting sigma-delta modulator employing flux quantisation.

Remarkable is how the digital-to-analog conversion and feedback happens implicitly. If an SFQ pulse (fluxon) is produced at the output when the clock pulse arrives, this is interpreted as a digital `1`. Simultaneously, an SFQ pulse of opposite polarity (anti-fluxon) is deposited into the LR loop, subtracting from the input signal. As this pulse is quantised, an exact amount of "signal"

is subtracted at this time. This suggests how quantum-accurate quantisation can be performed by superconducting analog-to-digital converters.

Note that the delta-sigma modulator from Figure 2.16 is a low-pass modulator. Replacing the integrator with a resonator enables band-pass modulation [64], which enables placing the band of interest at other locations on the spectrum.

Post-processing (for example a decimation filter) is required to extract a representation of the signal level from the single-bit output stream. The output frequency, f_r is twice the bandwidth recoverable B. The signal-to-noise ratio achievable by a general sigma-delta ADC is given by [58]

$$\mathrm{SNR} = \frac{3}{2}\left(\frac{2m+1}{\pi^{2m}}\right)(2^n-1)^2 R_o^{2m+1}, \tag{2.14}$$

where n is the number of quantiser bits and m is the order of the modulator. No straightforward way to realise superconducting modulators with $m > 1$ for band-pass ADCs has been reported to our knowledge (although a second-order low-pass modulator was reported by Herr [65]). The problem is that higher-order modulators generally require gain, which is achieved only non-trivially with superconductors, and is particularly problematic at high speeds. We are also not aware of any oversampling superconducting modulators that make use of multi-bit quantisers. We thus fix $m = 1$ and $n = 1$ for currently practical superconducting oversampling band-pass ADCs.

2.4.3 Discussion of ADC architectures

For this discussion, we take the flash ADC architecture and delta-sigma ADC architecture as representative of the two categories of superconducting ADCs examined, Nyquist-rate and oversampling ADCs respectively. We believe that fundamentally achievable key performance parameters do not differ significantly between different architectures of the same category.

From (2.14), reduced by fixation of n and m, we conclude

$$\mathrm{SNR} = \frac{3}{2}\left(\frac{3}{\pi^2}\right)R_o^3. \tag{2.15}$$

We now consider an example system. Say we require 10 MHz around the hydrogen line at $f_c = 1.42$ GHz. The read frequency is thus fixed at $f_r = 20$ MHz. If we can achieve a superconducting sigma-delta modulator which can be sampled at $f_s = 20.48$ GHz, which seems reasonable, the oversampling ratio is $R_o = 1024$.

According to (2.15), this enables SNR ≈ 85.1 dB, which is very respectable. If we require a bandwidth of 100 MHz instead, the SNR drops to around 55 dB.

A 4-bit flash ADC clocked at $f_s = 20.48$ GHz only achieves SNR ≈ 24 dB, but the bandwidth recoverable is $f_s/2 = 10.24$ GHz, potentially covering the entire band of interest of the SKA.

It is clear: a choice between flash and delta-sigma ADC imposes a trade-off between SNR (dynamic range) and bandwidth.

Since one of the documented objectives of the SKA is the investigation of the red-shift of celestial bodies, we reexamine this here in some detail. Redshift is typically determined by observing the hydrogen line, which nominally resides at 1.42 GHz. The redshift z is given by

$$z = \frac{f_{\text{emit}} - f_{\text{obsv}}}{f_{\text{obsv}}}, \tag{2.16}$$

where f_{emit} is the frequency emitted and f_{obsv} the frequency observed. When considering the hydrogen line, $f_{\text{emit}} = 1.42\,\text{GHz}$. The highest red-shifts measured so far are at $z = 8.6$, derived by employing gamma-ray bursts (rather than the hydrogen line). Observing similar redshifts by virtue of the hydrogen line (gamma-ray bursts are not observable with the SKA) would require $f_{\text{obsv}} \approx 0.148\,\text{GHz}$, an order of magnitude away from the nominal hydrogen line.

Emissions from the postulated "Population III stars" [66] are yet to be observed, but expected to have redshifts of $z > 20$. This would require $f_{\text{obsv}} \approx 0.067\,\text{GHz}$.

Naturally the exact value of z is unknown before measuring it, ideally one would define a band-of-interest that includes all possible ranges and a comfortable fringe at each boundary (say from 0.05 to 1.5 GHz) and observe it in its entirety simultaneously.

For a bandwidth of $\sim 1.5\,\text{GHz}$, (2.15) dictates that the SNR of a readily achievable superconducting delta-sigma ADC falls to $\sim 30\,\text{dB}$, still above that of the 4-bit flash ADC, but no longer with such a convincing margin. In fact, a 5-bit flash ADC would also achieve 30 dB, but allow recovery of a greater bandwidth.

Actually observing frequencies below several 100 MHz is somewhat impractical with the full dish array planned for the SKA, owing to requirements on the dimensions of the antenna, terrestrial interference, etc. Thus, the wideband antennas of the SKA dish array will probably not be particularly useful in recognising highly redshifted bodies by the position of their hydrogen line. However, (2.16) is true for all emitted radiation, hence the argument remains valid.

Doubtlessly, observing spectral lines of celestial bodies at narrow bandwidths of several 10s of MHz is valuable, perhaps once the redshift has been established by a coarser receiver. However, considering the objective of the SKA (observe bodies with various redshifts), tunability of a narrowband receiver seems a requirement.

We expect that a delta-sigma ADC will only be useful to the SKA if its high dynamic range can be taken advantage of. This requires that its bandwidth be narrow, probably around $10 - 100$ MHz, which in turn requires tunability

of the centre frequency (in our current opinion). We are unaware of robust methods available for tuning superconducting delta-sigma converters.

Thus, we choose to investigate further the superconducting flash ADC architecture and deem the superconducting delta-sigma architecture of lesser potential advantage to the SKA, at least for now. Changes in fabrication capabilities (clock speeds), hybrid technologies (gain), SKA requirements or other aspects will require re-examination of this issue.

2.4.4 Receiver proposition

A potential conventional receiver of an SKA antenna is schematically described in Figure 2.17a. Note that this receiver is really an extension of the receiver architecture of the KAT-7 telescope [33], but corresponds to general radio-astronomy receivers. The analog signal is received by the antenna, combined by the orthomode transducer (OMT) and passes immediately to the low-noise amplifier (LNA). After the LNA, the signal is transported in analog form to a digitisation station. An RF-over-fibre link may be employed to ensure that the signal does not degrade significantly, especially as the digitisation may happen far away from the antenna (for KAT-7, digitisation for all dishes is combined in a single location).

The signal is then mixed down to baseband, where it can be digitised. However, the band of interest normally exceeds the capabilities of available ADCs, thus the signal is channelised. Variation of the local oscillator frequency enables a section of the band of interest to be selected, which is then passed along to the ADC for digitisation. The bandwidth of the chunk of signal digitised at one time by the ADC is known as the instantaneous bandwidth. KAT-7 has an instantaneous bandwidth of $B_i = 256\,\text{MHz}$, for example. The channels can later be combined to provide a spectral image of the entire band of interest. This process can largely be automated by digital control circuitry such as a DSP.

Naturally, the possible conventional receiver shown in Figure 2.17a is a simplification of any actual receiver. A second mixer may be needed to mix the channel from intermediate frequency (IF) to baseband. On the other hand, the image frequency may be digitised directly, with further demodulation occurring digitally. Carefully choosing the image frequency should make this fairly straightforward even for contemporary technology.

Mixing here is an analog process, the details of which depend on many factors. It is by definition a non-linear operation. Furthermore, it relies on a local oscillator frequency, which must be carefully synthesised to combat phenomena such as drift or ageing. Also, it does of course expose the signal to further noise contamination.

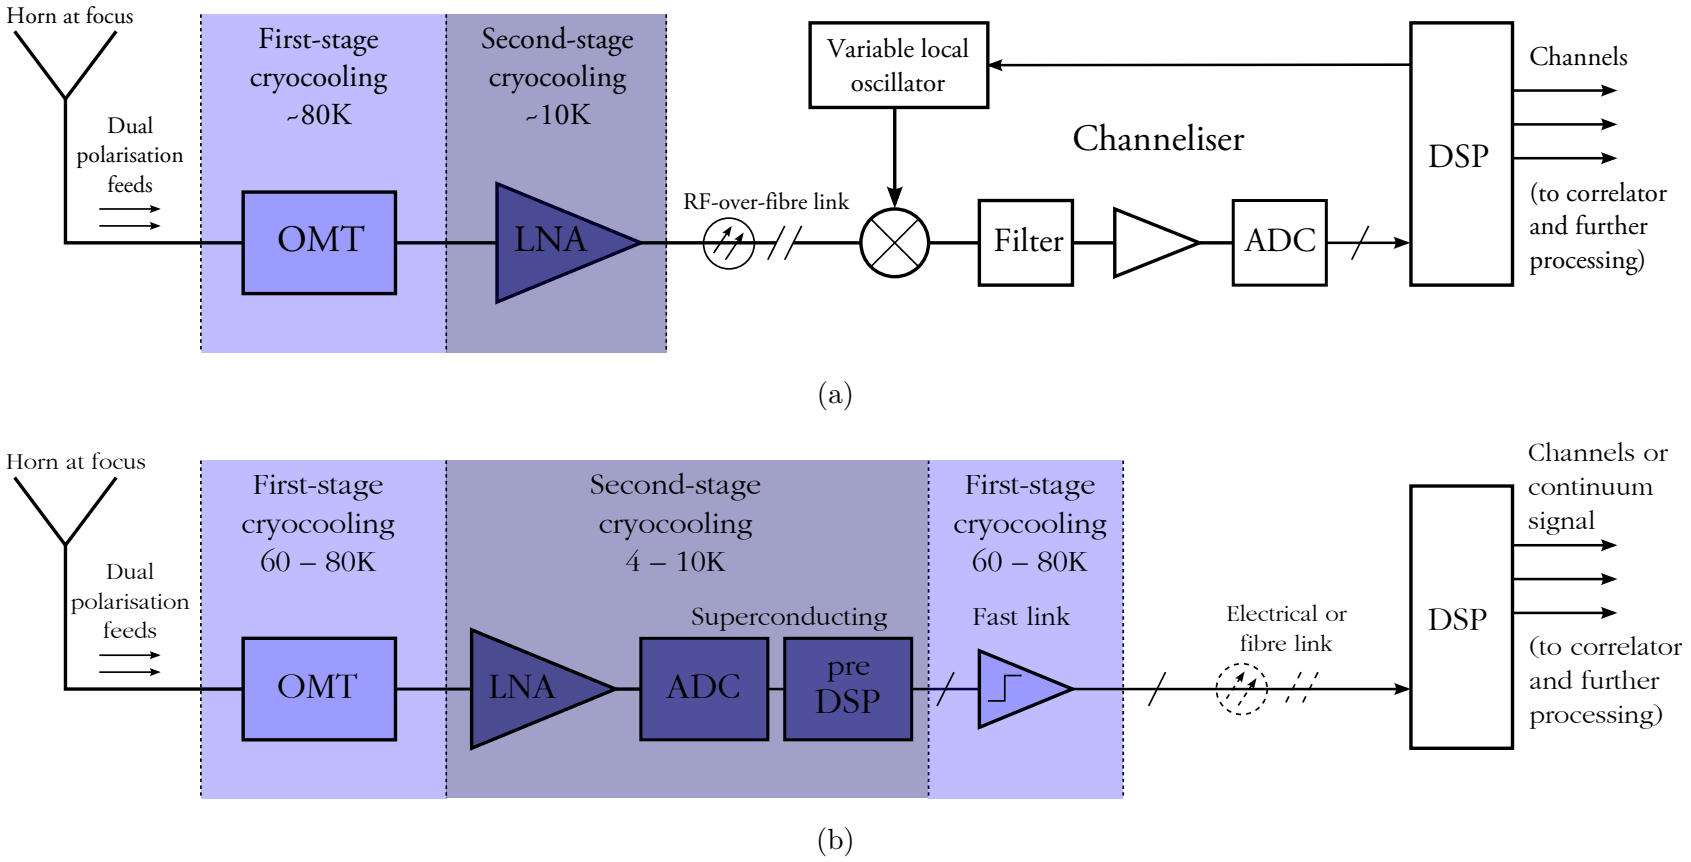

Figure 2.17: Possible conventional receiver (a) and proposed superconducting receiver (b) of an SKA antenna.

A superconducting receiver for the SKA is proposed in Figure 2.17b. Immediately noted is the absence of mixers. The superconducting ADC is placed directly behind the LNA, still in the cryocooled environment. A device that we call the "pre DSP" conditions the data from the ADC in such a manner that it can be processed by the digital backend. This device will almost certainly have to be superconducting due to the speeds required. An amplifier may be required to bridge the thermal barrier for the output bus. Once across the thermal barrier, the data may be transported to a different location, if required, where they may be processed. One aspect of such processing is the combination of the data with data from other receivers.

Due to its large bandwidth, employing a superconducting flash ADC enables digitisation of the entire band of interest instantaneously, obviating the need for a channelising setup. This shortens the receiver chain dramatically, as the ADC may be placed directly behind the LNA instead of far away in a different physical location, and fewer components are needed in the receiver chain.

The following advantages ensue:

1. Instantaneous digitisation of the entire band of interest ensures that transient events in the band (such as extra-terrestrial communication) are more likely to be captured.

2. A dramatically shortened receiver chain reduces noise contamination of the received signal. This is especially so as the signal is in the analog domain only in a controlled, cryocooled environment.

3. A far simpler receiver structure should reduce the maintenance effort substantially, ensuring greater uptime for science.

We also expect the following disadvantages:

1. The cryocoolers envisioned for the SKA may not be cold enough or powerful enough to support the superconducting hardware required.

2. Dynamic range achieved by flash ADCs is comparatively low, limited by thermal noise and/or fabrication tolerances. This could offset the gain obtained in terms of noise contamination.

3. The high data rates associated with the high instantaneous bandwidth may pose a problem for the digital backend, although this is not a disadvantage unique to a *superconducting* wideband digitiser.

Superconductors require cryocooling. The SKA receiver front-end will almost certainly be cryocooled to minimise thermal noise. The addition of superconducting circuitry thus seems plausible. Mature superconductor fabrication processes all rely on Niobium, which has a critical temperature of $T_c = 9\,\mathrm{K}$.

To ensure stability of circuit parameters and add some immunity to temperature variations, these integrated circuits are usually operated at a nominal temperature of $T_n = 4.2\,\mathrm{K}$.

It is currently unclear whether independent specifications for the SKA will require cryocooling to T_n. Specifications are hard to come by at this time, but figures of $\sim 10\,\mathrm{K}$ have been tossed about [30]. This may not be cold enough to enable the trivial addition of superconductors. It is possible that the additional effort of lowering the nominal temperature to $T_n = 4.2\,\mathrm{K}$ may be justifiable, especially in 10 years when cryocooling technology has advanced. There is also the possibility of employing a different superconductor, such as niobium-nitride, which has a critical temperature of $T_c = 17\,\mathrm{K}$. Currently such processes have been largely abandoned in favour of niobium, but advances could again be made in the next 10 years.

Power dissipation is problematic in cryogenic environments. The "cooling penalty" incurred by cryocooling may multiply power requirements at $T_n = 4.2\,\mathrm{K}$ by up to four orders of magnitude. Fortunately, the power dissipation of superconducting electronics is supremely low. We expect that the LNA will dissipate far more power than the superconducting components. The subject of power dissipation of superconducting circuits is discussed at length in Chapter 6.

A low dynamic range does not categorically prohibit the detection of faint signals, as this can be achieved by a good LNA and a sensitive ADC. However, bright celestial bodies may be positioned in close proximity to faint celestial bodies. In this case, high dynamic range may be required to discriminate between the two bodies. However, we expect that the angular resolution of the telescope is the dominant factor here, which is determined by parameters such as telescope baseline length and is independent of the dynamic range. This argument is briefly explained in Figure 2.18.

2.4.5 Dealing with the data

The downside of a high instantaneous bandwidth is a high data rate that must be processed (instantaneously). An n-bit flash ADC clocked at f_{clk}, for an instantaneous bandwidth of $B_i = \frac{1}{2} f_{\mathrm{clk}}$, has a data rate of

$$D_r = n f_{\mathrm{clk}} = 2 n B_i.$$

Continuing our assumption of an instantaneous bandwidth of $B_i = 10\,\mathrm{GHz}$, and assuming a 4-bit ADC, this requires a data rate per receiver of

$$D_r = 80\ \mathrm{Gbit\,s^{-1}}. \tag{2.17}$$

Clearly, this is a large figure, more than can be handled by any single link based on conventional technology (that we know of). Note that, for the entire Square Kilometre Array, this figure should be multiplied by the number

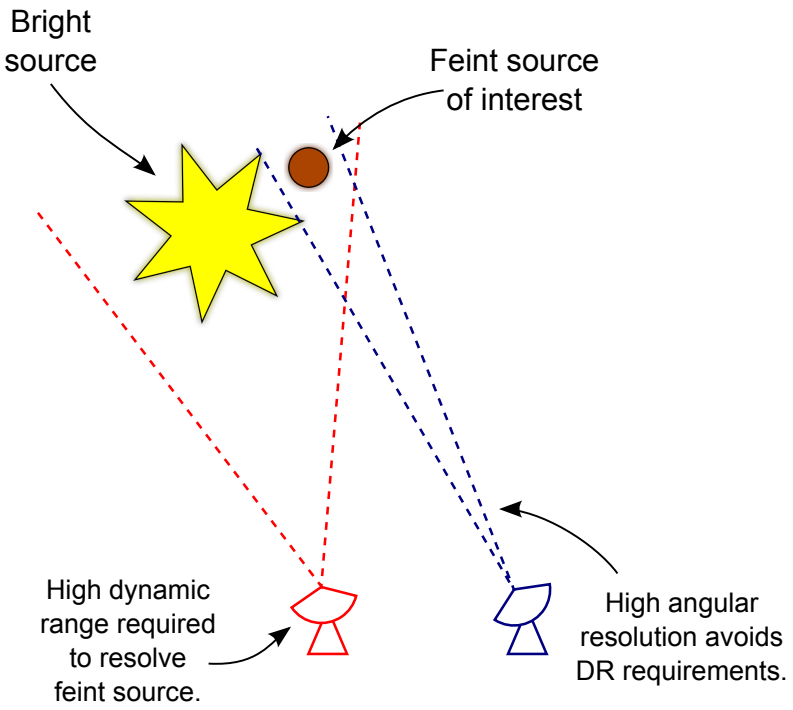

Figure 2.18: Angular resolution may obviate need for high dynamic range.

of concurrent receivers. Assuming that one receiver per dish is in operation at a time, and assuming 3 000 to 4 000 dishes, this corresponds to a total instrument-wide data rate of

$$(D_r)_{tot} \sim 300 \text{ Tbit s}^{-1}. \tag{2.18}$$

To put this into perspective, in 2011 the total worldwide bandwidth of submarine cables was $< 70\,\text{Tbit s}^{-1}$ [67], about a quarter of our projected raw telescope data output. Clearly, processing is a special challenge.

2.4.6 Technology challenges

No single link of conventional technology is likely to be able to process 80 Gbit s^{-1} today. Of course, this may have changed by the time the SKA is actually built. Nonetheless, we assume for now that we have to reduce the output data rate of the superconducting circuitry. The problem here lies in the disparity between clock speeds in superconducting and conventional technologies. This leaves us with two options:

1. Perform a substantial part of the data processing in superconductor technology, or

2. spread the data in such a way that they can be accepted and processed by conventional technology.

Processing of the raw data involves, simply put, correlation of different vectors and the FFT (the order is technically irrelevant). Two important performance indicators for correlators are: the integer length m of the vectors that can be correlated and the size n in bits of the entries in the vectors.

A correlator has been successfully implemented in superconducting technology before [68]. While this correlator benefited from the high f_{clk} of superconductor technology, the limited integration achievable with current processes confined m and n to 512 and 1 respectively. A vector of length 512 at $f_{\text{clk}} = 20\,\text{GHz}$ corresponds to a mere 25.6 ns of continuous observation, and that only at 1-bit resolution. This seems woefully insufficient for the SKA.

To our knowledge, a complete FFT processor has not been achieved in superconducting technology. The FFT is a highly parallelisable algorithm, which makes it eminently suitable for processes that can guarantee high integration, which current superconducting processes cannot, at least not when compared to semiconductor processes.

Although this would certainly be the preferable solution in terms of clock speeds and power dissipation, superconductor processes are probably, in their current form, not suitable for processing of the raw data of the SKA. Since both required steps are highly parallelisable, conventional technologies may scale sufficiently well to accommodate these requirements. However, the power dissipation of such a system may be prohibitive, a topic explored extensively in Chapter 6.

In the light of this, we will pursue the objective of spreading the raw data output from the ADC over several links, lowering the data rate of each link sufficiently to enable acquisition of the data by conventional technologies. This is accomplished by employing deserialisers. Our substantial work towards developing a suitable deserialiser is detailed in Chapter 6.

The state of superconducting technologies will certainly improve by 2024. Ideally, achievable integration will approach that of semiconductor processes and be sufficient for processing the raw data directly. Key in this endeavour will be the power dissipation of large circuits. We make a substantial leap in the right direction in Chapter 6.

2.5 Chapter summary and conclusion

The SKA is a radio telescope to be built in South Africa and Australia by around 2024. It will be the world's largest radio telescope and require technology that does not exist today. This chapter explored ways in which superconductors can help bridge some of the gaps in the technology.

1. Spread over some 3 000 dishes, the SKA will have an effective area exceeding $1\,\text{km}^2$.

2. The band of interest to the SKA starts at several 100 MHz, but may approach and exceed 10 GHz. Current receiver architectures see the band channelised and mixed-down before digitisation. This adds noise and means that transient events may be missed if they occur out of the current instantaneous band.

3. A receiver is characterised primarily by the key performance indicators of achievable dynamic range, recoverable bandwidth and sensitivity. Other characteristics exist, but we consider these less important here.

4. Nyquist-rate ADCs produce an output during every clock period, enabling recovery of frequency components up to half the clock frequency. Oversampling ADCs sample faster than the output sample rate, employing noise-shaping to improve dynamic range at the expense of bandwidth.

5. Superconducting electronics, such as RSFQ logic, are characterised by high clock speeds and low power requirements. Both of these properties would benefit the SKA due to its wide band of interest and demanding data processing requirements.

6. We expect a superconducting narrow-band receiver to be advantageous over conventional receivers only if it is sufficiently tunable. This is not easily achieved considering the nature of superconductors and the comparative "crudeness" of fabrication technologies.

7. The superconducting flash ADC architecture seems eminently suitable for the SKA. We expect its high instantaneous bandwidth to outrank in importance its comparatively low dynamic range.

8. The logistical challenges surrounding superconducting systems are profound, but overlap with some of the major requirements already characterising SKA specifications (such as cryocooled receivers).

Conclusion: Superconducting flash ADCs may provide an elegant method of digitising the entire band of interest instantaneously. RSFQ-based circuits should be capable of providing readout and elementary data-processing. The power advantage of SFQ electronics over semiconductor-based electronics may bring significant advantages at the post-processing stage.

2.6 Chapter research output

1. A conference poster will be presented at the *European Conference on Applied Superconductivity* in Genova, Italy in September 2013. The authors include the present author, his supervisors and a German collaborator.

M.H. Volkmann, C.J. Fourie, D. De Villiers, D.B. Davidson, W.J. Perold and T. Ortlepp, "Improving the Square-Kilometre Array with Superconductor Electronics,"

Abstract:
With 3000 dishes, totaling an effective area of 10^6 square metres, the Square-Kilometre Array (SKA) will be the world's largest radio-telescope. It will address frequencies from tens of MHz to tens of GHz, while exhibiting unparalleled sensitivity and myriad other superlative performance indicators. The frequency range of interest places the SKA signals squarely into the domain of recent and predicted superconducting analog-to-digital converters. Processing and imaging the raw data output of the SKA will require exascale high-performance computing power, another sector in which superconductor electronics are expected to underpin disruptive technologies in the immediate future.

Much of the technology and infrastructure required to achieve the touted specifications is not available today. This and the timescale of the SKA (with an estimated completion date of approximately 2022), suggests the ability for an applicable yet premature technology such as single flux quantum logic to evolve suitably to be of value. In previous work, we have concluded that superconducting electronics are an excellent fit for the SKA. In this work, we show how employing superconducting analog-to-digital converters will result in greater bandwidth/sensitivity, a prime performance indicator of the SKA. Furthermore, we outline the benefits of a superconducting high-performance computing backend.

2. A conference poster was presented at the *Superconductivity Centennial Conference* in Den Haag, Netherlands (2011).

M.H. Volkmann, C.J. Fourie, W.J. Perold, D.B. Davidson, D.I.L. De Villiers and H.C. Reader, Poster: "Superconducting circuits in the Square Kilometre Array radio telescope," *Superconductivity Centennial Conference*, Den Haag, Netherlands, September 2011. Available http://staff.ee.sun.ac.za/cjfourie/pdfs/Poster_MV_et_al_SKA.pdf (last accessed 28 August 2012).

Abstract:
The world's largest radio telescope, the Square Kilometre Array (SKA), will have up to 3,000 receiver dishes covering an effective one square kilometre at completion. The final site will be either in South Africa or Australia. Precursor telescopes are being built at both target locations. With an expected raw data output in high Tbps terms, the SKA will require telecommunication infrastructure and computing power that does not exist today, but should be available by 2020. For MeerKAT,

a South African precursor telescope to the SKA, receiver bands from 0.58 – 3 GHz and 8 – 14.5 GHz are of particular interest. Current plans see signals in these bands amplified, channelized, mixed down and then digitised. An all-digital front-end could simplify receiver structure and improve its performance. Semiconductor ADCs continue to make great progress and will likely find applications in the SKA, but superconductor ADCs benefit from higher clock speeds and quantum accurate digitisation. We investigate the current state- of-the-art in superconductor ADCs and software-defined radio (SDR) systems, attempt to define an SDR "dream system" for the SKA, and evaluate the practicality of designing all-digital superconductive front-ends for the SKA by 2020. The challenges for any candidate front-end system include: ultra-low radio frequency interference, manageable power consumption and overall costing which fits into a touted present day budget of 1000 Euros per installed system square metre.

Acknowledgements:
The authors thank EMSS Antennas, Stellenbosch, South Africa for detailed discussions on the cryogenic frontend of KAT-7 and MeerKAT (including pictures) and the design challenges facing SKA front-ends, and Hypres, Inc., New York for very helpful discussions on superconductive ADCs and field-tested systems.

This poster introduces the system laid out in this chapter, attempting to summarise its advantages and the challenges expected for development and implementation. The authors include the present author and his supervisor and co-supervisors, as well as other contributors from the Department of Electronic Engineering of the University of Stellenbosch.

Chapter 3

Addressing the Software Situation

3.1 Introduction

As in many engineering endeavours, the software toolset available to the designer is of great importance in shaping the landscape of superconducting electronics. Many engineering problems lend themselves well to computational analysis, such as simulation, automatic verification, and the analysis of sensitivity to tolerances. This is particularly true for many electronic engineering problems.

Many engineering practices, even in electronic engineering, precede the availability of mainstream personal computers. Mature fields are characterised by established design paradigms that are accepted and practised by the majority of the professionals active in the field. This provides a well-defined environment in which software can be developed and deployed with maximum effect.

For newer fields, such as that of superconducting microelectronics, established design paradigms do not yet exist. Although many design practices are simply borrowed from the similar paradigms established for more conventional microelectronics, fundamental differences between the two technologies make a one-to-one mapping impossible.

Some notable differences between superconducting and conventional microelectronics, and their effects, are summarised in Table 3.1.

Nonetheless, software for superconducting circuits is available. Much of it was custom-written for small audiences (such as research groups). Many tools rely on modified versions of the tools employed in conventional electronics. Some commercial efforts have been made in the past and go on today (such as WRSpice [69] and 3D MLSI [70]). Recently, NioCAD, a start-up located in Stellenbosch, has made significant efforts at developing a catch-all superconducting design software suite. This suite, NioPulse, has the purpose of

Table 3.1: Differences between superconducting and conventional microelectronics and their effects on software toolsets.

Difference	Effect
Dynamics governing Josephson junctions are fundamentally different from those governing transistors.	Modelling Josephson circuits with conventional SPICEs is computationally expensive.
Topologies of superconducting circuits are inductor-heavy, with capacitances playing a minor role.	The focus during layout verification is on the inductors, which is not the case for conventional circuits.
Superconductors have special electromagnetic properties, they do not act like, say, perfect electrical conductors in all situations.	Employing conventional EM-solvers results requires assumptions that do not always apply.
The frequencies (or bandwidths) targeted by superconducting circuits are often higher.	Software optimised for more conventional frequencies may not work (efficiently).

combining all aspects of superconducting circuit design in one package.

The design process of superconducting circuits, as well as the software currently employed, is explained in the next section. The limits of current approaches, as well as how we address these limits, are detailed in the remainder of the chapter.

3.2 Design process

Figure 3.1 outlines the typical process of superconducting design.

Moving from conceptual design to a working chip generally involves several iterations of certain steps. For example, although an initial simulation may show that a circuit is working as expected, further analysis may reveal that margins are narrow, or that the circuit is too sensitive to parasitics (affecting the ability to reliably manufacture the conceived circuit). Furthermore, extracting a circuit may reveal issues or problems not considered during original layout. This could require re-layout or even render a circuit unrealisable.

Naturally, experience and foresight should minimise such undesirable revelations late in the design process. Ideally, properties of great importance (such as physical circuit feasibility) should be understood early in the design phase.

3.2.1 High-level design

By *high-level design* we refer to the conceptualisation of a (presumably) realisable system with the purpose of solving a problem. Part of this process is

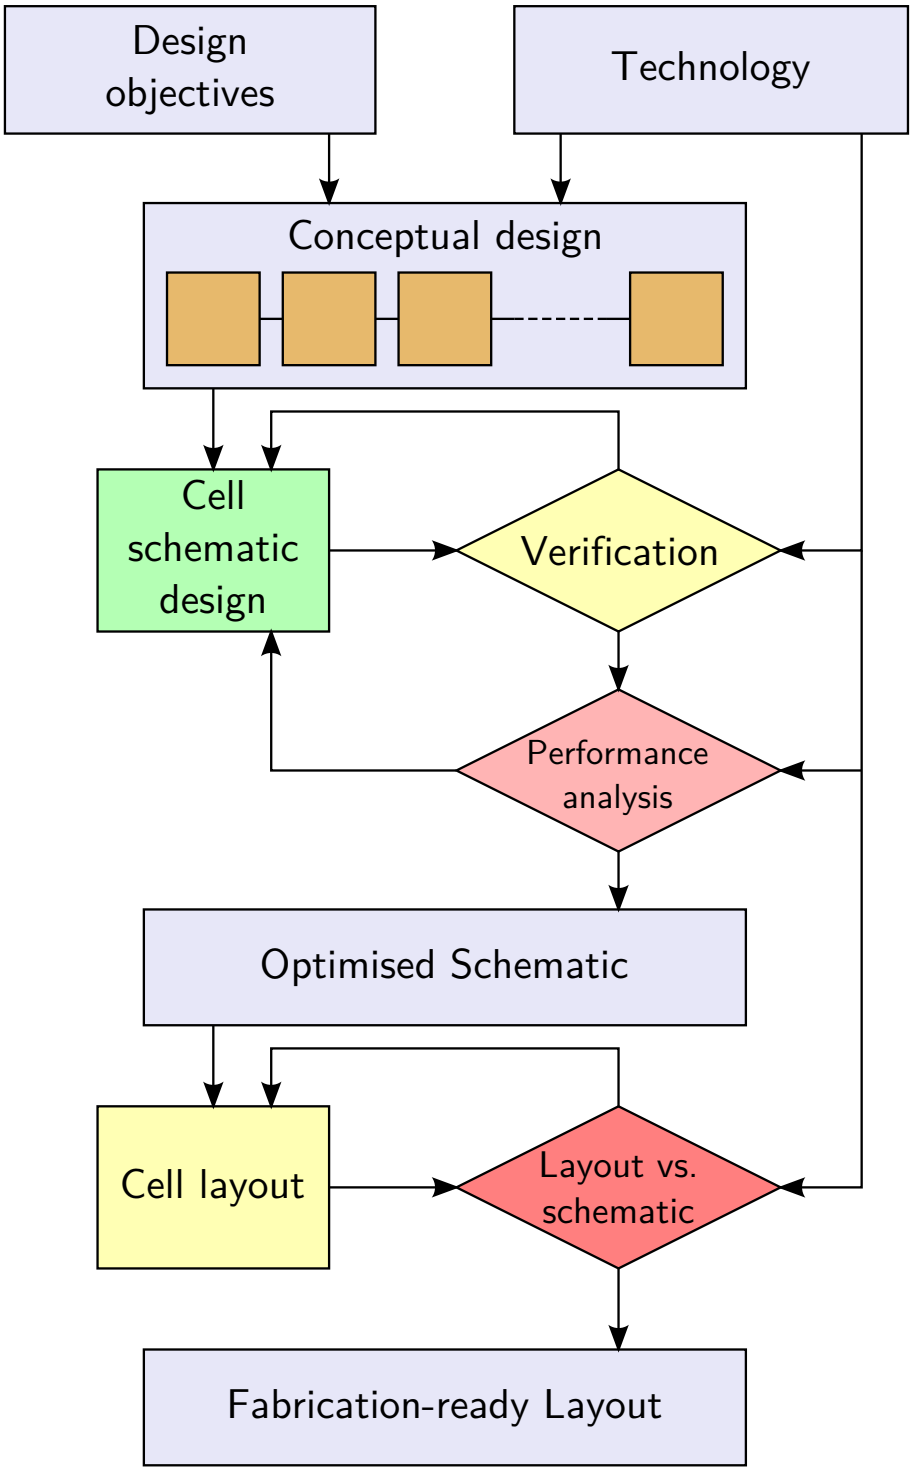

Figure 3.1: Typical design process for superconducting circuits

defining the problem that needs to be solved. High-level design should encompass the following actions (which cannot necessarily be neatly separated):

1. Identify the problem and clear objectives that need to be met to solve it.
2. Translate the objectives to numerical goals. These could include minimum values for key performance indicators or, ideally, a tailored figure of merit that clearly summarise the objectives of the problem.
3. Separate, if possible, design parameters from performance indicators, as depicted in Figure 3.2. Quantify, where possible, the relationship between key design parameters and performance indicators. This requires at least a rough model of the proposed design.

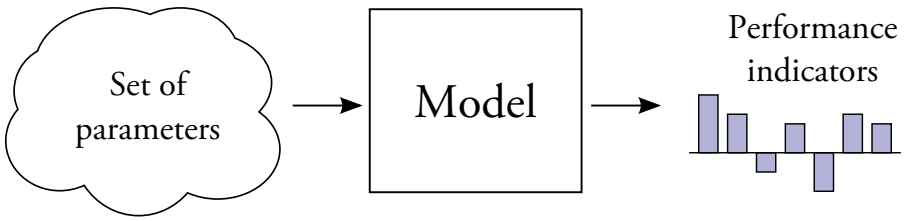

Figure 3.2: Ideal high-level design deliverables.

3.2.2 Divide and conquer

Devising a solution for a simple system is generally fairly straightforward. We define a *simple system* as one that can be designed and simulated without requiring other undesigned building blocks. A single circuit (with few or no subcells) usually fits this bill. From here onwards, we refer to such a system as a *cell*.

It is not prudent to design a complex system without first dividing it into cells. A cell corresponds roughly to a block in a system block diagram. Without such a block diagram, the innards of the system may not be completely understood. Correctly anticipating system response and best practices for design is difficult if the system is not broken down into easily comprehensible components. Breaking the system into cells decreases the likelihood of design errors, repetition or erroneous assumptions. Furthermore, complex, undivided systems are more difficult to analyse and the computational effort associated with such analyses is high.

Thus it is desirable to first divide a proposed design into cells, each of which can be designed and analysed separately. If the cells perform as expected, the system function can be inferred and, if possible, simulated.

Generally, a block diagram can be employed to show the relationship between cells and data in a system. The function of a system should be clearly evident from its block diagram. To achieve clarity, it may be necessary to group cells in block diagrams into subsystems. Each subsystem can be represented by a further block diagram if required.

Note that a block diagram is useful even if it seems more complex than the resulting circuit. Consider the block diagram from Figure 2.15 and compare it to the circuit achieving it in Figure 2.16. The circuit seems simpler than the block diagram, yet its operation follows more fluidly from the block diagram.

3.2.3 Circuit design

Circuit design is a well-understood field in electronic engineering. There are several components to circuit design, such as *circuit topology* and *parameter space*, as well as *response* and *verification*.

Circuit topology refers to the arrangement of circuit elements (often lumped elements, distributed elements, or subcircuits) and the connections between them. For example, the parallel combination of an inductor and a capacitor describes a resonator circuit, whereas the parallel combination of grounded Josephson junctions separated by inductors may refer to a discrete Josephson transmission line (JTL).

However, the topology alone does not govern the response of the circuit. Each circuit element generally has at least one property that must be set. Inductors have the *inductance* property, whereas Josephson junctions generally have more than one property (such as *critical current* and *parasitic junction capacitance*). Combining all settable properties of a circuit yields the parameter set. Combining the possible ranges of each parameter in the parameter set make up the parameter space of the circuit.

The response of a circuit refers to its behaviour in the presence of the environment (which is made up of other circuits and signals).

Circuits can generally be attributed an *input* and an *output*. The input refers to signals applied to the circuit that are independent of the circuit or its response. In contrast, the output refers to a subset of the response of the circuit. Generally the output of a circuit can be used as inputs to other circuits. It may also represent the system output.

Verification

Verification results from comparing the response of a circuit from a corresponding input to the circuit to the expected response. For example, an OR-gate

requires an output pulse for a pulse at any input, whereas an AND-gate requires an output pulse only if a pulse is applied at all inputs.

Circuit *simulation* in the presence of an input and an expected output template can be employed to achieve verification of the circuit operation. Circuit simulation relies on circuit topology, input, as well as the point in parameter space which the circuit occupies. This process is summarised in Figure 3.3.

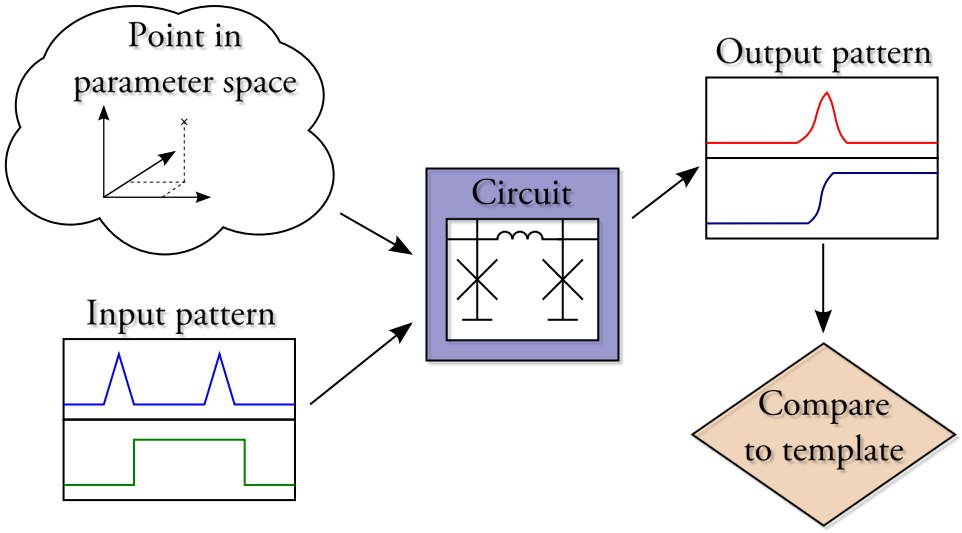

Figure 3.3: Circuit verification as part of circuit design.

Circuit verification can be employed for a plethora of valuable characterisations. Successful verification of a circuit does not generally alone ensure that it can be realised, as fabrication- and other tolerances may cause parameters to deviate from their nominal values. We define the *operational region* of a circuit as the subspace of the circuit's parameter space on which the circuit behaves as desired.

We assume that all circuit parameters are linearly independent, resulting in a dimensionality of the parameter space equal to the cardinality of the parameter set. This assumption could be challenged, with potentially rewarding outcomes, but we do not pursue this here. We believe that making this assumption covers the most general case. Hence, establishing the operational region using the verification tactics described above requires the evaluation of the circuit at suitably spaced discrete points in parameter space.

The DFF circuit from Figure 3.4 is characterised by the indicated topology as well as parameters

$$\mathbf{P} = \{I_{c1}, I_{c2}, L_1\}$$

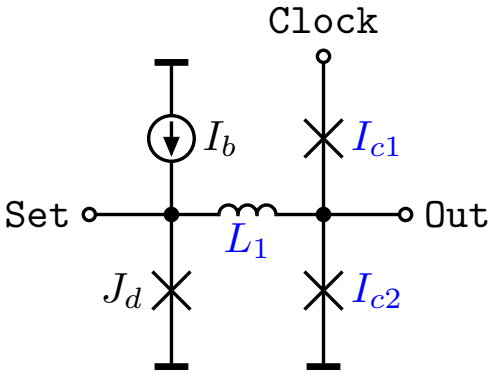

Figure 3.4: DFF circuit topology and parameters.

(we assume that other circuit properties are fixed). Our parameter space thus has dimensionality

$$D = |\mathbf{P}| = 3.$$

We can easily define an input and corresponding response that suggests correct operation of the circuit. Establishing the operational region over the full dimensionality of the parameter space requires, at N_p points per parameter, a total of $N_S = N_p^D$ simulations.

Clearly, N_S grows exponentially with $|\mathbf{P}|$, and polynomially of order D with N_p. As the correctness of the resulting operational region relies on a sufficiently high N_p, and worthwile circuits have $D > 1$, even $D \gg 1$, the computational effort required to confidently establish the region of operation is generally prohibitive.

Margin analysis

As a concession, we employ the *margin analysis* to establish a reasonable estimate of the operational region. In a margin analysis, each parameter p_i is swept while each other parameter $p_{j\neq i}$ is kept at its respective nominal value. The circuit is verified at N_p points in the sweep range, yielding the range(s) of the parameter at which the circuit behaves correctly. Whereas this can really only be formally interpreted as the operational region of the parameter *sub*space $\mathbf{P}_i$, characterised by p_i, this still provides useful information. In this case, the computational effort is lower, as only $N_S = N_p D$ simulations are required.

As a further concession, it is often assumed that the operational region of $\mathbf{P}_i$ is continuous and that the circuit works correctly at the nominal value of p_i (a reasonable assumption). This means that a binary search algorithm [71] can be employed to find the operational region of $\mathbf{P}_i$, lowering the computational effort further, as this requires merely $N_S \sim D \log N_p$ simulations.

In this case, the boundaries of the operational region of $\mathbf{P}_i$ are known as the *margins* of p_i, which are generally expressed as a percentage away from the nominal value v_i. For example, the margins [-30%,30%] correspond to an operational region of $[0.7v_i,1.3v_i]$. Such margins over the parameter set are generally depicted as a bar chart, as shown in Figure 3.5.

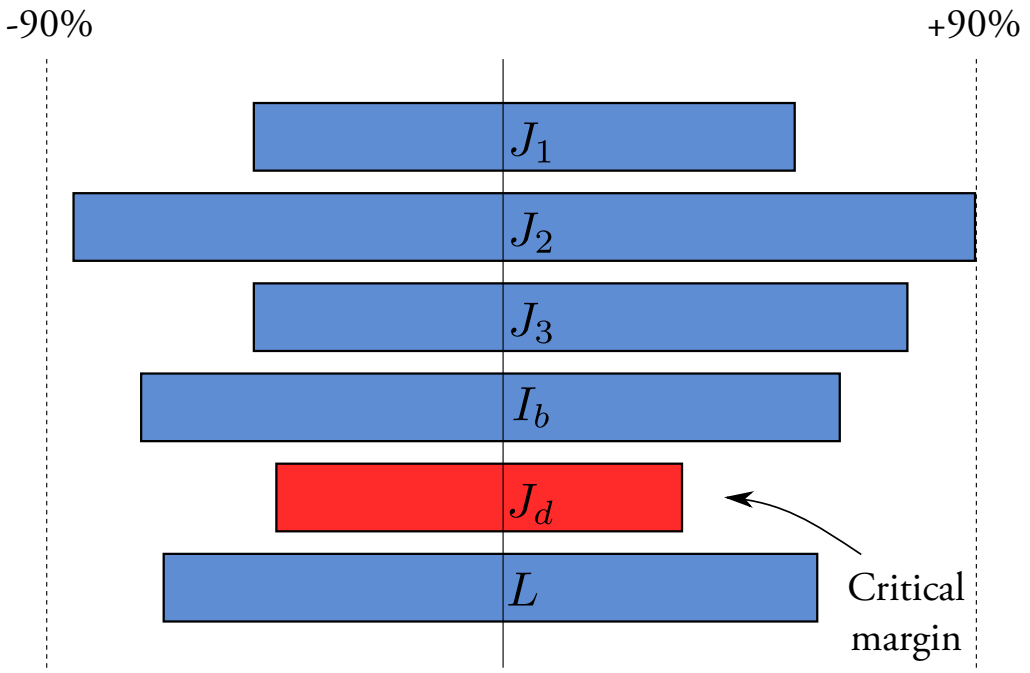

Figure 3.5: Results of a typical margin analysis. The critical margin is the lowest margin of the investigated circuit, and is often a good indication of where work on the circuit could be most rewarding.

Yield analysis

Another common analysis, the *yield analysis* returns an estimate of the probability of a circuit functioning correctly after fabrication. Naturally this requires some data from the fabrication process, such as tolerances of line widths, metal- and insulator thickness, critical current density and other quantities.

For a yield analysis, each parameter p_i is considered a random variable P_i. While, strictly, different probability density functions (PDFs) apply in practice, we find most natural the Gaussian PDF (motivated by the Central Limit Theorem [72] and the large number of factors characterising fabrication). Each P_i has as expected value the nominal value of p_i:

$$E[P_i] = v_i \quad \forall i.$$

Design rules often specify the tolerances that can be expected for certain kinds of parameters. For example, the design rules for the IPHT RSFQ1F process specify a maximum global spread of ±20% for the mean Josephson critical current density, and a maximum spread of ±10% for stripline inductors [13]. The on-wafer spreads for these parameters are ±15% and ±5 − −6% respectively. For the random variables P_i, the standard deviation, σ_{P_i}, is generally chosen at $\frac{1}{3}$ of the maximum expected spread, with some treatment to recognise both global and local variations. This ensures that, in 99.7% of cases, the drawn value of p_i will lie within the tolerances specified in the design rules.

The yield of a circuit is calculated by running N_S verifications, each of which relying on a parameter population drawn from the random variables. The ratio of successful verifications to N_S is the *statistical yield* of the circuit. If N_S is high enough, the statistical yield approaches the actual yield of the circuit with high confidence. This assumes, of course, that the specified fabrication tolerances are accurate.

Optimisation

An important procedure in circuit design is *optimisation*. Key circuit characteristics such as critical margins and yields can be employed to quantify the robustness of a circuit. Such a quantity is usually known as circuit *fitness*. Optimisation refers to the procedure employed to maximise circuit fitness.

Normally, only the parameter space of a circuit is considered during circuit optimisation (the assumption being that the chosen topology is most fit for purpose). A popular method of optimising RSFQ circuits is that of "iterative re-centering" [73]. A margin analysis is run on a circuit to identify the critical margins. Generally, critical margins are the narrowest margins found for any parameter in the parameter space. After the critical margins $[p_{i_{\min}}, p_{i_{\max}}]$ are found, a new nominal value for p_i is chosen as the centre of the range: $v_i = (p_{i_{\min}} + p_{i_{\max}})/2$. The margin analysis is now repeated, and a new critical margin, if any, identified.

If the critical margin is found for the same parameter p_i, the next critical margin may be considered. After iteratively centering the nominals of parameters in this manner, the margins and statistical yield of the circuit should increase somewhat, moving the circuit towards its optimal configuration.

This procedure can be automated, which has been demonstrated by proven heuristic algorithms such as Cowboy [74]. Iterative re-centring is not the only method of maximising circuit fitness. The centre-of-gravity method has been successfully employed in the past [75]. A good summary of modern automatic optimisation procedures has been published by Ortmann et al [76], but fails to identify a practical supreme optimisation method.

For now, as we deal with small circuits, we favour the manual procedure of iterative re-centering. Although requiring significant human input over time, it

fortifies insight into circuit operation, enabling better selection of parameters to be optimised.

3.2.4 Physical design

After a circuit has been optimised, it must generally be *laid out* for fabrication. Layout always occurs for a specific *process*, which refers to the fabrication procedure and materials involved. Before submission for fabrication, designers must verify that their layouts resemble the desired circuit. *Physical design* is an iterative procedure that relies on layout *extraction*, which maps layouts back to circuits.

Design rules

From a designer's point of view, processes are characterised by *design rules*, which are normally published by the foundry that operates the process. Such design rules generally specify constraints on the layouts that can be fabricated reliably, such as minimum dimensions of features, minimum spacing between features, and minimum overlap distances. It is the designer's responsibility to keep to the design rules, although many foundries run a design-rule check on submitted layouts and inform the customer of violations before fabrication.

Most superconducting processes of interest to us contain several layers. Features may be placed in different layers at the designer's discretion (within the design rules). Layers can roughly be categorised as conducting or insulating, with insulating layers separating conducting layers from each other. Vias, or holes in insulating layers, may be employed to connect conducting features in different layers to each other.

Junction and resistor design

The "tri-layer package" is a special layer that enables reliable fabrication of Josephson junctions. Consisting of a thin insulating layer sandwiched between two superconducting layers, superconducting charge carriers may tunnel across the insulating layer in the manner described in Section 2.3.2. The area of the trilayer package exposed to the electrodes connecting to it, A, defines the area of the Josephson junction achieved.

Process tri-layer packages have associated with them a critical current density J_c. The critical current of a laid-out Josephson junction is

$$I_c = AJ_c.$$

In fact, J_c is often used to identify the process. For example, Hypres Inc., a commercial foundry in Elmsford, NY, distinguishes between its $30\,\mathrm{A/cm^2}$ and $4.5\,\mathrm{kA/cm^2}$ processes (among others). The numerical designators refer to the value of J_c in both cases. (Currently we are not aware of any process

employing more than one tri-layer, although this would be very desirable in terms of parasitics of certain commonly-employed structures.)

The parasitic capacitance of a Josephson junction C_J was introduced in Section 2.3.2. Features in the tri-layer package consist of (super)conducting plates separated by an insulator (dielectric), making up a parallel plate capacitor. The capacitance C of a parallel-plate capacitor is given by

$$C = \frac{\epsilon A}{d}, \tag{3.1}$$

where ϵ is the permittivity of the dielectric, d is the distance between the plates and A is the area of the feature. As ϵ and d are fixed for a process, $C \propto A$. Process design rules generally specify a capacitance per unit area, C_0. Thus,

$$C_J = AC_0.$$

As junctions should normally be critically shunted (Section 2.3.2), external shunt resistors must be added in parallel with the junction, with the shunt resistance given by (5.15). The processes of interest to us all make available a resistive conducting layer for this purpose. A sheet resistance R_s is associated with this layer. The units of R_s are "Ω per square", or just Ω. A rectangular feature in the resistive layer has resistance

$$R = R_s \cdot \frac{\ell}{w},$$

with ℓ and w the length and width of the feature respectively, with ℓ taken as the dimension in the direction of the current.

Inductor design and extraction

Inductances are of critical importance in superconducting circuits (in particular SFQ circuits). However, the inductance of a feature is not as easily designed as its resistance. Although a quantity such as the "strip-line sheet inductance" is sometimes given in the process design rules [12], this is at best a rough guide that results from empirical measurements of structures laid out in a controlled environment. It is difficult to attribute an inductance to a single feature, as "inductance" is really a characteristic of loops.

As circuit theory is usually employed for superconducting cell design, the quantities of importance here are the *partial inductances* that best describe the mapping of a layout to the lumped-element representation of the circuit [**?**]. Here, *partial inductance* refers to that part of a loop represented by a lumped-element inductor in a circuit. "Extraction" refers to the process of determining such a mapping. Three inputs are required for the extraction process:

1. A *netlist* representing the inductor topology of the circuit to be achieved,

2. a layout designed to achieve the circuit, and

3. a *port-mapping* that binds key areas in the layout to corresponding points in the netlist.

A typical extractor builds an electromagnetic (EM) model of the layout and excites this model. The response of the EM model can then be analysed and mapped back to the netlist to determine the "best fit" of partial inductances that would elicit such a response. If the input netlist is a good representation of the input layout, the partial inductances calculated by the extractor can be a good to excellent indication of the circuit achieved. The quality of the fit is determined by many factors, such as the discretisation of the EM model and the assumptions made by the extractor.

Feedback and parasitics

Experienced designers are familiar with the processes in which they work and are able to achieve the correct circuit topology quickly. Nonetheless, extraction enables one to obtain a good to excellent indication of the point in parameter space that a layout occupies. This enables the designer to simulate the laid-out circuit and recognise its position in the operational region. Tweaking the layout and re-extracting allows the designer to position the layout at or near the desired optimum in parameter space.

Of importance here are the parasitics of a circuit. Grounding a junction, for example, generally incurs a *parasitic inductance* between the junction and ground. This usually depresses the parameter margins associated with the junction. If this parasitic inductance is not accounted for in the circuit topology, the extraction may fail or return inaccurate results.

3.2.5 Microwave design

Superconducting digital circuits may be clocked up to microwave and sub-microwave frequencies, up to tens of GHz in current processes (with 100s of GHz expected for future processes) [55, 56, 57]. This invites interesting possibilities for microwave structures that operate at those frequencies, such as resonators, transformers, transmission lines, filters, etc. The scale of the features in superconducting integrated circuits is such that these structures can be realised on-chip.

As superconducting processes generally have a groundplane and other layers in which structures can be placed, microstrip or stripline design practices may be employed to design such structures. Current processes, such as the standard Hypres $4.5\,\mathrm{kA/cm^2}$ process, provide for up to four layers, allowing for complicated structures such as multi-order filters. Future processes should have more layers, opening up further possibilities.

However, designers often resort to analytic or empirical methods alone for designing such structures. In conventional microwave or RF/HF-electronics, analytically designed structures are often verified by EM-simulations. This practice is rarely employed for superconducting designs. In the present author's opinion, this is due to a dearth of suitable EM-software for thin-film superconductors. Although full-wave EM-solvers that support superconductivity do exist, these are primarily designed for bulk superconductors, where the dimensions of the features exceed the supercurrent penetration depth λ.

3.3 Capabilities and limitations of available software

3.3.1 Spices

The workhorses of electronic circuit simulation are found in the form of SPICEs. A SPICE (Simulation Program with Integrated Circuit Emphasis) is a general-purpose analogue circuit simulator that allows prediction and verification of the response of electronic circuits [77]. SPICEs are invaluable in circuit design, a fact no less true in superconducting circuit design.

Most SPICEs take the form of an iterative differential equation solver. The input to a SPICE is a netlist (or SPICE deck) that contains descriptions of circuit elements (SPICE cards) and how they are connected to each other. The SPICE uses this netlist to construct a mathematical model of the circuit, which it goes on to solve in the presence of boundary values (signals). Many conventional SPICEs allow solving circuits in the time and frequency domain.

The output of a SPICE usually takes the form of vectors of electrical quantities, such as voltage or current, in all areas of the circuit netlist. The vectors may be compared to a time or frequency axis, which is also often output by the SPICE (but can otherwise be inferred). Some SPICEs employ dimensionless quantities instead, which may require some post-processing to convert to the desired units.

Non-superconducting SPICEs

Many conventional SPICEs exist. A SPICE usually provides the capability of simulating ideal elementary circuit elements, such as inductors or capacitors. These correspond well to superconducting inductors and capacitors. Different strategies are employed for modelling more complicated devices like transistors, operational amplifiers, etc. One approach is to synthesise the function of the required device as a subcircuit from the elementary circuit elements. Generally dependent sources are a requirement for this (to model phenomena like gain). Note that practically all SPICEs allow the hierarchical modelling of circuits, where subcircuits may be employed to essentially arbitrary depths. Some

SPICEs support "macro-models", which attempts to mathematically model the desired function without the need for complicated hierarchical subcircuits. This can significantly reduce the computational effort required.

Conventional SPICEs, usually meant for semiconductors, do not support the Josephson junction, a crucial element in superconducting circuit design. As described above, a Josephson junction can be "assembled" as a subcircuit model. This approach has the distinct disadvantage of substantial computational effort.

The reason for this lies in the typical models employed by conventional SPICEs. The choice of independent and dependent variables in the differential equations results in a model that is straight-forward to solve for conventional circuits. However, as the Josephson effect is specified in a quantity foreign to conventional SPICEs, the superconducting phase $\phi(t)$, it must first be rewritten.

JSIM and JSIM_n

The *Josephson simulator*, or *JSIM*, was specifically written by Fang and Van Duzer for simulation of Josephson junctions [78]. It avoids the slowdown experienced by conventional SPICEs by choosing its dependent and independent variables differently, resulting in lower-order differential equations for Josephson junctions. This means, of course, that inductors and capacitors must be modelled differently as well. This is really only a cosmetic change, as the same output quantities (currents and voltages) can be synthesised as in conventional SPICEs. In addition, the superconducting phase drop across a junction can be output.

The drawback of this approach is that other semiconductor devices like transistors or operational amplifiers are not readily simulated in JSIM. For circuits without semiconductors this is clearly not a problem. As our application as currently envisioned does not rely on semiconductors in the core circuitry, a simulator like JSIM is suitable for this project (note that JSIM still supports resistors). Conveniently, JSIM is also an open-source tool [78].

It should be noted here that JSIM supports only the RCSJ model of the Josephson junction (Section 2.3.2). We do not think that this is a problem for our application, but it does have the side-effect that reliable frequency-domain simulations are not possible. JSIM only supports transient analyses, whereas many conventional SPICEs allow other analyses (such as AC analyses). One might argue that the frequency response can be extracted from the time-domain data by employing algorithms like the Fast-Fourier-Transform (FFT), but this is cumbersome for two reasons: Firstly, the required frequency sweeps are somewhat difficult to apply with JSIM's time-domain sources and secondly, long simulations must be run to obtain enough time-domain data for a precise frequency-domain representation.

As will become apparent later, this drawback does affect our simulation procedure.

Satchell expanded JSIM to support modelling thermal noise [73, 43], yielding JSIM_n. This is our simulator of choice for this application.

PSCAN

PSCAN was developed specifically for superconducting circuits including Josephson junctions [79]. Many different simulations are supported, including transient analyses and margin analyses. Since its inception, many improvements were made, and a toolset has grown around it (sometimes termed *Rodeo*). An important feature of this toolset is the hierarchical circuit verification tool termed SFQHDL (Single flux quantum hardware description language), which reliably verifies circuits irrespective of the input pattern applied [74].

The simulator is still extensively used today, but serious limitations hamper its applicability. For example, the maximum number of junctions supported is only of the order of several hundred, a constraint that hampers even moderately-sized modern circuits. Furthermore, the simulation of thermal noise (essential for our application) is not supported natively.

WRSpice

WRSpice is a commercial product developed by Whiteley Research, Inc. in Sunnyvale, CA [69]. This is an advanced SPICE, which supports Josephson junctions as well as semiconductor devices. It also supports a form of parametric modelling and provides a scripting interface that allows for custom analyses (like margin and yield analyses). Also, at least compared to JSIM, more analyses and junction models are supported. Furthermore, WRSpice supports simulation of thermal noise.

We decided against using WRSpice primarily because it seems to be slower than JSIM. Furthermore, it is not open-source or free and the installation available to our department is outdated and largely defunct. Finally, we do not have a wealth of experience with WRSpice and already trust JSIM. One of the reasons for this is the comparatively patchy support for WRSpice in past and current versions of NioPulse.

3.3.2 Layout tools

Layout tools generally refer to software products that enable the layout of physical circuits in such a form that they can be shipped for fabrication. Generally the feature-set required for superconductor layout is equivalent to or a subset of that required for semiconductor layout, as the format that the foundry requires is similar for both technologies. However, sophisticated layout tools also provide verification capabilities. Many aspects of the verification process

(such as extraction) differ greatly between semiconductor and superconductor circuits.

LASI

The layout tool of choice for some superconductor design outfits (such as the IPHT in Jena, Germany [80]), *LAyout System for Individuals* (*LASI*) [81] is a capable and robust software product that enables layouts for any foundry that accepts the GDSII file format (we do not know of any that do not). This includes semiconductor foundries. We employ LASI 7 for superconductor layouts, as it interfaces well with our extractor of choice (InductEx) and provides other verification capabilities (such as design-rule checking and resistance extraction).

We lament the lack of shortcut keys (known to us) for many of LASI's key commands, but have remedied this situation by employing an AutoHotKey [82] script to map single-key taps to LASI commands (such as "m" for the *Mov* command).

LayoutEditor

The feature set of the commercial product *LayoutEditor* [83] suggests that it was designed with semiconducting circuits in mind. However, the recent addition of an interface to InductEx, a superconducting inductance extractor, suggests its suitability for adaptation to support superconducting circuits as well. Like LASI, it supports layout independent of the technology targeted (be it semiconducting or superconducting). We have not had time to extensively evaluate LayoutEditor in this context.

3.3.3 Extractors

An extractor is employed to extract circuit parameters from layouts and map these back to a schematic (generally fed in the form of a netlist). The selection of useful extractors is small. Others exist, but have limited applicability here [84].

LMeter

The open-source tool *LMeter*, written by Bunyk [85], has been the tool of choice for superconducting circuit designers since its release in the early 1990s. Useful for most RSFQ cells from that era, LMeter is still the most widely used extractor in our field today. Usable toolchains exist for using LMeter with Cadence, XIC and NioPulse, making it the extractor of choice for many designers. However, LMeter does not support groundplane holes or full-3D extraction, which may cause inaccuracies near vias. As groundplane holes are commonplace in modern cells, we refrain from using LMeter.

InductEx and FastHenry

Fourie's *InductEx* has been around since 2004, but has recently been rewritten and improved dramatically [84]. A full 3D extractor which uses a modified version of FastHenry [86, 87] as backend, InductEx is, in our opinion, the most accurate inductance extractor currently available for superconducting thin-film structures. Calibrations have been performed for contemporary fabrication processes [88]. Furthermore, it interfaces well with LASI 7 layouts, although some manual steps are required.

Full-3D support increases computation time compared to, say, LMeter, but the limitations regarding groundplane holes and vias are removed. This makes InductEx suitable for most SFQ-type cells, as the capacitance can generally be neglected at the frequencies and scales important to our applications. Note that FastHenry is a magnetoquasistatic solver not suitable for full-wave analysis (displacement currents are assumed to be zero). This limitation is generally of little consequence in digital superconducting circuit design, but renders it unsuitable for microwave designs.

Empirical results have confirmed the correctness of InductEx for our target processes [84, 89].

Recent versions of InductEx support resistance extraction [90], but since simple resistors are easily designed by straightforward analytical methods, and our needs do not extend beyond simple resistors, we do not make use of this feature. However, we do utilise the ability of InductEx to calculate the inductance of resistive features, such as later in Section 5.5.

ImpEx (NioPulse)

NioPulse bundles an extractor named *ImpEx*, which can extract inductances and resistances. It also employs a modified version of FastHenry [91] as backend solver. ImpEx is built into NioPulse, it is not accessible from outside NioPulse. Hence, the layout to be extracted as well as the circuit to which it should be mapped must be present in a NioPulse project before extraction can occur.

In our experience, ImpEx has been buggy and the results somewhat inconsistent. However, a lack of instructions in the user manual may mean that we are not using it in the intended way. As we are not using NioPulse for layout and verification, ImpEx is not useful for us in the context of this application.

3.3.4 Suites

The software suites considered here are meant to be used as a comprehensive solution to all aspects of superconducting circuit design, from a high-level conceptual or logical design to physical design and extraction. Several software suites are available for superconducting circuit design.

XIC Tools

XIC is the graphical user interface that combines WRSpice with physical circuit design capabilities. It is developed by Whiteley Research, Inc. [69] and commercially available. XIC provides schematic as well as physical layout capabilities. Schematics and layouts can be linked for straightforward storage and simulation.

XIC does not provide a particularly useful way of organising circuits into projects and libraries (this must be done by the designer by making use of the file system). We lament the absence of such a tool, as the complexity of circuits and number of revisions grows rapidly even for small projects. The layout-versus-schematic (LVS) implementation for superconducting circuits is incomplete, as XIC does not supply a good interface to an extractor. However, XIC provides an excellent "path-extraction" tool that allows for continuity checking in physical circuits.

XIC is not meant for superconductors alone and provides extensive support for semiconductor circuit design.

As JSIM_n is our simulator of choice, and XIC supports WRSpice only, we have decided against using XIC for most purposes. Furthermore, it is not freely available (like LASI) and does not yet interface well with our extractor of choice (InductEx). We do, however, employ XIC for verification of circuits, making liberal use of the path-extraction tool.

Cadence

The suite collectively referred to as *Cadence* is the most fully-featured software suite mentioned here. Cadence actually refers to the name of the developer, Cadence Design Systems, Inc. [92].

Cadence' tools are developed for semiconductor circuit design and, in off-the-shelf form, not particularly suitable for superconducting circuit design. However, several superconducting design outfits have contributed to Cadence to add the features required for superconducting circuit design. These contributions take the form of models in the simulation engine, scripts to interface with other tools (such as extractors) and scripts for custom analyses. As neither Cadence nor these contributions are available to us at this moment, we do not use it for our designs.

NioPulse

NioPulse, developed by Stellenbosch start-up NioCAD [93], is the only comprehensive software suite that we know of that specialises exclusively in the design of superconducting circuits. The NioPulse feature-set includes [91]:

- Design, analysis, optimisation, layout and verification of electronic circuits such as SFQ circuits,

- closed-cycle design capability for analog and digital circuits,
- reduced-complexity design cycle,
- a generic and adaptable environment,
- a unified graphical user interface,
- version control and library features, and
- a powerful scripting environment.

The objective of NioPulse is to facilitate the entire design process, from conceptualisation over circuit design, analysis and optimisation to layout, extraction and verification.

After an extensive evaluation of NioPulse (performed, in part, by the present author for Hypres, Inc.), we conclude that it is still of beta-quality and not yet suitable for the entire design process. We currently limit our use of NioPulse to schematic capture and rudimentary circuit simulations. Whereas the ideas behind NioPulse are sound and progressive, the software suite in its current form is unsuitable for certain purposes for the following reasons:

- Operations do not scale well; even moderately large circuits and layouts cause memory-usage to balloon, resulting in thrashing,
- the scripting environment is cumbersome, making the default feature set the only one readily accessible,
- while a critical part of the shipping-out process, exporting layouts to the GDSII file format is untested,
- parametric modelling is not supported in any useful form,
- the library system and version control system is not suitable for projects of moderate or large size,
- the design-rule checker is not operational, while the electrical rule checker is buggy, and
- extractions are unverified and difficult to reproduce.

Until the identified limitations are addressed, we will continue to use NioPulse for simulation purposes only. At the cut-off time for this text, the limitations have not been substantially addressed in the implementation.

3.4 Addressing the limitations

After reviewing and evaluating available software solutions, we identify the following serious limitations of available software tools:

1. *Parametric modelling*, a crucial feature for the design of circuits of even moderate complexity, is not supported for our simulator of choice, JSIM_n. Even NioPulse, a comprehensive software suite that interfaces with JSIM_n, does not support parametric modelling. JSIM_n only supports netlists with fixed absolute values, whereas NioPulse requires those values to be specified during schematic capture. Parametric modelling in some form is supported by WRSpice, but this simulator is substantially slower than JSIM_n. Limited support is available for PSCAN, but this simulator has many limitations.

2. Scaling up to *large circuits* is not readily achievable with available software. Whereas JSIM_n can certainly handle large circuits, it is difficult to verify these with available tools. In NioPulse, for example, simulations of large circuits causes memory overruns and crashes.

3. *Custom analyses* with JSIM_n are difficult to perform without extensive scripting, for which JSIM_n provides no native support. NioPulse does provide a scripting environment, but this is currently poorly documented and cumbersome to use. Also, the inability to scale remains.

4. *Full-wave EM-solvers* specifically minted towards superconducting thin-film structures are not available at present, to our knowledge.

While our choice of layout editor (LASI 7) requires a manual interface to our extractor of choice (InductEx series 4.1 and 4.2), we do not regard this as a serious limitation here. The lack of a suitable library and project management database tool is somewhat frustrating, but is unlikely to introduce errors due to the moderate scale of our research-scale projects. More important to us is the robustness of the layout editor and extractor, for which ample evidence supports LASI 7 and InductEx 4.

The remainder of this section details our attempts to address the listed limitations. We successfully address Items 1 to 3, whereas Item 4 will be the topic of further research by the present author and the group of which he is a member (outside the scope of this work).

3.4.1 Modifications to existing tools

Many software tools that perform tasks similar to our needs already exist from the semiconductor world. Tasks such as schematic capture, cell layout and plotting are not unique to the superconducting design paradigm. Hence, tools for these tasks can be borrowed from other, well-established design paradigms.

JSIM_n

JSIM_n is a console utility, meaning that it has no graphical user interface. The output of JSIM_n is usually written to the standard output, whereas the input is read from the standard input or a netlist file (which is a text file that adheres to a specific format substantially common to all SPICEs).

The output consists of a tabulated dataset where each line represents a point in simulation time. Each line consists of space-separated records, where the first record is a floating-point number stating the point in simulation time that the record pertains to. Each subsequent record is a floating-point number representing the level of an investigated variable (such as voltage, current or phase) at the specified simulation time.

As we anticipated the requirement of analysing SPICE simulation outputs in `MATLAB` and python, specifically pylab, we decided to modify JSIM_n to write the output data to a `MATLAB`-compatible file when passed the `-m <file.dat>` command-line switch (well-known libraries to read `MATLAB` files are available for python). The dataset is written to the file as described above, but a header-line is added to allow the post-processor (`MATLAB` or python) to gather the data into variable vectors when reading in the file.

After reading such a `MATLAB`-compatible file, the `MATLAB` or python workspace is populated with the vector `time` and a vector for each examined variable (such as vector `L1_I` for the current through inductor L_1). Each vector read from the file is the same length. This allows us to use, say, `MATLAB`'s `plot` function to plot the simulation traces.

An intended usage of our modified version of JSIM_n is the command

```
jsim_n -m output.dat netlist.cir,
```

after which the command `importfile output.dat` may be executed in `MATLAB`. The `importfile` function in `MATLAB` was synthesised from the `ImportData` function, and similar functions are available in python.

Schematic capture: gEDA for superconducting circuits

An open-source tool popular under Linux users, gEDA is a framework of software products to do with electronic design automation (EDA). Among other tools, these products include `gschem`, a schematic capture utility, as well as `gnetlist`, a netlister that takes `gschem` schematics as input and produces SPICE netlists after running a backend script.

The gEDA framework is very flexible. Although the framework was certainly not written for superconducting circuits, the open-source nature of the tools allowed for straightforward extension of the toolset to suit our needs. We made two significant changes to gEDA: support for JSIM_n and support for parametric modelling.

The backend script of interest distributed with gEDA is known as `spice-sdb`, which outputs SPICE netlists. We modified the script to ensure JSIM compatible output in the presence of a command-line flag passed to `gnetlist`. Some of the applicable code is listed in Listing B.1.

To support parametric modelling, we modified the framework to enable output of *netlist templates*. Such templates would then be processed by a separately written tool (`apply`, described in Section 3.4.2). Figure 3.6 shows our modifications at work: when values are specified for the circuit elements, they become fixed in the netlist (Figure 3.6a). However, when the values are omitted, a template variable is added instead (Figure 3.6b). In our current implementation, the template variable is taken as the name of the component in question. As an example, the code achieving the parametric template for an inductor is shown in Listing B.2.

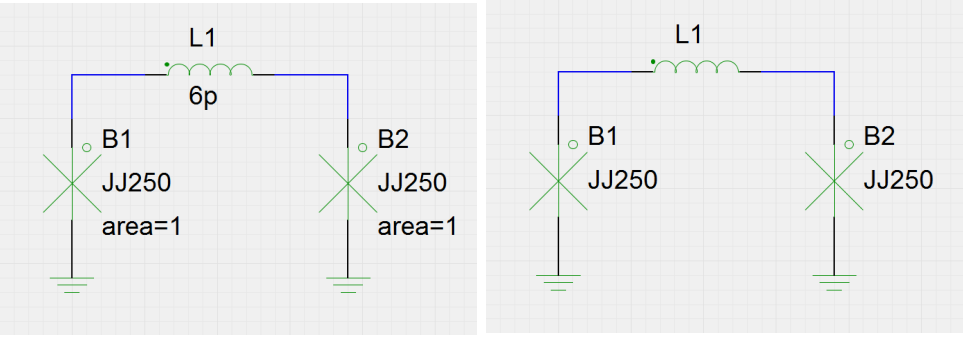

(a) (b)

Netlist produced:

```
B1 1 0 JJ250 area=1
B2 2 0 JJ250 area=1
L1 1 2 6p
```

Netlist produced:

```
B1 1 0 JJ250 area=@{B1}
B2 2 0 JJ250 area=@{B2}
L1 1 2 @{L1}
```

Figure 3.6: Parametric modelling support in gEDA: when values are specified, they are fixed in the netlist (a), whereas their omission sees them replaced by variables (b).

After these modifications, `gschem` became our schematic capture tool of choice.

FastHenry

As the backend to InductEx, FastHenry execution takes up the majority of processor clock cycles during a typical design process. To extract a typical layout, InductEx discretises the passed structure into current "filaments", which FastHenry employs to generate a complete list of currents (complex quantities) that map to each filament, given the excitation of certain terminals of the structure (by means of a port voltage).

For each "port" that the structure possesses, FastHenry calculates the solution to a large set of linear equations, for which it employs the well-understood *generalised minimal residual method* (GMRES) developed by Saad and Schultz in 1986 [94]. This process resembles the iterative solution of an equation of the form

$$Ax = b,$$

with $A \in \mathbb{C}^{m\times n}$, $x \in \mathbb{C}^n$ and $b \in \mathbb{C}^m$. Here, x is the desired solution (which relates to filament currents) and b is the relevant excitation (which relates to whatever port is excited). A derives from the geometry of the discretised structure passed to FastHenry.

The conventional sequence of operations that FastHenry performs is depicted in Figure 3.7, whereas it is explained in detail in Kamon's MS thesis [86]. From the supplied N-port geometry, the matrix A is constructed. Next, excitation vectors $b_1, b_2, \cdots, b_N$ are constructed, concluding the preparation process at time τ_{prep}.

FastHenry includes a number of built-in preconditioners which speed-up this process considerably for matrices A that meet certain conditions. Unfortunately, construction of the preconditioner does not scale well in time, and currently we do not employ the preconditioner for moderate to large structures and do not consider it further.

The bulk of the time is now spent on GMRES iterations. Sequentially, the solution x_i is found from A and b_i.

Fortunately it is possible from a mathematical/computational standpoint for each solution for x_i to run in parallel. We have modified the FastHenry code made available by Whiteley [87], which includes his modifications for enabling the analysis of superconducting networks, to enable this. We employed the popular OpenMP libary [95] to facilitate cross-platform parallelisation. Issues that needed solving were the conversion to 64-bit architecture on Windows, which was not supported but is required to address large amounts of memory, and special memory management for some data structures in the fast-multipole accelerator employed by FastHenry, which was not immediately suitable for parallelisation. Our `fasthenry` program takes an additional command-line flag, '`-j p`', where `p` is the number of threads to be employed when performing the solutions x_i. This enables the sequence of operations depicted in Figure 3.8.

To obtain an understanding of the benefits and drawbacks of our parallel FastHenry, we now make the following (optimistic) assumptions:

1. Each GMRES solution $x_i = \texttt{gmres}(A, b_i)$ requires the same resources in terms of memory and time, M_{sol} and τ_{sol}.

2. Overhead in switching thread contexts is negligible.

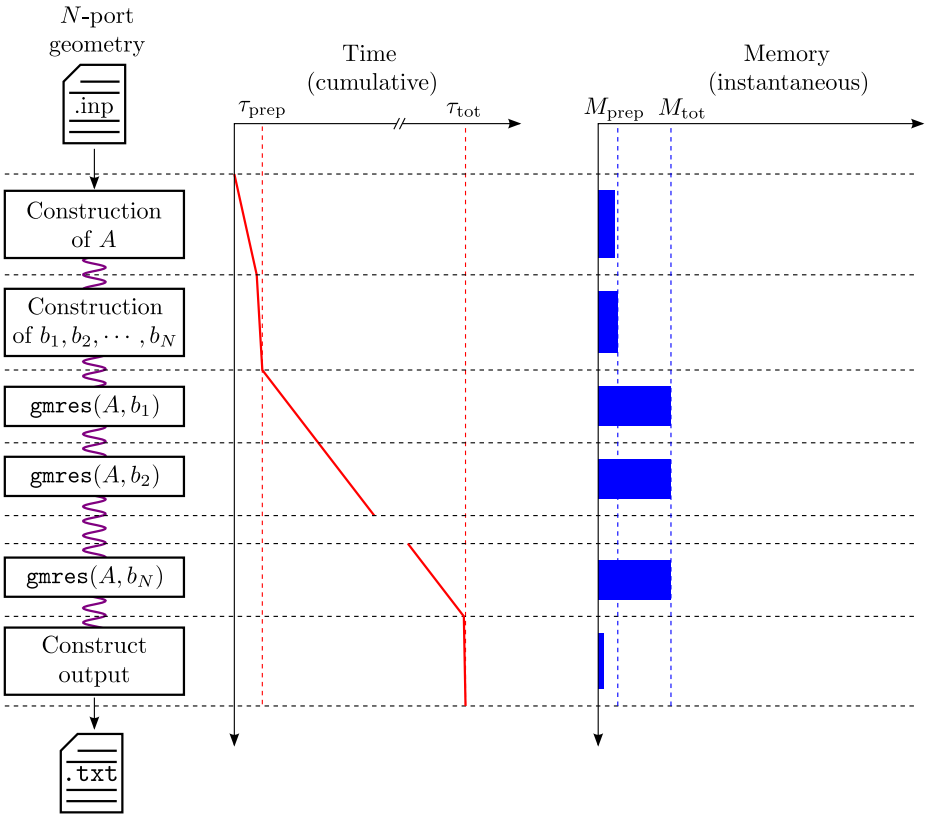

Figure 3.7: The conventional FastHenry process (without preconditioning).

3. Execution times and memory requirements for operations other than preparation and solution are negligible.

Suppose there are q cores in the system and p threads allocatable to FastHenry, we obtain the resource requirements reported in Table 3.2.

For the designer, it is generally most important that the extraction run finishes quickly, that is, τ_{tot} is minimised. As the GMRES runs were the only part of the algorithm that was parallelised, and there are N GMRES runs, the maximum conceivable speed-up of that component is N. On the other hand, any algorithm running on a q-core system can achieve a maximum conceivable speed-up of q (compared to a single-core system). Generally, for our designs, $N > q$, so that the limiting factor is the hardware, not the algorithm.

The best case, a speed-up of q, is revealed by Table 3.2 to be achieved when $p \geq q$ and p divides N (which requires $a = 0$). Since contemporary operating systems generally enable thread-counts far greater than any N conceivable for our purposes, this best case can always be achieved by choosing $p = N$.

This assumes sufficient memory, however. If the required memory M_{tot} exceeds available physical memory, frequent paging (thrashing) will occur, imposing a penalty on execution time that would almost certainly exceed the

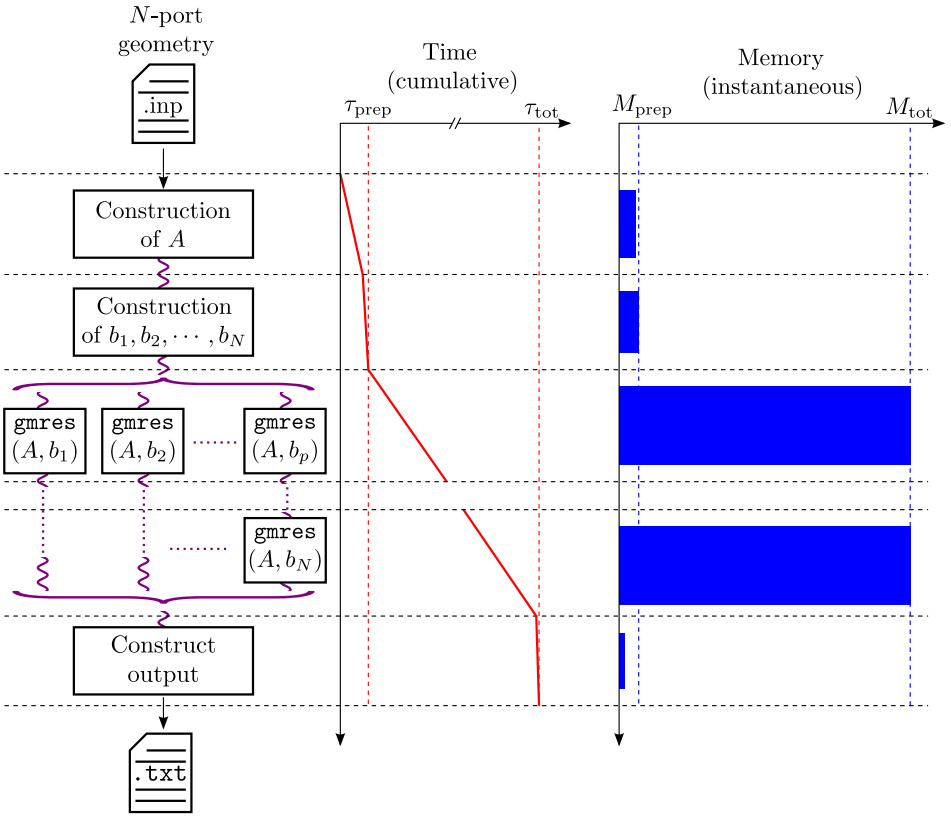

Figure 3.8: Our parallel FastHenry process (without preconditioning).

Table 3.2: Resources required for FastHenry execution.

		Time τ_{tot}	**Memory** M_{tot}
Conventional		$\tau_{\text{prep}} + \tau_{\text{sol}} N$	$M_{\text{prep}} + M_{\text{sol}}$
Memory	$p \leq q$	$\tau_{\text{prep}} + \tau_{\text{sol}} \left\lceil \frac{N}{p} \right\rceil$	$M_{\text{prep}} + pM_{\text{sol}}$
	$p > q$, $N = kp$[1]	$\tau_{\text{prep}} + \frac{\tau_{\text{sol}}}{q} N$	
	$p > q$, $N = kp + a$[2]	$\tau_{\text{prep}} + \frac{\tau_{\text{sol}}}{q} \big(kp + \max(a, q)\big)$	

[1] Here $k \in \mathbb{N}$, that is, p divides N.

[2] Here $k, a \in \mathbb{N}, a < p$, most general case.

```
$ time fasthenry -p off test.inp
<fasthenry output>
real    22m43.377s
user    22m37.125s
sys     0m1.652s
```

(a)

```
$ time fasthenry-par -p off test.inp -j 7
<fasthenry output>
real    6m36.427s
user    30m17.590s
sys     0m2.384s
```

(b)

Figure 3.9: FastHenry execution time measurement for the conventional (a) and parallel (b) output. M_{prep} and τ_{prep} are approximately 500 MiB and 80 s respectively. M_{sol} is approximately 750 MiB. Of importance to the designer is the `real` entry, which best describes elapsed wall-clock time.

benefit gained from parallelising the algorithm. Consequently, the maximal useful number of threads p is clearly bounded by the available physical memory.

Considering this bound, it may or may not be possible to satisfy the conditions for the best-case speed-up for a given design. If it is not possible, the greatest speed, at the lowest memory use possible for that speed, for a given N, is achieved by choosing the $p \geq q$ that minimises $N - a$. If p is bound to below q, p should be maximised.

To illustrate these constraints, example commands and corresponding outputs are recorded in Figure 3.9. The input geometry, `test.inp`, stems from an RSFQ cell and consists of approximately 26 000 filaments. Furthermore, $N = 7$ and the execution occurred on a 4-core machine (Core i7 series), so $q = 4$. Physical memory of 8 GiB is available, which suffices for the best-case scenario, hence we provide $p = 7$.

The achieved speed-up (in wall-clock time) is 3.44, which agrees well with our estimates. The speed-up of the solutions alone is 4.06, higher than the predicted maximum bound imposed by $q = 4$. This is due to hyperthreading, which seems to provide some benefit here that exceeds even the penalties imposed by the invalidity of the optimistic assumptions made earlier.

We do not see an impediment to parallelising the preparation and preconditioning process as well, but deem this out of scope of this work.

3.4.2 Purpose-written tools

The modifications described in the last subsection allow us to use gEDA, JSIM_n and `MATLAB` to design, simulate and evaluate superconducting circuits. To enable parametric modelling, circuit verification and common analyses, we have developed three executables: `apply`, `verify` and `analyse`.

All of these executables are written from scratch in C++. The compiler employed is `MinGW` on Windows and `gcc` on Linux and Mac. The `boost` li-

braries [96] are employed liberally and the popular OpenMP [95] library is employed for parallelisation.

Parametric modelling - `apply`

As already discussed, modifying the gEDA framework enables us to write netlist templates that contain expressions. As JSIM_n only supports fixed values in the input netlist, we have developed the executable `apply` that is able to parse those expressions and replace them with fixed values. The set of values, termed the *population*, is supplied in a separate file.

Netlist templates contain expressions that are evaluated from the population file. The expressions that are to be replaced are of the form `@{<expression>}`. The entire token, including the "`@{·}`", is replaced by the value of the expression. For example, the line

```
L1 1 2 @{Lstor/2}
```

will be replaced by

```
L1 1 2 3e-12
```

when it is stated elsewhere that `Lstor = 6p`. The entire grammar of the expression tree that is parseable by `apply` is stated in Listing 3.1 in the Extended Backus-Naur Form (EBNF) [97]. Note that the grammar is case insensitive. Some conventions are taken from other common grammars. A variable name may contain alphanumeric characters and the underscore character, but must start with a letter. Common SPICE value expressions such as "10GHz" or "6.2p" are supported.

Listing 3.1: Grammar of expression tree parseable by `apply`.

```
1  expression      = term { ( "+" | "-" ) term }
2
3  term            = expnt { ( "*" | "/" ) expnt }
4
5  expnt           = factor [ "^" factor ]
6
7  factor          = [ "-" | "+" ] factor
8                  | group
9                  | spice_val
10                 | var_name
11
12 group           = "(" expression ")"
13
14 var_name        = ["a"-"z"] { "a"-"z" | "0"-"9" | "_" }
15
16 spice_val       = <floating point number>
17                   [ "a" | "f" | "p" | "n" | "u" | "m" | "k" | "meg" |
                         "g" ]
18                   [ "ohm" | "hz" | "h" | "f" | "a" | "v" | "s" ]
```

An example of a population file is displayed in Listing 3.2. The file consists of lines that start with the name of a variable, followed by attributes of the `key=value` form. The `nom` attribute is employed to set the nominal value of the variable. This is the only attribute read by `apply` (our other tools, described later, make use of more attributes). The nominal values of variables, as specified in the population file, are employed to evaluate the expressions in the netlist template. Note that expressions may also be used as `value`s.

Listing 3.2: Example of a population file readable by `apply`.

```
% This is a comment
L1      nom=3p
B1      nom=0.5
Lstor   nom=6p
L2      nom=Lstor/2
```

The usage of `apply` is exemplified in the command

```
apply netlist-template.cir population.pop | jsim_n -m output.dat,
```

where "|" is the pipe operator that connects the standard output of `apply` to the standard input of `jsim_n`. Note that the omission of a netlist file from the call to `jsim_n` tells JSIM_n that it should read the netlist from the standard input. Hence, `apply` outputs a JSIM_n-compatible netlist synthesised from a netlist template and a value population.

One may also use `apply` by issuing the command

```
apply netlist-template.cir population.pop netlist.cir,
```

in which case the output is written to the file `netlist.cir` instead of the standard output.

There is one special variable that `apply` recognises: `Tnoise`. If the variable `Tnoise` is defined in the value population, its nominal value is assumed to be the noise temperature (in Kelvin) of the environment in which the circuit is simulated. Akin to the procedure described in the JSIM_n user manual [98], current noise sources that simulate Johnson-Nyquist noise [99, 100] are added in parallel with the resistors found in the netlist.

Level verification - `verify`

To enable verification of circuits, we have developed the tool `verify` to compare the simulation output of JSIM_n to an expected template. The template takes the form of a *specification* file, like the one from Listing 3.3.

Listing 3.3: Example of a specification file readable by `verify`.

```
time            B1_P            R1_V
0               2.0944          10e-6
100e-12         z               0
```

```
4  150e-12          1                0
5  200e-12          2                0
6  250e-12          3                d
```

The first line in the specification file is the header line. It begins with the word "time", followed by a whitespace-separated list of variable names. These variables correspond to output variables of JSIM_n. The second line contains the *range specifier* for each variable. The first record is ignored, each subsequent record corresponds to the allowable deviation (in each direction) of each subsequently specified variable level.

Each subsequent line starts with a simulation time index, then specifies the level that each variable should have (within the allowable range) at that time. So, the specification from Listing 3.3 requires the variable `R1_V` (the voltage across resistor R_1) to be within $\pm 10\,\mu V$ of 0 at each specified time index before or on 200 ps.

There are several special cases for such level specifications. For junction phase variables (ending in "_P"), the numerical value provided should be an integer which specifies the multiple of 2π that the phase level should have at the relevant time index. Usually this corresponds to the number of junction switches that have occurred prior to the relevant time-index. In Listing 3.3, variable `B1_P` could refer to a clock junction that switches once every 50 ps. Note that the range specifier remains in (dimensionless) phase units.

The letter `z` "zeroes" the level of the variable at that time (that is, it subtracts the value of the variable at that time from each subsequent variable level check). This is useful for certain configurations where ramping up the bias current may cause junctions to switch in an unpredictable (but ultimately irrelevant) manner. Using the `z` entry, the phase accumulated in this process can be excised from future evaluations.

The letter `d` indicates a "don't care" entry. The level of the variable in question is not evaluated at the relevant time index. Hence, according to the specification in Listing 3.3, the level of variable `R1_V` is not deemed important for circuit operation at 250 ps simulation time.

The intended usage of `verify` is summarised in the command

```
verify spec.sp -jsim netlist.cir
```

where `spec.sp` is the specification file and `netlist.cir` is the netlist file that is to be passed to JSIM_n. This command starts a `verify` process which, in turn, starts a `jsim_n` process. The `verify` process connects to the standard input and standard output of `jsim_n`. The netlist is passed to `jsim_n`, which begins to simulate the netlist as soon as it has been passed. While simulating, `jsim_n` outputs the simulation results to the standard output, which is absorbed by parent process `verify`. The output is compared to the specification as it arrives. When the specification is violated, `jsim_n` is immediately terminated by `verify`, which subsequently returns an error. If the specification is not

violated, and each time index is consumed, `jsim_n` is terminated (if it has not already) and `verify` outputs that the circuit has passed verification. In this way, the shortest possible simulation time is achieved, as failures are identified immediately and simulations only run to their conclusion if required.

Our `verify` tool can also accept `jsim_n` output on its standard input. The command

```
apply netlist-template.cir population.pop | jsim_n | verify spec.sp
```

first applies a population to a netlist template, then feeds the template to `jsim_n`, the standard output of which is attached to `verify`'s standard input. Note that two pipes are constructed and bound by the shell, eliminating the need for disk access and keeping the population separate from the template, consistent with parametric modelling practice.

The operation of `verify` is depicted in Figure 3.10. The above command is tested on the JTL from Figure 3.10a. Voltage source V_1 generates three SFQ pulses that, during normal operation, should propagate along the JTL. Listing 3.3 correctly describes this behaviour, specifying the phase across B_1 and the voltage across R_1. The only parameter in question is the bias current I_b, applied at current sources I_{b1} and I_{b2}.

Figure 3.10b depicts the operation for the bias current of $I_b = 350\,\mu\text{A}$. The three pulses propagate correctly. Figure 3.10c, however, depicts failure for $I_b = 50\,\mu\text{A}$. Only two pulses make it through, the third one is stored in the JTL. This is incorrect according to Listing 3.3, which is picked up by `verify`.

Analysis - `analyse`

As we are now able to parametrically specify our circuits, and verify them in an efficient manner, we are able to do more sophisticated analyses. Our tool `analyse` can be employed for customisable margin and yield analyses. Care was taken to write `apply` and `verify` in a modular manner, so that the code can be re-used for `analyse`.

Parameters are set up for margin analysis in the population file. In addition to the nominal value (`nom` attribute), each parameter may be adorned with the attributes listed in Table 3.3. Attributes `min` and `max` constrain the range that is to be investigated during the analysis, whereas `prec` is employed to set the precision of the analysis. For `prec=0.01`, the margins are calculated with a precision of at least 0.01 times the nominal. The `ma` attribute is employed to set which kind of margin analysis to run. A value of `n` runs a normal margin analysis relying on the binary search method. A value of `e` runs an *exhaustive* margin analysis, which does not rely on the binary search method but simply discretises the investigated range into enough points to guarantee the required precision and evaluates the circuit at every point. In this way,

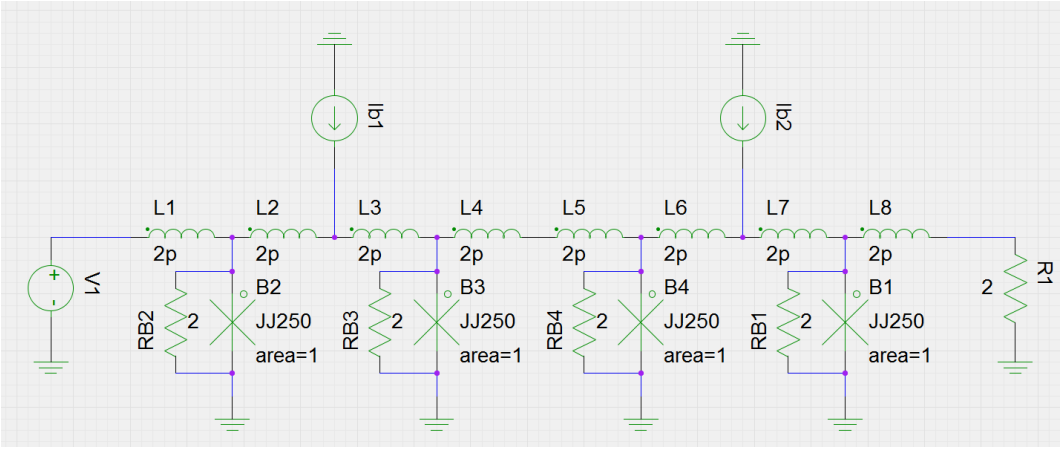

(a)

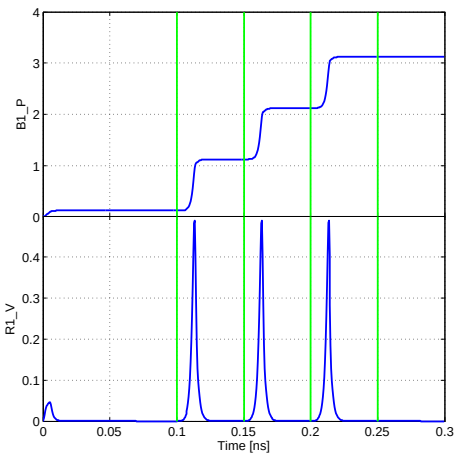

(b)

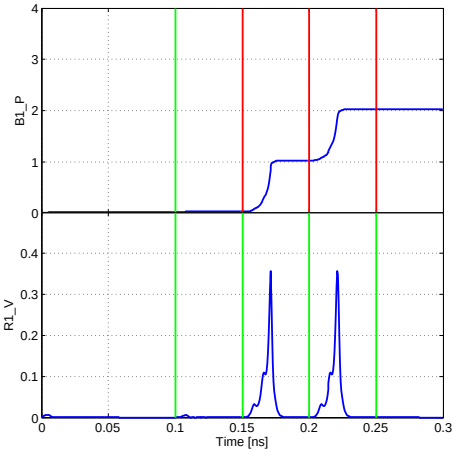

(c)

```
verify output:
1e-010 :  z.
1.5e-010 :  ..
2e-010 :  ..
2.5e-010 :  .d
+SUCCESS Successfully verified!
```

```
verify output:
1e-010 :  z.
1.5e-010 :  x
-ERROR Not verified.
```

Figure 3.10: Our tool `verify` operating on a circuit (a), for the successful (b) and an unsuccessful (c) case. The specification is from Listing 3.3.

Table 3.3: Parameter attributes for margin analyses.

Key	Value	Description
`min`	floating-point	Minimum value (multiple of nominal).
`max`	floating-point	Maximum value (multiple of nominal).
`prec`	floating-point	Precision of margins.
`ma`	{`n`,`e`,`i`}	Type of analysis.
`depth`	integer	Descent-depth when updating parameters.

multiple regions of operation can be found, if they exist. A value of `i` ignores the relevant parameter during the margin analysis.

The integer attribute `depth` can be used to set how deep the expression tree should be evaluated when updating the parameter. When `depth=0` for parameter `X`, then no parameters relying on parameter `X` will be updated when `X` is varied during the margin analysis. When the `depth` attribute is omitted entirely, or set to -1, then all parameters dependent on `X`, as well as their dependents, will be updated recursively. If `depth` is set to a positive integer, the number of descent steps taken when updating parameters is equal to that integer.

To illustrate the working of `analyse`, we employ the DFF test circuit displayed in Figure 3.11. Three clock pulses are applied, with a data pulse applied between the first and second clock pulses. Correct operation requires junction switching behaviour as captured in the specification of Listing 3.4.

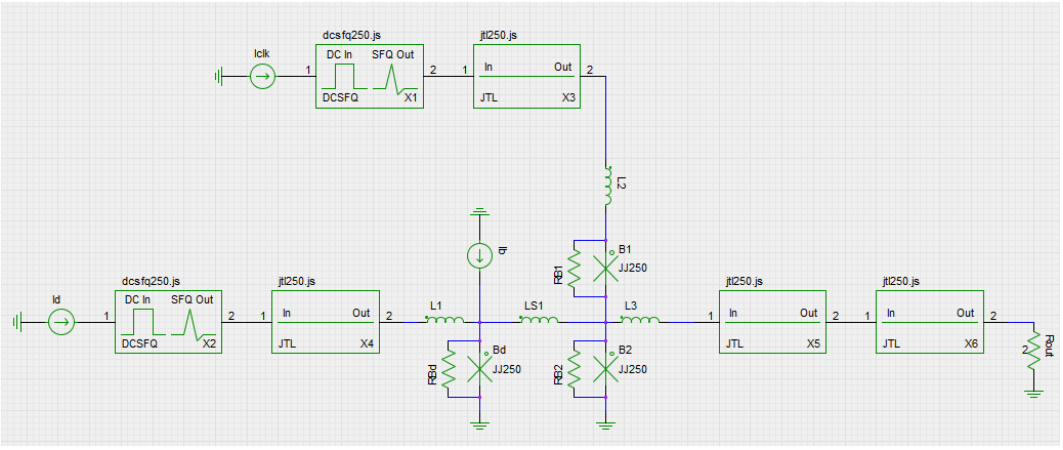

Figure 3.11: DFF in test environment to show usage of `analyse`.

Listing 3.4: Specification for example DFF test circuit.

```
1 time            BD_P          B1_P          B2_P
2 0               2.0944        2.0944        2.0944
3 50e-12          z             z             z
```

```
4  100e-12          0                0                0
5  150e-12          0                1                0
6  200e-12          1                1                0
7  300e-12          1                1                1
8  400e-12          1                2                1
```

The population file for the example DFF is shown in Listing 3.5. The parameters to be fixed (not investigated during margin analysis) have attribute `ma=i`. The default value for those parameters where it is not specified is `ma=n`. Other defaults are `min=0.1, max=1.9` and `prec=0.01`.

Listing 3.5: Population for example DFF test circuit.

```
1  % Parameters of interest
2  LS1      nom=6p
3  Ib       nom=150u                  min=0    max=3
4  B1       nom=175u/Ic0                       depth=0
5  B2       nom=250u/Ic0                       depth=0
6  Bd       nom=250u/Ic0                       depth=0
7  betac    nom=1            min=0.1 max=5            prec=0.1
8  B#                        min=0.2 max=2
9
10 % Inductors
11 L1       nom=2p           ma=i
12 L2       nom=2p           ma=i
13 L3       nom=2p           ma=i
14 % Shunt resistors
15 RB1      nom=((phi0*betac)/(2*pi*Ic0*C0*B1^2))^0.5        ma=i
16 RB2      nom=((phi0*betac)/(2*pi*Ic0*C0*B2^2))^0.5        ma=i
17 RBd      nom=((phi0*betac)/(2*pi*Ic0*C0*Bd^2))^0.5        ma=i
18
19 % Junction model
20 Ic0      nom=250u                  ma=i
21 C0       nom=330f                  ma=i
```

Note in particular the `depth=0` adornment of the variables `B1, B2` and `Bd`. These variables represent junction areas, which are used in the expressions for their shunt resistances `RB1, RB2` and `RBd`. As we do not want the shunt resistances varying with the junction areas in our margin analysis, we specify the descent-depth as 0.

On the other hand, the `depth` attribute is omitted from the variable `betac`. Hence, the shunt resistances, which are specified in terms of `betac`, will vary when `betac` is swept across the investigated range, as is our intention.

Also note the variable `B#`. This is a special variable with no nominal value. The `#` character is a wildcard, meaning that `B#` refers to all variables starting with "`B`". This means that for the investigation of `B#`, all these variables are swept and the proportional margins established this way.

The command to run the margin analysis is

```
analyse -ma netlist-template.cir specification.sp population.pop
```

where "`-ma`" is the directive indicating to `analyse` that a margin analysis is to be performed. The output of the command for our test DFF circuit is reproduced in Listing 3.6.

Listing 3.6: Output of margin analysis for DFF test circuit.

```
Copyright (C) 2011-2012 Stellenbosch University (20120815)
Licensed under the Boost Software License - Version 1.0
Compiled with OpenMP threading support: [yes: 4 processors]

Performing margin analysis.

7 parameters will be investigated.
Par              Nominal           Min           Max        Min%        Max%
================================================================================
B1            7.000e-001    7.492e-002    1.094e+000      -89.30       56.25
Bd            1.000e+000    1.703e-001    1.710e+000      -82.97       71.02
B2            1.000e+000    1.000e-001    1.394e+000      -90.00       39.37
B#                                                      -65.62       67.97
betac         1.000e+000    1.000e-001    5.000e+000      -90.00      400.00
LS1           6.000e-012    3.384e-012    1.009e-011      -43.59       68.20
Ib            1.500e-004    0.000e+000    3.645e-004     -100.00      142.97
[================================================] 7/7 (100.0 %)
Margin analysis completed.
Name      Margins
B#          [...=======|=========...]                                  -65.6,  68.0
B1         [===========|=======....]                                   -89.3,  56.2
B2         [===========|=====......]                                   -90.0,  39.4
Bd         [..=========|=========..]                                   -83.0,  71.0
Ib        [============|==================......]                     -100.0, 143.0
LS1        [......=====|=========..]                                   -43.6,  68.2
betac      [.==========|================================================]   -90.0, 400.0
```

Our tool `analyse` can also perform yield analyses. A control variable named `.runs` must be added to the population for the yield analysis. Its nominal value is an integer and specifies the number of runs that should be performed. In addition, each parameter that is to be part of the yield analysis must be adorned with the attribute `spread`, of which the value corresponds to 3 times the standard deviation σ that pertains to the parameter.

The directive for a yield analysis is `-ya`, making the command to be issued

```
analyse -ya netlist-template.cir specification.sp population.pop.
```

Note that the only difference between this command and the command for the margin analysis is the directive, all relevant information must be contained in the `population.pop` file. More than one directive can be passed to `analyse` in the same call, in which case they will be processed sequentially.

A special kind of yield analysis is the *Bernoulli analysis*, which maintains the parameters at their nominal values but performs a number of trials, as in the yield analysis. The number of trials must be specified as the (integer) nominal value of the `.trials` control variable. Naturally, if the parameters remain fixed, the deterministic outcome of each trial must always be the same, ensuring that the reported "yield" must either be 0 or 100%.

The Bernoulli analysis thus only makes sense in the presence of non-deterministic influences such as thermal noise. In that case, the "yield" reported by the Bernoulli analysis corresponds to the probability of the circuit working as per the specification, or the inverse of the probability of the noise ruining circuit operation. The directive to perform a Bernoulli analysis is `-ba`.

All of the analyses described require many JSIM_n simulations that are essentially independent of each other. For the yield- and Bernoulli analyses, there is no interdependence between the different simulations, each is a self-contained unit of work. For the (binary search) margin analysis this is not the case, as the outcome of one simulation determines which circuit is to be simulated next. However, the margin searches for different parameters are independent of each other.

These problems belong to the class of *embarrassingly parallel* problems [101]. As virtually all modern personal computers feature multi-core architectures, it is useful to enable parallel simulations. The `analyse` tool supports parallel execution of independent problems, employing dynamic thread-allocation to minimise the time required. This is useful because the units of work that are performed require different amounts of computational effort (it takes less time for a circuit to fail verification than to pass it). We employed the well-known OpenMP libraries [95] to facilitate parallel execution.

Note that there are no fundamental barriers to developing this tool further and executing it on a distributed computing system. With modern cloud-computing capabilities, this can furnish dramatic advantages in terms of scaling.

Audience

When these tools were written, they were intended only for the present author's use, as he needed the capabilities of the toolset to perform his research. From 2013 onwards, however, they were used at the University of Stellenbosch exclusively for superconducting circuit design, both by students taking relevant courses and by researchers.

3.5 Chapter summary and conclusion

The peculiarities of superconducting circuit design, and how these impose special requirements on the software toolset, were listed and analysed in this chapter. Capabilities of available software were investigated and pitted against designer requirements, where several limitations were identified. We addressed many of these limitations. We further summarise our work and insights developed as follows:

1. Fundamental differences between superconducting and electronic circuits make a one-to-one mapping between superconducting and conventional software tools impractical.

2. Available purpose-written software is limited. Such limitations arise from the limited audience that requires this software, or from clumsy adaptations of software intended for other purposes.

3. Parametric modelling is eminently suitable for superconducting circuits, and practically a requirement for our application. No software available to us currently supports this, except in very limited forms.

4. We have written a new set of software tools that attempts to avoid the pitfalls of available software. The audience is still small, but we hope that once the toolset matures, it will be attractive to other designers. Capabilities include parametric modelling, margin and yield analyses and the parallel execution of simulations. Schematic entry is performed with an adapted open-source tool.

5. The problem of full-wave numerical EM analysis has not been solved yet. More research is required.

The software that we use in this project, resulting from our extensive evaluation and remedial efforts, is summarised in Table 3.4.

Conclusion: Our modified and extended software toolset equips us to perform the designs required for our application. We expect that full-wave EM-solvers are not necessary for our purposes.

Table 3.4: Software used for project.

Purpose	Software tool	Reasons
Schematic capture	gEDA	Flexible, extensible, fast, scales well, open-source
	NioPulse	Capable, university has a license, modern
Simulation	JSIM_n	Fast, scales well, can simulate thermal noise, open-source
Layout	LASI 7	Established, capable, suitable format supported, free
Extraction	InductEx, Parallel FastHenry	Established, capable, modern, accurate, free
Analysis	MATLAB, python	Capable, university has a license
	devloped-for-purpose tools	Capable, fast, scale well, flexible

3.6 Chapter research output

1. A journal article was published in the international journal *IEEE Transactions on Applied Superconductivity*, in conjunction with a conference poster exhibited on 9 October 2012 in Portland, OR, USA at the Applied Superconductivity Conference '12.

   C.J. Fourie and M.H. Volkmann, "Status of superconductor electronic circuit design software,", *IEEE Transactions on Applied Superconductivity*, vol. 23, no. 3, 1300205, June 2013.

   Abstract:
   More than a decade has elapsed since the publication of the last thorough evaluation of the global state of superconductor electronic (SCE) design software, and seven years since the publication of the 2005 SCENET roadmap. In this work we discuss the progress made to date on SCE design software, present a critical analysis of the capabilities of the software available today, and attempt to lay the foundation of a roadmap for SCE design software to complement published hardware and technology roadmaps. The discussion includes design techniques, circuit optimizers, logic simulators, inductance calculators, simulation engines, full-wave EM solvers and the NioCAD project. We also discuss requirements not yet met.

Chapter 4

Superconducting Flash ADC Architecture

4.1 Introduction

Superconducting flash ADCs are not new devices. However, we are not aware of literature that develops the concepts surrounding them in a rigorous and accessible manner. Hence, the work presented in this chapter has the following objectives:

1. Develop and describe the operation of superconducting periodic comparators, obtaining a useful model,
2. identify important model parameters and constraints,
3. identify key performance indicators that quantify model response to achievable parameters,
4. extend the model to include non-idealities due to thermal noise, and
5. investigate the response of the model to the magnitude of the non-idealities.

Achieving these objectives should equip us with the insight necessary to move on to designs of actual circuits and understand the fundamental limitations of the technology.

4.2 A multi-threshold comparator

4.2.1 A valuable response

Consider the parallel combination of an inductor and a Josephson junction, as depicted in Figure 4.1. It will be of benefit to understand the response of this

circuit, as will become evident later. We thus attempt to glean an intuitive understanding of the circuit. To achieve this, we first consider the quantities describing the state of the system represented by this circuit. The quantity known as superconducting phase, or ϕ, was introduced in Section 2.3.2. The only other quantity required to, in conjunction with ϕ, unambiguously describe the state of a superconducting element, is the electric current I. Some of the explanations from this subsection were adapted from Wetzstein [**?**].

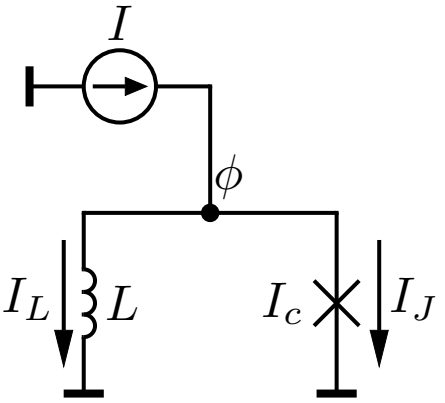

Figure 4.1: Parallel combination of Josephson junction and inductor.

From the circuit it follows that the superconducting phase drop across the junction is necessarily equal to that across the inductor (phase and voltage are analagous in this sense). The first Josephson equation, introduced in Section 2.3.2 and simply repeated here, states that the current I_J through a Josephson junction relates to the phase drop ϕ across it as

$$I_J = I_c \sin \phi \quad , \tag{4.1}$$

where I_c is the critical current of the junction. Note that, to simplify this exercise, we consider only the static case. This equation holds as long as the junction remains in the superconducting state, which requires that $|I_J| \leq I_c$. As this neatly ties the quantities phase and current together for the junction, we next derive the corresponding relationship for the inductor.

The voltage v across an inductor relates to the transmitted current I_L as

$$v(t) = L\frac{di_L(t)}{dt} \quad . \tag{4.2}$$

From (4.2) and the superconducting phase evolution equation (2.4), we conclude that

$$L\frac{di_L(t)}{dt} = \frac{\Phi_0}{2\pi}\frac{d\phi(t)}{dt} \quad ,$$

which, after integrating on both sides, yields

$$i_L(t) = \frac{1}{L}\frac{\Phi_0}{2\pi}\phi(t) + I_0 \quad ,$$

where I_0 is the constant of integration. We presume the absence of an external magnetic field, setting $I_0 = 0$. Hence, we establish the phase-current relationship in an inductor as

$$I_L = \frac{\Phi_0}{2\pi L}\phi \quad . \tag{4.3}$$

Note how, when arranged differently, (4.3) becomes $\phi = Ai_L$ (with some positive constant A), which has the exact same form as Ohm's law. The relationship between current and phase for an inductor is a linear one, congruent to the relationship between current and voltage for a resistor.

Looking back at Figure 4.1, we now understand that as ϕ grows, the current in the inductor grows linearly, whereas the current through the junction oscillates between $\pm I_c$. This is qualitatively illustrated in Figure 4.2.

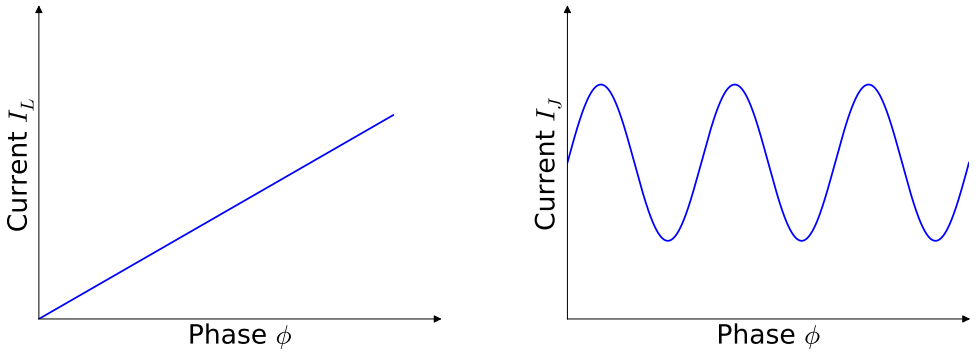

(a) (b)

Figure 4.2: Qualitative response of parallel inductor (a) and Josephson junction (b).

To illustrate the response of the parallel combination rather than the individual elements, we now plot the $I\phi$-curve of the circuit. The total current I through the parallel combination results from (4.1) and (4.3), namely

$$I = I(\phi) = \frac{\Phi_0}{2\pi L}\phi + I_c \sin\phi. \tag{4.4}$$

This relation is illustrated in Figure 4.3, for a large and a small value of I_c. The difference between the two traces is significant: for sufficiently small I_c, the gradient of $I(\phi)$ is never negative, whereas for large I_c the gradient is negative for certain values of ϕ.

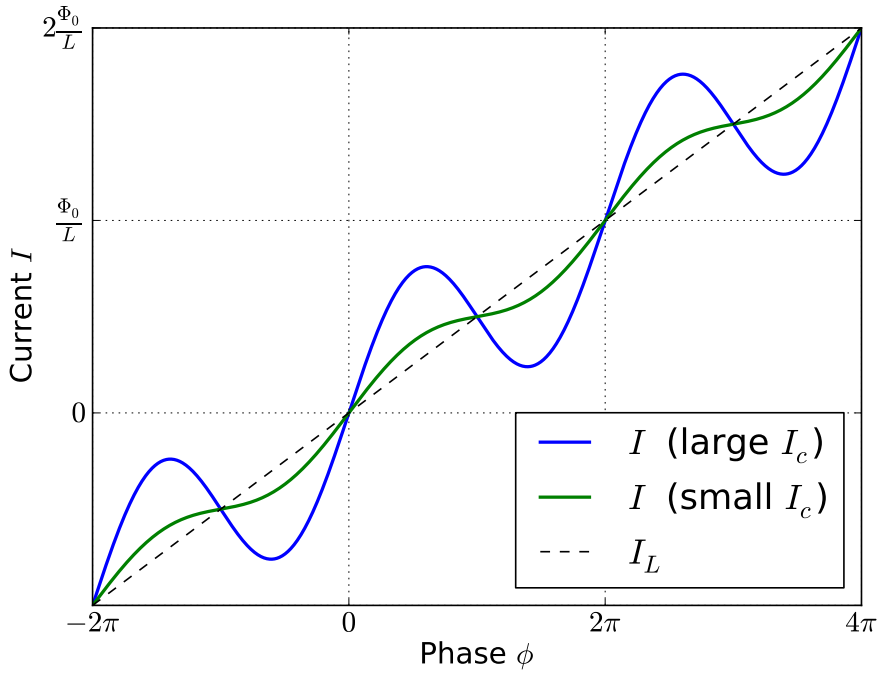

Figure 4.3: A quantitative $I\phi$-curve of the circuit from Figure 4.1. The dashed line is $I_L(\phi)$, included for clarity.

The significance of this difference becomes clear when we consider a practicality: in a real-world application, the current I is the quantity under direct control, whereas the phase ϕ is not. Since we are targeting a real (if somewhat exotic) application, it would be of greater value if we could obtain ϕ as a function of I, that is, invert $I(\phi)$ (4.4). This is only possible for sufficiently small I_c, as only then the function $I(\phi)$ is injective. Note that monotonicity of $I(\phi)$ is a sufficient condition for injectivity.

To translate this condition into a constraint on design parameters (L and I_c), monotonicity is ensured by keeping the gradient of $I(\phi)$ always positive, by requiring

$$I'(\phi) = \frac{\Phi_0}{2\pi L} + I_c \cos\phi \geq 0 \quad \forall\phi,$$

which translates to

$$I_c \cos\phi \geq -\frac{\Phi_0}{2\pi L} \quad \forall\phi.$$

Since the right hand side of this inequality is a negative constant, and the left hand side can minimally be $-I_c$ this reduces to

$$I_c \geq \frac{\Phi_0}{2\pi L}. \tag{4.5}$$

This condition can be expressed as a constraint on L as

$$L \leq \frac{\Phi_0}{2\pi I_c}. \tag{4.6}$$

To illustrate the difference between the injective and non-injective regimes, an axes-inverted version of Figure 4.3 is shown in Figure 4.4, first for sufficiently small- and then for large I_c. Under the assumption that our static model holds up during sufficiently slow dynamic variations (later simulations will provide evidence supporting this assumption), we can trace the state of the system as the variable I changes with time.

In the injective regime (Figure 4.4a), the response of the system to a varying I follows trivially from the trace: as I increases, ϕ increases monotonically and continuously, the point in the state space moves upward along the trace. If I reaches a maximum and then decreases, ϕ decreases as well, the point in the state diagram traverses the reverse of the route it took while I was increasing.

In the non-injective regime (Figure 4.4b), the increase of ϕ with increasing I is still monotonic, but no longer continuous. Instead, the journey of the point in the state diagram can be described as progressions along smooth, piece-wise continuous trajectories, punctuated by "jumps" across portions of the trace. More specifically, as I increases from zero, the state moves along the trace until **A**. Since I is forced by a current source, a continuing increase in I means that the state must "jump" across the "no-go area" to **B**, before ϕ can continue to increase smoothly (until the next jump).

A major difference between this regime and the injective regime becomes apparent when I decreases again, after having increased to, say, **C**. At first, the state moves along the route it took up to **C**, but it is not compelled to jump back to **A** at **B**, rather it traverses the trace until **D**, beyond which it performs the jump to **E** and proceeds in that fashion along the path indicated. It is clear that the downward path of the state is not the reverse of the upward path, the system is hysteretic in this regime.

Our model developed here does not describe what happens to the state variables during the jumps. For now we assume that each jump is characterised by a fast and monotonic increase in ϕ. In fact, each jump is accompanied by an excursion of the Josephson junction outside the superconducting regime, precipitated by an increase of $|I_J|$ beyond I_c. When the junction is critically damped (as described in Section 2.3.2), its phase does indeed experience a fast, monotonic jump during such an excursion — the junction switches.

4.2.2 A periodic response

While the $I\phi$ relationship explored above is clearly not periodic, Figure 4.3 betrays a periodic component. Indeed, when considering I the independent variable, and in the case where the Josephson junction never leaves the superconducting regime, ensured by (4.5), the current through the Josephson

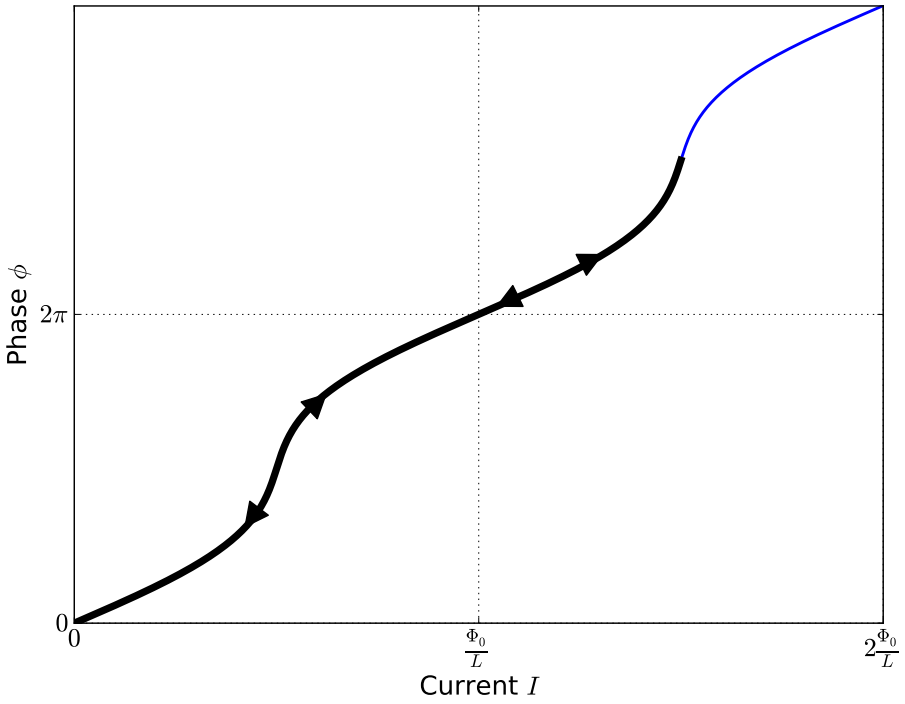

(a) State response of the system in the injective regime.

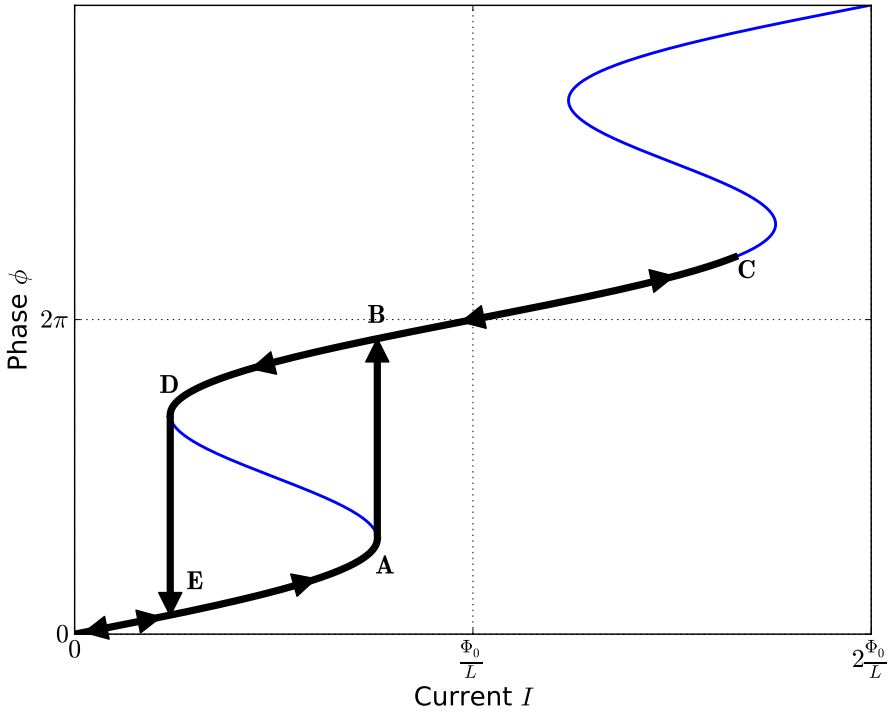

(b) State response of the system in the non-injective regime.

Figure 4.4: State response of the parallel inductor-junction combination.

junction I_J must remain confined to $\pm I_c$. Considering the monotonic increase of ϕ with I together with the dc Josephson effect (4.1), it is clear that I_J will oscillate periodically with a linearly varying I.

To formalise this relationship $I_J(I)$, one could explore solving (4.4) for ϕ and simply inserting the result into (4.1). However, (4.4) cannot be solved for ϕ due to the presence of the transcendental function sin. Instead we start from Figure 4.1, which yields

$$I_L = I - I_J,$$

Rearranging (4.3), we obtain

$$\phi = \frac{2\pi L}{\Phi_0} I_L,$$

or

$$\phi = \frac{2\pi L}{\Phi_0}(I - I_J).$$

Inserting this result into (4.1) yields

$$I_J = I_c \sin\left(\frac{2\pi L}{\Phi_0}(I - I_J)\right). \tag{4.7}$$

Naturally (4.7) is still transcendental and cannot neatly be solved for I_J, but assuming restriction to the injective case, several properties about the relationship it describes can be concluded:

- I_J is a function of I, though not an invertible one.
- The range of $I_J(I)$ is confined to the interval $[-I_c, I_c]$, while the domain of $I_J(I)$ comprises $\mathbb{R}$.
- $I_J(I)$ is indeed periodic, with period

$$I_m = \frac{\Phi_0}{L}. \tag{4.8}$$

Although (4.7) is transcendental, as it describes a function there is a unique output for every input value I. We can find these outputs by using, for example, the Newton-Raphson method [102]. Employing this method in SciPy [103], a graphical representation of (4.7) is obtained and depicted in Figure 4.5. $I_J(I)$ resembles a sinusoidal function with a "sawtooth distortion". The listed properties are clearly visible.

Figure 4.5 reveals another property:

- The range of I_J is equally divided between $I_J \geq 0$ and $I_J \leq 0$ over a linearly varying I.

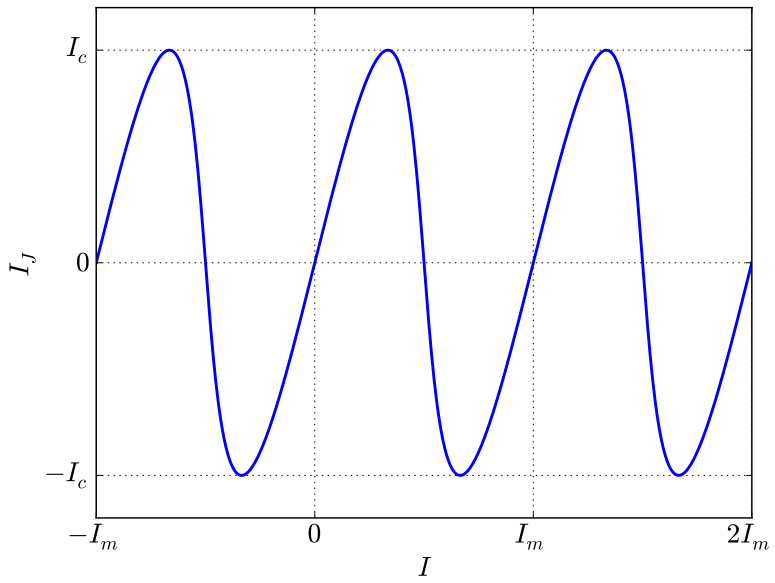

Figure 4.5: Static response of I_J as a function of the independent variable I in the injective case.

This property is important, as it allows us to divide the domain of the function into equally-spaced bins when discriminating between $I_J > 0$ and $I_J < 0$. This property is what makes the one-junction SQUID from Figure 4.1 a *quantiser*. The one-junction SQUID was first described for this purpose by Ko and Van Duzer in 1988 [104].

4.2.3 The quasi-one-junction SQUID multi-threshold comparator

A *discriminator* as mentioned in the previous subsection can be realised in the form of a *comparator*. A comparator converts an analog input signal to a digital output which discriminates between two states: 1) input signal below threshold and 2) input signal above threshold.

Figure 4.6 shows a conceptual implementation of the described assembly. The quantised current signal, produced by the quantiser described in the previous section, is fed into a comparator with threshold set at $I_{\text{th}} = 0$. The output of the comparator is now a classification of the input signal into one of two "bins", designated `1` and `0`. This is a relative digital representation of the input signal.

A comparator that takes current as input signal is realised by a *decision-making pair* (DMP) of Josephson junctions. Such a DMP is schematically depicted in Figure 4.7.

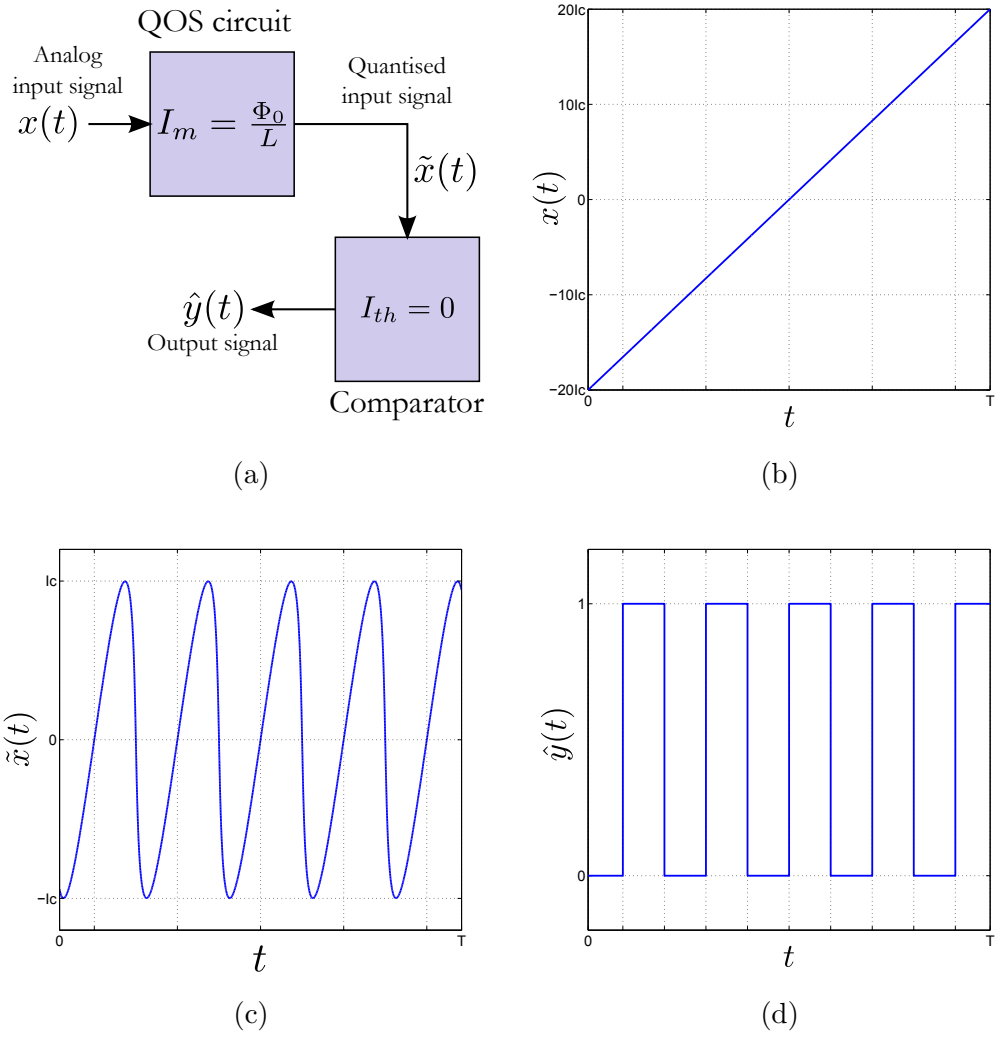

Figure 4.6: Abstract MTC, a multi-threshold comparator with binary output (a). A first-order input signal $x(t)$ (b) is quantised by the QOS circuit, which produces $\tilde{x}(t)$ (c). The quantised signal is fed into a comparator with threshold at zero, yielding coded output signal $\hat{y}(t)$ (d).

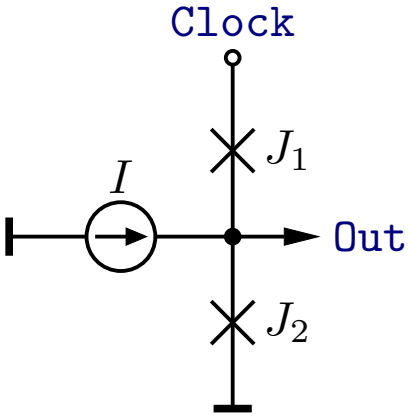

Figure 4.7: A *decision-making pair* made from two Josephson junctions.

The current signal enters the DMP at the point where the two junctions are connected. To perform the discrimination operation, the comparator must be interrogated by an SFQ pulse, which is injected at the `Clock` terminal. The corresponding 2π phase drop over the comparator must be equalised by the switching of exactly one of the comparator junctions.

The critical currents of the junctions, as well as other circuit parameters, are designed in such a way that J_1, the upper junction, switches when the input signal I is below the comparator threshold. In this case, the 2π phase drop is equalised by a 2π increase across J_1, and the SFQ "escapes". When the input signal is above the comparator threshold, however, the lower junction J_2 switches instead, providing the equalising 2π increase in phase. In this case, an SFQ pulse occurs at `Output`.

In this way, the binary output signal $y(t)$ from Figure 4.6 can be generated. This assembly is known as a *multi-threshold comparator* (MTC) as it changes its output across multiple thresholds, as opposed to the more conventional *single-threshold comparator* ubiquitous in electronics.

The full circuit schematic is shown in Figure 4.8. This circuit is known as the *quasi-one-junction SQUID* (QOS) comparator [104], as the additional junction J_2 does not participate significantly in the response of the SQUID. Signal names and circuit parameters are redefined here in an effort to progress from a conceptual discussion of the core components towards a real system. The important circuit parameters are listed and described in Table 4.1.

4.3 The superconducting flash ADC

4.3.1 An n-bit superconducting flash ADC

The periodic response of superconducting SQUIDs can be exploited for analog-to-digital conversion using a concept first proposed by Zappe in 1975 [52].

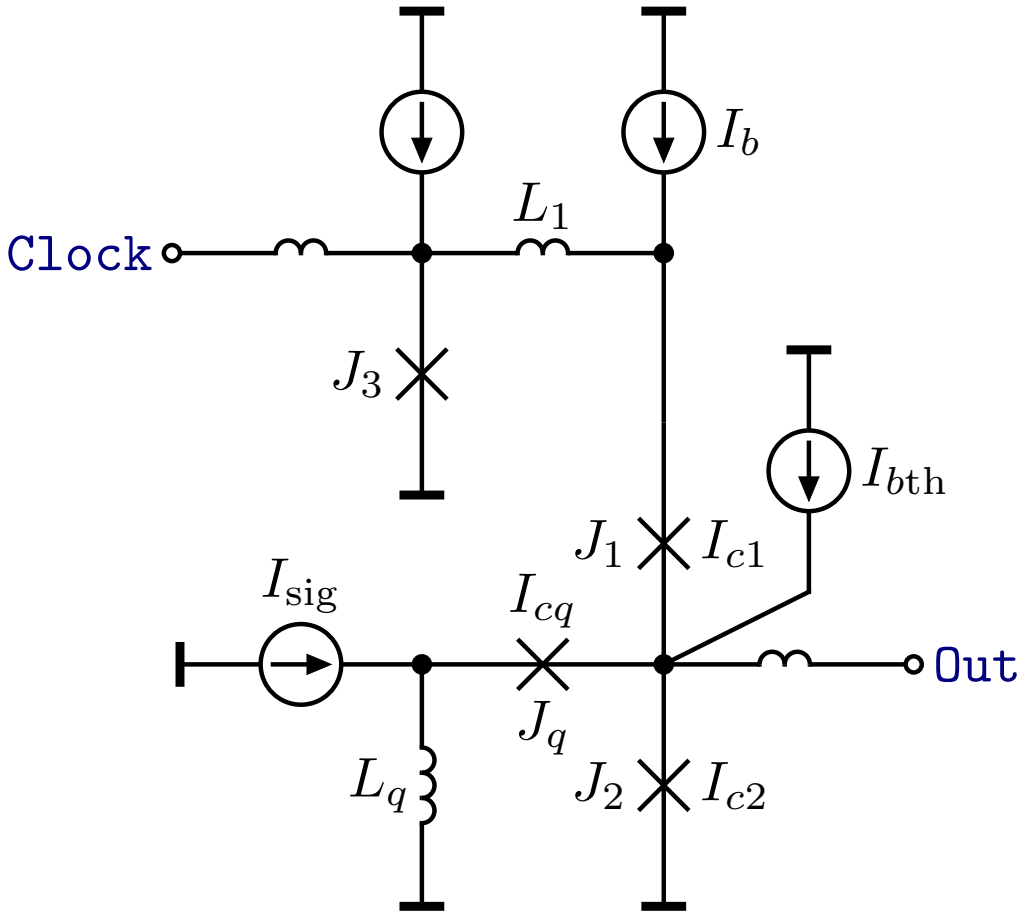

Figure 4.8: Schematic of a QOS multi-threshold comparator.

An ideal MTC as described in Section 4.2.3, with period I_m, divides an input signal I_{sig} into bins of `1` and `0`, each of width $\frac{1}{2}I_m$, interpretable as a relative digital representation of the signal. This representation is shown in Figure 4.9. It is important to note that simply crossing the threshold does not produce a measurable change at `Output`, for this the MTC must be interrogated. Generally such an interrogation occurs by means of a clock signal that periodically samples the output. We define the term *interrogation potential*, which is a measure of the digital output a circuit would produce on a clock event at the time under consideration.

If I_{sig} is confined to $I_0 \pm \frac{1}{2}I_m$, with I_0 a threshold, the MTC can be interpreted as a 1-bit ADC. Naturally, for that application, the MTC is not really necessary, a single-threshold comparator would suffice. Next, however, we consider a resistive divider feeding n MTC cells, each with a fraction of I_{sig}.

Figure 4.10 shows just such an assembly with $n = 4$. The combined digital output of the assembly is outlined in Figure 4.10b. It is clear that the output is a digital representation of I_{sig}. When confined to a certain range, indicated

Table 4.1: Important parameters of the QOS multi-threshold comparator

Parameter	SI Quantity [SI Unit]	Description
	Design parameters	
I_{cq}	Current [A]	Critical current of the quantising junction J_q.
L_q	Inductance [H]	Quantising inductance of the QOS quantiser.
I_{c1}	Current [A]	Critical current of upper comparator junction J_1.
I_{c2}	Current [A]	Critical current of lower comparator junction J_2.
L_1	Inductance [H]	Inductance of the DMP interrogation loop J_3, L_1, J_1, J_2.
	Experimental parameters	
I_{sig}	Current [A]	Input signal.
$I_{b\text{th}}$	Current [A]	Comparator threshold dc adjust current.
I_b	Current [A]	Decision-making pair dc bias current.

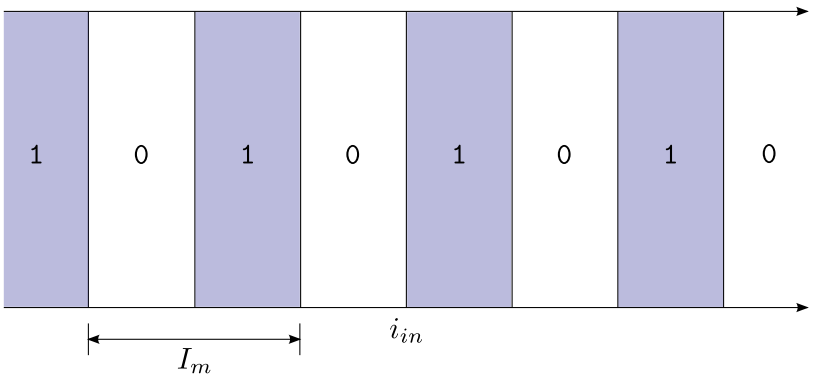

Figure 4.9: Bins of interrogation potential into which I_{sig} is divided by a single, ideal MTC.

on Figure 4.10b, the output represents an n-bit, digital representation of I_{sig}. Note that the output is coded in Gray code [105], which, akin to binary code, has a one-to-one mapping to decimal representation, as reproduced in Table 4.2. One benefit of Gray code over, say, binary code, is that each transition from one digital level to the next is characterised only by a single change (bit flip) in the output word. This makes the effect of switching errors, to be analysed later, more manageable.

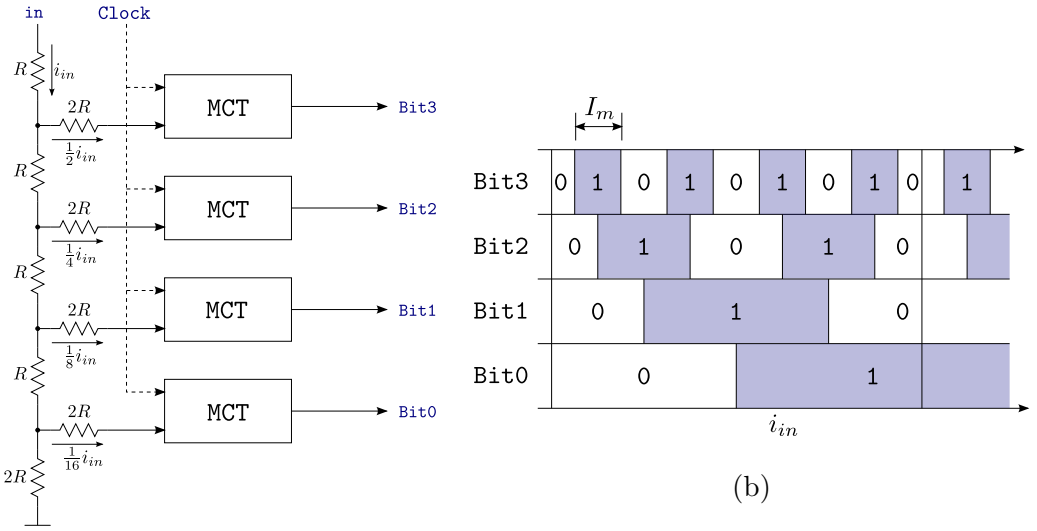

Figure 4.10: Schematic of ADC4, a 4-bit superconducting flash ADC (a), and the interrogation potential of the four comparators as I_{sig} varies (b).

This flash ADC concept, which allows clocking at SFQ-logic speeds, requires only n comparators. Known semiconductor flash ADCs use $O(2^n - 1)$ comparators instead, making them unsuitable for high resolutions. However, it should be stated here that the integration density achievable in current superconducting foundries does not permit a particularly high-resolution implementation of this superconducting concept either. Furthermore, signal strength considerations may make implementation in a future highly-integrated process difficult. Nonetheless, the relative simplicty of the concept is attractive and has withstood experimental scrutiny in current and older technology [11, 61, 60].

4.3.2 Achievable signal-to-noise ratio for an ideal flash ADC

Generally speaking, an *ideal* ADC should produce an exact digital representation of the input signal $x(t)$, so that the original signal could be reconstructed

Table 4.2: 4-bit Gray-to-decimal conversion table

Gray code	Decimal	Gray code	Decimal
0000	0	1100	8
0001	1	1101	9
0011	2	1111	10
0010	3	1110	11
0110	4	1010	12
0111	5	1011	13
0101	6	1001	14
0100	7	1000	15

perfectly by an ideal digital-to-analog converter (DAC).

Nyquist states that sampling a signal at a frequency F_s higher than twice the signal bandwidth results in perfectly preserved information, that is, the signal could be reconstructed in its entirety by a suitable interpolation filter. So far, an ideal ADC may still be possible for a band-limited signal.

However, this would require perfect digitisation of the signal level at periods $\frac{1}{F_s}$, where the signal level is (generally, and clearly in the case of the signals the SKA will receive) continuous. If one assumes that a sampling mechanism exists that can achieve arbitrary precision, this would also require an arbitrary amount of memory and thus would not scale well with increasing precision. This is clear if one considers that the signal level will traverse an infinite amount of irrational numbers, which cannot be exactly represented in finite memory or transmitted digitally in finite time.

Hence, quantisation necessarily introduces an error into the digitisation of the signal. Note that an error is not usually attributed to the process of *coding*, since it is, in practice, adapted to fully accommodate the chosen quantisation scheme.

Correspondingly, we define an *ideal* ADC as one that introduces only quantisation noise and no other loss of information. We further define the continuous input signal as $x(t)$ and the discrete signal

$$x[k] = x\left(\frac{kt}{F_s}\right), \quad k \in \mathbb{Z}_{\geq 0} \tag{4.9}$$

as a perfect sample of $x(t)$ for the causal region $t \geq 0$.

Quantisation error

The error introduced by the quantisation process is generally modelled as $e_q[k]$, so that

$$x_q[k] = x[k] + e_q[k], \tag{4.10}$$

with $e_q[k]$ termed the *quantisation noise* [106].

A key performance indicator for any ADC is the *signal-to-quantisation-noise ratio* (SQNR), defined as

$$\text{SQNR} = \frac{P_x}{P_q}, \tag{4.11}$$

where P_x is the signal power (of the desired signal) and P_q is the noise power.

To clarify, SQNR refers to the power ratio between any signal and the corresponding quantisation noise. We are more interested in the *maximum* SQNR achievable from an ADC, which imposes an upper bound on its dynamic range.

We continue our analysis from a probabilistic viewpoint and draw on Lathi's explanation [106], although others might find other approches more intuitive. We choose this path because non-idealities considered later are readily modelled in probabilistic terms.

Signal and noise power are defined as

$$P_x = E\left[x^2[k]\right], \quad P_q = E\left[e_q^2[k]\right],$$

respectively, where $E[\cdot]$ is the expected value function.

Quantisation by truncation, for unipolar signals

In the previously introduced conceptual ADC4 depicted in Figure 4.10, sampling and quantisation happens simultaneously. Continuous-time input signal $x(t)$ is quantised by *truncation* when the ADC is triggered at t_1 (by applying an SFQ pulse to the `Clock` terminal). The decision levels partitioning the input signal into bins correspond to the quantisation levels assigned, thus the produced output word of the quantiser is a truncated version of the instantaneous input level $x(t_1)$.

For an n-bit flash ADC trivially extrapolated (conceptually) from ADC4, there are $M = 2^n$ quantisation levels spaced uniformly apart. The quantiser step size, that is, the difference between two successive quantisation steps, is

$$\Delta = \frac{1}{2} I_m. \tag{4.12}$$

First, we consider the flash ADC only for unipolar input signals, so that its input range is confined to

$$0 \leq x(t) \leq I_{\max},$$

with

$$I_{\max} = M\Delta = 2^{n-1} \cdot I_m.$$

This quantisation scheme is summarised in Figure 4.11.

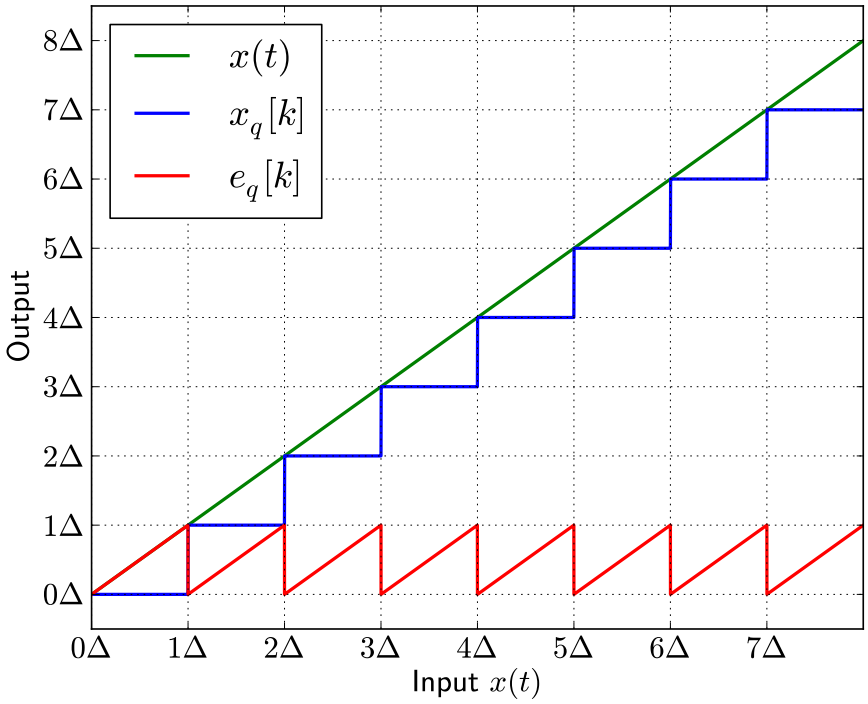

Figure 4.11: Quantisation by truncation, a natural scheme for unipolar inputs and the flash ADC architecture. Here shown for $n = 3$.

When $x(t)$ remains within this input range, the quantisation error is limited to

$$|e_q| \leq e_{\max} = \Delta,$$

as the truncation process maps inputs close to, but not yet exceeding a decision level to the next lowest decision level, introducing an error that is maximally one decision level (Δ) large.

Note that the wrap-around nature of this ADC is such that inputs outside the range are coded to levels within the range, introducing greater error. Thus, inputs outside the range should be prevented.

Note also that Δ corresponds to the quantity *least-significant bit* (LSB).

To analyse the SQNR of the 4-bit design proposed in Figure 4.10, a statistically significant distribution of input signals could be considered. The error signal magnitude is generally assumed to be equally likely to take on any value between 0 and $e_{\max}$. This is a valid assumption for many real random signals, but can be artificially ensured by the intentional addition of noise to the input signal (which can later be subtracted in the digital domain). This practice is known as *dithering* [107], and will probably be employed in the SKA.

This assumption enables definition of the quantisation error by a uniform probability density function (PDF),

$$f_{e_q}(z) = \begin{cases} \frac{1}{\Delta} & \text{if } 0 \leq z \leq \Delta \\ 0 & \text{otherwise.} \end{cases}$$

which is depicted in Figure 4.12a.

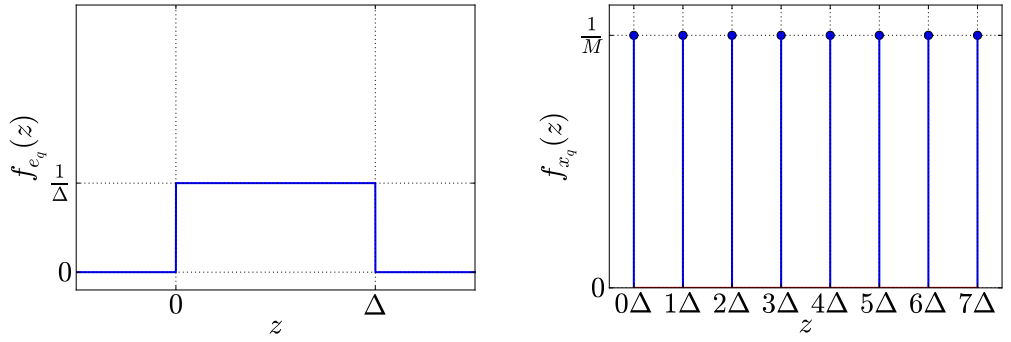

(a) (b)

Figure 4.12: Uniform distribution of the continuous random variable e_q (a) and the discrete random variable x_q (b) for quantisation by truncation.

The expected value of the quantisation error is, trivially,

$$\mu_{e_q} = E[e_q] = \frac{1}{2}\Delta.$$

The average power of the error signal is

$$\begin{aligned} P_e &= E[e_q^2] \\ &= \frac{1}{\Delta}\int_0^{\Delta} z^2 dz \\ &= \frac{1}{3}\Delta^2 \\ &= \frac{I_m^2}{12}. \end{aligned} \tag{4.13}$$

Continuing in this vein, we consider $x_q[k]$ to be a discrete random variable that is equally likely to take on any ADC output value. For the simple coding proposed, it thus assumes the uniform probability mass function (PMF)

$$f_{x_q}(z) = \begin{cases} \frac{1}{M} & \text{if } z \in \{0, \Delta, 2\Delta, \ldots, (M-1)\Delta\} \\ 0 & \text{otherwise,} \end{cases}$$

which is depicted in Figure 4.12b.

The expected value of x_q is thus

$$\begin{aligned} \mu_{x_q} = E[x_q] &= \sum_{i=0}^{M-1} (i\Delta) \cdot f_{x_q}(i\Delta) \\ &= \sum_{i=0}^{M-1} i\Delta \left(\frac{1}{M}\right) \\ &= \frac{1}{M}(0 + 1 + 2 + \cdots + (M-1))\Delta \\ &= \frac{1}{2}(M-1)\Delta \end{aligned}$$

and the signal power P_x follows as,

$$\begin{aligned} P_x &= E[x_q^2] \\ &= \frac{1}{M} \sum_{i=0}^{M-1} (i\Delta)^2 \\ &= \frac{\Delta^2}{M} \sum_{i=0}^{M-1} i^2 \\ &= \frac{\Delta^2}{M} \left[\frac{(M-1)(M)(2M-1)}{6}\right] \qquad \text{by Faulhaber's formula [108]} \\ &= \frac{\Delta^2}{6}(M-1)(2M-1) \\ &= \frac{I_m^2}{24}(2^n - 1)(2^{n+1} - 1). \end{aligned} \tag{4.14}$$

Note that P_e is independent of n, whereas P_x grows exponentially with n.

The SQNR achieved by this quantisation scheme follows from inserting (4.13) and (4.14) into (4.11), which yields

$$\text{SQNR} = \frac{1}{2}(2^n - 1)(2^{n+1} - 1). \tag{4.15}$$

Quantisation by rounding, for bipolar signals

The quantisation scheme by truncation discussed for unipolar signals is not optimal, it is simply a natural way to interpret the output of the conceptual flash ADC. Since the SKA will be receiving bipolar signals, the quantisation scheme is modified to the one depicted in Figure 4.13, quantisation by *rounding*. Note that the output signal is coded so that a zero input corresponds to the middle of the output range, and that an offset may have to be artificially introduced to achieve this. Since these are current signals, and currents easily sum, this can be achieved easily.

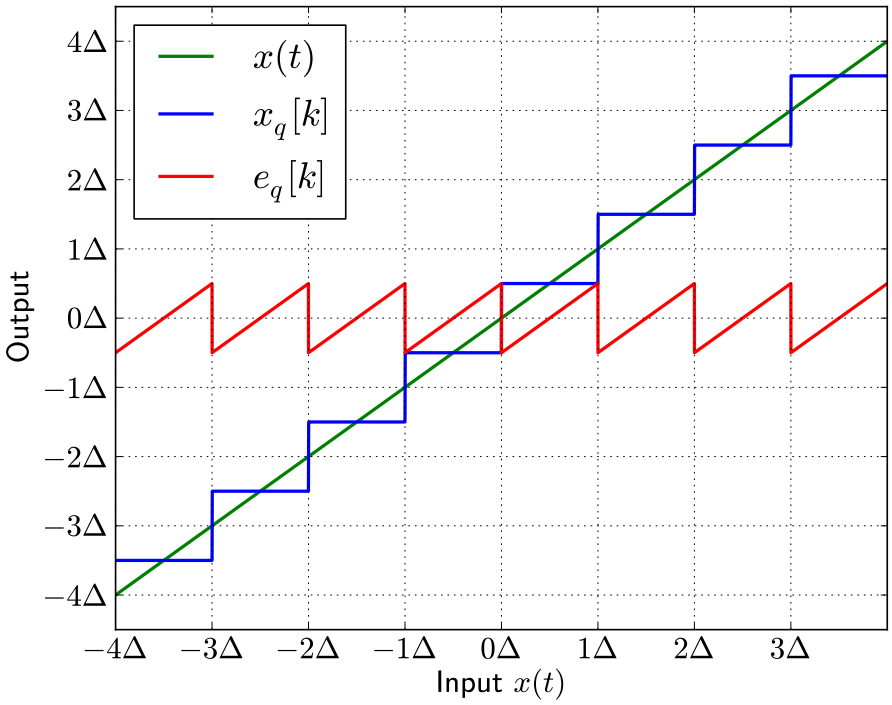

Figure 4.13: Quantisation by rounding, with modification for bipolar input signals. Here shown for $n = 3$.

In contrast to quantisation by truncation, quantisation by rounding as depicted in Figure 4.13 has the clear advantage that the maximum quantisation error is limited to

$$|e_q| \leq e_{\max} = \frac{1}{2}\Delta,$$

not Δ. We continue to make the assumption that each input signal level is equally likely, which allows us to characterise both e_q and x_q as uniformly distributed random variables. For e_q, the PDF is

$$f_{e_q}(z) = \begin{cases} \frac{1}{\Delta} & \text{if } -\frac{1}{2}\Delta \leq z \leq \frac{1}{2}\Delta \\ 0 & \text{otherwise,} \end{cases}$$

as depicted in Figure 4.14a. Analogous to the unipolar case, the average power of the error signal is

$$\begin{aligned} P_e &= E[e_q^2] \\ &= \frac{1}{\Delta}\int_{-\frac{1}{2}\Delta}^{\frac{1}{2}\Delta} z^2 dz \\ &= \frac{1}{12}\Delta^2 \\ &= \frac{I_m^2}{48}. \end{aligned} \tag{4.16}$$

The desired signal, x_q, has the probability mass function

$$f_{x_q}(z) = \begin{cases} \frac{1}{M} & \text{if } z \in \{-\frac{M-1}{2}\Delta, \ldots, -\frac{1}{2}\Delta, \frac{1}{2}\Delta, \ldots, \frac{M-1}{2}\Delta\} \\ 0 & \text{otherwise,} \end{cases}$$

which is shown in Figure 4.14b. The average power of the desired signal is

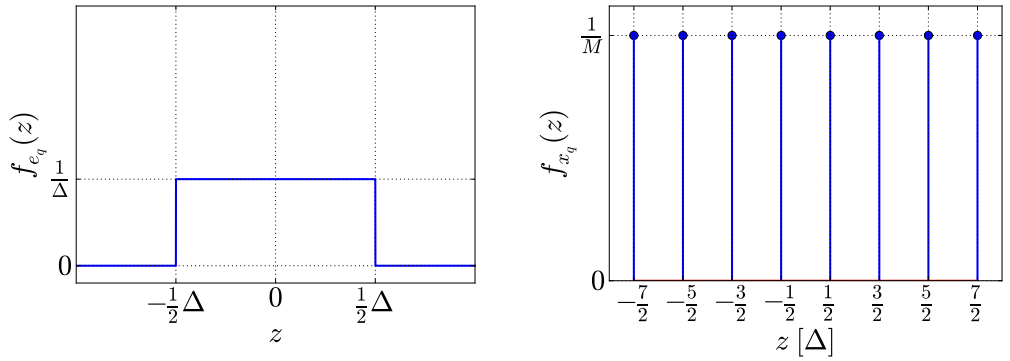

(a) (b)

Figure 4.14: Uniform characteristics of the continuous random variable e_q (a) and the discrete random variable x_q (b) for quantisation by rounding.

$$
\begin{aligned}
P_x &= E[x_q^2] \\
&= \frac{1}{M}\sum_{i=0}^{M-1}\left(\left(i-\frac{M-1}{2}\right)\Delta\right)^2 \\
&= \frac{\Delta^2}{M}\sum_{i=0}^{M-1}\left(i-\frac{M-1}{2}\right)^2 \\
&= \frac{\Delta^2}{4M}\sum_{i=0}^{M-1}(2i-(M-1))^2 \\
&= \frac{\Delta^2}{M}\sum_{i=0}^{M-1}\left(i^2-i(M-1)+\frac{1}{4}(M-1)^2\right) \\
&= \frac{\Delta^2}{M}\left[\frac{(M-1)^3}{3}+\frac{(M-1)^2}{2}+\frac{M-1}{6}-\frac{M(M-1)^2}{4}\right] \qquad \text{by Faulhaber's formula [108]} \\
&= \frac{\Delta^2}{M}\left[\frac{M^3}{12}-\frac{M}{12}\right] \\
&= \frac{\Delta^2}{12}(M^2-1) \\
&= \frac{I_m^2}{48}\cdot(2^{2n}-1).
\end{aligned}
\tag{4.17}
$$

To obtain the SQNR achievable by this quantisation scheme, (4.17) and (4.16) are inserted into (4.11), which yields

$$\text{SQNR} = 2^{2n}-1. \tag{4.18}$$

4.3.3 Comparator grey-zone: a significant non-ideality

The effect of thermal noise: grey decisions

So far, the 2-junction DMP employed as discriminator in the multi-threshold QOS comparator was assumed to have a fixed threshold I_{th}, with, after a clock event, any input current $I > I_{\text{th}}$ guaranteed to produce an output SFQ pulse and any input current $I < I_{\text{th}}$ guaranteed not to. This is illustrated in Figure 4.15a.

This is not a complete model of DMP behaviour in practice. Thermal noise imposes *threshold jitter* onto the system and thus distorts the probability distribution from Figure 4.15b. Note that this is an important concern even at the low typical operating temperatures of 4.2 K of these (mature niobium technology) circuits.

Real Josephson junctions necessarily have, in parallel with the Josephson element, a resistive channel (this was discussed in Section 2.3.2). In practi-

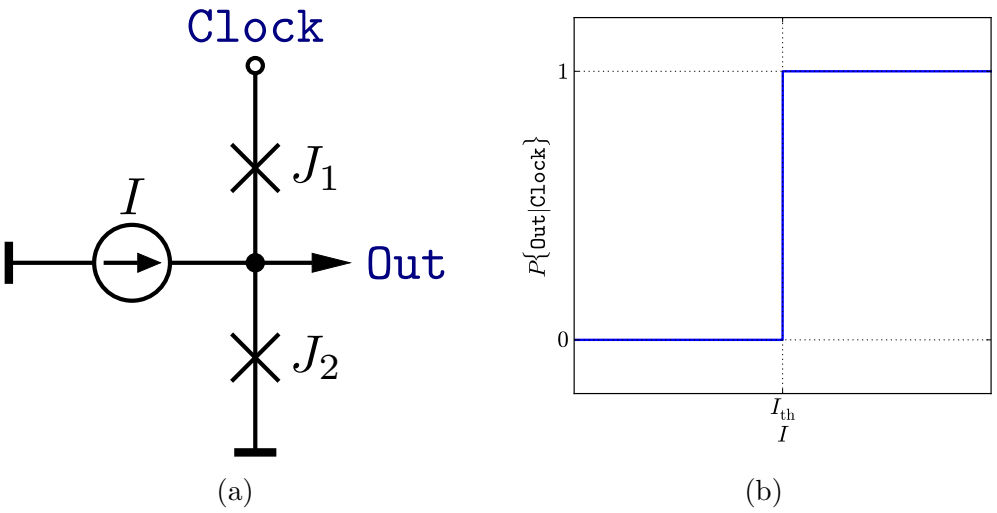

(a) (b)

Figure 4.15: An ideal DMP with threshold I_{th} (a) and the associated switching probability as a function of I (b).

cal circuits of current technology, this channel is often modified through the addition of an external shunt resistor.

The resistive channel exhibits black-body radiation that can be modelled as a noise source, such as I_N in Figure 4.16. The noise of a black-body radiator is white Gaussian, thus the expected value of I_N is zero and the power spectrum $S_{NN}(f)$ is flat across all frequencies, or

$$E[I_N] = 0, \quad S_{NN}(f) = \frac{N_0}{2}. \tag{4.19}$$

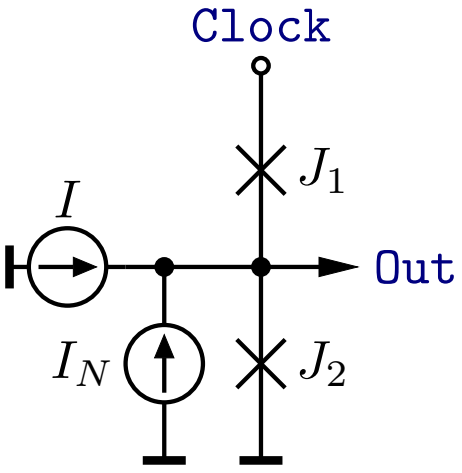

Figure 4.16: A noisy DMP with threshold I_{th} and random noise source I_N.

We now consider the current intentionally applied to the comparator I as a deviation from the comparator threshold,

$$I = I_{\text{th}} - \Delta I,$$

then the current entering the comparator is

$$I_{\text{in}} = I + I_N = I_{\text{th}} - \Delta I + I_N.$$

As the comparator has an interrogation potential of `1` if and only if $I_{\text{in}} > I_{\text{th}}$, or

$$I_{\text{th}} - \Delta I + I_N \geq I_{\text{th}},$$

the comparator will produce an output, in response to a clock event, if and only if, at the time of the event, the samples (i_N,i) of (I_N,I) are such that

$$i_N \geq \Delta i.$$

Note that we here make the assertion that, at threshold, the comparator has an interrogation potential of `1`, not `0`. This is done out of convenience. As these are stochastic quantities, special consideration for equality is not necessary.

We conclude immediately that, if $\Delta i = 0$, or $i = I_{\text{th}}$, the comparator will switch with a probability of 50%, since

$$P\{I_N \geq 0\} = P\{I_N \geq E[I_N]\} = 0.5. \tag{4.20}$$

Since I_N is white Gaussian, with expected value 0, its cumulative distribution function is given by

$$P\{I_N \leq z\} = \Psi_{I_N}(z) = \frac{1}{2}\left[1 + \text{erf}\left(\frac{z}{\sqrt{2\sigma_N^2}}\right)\right], \tag{4.21}$$

where σ_N^2 is the variance of I_N and erf is the error function, defined as

$$\text{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-u^2} du.$$

The variance σ_N^2 for the Gaussian noise signal corresponds to the noise power, which, when including (4.19), is defined as

$$\sigma_N^2 = P_n = N_0 B,$$

where B is the system bandwidth.

We now define the term *switching probability* of the comparator, which refers to the probability of the comparator producing an output SFQ pulse in the event of a clock pulse, $P\{\texttt{Out}|\texttt{Clock}\}$. This corresponds to switching

of the lower junction J_2. Note that this is somewhat of a misnomer, since switching happens in either case in exactly one of the comparator junctions. It is, however, the conventional term for this quantity.

From (4.20) and (4.21), it follows that

$$\begin{aligned} P\{\texttt{Out}|\texttt{Clock}\} &= P\{I_N \geq \Delta i\} = \Psi_{I_N}(\Delta i) \\ &= \frac{1}{2}\left[1 + \operatorname{erf}\left(\frac{\Delta i}{\sqrt{2\sigma_N^2}}\right)\right], \end{aligned} \tag{4.22}$$

which enables us to plot the switching probability against the intentional input current I in Figure 4.17.

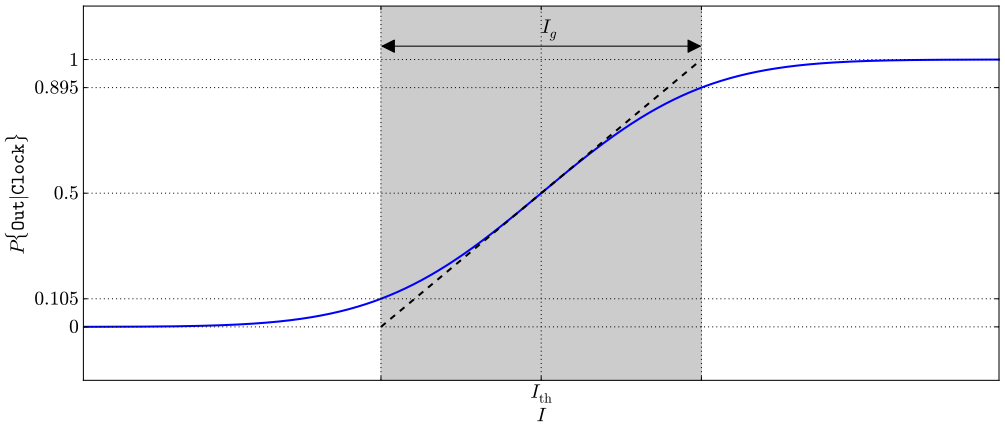

Figure 4.17: Switching probability of a noisy DMP in response to a clock event. The grey patch indicates the grey-zone.

As *grey-zone* of the comparator is known the portion of the abscissa between the $P = 0$ and $P = 1$ intersection of the tangent of the switching probability curve at $I = I_{\text{th}}$ [109]. The width of the grey-zone we term I_g. Henceforth, we will not expressly distinguish between "grey-zone" and "width of grey-zone", as the context should be clear.

Differentiation of (4.22) yields

$$\Psi'_{I_N}(\Delta i) = \frac{1}{\sqrt{2\pi\sigma_N^2}} e^{-(\Delta i)^2},$$

which, evaluated at $\Delta i = 0$ ($I = I_{\text{th}}$), yields the gradient of the tangent

$$\Psi'_{I_N}(0) = \frac{1}{\sqrt{2\pi\sigma_N^2}} = \frac{\Delta P}{I_g}.$$

As, by definition, $\Delta P = 1$, the grey-zone is

$$I_g = \sqrt{2\pi\sigma_N^2} = \sqrt{2\pi P_n}. \tag{4.23}$$

Hence, the grey-zone width grows like the square root of the power of the thermal noise coupling into the decision-making pair. Note, however, that here we assume an instantaneous switching process. This assumption is challenged later in Section 5.4.7.

The effect of grey decisions: grey comparators

We recall from Figure 4.8 that the DMP is employed to discriminate between two partitions of the range of the quantised signal I_q, namely $I_q < I_{\text{th}}$ and $I_q \geq I_{\text{th}}$. When the DMP is grey, threshold jitter is imposed on all the thresholds of the multi-threshold comparator, as illustrated in Figure 4.18.

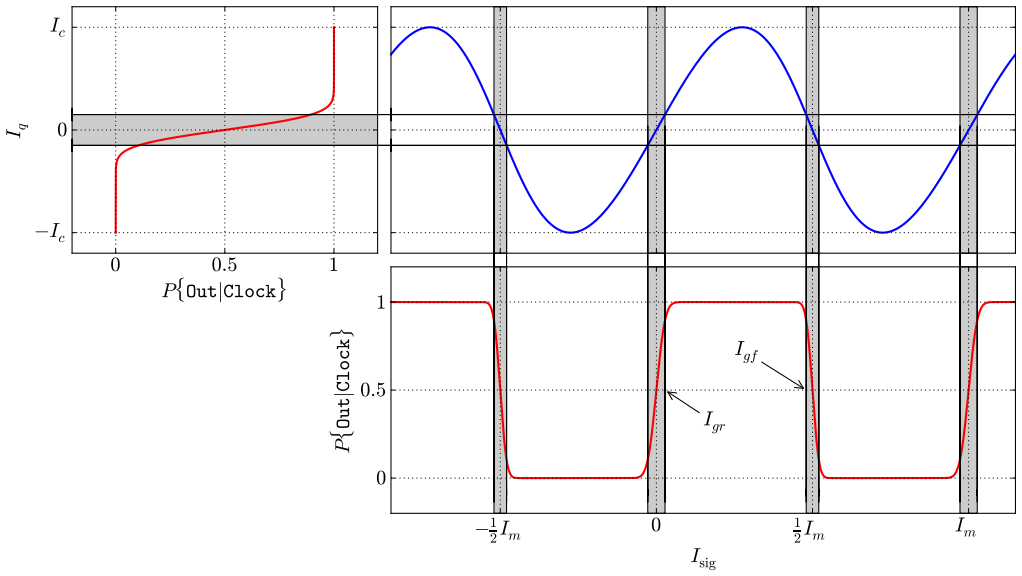

Figure 4.18: A grey DMP imposes grey-zones on every threshold of the QOS-based multi-threshold comparator.

Within the period I_m, the QOS comparator has one rising and one falling threshold. We term the grey-zones associated with these thresholds I_{gr} and I_{gf} respectively. Figure 4.18 suggests that these grey-zones grow as I_g does. To distil this relationship, we start by implicitly deriving $I_q(I_{\text{sig}})$, first defined as I_J in (4.7), obtaining

$$I'_q = \frac{2\pi L I_c}{\Phi_0}(1 - I'_q)\cos\left(\frac{2\pi L}{\Phi_0}(I_{\text{sig}} - I_q)\right).$$

Since the grey-zone applies only in a small zone around the threshold, which implies $I_q \approx 0$, the gradient of I_q at the threshold is

$$I_q'(0) = \beta_L(1 - I_q'(0)) \cos\left(\frac{2\pi L}{\Phi_0} I_{\text{sig}}\right),$$

with

$$\beta_L = \frac{2\pi L I_c}{\Phi_0}. \tag{4.24}$$

Clearly β_L is a non-negative quantity. In real circuits, which require finite non-zero spacing, β_L is strictly positive. Note that the injectivity condition, first stated in (4.5), further imposes

$$\beta_L < 1. \tag{4.25}$$

At threshold,

$$I_{\text{sig}} = \frac{k}{2} I_m = \frac{k}{2} \frac{\Phi_0}{L}, \quad k \in \mathbb{Z},$$

which facilitates

$$I_q'(0) = \beta_L \left(1 - I_q'(0)\right) \cos(k\pi), \quad k \in \mathbb{Z}.$$

This expression further simplifies to

$$I_q'(0) = \pm\beta_L \left(1 - I_q'(0)\right),$$

which enables the solution for $I_q'(0)$, at

$$I_q'(0) = \frac{\beta_L}{1 + \beta_L} \quad \vee \quad I_q'(0) = \frac{\beta_L}{\beta_L - 1}.$$

As β_L is strictly positive and further constrained by (4.25), the two solutions for $I_q'(0)$ include one negative and one positive value, as expected from Figure 4.18.

If we assume that the grey-zone I_g is small, $I_g \ll I_c$, then, around threshold, where the grey-zone applies, the gradient of I_q is linear, which enables the expression of I_{gr} and I_{gf} in terms of I_g as

$$I_{gr} = I_g \left(\frac{1}{\beta_L} + 1\right), \qquad I_{gf} = I_g \left(\frac{1}{\beta_L} - 1\right). \tag{4.26}$$

The effect of grey comparators: discrimination errors

Grey comparators clearly must influence the fidelity of the signal digitised by a real ADC. Also, all other quantities being equal, the effect of the grey-zone should become more pronounced with greater SQNR, as this requires a higher number of bits n and, in turn, feeding feinter signals into grey comparators.

We assume that a well-designed multi-threshold comparator experiences switching only immediately after a clock event, and that such switching is well-defined for all clock events (that is, for every clock event, exactly one comparator junction switches exactly once and does not affect the steady-state of the system). The validity of this assumption will be tested later in Section 5.4.7. Under this assumption, the results of each clock event are independent from all other clock events.

A single-sweep high-level simulation with non-zero I_g yields ADC output against input as depicted in Figure 4.19. Clearly, the error signal is no longer bounded by $\pm\frac{1}{2}\Delta$, suggesting a depression in achievable SNR. This is examined in detail next.

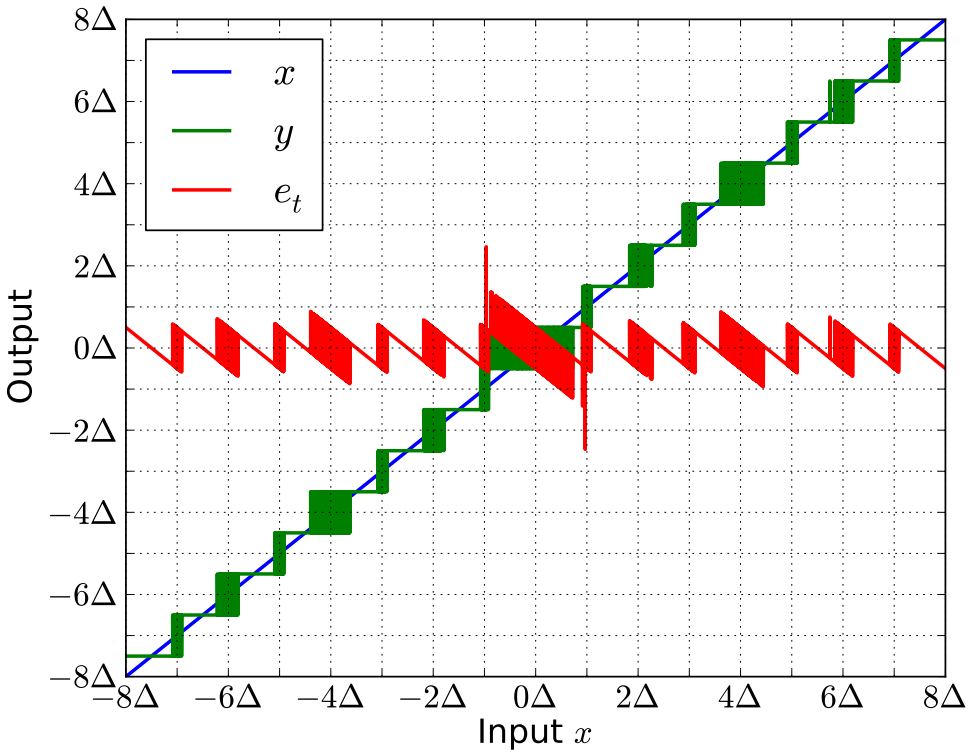

Figure 4.19: Single sweep of $N = 100\ 000$ samples across the full range of a 4-bit flash ADC based on grey QOS comparators.

4.3.4 Achievable SNR for a grey flash ADC

Definition of quantities

A non-zero comparator grey-zone I_g results in depression of achievable SNR for an ADC built from such comparators. Discrimination errors within the multi-threshold comparator add to the quantisation errors inherent to ADC operation, increasing the power of the error signal P_e. The error model from (4.10) is expanded to accommodate this modification, resulting in

$$y[k] = x[k] + e_t[k] = x[k] + e_q[k] + e_d[k],$$

where $e_t[k]$ is the total error and $e_d[k]$ is the discrimination error. The quantisation error $e_q[k]$ remains as defined in (4.10). Note that $x_q[k]$ is no longer simply a function of $x[k]$ but depends also on the comparators' response to the noise injected locally at each comparator. Discrimination errors result in incorrect bits of the output word, which correspond to an error that must be an integer multiple m of Δ, where m is confined to $[-M, M]$, or

$$e_d[k] = m\Delta, \qquad m \in \{0, \pm 1, \pm 2, \cdots, \pm M\}.$$

Composition of the error signal

To facilitate brevity, we start from the simulation depicted in Figure 4.19. Sorting the total error e_t into bins, the histogram in Figure 4.20 is constructed.

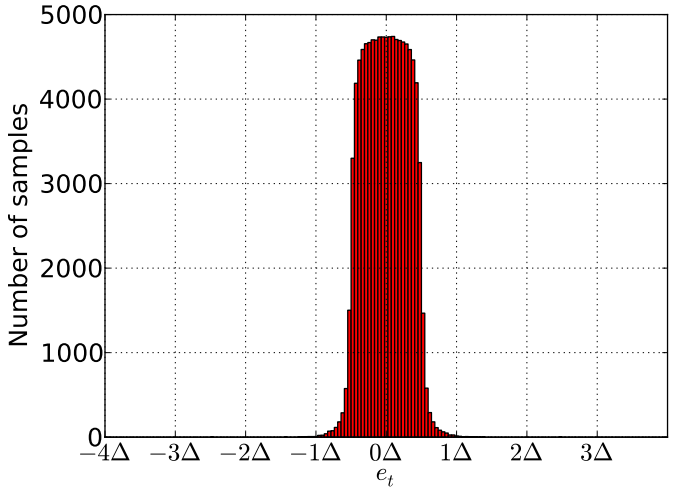

Figure 4.20: Distribution of total error e_t for the results from Figure 4.19.

To further illustrate the error contribution, the ideal output signal $x_q[k] = Q_r(x[k])$ is constructed and, from this, the quantisation error $e_q[k]$. This enables the neat separation of the two components of $e_t[k]$ into $e_q[k]$ and

$$e_d[k] = e_t[k] - e_q[k].$$

The appropriately constructed histograms from the signal samples from Figure 4.19 are depicted in Figure 4.21.

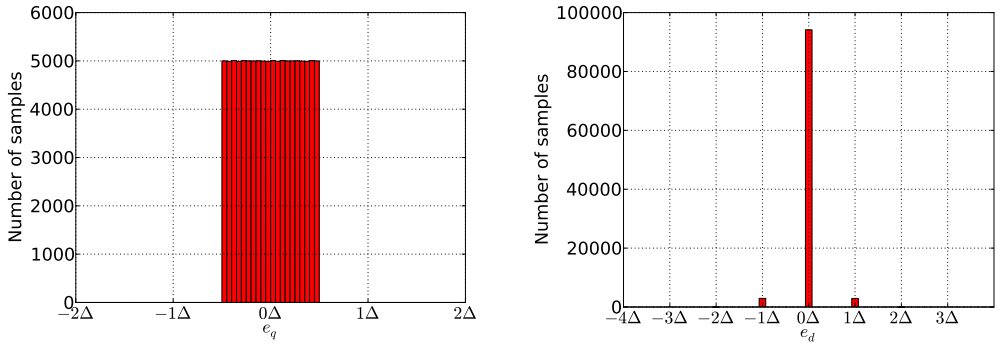

(a) (b)

Figure 4.21: Distributions of the components of the error signal confirm that e_q is still uniformly distributed (a) and e_d is discrete (b). This is for $N = 100\,000$ samples and $n = 4$ bits.

To better understand qualitatively the error signal distribution for an ADC built from grey comparators, we repeat the experiment for $n = 7$ bits and 4 times the I_g, with results shown in Figure 4.22.

To obtain an analytical expression for the SNR achievable by a grey ADC, an expression for the power of the error signal, $P_e = E[e_t^2]$ would be required. Assuming that e_t has a zero mean, this reduces to

$$P_e = \sigma_{e_t^2} = \sigma_{(e_q+e_d)^2} = \sigma_{e_q}^2 + \sigma_{e_d}^2 + 2\sigma_{e_q e_d}, \tag{4.27}$$

where $\sigma_{e_q e_d}$ is the *statistical covariance* of e_d and e_t. When e_d and e_t are independent, $\sigma_{e_d e_t} = 0$.

Unfortunately, an inspection of Figures 4.21 and 4.22 suggests that the components e_d and e_q are not independent. A necessary condition for statistical independence of two random variables is that the probability distribution of the sum of the two variables is equal to the convolution of the probability distributions of the two individual variables. This is clearly not the case here, as the convolution of a rectangularly-shaped function with any function consisting only of Dirac impulses would result in a distribution made from several adjacent (and perhaps superimposed) rectangular shapes.

To confirm this, the covariance matrices for the sets of samples from e_q and e_d are found for both experiments (using the `cov` function in numpy [103]):

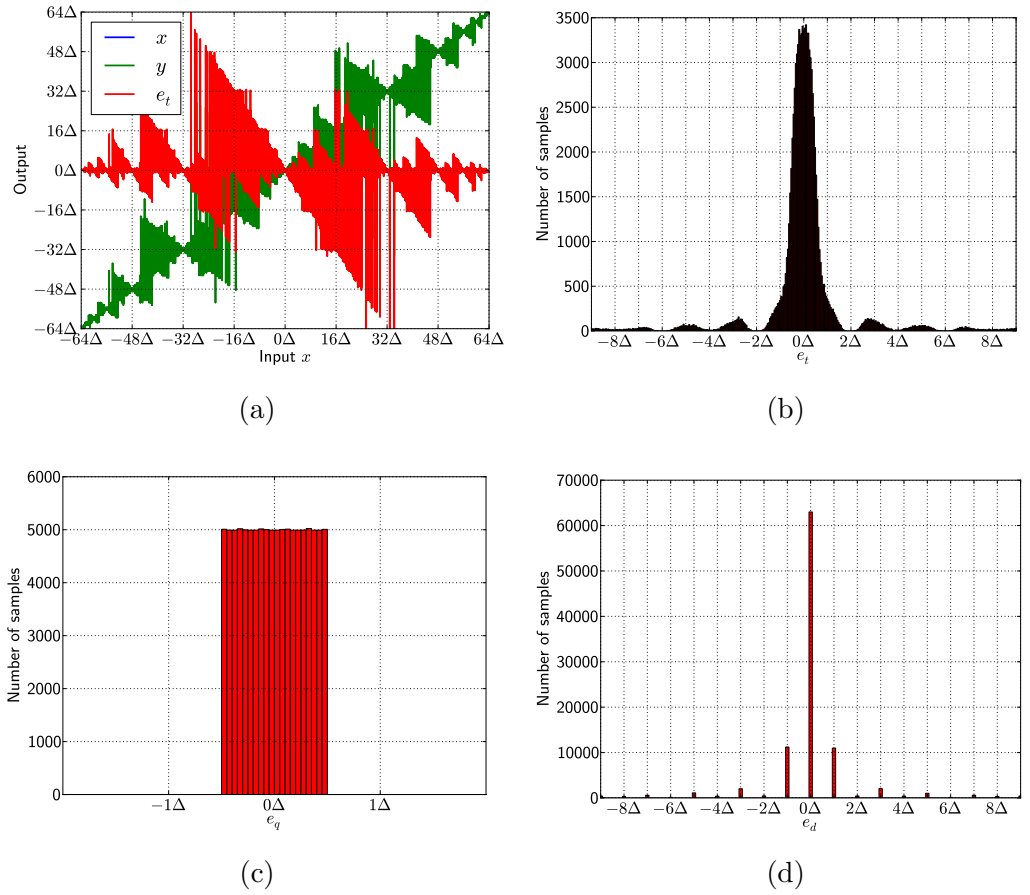

Figure 4.22: Components of the error signal for $N = 100\ 000$ samples and $n = 7$ bits.

$$(\Sigma_{e_q e_d})_{n=4} = \begin{bmatrix} 0.0833\Delta^2 & -0.813\Delta\sigma_N \\ -0.813\Delta\sigma_N & 80\sigma_N^2 \end{bmatrix},$$
$$(\Sigma_{e_q e_d})_{n=7} = \begin{bmatrix} 0.0833\Delta^2 & -0.965\Delta\sigma_N \\ -0.965\Delta\sigma_N & 3480\sigma_N^2 \end{bmatrix}$$

We stress that the matrices obtained here result only from one experiment each can not necessarily be attributed a high-confidence value. However, the covariance of e_q and e_d, shown in the off-diagonal entries of $\Sigma_{e_q e_d}$, is clearly non-zero and significant, confirming that e_q and e_d are not statistically independent. Hence, all three terms on the right side of (4.27) are significant and must be accounted for for in any prediction of P_e.

Although finding an approach to solving for $\sigma_{e_d}^2$ and $\sigma_{e_q e_d}$ analytically is certainly possible, a closed-form solution is probably not possible (for example, no closed-form solution for erf is known). We have found that an analytical

approach would involve substantial numerical computation. For this reason, we deem an analytical approach to predicting P_e unsuitable and prefer to use statistical methods for the prediction of SNR for grey ADCs.

SNR prediction of grey ADCs by statistical methods

The advantages of statistical methods lie mainly in the flexibility they enable, whereas the disadvantages lie mainly in the computational effort required to achieve arbitrary confidence levels.

The quantities of interest are the statistical random variables Y and E_t, corresponding to the discrete-in-value output of the grey ADC and the associated total error referred to the random input signal X respectively.

An arbitrary number of samples Y and E_t can be generated for an arbitrary distribution of X. The program employed for this was implemented in *python* [110], using the popular *scipy* [103] library in the *pylab* environment. A diagram of the program `adcn_sample` is shown in Figure 4.23. The code is reproduced in Listing B.3.

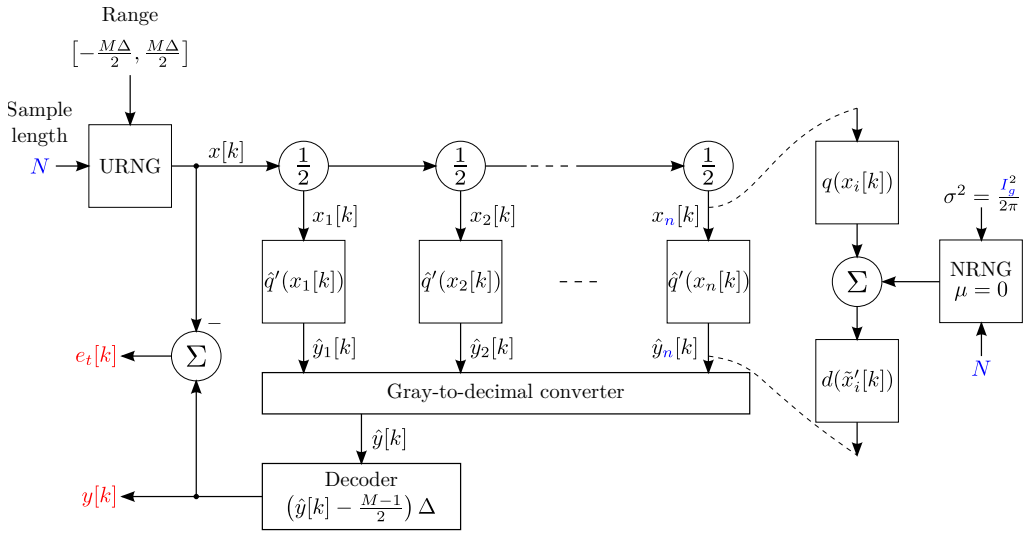

Figure 4.23: Program `adcn_sample`, an ADC sample generator for statistical evaluation of ADC performance. The input paramters (n, N, I_c, I_g, Δ) are the number of comparators, number of samples, critical current of the quantisation loop, absolute grey-zone and LSB respectively. Output signals $(y[k], e_t[k])$ are discrete signals of length N representing sets of random samples of the ADC output and total digitisation error respectively. URNG and NRNG refer to uniform and normal random number generators respectively. Note that every comparator $\hat{q}'(\cdot)$ has an independent NRNG.

From the input sample length N, the number of quantisation levels $M = 2^n$ and the LSB Δ, a uniform random number generator generates the input vector $x[k]$ of length N. The input vector is divided by powers of two to generate n input vectors $x_i[k]$, which are to be fed to each multi-threshold comparator as described in Figure 4.10. The quantisation operation $q(\cdot)$ is implemented in function `quant`, which takes as parameters L, I_c and the vector of input samples $x_i[k]$ and provides as output the vector of quantised samples $\tilde{x}_i[k]$. The noisy discrimination operation $d'(\cdot)$ is implemented in function `ndisc`, which takes as parameters I_g and the vector of quantised samples $\tilde{x}_i[k]$ and provides as output the bit vector $\hat{y}_i[k]$. The function `Gray2dec` takes as parameters all bit vectors $\hat{y}_i[k]$ and provides as output the coded output vector $\hat{y}[k]$. The function `decode` takes as parameters the coded vector $\hat{y}[k]$, the number of bits n and the LSB Δ and provides as output the vector of digitised samples $y[k]$. Finally, the total error vector $e_t[k]$ is generated from the input vector $x[k]$ and the output vector $y[k]$ through subtraction.

To calculate the SNR achieved during one experiment, the quantities P_y and P_e must be determined. This is achieved respectively by the formulas

$$P_y = \frac{1}{N}\sum_{i=1}^{N} y^2[i], \quad \text{and} \quad P_e = \frac{1}{N}\sum_{i=1}^{N} e_t^2[i],$$

from which, in turn, the achievable SNR is estimated as

$$\widehat{\text{SNR}} = \frac{P_y}{P_e}.$$

A Monte Carlo simulation involves repeating this estimate a number of times and so obtaining an impression of its distribution. In this way, a confidence interval for the true achievable SNR can be obtained. This is done in program `snr_estimator`, which relies on `adcn_sample` to generate the observations.

Program `snr_estimator` takes as input a tuple of ADC parameters

$$\theta = (n, \Delta, I_c, I_g),$$

and a tuple of methodological parameters

$$\psi = (N, \alpha, \beta),$$

where $1 - \alpha$ is the desired confidence with which the true achievable SNR lies within $(1 - \beta, 1 + \beta)$ multiplied by the estimate. That is, if $\alpha = 0.05$ and $\beta = 0.01$, then a $\pm 1\%$ of the estimate interval around the estimate will contain the true SNR with a confidence of 95%.

To discuss how to obtain the confidence interval, we first investigate the distribution of estimator $\widehat{\text{SNR}}$. A histogram of the estimates $\widehat{\text{SNR}}$ achieved from 10 000 repetitions of `adc_sample` (with an arbitrary but practical example

parameter set) is shown in Figure 4.24a. The distribution does not look normal, as a longer tail appears to the right of the bulk than to the left. To obtain a normal distribution, a logarithmic transformation seems appropriate. Conveniently, SNR values are generally stated in dB, which suggests the transformation

$$\mathtt{SNR}_{\mathrm{dB}} = 10 \log_{10} \mathtt{SNR}.$$

Indeed, the distribution after this logarithmic transformation, as illustrated in Figure 4.24b, looks approximately normal.

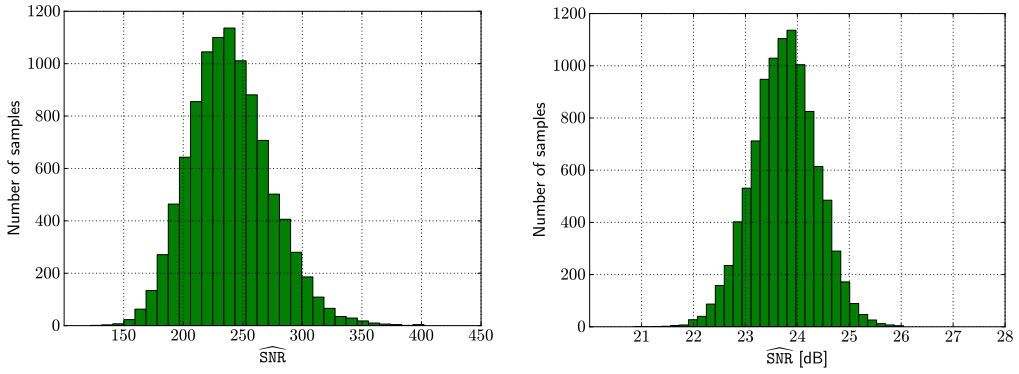

(a) (b)

Figure 4.24: Distribution of $\widehat{\mathtt{SNR}}$ before (a) and after its transformation to the dB scale (b).

Without submitting a rigorous proof here, we believe that the distribution of the transformed estimator is close enough to normal to enable determination of statistical confidence by the well-known formula predicting confidence intervals for Monte Carlo analyses,

$$P\left(\hat{\mu} - \frac{z_{\alpha/2}\hat{\sigma}}{\sqrt{R}} \leq \mu \leq \hat{\mu} + \frac{z_{\alpha/2}\hat{\sigma}}{\sqrt{R}}\right) = 1 - \alpha, \tag{4.28}$$

where $\hat{\mu}$ is the mean of all estimator samples $\widehat{\mathtt{SNR}}$, μ is the actual achievable $\mathtt{SNR}$, $\hat{\sigma}$ is the square root of the variance of the estimator samples, $z_{\alpha/2}$ is the quantile of $\alpha/2$, and $R \in \mathbb{Z}$ is the number of repetitions. Note that $z_{\alpha/2}$ is obtained by inverting the cumulative distribution function of a normally distributed variable with zero mean and unity variance:

$$z_{\alpha/2} = \left|\Psi^{-1}\left(\frac{\alpha}{2}\right)\right| = \left|\sqrt{2}\,\mathrm{erf}^{-1}(\alpha - 1)\right|.$$

The code for program `snr_estimator` can be found in Listing B.4, but its operation is summarised in Algorithm 1.

Algorithm 1 `snr_estimator`

empty list `SNRdBs`
$z \leftarrow |\Psi^{-1}(\alpha/2)|$
$R \leftarrow 0$
repeat
  run `adcn_sample`; $R \leftarrow R + 1$
  find $\widehat{\text{SNR}}$ in dB and append it to `SNRdBs`
  $\hat{\mu} \leftarrow$ `mean(SNRdBs)`
  $\hat{\sigma} \leftarrow \sqrt{\texttt{var(SNRdBs)}}$
  $\epsilon \leftarrow z\hat{\sigma}/\sqrt{R}$
until $\epsilon \leq \beta\hat{\mu}$
return $\hat{\mu}$

To test a part of this algorithm, we compare its output to the predicted value for an ideal flash ADC, (4.11). The results of this experiment are shown in Figure 4.25 and mutually support the correctness of (4.11) and the devised algorithm.

4.3.5 ADC performance

Several key performance indicators for ADCs were discussed in Section 2.4.1. Here we concentrate on two: achievable SNR and sensitivity. As sensitivity ν is defined the span of all ADC values, or

$$\nu = 2^n \Delta. \tag{4.29}$$

The sensitivity is a measure of the signal power required to operate an ADC optimally. Note that, in our case, the quantity ν is an electrical current and is thus measured in A. For different purposes, one might consider an alternate definition proportional to A^2.

It is desirable from a power perspective to obtain as low a value for ν as possible.

Effect of I_g on SNR

With the new ability to predict achievable SNR over the parameter space spanned by θ, we first investigate the effect of the absolute value of the grey-zone. For this, we make some practical choices for other parameters, which are summarised in Table 4.3.

The results of the experiment are shown in Figure 4.26. Note that the SNR values reported are conservative as they are the lower bound of the confidence interval defined by $\alpha = 0.01$, $\beta = 0.01$.

It is clear that even moderate grey-zones can cause significant depression of achievable SNR. A useful boundary to define is the 3 dB depression point, as

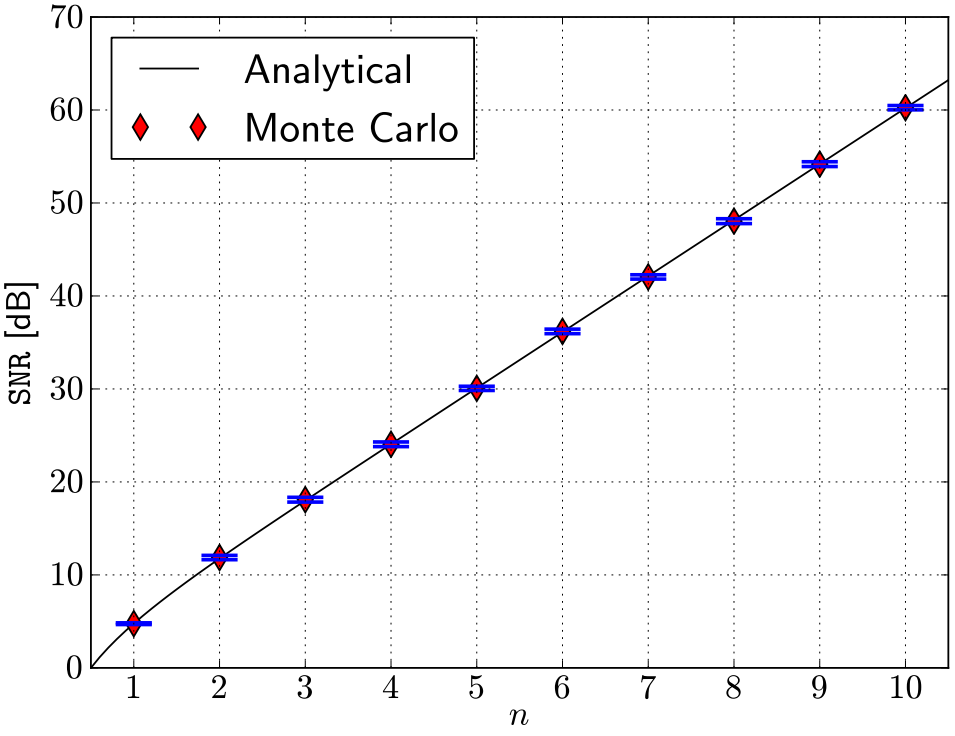

Figure 4.25: Output of `snr_estimator`, compared against (4.11), for $I_g = 0$, $\alpha = 0.001$, $\beta = 0.02$ (and other arbitrary practical parameters). The error bars indicate the interval (maximum 2% away from the plotted point) within which the true achievable SNR lies with 99.9% confidence.

Table 4.3: Parameter choices for I_g experiment.

Parameter	**Value**	**Comment**
n	$\{2, 4, 7\}$	Practical and optimistic ADCs [61].
I_c	$125\,\mu A$	Design rules of real processes.
Δ	$2\,mA$	Physically achievable values for L.
I_g	$[0, 20\,\mu A]$	From ideal to very grey.

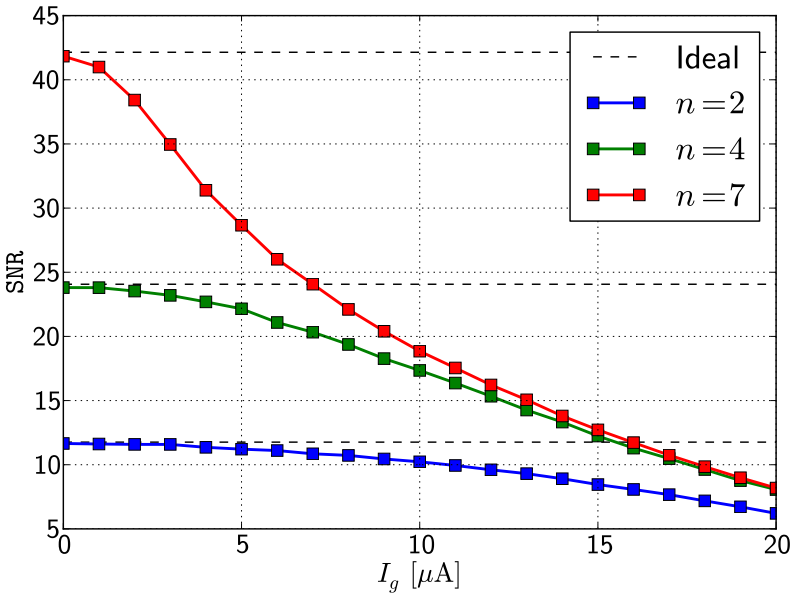

Figure 4.26: Simulation of relationship between I_g and conservatively predicted SNR for several values of n.

this is where the ADC has lost half an effective bit. Above this boundary, one could claim that the achievable SNR is still closer to its ideal achievable SNR than to the ideal achievable SNR of an ADC with fewer bits. We thus define the quantity $(I_g)_{3\,\mathrm{dB}}$ as the DMP grey-zone at which the 3 dB depression point is reached.

By employing a binary search algorithm, the value of $(I_g)_{3\,\mathrm{dB}}$ is found for several values of n, with other parameters the same as in Table 4.3. The results of this experiment are shown in Figure 4.27.

Addressing the MSB Catastrophe

Note that a critical component is the comparator synthesising the most significant bit, `bit0`, as this comparator is fed the lowest signal amplitude and thus its grey-zone is "smeared out" across the widest range of the input signal. Over the full ADC input range, every comparator except the MSB comparator have an even number of transitions, whereas the MSB comparator has exactly one transition. In the conventional coding method introduced in Figure 4.10, the single transition of the MSB comparator is across a rising threshold. This is illustrated in Figure 4.28a.

From (4.26) it is known that, for practical values of β_L, the falling threshold grey-zone is always narrower than the rising threshold grey-zone. Hence, the single transition of the MSB comparator should be across a falling threshold. By injecting an offset current into each comparator, this can easily be

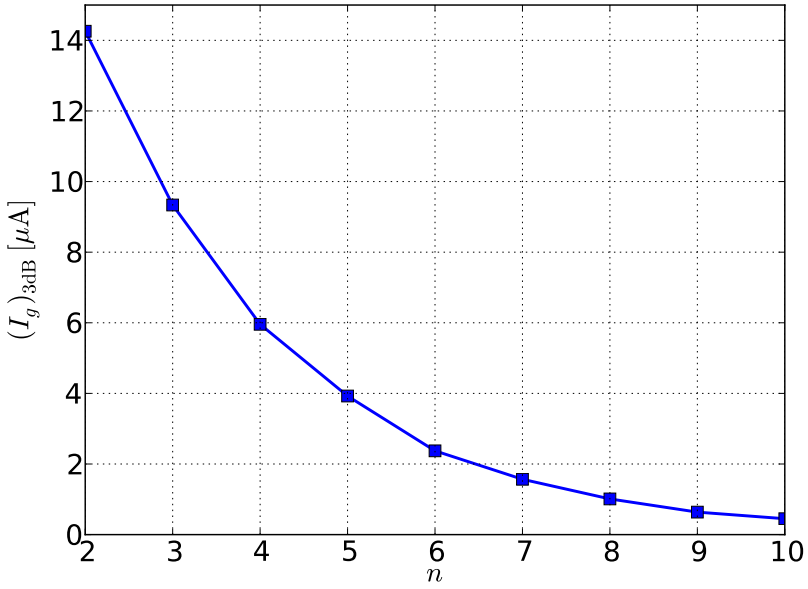

Figure 4.27: The DMP grey-zone which depresses by 3 dB the achievable SNR of a QOS-based n bit superconducting flash ADC.

achieved, resulting in the output shown in Figure 4.28b. Here the output is in inverted Gray code, which must naturally be respected when decoding the output signal.

To investigate the effect of this improved method of biasing, the 3 dB depression point is re-calculated for the ADC parameters investigated previously. Note that, for these parameters, $\beta_L \approx 0.2$, which means that $I_{gf} \approx \frac{2}{3} I_{gr}$. For this modest improvement, the achieved 3 dB depression points are noticeably better than for the conventional architecture, as shown in Figure 4.29.

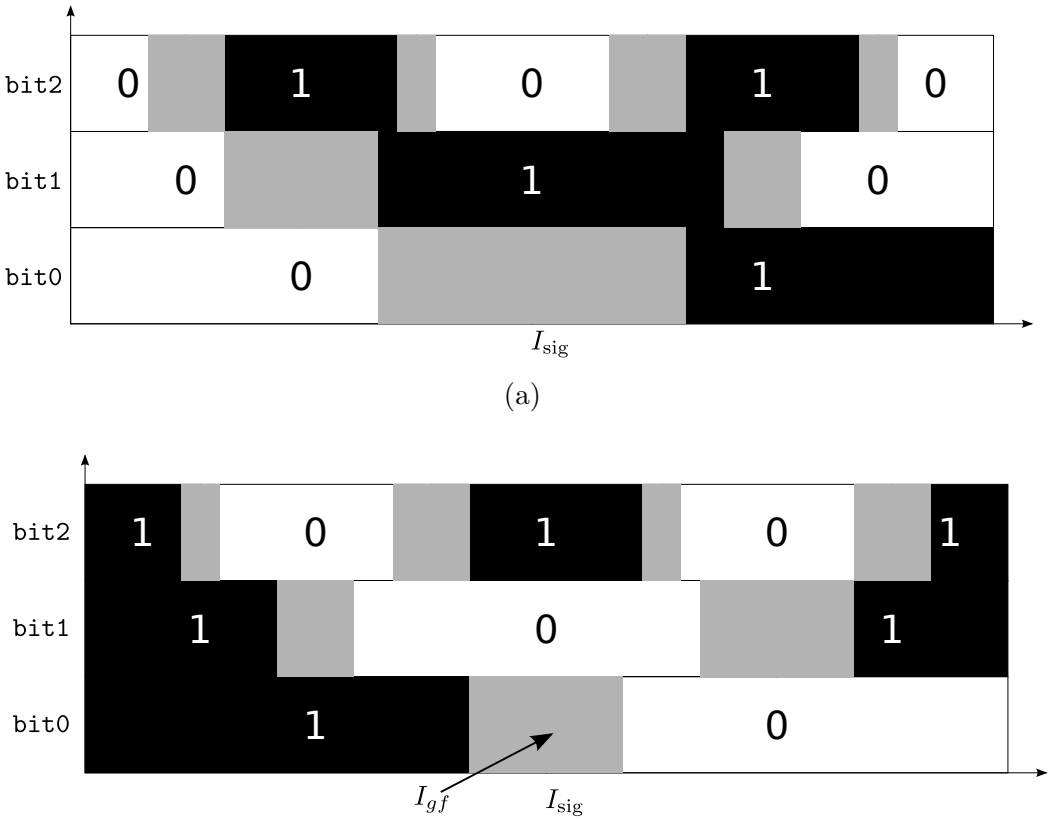

Figure 4.28: Conventional (a) and improved (b) methods of biasing the comparators in a superconducting flash ADC (demonstrated here for $n = 3$). The improved method exhibits a lower grey-zone across the MSB transition as the conventional method.

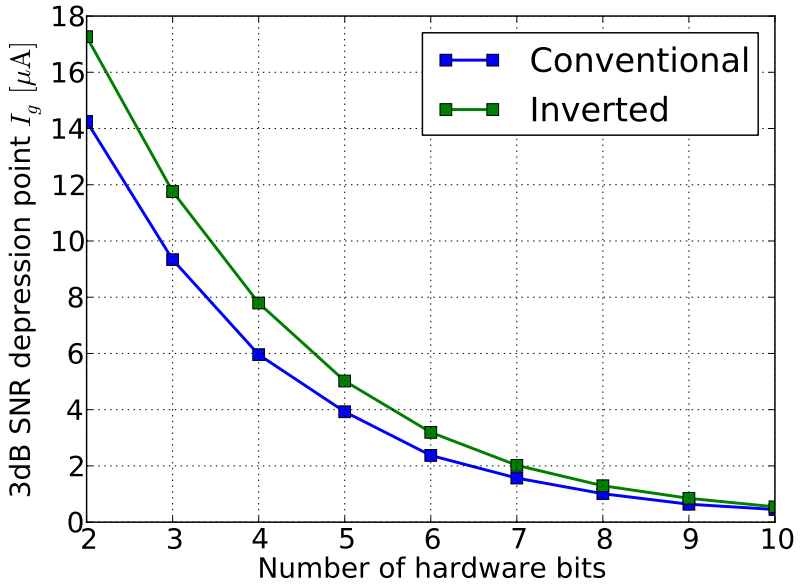

Figure 4.29: Results of the repeated experiment for the improved architecture, compared to the conventional method.

The price of ν

A low ν requires a low Δ for any given n. Also, from (4.29) we know that Δ decreases exponentially with increasing n for any given ν.

Given (4.29), (4.12) and (4.24), it follows that

$$\nu = \frac{I_m}{2}2^n = \frac{\Phi_0}{L}2^{n-1} = \frac{\pi I_c}{\beta_L}2^n.$$

As $\beta_L \in (0, 1)$, we establish a lower bound on ν

$$\nu_{\min} = \pi I_c 2^n. \tag{4.30}$$

This is an important limit, as small I_c values are contingent on a well-defined fabrication process. For a popular contemporary RSFQ process [48], the minimum achievable I_c is 125 µA. Hence, the sensitivity ν of, for example, a 4-bit flash ADC fabricated in this process is bounded from below by approximately 6.28 mA.

While a low sensitivity is certainly desirable, the effect of ν on achievable SNR should be understood. That such an effect is probably not negligible follows from (4.29) and (4.26), which link ν to the effective grey-zones of the ADC transitions.

The DMP grey-zone I_g is an absolute quantity, the result of thermal noise, and does not change significantly with ADC parameters θ. A decrease in ν, however, is associated with a decrease in signal power, suggesting that the grey-zone becomes more significant.

The effective grey-zones of the QOS comparator for different values of β_L, from (4.26), are shown in Figure 4.30. Thus, with increasing β_L, or signal power, the effective grey-zones of the comparator transitions decrease.

Considering Figure 4.28, however, it is clear that it is not the absolute value of the effective grey-zones that is significant, but rather the proportion of the traversed signal range that is grey. We term this proportion the comparator *greyness* γ and define it as

$$\gamma = \frac{I_e}{I_m}, \tag{4.31}$$

where

$$I_e = \frac{I_{gr} + I_{gf}}{2}$$

is the average effective grey-zone across all comparator thresholds.

Note that

$$I_e = \frac{I_{gr} + I_{gf}}{2} = \frac{(\frac{1}{\beta_L} + 1)I_g + (\frac{1}{\beta_L} - 1)I_g}{2} = \frac{1}{\beta_L}I_g,$$

which reduces (4.31) to

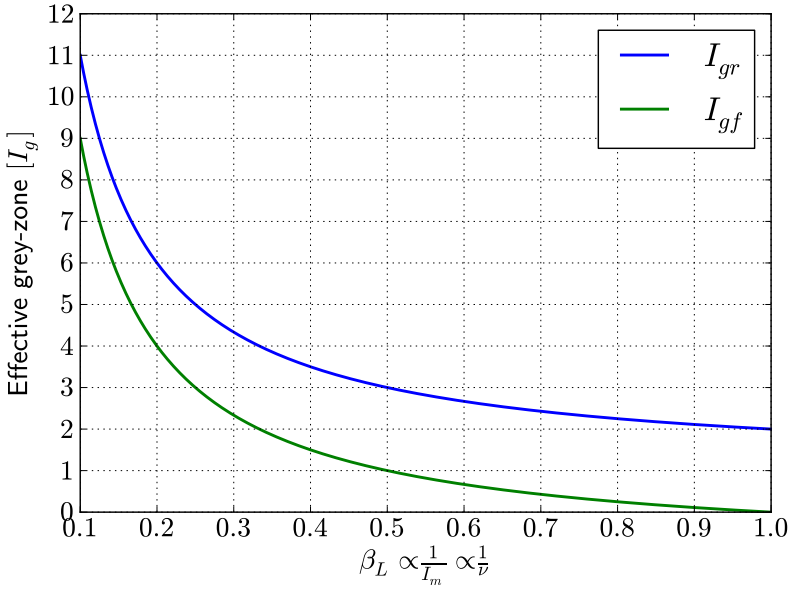

Figure 4.30: Effective grey-zones of comparator transitions varying with β_L.

$$\gamma = \frac{1}{\beta_L}\frac{I_g}{I_m}.$$

Using (4.8) and (4.24), we conclude that $I_m = \frac{2\pi I_c}{\beta_L}$ and, hence,

$$\gamma = \frac{I_g}{2\pi I_c}.$$

Hence, $\nu_{\min} \propto I_c$ and $\gamma \propto \frac{1}{I_c}$, which implies that a lower $\nu_{\min}$ requires a higher γ. However, for a given I_c, γ is independent of I_m and independent of ν. This is an important insight, as it suggests that once I_c is fixed, minimal ν is not detrimental to achievable SNR.

To support this suggestion, we employ our statistical program to determine the relationship between SNR and ν for various configurations. The results of this simulation are shown in Figure 4.31. Note that we employ our new, optimised biasing and coding method shown in Figure 4.28b.

The response of achievable SNR to ν is not entirely flat for all configurations. For large values of I_g (beyond the 3 dB depression point), a low maximum is visible near $\nu \approx 2\nu_{\min}$. However, within the band of interest, above the 3 dB depression point, we are confident that, all other parameters being equal, the choice of ν does not significantly affect achievable SNR.

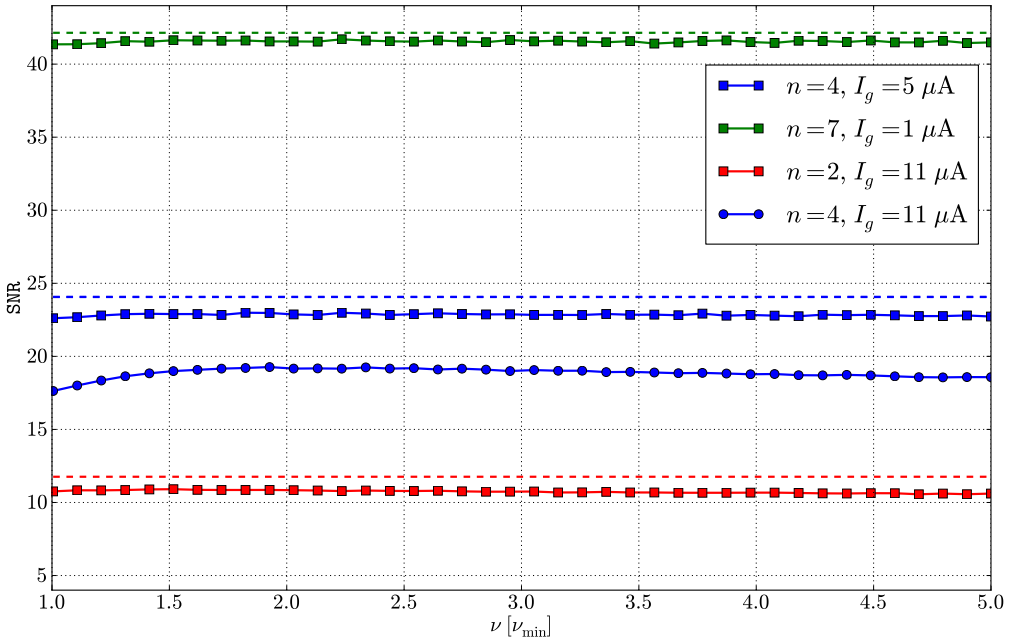

Figure 4.31: The achievable SNR as ν varies.

4.4 Chapter summary and conclusion

In this chapter, a model of the n-bit ideal and grey flash ADC, based on the QOS-comparator, was developed. We summarise our related and performed work and obtained insights as follows:

1. Subject to certain parameter constraints, a one-junction-SQUID can be employed to quantise a current, having an almost sinusoidal response w.r.t. the input signal (a "sawtooth" distortion of the sinusoid is apparent).

2. Combining this quantising circuit with a balanced Josephson comparator yields a quasi-one-junction SQUID multi-threshold comparator that, properly biased and clocked, divides the input current into alternating bins of equal width, labeled `1` and `0`.

3. Employing n parallel identical comparators to quantise chunks of the input signal divided by a $R/2R$ ladder yields an n-bit ADC with Gray-coded output. An ADC of this kind is called a *flash ADC* and belongs to the Nyquist-rate ADC family.

4. Such an ADC is limited in bandwidth by the maximum clock frequency of the comparators, and limited in signal-to-quantisation-noise ratio (SQNR) by n.

5. Thermal noise must be considered, even at the low temperatures required for superconductivity. Noise sources impose *threshold jitter* on the balanced Josephson comparator, resulting in a significant *grey-zone* around each threshold.

6. The comparator grey-zone significantly depresses ADC SNR. For a 4-bit ADC, a grey-zone of $\sim 10\,\mu A$ can depress SNR beyond the important 3 dB-depression point.

7. Minimizing the effect of the grey-zone is more difficult as n increases. We have devised a flexible model that is able to predict the depression of ADC SNR given a certain grey-zone.

8. We have devised a new architecture for superconducting flash ADCs that can improve the grey-zone permissible.

Conclusion: A superconducting flash ADC as proposed in this chapter should indeed be capable of digitising an astronomical signal, that is, a signal that might be received by the SKA. The maximum signal bandwidth will be limited by the comparator clock frequency (according to Nyquist's sampling theorem [24]), whereas the signal-to-noise ratio (or dynamic range) will be limited by the number of comparators n and, crucially, the grey-zone achievable.

4.5 Chapter research output

Research outputs resulting from work done in this chapter are summarised at the end of Chapter 5 (Section 5.8), as they are best viewed in context of the work presented in that chapter.

Chapter 5

ADC Design

5.1 Introduction

The last chapter related our efforts at developing a high-level mathematical model of a superconducting flash ADC, and to investigate how this model responds to certain parameters. A particularly important quantity at the comparator level was identified: the comparator grey-zone I_g, a result predominantly of thermal noise.

In this chapter we apply our insights to design comparators at the circuit level and verify their operation through simulation and experimental verification. The work related in this chapter has the following objectives:

1. Design multi-threshold comparators at the circuit level and verify their operation through simulations,
2. lay out several comparators and re-simulate their extracted circuits to obtain an understanding of their working the in the presence of parasitic components,
3. design and carry out a cryogenic experiment to measure the grey-zone or an equivalent quantity of real comparators fabricated from our designs,
4. use measured quantities as well as our previous model to predict achievable performance metrics of ADCs built from such comparators, and
5. integrate the newly found or confirmed knowledge to generate a comparator design optimal for our application.

The multi-threshold comparator is a key component of any superconducting flash ADC. Identifying a suitable, realisable comparator design will be a significant step towards fabricating a full ADC for the SKA.

Throughout this chapter, we will be designing our circuits for the $J_c = 1\,\mathrm{kA/cm^2}$ RSFQ1G process [48] supplied by the Institute of Photonic Technology in Jena, Germany.

5.2 Designing a QOS comparator

This section deals with a design of the multi-threshold comparator required for a superconducting flash ADC. Such a comparator consists of a quantising circuit, essentially mapping the input signal into the correct "bin" (`1` or `0`), as well as a discriminator, which detects the bin into which the input signal has been mapped.

A high-level model of such a multi-threshold comparator based on the QOS quantiser was investigated extensively in Chapter 4. One might conclude that the next stage in the design would be a circuit-level investigation of possible QOS comparators. This is valuable, but to design a realisable circuit, the fabrication process must be understood and respected. It is for that reason that we follow the following design procedure for the multi-threshold comparator:

1. Design and lay out a nominal single-threshold comparator (STC), which, or a version of which, will eventually be placed behind the QOS quantiser.
2. Extract circuit parameters from the STC and thus develop and understanding of the properties of the process. Estimate the parasitic load that would stress the QOS quantiser.
3. Design the quantiser in the presence of the parasitic load.
4. Lay out the quantiser, attach it to an STC to form an MTC.
5. Extract and optimise the MTC cell.

General design procedures and the software toolset employed to design and lay out these circuits were described in detail in Chapter 3.

5.2.1 Discrimination circuit: the single-threshold comparator

In SFQ electronics, a single-threshold comparator is easily realised by a two-junction decision-making pair (DMP). This concept was first introduced in Section 2.3.3. In addition to the DMP that performs the comparator function, a comparator needs to provide a means to be clocked and output its data.

Clocking a comparator is straightforward and follows from the functionality of the DMP. However, (deterministic) switching errors can occur when the requirements of surrounding circuitry are not taken into account. This is true for any cell (see Section 2.3.3), and for this reason the input and output of the comparator should receive careful attention as well.

The chosen circuit topology for the single-threshold comparator is depicted in Figure 5.1. This topology is fairly well-known [9], although details are the subject of some current research [111, 112, 113]. The balanced Josephson

comparator is made from junctions B_1 and B_2. Junction B_c shapes the clock pulse arriving at `Clk` before applying it to the comparator, whereas B_s shapes the output pulse produced by B_2 before transmitting it on to `Out`.

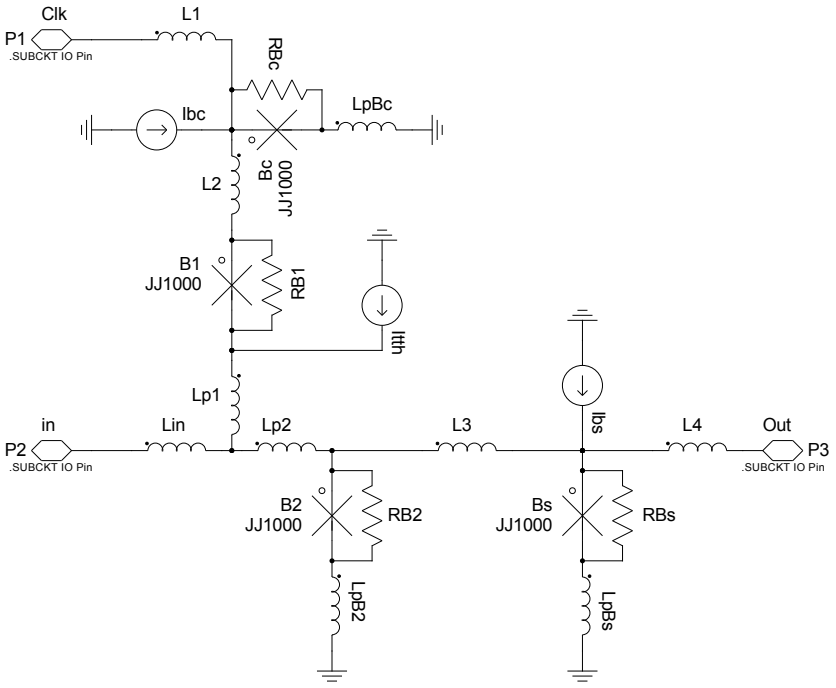

Figure 5.1: The single-threshold comparator.

This comparator discriminates a current signal fed into `in`. The current I_{tth} is used to modify the threshold of the comparator. For the ADC it is useful if the threshold is at $I_{\text{th}} = 0$ in the full configuration. Although this could potentially be achieved by design in the absence of I_{tth}, manufacturing tolerances require that this threshold be *in situ* modifiable.

The corresponding layout for the single-threshold comparator is depicted in Figure 5.2. The extracted values for the layout are summarised in Table 5.1. Note that extraction was performed with InductEx [84], with the supplied process data `i1k.ldf` [114].

From Table 5.1 we draw several conclusions.

First, while generally undesirable, parasitics cannot be avoided. The metal connection between the junction active area and ground plane, for example, creates or increases a current loop and an associated (parasitic) inductance. The area of this loop can be minimised, but its lower bound is constrained by design rules and process dimensions.

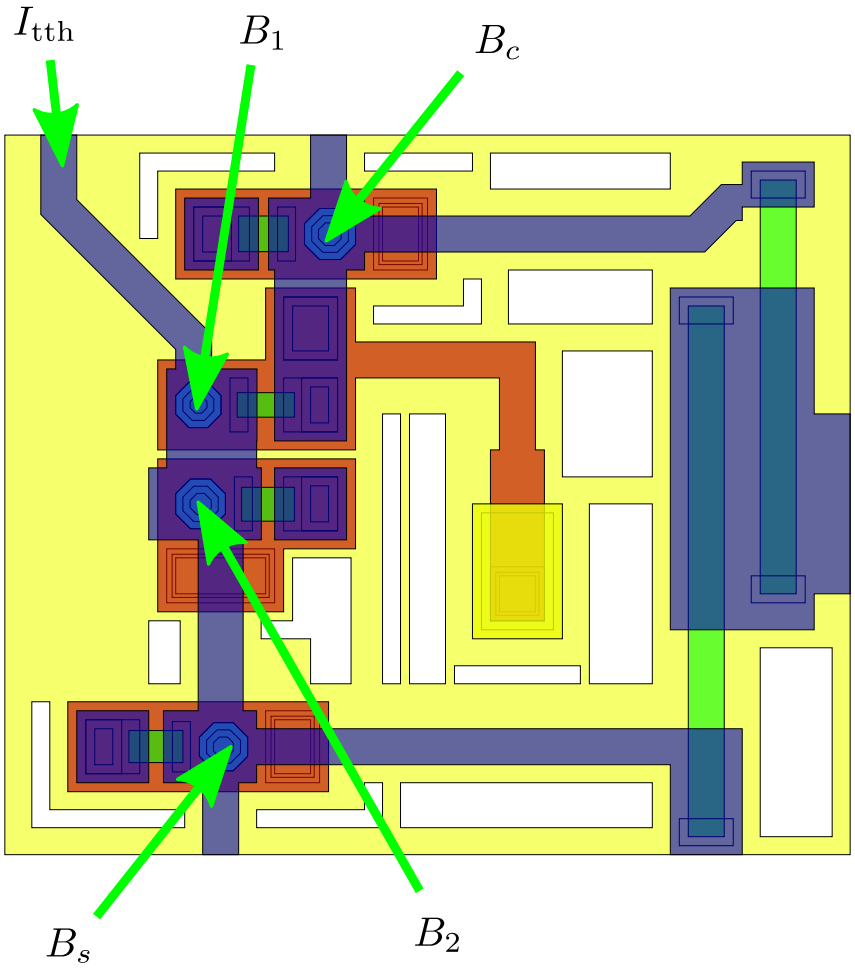

Figure 5.2: An example single-threshold comparator layout.

In all processes available commercially [12, 48] and known to us [115, 116], the decision-making pair made from B_1 and B_2, for example, can only take two substantially different forms. These are demonstrated in Figure 5.3.

The crucial difference between the two designs is the direction of current through the junctions. Whereas Josephson junctions are entirely symmetrical w.r.t. their terminals (and this symmetry is substantially preserved in real junctions), the parasitic inductances characterising the different directions are substantially different.

For Option 1, the "bridge" between the two junctions occurs in Layer M2, which is the layer furthest away from the ground plane. The loop area of Loop 1 is thus comparatively large. The circuit element L_{p1} corresponds most closely to this bridge, and thus exhibits a large inductance.

Option 2 allows the bridge to occur in Layer M1, substantially reducing the value of L_{p1}. However, the connection from B_2 to Ground must now include

Table 5.1: Extracted inductances and other circuit parameters for the single-threshold comparator.

Parameter	Desired	Extracted	Parameter	Desired	Extracted
L_1	2 pH	1.94 pH	L_{pBs}	0	0.122 pH
L_2	2.8 pH[1]	1.92 pH	I_{cB1}	250 µA	—
L_3	4 pH	3.89 pH	I_{cB2}	250 µA	—
L_4	2 pH	2.20 pH	I_{cBc}	250 µA	—
L_{in}	0	0.443 pH	I_{cBs}	250 µA	—
L_{p1}	0[1]	0.935 pH	β_{cB1}	1	—
L_{p2}	0	0.245 pH	β_{cB2}	1	—
L_{pB2}	0	0.032 pH	β_{cBc}	1	—
L_{pBc}	0	0.113 pH	β_{cBs}	1	—

[1] The desired value for L_2 is actually the sum of L_2 and L_{p1}, as they are in series from the perspective of the interrogating current.

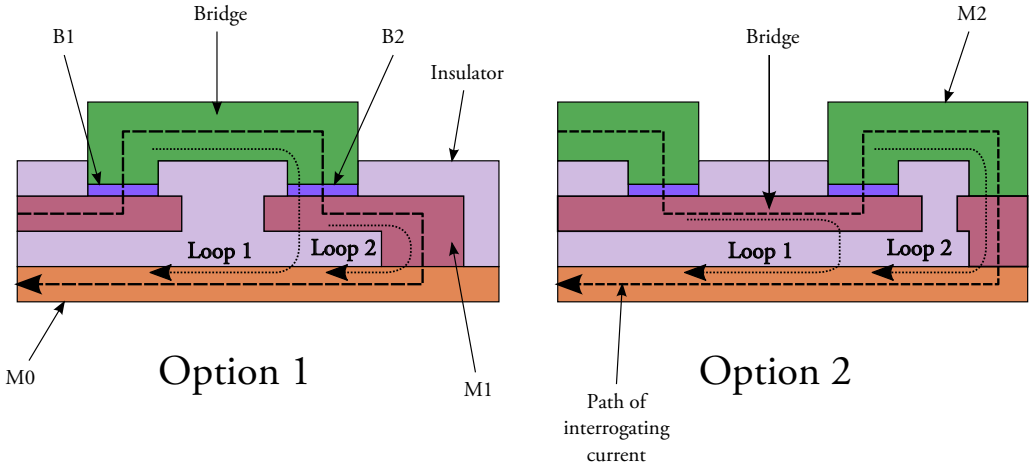

Figure 5.3: Two options for decision-making pair layout: minimisation of Loop 1 or minimisation of Loop 2. Illustration is neither to scale, nor is the relief of the layers indicated correctly.

metal in Layer M2, which increases the loop area of Loop 2. This increases the value of parasitic inductance L_{pB2}.

As will be shown later, minimisation of L_{pB2} is more important, hence Option 1 (employed in the layout of Figure 5.2) will remain the design of choice.

Note that, for almost all purposes, we embrace the concept of partial inductance, which allows division of the inductance of an inductive loop across several edges in a circuit topology and enables modelling of mutual coupling between such edges rather than between loops. The correctness of this is predicated on a number of assumptions (for example the assumption of magnetoquasistaticity), but we do not go into detail in this restricted format.

5.2.2 QOS circuit

The QOS circuit is a useful quantiser first introduced in Section 4.2.3. In the ideal case, the QOS quantiser resembles exactly the circuit first depicted in Figure 4.1. However, especially when fed into a comparator such as the one described in Section 5.2.1, parasitic components must be taken into effect.

The effect of a parasitic inductive load

To this end, we investigate the circuit of Figure 5.4. We model the equivalent inductance loading the quantiser as parasitic inductance L_p.

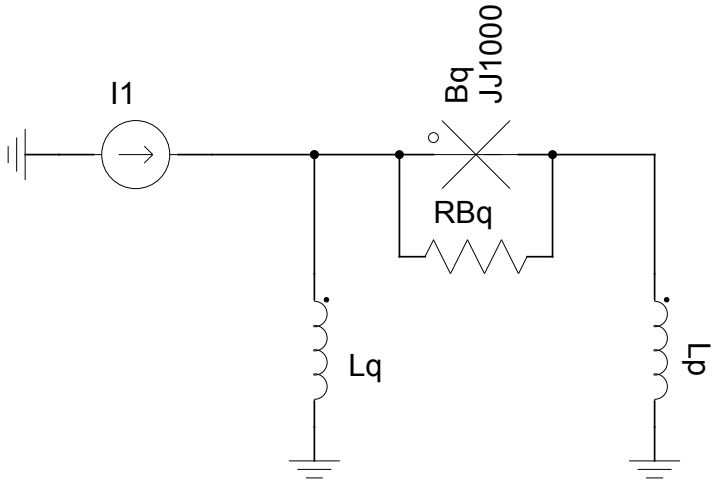

Figure 5.4: A QOS-quantiser with parasitic load L_p.

Simulations of this QOS-quantiser are depicted in Figure 5.5. From these, it is clear that the parasitic load matters least when L_q is low, as the response for $L_q = 0.4\,\text{pH}$, investigated second on Figure 5.4, is most uniform for $L_p = 0.5\,\text{pH}$

and $L_p = 1\,\text{pH}$. For the large parasitic $L_p = 6\,\text{pH}$, no investigated value of L_q is low enough to prevent significant distortion, whereas for the large quantising inductance $L_q = 2\,\text{pH}$, no investigated value of L_p is low enough.

It is expected that a parasitic load has undesirable effects, but since this is unavoidable and ostensibly significant, the effect should be quantified. We recall that I_{cq} imposed a maximum bound on L_q (4.6). Factoring L_p into the system turns out to be straightforward.

Figure 5.4 includes the loop L_q, B_q and L_p, forming the one-junction SQUID. Before, we have neglected the parasitic inductance L_p and have established the constraint (4.6).

To include L_p, (4.4) is applied to Figure 5.4 and expanded to

$$I_1(\phi) = \frac{\Phi_0}{2\pi L_q}\phi + I_{cq}\sin\left(\phi - \phi_p\right),$$

where ϕ is the phase drop atop L_q and ϕ_p is the phase drop across L_p (previously assumed zero), and is given by

$$\phi_p = \frac{2\pi L_p}{\Phi_0}(I_1 - I_L),$$

with I_L the current through L_q. Quantity ϕ_p reduces to

$$\phi_p = \frac{2\pi L_p}{\Phi_0}I_1 - \frac{L_p}{L}\phi,$$

which facilitates

$$I_1(\phi) = \frac{\Phi_0}{2\pi L_q}\phi + I_{cq}\sin\left(\phi - \left(\frac{2\pi L_p}{\Phi_0}I_1 - \frac{L_p}{L}\phi\right)\right),$$

and, consequently,

$$I_1(\phi) = \frac{\Phi_0}{2\pi L_q}\phi + I_{cq}\sin\left(\phi\left(1 + \frac{L_p}{L_q}\right) - \frac{2\pi L_p}{\Phi_0}I_1\right).$$

Recalling that injectivity requires $I_1'(\phi) \geq 0\ \forall\phi$, we derive (implicitly) from $I_1(\phi)$ to obtain

$$I_1'(\phi) = \frac{\Phi_0}{2\pi L_q} + I_{cq}\cos\left(\phi\left(1 + \frac{L_p}{L_q}\right) - \frac{2\pi L_p}{\Phi_0}I_1\right)\cdot\left(1 + \frac{L_p}{L_q} - \frac{2\pi L_p}{\Phi_0}I_1'(\phi)\right),$$

and conclude, since $\cos(x) \geq -1\ \forall x$,

$$\frac{\Phi_0}{2\pi L_q} \geq I_{cq}\left(1 + \frac{L_p}{L_q} - \frac{2\pi L_p}{\Phi_0}I_1'(\phi)\right).$$

This further reduces, in the most restrictive case (as $I_1'(\phi) \geq 0$), to

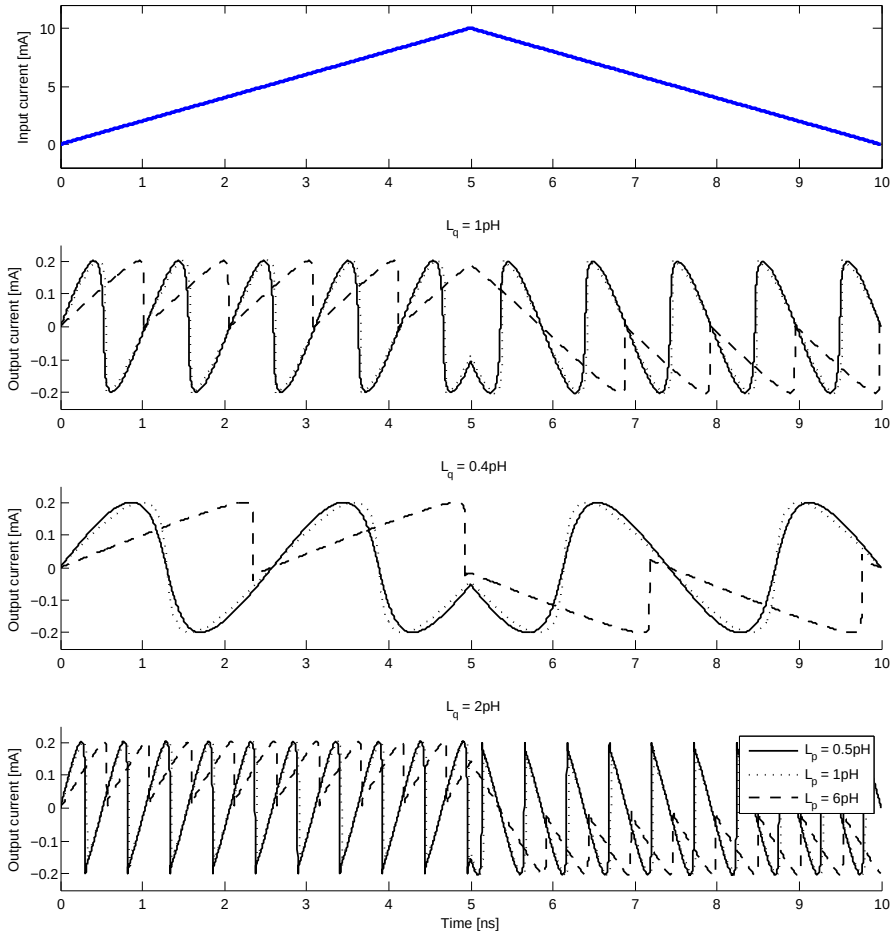

Figure 5.5: Response of QOS-quantiser with parasitic output

$$\frac{\Phi_0}{2\pi L_q} \geq I_{cq}\left(1+\frac{L_p}{L_q}\right),$$

from which we formulate the new constraint

$$L_q + L_p \leq \frac{\Phi_0}{2\pi I_{cq}}. \tag{5.1}$$

Note that this can be written as

$$L_q + L_p \leq (L_{B_q})_{\min},$$

where $(L_{B_q})_{\min}$ is the minimal Josephson inductance of quantising junction B_q.

Importantly, note that the design equation for I_m, (4.8), remains unchanged.

For a given I_{cq}, the choice of L_q (or I_m) thus imposes an upper bound on the allowable parasitic load L_p. Exceeding this upper bound leads to distortion of the quantised current, visible in Figure 5.5.

Sensitivity vs. injectivity

It is established [117] that QOS-quantiser fabricatability must be traded off against sensitivity. The previously derived constraint (4.6), expanded to (5.1), backs this up. To allow a parasitic load L_p, a low quantising inductance L_q is required. But a lower L_q in turn implies a higher I_m, a higher Δ and a higher ν.

For now, we concentrate on achieving sufficient fabricatability, since there is a hard limit, set by (5.1). In fact, the sensitivity ν is less of an issue for our application, as the LNA after the antenna should provide sufficient gain (also, further pre-ADC amplification is not ruled out, as the noise figure of the system is primarily determined by the first device in the receiver chain, the LNA). Thus, we might endeavour to choose an arbitrarily low L_q to relax the upper bound on L_p as much as possible.

There is, however, also a lower bound for L_q. To achieve a 4-bit ADC, for example, the comparator synthesising the LSB will have to respond linearly to inputs ranging from $-4I_m$ to $4I_m$ (see Figure 4.19). The maximum input current will thus be

$$I_{in} = 4I_m. \tag{5.2}$$

Superconductors are characterised by their critical current I_c. Exceeding this critical current kicks the material out of superconductivity. If this should happen on an IC, the following sequence of events is common:

1. A small "bottleneck" area in a superconducting line is kicked out of superconductivity.

2. A resistance appears in the superconducting line, immediately resisting the flow of further current. However, the line has an inductance, which may be very large (for example in bias lines or, likely in this case, the line carrying the ADC input current). This inductances strives to maintain the current for a period of time, *joule heating* the area proportional to I^2R.

3. Heat travels to surrounding areas, which is especially problematic on densely populated structures like ICs.

4. In a cryocooled environment, even a small increase in heat can raise the temperature of nearby superconducting structures above T_c, kicking them out of superconductivity too.

5. This precipitates a chain reaction, as more current-carrying structures are now likely to become resistive, adding further heat to the area. Through this positive feedback, the entire IC could transition into resistive mode, lose functionality and potentially be damaged.

6. This process is not automatically self-resetting, as an IC is generally fed by current sources. If no alternative path is available (after a superconducting path has vanished), as may very well happen when the entire IC is kicked out of superconductivity by the chain reaction described above, the current and associated joule heating will be maintained indefinitely if not explicitly prevented from doing so.

Clearly, this should be avoided at all cost. An upper bound on the maximum ADC input current is thus imposed, which, in turn, imposes a lower bound on L_q.

To quantify this, we turn to the process design rules. Thin-film superconductors, used in IC manufacture, have a fairly well-defined critical current that depends on the thickness of the thin-film as well as the width of the conducting line. The design rules for Hypres's 4.5 kA/cm^2 process specify thicknesses and critical currents as summarised in Table 5.2.

Table 5.2: Critical currents per unit line width of typical thin-film Nb superconductors [12]

Layer	**Thickness** [nm]	**Critical current** [mA/µm]
M0	100 ± 10	20.0
M1	135 ± 10	30.0
M2	300 ± 20	50.0
M3	600 ± 50	70.0

Whereas the IPHT design rules do not specify critical currents for their metal layers, they do specify their thicknesses, which are reproduced in Table 5.3.

Table 5.3: Niobium layer thicknesses of IPHT process [13]

Layer	**Thickness** [nm]
M0	200
M1	250
M2	350

From Tables 5.2 and 5.3 we conclude that assuming a critical current of $I_c = 30.0\,\text{mA}$ per µm line width is safe for all IPHT metal layers, since all IPHT metal layers are thicker than M1 in Hypres's process, with which this critical current density is associated. Note that the superconducting metal employed by the two processes is the same, pure Niobium.

A complicating factor is that the thickness of conductors in a layer is not uniform across a typical IC. When a conductor passes over a step, for example, it constricts. Hypres's design rules warn that such a step may reduce the critical current of the feature by "more than 50%" [12]. In the light of this, we feel that a factor 3 reduction of the expected critical current in a typical chip is conservatively appropriate (note that total critical current of a line is only determined by the weakest point and is thus not cumulatively affected by successive steps).

Resolving to keep all features that conduct large currents at least 10 µm wide, we finally conclude the maximum current safely transmissible into the comparator as

$$(I_{\text{in}})_{\max} = \frac{30.0\,\text{mA}/\mu\text{m} \times 10\,\mu\text{m}}{3} = 100\,\text{mA}. \tag{5.3}$$

Considering (5.3) and (5.2), as well as the previously derived (4.8), we conclude that

$$L_q \geq \frac{4\Phi_0}{(I_{\text{in}})_{\max}} = 82.71\,\text{fH}. \tag{5.4}$$

Repeating the simulation with an L_q of 0.2 pH (incorporating, once again, a safety margin), we achieve the response depicted in Figure 5.6.

To compare the response from Figure 5.6 to our expectations: the minimum Josephson inductance of B_q is, at $I_{cq} = 200\,\mu\text{A}$

$$(L_{B_q})_{\min} = \frac{\Phi_0}{2\pi I_{cq}} \approx 1.65\,\text{pH},$$

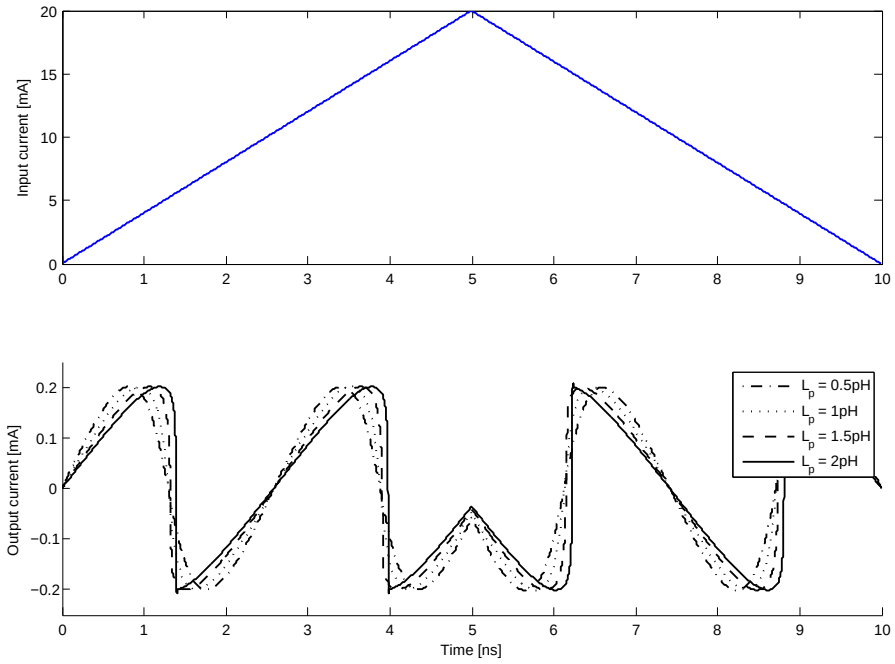

Figure 5.6: The response of an inductively-loaded QOS-quantiser with low L_q.

which requires that the parasitic inductance should not exceed

$$L_p < (L_{B_q})_{\min} - L_q \approx 1.45\,\text{pH}.$$

This is mirrored in Figure 5.6, as $L_p = 1.5\,\text{pH}$ appears right on the edge of linearity, and $L_p = 2\,\text{pH}$ distinctly forces the quantiser beyond the linear regime (into the hysteretic regime).

Minimising the inductive load

If connected to the comparator from Section 5.2.1, L_p would roughly equal the series combination of L_{in}, L_{p2}, L_{pB2} and the inductance of B_2. Although the inductance of B_2 is non-linear, varying with the current through B_2 (I_2), we can make the assumption that the ratio I_2/I_c will not exceed 0.7 when the comparator lies dormant (unclocked). The comparator switching event is a transient event, the effects of which will be investigated later.

Hence, we can assume that

$$L_{B_2} < \frac{\Phi_0}{2\pi 250\ \mu\text{A}} \frac{1}{1 - 0.7^2} \approx 2.58\,\text{pH}.$$

Combining this result with the extracted values from Table 5.1, we thus conclude the parasitic inductance L_p to approach

$$L_p \approx L_{\text{in}} + L_{p2} + L_{pB2} + L_{B_2} \lesssim 3.3\,\text{pH}.$$

This is clearly too high. As our earlier investigation shows (Figure 5.6), even for an L_q approaching the lower bound, an L_p greater than 1.45 pH causes distortion that affects quantiser linearity and thus, presumably, lowers achievable SNR.

Inductance extraction, like any CEM modelling endeavour, is a process that requires skill and finesse. It is a parametric process which relies on several modelling quirks and tricks to produce useful results. Inductors that end in a layout boundary (tail-inductors) require terminals to be placed at that boundary. Terminal modelling relies on several approximations, making extractions of such tail-inductors and adjacent structures less accurate, in practical cases, than structures further away.

For this reason, we combine the designed STC with a QOS quantiser, laid out to respect the previously derived $L_q = 0.2\,\text{pH}$. We take care to minimise the physical distance between the quantiser junction and the STC. The process design rules impose a non-negotiable minimum distance here. The resulting MTC layout is depicted in Figure 5.7. The corresponding schematic and extracted parameters are depicted in Figure 5.8 and Table 5.4 respectively.

Table 5.4: Parameters of QOS-based multiple threshold comparator with minimal L_p.

Parameter	Desired	Extracted	Parameter	Desired	Extracted
L_1	2 pH	1.93 pH	I_{cB1}	250 µA	—
L_2	1.9 pH	1.90 pH	I_{cB2}	250 µA	—
L_3	4 pH	3.87 pH	I_{cBc}	250 µA	—
L_4	2 pH	2.19 pH	I_{cBs}	250 µA	—
L_{in}	0	1.03 pH	I_{cq}	250 µA	—
L_{p1}	0	0.969 pH	β_{cB1}	1	—
L_{p2}	0	0.221 pH	β_{cB2}	1	—
L_{pB2}	0	0.035 pH	β_{cBc}	1	—
L_{pBc}	0	0.114 pH	β_{cBs}	1	—
L_{pBs}	0	0.125 pH	β_{cq}	$\gg 1$	—

For this complete comparator layout, the effective L_p loading the QOS-quantiser still consists of the series combination of L_{in}, L_{p2}, L_{pB2} and the Josephson inductance L_{B_2}. However, we believe our extraction of those key components (Table 5.4) to be more accurate. Thus,

$$L_p \approx L_{\text{in}} + L_{p2} + L_{pB2} + L_{B_2} \lesssim 3.9\,\text{pH}. \tag{5.5}$$

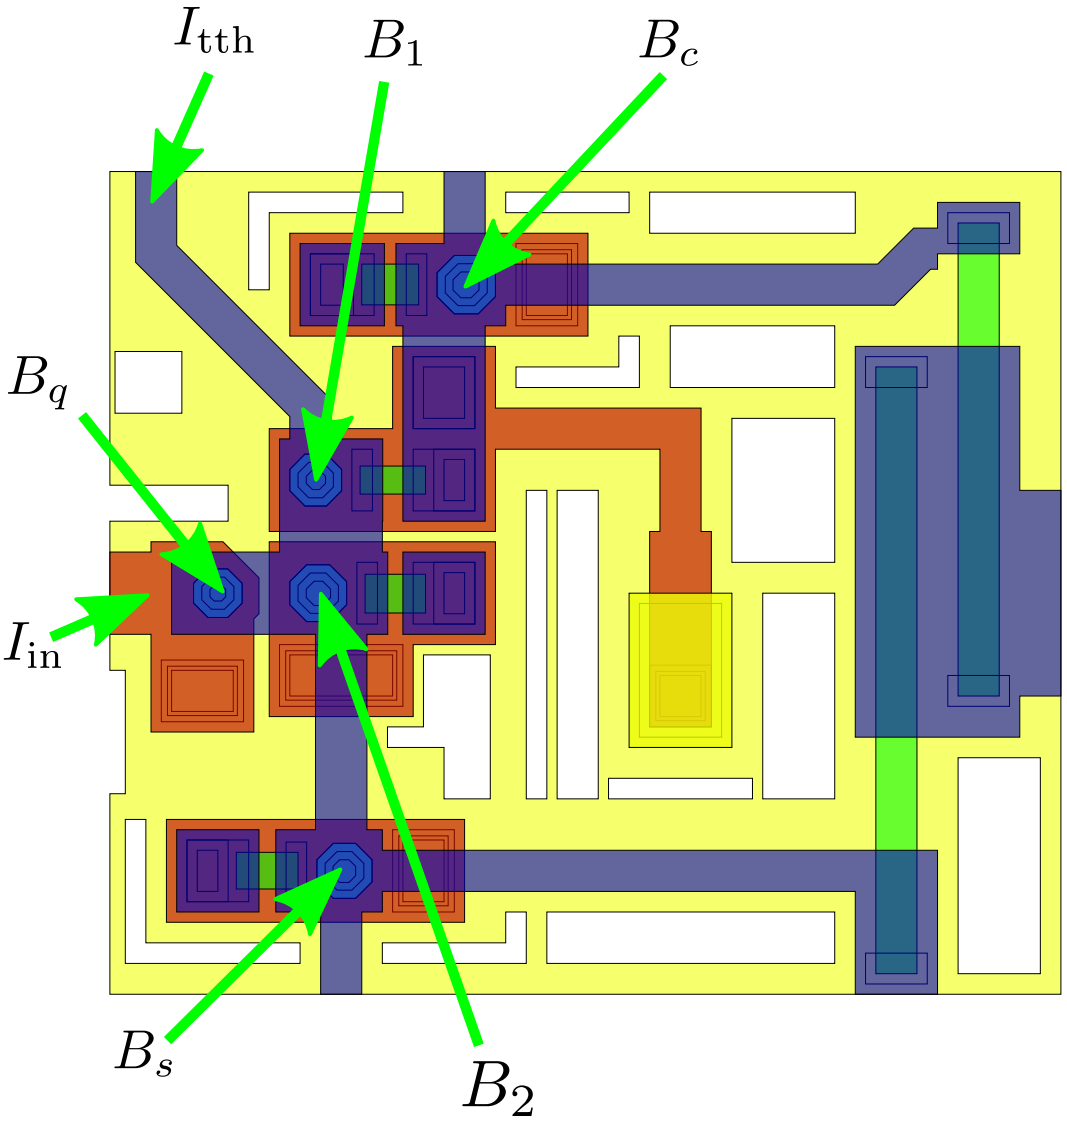

Figure 5.7: Layout of a QOS-based MTC with minimal L_p. The input current enters from the left.

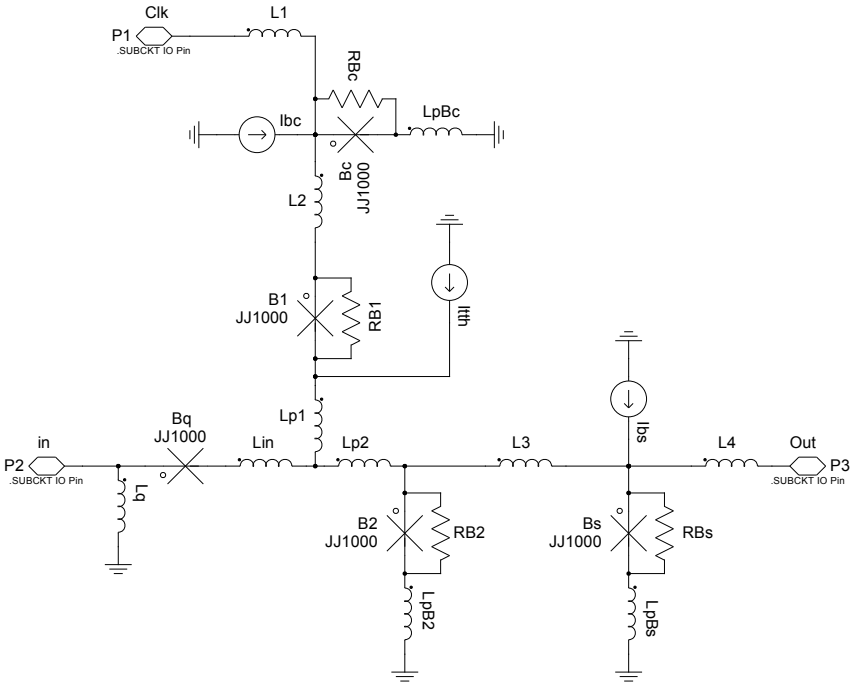

Figure 5.8: Schematic of a QOS-based MTC with minimal L_p.

which is even higher than estimated before. A simulation of the unclocked QOS-quantiser-based multi-threshold comparator is depicted in Figure 5.9. It is clear that L_p is still too high to prevent the duty-cycle of the QOS-quantiser from being skewed. While the input current is rising, the quantised current does reach $+I_{cq}$, but never quite reaches $-I_{cq}$, whereas an analogous response is exhibited for the falling input current. This means that the distances between thresholds cannot be set at $\frac{I_m}{2}$ for both rising and falling input signals, necessarily distorting the digital output.

One way to remedy this is to reduce I_{cq}. Process design rules impose a lower bound of $I_c > 125\,\mu\text{A}$, at which we fix I_{cq}.

The response now is as shown in Figure 5.10. The quantised current now varies correctly between $\pm I_{cq}$ for both rising and falling input signals.

The correct output could be somewhat surprising to the reader, after all the minimum Josephson inductance of a 125 µA junction is still only approximately 2.63 pH, not exceeding the value stated in (5.5). This is because we interpreted the loop inductance conservatively. Naturally, the final term in (5.5), L_{B_2}, is only one path the current I_q can take, hence the actual loop inductance is bound to be smaller (there are paths through B_1 and through B_s). Predicting this inductance is difficult due to the non-linear nature of the paths, hence we believe a circuit-simulation approach to be the correct one here.

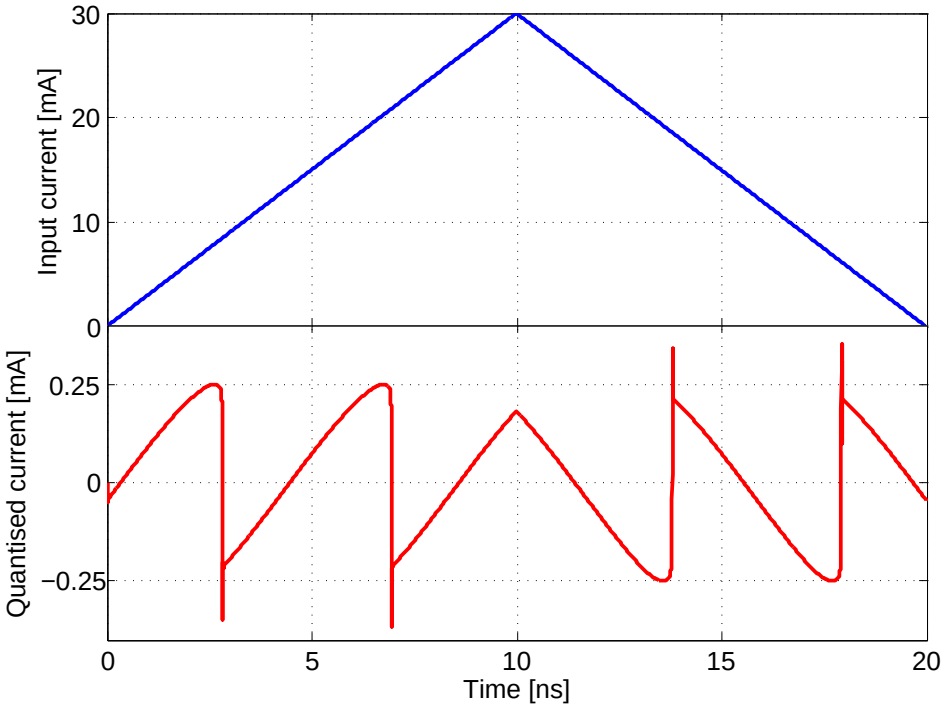

Figure 5.9: Response of the quantiser in the unclocked, preliminary QOS-based MTC.

Applying the clock

So far in our investigation and design procedure we have ignored the impact of the comparator switching. Such switching is the result of a clock pulse arriving at the comparator junction through terminal `Clk`.

When the upper comparator junction (B_1) switches, the SFQ pulse arriving from the clock escapes, it leaks out through the junction. When the lower comparator junction (B_2) switches, its phase increases by 2π. This phase increase falls across the quantising loop, B_q, L_q, L_{p1} and L_{p2}, as well as the output shaping loop, L_3 and B_s.

Since B_s is biased, and L_3 designed accordingly, the phase increase across B_2 leads to a corresponding 2π phase jump of B_s, an output pulse is produced. In the quantiser, the phase increase deposits current into the storing loop. This is undesirable, as it subtracts from the quantised current, introducing an error that depends on the loop inductance. Instead, the loop should be designed so that B_q switches whenever B_2 switches. This equalises the phase difference, introducing only a transient effect that will have settled before the following clock pulse (this imposes an upper bound on the clock frequency).

After applying a 20 GHz clock to the QOS-based MTC, general operation

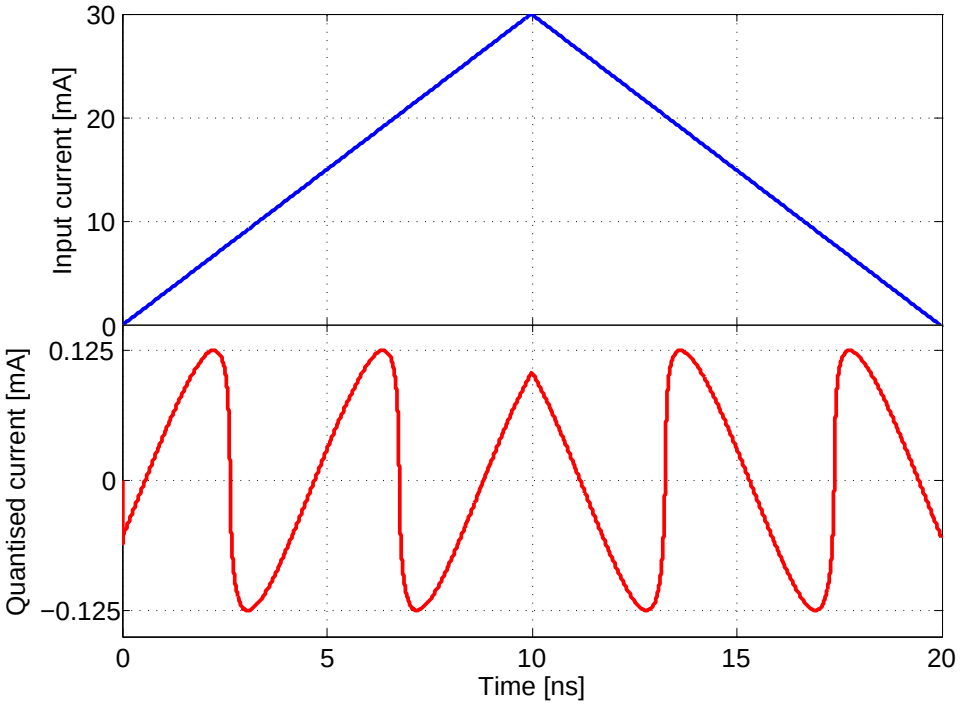

Figure 5.10: Response of the quantiser in the unclocked, preliminary QOS-based MTC after reducing I_{cq}.

of the comparator is observed, although some undesirable behaviour is seen at the comparator thresholds (related to SFQ-bunching and slow recovery of switching junctions, such as B_2). We make the following adjustments to compensate:

- Increase I_{cB2} and I_{cBc}. The larger critical current of B_2 gives it more punch, which, in turn, enables safer switching of B_q and B_s. The increased critical current also enables satisfaction of the $\sqrt{2}$ recommendation (which is explained in Section 2.3.3). Some inductances and bias currents were optimised to reflect these changes. The larger driver junction B_c is required to correctly drive the larger comparator junction.

The output of the clocked QOS-based multi-threshold comparator, after final modifications, is depicted in Figure 5.11. Clearly, the comparator now works as expected (at least with the nominal parameter population). Repeated simulations show correct operation for speeds somewhat exceeding 20 GHz, matching expectations for the IPHT process.

The final layout of the QOS-based multi-threshold comparator designed in this section, as well as corresponding schematics and extracted parameters, is summarised in Section A.1.1.

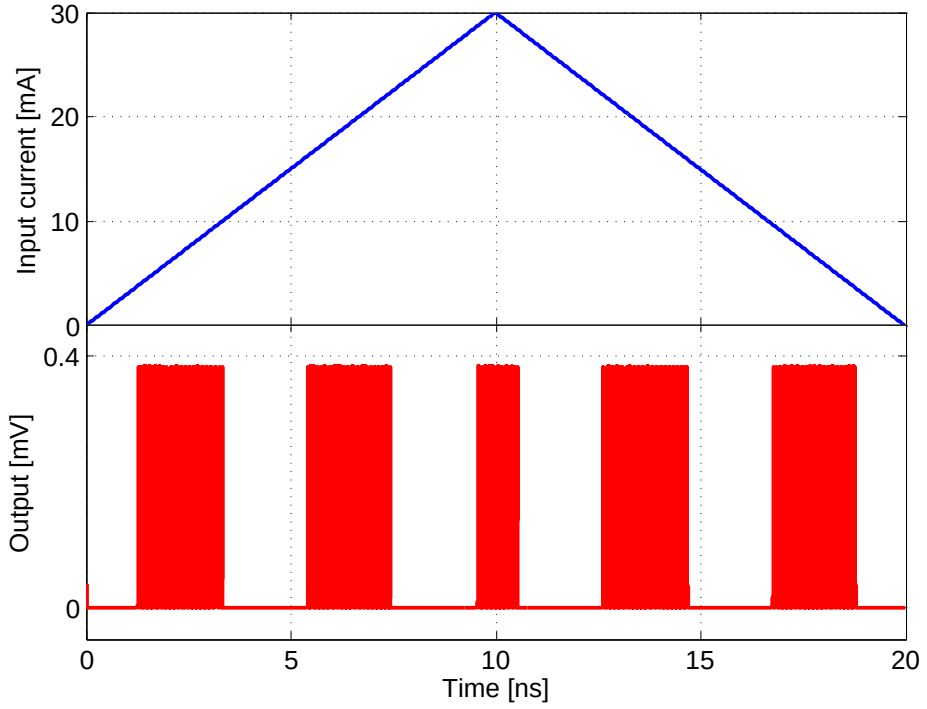

Figure 5.11: Simulated response of clocked QOS-based MTC. The solid rectangular features represent SFQ pulses in quick succession.

5.2.3 Revisiting comparator greyness

Some clarification is in order after the inclusion of parasitics. The quantised current trace in Figure 5.10 suggests that the relation $\beta_L \ll 1$ does not hold here, that is, a significant saw-tooth distortion characterises the quantised current signal. This is to be expected, as injectivity was guaranteed only by decreasing I_{cq} to the minimum allowable. Note that β_L, which is instrumental in determining injectivity, as defined in (4.25), must be expanded to include parasitics, yielding

$$\beta_L = \frac{2\pi I_c(L_q + L_p)}{\Phi_0}. \tag{5.6}$$

We also define β_{L_q} as

$$\beta_{L_q} = \frac{2\pi I_c L_q}{\Phi_0}. \tag{5.7}$$

It is clear that $\beta_{L_q} \leq \beta_L$ and only in the absence of parasitics, $\beta_L = \beta_{L_q}$. Reinvestigating the derivation of sensitivity in Section 4.3.5 allows us to refine its definition to

$$\nu = \frac{\pi I_c}{\beta_{L_q}} 2^n,$$

meaning that it is impossible to reach $\nu_{\min}$, as previously defined in (4.30) in the presence of parasitics (which is unsurprising).

Note that the presence of parasitics also changes the absolute value of I_{gr} and I_{gf}, to

$$I_{gr} = I_g \left(\frac{1}{\beta_L} + 1\right)\left(1 + \frac{L_p}{L_q}\right), \quad \text{and} \quad I_{gf} = I_g \left(\frac{1}{\beta_L} - 1\right)\left(1 + \frac{L_p}{L_q}\right),$$

but the definition of the (relative) quantity γ remains unchanged.

So, the addition of parasitics penalises ν, but, at first glance, leaves γ, and thus achievable SNR, unchanged. However, if one compensates for ν by, say, decreasing I_{cq}, γ is indeed penalised. In the QOS design related above, the (process-determined) minimal value for I_{cq} was employed, penalising γ and, consequently, achievable SNR.

Whether this penalty is significant will be determined later, for now we explore a recently proposed architectural remedy to this tradeoff. The remedy was proposed by Maruyama and Suzuki et al. in 2009 [10]. Examination of the quantised current trace in Figure 5.10 shows that the falling gradient is far steeper than the rising gradient. Viewed in conjunction with Figures 4.18 and 4.30, it is clear that most of the greyness results from the grey-zone associated with the rising threshold.

Clearly it would be beneficial to have a steep gradient for both transitions, a concept embodied in the *complementary quasi-one-junction-SQUID (CQOS)* comparator.

5.2.4 Complementary QOS comparators

A schematic representation of the Complementary QOS (CQOS) multi-threshold quantiser is depicted in Figure 5.12. Two identical QOS quantisers are fed with $i_{\text{in}}(t)$ and $-i_{\text{in}}(t)$ respectively. The lower quantiser is offset by applying a dc current $I_{\text{ph}} = \frac{I_m}{2}$.

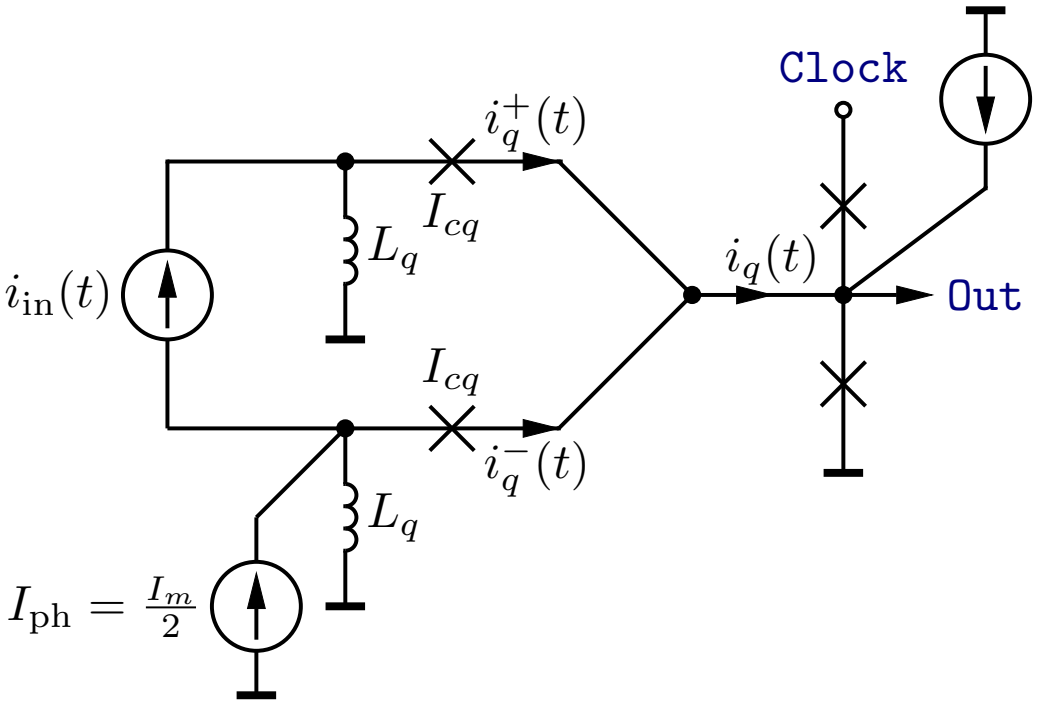

Figure 5.12: The Complementary QOS multi-threshold comparator [10, 11].

Figure 5.13 compares the response of the quantisers from Figure 5.12. I_{cq} is fixed at the process-dependent minimum: $I_{cq} = 125\,\mu\text{A}$. Three values of β_L are investigated, two of which satisfy the condition for injectivity.

Note first, how, at all thresholds in all cases, the CQOS quantiser seems to produce at least as steep a gradient as the QOS quantiser. This is the objective. At $\beta_L \lesssim 1$, the middle plot, the QOS quantiser produces a signal with high sawtooth distortion, as expected. The CQOS quantiser, however, produces a signal that, while not sinusoidal either, has symmetric threshold crossings, each of which is steeper than characterising the QOS-produced signal.

Some literature suggests that the CQOS comparator can be operated when $\beta_L > 1$ [10]. One advantage of employing the CQOS comparator here is that it

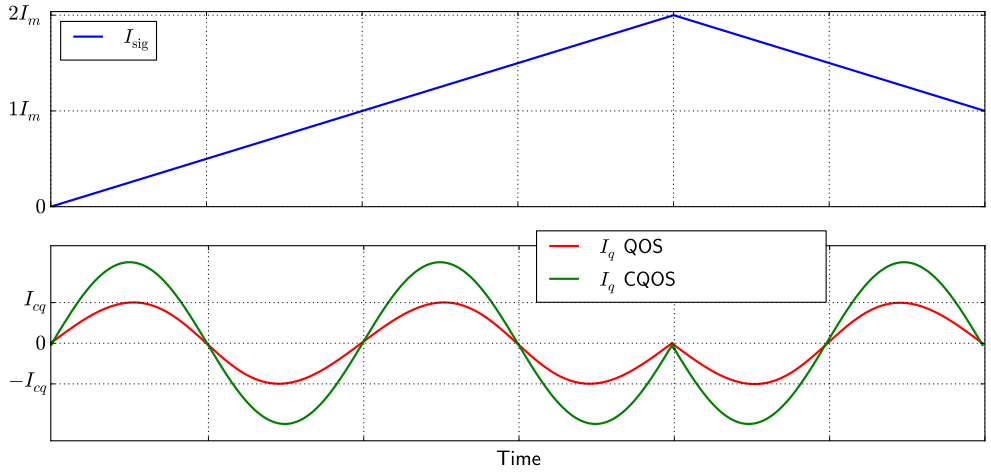

(a)

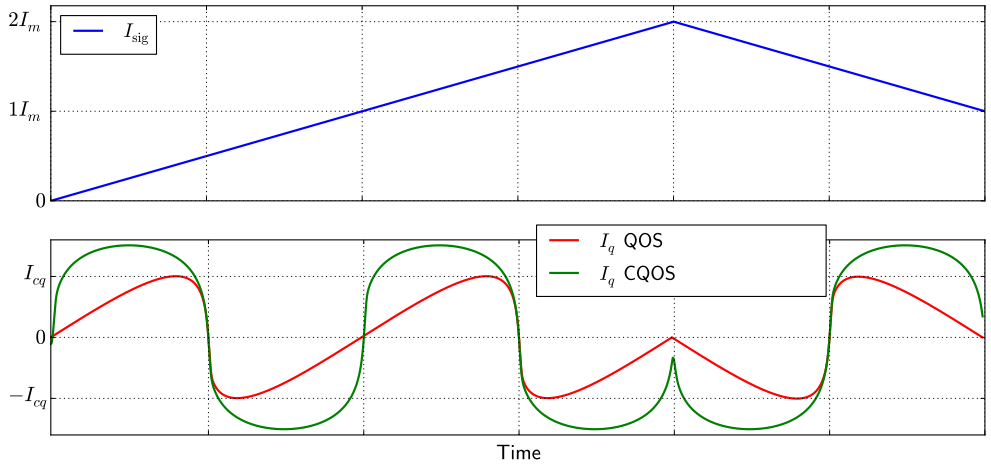

(b)

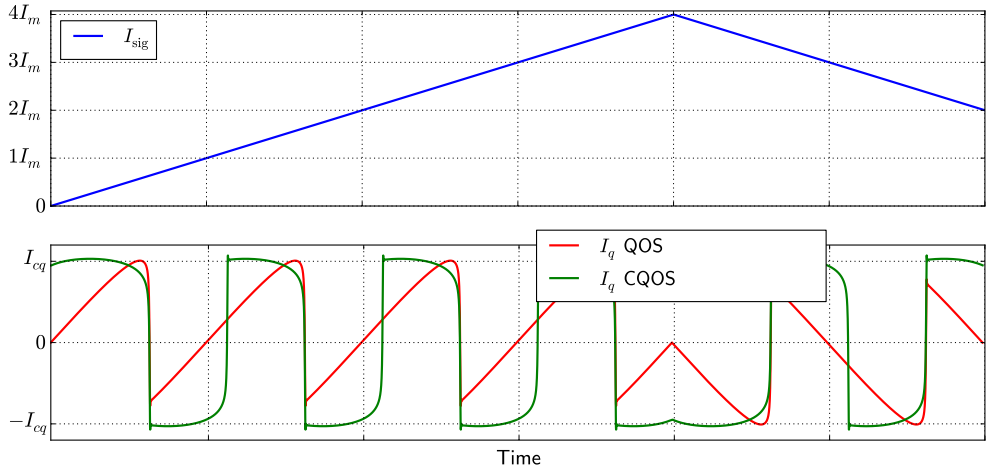

(c)

Figure 5.13: Response of equivalently parameterised QOS and CQOS quantisers. Schematics were depicted in Figures 4.8 and 5.12. The traces are simulations for $I_c = 125\,\mu A$ and $\beta_L = 0.1$ (a), $\beta_L = 0.95$ (b), and $\beta_L = 2.0$ (c).

continues to have a 50% duty-cycle, meaning that equal amounts of the signal period are spent above and below threshold for both rising and falling signals. This is clearly not true for QOS comparators (a property sometimes known as the *duty-cycle problem*). Unfortunately, however, a 50% duty-cycle alone is not sufficient for our application. The CQOS quantiser exhibits hysteresis when $\beta_L > 1$ (as does the QOS quantiser, due to the unsatisfied injectivity condition), as is betrayed by the unequal distances on the abscissa between the signal peak and the two proximal threshold transitions. There is clearly a "dead-zone" after the signal peak before "normal operation" returns. A similar "dead-zone" is visible near the beginning of the trace.

The simulations from Figure 5.13 did not include sampling of the quantised current. To realise the sampler, the Josephson decision-making pair from `MC1`, as well as its clock input and the output shaping circuitry, are re-employed, although some modifications were required to enable the comparator to switch both quantising junctions. We term the resulting multi-threshold comparator `MC2`. The circuit schematic and an excerpt from the layout are depicted in Figure 5.14.

5.3 Benchmarking comparators

The grey-zone of a comparator, defined in Section 4.3.3, refers to the area around its threshold(s) where the decision of the comparator is uncertain. A phenomenon usually attributed to thermal noise, a wide grey-zone is undesirable, as it diminishes ADC SNR (see Figure 4.26).

So far we have discussed the comparator grey-zone as a consequence of the DMP grey-zone, I_g. However, I_g is not a directly measurable quantity for the comparators already discussed. I_{gr} and I_{gf}, however, as well as I_m, are measurable quantities, from which γ is easily derived, as will be demonstrated later. We have concentrated on thermal noise as the source of the DMP grey-zone. Whereas this probably accounts for most of γ, there are surely other sources of uncertainty. It is our hope that those uncertainties can be modelled in terms of additive random variables with Gaussian distribution, in which case our model may continue to prove useful.

Hence, in this section, we evolve the problem from an experimental point of view, so that we are able to design useful experiments later.

5.3.1 Considerations for practical comparators

QOS multi-threshold comparator grey-zones were the subject of some recent research [117, 118, 119], albeit not for this application. The salient points that may pertain to our application are summarised as follows:

- The grey-zone of any comparator is inversely proportional to its decision delay, τ. An intuitive argument for this is that the longer a comparator

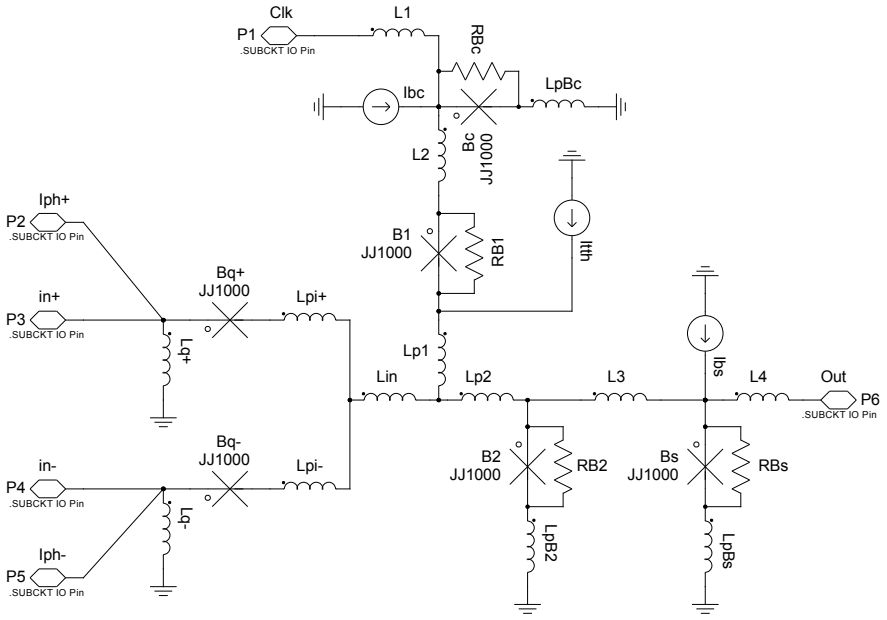

(a)

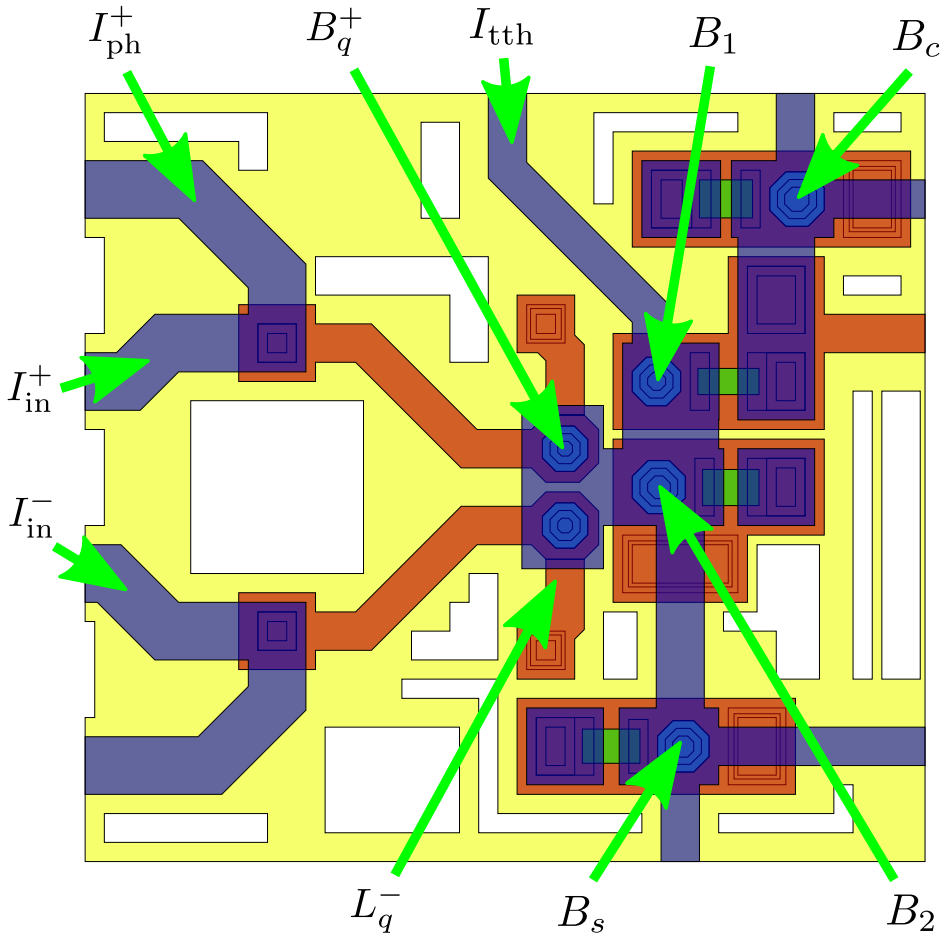

(b)

Figure 5.14: Excerpts of circuit schematic (a) and layout (b) for MC2 (for full details, see Section A.1.2).

has to make a decision, the less likely an error is. The non-zero decision-delay places an upper bound on the maximum clock frequency.

- Maximum comparator operating frequency depends on the characteristic frequency associated with the Josephson junction. This concept was introduced for arbitrary circuits in Section 2.3.2.

In general, we will thus have to be prepared to trade maximum speed of operation against comparator grey-zone.

For the particular process we are targeting for our comparators, the IPHT RSFQ1G process [48], we can make some observations. A typical junction of, say $I_c \sim 250\,\mu\text{A}$, has a shunt resistor of $R \sim 1\,\Omega$. This is to critically damp the junction, so that the characteristic frequency is highest.

As SFQ pulses are roughly triangular, the width of an SFQ pulse is approximately

$$t_p \sim \frac{2\Phi_0}{I_c R} \approx 16.6\,\text{ps}.$$

At first glance this may suggest possible operating frequencies of around 60 GHz, but for our multi-threshold comparators this is a clear over-estimation. The pulse-width is a rough indication pertaining only to a single junction. If pulses are transmitted in a transfer element (such as the one connecting the comparator to the output shaping junction B_s), a good rule of thumb is to make sure that a pulse has traversed the entirety of the transfer element before the next pulse enters it. Whereas this is more a recommendation than a hard rule, the effects of flux quanta "bunching" are somewhat difficult to predict if the pulses are random. At the very least, a slow-down is expected, as flux "back-pressure" temporarily converts transfer elements to storage elements, increasing the time it takes to transmit pulses (back-pressure can be employed intentionally [120]). At worst, bunching and reflecting flux quanta effect a correlation of successive comparator outputs, introducing the comparator *memory effect* [113].

Furthermore, in the case of the two comparators already designed, the quantiser junctions, which have to switch to equalise comparator phase-jumps, are not shunted externally. The normal resistance associated with such a junction is

$$R_n \sim 100\,\Omega.$$

This is about two orders of magnitude higher than the shunt resistance required for critical shunting, implying a corresponding decrease in the characteristic frequency of the junction. Whereas this does not imply a reduction also by two orders of magnitude, the effect is bound to be significant.

As we want to achieve an ADC bandwidth of, at least, 10 GHz, the maximum operating frequency of the comparator should be bounded by

$$f_{\max} \geq 20\,\text{GHz}.$$

We thus deem a comparator that satisfies this criterion as adequate for our purposes.

5.3.2 Finding comparator grey-zones

To obtain an impression of our comparators' non-ideal response, we demonstrate by means of SPICE simulations a switching probability in terms of a dc input current. The test-bed employed for these simulations is depicted later in Figure A.15. We take care to design this test-bed so that it is transferable to an experimental context.

First, we investigate the grey-zone of `MC1`. To estimate the switching probability, we employ Bernoulli trials, as demonstrated by, for example, Ebert and Ortlepp [118, 111]. A parameterised SPICE simulation is set up in which N_{tot} clock pulses are applied to the comparator in the presence of a dc input signal $i_{\text{in}}(t) = I_{\text{in}}$. Every output pulse produced is considered a "hit", whereas the absence of such a pulse after the clock pulse is considered a "miss".

The number of hits N_{pass} is counted and the switching probability is estimated to be

$$P\{\texttt{Out}|\texttt{Clock}\} = \frac{N_{\text{pass}}}{N_{\text{tot}}}. \tag{5.8}$$

If we employ this technique for discrete values of I_{in} swept over a range where we expect the threshold to be, we obtain a switching probability plot akin to that predicted in Figure 4.17. We can then fit the expected probability curve (4.22) to it and extract the actual (simulated) threshold and grey-zone from the fit.

To obtain the fit for the rising curve (a `0` to `1` threshold), we employ a well-understood, multi-dimensional function minimisation algorithm known as the Nelder-Mead Simplex algorithm [121], as facilitated by SciPy [103]. Employing (4.22) and (4.23), a function is constructed for the ideal switching probability,

$$f_p(I_g, I_{\text{th}}, I_{\text{in}}) = \frac{1}{2}\left(1 + \text{erf}\left(\frac{\sqrt{\pi}}{I_g}(I_{\text{th}} - I_{\text{in}})\right)\right). \tag{5.9}$$

Then, after synthesising vector $p[k]$ of switching probabilities for input currents $x[k]$, the error function to be minimised is constructed, as

$$f_{\text{err}}(I_{\text{th}}, I_g) = \sum_{k=1}^{N} \left[f_p(I_g, I_{\text{th}}, x[k]) - p[k]\right]^2. \tag{5.10}$$

The Nelder-Mead algorithm then minimises the function, yielding a least-squares best fit of the quantities I_{th} and I_g. This fitting procedure is easily adjusted for the falling thresholds (`1` to `0` threshold).

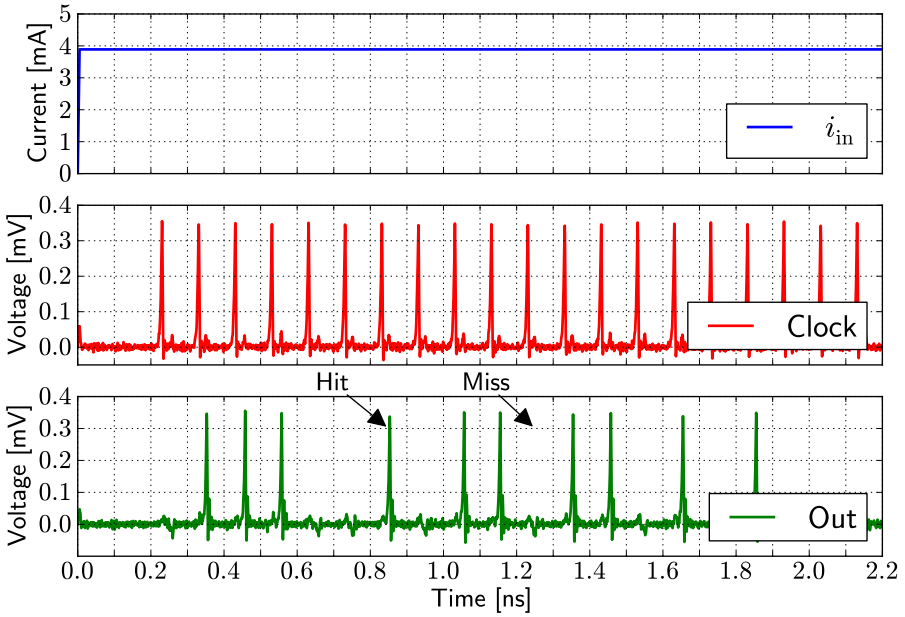

Figure 5.15: Demonstration of how some `Clock` events produce an `Out` event and others do not for `MC1` at I_{in} near threshold.

A trial fit for one rising and one falling threshold of `MC1` is depicted in Figure 5.16. Visual inspection suggests that the fit is good, but we rely on the $N_{\text{tot}} = 10\,000$ pulses generally reported as sufficient in the literature [118, 117].

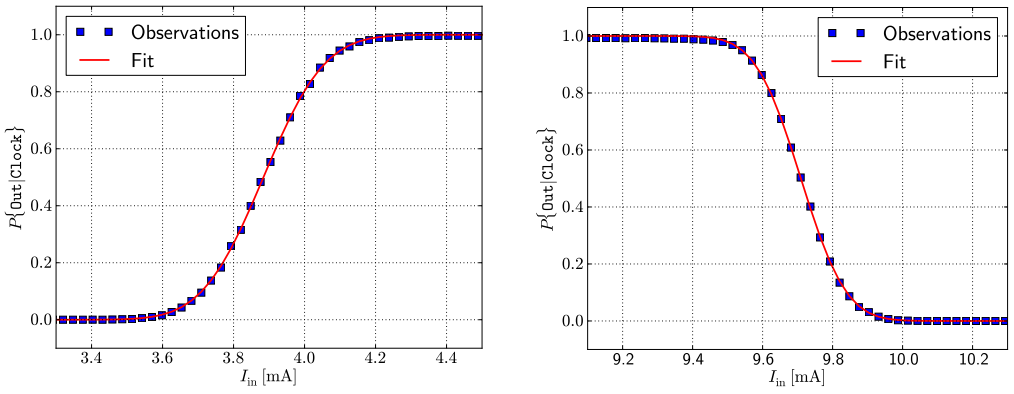

(a) (b)

Figure 5.16: An illustration of the fitting procedure employed to determine threshold and grey-zone of an MTC transition. The example investigates a rising (a) and a falling (b) threshold of `MC1`.

So far, we have designed two different multi-threshold comparators, `MC1` and `MC2`. The results of investigating their grey-zones over their first five periods I_m by means of SPICE simulations are summarised in Table 5.5.

Table 5.5: Grey-zones as estimated from SPICE simulations.

	MC1		MC2	
Threshold	I_{th}[mA]	I_{gr}, I_{gf}[µA]	I_{th}[mA]	I_g[µA]
1st rise	3.95	333.5	0.194	12.3
1st fall	9.71	260.6	1.00	11.6
2nd rise	16.3	333.7	1.77	12.7
2nd fall	22.1	260.2	2.58	12.2
3rd rise	28.7	334.3	3.35	12.3
3rd fall	34.5	259.5	4.16	12.4
4th rise	41.1	336.2	4.93	12.4
4th fall	46.9	259.0	5.74	12.3
5th rise	53.5	336.1	6.51	12.6
5th fall	59.2	262.7	7.32	11.1
avg rise		334.8		12.5
avg fall		260.4		11.9

From averaging the determined I_{th} values for all rising thresholds, we conclude that for `MC1`, $I_m \approx 12.38\,\text{mA}$, and for `MC2`, $I_m \approx 1.578\,\text{mA}$. Both values agree exactly with our prediction, which results from (4.8) and the extracted values of L_q for `MC1` and `MC2`. Whereas the dataset is too small to resolve a meaningful standard deviation, the observed values for I_m are all within less than 1% of the mean.

This by itself is not further remarkable, after all we have obtained much evidence for the validity of (4.8). However, it does inspire confidence in our fitting method and observed data. We also observe that the rising threshold grey-zone for `MC1` is greater than the falling threshold grey-zone, as predicted, whereas the grey-zones for `MC2` are far closer to equally wide. If we extrapolate from this confidence, we make an important observation: the average grey-zone (for rising and falling thresholds) of `MC2` is substantially smaller than that of `MC1`, even when taking the difference of I_m into account.

The important measure here is the greyness γ, the proportion of the period I_m that is grey. For our small test set, we obtain the following values for rising (and falling) thresholds:

$$\texttt{MC1}: \quad \gamma \approx 0.0270\ (0.0210), \qquad \texttt{MC2}: \quad \gamma \approx 0.0079\ (0.0076).$$

Taking the average, which is reasonable as each period is bordered by one rising and one falling threshold, we obtain

$$\texttt{MC1}: \quad \gamma \approx 0.024, \qquad \texttt{MC2}: \quad \gamma \approx 0.0077.$$

This represents a reduction by approximately factor ~ 3 for MC2 over MC1, certainly significant!

While the grey-zone of Josephson decision-making pairs has been investigated extensively before [118] and recently [111], and the grey-zone of QOS-based comparators has been investigated for a different application [117, 119], we are unaware of efforts made specifically to measure the grey-zones of CQOS-based comparators.

5.4 Experimental verification of comparator performance

5.4.1 Parameters influencing the grey-zone

Besides the key design parameters derived, such as L_q, I_m and I_{cq}, other parameters influence the behaviour of the comparators investigated so far. Once the quantiser period is fixed, and assuming the threshold is set up correctly, the higher-order behaviour of the Josephson junctions in the circuit should be considered. These are determined mainly by the damping factor β_c of the junctions, which, in a fixed process, is determined by the junction area and corresponding shunt resistance (see Section 2.3.2). Note that the shunt resistance also determines the character and effect of thermal noise in the system.

Since shunt resistors determine the nature of junction switching behaviour, as well as the locality and intensity of noise sources, we now focus our investigation on their arrangement.

For both MC1 and MC2, the quantising junctions B_q, B_{q+} and B_{q-} were left unshunted. There were several reasons for this choice:

- Resistors are noisy, potentially adding to the width of the grey-zone,
- non-shunted junctions are smaller and thus impose smaller parasitics, and
- the analog behaviour of the quantiser does not necessarily require this junction to switch.

Switching of an underdamped ($\beta_c \gg 1$) junction is accompanied by oscillations in ϕ_J (see Section 2.3.2). These oscillations are undesirable as the corresponding current oscillations cause the comparator threshold to fluctuate.

Both junctions of the Josephson decision-making pair (B_1, B_2) were critically damped, in MC1 as well as MC2. This is common practice for DMP

junctions, as one of the DMP junctions switches during each clock period, substantially contributing to the transient behaviour of the circuit. Critical damping is desirable, as it minimises higher-order effects that are difficult to predict at the design stage. In the case of multi-threshold comparators in particular, the threshold should remain constant over time, as a drift or oscillation would diminish the achievable SNR.

5.4.2 Experimental structures

To fortify our understanding of the Complementary QOS comparator, it was decided to investigate the effect of junction damping. While the ideal Josephson junction model has proven itself suitable in developing an understanding of the quantising circuit, the character of the discriminator requires, at least, a second-order differential equation to model the most significant effects.

As we already have a mechanism to determine comparator grey-zones by SPICE simulation, we will employ this setup to gather data on three more comparators: `MC3`, `MC4` and `MC5`. This set was chosen to populate one 5×5 mm chip made available to this experiment.

MC3 — critically damped B_{q1} and B_{q2}.

Reducing β_c of B_{q1} and B_{q2} to 1 by employing the corresponding shunt resistance reduces the tail oscillations accompanying the switching of the quantiser junctions. As these junctions switch when B_2 switches (sampling), the hope is that critically damping these junctions will reduce the length of transient effects, possibly improving the maximum operating frequency achievable.

MC4 — underdamped B_2.

Here the quantising junctions are kept critically damped, but the shunt resistor is removed from B_2, leaving it underdamped. An underdamped junction has a higher bandwidth and thus absorbs more k_BTB noise, possibly increasing threshold uncertainty. The hope is that greater speeds could be employed.

In fact, the lack of a convenient sample-and-hold circuit arguably is a disadvantage of superconducting analog-to-digital converters, as such a circuit would free up the entire clock period to make a decision about the signal, potentially significantly reducing the noise contribution.

MC5 — co-damped B_1 and B_2.

The tradeoff between comparator speed and grey-zone is a known property of Josephson comparators. Ortlepp proposed minimising the grey-zone by employing the damping scheme depicted in Figure 5.17 [111].

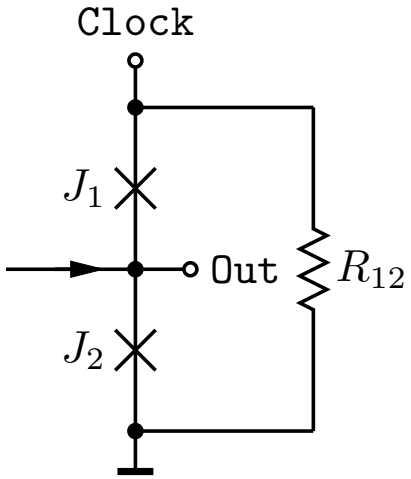

Figure 5.17: Co-damping to maximise speed/grey-zone figure of merit.

This strategy is employed in structure MC5. The damping resistance is chosen as the shunt resistance that would critically damp B_1 if connected singly.

5.4.3 Verification structures

In addition to MC1 and MC2, the comparators MC3 through MC5 were also laid out in the IPHT RSFQ1G process, completing our set of experimental structures. These structures are summarised in Table 5.6.

Layouts and circuit parameters are detailed in Section A.1. The chip containing the experimental structures is depicted in Figure A.16. In Figure 5.18 we depict the test-setup of one comparator, in this case MC3. Figure 5.19 depicts a photograph of the testbed for MC1. Standard FLUXONICS DCSFQ converters, SFQDC converters and FLUXONICS JTLs [?] are employed to connect to the SFQ interfaces of the comparator.

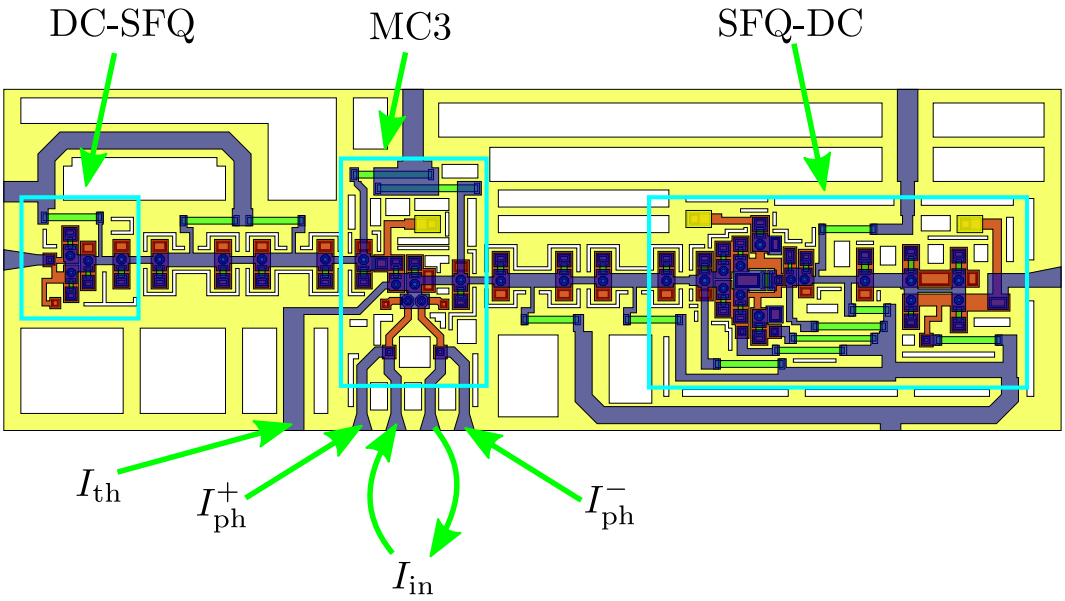

Figure 5.18: Test structure to evaluate comparator `MC3`. This particular structure is located in the South-West of `mvv2-gz-inv`.

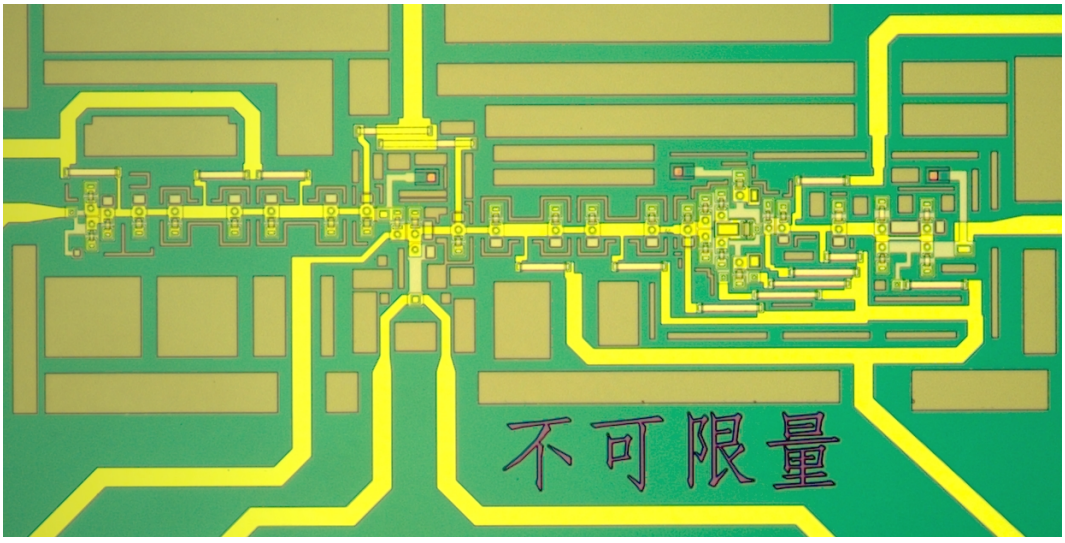

Figure 5.19: Microphotograph of test structure to evaluate comparator `MC1` (also see Figure 5.18).

Table 5.6: Structures developed to investigate various comparators

Structure	**Location**[1]	**Quantiser Type**	L_q[2]	**Quantising JJ(s)**[3]	B_1[4]	B_2[5]
Investigated CQOS comparators – `mvv2-gz-inv`						
`MC1`	South Centre	QOS	very low	$\beta_c \gg 1$	$\beta_c \approx 1$	$\beta_c \approx 1$
`MC2`	South-East	CQOS	medium	$\beta_c \gg 1$	$\beta_c \approx 1$	$\beta_c \approx 1$
`MC3`	South-West	CQOS	medium	$\beta_c \approx 1$	$\beta_c \approx 1$	$\beta_c \approx 1$
`MC4`	North-East	CQOS	medium	$\beta_c \approx 1$	$\beta_c \approx 1$	$\beta_c \gg 1$
`MC5`	North-West	CQOS	medium	$\beta_c \approx 1$	shunted together	

[1] On the GDSII layout.
[2] Grounded inductance of the quasi-one-junction SQUID(s).
[3] Quantising junction of the quasi-one-junction SQUID(s).
[4] Upper junction of decision-making pair.
[5] Lower junction of decision-making pair.

5.4.4 Cryogenic experiment

Test setup

The experimental test setup employed to measure the grey-zone of the designed `MTC` iterations is depicted in Figure 5.20. This test setup was conceived and constructed at the Institute for Photonic Technology (IPHT) in Jena, Germany, where such equipment is generally available. In fact, much of an existing setup could be converted for this experiment, only the FPGA gateware and the Labview interface that controls the FPGA had to be customised for this experiment. The present author is indebted to the staff at the Department of Quantum Dectection of the IPHT, particularly to Dr. Olaf Wetzstein, for helping to devise the setup and implementing some of the required functionality.

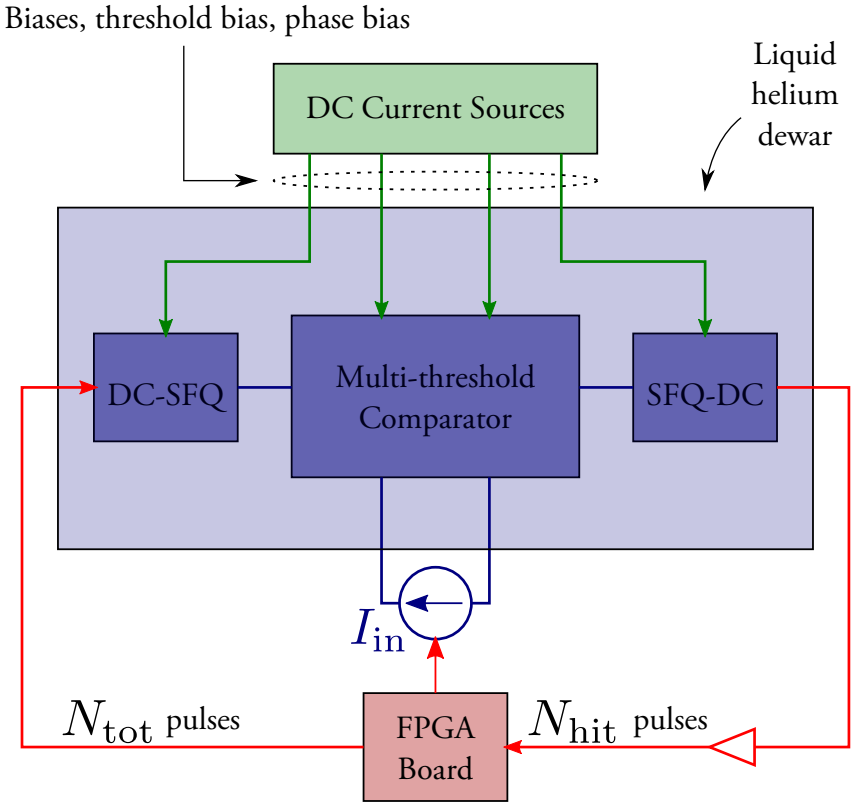

Figure 5.20: Experimental test setup to measure the grey-zone of various multi-threshold comparators.

For each test structure (`MC1` through `MC5`), the following experiment is devised:

1. Bias currents for the comparator and the periphery (DC-SFQ and SFQ-DC converters) are set up by hand.

2. Threshold and phase biases I_{tth} and $I_{\text{ph}-}$ are set up by hand to obtain the correct period I_m. The period is determined by clocking the comparator and displaying the output on an oscilloscope so that real-time changes can be detected.

3. The FPGA is employed to vary I_{in} over several I_m cycles of the comparator. For each step in the investigated range of I_{in}, the FPGA applies N_{tot} pulses to the input SFQ-DC converter, which interrogates the comparator clock. The FPGA then counts the number of pulses N_{hit} (hits) that exit from the comparator, asserting that at most one pulse is received per supplied clock pulse. The switching probability for that particular input current I_{in} is then estimated as in (5.8).

4. Once the switching probability has been established over a sufficiently large range, the grey-zone can be determined by curve-fitting, where the method described in Section 5.3.2 is employed.

Note that the input current I_{in} is supplied by a floating current source. This is required for structures `MC2` to `MC5`, as a differential input is required for the CQOS comparators.

Fabrication

Several wafers that contained the designed chip, `mvv2-gz-inv`, were manufactured within the research time allocated to this work.

IC fabrication generally occurs in runs, where each run includes the following steps:

1. Designers send in their layouts. A layout for the entire wafer is assembled and a design-rule check is run to verify that designers kept to the rules. If they did not, they are informed of this and given a chance to make corrections.

2. The layout is sent away to the mask manufacturer. After this step, no changes can be made to the layout. The mask manufacturer makes one mask for each layer. This is expensive.

3. The set of masks is sent back to the foundry, where wafers can now be manufactured. Normally there are many steps involved in making a wafer, but at the very least for each layer a material must be deposited, then exposed to a photoresist through the mask, and finally the exposed parts etched away.

4. After all layers are deposited and etched, including the contact pads, the wafer is cleaned and diced into chips.

5. Specially included diagnostic chips are measured and compared against known results to assess the quality of the wafer. Common parameters to be identified include the critical current density of the tri-layer, J_c, the critical current of several Josephson junctions (especially very small junctions, which are most challenging to manufacture accurately) and the resistances, capacitances and inductances of certain structures. Common diagnostic structures to facilitate this are SQUIDs, Josephson junction arrays and large capacitive plates.

6. If diagnostics yield parameters that are outside specifications, the wafer is discarded and this process repeated from Step 3. Sometimes, diagnostics are also performed during fabrication, before all layers are deposited and etched.

Note that the same mask can be "stepped" across the wafer, enabling multiple copies of a chip to be produced on the same wafer.

For our experimental chip `mvv2-gz-inv`, a good wafer, with measured diagnostic parameters within specifications, was made available. This made it possible to carry out the experiment described in Figure 5.20.

5.4.5 Results

A first look at cryogenic results

After a wafer was fabricated to specification and diagnostic structures had been tested successfully, the comparators were subjected to the cyrogenic experiment described earlier in this subsection. First, comparator operation was confirmed visually by means of oscilloscope traces, an example of which is depicted in Figure 5.21. Here the oscilloscope is triggered on the periodic triangular input signal $i_{\text{in}}(\text{t})$, and the clock frequency f_c is an integer multiple of the frequency of $i_{\text{in}}(\text{t})$. The clock signal and the input current control signal were both generated by the same arbitrary-wave generator (from the Agilent 33500B series [122]), which supplies two internally synchronised channels.

The grey areas of the comparators are not conveniently visible on the traces from Figure 5.21, they correspond to the range around each threshold where a `Clock` event may or may not result in an `Out` event. Since this is a measure of probability, which we have defined in terms of frequency of output events (N_h) given clock events (N_{tot}), a sizeable number of clock pulses must be applied and evaluated for different grey input current levels.

Employing the "visual confirmation" method enabled "tuning" the comparator bias currents in the presence of a continuous clock signal to achieve the desired response. After this was done, the FPGA bit-error measurement

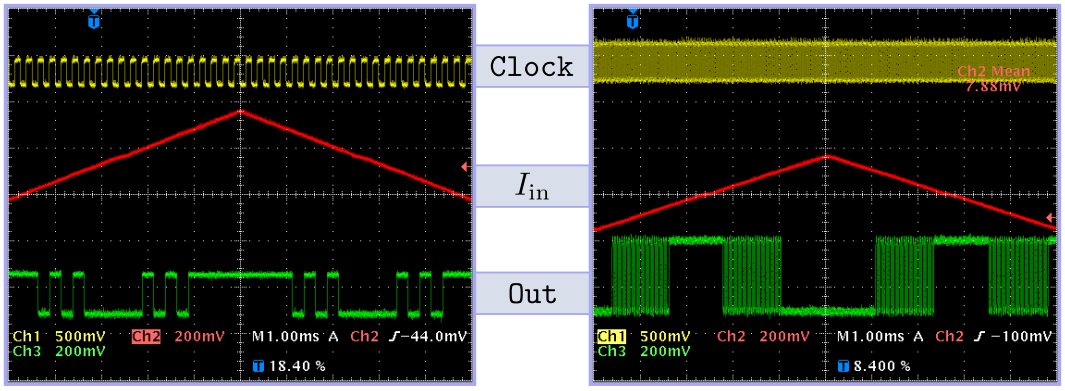

Figure 5.21: Correct operation of a fabricated comparator (here `MC4`), with a comparatively low (left) and high (right) frequency clock (both in the $\sim$ 10 kHz range). Every rising edge of the `Clock` signal signifies one clock event and every rising and falling edge of the `Out` signal signifies one output event. (Note that the experiments were not recorded at the same time or necessarily the same operating point, which explains the differences in signal level.)

setup was connected, replacing the arbitrary-wave generator and receiving a copy of the `Out` signal for evaluation.

In this way, $I_{\text{in}}$was swept over the range $[0, 4\,\text{mA}]$ and the switching probability was measured for a number of points in the rangeis. Bias currents were further adjusted until the areas between transitions were roughly equal. A deviation of $0.01I_m$ from the ideal value of $0.5I_m$ was deemed acceptable as the distance between two transitions.

After the likely thresholds were read off from a rough sweep, fine sweeps around the relevant areas were conducted, applying $N_{\text{tot}} = 10\ 000$ pulses for each investigated value of I_{in}. The switching probability was evaluated for each point, yielding a curve from which threshold I_{th} and grey-zone I_g could be determined by curve fitting. Two examples of such fits are shown in Figure 5.22.

While no formal goodness-of-fit analysis was carried out, visually the fits look decent and it looks as though the choice of fitting function (5.10) is a good one. Two additional examples of experimentally determined results are shown in Figure 5.23. is

The measured curves for `MC5` suggest that, at least for that comparator, assumptions made when deriving comparator switching behaviour, derived in Section 5.3, are not entirely accurate. A clue is found in our hasty simplification of all noise sources into one, and our simplification of switching to an instantaneous stochastic process. In real Josephson balanced DMPs, as first shown in Figure 4.15a, junction switching is a dynamic process where the dynamics of both junctions play a role. While it is desirable to have the balanced comparator act as a monolithic device, and this assumption is suitable for many applications (backed up by much empirical evidence, especially in

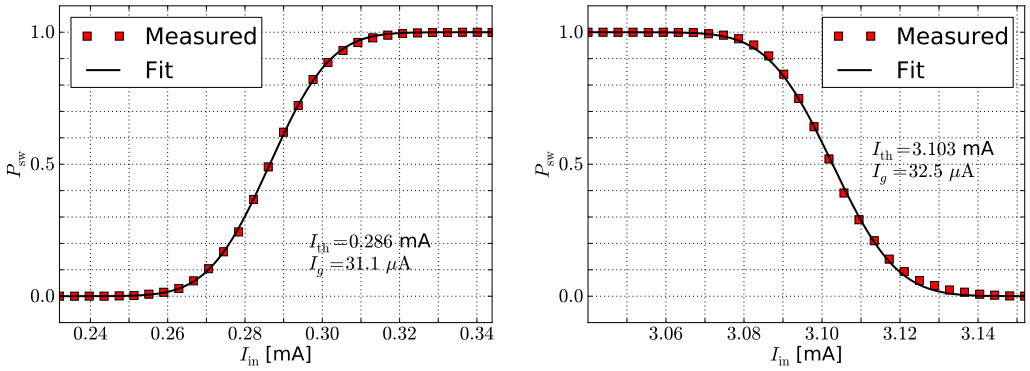

(a) (b)

Figure 5.22: Measured switching probabilities and fit curves for the first (a) and fourth (b) measured transition of `MC3`. Both the threshold and the grey-zone can be determined in this way. Visually the fits look good.

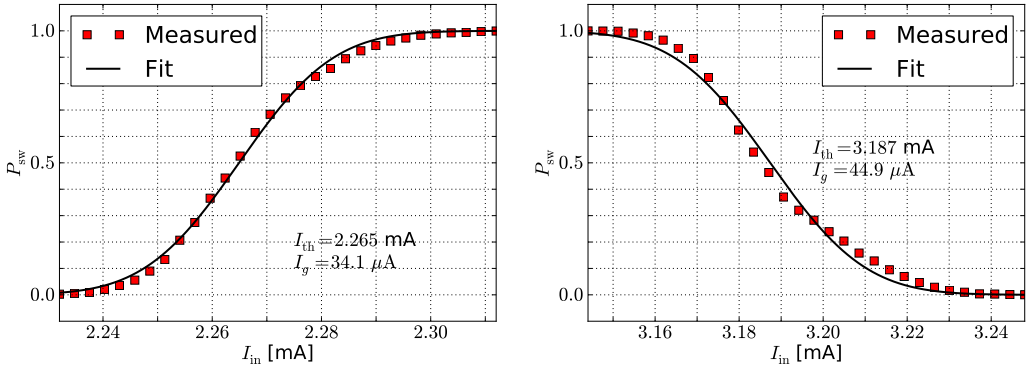

(a) (b)

Figure 5.23: Additional measured switching probabilities and fit curves, here for comparator `MC5`, for the third (a) and fourth (b) transitions. Visually the fits do not look as good as in the case of `MC3` (Figure 5.22).

the domain of digital circuits), the evidence presented in Figure 5.23 suggests re-examination of the fitting function.

Correction of fitting function

The dynamic response of single noisy Josephson junctions has been extensively studied, for example by Ambegaokar [123] and Filippov [109]. Derivations generally assume fluctuations consistent with Brownian motion, which results, at least when simplified to the static case, in the behaviour we had assumed in this work.

We believe it outside the scope of this work to derive the general solution for the dynamic response of the decision-making pair, but summarise the following insights from our survey of known work:

1. The two junctions in the DMP, while certainly not independent (in fact their behaviour is strongly coupled), do have dynamic response components which are separate from each other. Nonetheless, one junction normally dominates the response, with the other junction eventually following suit.

2. Stochastic switching errors, such as both junctions or neither junction switching under interrogation, are possible [118], with corresponding error propagation through the circuit.

3. Given these insights, static models are remarkably good when employed within parameter boundaries that, fortunately, include practical circuits.

Hence, it is likely that each Josephson junction in the decision-making pair can be modelled as a static comparator, with each contributing to the switching probability but no change to the outcome space possible (either an `Out` event occurs or it does not). While switching errors can be problematic, they can be accommodated at the circuit level, through SFQ buffers for example, and thus are not able to effect a decision outside the described outcome space. Hence, we assume that they can be absorbed into our model. We thank Dr Thomas Ortlepp for very helpful suggestions regarding this matter.

We thus construct a new sigmoid function for $P\{\texttt{Out}|\texttt{Clock}\}$ consisting of the normalised superposition of two noisy comparators:

$$
\begin{aligned}
f_{p_2}(I_{g_1}, I_{\mathrm{th}_1}, I_{g_2}, I_{\mathrm{th}_2}, I_{\mathrm{in}}) &= \frac{1}{2} f_p(I_{g_1}, I_{\mathrm{th}_1}, I_{\mathrm{in}}) + \frac{1}{2} f_p(I_{g_2}, I_{\mathrm{th}_2}, I_{\mathrm{in}}) \\
&= \frac{1}{2} \\
&+ \frac{1}{4} \operatorname{erf}\left(\frac{\sqrt{\pi}}{I_{g_1}} (I_{\mathrm{th}_1} - I_{\mathrm{in}}) \right) \\
&+ \frac{1}{4} \operatorname{erf}\left(\frac{\sqrt{\pi}}{I_{g_2}} (I_{\mathrm{th}_2} - I_{\mathrm{in}}) \right).
\end{aligned} \tag{5.11}
$$

and the corresponding error function

$$
f_{\mathrm{err}_2}(I_{g_1}, I_{\mathrm{th}_1}, I_{g_2}, I_{\mathrm{th}_2}) = \sum_{k=1}^{N} \left[f_{p_2}(I_{g_1}, I_{\mathrm{th}_1}, I_{g_2}, I_{\mathrm{th}_2}, x[k]) - p[k] \right]^2. \tag{5.12}
$$

Minimisation of the error function f_{err_2} yields the, visibly improved, fits for `MC5` displayed in Figure 5.24. Other fits look similarly good and are omitted here for brevity.

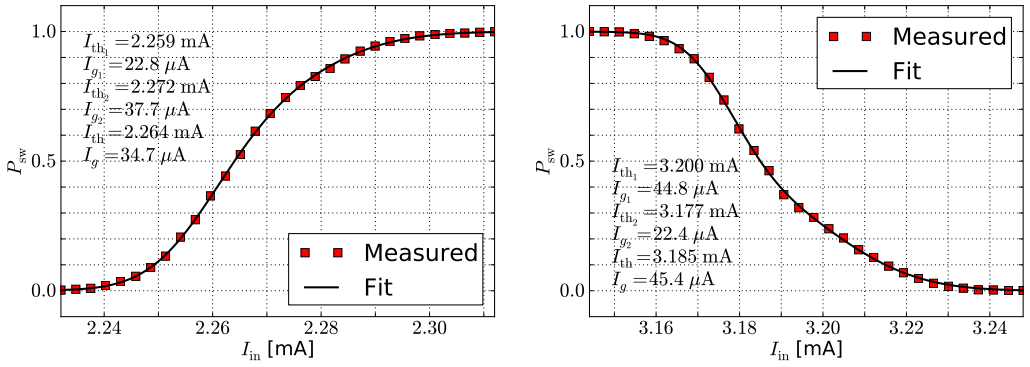

(a) (b)

Figure 5.24: Measured switching probabilities and improved fit curves, here for comparator `MC5`, for the third (a) and fourth (b) transitions.

Measured and derived quantities

Considering the quality of the fits achieved for our improved function f_{p_2}, we feel comfortable recording and reporting the observed quantities in Table 5.7. Note that comparator `MC1` was not evaluated by experiment in this way, which is explained later in Section 5.4.6.

Table 5.7: Measured performance metrics of the first batch of fabricated CQOS comparators.

	T[1]	I_{th_1}	I_{g_1}	I_{th_2}	I_{g_2}	I_{th}	I_g	I_m	γ
MC2	1	0.302 mA	23.5 µA	0.295 mA	29.4 µA	0.299 mA	27.7 µA		
	2	1.057 mA	25.1 µA	1.053 mA	43.5 µA	1.056 mA	34.4 µA		
	3	1.824 mA	29.0 isµA	1.816 mA	34.5 µA	1.821 mA	33.3 µA		
	4	2.627 mA	48.4 µA	2.635 mA	29.2 µA	2.632 mA	39.9 µA		
	Rise						30.5 µA	1.522 mA	2.00%
	Fall						37.2 µA	1.576 mA	2.36%
	Avg						33.6 µA	1.549 mA	2.17%
MC3	1	0.287 mA	37.2 µA	0.286 mA	25.4 µA	0.286 mA	31.1 µA		
	2	1.227 mA	39.5 µA	1.219 mA	27.1 µA	1.222 mA	34.8 µA		
	3	2.180 mA	23.6 µA	2.182 mA	35.8 µA	2.181 mA	29.7 µA		
	4	3.107 mA	38.3 µA	3.099 mA	24.6 µA	3.102 mA	32.8 µA		
	Rise						30.4 µA	1.895 mA	1.60%
	Fall						33.8 µA	1.880 mA	1.80%
	Avg						32.1 µA	1.888 mA	1.70%
MC4	1	0.974 mA	28.6 µA	0.987 mA	27.7 µA	0.980 mA	31.8 µA		
	2	1.924 mA	32.6 µA	1.917 mA	28.5 µA	1.921 mA	31.8 µA		
	3	2.844 mA	28.9 µA	2.848 mA	34.0 µA	2.846 mA	31.8 µA		
	4	3.835 mA	35.7 µA	3.829 mA	24.1 µA	3.832 mA	30.7 µA		
	Rise						32.4 µA	1.866 mA	1.74%
	Fall						31.3 µA	1.912 mA	1.64%
	Avg						31.8 µA	1.889 mA	1.68%
MC5	1	0.380 mA	23.2 µA	0.390 mA	36.7 µA	0.384 mA	32.7 µA		
	2	1.296 mA	23.2 µA	1.317 mA	45.3 µA	1.303 mA	45.0 µA		
	3	2.259 mA	22.8 µA	2.272 mA	37.7 µA	2.264 mA	34.7 µA		
	4	3.200 mA	44.8 µA	3.177 mA	22.4 µA	3.185 mA	45.4 µA		
	Rise						33.7 µA	1.880 mA	1.79%
	Fall						45.2 µA	1.882 mA	2.40%
	Avg						39.5 µA	1.881 mA	2.10%

[1] Number of threshold discovered in investigated range of I_{in}.

The measured thresholds yield an I_m that is as expected from parameter extractions and (4.8). The experimental I_m is derived from the combined threshold value, which we define as the value of $I_{\text{in}}$where $f_{p_2} = 0.5$, or

$$I_{\text{th}} = \frac{I_{\text{th}_1} I_{g_2} + I_{\text{th}_2} I_{g_1}}{I_{g_1} + I_{g_2}}.$$

The combined grey-zone for each transition is determined numerically. Since f_{p_2} is a sigmoid function, which is monotonically non-decreasing, there is one contiguous interval of I_{in}, which maps onto an interval of f_{p_2} that corresponds to the grey interval of f_p, roughly $[0.105, 0.895]$. Through zero-finding, here using the Newton-Raphson method as facilitated by SciPy [103], the corresponding boundaries on the I_{in} interval are established.

Note that, using these definitions, the threshold no longer necessarily lies in the center of the grey-zone.

Interpretation of results

We first note that the measured grey-zones exceed the simulation for `MC2` reported in Section 5.3.2 substantially. This was unexpected, but was subsequently explained by the likely violation of the injectivity condition (see later Section 5.5), which prevented placing the threshold at the optimal position.

Since we are confident of our fits, the estimated grey-zones and thresholds from Table 5.7 fully describe the estimated switching probability curve, we no longer need to model the quantiser operation explicitly (which saves on computation time, as finding the roots of (4.7) is no longer required). The function $\hat{q}'(\cdot)$ from Figure 4.23 was modified to employ f_{p_2} to synthesise the switching probability curve from threshold, period and grey-zone data, and model the multi-threshold comparator in this way. Each comparator was deemed fully described by the parameter tuple

$$\bar{\rho} = (I_{\text{thr}_1}, I_{gr_1}, I_{\text{thr}_2}, I_{gr_2}, I_{\text{th}f_1}, I_{gf_1}, I_{\text{th}f_2}, I_{gf_2}),$$

where I_{thr_i} and I_{gr_i} describe each rising threshold and $I_{\text{th}f_i}$ and I_{gf_i} describe each falling threshold (for $i = 1, 2$). For the purposes of our experiment, an ADC was deemed fully described by the paramter tuple

$$\rho = (n, \bar{\rho}),$$

with n remaining the number of hardware bits in the ADC.

To predict the SNR achievable by our measured comparators, we employed the worst values recorded in Table 5.7 for each comparator when populating $\bar{\rho}$. The code we employ for simulation of comparator operation is listed in Listing B.5. We make the assumption that the grey-zones are small enough that two neighbouring thresholds do not overlap. This is easily the case in the five comparators in our experimental set (were this not the case, the comparators

would probably be of little use for our intended application). The results are depicted in Figure 5.25.

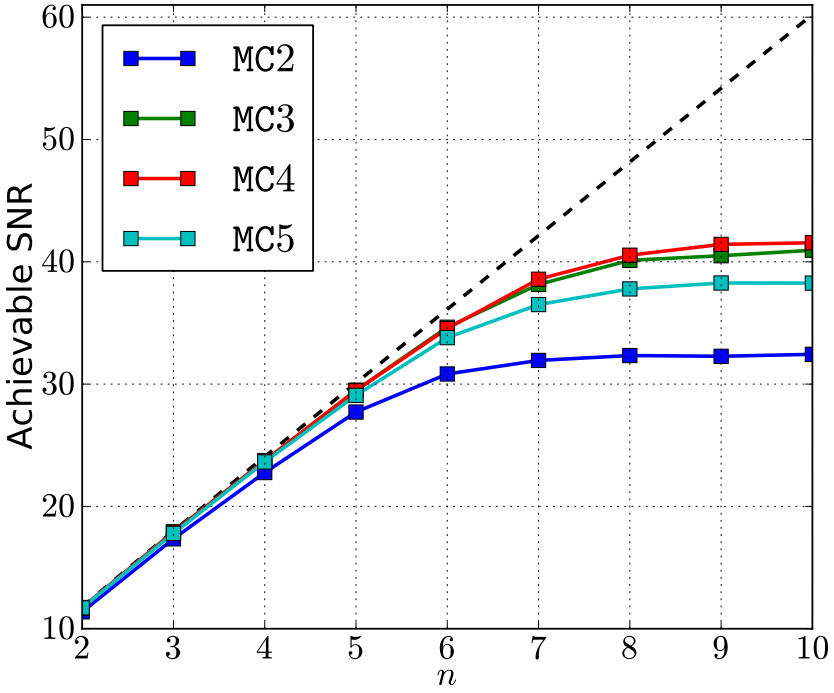

Figure 5.25: Estimated achievable SNR for the comparators in our test set.

The results are interesting, if not very surprising. `MC2` performs worst, exhibiting marked SNR depression (compared to the ideal value) even at $n = 4$, and flattening out at an SNR of approximately 32.4 dB or ~ 5.4 effective bits. The next-worst comparator is `MC5`, showing a marked drop from ideal SNR at $n = 5$ and flattening out at approximately 38.3 dB or 6.4 effective bits. `MC3` and `MC4` perform similarly, dropping markedly from the ideal value at $n = 6$ and flattening out at approximately 40.9 dB and 41.6 dB, or 6.8 and 6.9 effective bits respectively. This trend is mirrored in the recorded values for greyness γ in Table 5.7. The comparatively large gap between `MC5` and `MC2` is compounded by the fact that the rising threshold greyness of `MC2` is far worse than that of `MC5`, where the rising threshold determines the MSB.

For `MC5`, we have expected better results than recorded. This is because the co-shunted DMP arrangement has yielded significant improvements in comparator grey-zone in recent experiments by foremost international researchers [111]. We believe that we have not reached better grey-zones because our co-shunt, which should exhibit a small impedance, is burdened by a large inductive component due to design rules (the trade-off is akin to that described in Figure 5.3). For a Josephson characteristic frequency of ~ 100 GHz, the difference

between a 1 pH (desired) and a 7 pH (likely achieved) inductive component is more than 3.5 Ω, which alone exceeds the desired shunt resistance by a factor of three. A different process should improve this trade-off, potentially yielding a better comparator.

We conclude from our experiment that, when considering the grey-zone alone, all of our designed comparators should be useful in at least a 4-bit configuration, with `MC2` only just making the cut. Since other factors will probably depress achievable SNR (calibration mismatch, resistor-ladder tolerances (see Section 5.6.1), signal delay (see Section 5.6), clock jitter, etc.), it is prudent to choose a better-performing comparator. These results suggest `MC3` or `MC4` as the best choice.

5.4.6 An additional constraint: comparator range

For correct ADC operation, the comparator synthesising the LSB must correctly respond to an input signal in the range

$$I_{\text{in}} \in [-2^{n-3} I_m, 2^{n-3} I_m],$$

which corresponds to $[-2I_m, 2I_m]$ for a 4-bit ADC. The other comparators must respond linearly to a lower range.

This criterion was not explicitly respected during comparator layout, which led to an unexpected discovery during experimental verification, the essence of which is qualitatively captured in the trace from Figure 5.26.

While the criterion seems met for a 4-bit ADC, even a 5-bit ADC is probably not realisable with this comparator, due to the deterioration of the comparator response at input currents with high absolute value. This effect also led to our decision not to evaluate `MC1` experimentally, as the large currents required for the insensitive `MC1` would almost certainly suppress correct operation.

We believe that this phenomenon is a direct consequence of the input current flowing underneath the quantising junctions, as shown in Figure 5.27. Josephson junctions are susceptible to magnetic fields, and post-experiment research has revealed that even currents as small as $\sim 10\,\text{mA}$, when flowing close enough to the trilayer package, can have a significant to catastrophic effect on the junction. This effect manifests itself in a suppression of junction I_c.

For large values of I_{in}, the substantial part of the input current I_{in} flows through the quantising inductances L_{q1} and L_{q2}. These were laid out in such a way, that currents flowing through them also flow close to or directly underneath B_{q1} and B_{q2}. This should be avoided in future layouts.

To our knowledge, this effect has not been reported in literature concerning QOS or CQOS comparators.

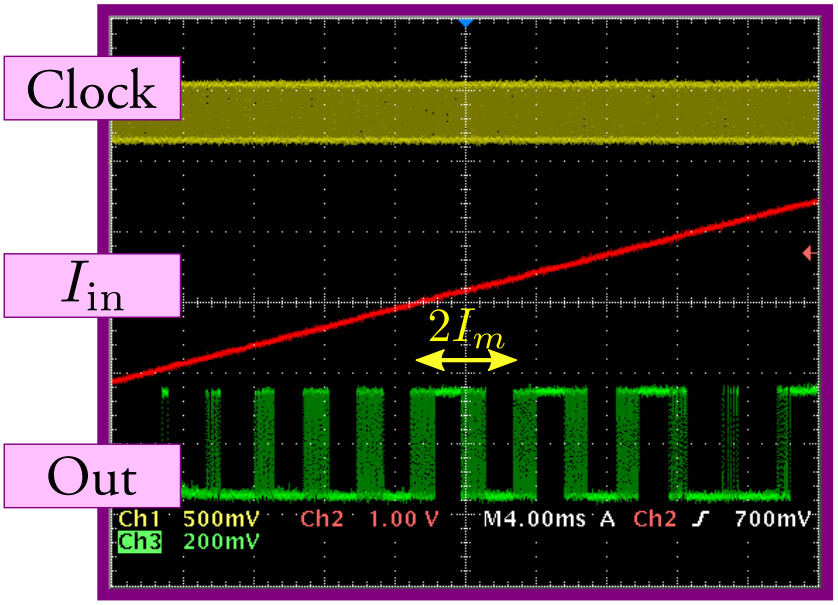

Figure 5.26: Comparator response of MC4 over multiple I_m. While $[-2I_m, 2I_m]$ seems correctly represented, outside of approximately $[-3I_m, 3I_m]$ the response is erroneous.

The trend in modern fabrication processes is to place a number of additional wiring layers under the ground plane (the ground plane is the pale yellow feature in Figure 5.27). Once we have access to such a process, care should be taken to route large currents in such a layer to place the ground plane (or a dedicated shield) between the large currents and proximal junctions.

5.4.7 Frequency of operation

The greyness, one important benchmark of our comparators, has been investigated experimentally. The second benchmark, the maximum frequency of operation $f_{\max}$, can, unfortunately, not be measured with our experimental setup.

Maximal frequency of operation is not straightforward to define. One definition could include the frequency at which the charge-carrying Cooper pairs break up. A second definition might include the clock frequency above which the comparator offers no output for any useful input. Clearly, both of those definitions yield an optimistic bound.

Recent work by Ortlepp et al [113] has revealed that the output of a Josephson decision-making pair can be correlated even when the switching probability is exactly 0.5. This is not desirable, as under those conditions the thermal noise model related in the previous chapter requires the output to be random, with the additional constraint that no output (1 or 0) is more likely than the other

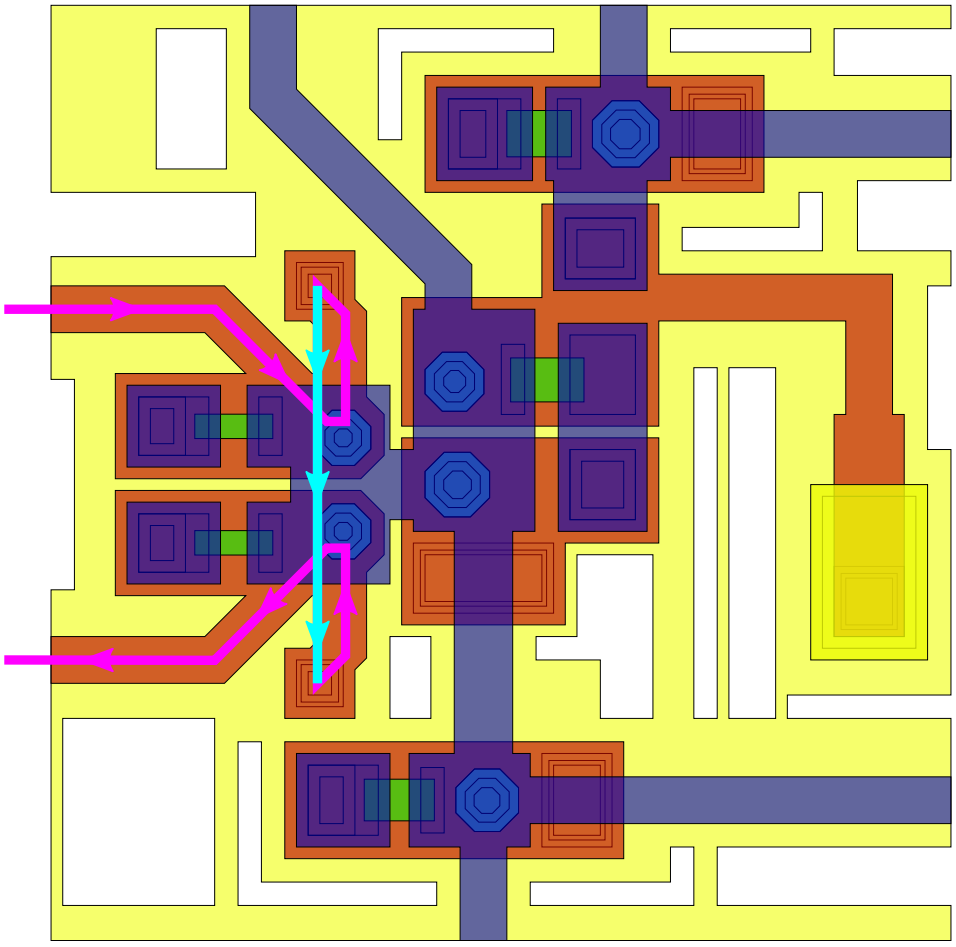

Figure 5.27: Input current flow for large I_{in} through `MC5`. The current flows underneath the quantising junctions, which probably accounts for the observation depicted in Figure 5.26.

(uncorrelated).

To understand this non-ideality better, we formalise it. The output bit sequence of a comparator $\hat{y}[k]$ consists of N elements drawn from $\{0, 1\}$. We define the correlation of the output bit vector with lag 1 as

$$c_1 = \frac{1}{N-1} \sum_{i=2}^{N} 2(\hat{y}[i] - 0.5) \cdot 2(\hat{y}[i-1] - 0.5), \tag{5.13}$$

where the $2(\cdot - 0.5)$ transformation is applied to each entry to map `0` entries to -1 and `1` entries to 1. This ensures that each entry contributes equally.

The quantity c_1 depends on a sum of terms, where each term depends on a pair of output signals. Since there are two possible entries, there are four possible pairs. Each pair, its contribution to the sum in c_1, as well as the probability of its occurrence, is listed in Table 5.8.

Table 5.8: Possible output sequence entry pairs and corresponding probabilities.

Pair	Probability, $p = P_{\text{switch}}$	Contribution to sum in c_1
(1,1)	p^2	1
(0,0)	$(1-p)^2$	1
(1,0)	$p(1-p)$	-1
(0,1)	$(1-p)p$	-1

The probabilities listed in Table 5.8 assume uncorrelated output. Considering Table 5.8 and (5.13), the ideal c_1 for a DMP can be written as a function of its switching probability p:

$$\begin{aligned} c_1(p) &= 1 \cdot P\{(1,1)\} \;+\; 1 \cdot P\{(0,0)\} \\ &+\; (-1) \cdot P\{(1,0)\} \;+\; (-1) \cdot P\{(0,1)\} \\ &= \; p^2 + (1-p)^2 - p(1-p) - (1-p)p \\ &= \; 4\left(p - \frac{1}{2}\right)^2 . \end{aligned}$$

Figure 5.28 shows a simulation of MC3 at 10 GHz and 22 GHz, with both the ideal and observed c_1 shown. Clearly the observed quantity deviates significantly from the ideal quantity, with the deviation worse at the higher frequency. Also, the deviation seems worst around the threshold.

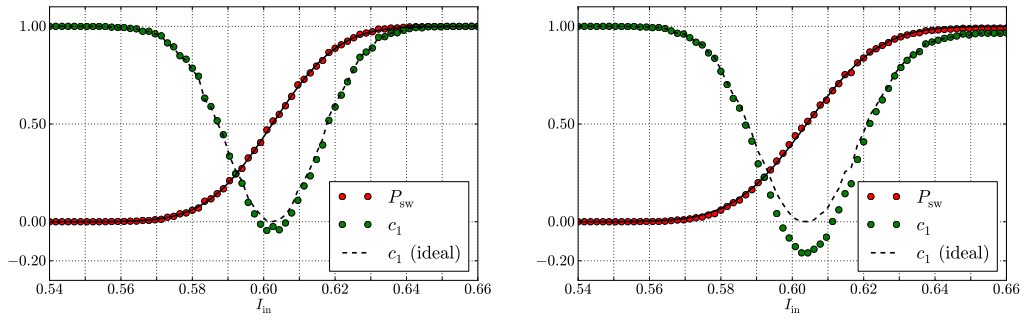

(a) (b)

Figure 5.28: Simulated switching probabilities and correlation, here for comparator MC3, at 10 GHz (a) and 22 GHz (b).

Within parts of the grey-zone, the observed c_1 is, in both cases, less than the ideal value, suggesting that alternating output bit pairs $(1,0)$ or $(0,1)$ (which contribute negatively) are favoured. This leads to the following explanation:

1. A decision, either `1` or `0`, has associated with it a certain transient response in the DMP. This response takes time to complete. This *decision time* is a random variable, the exact distribution of which depends on the input current I_{in} as well as, at least, fluctuations due to thermal noise. Intuitively, the mean decision time is longest when the signal is closest to threshold, as then the junction switching event is precipitated by a current closest to its respective critical current.

2. If the next decision is forced (by an incident clock pulse) before the transient response has died down sufficiently, the opposite outcome is favoured.

3. This introduces what we term the *comparator memory effect*, which captures that the comparator output depends on its history. This is undesirable and also effectively leads to a wider grey-zone.

Naturally, this memory effect is dependent on the clock frequency f_c. At high frequencies it is more likely for a decision to be influenced by a previous decision. We believe that this memory effect would cause comparator grey-zones to widen at higher frequencies.

In addition, when the output bits of a comparator are correlated, a loss of information occurs. This can be quantified in different ways. We choose to define the number of corrupted output bits as $N|\Delta c_1|$, where Δc_1 is the difference between observed and ideal c_1, which leads to the definition of maximum useful data rate

$$f_{\max} = f_c \left(1 - |\Delta c_1|_{\max}\right),$$

where $|\Delta c_1|_{\max}$ is the maximum of $|\Delta c_1|$, generally observed near the comparator threshold.

A conservative approach would be to claim that the lowest maximum achievable data rate implies the bandwidth recoverable from the digitised signal. This is conservative because this maximum data rate is only substantially lower than the clock frequency when the comparator is in its grey-zone, outside the grey-zone corruption of bits is very rare.

Applying this definition, the condition

$$f_{\max} > 20\,\text{GHz} \tag{5.14}$$

should be ensured for each comparator. However, any increase of f_c places greater effort on the backend, in which initial processing must occur on all bits produced.

The quantities measured in Figure 5.28b yield $f_{\max} = 18.5\,\text{GHz}$. Even when increasing f_c to 24 GHz, the observed $f_{\max}$ increases only to 19.0 GHz. Hence, the condition (5.14) does not seem achievable for our fabricated comparators, at least not over a fabricatable range of parameters. Even if it is reachable for

a certain corner of the parameter space, a substantial penalty would be levied onto the grey-zone as well as the backend.

Fortunately, this is readily remedied by increasing J_c. Circuit speed generally increases as $\sqrt{J_c}$ [9]. A well-known, readily available commercial foundry exists, Hypres Inc., which supplies a $J_c = 4.5\,\mathrm{kA/cm^2}$ process (which we make use of for a different application in Chapter 6. A simulation of `MC3` with $J_c = 4.5\,\mathrm{kA/cm^2}$ yields the c_1 depicted in Figure 5.29 for $f_c = 21\,\mathrm{GHz}$. Here, $f_{\max} \approx 20.2\,\mathrm{GHz}$, which satisfies (5.14) without placing as much strain on the backend as for $J_c = 1\,\mathrm{kA/cm^2}$.

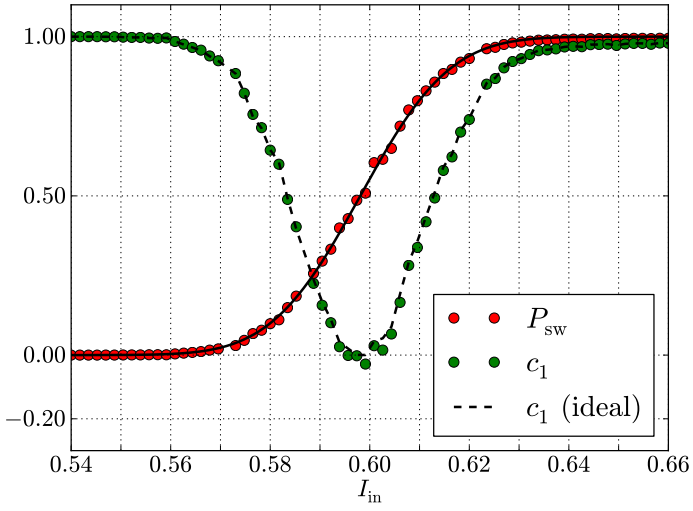

Figure 5.29: A simulation of `MC3` adjusted for $J_c = 4.5\,\mathrm{kA/cm^2}$ at $f_c = 21\,\mathrm{GHz}$. The observed $f_{\max} = 20.2\,\mathrm{GHz}$.

In this context, an increase in J_c has at least one detrimental effect (there are also technological difficulties). In Chapter 4 it was found that $\nu \propto I_c$, meaning that a greater J_c makes it more difficult to achieve low ν. Fortunately, the minimal I_c in Hypres's $4.5\,\mathrm{kA/cm^2}$ process is lower even than that in the IPHT's $1\,\mathrm{kA/cm^2}$ process, avoiding this concern for our application.

5.5 Redesigned Comparators

During design, fabrication and experimental operation of our comparators, we have explored the limits of the theory and the fabrication process, attempting to quantify them and learn their effect on eventual comparator operation. We have also run into several issues that should be addressed before further experimentation. In this section we relate our efforts toward this goal.

5.5.1 Summary of issues

Injectivity

The injectivity condition (4.25) was probably violated for our measured comparators. There are two reasons for this:

1. A design-rule miscommunication (compare design rules [13] and therein cited design kit) forced a last-minute change of the quantising junctions from $I_c = 100\,\mu A$ to $I_c = 125\,\mu A$, without time to redesign the comparator. This suppressed the available loop inductance by 0.66 pH, and since the design was such that β_L is as high as possible without violating injectivity, it is likely that the condition was violated in the fabricated circuits. We have not found a way to reliably measure this (any introduced hysteresis was too small to be measured with our setup), but it was not straightforward to achieve correct comparator operation (equidistant transitions, low grey-zones) trivially, fine-tuning of the bias currents was required, suggesting that the quantising junctions were in the switching regime and a sub-obtimal gradient over the threshold was achieved.

2. Measured comparator periods I_m were slightly lower than simulated, suggesting a small increase in L_q from the extracted and simulated value. While this would not have been problematic in and of itself, it compounds the problem related in 1, especially since it suggests that parasitics were also slightly higher than predicted by extractions.

Magnetic field effects

A clumsy layout choice has exasperated the problem of critical current suppression brought on by magnetic field effects. We note here that layouts published in the literature appear to have made similar choices (but do not cite them here). The reasons for this choice, besides inexperience on the part of the present author, are twofold:

1. Design rules and the superconducting properties of niobium make layout of small inductors over large distances problematic. In fact, this is a suggested reason for the failure of high-T_c superconductors to play a larger role in superconducting electronics [43]. Here this property makes it difficult to realise the quantising inductance L_q by, at the same time, carrying the large current (which flows primarly through this inductance) away from the active junctions.

2. For similar reasons, it is difficult to place the node at which the quantising junction and the quantising inductance meet far away from the actual junction (without punishing increases to the parasitic component of the

loop inductance), forcing a large amount of current to make a change of direction near the Josephson junction.

High-speed measurements

There is a significant and non-trivial relationship between comparator output correlation (and other performance metrics) and clock frequency. Literature suggests that substantial effects should be measured at high frequencies [112, 113]. Our first batch of experimental comparators was not suitable for high frequency testing, as this requires either high-frequency room temperature equipment, which is very expensive and has many limitations, or a significantly larger on-chip testbed, which is not desired for first tests due to its potential effects on the results. Preceding a high-speed experiment with a safer low-speed experiment is common practice in the community of applied superconducting electronics.

5.5.2 Next experiments

Desired observable quantities

For the next experimental suite of comparators, we desire the following observable quantities:

1. Switching-probability curve for both rising and falling thresholds, also at high frequencies,
2. correlation (c_1) curve for both rising and falling thresholds, also at high frequencies, and
3. comparator operation at high frequencies across a large range of I_{in}.

New layouts

A second suite of comparators was laid out and is, at the cut-off time for this text, in fabrication. This suite consists of one QOS and one CQOS comparator, resulting from the following design choices:

- The chosen β_L for both comparators is identical, designed at 0.9.
- A quantising inductance of $L_{q(i)}$ of approximately 0.8 pH was chosen for both comparators, yielding a designed I_m of ~ 2.5 mA, somewhat higher than that chosen for the previous CQOS comparators (`MC2` through `MC5`) and much lower than that chosen for the previous QOS comparator `MC1`.

One objective of these designs is to obtain a direct comparison of QOS and CQOS comparators. The layout for `MX`, the QOS comparator, is depicted in

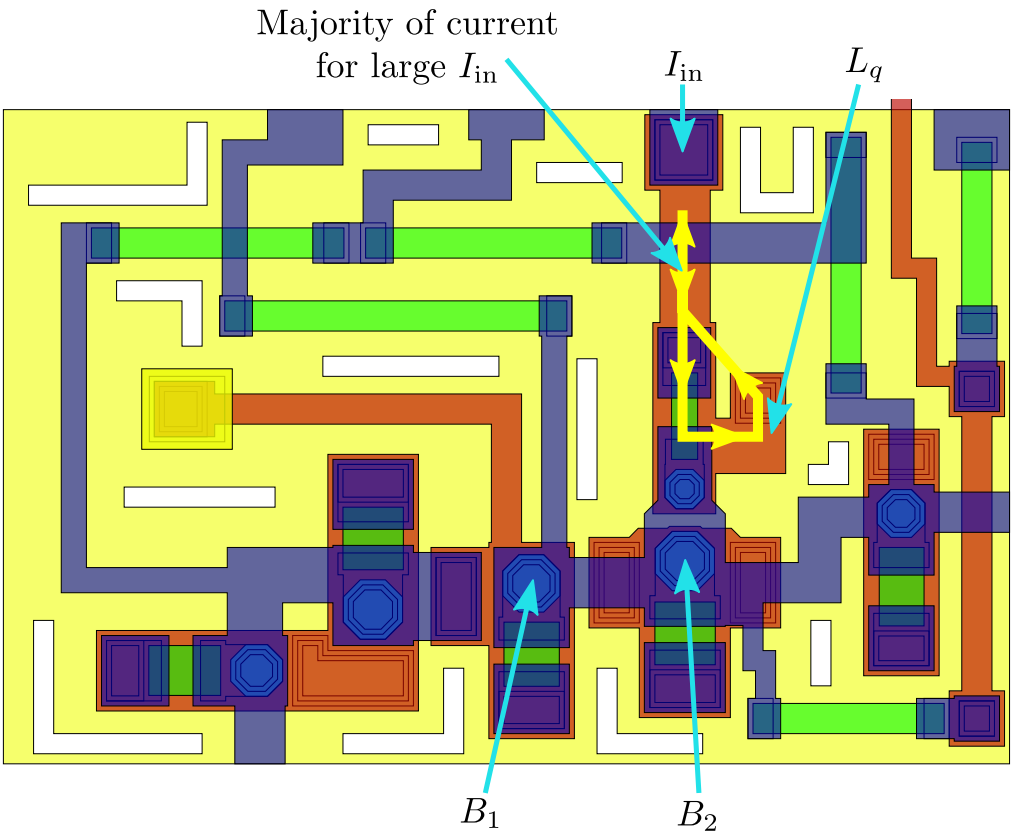

Figure 5.30: MX cell layout, new experimental QOS comparator, currently in fabrication. Details, such as extracted parameters, can be found in Section A.1.6.

Figure 5.30. The layout for MCX, the CQOS comparator, is depicted in Figure 5.31. The layout of the testbed is reproduced in Figure 5.32.

The testbed is shared by both experimental comparators. They are connected "in series", so that an input pulse first clocks MCX, and then its output pulse clocks MX. This enables testing both comparators in a substantially equal environment, while minimizing pinout cardinality. When testing MX, the switching probability of MCX must be fixed at 100%, which is generally easily achievable. For testing MCX, MX must be biased in an analogous way.

The testbed supports both a slow and a fast clock. Its operation is explained by means of the signal traces in Figure 5.33. The fast clock is triggered by a standard signal applied to a DC-SFQ converter (which repeats at a low frequency). The fast clock takes the shape of a train of 4 SFQ pulses. The 4-SFQ-pulse train is generated by a "ladder multiplier", which consists of 4 splitters and 4 mergers. The distance in time between those SFQ pulses (the inverse of the frequency of the high-speed clock) can be adjusted by means of I_{lad}.

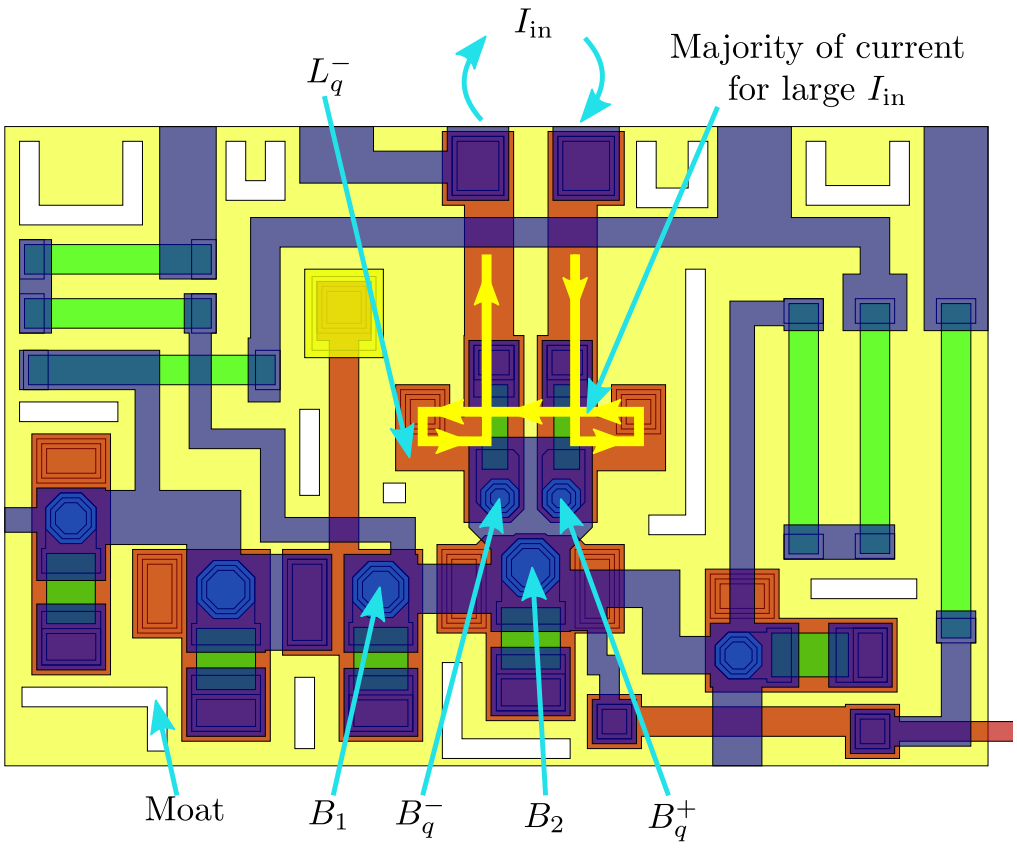

Figure 5.31: `MCX` cell layout, new experimental CQOS comparator, currently in fabrication. Details, such as extracted parameters, can be found in Section A.1.7.

We foresee the following experiment, which requires only low-speed room temperature equipment:

1. Comparator `MX` is biased into transmit mode, where it has a switching probability of 1. Now the response of `MCX` can be measured.

2. A desired point of operation for `MCX` is obtained roughly, by means of varying biases and observing oscilloscope traces such as in Figure 5.21. Fine-tuning is achieved by means of plotting a switching probability curve over I_{in} with a setup akin to that shown in Figure 5.20. An analogous process was employed in our previous cryogenic experiment.

3. A discrete sweep of I_{in} is performed where, for each I_{in}, a discrete sweep of I_{lad} is performed, for each element of which:

   a) The fast clock is triggered once (any immediate output is ignored). A string of 4 comparator output bits is now stored in the 4-bit shift

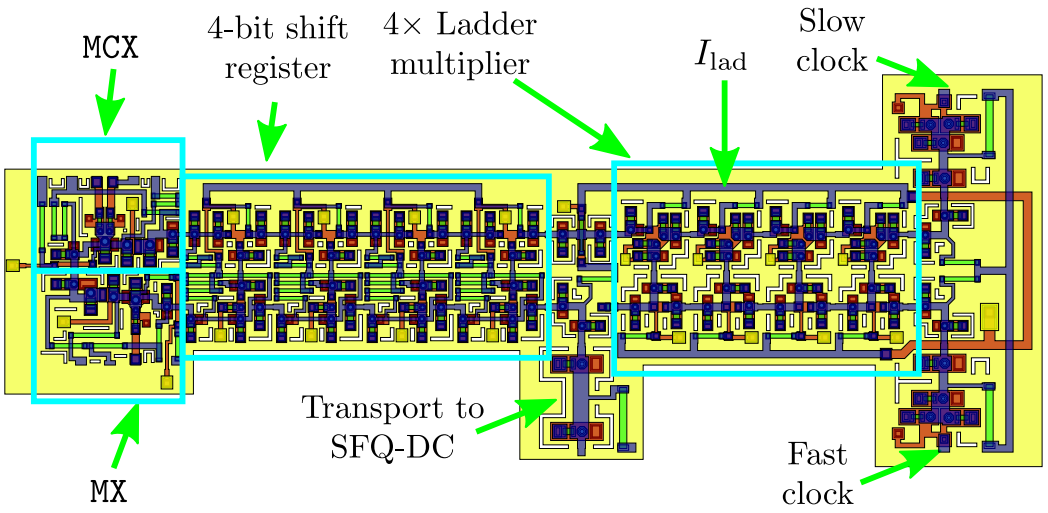

Figure 5.32: Testbed for `MC` and `MCX`, currently in fabrication. Both comparators share a high-speed testbed, due to limited available chip pads.

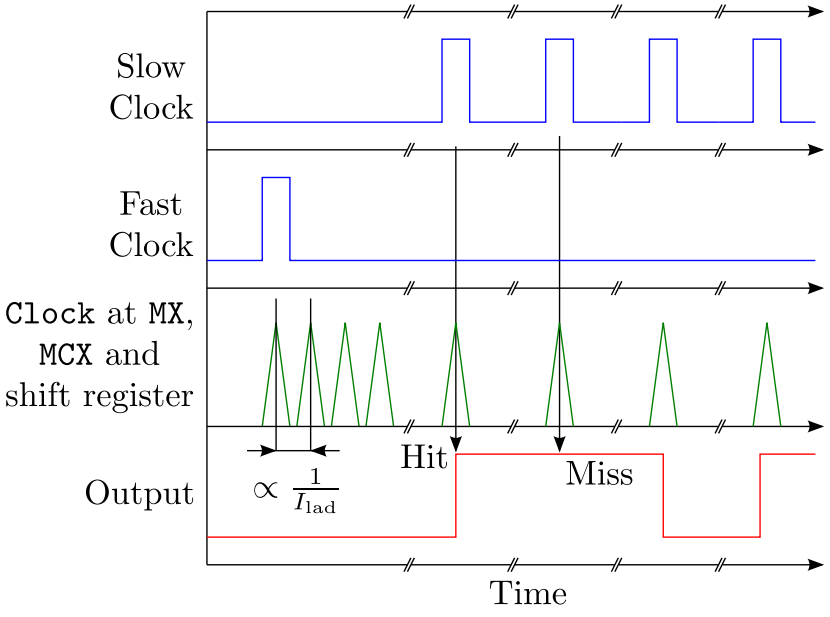

Figure 5.33: Traces of possible experimental signals for the testbed from Figure 5.32. Blue signals are input signals, green signals are (unobservable) internal signals and red signals are observable output signals.

register.

b) The slow clock is triggered 4 times, reading out the shift register, observing and recording the output each time.

c) Repeat this $\frac{N_{\text{tot}}}{4}$ times.

4. For each I_{in} and I_{lad}, the switching probability is calculated by comparing the number of hits N_{hit} to N_{tot}, as done in our previous experiment.

5. For each I_{in} and I_{lad}, the correlation c_1 is calculated by comparing successive bits in the output sequence. Each 4-bit output word then generates 3 entries of equation 5.13.

6. Switching probability and c_1 curves can now be plotted for each I_{lad}, or frequency. Function f_{p_2} is fit to obtain the parameters $\bar{\rho}$ describing the switching probability.

7. From $\bar{\rho}$ and the (minimum) c_1, bounds on achievable SNR and maximum operating frequency follow.

8. The experiment is repeated for `MX`, with `MCX` biased in transmit mode.

5.6 Applying the signals

While the comparators are the key components of any superconducting flash ADC, other components of it are far from trivial. Distributing, scaling and applying the signals required to operate the ADC correctly is an intricate problem with competing constraints. In this section we present a quick overview of the important considerations and estimate some of the design constraints involved.

5.6.1 Resistive divider

A 4-bit ADC based on a multi-threshold comparator was depicted in Figure 4.10. A resistive divider is employed to divide the input signal so as to apply the correct proportion of the signal to the correct comparator.

Such a resistive divider is easily realised on an integrated circuit, as the superconducting processes we investigate also provide a layer of material that is Ohmic at the targeted operating temperature. This layer is provided specifically to enable reproducible resistors. In every damped junction (at least for low-T_c processes), resistive structures in this layer are employed to connect an appropriate damping resistor in parallel with the junction.

Physically these resistors resemble normally-conducting microstrip lines over a superconducting ground-plane. A sheet resistance R_s is associated with the resistive layer employed. Any rectangular structure in that layer then exhibits a resistance of

$$R = \frac{\ell}{w} R_s, \tag{5.15}$$

where ℓ and w are the length and width of the structure respectively, and the length is the dimension in the direction of the current.

A resistor ladder as depicted in Figure 4.10 is thus easily realised as a corresponding network of such resistive features. The input impedance, as seen by the applied analog signal current, is simply R. The choice of R does not affect the functionality of the resistor ladder in circuit theory terms, at least in the static case. However, the following concerns should be borne in mind:

- Heat dissipation in the resistor ladder scales linearly with R (as the required current depends on comparator parameters),
- impedance matching between the ADC and the arrival medium of the signal will depend on R, and
- the length ℓ of the resistive features scales linearly with R (for a fixed choice of w).

In terms of impedance matching, one useful choice is $R = 50\,\Omega$, as $Z_0 = 50\,\Omega$ is a typical characteristic impedance of many coaxial cables. Coaxial cables are a popular choice of transmission medium for high-frequency signals.

For our chosen architecture, the full range of the ADC is spanned by ν, with a maximum input current amplitude of $I_{\max} = \frac{1}{2}\nu = 2^{n-2}I_m$. Power dissipation in the resistor ladder is thus bounded by

$$P_{\text{lad}} \leq I_{\max}^2 R = 2^{2(n-2)} I_m^2 R.$$

For a 4-bit ADC based on MC3, with R = $50\,\Omega$, this amounts to $\sim 4\,\text{mW}$. While certainly significant in a cryocooled environment, this is entirely acceptable in contemporary cryocoolers.

The same ADC based on MC1, however, will dissipate up to $P_{\text{lad}} \sim 150\,\text{mW}$ in the resistor ladder. This already exceeds the maximum heat lift of small-scale cryo-coolers and could thus dictate the power-budget of the system. Furthermore, although cryo-coolers with greater heat lift exist, the resistor ladder may still pose a problem if it is concentrated in a fairly small area, as in that case the required heat lift can possibly not be realised across it. Careful analysis of this requirement is necessary for actual implementation of such a resistor ladder.

Furthermore, care should be taken to minimize resistor temperature fluctuations as the signal power fluctuates, especially as different fractions of the signal are conducted by different resistors.

Note that a reduction of R is certainly possible when employing more sophisticated impedance-matching practices.

5.6.2 Compensation for signal lag

A physical resistor ladder can be assembled from chunks of identical resistive features of resistance R, interconnected by superconducting metals. This is straightforward. The design rules for the IPHT RSFQ1F process recommend a width of $w = 10\,\mu\text{m}$ for current-carrying resistors. The sheet resistance is specified as $R_s = 1\,\Omega$. For $R = 50\,\Omega$, this requires an R-chunk to be

$$\ell = 500\,\mu\text{m} \tag{5.16}$$

long, according to (5.15). A 4-bit resistor ladder would thus comfortably fit onto a $5 \times 5\,\text{mm}^2$ IC.

However, as this resistor ladder represents a transmission line network, delays are imposed on parts of the signal. If the signal at the input of the ladder is described by $x(t)$, the signal at another point in the network is better described by $x(t + \tau_d)$, where τ_d is the time it takes the signal to propagate from the input.

The electromagnetic properties of the insulators (dielectrics) on the IC are such that signals propagate at approximately $\frac{1}{3}c_0$, with c_0 the speed of light

in a vacuum. The signal arrives at the LSB comparator first. Considering the chunk length estimation from (5.16), the signal arrives $\tau_{dc} \approx 5\,\text{ps}$ later for each subsequent bit. For a 4-bit ADC, the signal lag would thus be $3\tau_{dc} \approx 15\,\text{ps}$ between the first and last comparator. The maximum frequency in our band of interest, $f_{max} = 10\,\text{GHz}$, has a corresponding period of 100 ps. Hence, this signal lag is certainly significant.

The clock is distributed to the comparators by a JTL network. This is illustrated in Figure 5.34. The speed with which an SFQ pulse propagates across a JTL depends on the ratio between the bias current I_b and the critical current I_c of the junctions employed (see Section 2.3.2). Hence, if the JTL section delivering the pulse to a comparator is isolated from the remainder of the JTL network, that section can be slowed down by lowering its bias voltage. Figure 5.34 shows each comparator fed by a separately biased JTL, allowing the phase of the applied clock to be adjusted. The required adjustment factors should easily be achievable.

One may, of course, formally calculate the correct bias resistors for the different JTL sections to compensate for signal lag. However, as fabrication is subject to tolerances, it is recommended to bias the JTLs separately. This allows one to adjust the compensation factors in situ, maximising achievable SNR.

Naturally, this increases the pinout of the ADC system. Whereas this may be acceptable for testing purposes, the increased bias lines increase the load on the cryocooler.

5.7 Chapter summary and conclusion

In this chapter, several QOS and CQOS-based comparators suitable for the flash ADC architecture examined in the previous chapter were designed and experimentally tested.

1. QOS quantisers work as modelled in Chapter 4. The sampling function introduces transient behaviour in the form of SFQ pulses, but these escape with minimal effect through the quantising junction.

2. Process characteristics limit the sensitivity achievable by QOS quantisers. Whereas this may not be a problem in terms of available signal power, it could impose unacceptable heat loads through the resistor ladder and may impose the need to shield active Josephson junctions from large input currents.

3. Complementary quantisers result in a steeper gradient of I_q across the comparator threshold, lowering the effect of threshold jitter and, consequently, improving the effective grey-zone. As a consequence, the greyness of CQOS comparators is better than that of QOS comparators. In

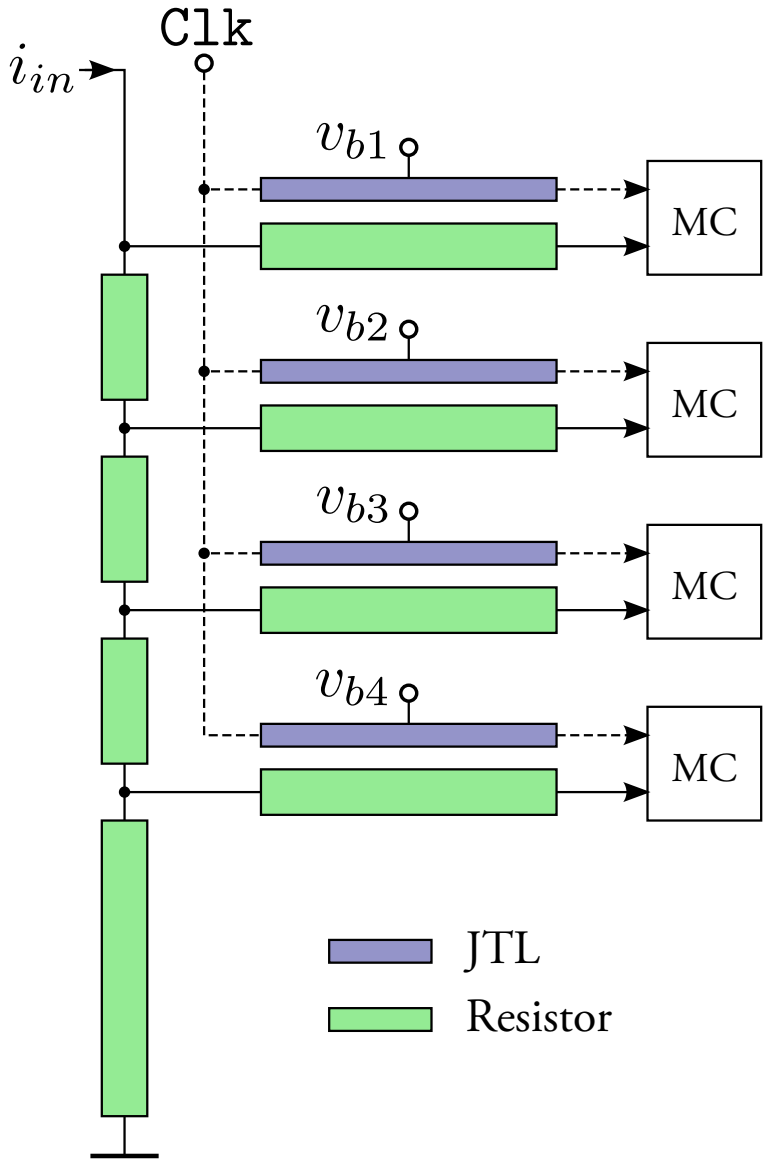

Figure 5.34: Applied signal for parallel comparators.

theory, this allows for a lower I_m before the grey-zone becomes problematic, but a lower bound is imposed onto I_m by parasitics.

4. Co-damping the Josephson comparator, as first reported by Ebert [118], may be effective in lowering the grey-zone while maintaining sufficient comparator switching speed, but relies on suitable process characterisitics, which we have not yet had access to.

5. Our experimental comparators did not perform quite as well as expected, but still outperform any expected QOS comparator (see curves in Figures 4.26 and 4.29). The discrepancy between simulated and experimental results has been explained and should prove remedied in subsequent experiments. We note that experimentally confirmed grey-zones of QOS comparators have recently been reported [119], but do not describe QOS comparators useful for our application.

6. The $J_c = 1\,\text{kA}$ process probably does not suffice to reliably digitise a 10 GHz signal (though this is a near miss), but employing a (readily available) $J_c = 4.5\,\text{kA}$ process should remedy this with conservative estimates showing that less than a 5% increase over the ideal clock frequency of 20 GHz required.

7. The resistor ladder employed imposes a significant heat load, but this still seems acceptable for the experimentally verified comparators.

8. Signal lag imposed by the resistor ladder can be compensated by separating the bias supply of sections of the clock JTL network.

Conclusion: QOS and CQOS comparators can be realised in the processes targeted for this application. The achieved greyness exceeds the requirements for a 4-bit ADC, even for our first round of fabricated comparators, with ample room for improvement expected. Process parasitics impose a lower bound on the comparator period I_m of $\sim 1.5\,\text{mA}$, suggesting $\nu \sim 12\,\text{mA}$ for a 4-bit ADC. We believe our second round of comparators, with $I_m \sim 2.5\,\text{mA}$ will perform better and thus be more suitable for the SKA, yielding $\nu \sim 20\,\text{mA}$. Input impedances from several Ω up to $\sim 100\,\Omega$ should be readily achievable. Readily available contemporary processes should achieve the required clock frequencies.

5.8 Chapter research output

1. An article in the international journal *IEEE Transactions on Applied Superconductivity* is in preparation. This article will be titled "Performance metrics of CQOS comparators" and will be submitted later in 2013.

2. A conference poster was exhibited on 8 July 2013 at the International Superconducting Electronics Conference in Cambridge, MA, USA. The authors include the present author, his supervisor, as well as staff at the Institute of Photonic Technology in Jena, Germany.

   M.H. Volkmann, T. Ortlepp, O. Wetzstein, R. Stolz and C.J. Fourie, "The grey-zone of CQOS Comparators (with experimental results),"

   Abstract:
   Complementary Quasi-One-Junction-SQUID (CQOS) comparators address the sensitivity bound which characterises Quasi-One-Junction SQUID (QOS) comparators. Multi-threshold comparators, such as QOS and CQOS comparators, are the key component of superconductor flash ADCs. Such ADCs are characterised by high bandwidth, but low sensitivity. A more sensitive multi-threshold comparator is thus valuable. In previous work, we have investigated the effects of circuit topology and parameters on the gray-zone of CQOS comparators and have defined as figure of merit their grayness, which is a measure for the suitability of a multi-threshold comparator for flash ADCs. In this work, we experimentally verify our conclusions by measuring the gray-zone of CQOS comparators manufactured at the FLUXONICS foundry.
   Acknowledgements:
   This work was supported by the Square-Kilometre Array Project.

   An article is about to be published on IEEE Xplore (not peer-reviewed) in conjunction with a poster that was exhibited at the International Superconductor Electronics Conference '13 in Cambridge, MA, USA, on 8 July 2013. The authors include the present author and collaborators from Germany and Japan.

   T. Ortlepp, M.H. Volkmann and Y. Yamanashi, "Memory Effect in Balanced Josephson Comparators,"

   Abstract:
   The performance of a balanced Josephson comparator is measured by its gray-zone and its maximum operation frequency. A typical effect at high clock frequencies is the correlation of output bits with their predecessors, the comparator develops a memory. This is undesirable, as it imposes an upper limit on the effective data rate achievable by the comparator. In this work, we define the memory effect of a Josephson comparator and study its influence on the maximum effective data rate.

3. A conference poster was exhibited on 8 October 2012 at the Applied Superconductivity Conference in Portland, OR, USA. The authors include

the present author, his supervisor and co-supervisors, as well as staff at the Institute of Photonic Technology in Jena, Germany.

M.H. Volkmann, C.J. Fourie, O. Wetzstein, J. Kunert, D.B. Davidson, W.J. Perold and H.-G. Meyer, "Gray-zone of Complementary Quasi-One Junction SQUID Comparators,"

Abstract:
Superconducting flash ADCs have long been an envisioned application of superconducting electronics. Whereas manufacturing tolerances generally limit the performance metrics achievable by flash ADCs, the domains of high bandwidth and moderate to high sensitivity applications remain targets of this architecture. Most published implementations rely on the quasi one-junction SQUID (QOS) multi-threshold comparator, which, in real layouts, imposes a trade-off between linearity and sensitivity. Furthermore, the comparator gray-zone determines the lowest significant bit achievable.
The recently introduced complementary QOS (CQOS) comparator addresses the linearity-sensitivity trade-off. The potential for achieving high bandwidth combined with high sensitivity suggests applications in the front-end of radio astronomy receivers. We attempt to identify the advantages that CQOS designs could offer to an experimental digital front-end for the Square-Kilometre Array. In particular, we study the gray-zones achieved by different CQOS implementations. This paper will discuss experimental results for circuits fabricated at the FLUXONICS foundry.
Acknowledgements:
This work was in part supported by the South African National Research Foundation, Grant No. 78789.

Chapter 6

Ultra-Low-Power Data Transport

6.1 Introduction

We are now confident that we can build the key components of a superconducting ADC of substantial benefit to the SKA. The output of the ADC is in the form of SFQ pulses, which cannot be detected reliably by conventional technology (at room temperature). Responsible for this are the high bandwidth and low energy content of the output pulses. Both make it impractical to bridge the thermal barrier (get the signal out of the cryocooler). Even a cryogenic amplifier based on, say, CMOS technology is unsuitable for single SFQ pulses.

Special interface circuits have been developed to enable transmission of output bits across the thermal barrier. These can achieve the goal, but have some limits. They are further discussed in Section 6.3.3.

In this Chapter, we discuss the development of a novel logic family, termed "eSFQ" logic, to achieve the following goals:

- Provide a means to distribute ADC output to room-temperature devices, and
- develop a low-power logic platform suitable for eventual backend processing of the SKA data.

It is noted here that much of the equipment, consumables and work was funded by Hypres, Inc. of Elmsford, NY, USA, who also holds a substantial interest in ultra low-power logic families. Hypres is credited with the original idea of eSFQ logic.

6.2 Motivation

6.2.1 Need for deserialisation

Even if the output from an SFQ circuit can be made to cross the thermal barrier, the high output sample rate of the ADC architecture we are targeting presents a problem for conventional technology expected to acquire and process the ADC output. Single-bit input channels to modern FPGAs are capable of several gigabit per second [124]. The clock frequencies of our proposed systems exceed this capability by an order of magnitude, as shown in Figure 6.1.

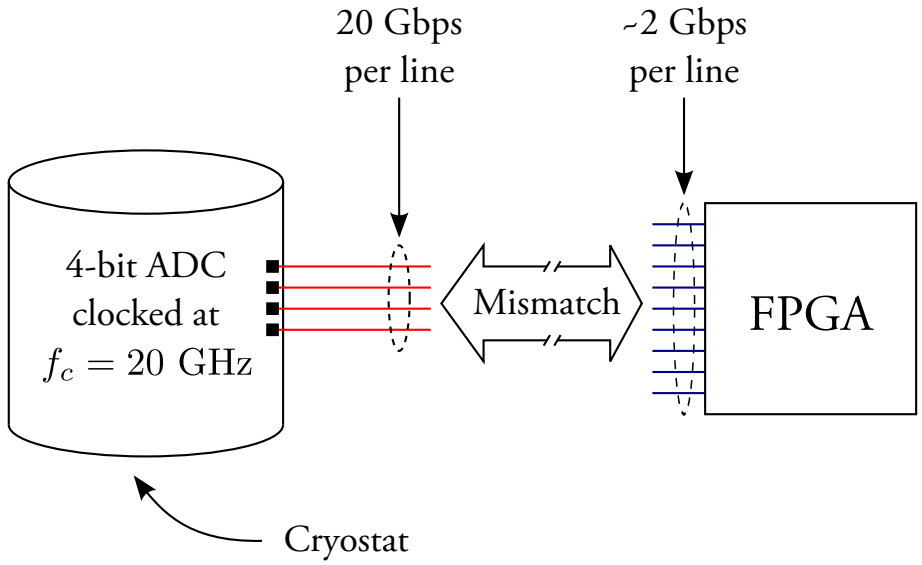

Figure 6.1: Inability of conventional technology to acquire ADC output.

In fact, modern FPGAs are capable of achieving multiple giga*byte* per second when employing parallel channels with certain encodings (like 8-to-10 encoding [125]). However, performing such an encoding on the superconducting chip is currently impractical due to integration limitations.

Conventional technology achieves far superior integration levels, even optimistic projections for superconducting technology do not dare predict beyond 10 million junctions per $5 \times 5\,\text{mm}$ chip on a five-year timescale. High integration means that, although the maximum acquisition rate of a single link is comparatively low, many links are available. It is conceivable that FPGAs of the SKA era will be able to acquire a digital signal of sufficient bandwidth. Even if not, several currently available FPGAs [124] could be combined in parallel to achieve this goal. This places a burden on the superconducting technology to spread out the signal over sufficiently many links. Our efforts at this are summarised in Section 6.5.

6.2.2 Need for low power dissipation

In current research, much scrutiny concerns the power dissipation of digital logic circuits. Ever greater integration and ever greater demands for computing power are steadily advancing high-performance computing towards the coveted exa-scale (10^{18} operations per second) boundary.

The SKA, at least according to current specifications (see Section 2.2.3), requires exa-scale computing power to process the raw data output. One major concern is the power requirements of such a system. With contemporary technology, projections place these requirements in the high 10s to over 100 MW [21, 126]. This is a major quantity of power and will incur considerable expense in terms of infrastructure and, arguably, environmental stock.

Perhaps even more importantly, distributing such a large amount of power brings with it immense EMC concerns. This is especially pertinent in the vicinity of an instrument as sensitive as the SKA. It may be possible to locate the central processing facility sufficiently far away from the SKA receivers to avoid this, but that will impose greater requirements on the data transport network. This trade-off should be closely examined, but that is outside the scope of our work.

It is clear that ultra-low-power digital circuits are of great value to the SKA. The low switching energy of Josephson junctions make digital logic based on Josephson technology eminently suitable for ultra-low-power digital logic.

Early and current RSFQ logic cells, however, severely dilute this benefit by employing a lossy bias-current distribution network. We exploit the need for a large (compared to the ADC comparators) deserialiser as a platform to develop first gates of a new, ultra-low-power logic family, eSFQ. The potential applications of eSFQ logic are profound, extending beyond even the hyper-ambitious context of the SKA.

Our efforts and successes in this crucial endeavour are presented in the remainder of this chapter.

6.3 RSFQ power requirements

RSFQ logic was invented at a time where power dissipation concerns played second fiddle. The focus was on speed. The low-power aspects of the logic family were known, but not developed extensively. Care was taken to achieve circuit stability and high clock frequencies with minimal fabrication load (number of junctions).

This mindset no longer exists in serious large-scale digital logic applications. Today power dissipation per unit performance is a crucial figure of merit [126].

6.3.1 Categories of power dissipation

In SFQ circuits including all types of RSFQ and RQL [127], digital information is encoded, processed, stored and transported by means of single flux quanta (SFQs). Physically, SFQ circuits comprise a network of active Josephson junctions, and passive lossless and low dispersion superconducting strips forming inductors and microstrip lines. The energy is spent when an SFQ traverses a Josephson junction, causing a discrete 2π superconducting phase slip. This process is equivalent to the regeneration of a quantized voltage pulse $\int v(t)dt = \Phi_0$. The energy consumed during this event is dissipated in the R of the Josephson junction (see the RCSJ model introduced in Section 2.3.2) and requires replenishment for continuous circuit operation.

In RSFQ, energy replenishment is achieved by the injection of dc bias currents at strategic locations. The bias currents injected at the injection point typically do not ever change during circuit operation, although they may be "rerouted" by the deposition of flux quanta to facilitate decision-making, as is done in the DFF gate (see Figure 6.2, Section 2.3.3).

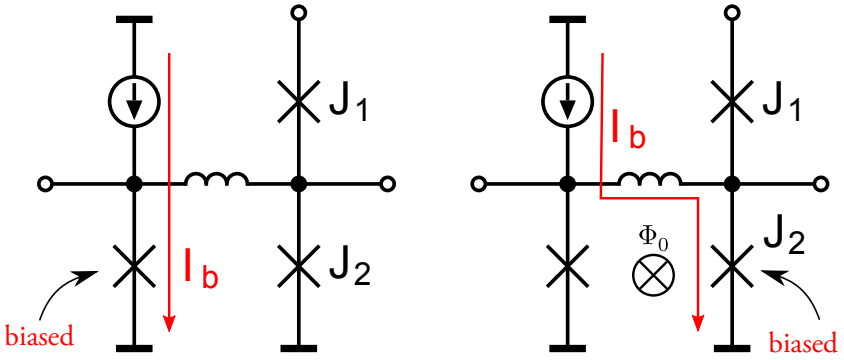

Figure 6.2: Deposited flux quanta reroute bias currents, leading to different biased junctions in a non-storing (a) and storing (b) DFF.

In RSFQ logic circuits, power dissipation can be classified into two categories: *static dissipation* and *dynamic dissipation*. Static dissipation occurs whether the circuit is active or quiescent, whereas dynamic dissipation occurs only when the circuit is active. An active circuit refers to one that currently has flux quanta propagating through it, whereas in a quiescent circuit they are fixed in space (not propagating through the biased circuit, but possibly stored somewhere).

Static dissipation occurs in the bias network, as bias currents are always flowing. Static dissipation may also occur in the logic circuitry, but this is rare

for superconducting logic families, as, in the quiescent state, currents in the logic circuitry flow only across superconductors.

6.3.2 Resistor networks

RSFQ logic, unlike some other superconducting logic families, relies on dc bias currents. These are employed to ratchet up the phase of some grounded junctions to subcritical level, so that they switch in the presence of an applied SFQ pulse. From the earliest days of RSFQ logic, dc bias currents were divided and distributed by a resistor bias network. Superconductor IC technology does not provide a known elegant method of realising a current source, thus the current is injected from off-chip. An RSFQ circuit template biased in this way is depicted in Figure 6.3.

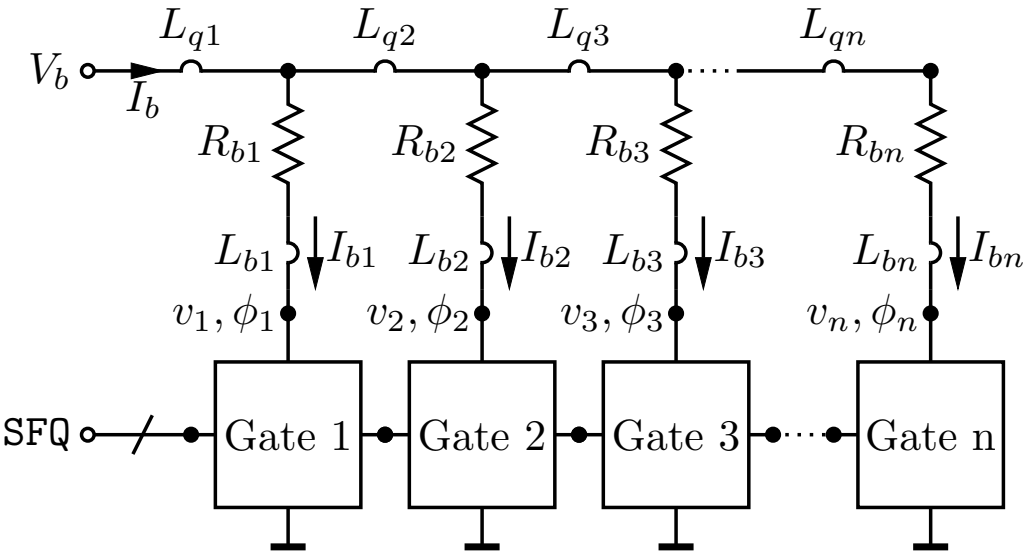

Figure 6.3: Conventional, resistive RSFQ biasing.

One of the design parameters of RSFQ logic is the bias voltage V_b. It is generally common across the entire integrated circuit. As the bias source is off-chip, the path from the source to the chip is resistive, incurring a voltage drop from the source to the chip. As typical bias voltages are of the order of several mV, such a voltage drop is non-negligible. Hence, a voltage-source cannot be used to reliably bias the chip from off-chip. Instead, the total bias current required for the biased circuitry is injected at terminal V_b from off-chip. The resistors divide the current in the desired manner and V_b takes on the desired bias voltage.

This is straightforward to design for when the logic circuitry is quiescent, as then $v_i = 0\ \forall i$ (superconductors incur no voltage drop). When the logic circuitry is active, however, $v_i > 0$. If all v_i were equal, then the voltage drop

across the resistors would still be V_b (as the chip is biased by a current source), but this is not true in the general case. Voltages v_i occur because of flux quanta propagating in the gates. The associated voltage pulses, $\int v(t)dt = \Phi_0$, occur at nodes v_i. If one flux quantum traverses node v_i during every clock period $\frac{1}{f_{\text{clk}}}$, then the average voltage at the node is $v_i = \Phi_0 f_{\text{clk}}$. As the clock precipitates all decisions in RSFQ logic, no flux quanta cannot propagate at a rate higher than one per clock period, ensuring that

$$v_i \leq (v_i)_{\max} = \Phi_0 f_{\text{clk}} \ \forall i. \tag{6.1}$$

Considering typical clock frequencies (tens of GHz), we conclude that $v_i < 100\,\mu\text{V}$. This means that $V_b \gg v_i$, ensuring that the dc bias current distribution is not skewed significantly when the fed circuitry is active. In fact, this is one of the constraints that imposes a lower bound on V_b [128, 129].

Static dissipation of conventional RSFQ logic occurs in the resistive bias network. Total static dissipation P_s is given by

$$P_s = \sum_i V_b I_{bi} = V_b I_b. \tag{6.2}$$

Dynamic dissipation of conventional RSFQ logic is given by

$$P_d = \sum_i v_i I_{bi} \leq I_b (v_i)_{\max} = I_b \Phi_0 f_{\text{clk}}. \tag{6.3}$$

Thus, the ratio of static to dynamic power dissipation is

$$\frac{P_s}{P_r} = \frac{V_b}{\Phi_0 f_{\text{clk}}}. \tag{6.4}$$

For typical bias voltages (several mV) and clock frequencies (tens of GHz), this ratio ranges from one to two orders of mangitude. Clearly, addressing static dissipation is paramount in achieving ultra-low-power dissipation.

6.3.3 Interfaces

SFQ pulses have a low energy content. A 20 GHz pulse train, generated by a junction biased with 200 µA, has power $P \sim 8\,\text{nW}$, or ~ -50 dBm. This is a low figure, and the low amplitudes and high bandwidths of the pulses make it difficult to amplify them. Furthermore, the generating junction is shunted with resistors of the order $R \sim 1\,\Omega$, which makes impedance-matching with subsequent non-Josephson devices difficult.

Likharev and Semenov proposed a limited solution, the SFQ-DC converter, which is schematically depicted in Figure 6.4. An SFQ pulse arriving at `Set` "triggers" the device, depositing a flux quantum into loop J_1, L, J_3. A current is tapped off into J_5, adding to current I_{b2}. This places J_5 into continuous resistive mode, corresponding to a dc voltage appearing at `Out`. Rapid decay

of the dc voltage occurs when an SFQ pulse arrives at Reset, placing the circuit into its original state. Junction J_6 is required to prevent pulses from J_5 back-propagating through the network.

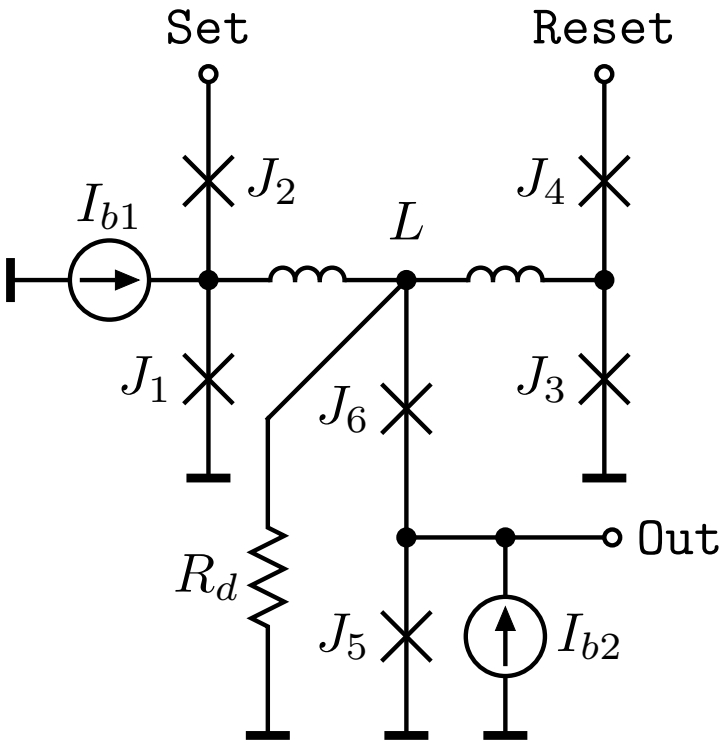

Figure 6.4: Schematic of SFQ-DC converter proposed by Likharev and Semenov [9].

Whereas this SFQ/dc converter outputs a dc signal that can be measured across the thermal barrier, the amplitudes achievable are still of the order sub-mV. The low output impedance makes power delivery into high-speed amplifiers difficult. Instrumentation amplifiers (like the Texas Instruments INA129 [130]), can be employed to amplify the output at low speeds, at which point it can be displayed on an oscilloscope or acquired by an FPGA. This is a staple experimental observation method in the industry.

Note that often in practice a variant of this SFQ-DC converter is employed. This variant only has one input, Toggle, which is followed by a toggle-flip-flop that alternately directs pulses to Set and Reset. In that case, each arriving pulse either switches the output on or off. The amplified output signal must then be interpreted accordingly.

For high-speed conversion from SFQ to conventional logic, however, special SFQ circuits must be employed. There are two solutions that we are aware of, both relying on stacks of Josephson junctions which are all toggled into and out of the voltage state [131, 132, 133]. These circuits can deliver some

power into a high-bandwidth room-temperature amplifier, enabling readout of signals until frequencies of the order ~ 5 Gbs^{-1}. Note that these circuits are not trivially implemented and take up considerable chip area. In certain cases, an ac bias must be applied.

6.4 Lowering RSFQ power consumption

Several methods of reducing the static power dissipation of RSFQ circuits have been extensively investigated. Other methods that depart further from the concept of RSFQ logic, such as RQL, have also been demonstrated [127].

6.4.1 LR-biasing

A straightforward approach to reduce the dominant static power dissipation would be to reduce V_b as much as possible. However, bringing V_b closer to $(v_i)_{\max}$ makes the bias current distribution less stable, requiring a corresponding increase in L_{bi}. This approach, known as "LR-biasing" was first introduced in 1997 and recently was further explored [134, 135, 129]. While at first glance this method may seem to support an arbitrarily low static power reduction, the L/R time constant of the branch imposes an upper bound on the frequency of operation, decreasing the effectiveness and applicability of LR-biasing.

6.4.2 ERSFQ

An examination of the IV-curve of the Josephson junction (see Figure 6.5a) leads to the observation that, at the transition point between the two regions of operation, the junction transmits a current of $i = I_c$ at zero voltage. This suggests that the junction may be employed as a limiting device to set up the required bias current level for a gate. This is done in Figure 6.5b, a concept known as ERSFQ biasing.

Injecting a total bias current $I_b = \sum_i I_{ci}$ forces each junction to transmit its critical current, which is designed to the desired I_{bi} of each branch:

$$I_{ci} = I_{bi} \ \forall i.$$

This enables current-limiting with zero voltage-drop across the junction, essentially achieving $P_s = 0$. However, when the biased gates are active (switching), the voltage at the bias injection points is non-zero. To maintain the correct region of operation of the current-limiting junctions, the voltage on the bias-line must at least equal the voltage at the bias injection points. In ERSFQ circuits, this can be achieved by connecting a so-called *feeding Josephson transmission line (JTL)* to the bias line as depicted in Figure 6.5. If the feeding JTL is connected to the same clock as the fed gates, the bias injection

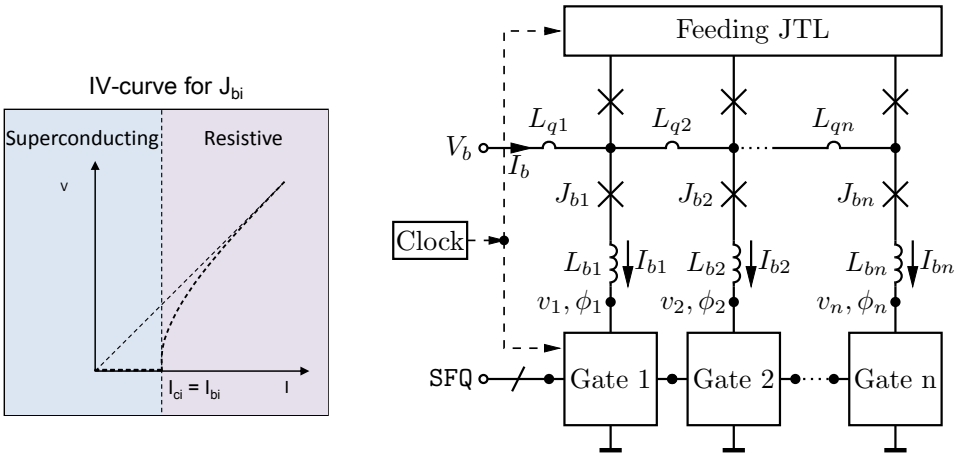

(a) (b)

Figure 6.5: The current-limiting properties of the Josephson junction (a) can be exploited to achieve the desired bias current distribution (b).

points on the JTL experience a 2π phase slip at every clock period. This ensures that the voltage on the bias line is kept at $v \approx \Phi_0 f_{\text{clk}}$, which we know from (6.1) to match or exceed the voltage at all other bias injection points. Hence, the voltage across the biasing branches is kept at or just above zero, maintaining the correct operational region of the limiting junctions.

Although the average current transmitted by the limiting junction is equal to its critical current, the instantaneous current may deviate, especially when the gates are active. To reduce this effect, an averaging inductor L_b must be employed. For typical required bias currents (several 100 µA), guaranteeing a deviation of $\Delta I_{bi} < 5\%$ requires a (comparatively large) inductance of $L_b \approx 400\,\text{pH}$ per branch. This translates to a larger area required for ERSFQ gates as compared to conventional RSFQ gates. Nonetheless, complex ERSFQ circuits have been successfully demonstrated [136].

6.4.3 The answer: eSFQ

While ERSFQ successfully minimizes static power dissipation with minimal modification to existing RSFQ gates, the addition of large averaging inductors and feeding JTLs bring about certain design challenges. These inconvenient requirements can be removed in the recently proposed eSFQ logic [25] (admittedly at some cost, mostly measured in design time). The eSFQ bias network (Figure 6.6) is topologically similar to the ERSFQ bias network. The chief difference lies in the size of the limiting inductor L_{bi} and the absence of a feeding JTL. Instead, the function of the feeding JTL is accomplished by the SFQ clock distribution network, which is largely a JTL network.

To enable biasing in this manner, care must be taken to ensure that the

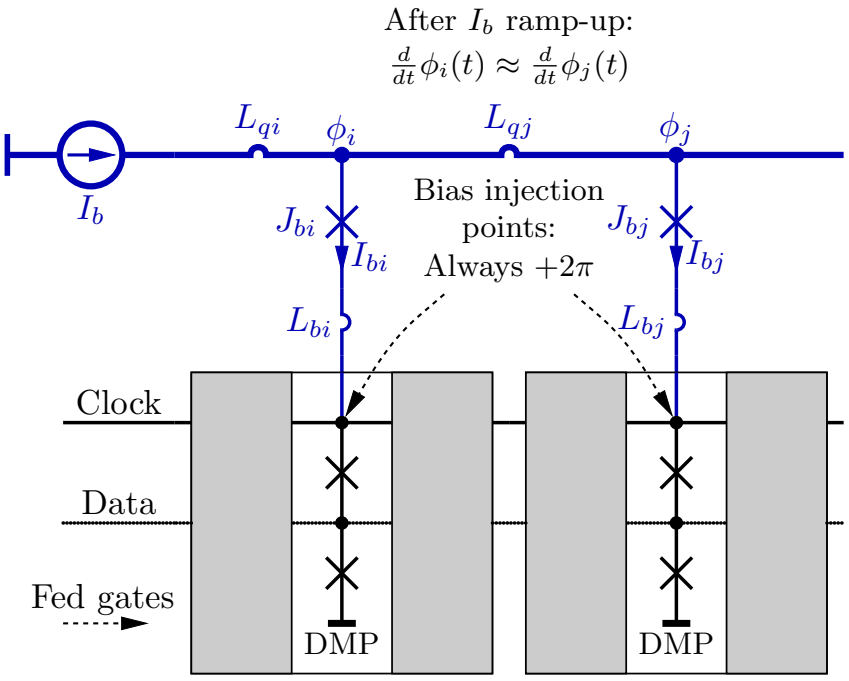

Figure 6.6: eSFQ biasing principle: Exploiting the current-limiting properties of the Josephson junction without the need for a large L_b is possible when the phase gradients between different bias terminals remain constant ($\dot{\phi}_i(t) \approx \dot{\phi}_j(t)\ \forall i, j$). This can be achieved by ensuring that all bias terminals are connected to the clock net, as each point on the clock net experiences the same phase shift ($+2\pi$) during each clock period. Most RSFQ gates contain a decision-making pair (DMP) which is interrogated by the clock signal, making for a suitable bias terminal.

limiting junctions do not switch during circuit operation. This obviates the need for a large averaging inductor L_b. To ensure this, the junction phases must be kept at sub-critical level during circuit operation. This requirement is met trivially while the fed circuitry is quiescent, as the phase at each bias injection point remains constant. Therefore, the phase across the limiting junctions, once established during ramp-up of the bias current, remains undisturbed (at sub-critical level).

When the fed circuitry is active, however, the phase at the bias injection points does not remain constant. Switching events in the fed circuitry bring with it phase increments of 2π. The key observation here is that for a general RSFQ-type circuit, these phase increments may or may not occur at the bias terminals, depending on data SFQ propagation. So, usually, in any given clock period, the phase at a bias injection point either jumps by 2π, or remains constant. Since data are generally random, different bias injection terminals accumulate different total phase, creating a changing phase gradient between terminals over time. This phase imbalance is acceptable in conven-

tional RSFQ, in which the normal metal bias resistors leak out the accumulated magnetic flux and keep phases equalised. Were these to remain unequalised, a parasitic supercurrent would flow across the phase gradient, skewing the bias current distribution. In *LR*-biased RSFQ circuits, which have significantly larger bias inductors, the LR time constant limits the speed of this phase balancing, thereby effectively limiting the clock speed. In ERSFQ circuits, phase equalization is achieved by compensating 2π phase increments occurring once the parasitic current increases the bias of a gate beyond the critical current of the gate limiting bias junction. In all these different RSFQ circuit approaches, the phase imbalances equalise in an asynchronous manner, which adds some degree of uncontrollable variations to dc bias currents. These contribute to the time jitter and limit ultimate circuit performance.

In contrast, the eSFQ approach removes this limitation by eliminating the very source of phase imbalance — the data dependent 2π phase increments at the bias injection terminals. In eSFQ, the biasing network is designed so that in each period, the phase at *all* bias injection terminals goes through the same change (Figure 6.6). In this way, the phase across each limiting junction does not change after the initial ramp-up of bias currents, whether the active circuitry is switching or not. Although conventional RSFQ gates are not designed with this requirement in mind, many RSFQ cells lend themselves well to conversion to eSFQ, as most RSFQ cells are clocked, requiring that a phase increment occur somewhere in the cell during each period. We also submit a preliminary solution to unclocked cells in Section 6.7.

6.5 Designing eSFQ data transport cells

The introduction of this chapter made the requirement of deserialisation clear. We now know of a separate constraint: high-speed interfaces between SFQ and conventional logic families are limited to approximately 2 Gbps, at least for current technologies. Although technologies are bound to improve over the SKA timeline, deserialisation still seems a likely requirement.

6.5.1 The objective of deserialisation

Deserialisation, also known as *serial-to-parallel conversion* or *demultiplexion*, is the process of directing a bit from an input to one of several outputs, where the selected outputs rotate in a predictable manner. A 1-to-n deserialiser takes an n-bit chunk of a serially applied input and provides it on n outputs, one bit per line. The main objective of deserialisation, at least in our case, is the reduction of data rate it achieves. A serial stream of 16 Gbps fed into a 1-to-16 deserialiser results in 16 streams of 1 Gbps. A 1-to-8 deserialiser is conceptually depicted in Figure 6.7.

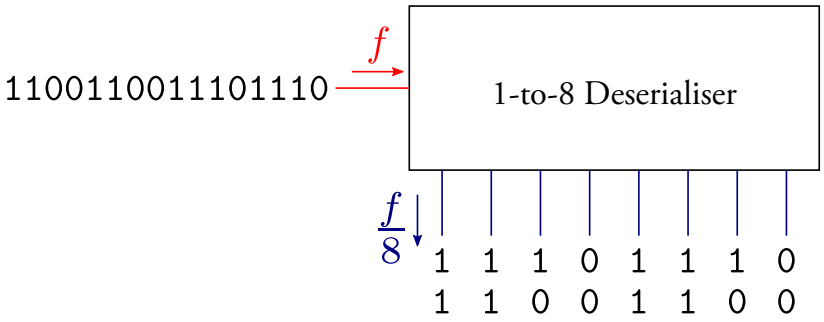

Figure 6.7: A 1-to-8 deserialiser reduces the data rate by factor 8.

Deserialisation makes no assumption about the nature of the serial input, it simply converts it to several parallel streams. If the serial nature of the stream is important, this must, of course, be respected in further processing.

RSFQ deserialisers, as shown conceptually in Figure 6.8 are practical, fairly well-known devices. Pulses arriving serially are directed to consecutive outputs. Since the absence of a pulse carries much significance (coding the digital 0), an absence must also be captured and output. For this reason, a clock must be employed. In addition, synchronisation is employed to ensure that all output bits appear at the same time (while this is not strictly necessary, it makes post-serialisation acquisition and processing of data easier).

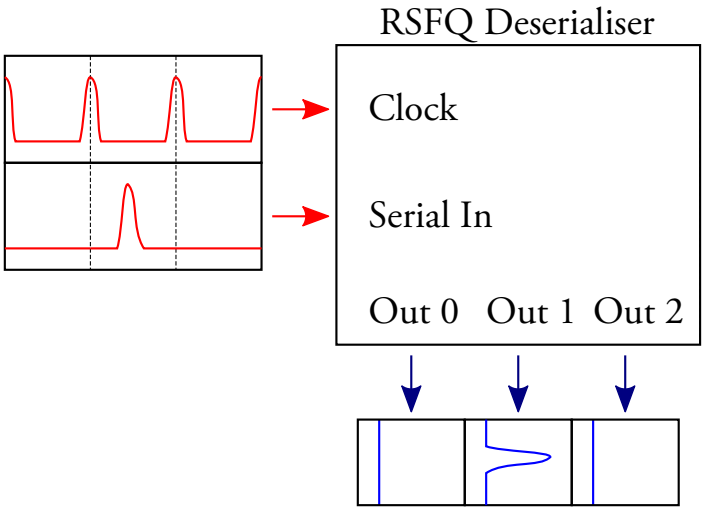

Figure 6.8: Function of an RSFQ deserialiser. The serial input is distributed across several outputs.

For our targeted 4-bit flash ADC, clocked at 20 GHz, the serial output rate per bit is 20 Gbps. Considering the constraints illuminated so far, we believe

that employing a 1-to-16 deserialiser for each bit will reduce the data-rate sufficiently to be reliably detectable by conventional technology 6.9. Similar approaches have been employed in the past [131].

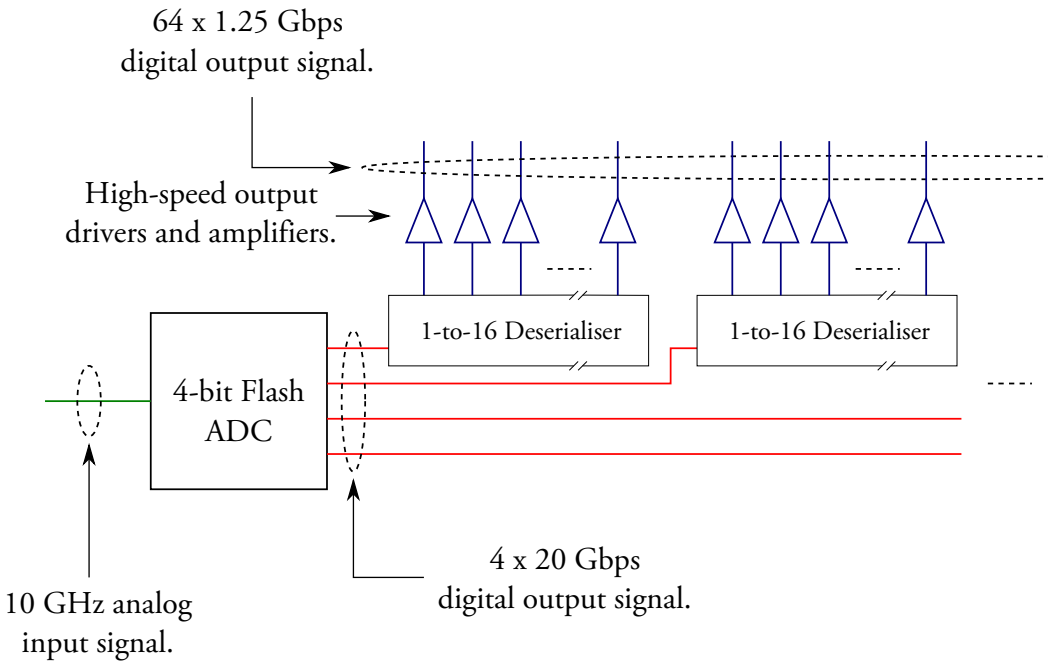

Figure 6.9: Concept of making ADC output available to conventional technologies.

6.5.2 Converting RSFQ to eSFQ

Before synthesising an eSFQ deserialiser we consider the extent of the modifications required to convert existing RSFQ gates to eSFQ biasing. Most RSFQ circuits are generally well-suited to conversion to eSFQ. Consider, for example, a JTL transmitting the clock signal. As each junction in the JTL switches during each clock period (exhibiting a 2π phase slip), bias current injection can occur at arbitrary locations. RSFQ logic gates rely on serially connected pairs of Josephson junctions — Decision-Making Pairs (DMPs) — to perform a function and control output. In a DMP, exactly one junction switches whenever an interrogating SFQ pulse is applied. Whichever junction switches, the sum of the phases across the DMP increases by 2π. As, generally, the interrogating pulse is a clock signal, the phase atop the DMP meets the requirements of an eSFQ bias injection terminal.

The conversion from RSFQ to eSFQ of a ubiquitous RSFQ cell, the D-cell (D flip-flop), is depicted in Figure 6.10. Conventionally, the D-cell is biased so that it initially stores a 0. Such biasing occurs on junction J_1, which

only experiences a 2π phase-increment when a pulse arrives at `In`. However, exactly one pulse arrives at `Clock` during every clock period, ensuring a 2π phase increment across the DMP. Hence, for eSFQ, the bias injection point is moved to the DMP. As a side-effect, the converted D-cell stores a logic `1` after initial bias ramp-up.

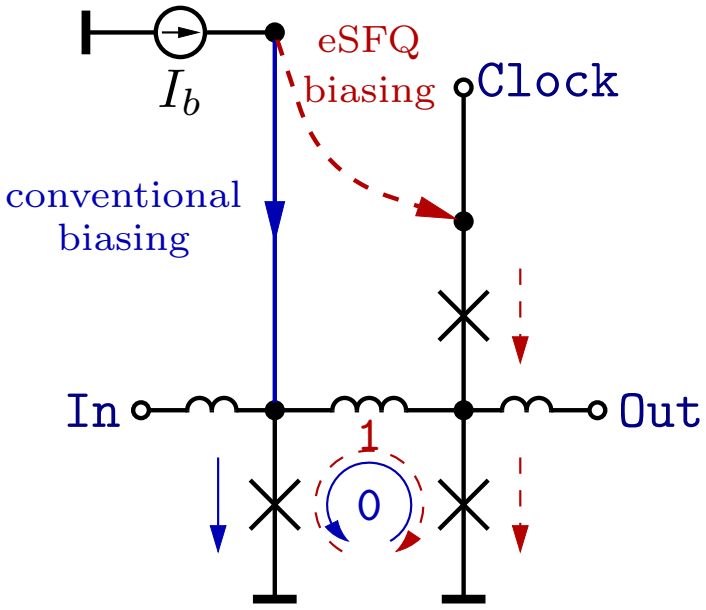

Figure 6.10: Conversion of D-cell from RSFQ to eSFQ.

In an analogous manner, most RSFQ cells can be converted to eSFQ. Note that certain RSFQ cells do not lend themselves as intuitively to conversion to eSFQ, such as the T flip-flop (TFF), which is not clocked and thus, in conventional form, does not have a bias injection point suitable to eSFQ. In this case, the design could be modified to the unbiased design, known as the "supply-free" design. Pulses would propagate through the circuit in ballistic mode [137]. There are limitations to this, but they could be overcome by, say, "sandwiching" supply-free cells between biased cells (such as JTLs) that regenerate pulses. We also propose another method to deal with this problem in Section 6.7.

6.5.3 Shift registers

Shift-registers are very natural RSFQ logic cells, as RSFQ is a *sequential logic family* rather than a *combinational logic family*, such as CMOS. A number of RSFQ shift registers designs were developed in the past, but the most robust and widespread design comprises DFFs with an integrated SFQ clock network [48]. We chose this design for implementation in eSFQ logic. Two different eSFQ versions of these shift registers were implemented: a straight-forward

conversion from RSFQ termed "eSR" and a version with an additional magnetic flux bias, "MeSR".

eSFQ shift register: eSR

The schematic of our first shift register cell, eSR, is depicted in Figure 6.11a. It can be topologically partitioned into two sections: *clock* and *data*. Junctions J_{c1}, J_{c2} and the DMP J_1, J_2 make up the clock section. This part of the cell transmits the clock and applies it to the DMP as an interrogator. Note that we employed the counterflow clocking scheme (see Figure 6.11c), as this has generally yielded higher margins in RSFQ shift registers [138]. If J_1 switches, the clock pulse is simply transmitted from CIn to COut. If J_2 switches instead, an output pulse exits at DOut in addition to the transmitted clock pulse. In either case, the phase at node x, ϕ_x, increases by 2π.

Node x is the only bias current injection point for eSR. The distribution within the cell depends mainly on the inductors L_2, L_3 and L_5 (with Josephson inductances and parasitics playing a minor role).

Junctions J_d and J_2 make up the data section. After bias current ramp-up, J_2 is biased, whereas J_d is not, which corresponds to the cell storing a 1. After the first clock signal, the bias current redistributes to J_d, which corresponds to the cell storing a 0. A pulse appears at DOut, representing the initially stored 1 as depicted in Figure 6.11b.

Margins of operation of this circuit were determined for a 4-bit shift register configuration. While a 4-bit configuration is not very useful in practice (although, see Section 5.5), we feel that it is an acceptable compromise between comprehensiveness and computational tractability. The critical parameter was identified as the critical current of junction J_1, the upper (escape) junction of the DMP. One of the reasons for this is the injection of bias current through the DMP as required in accordance with the eSFQ biasing scheme. The difference between the biased and unbiased DMP is evident from the phases of the DMP junctions shown in Figure 6.12.

The grounded junction in a DMP (J_2) switches when it is biased, whereas the escape junction (J_1) switches when the grounded junction is not biased. The escape junction J_1 is, conventionally, not biased. Hence, when J_2 is unbiased as it is in case of RSFQ (ERSFQ) circuits, both junctions have a phase near zero, and as J_1 has a lower critical current, and is closer to the source of the interrogating pulse, it switches in the presence of an interrogating pulse. When J_2 is biased, its phase is nearly critical, with the phase of J_1 remaining near zero, making J_2 the switching junction when the DMP is interrogated. When biasing through the DMP as in eSFQ, the escape junction J_1 is permanently biased. This fixes its operating point phase at greater than zero, increasing its switching affinity, particularly as it is closer to the source of the interrogating pulse (the Clock node). This reduces the difference between the steady-state phases of J_1 and J_2 when J_2 is biased. In this case, the increased

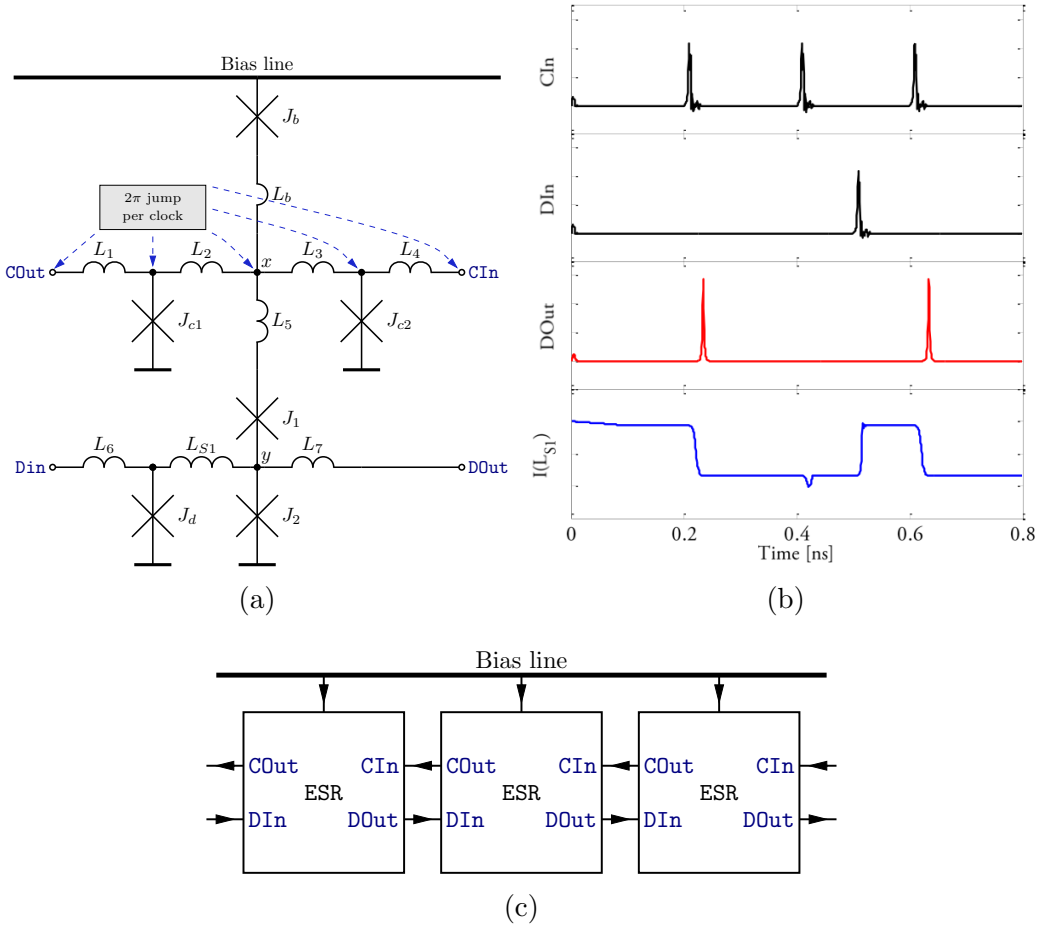

Figure 6.11: eSR - eSFQ shift register cell: schematic (parasitics omitted) (a), simulation (b) and in a typical configuration (c).

switching affinity of J_1 is undesirable (as J_2 should then switch) and results in lower parameter margins for J_1.

In order to improve the margins of J_1 in eSR, it was designed in an underdamped configuration with $\beta_c \sim 0.25$. An underdamped junction exhibits lower switching speeds and J_1 is thus less likely to switch before J_2 when the DMP receives the interrogating pulse. With the underdamped J_1, a 4-bit configuration of eSR achieved critical margins of $\pm 24\%$ and bias margins of $\pm 34\%$, with bias margin relating to the bias of the entire 4-bit test structure and critical parameter the area of J_1 for all four eSR cells simultaneously.

Corrective bias: MeSR

Although underdamping J_1 achieves the goal of increased parameter margins, it has the undesired side-effect of increasing data-dependent clock skew. It takes J_1 longer to switch than J_2, which means that when the shift register

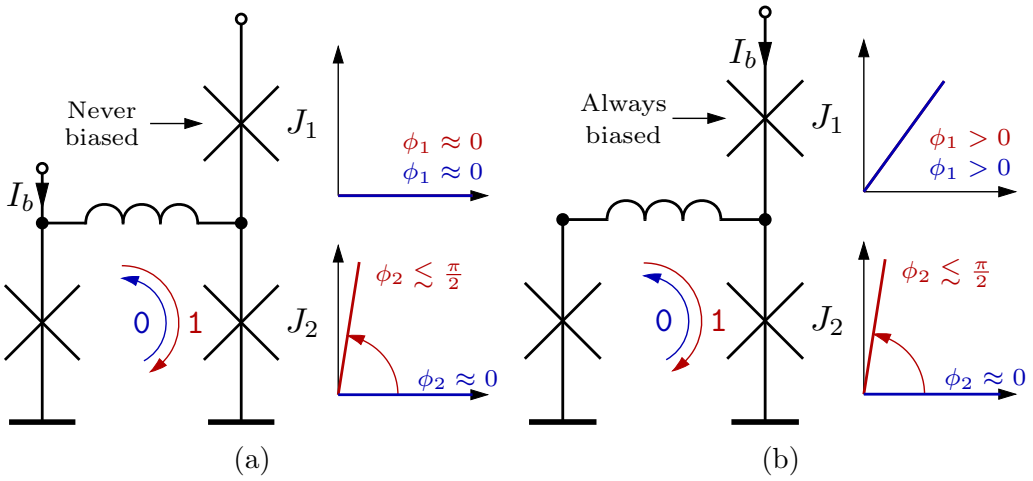

Figure 6.12: Comparison of DMPs in RSFQ (ERSFQ) and eSFQ. Injecting the bias current in the conventional DFF leaves J_1 unbiased (a). Moving the injection point to the DMP forces the phase in both DMP junctions in the same direction during ramp-up (b).

stores primarily `1`s, the clock propagates through faster than when the stored bits are primarily `0`s. Although this effect is also seen in conventionally biased shift registers, due to an underbiased escape junction J_1, it is somewhat more pronounced in `eSR`.

From (2.8) we know that, compared to critical damping, $\beta_c \sim 0.25$, as used above, doubles the characteristic time of J_1. This potentially halves the maximum clock frequency achievable.

Hence it is desirable for all junctions to be equally damped at $\beta_c \sim 1$ and therefore have the same switching speed, achieving the maximum frequency for a typical 4.5 kA/cm^2 critical current density. To accomplish this without punishingly narrow parameter margins, we investigate several options.

For correct operation, an interrogating pulse must not cause J_1 to switch when J_2 is biased. At first glance, keeping the critical current of J_2 low enough should achieve this. But consider Figure 6.11a: when the bias current enters at Node x, some travels down the DMP, biasing junction J_1. After crossing J_1, the bias current divides again at node y. The alternative path to ground through J_d means that some of the bias current leaks away from J_2, ensuring that J_2 always receives less bias current than J_1. Lowering the critical current of J_2 increases its Josephson inductance, which exacerbates the leakage effect.

A magnetically introduced corrective flux bias was employed to solve the leakage problem, resulting in cell `MeSR`, depicted in Figure 6.13. The dc flux bias, introduced through L_f, forces the current in the storage loop to redistribute as intended, opposing the leakage effect. In this way, the phase offset of J_2 (as a result of the eSFQ bias current) can be modified. During circuit optimization, it became apparent that using the flux bias to redirect initial bias

current from J_2 to J_d was most effective at maximizing parameter margins. A potential advantage of this is that shift registers based on MeSR initially store a 0 which aligns well with conventional RSFQ shift registers.

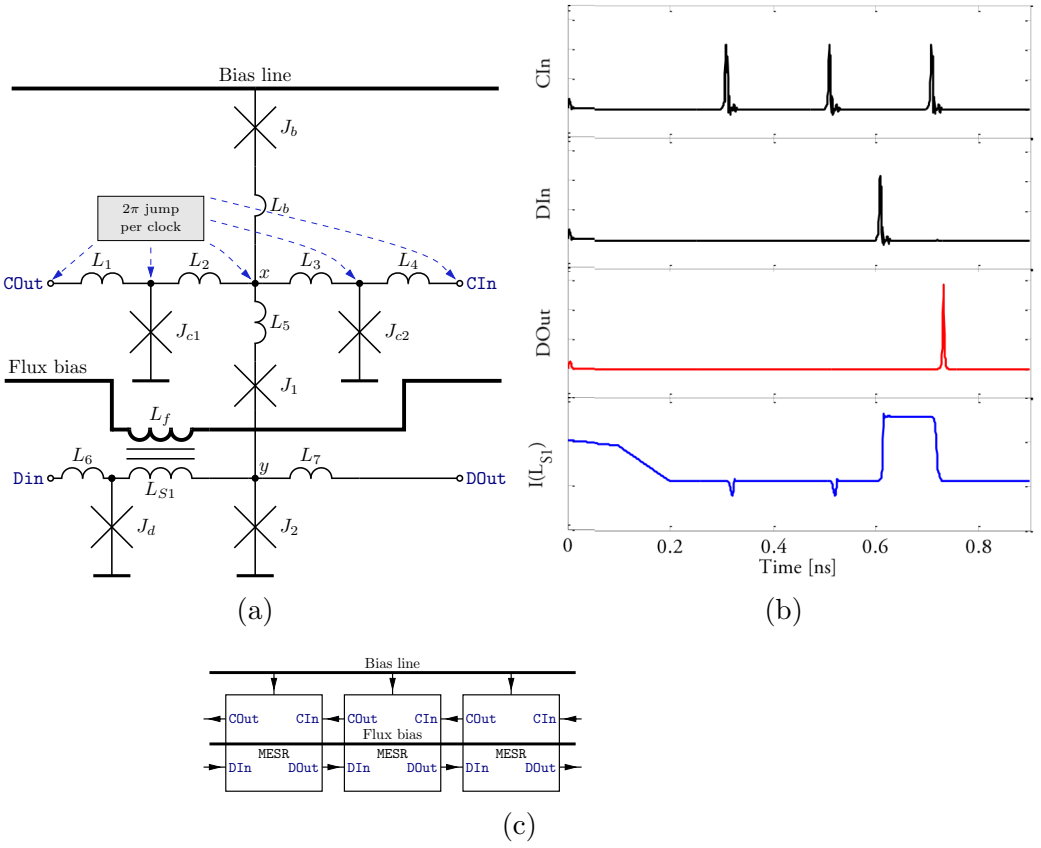

Figure 6.13: MeSR - eSFQ shift register cell with corrective flux bias: schematic (parasitics omitted) (a), simulation (b) and in a typical configuration (c).

A further advantage of the flux bias manifests itself in reduced bias current requirements in terms of injected bias current (and corresponding decrease in dynamic power dissipation). As undesired leakage can be avoided and the desired balance in the storage loop be established using the flux bias, less bias current needs to be injected at the injection point x. Initial 0 storage and reduced bias current requirements are shown in Figure 6.18. Note that one flux bias line is required to bias the entire shift register, irrespective of its length. For a 4-bit MeSR-based shift register, a critical margin of ±27%, and bias margin of ±36% were achieved.

The presence of the additional flux bias line might appear as a complication. In reality, the bias does not have to be introduced this way. A constant flux bias can be implemented in a variety of ways ranging from a small superconducting

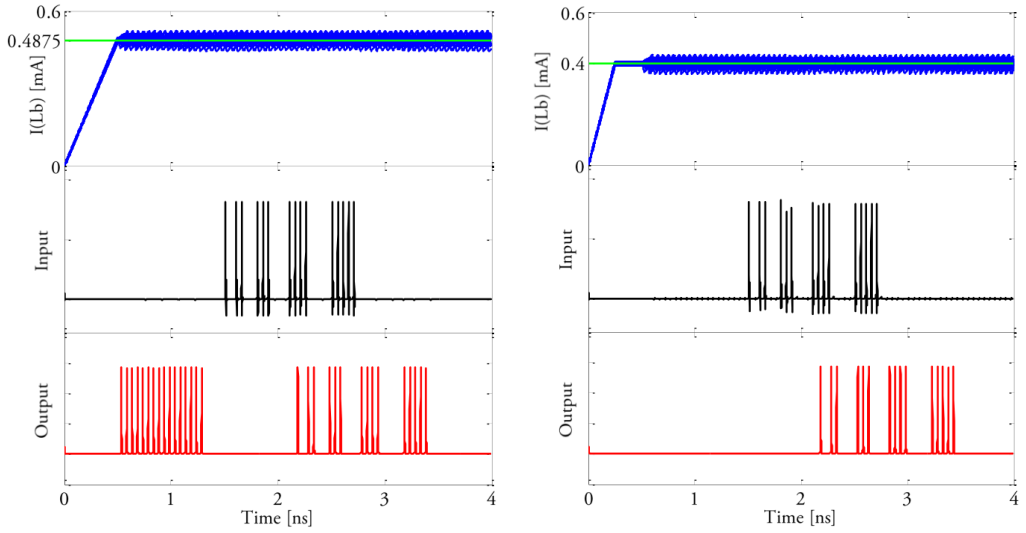

Figure 6.14: Results of simulation with a 20 GHz clock for 1.0 aJ/bit `eSR`- (a) and 0.8 aJ/bit `MeSR`-based (b) 16-bit eSFQ shift registers at bias current corresponding to the centre of the operational region. Bias currents distribute correctly, with accetable distortion through switching events. For the eSR-based shift register, 16 output bits are immediately observed after starting the clock, before any input is applied. In both cases, the input pattern is reproduced at the output. Lower bias current requirements are evident for the magnetic flux-biased shift register (at higher margins).

loop with frozen-in SFQ to a π-junction implemented using superconducting-ferromagnetic-superconducting (SFS) Josephson junctions [139, 140, 141, 142]. Even for conventional RSFQ circuits, improved operational margins, bit-error rates, and gate memory non-volatility were reported [143, 144, 145].

As (6.3) dictates, the reduction of bias current directly translates into the reduction of dynamic power dissipation. This makes the magnetic bias approach especially valuable. As the magnetic flux bias is a passive, non-switching, superconducting element, it does not contribute to power dissipation. As is evident from the results of simulations for a 20 GHz clock (Figure 6.14), the `eSR` shift register consumes $\sim$ 1.0 aJ/bit, while the `MeSR` shift register consumes $\sim$ 0.8 aJ/bit. These energies correspond to the centre of the bias current operational region. At the lower limit, the energy per bit operation reaches $\sim$ 0.5 aJ/bit.

6.5.4 An eSFQ deserialiser: eDES

Two approaches have proven popular for RSFQ deserialisers: the binary tree architecture [146] and the shift-and-dump architecture [147]. We have chosen

the latter approach, as it has proven more practical due to its high modularity and simple timing requirements. Our eSFQ deserialiser cell is based on a dual-port D-flip-flop, or D^2-cell [148], which derives from the B flip-flop [149]. The D^2-cell is similar to the D-cell (see Section 2.3.3), but the storage loop has two decision-making pairs connected in parallel, providing two read-out ports. One port is intended for serial propagation of data along the deserialiser, the other for parallel readout. An n-bit deserialiser contains n D^2-cells, dividing a serial stream of bits into n parallel streams.

The designed eSFQ deserialiser cell, `eDES`, is depicted in Figure 6.15. The two readout ports are topologically symmetrical, both achieving destructive readout of stored flux. Note the additional escape junction in each readout arm (J_{der}, J_{des}). The deserialiser cell contains two DMPs, suggesting two bias injection points. As in `MeSR`, a flux bias is employed in the data section to achieve the desired bias current distribution between J_d, J_{2r}, J_{2s}. All junctions were designed to be critically shunted. When correctly biased, `eDES` stores a `0` after ramp-up.

Two clocks thread each deserialiser cell. The symmetry of the cell and size of the limiting junction means that the per-bit switching energy required by the shift operation as well as the read operation is comparable to that of the `MeSR`-based shift register. For normal operation the ratio of the clock frequencies depends on the length of the deserialiser. A per-bit switching energy is thus not meaningfully ascribed to the deserialiser cell, but for long deserialisers the per-bit switching energy of the deserialiser approaches that of the `MeSR`-based shift register.

6.6 Experimental verification

6.6.1 The chips

In order to investigated eSFQ logic experimentally, the `eSR`, `MeSR` and `eDES` cells were laid out and circuits re-optimised to account for the extracted layout parasitics. The full circuit diagrams, layouts and extracted parameter values are shown in Section A.5.

Next, three consecutive suites of tests were performed. Each suite consists of layout, manufacture and test of a batch of chips.

Suite I — functional tests

The objectives of this suite of tests were as follows:

1. Establish functional correctness of the designed eSFQ cells,
2. find bias margins and establish whether they are high enough, and
3. investigate the effect of different bias line fan-in configurations.

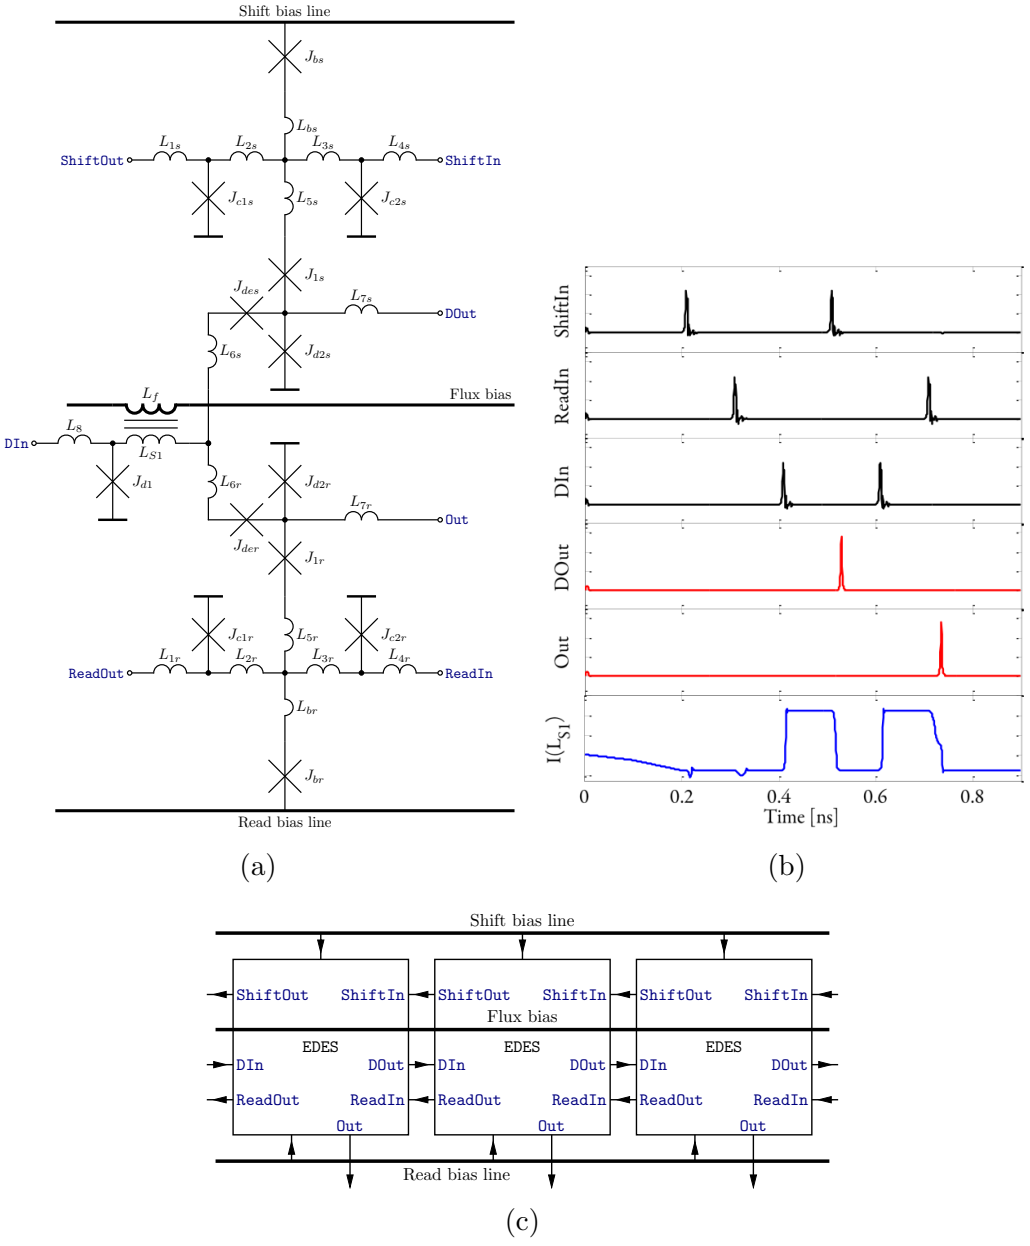

Figure 6.15: `eDES` - eSFQ deserialiser cell: schematic (parasitics omitted) (a), simulation (b), in a typical configuration (c). The flux bias is ramped up from 0.1 ns to 0.2 ns in the simulation from (b), as evident from the bottom trace. Initially, eDES is non-storing, but when a pulse arrives at `DIn`, it is stored, readable by both a `Shift` or a `Read` pulse.

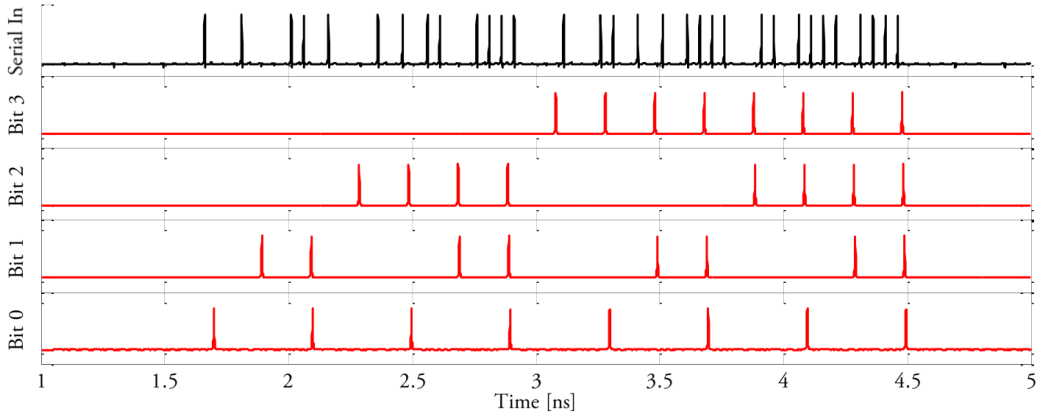

Figure 6.16: Simulation of a 4-bit eDES-based deserialiser at 20 GHz. Deserialisation is clearly evident, the serial input signal (a 4-bit count from 0 to 15) is parallelised, resulting in 4 output streams.

To these ends, several eSFQ shift registers of 16- and 32-bit length, as well as deserialisers of 4-, 8- and 16-bit lengths were assembled. Five 5×5 mm^2 chips were populated with 12 structures in total and manufactured in Hypres's Niobium 4.5 kA/cm^2 process [49]. Microphotographs of two fabricated chips are shown in Figure 6.17. The full layouts of all five chips are repeated in Section A.6.

Microphotographs of fabricated `eSR` and `MeSR` circuits are compared in Figure 6.18. As evident, they differ in the shunt resistor of escape junction J_1 in order to achieve overdamping in `eSR` and critical damping in `MeSR`. For `MeSR`, the magnetic flux bias is delivered by a superconducting line under the cell storage inductors to magnetically apply the required corrective bias. Microphotographs of the deserialiser test structures are depicted in Figure 6.19. In order to protect circuits from flux trapping, ground plane moats [150] were placed around cells and further ground plane holes on unused chip area.

To concentrate design effort on the eSFQ demonstrator cells, and to minimise the probability of failure in the peripheral circuits, existing conventional RSFQ cells from the Hypres cell library were employed as a testbed. These comprise standard interfaces to room-temperature circuitry, such as the DC-SFQ and SFQ-DC converters [9]. Figure 6.19b shows an RSFQ testbed made up of these standard library RSFQ cells. The test chips also contain standard diagnostic circuits for fabrication process control. These include Josephson junction arrays that can be measured to determine process parameters like critical current density. These diagnostic structures are visible in Figure 6.17.

Thus, in addition to laying out different combinations of the basic designed cells, these were placed in a variety of different bias configurations. The designed bias lines are characterised by the cell-to-cell inductance of the line, L_q, as well as the line-to-cell limiting inductance, L_b. Bias lines of two different

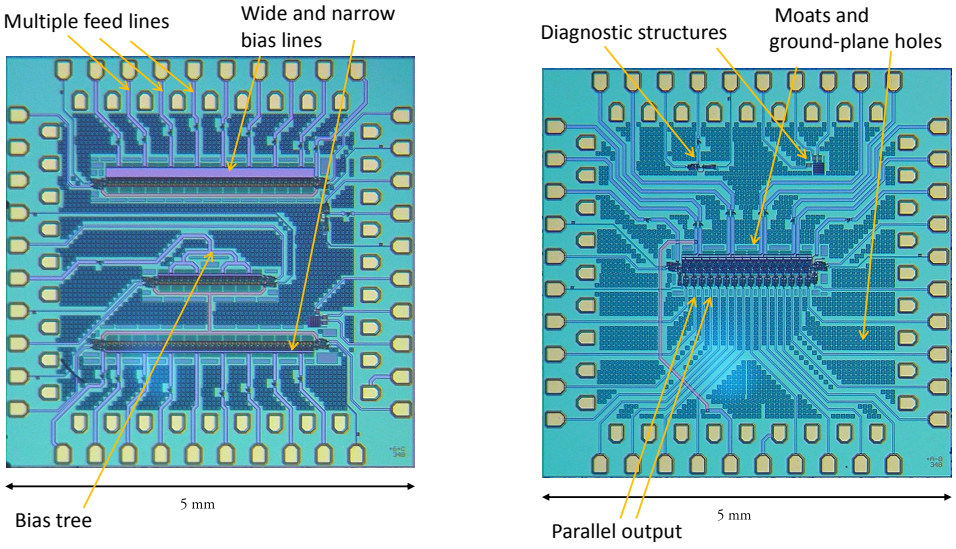

(a) (b)

Figure 6.17: Microphotographs of two of the five $5 \times 5\,\mathrm{mm}^2$ ICs fabricated with Hypres's 4.5 $\mathrm{kA/cm}^2$ process, one with three shift registers (Section A.6.2) (a), and one with a 16-bit deserialiser (Section A.6.5) (b). Each bias line is fed from multiple contact pads, which enables experimental investigation of different bias fan-in configurations.

widths (narrow: $L_q \approx 1.5$ pH, wide: $L_q \approx 0.5$ pH) were laid out and combined with three lengths of limiting inductor (short: $L_b \approx 10$ pH, medium: $L_b \approx 50$ pH, long: $L_b \approx 150$ pH).

For each laid-out structure, bias fan-in was conservatively high. Each structure has its own bias line, which is shared by all cells in the structure (deserialiser structures each have two bias lines, one for the read- and one for the shift operation). One bias pin was allocated to every four cells in a structure. This enables comprehensive investigation of different bias-current fan-in configurations (biasing from all available pins or biasing from one pin only, for example). Figure 6.17 depicts two examples of chip layout, illustrating the high number of bias-current pins. To further the breadth of this investigation, one structure was equipped with a bias-current divider that binds the four bias line entry points of the 16-bit shift register to a single pin.

Suite II — high-speed tests

After functional verification, the next logical steps are in high-speed tests. Suite II contains four chips, three of them 5×5 mm^2 and one of them 10×10 mm^2. The three smaller chips contain 16- and 32-bit shift registers, whereas the larger chip contains two 16-bit deserialisers. Figures 6.20a and 6.20b depict microphotographs of a shift register and deserialiser chip respectively.

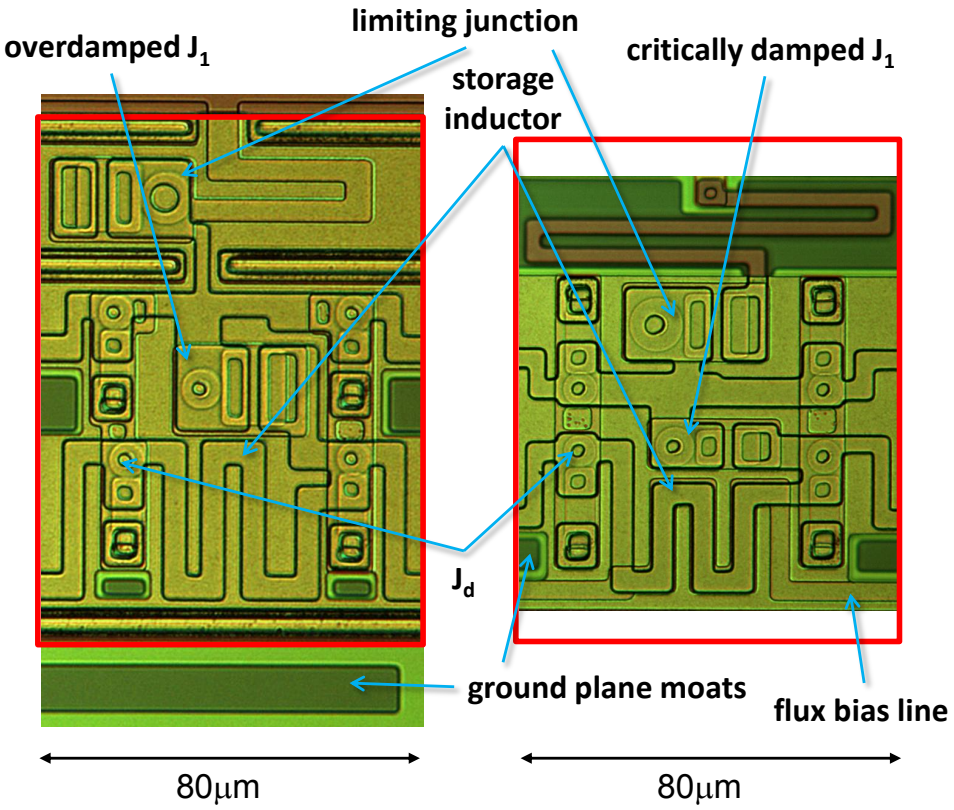

Figure 6.18: Photos of eSFQ shift register cells: eSR with overdamped J_1 (left) and MeSR with critically damped J_1 and a flux bias line inductively coupled to cell storage inductor (right). Cell sizes (indicated by red boundary): eSR: 80×110 µm^2, MeSR: 80×105 µm^2. These dimensions do not include the bias line.

Most design decisions from Suite I are repeated here, except that instead of a typical output SFQ-DC monitor (see Section 6.3.3), a high-speed output driver developed by Kirichenko [151] was employed for the shift-register data output. This output driver works at up to 2 Gbps and delivers enough power to bridge the thermal barrier. The signal, however, still only has ~ 5 mV of swing (which is about an order of magnitude greater than that achieved by a typical SFQ-DC monitor), and must thus be amplified further at room temperature.

More details of the structures designed can be found in Section A.7.

Suite III — further high-speed tests

While testing Suite II chips, several desirable improvements became evident:

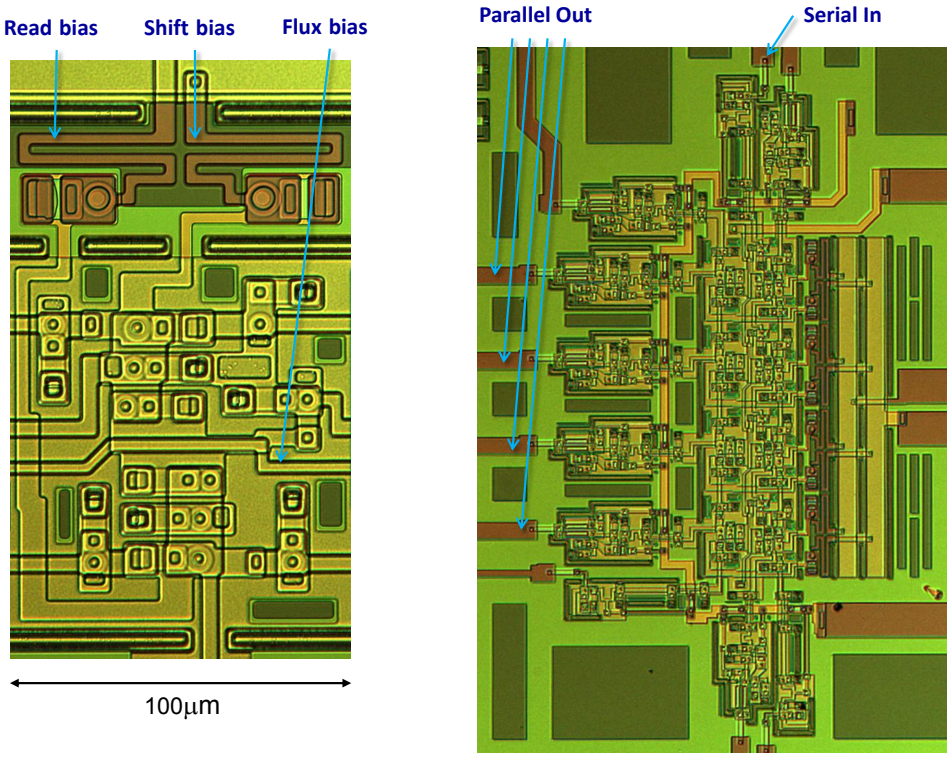

(a) (b)

Figure 6.19: Photographs of deserialiser test structures: the deserialiser base cell `eDES` (a), and a corresponding 4-bit eSFQ deserialiser (b). Bias lines and some key devices are indicated. A set of DC-SFQ and SFQ-DC converters enable the interface with room temperature electronics.

1. While shorter shift registers (up to 32-bit) are valuable demonstrators, larger shift registers are more practical and have the added benefit of adding a greater delay to the signal, which should make for easier experimental verification.

2. Since $\sim$ 2 Gbps output drivers are employed, all data signals are limited in bandwidth (achievable clock frequencies are an order of magnitude higher). While this does not necessarily limit clock speeds, a means of testing faster data signals is desirable.

3. It is not possible to apply broadband data signals (bandwidths greater than 2 Gbps), due to the non-flat transfer function of experimental signal transmission lines (mainly co-axial cables).

4. There were no single chips containing both `eSR`- and `MeSR`-based shift registers, making a direct comparison difficult.

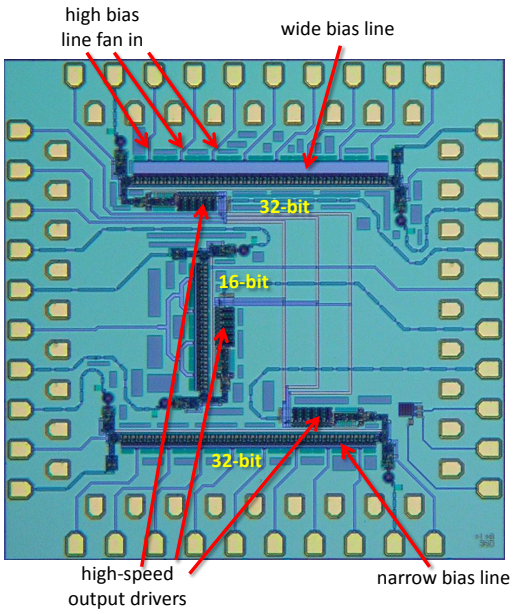

(a)

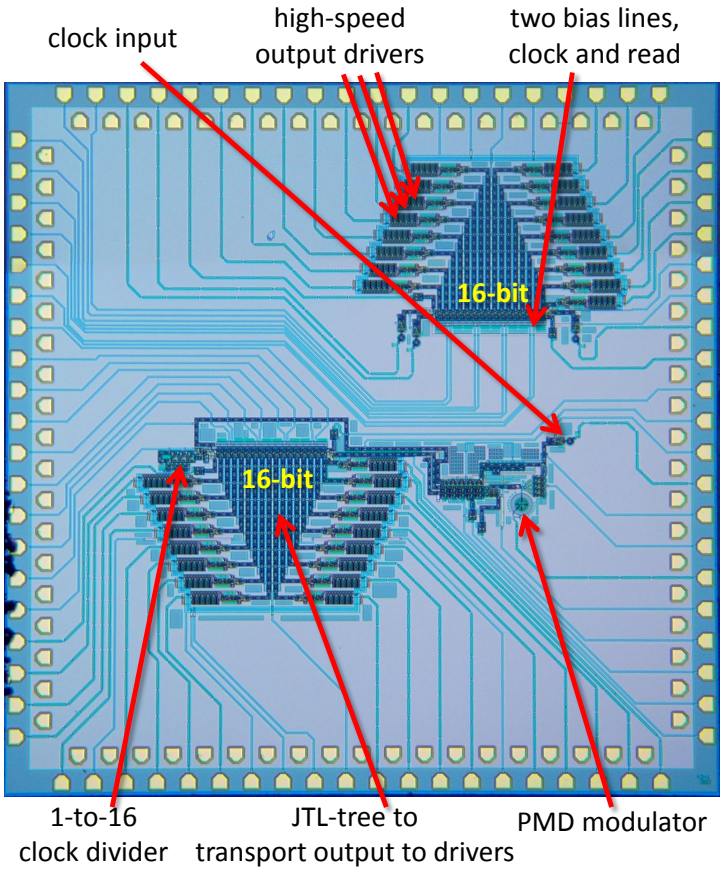

(b)

Figure 6.20: Photographs of $5 \times 5\,\mathrm{mm}^2$ shift register (a) and $10 \times 10\,\mathrm{mm}^2$ deserialiser (b) chips suitable for high-speed testing.

5. The high-speed derserialiser test setup, while innovative, is complicated and displayed signs of instability during preliminary testing. Furthermore, a design error was uncovered that limited one degree of freedom.

Consequently, Suite III was designed, laid out and fabricated. An example chip from Suite III is reproduced in Figure 6.21.

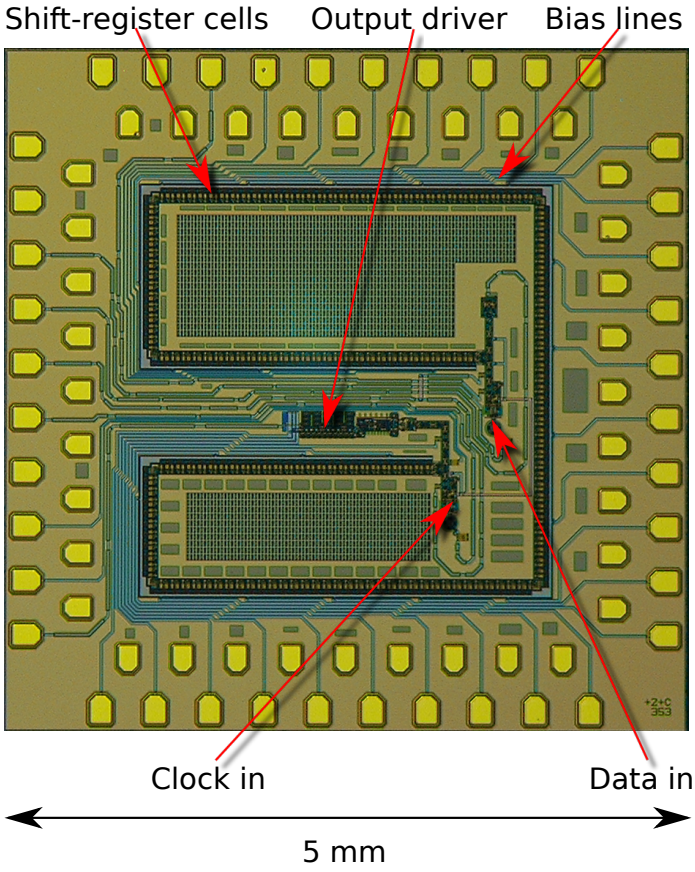

Figure 6.21: Microphotograph of 184-bit eSR-based shift register test chip.

Detailed designs and layouts are listed in Section A.7, here we only summarise how we addressed the identified insufficiencies of Suite II:

1. Large (184-bit) shift registers were laid-out.

2. All data lines in shift registers were given a means to measure their average-voltage, as illustrated in Figure 6.22. While this does not enable a high-fidelity measurement of the data, some predictions can be validated in this way.

3. A means to modulate (mix) the data signal was provided to avoid the limitations the data channel transfer function. This should enable application of high-speed data signals (greater than 2 Gbps), though the shape of the signal would be limited to the comb shown in Figure 6.23. The voltage-controlled switch required is schematically depicted in Figure 6.24.

4. One chip, [...], was designed to include one eSR- and one MeSR-based shift register, both 32 bits in length, to facilitate a direct comparison.

5. The deserialiser testbed was adjusted slightly, with emphasis placed on synchronisation of the high-speed signal. The uncovered design error was corrected.

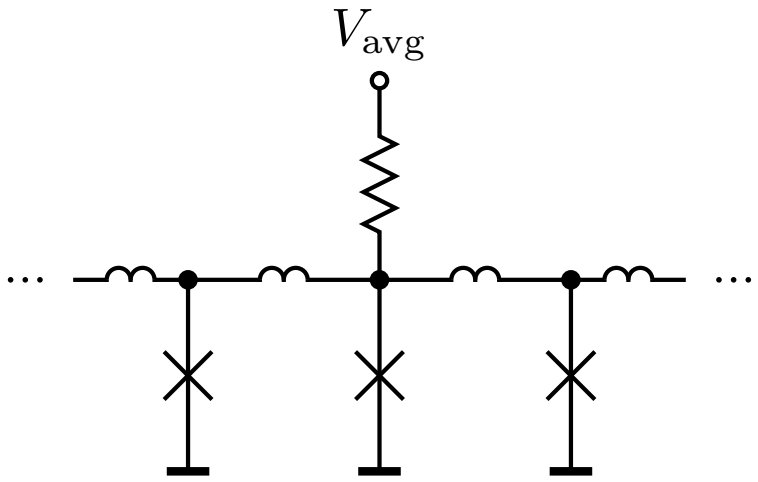

Figure 6.22: An average voltage pad is connected to SFQ lines, providing the means to measure the average (dc) voltage on the line, where $V_{\text{avg}} = \Phi_0 f$, with f the frequency of pulse incidence on the line.

6.6.2 The test setup

Functional tests

Functional tests were performed with an Octopux [152] test setup in Hypres' test lab. The tests were not performed by the present author, who designed them, as he was not present at the time when the chips were completed. Instead, they were performed by A. Sahu, a member of Hypres's staff. The present author is indebted to him.

The Octopux test setup is conceptually depicted in Figure 6.25. By means of LISP scripts (and a real-time interface based on LISP), patterns can be

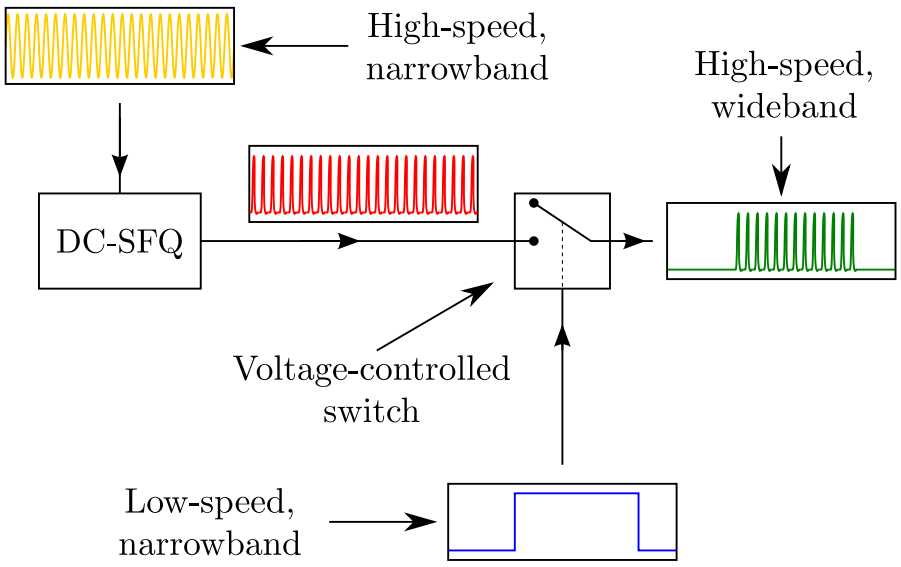

Figure 6.23: A crude mixer that enables synthesis of a wideband data signal (of experimental value) from two narrowband signals (easily transmitted) by means of an on-chip voltage-controlled switch.

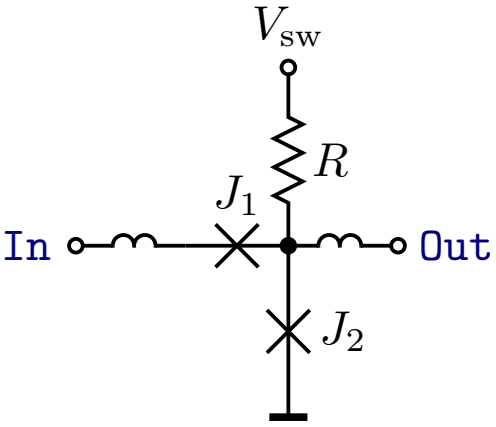

Figure 6.24: Schematic of a voltage-controlled switch as employed in the mixing concept introduced in Figure 6.23. $V_{\mathrm{sw}} = 0$ leaves J_2 unbiased, so any SFQ pulse arriving at `In` promptly escapes through J_1. When a sufficient voltage is applied at V_{sw}, J_2 is biased sufficiently and pulses arriving at `In` are propagated to `Out`.

programmed that are then applied to the chip by a purpose-built hardware platform. Responses are measured by the same platform and made available in the LISP environment. Octopux can also apply and vary bias currents and, by verifying responses to patterns, determine bias margins. An algorithm akin to that employed by our command-line tool `analyse` (see Section 3.4.1) was employed to determine the margins expediently. Naturally, parameter margins cannot be experimentally verified in Octopux in this way.

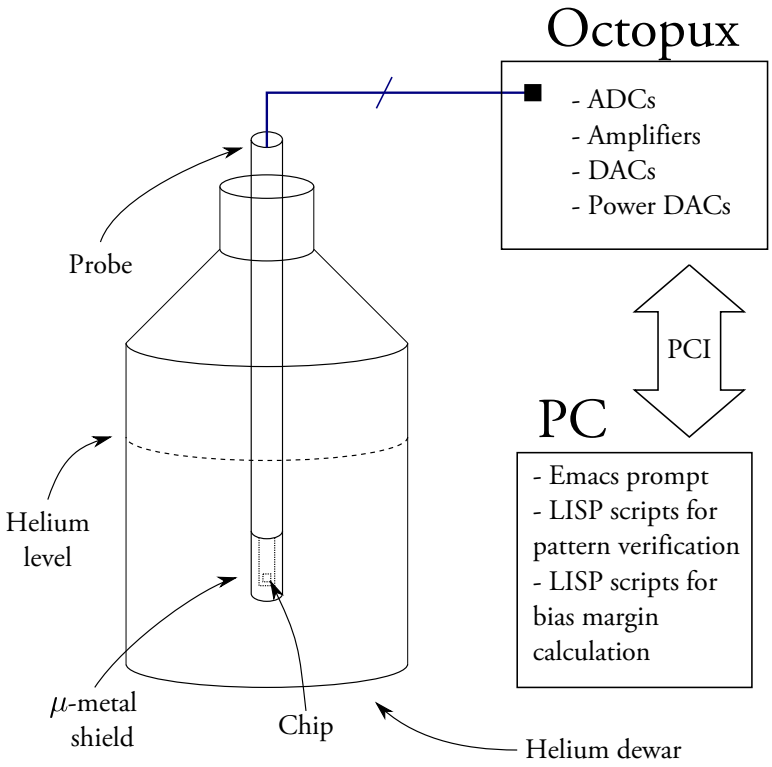

Figure 6.25: Octopux test setup employed for functional tests.

Note that this Octopux setup is not suitable for high-speed tests. The objective here was simply to confirm the correct operation of the designed cells by applying patterns and verifying responses. Correct operation of the shift register structures was established by feeding in a bit pattern and verifying its transmission with the correct delay (in terms of clock events). Correct n-bit deserialiser operation was established by feeding in a pattern of length n with the shift clock, applying a read pulse, and then verifying the parallel readout against the input pattern. This process was repeated several times to verify deserialiser operation.

High-speed tests

The high-speed tests were meant mainly as a confirmation of high-speed functionality and an evaluation of the test setup (as the present author had no experience on the test setup before). The first high-speed test setup employed for the shift-register structures is depicted in Figure 6.26.

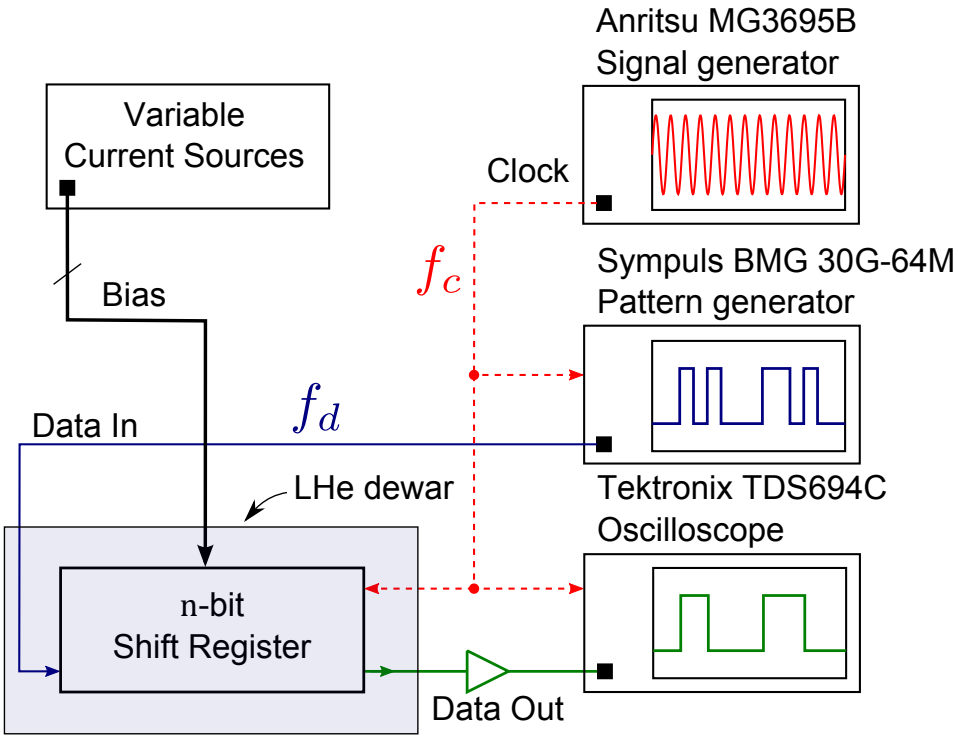

Figure 6.26: High-speed test setup for shift registers in Suite II.

To observe correct operation of the shift register, the following experiment was performed:

1. A data pattern generated by the high-speed pattern generator is applied to the shift register with frequency f_d. That is, each bit, whether `1` or `0`, is separated by a time period of $1/f_d$. The actual pattern is a broadband signal.

2. The pattern is shifted through the shift register with the high-speed clock generated by the signal generator. The clock frequency is f_c, an integer multiple of f_d. This ensures that the pattern is reproduced correctly at the output. The output is routed through a high-speed amplifier (with a gain of ~ 40 dB and a cutoff frequency of approximately 4 GHz) and displayed on a high-speed oscilloscope. Visual inspection of the output

signal is deemed sufficient as a means of testing whether the pattern is reproduced correctly.

3. The oscilloscope trigger is synchronised with the clock, but actual triggering is performed by the "frame clock", which is the period of the entire data pattern. On the oscilloscope screen, the output signal is thus displayed a time period τ_{off} after the trigger. For a certain base value of f_c, $(f_c)_1$, this offset is recorded as $(\tau_{\text{off}})_1$.

4. Clock frequency f_c is now doubled, but f_d does not change. As the data are shifted through the shift register faster now, the pattern should occur at a lower time offset from the trigger, although the shape of the pattern should not change. This can be confirmed by visual inspection again. By comparing the offset of the doubled f_c, $(\tau_{\text{off}})_2$, to the base clock offset, we can obtain the number of bits of the shift register n by evaluating

$$n = \frac{(\tau_{\text{off}})_1 - (\tau_{\text{off}})_2}{\frac{1}{(f_c)_1} - \frac{1}{(f_c)_2}}. \tag{6.5}$$

5. This relationship can be generalised for other clock frequencies, as long as f_c remains a multiple of f_d (in fact, an analogous relationship remains true even if not, but then the pattern may not be reproduced exactly):

$$n = \frac{\Delta\tau_{\text{off}}}{\Delta\frac{1}{f_c}}. \tag{6.6}$$

6. When the pattern is reproduced correctly and the shift is correct, the shift register is deemed to function as expected. This approach also allowed us to investigate bias margins by manually adjusting the bias currents until the output degrades.

One limitation of this test is the inability to apply data signals as fast as the clock signal. To work around this problem, a second, complementary test setup was devised. This test setup is depicted in Figure 6.27.

The second test was conducted as follows:

1. A high-speed sinusoidal signal with frequency f_c is generated and employed as shift register clock. Furthermore, the clock signal serves as input to the on-chip voltage-controlled switch mixer introduced in Figure 6.23.

2. A low-frequency square-wave signal with frequency f_m is generated and serves as control for the voltage-controlled switch. In this way, a "comb" signal is produced and applied to the data input of the shift register.

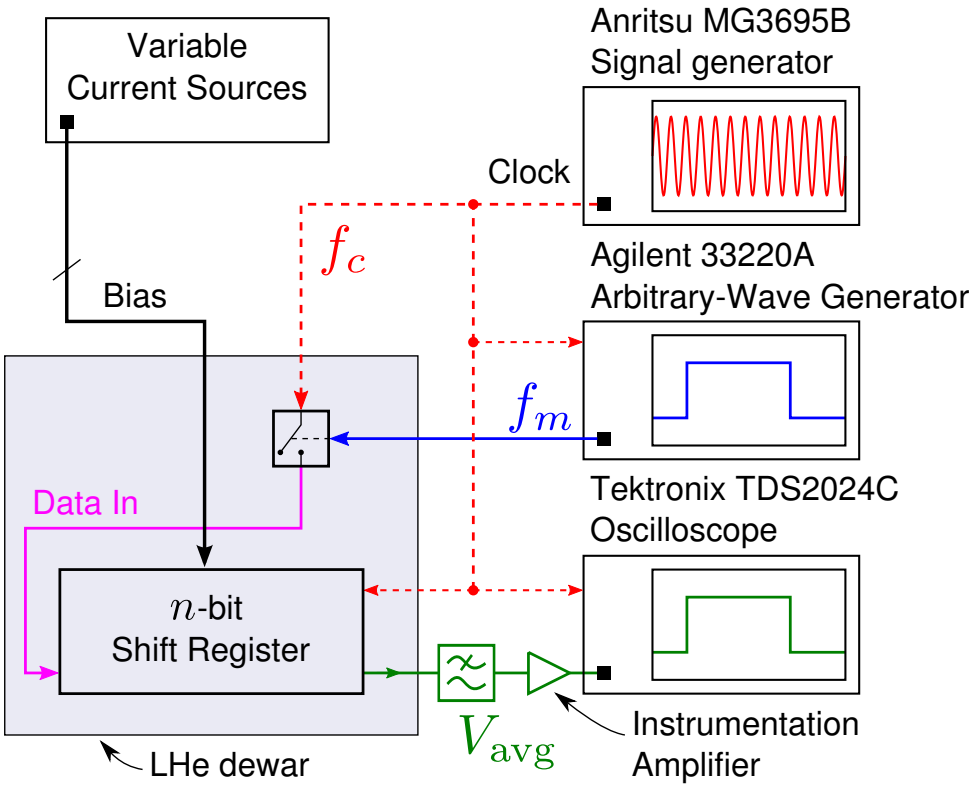

Figure 6.27: High-speed test setup for shift registers with average voltage pin and on-chip modulator.

3. The individual output data of the shift register, arriving with frequency f_c, are discarded. Instead, the average voltage on the output data line is displayed on an oscilloscope after low-pass filtering and amplification with well-known gain A. Cut-off frequency of the filter is chosen at several multiples of f_m (but far below f_c) to ensure that substantial harmonics of f_m are included and the square shape is maintained.

4. The oscilloscope then displays the envelope of the data signal output by the shift register. During correct operation, the output resembles the envelope of the comb signal, with the voltage of the crest of the comb at

$$V_{\text{crest}} = A\Phi_0 f_c. \tag{6.7}$$

5. Bias margins are obtained by varying the biases and observing the range over which the envelope is stable and (6.7) is statisfied.

While analogous tests could be devised for the deserialiser, this would be cumbersome as the number of output amplifiers and oscilloscope channels required would be high. In fact, we do not know a way to avoid the requirement of n amplifiers, but a logic analyser can be employed to acquire the amplified

output (instead of, for example, oscilloscope channels). Once acquired and in the digital domain, software can be employed to decode the output and check whether it conforms to expectations. This does not necessarily have to occur in real time.

It is difficult to apply an arbitrary high-speed pattern greater than 1 – 2 Gbps to an on-chip DC-SFQ converter. As the transfer impedance of the cable used to transmit the pattern generally increases with frequency, the transmitted lower-frequency components of such a (broadband) pattern have greater amplitudes than the transmitted higher-frequency components. The DC-SFQ circuit is a threshold circuit, but normally has several amplitude thresholds where a pulse is produced [9]. If the discrepancy in attenuation across the band is too high, it is not possible to find a transmitted signal amplitude at which, at the dc/SFQ input, each transition from low to high (or high to low) amplitude traverses only one threshold. Whereas this could potentially be remedied by distorting the transmitted signal accordingly, this is by no means trivial and was not considered practical for our tests.

Instead, an existing delta modulator [153, 154] was placed ahead of the deserialiser. This modulator has the purpose of generating a predictable high-speed pattern in the presence of a clock (clocks are narrowband signals and can thus be transmitted relatively easily). The modulator is fed with a (low-speed) analog input signal that determines the pattern it produces and feeds into the deserialiser. The deserialised data are then sent through a bank of amplifiers and routed to a logic analyser, where they can be picked up and demodulated. If the demodulated signal corresponds to the signal applied to the modulator, correct deserialiser operation is implied.

This deserialiser test setup is summarised in Figure 6.28.

6.6.3 Results

Short shift registers

Figure 6.29 depicts an example of the measured correct functional test patterns of the 16-bit eSFQ shift registers. Tests were exhaustive, that is, not only uniform clock and data patterns were employed, but also randomly generated ones. All tested shift registers passed functional testing.

To determine the bias margins of investigated structures, random 200-bit test patterns were applied to the devices under test for various bias currents. Various bias-current feeding configurations were investigated, most notably biasing from one pin only (repeated for each available pin), and biasing from all pins simultaneously. Only the largest identified continuous region of operation was captured. All tested shift registers passed functional testing for all tested patterns. The measured results of the bias current margin investigation are listed in Table 6.1.

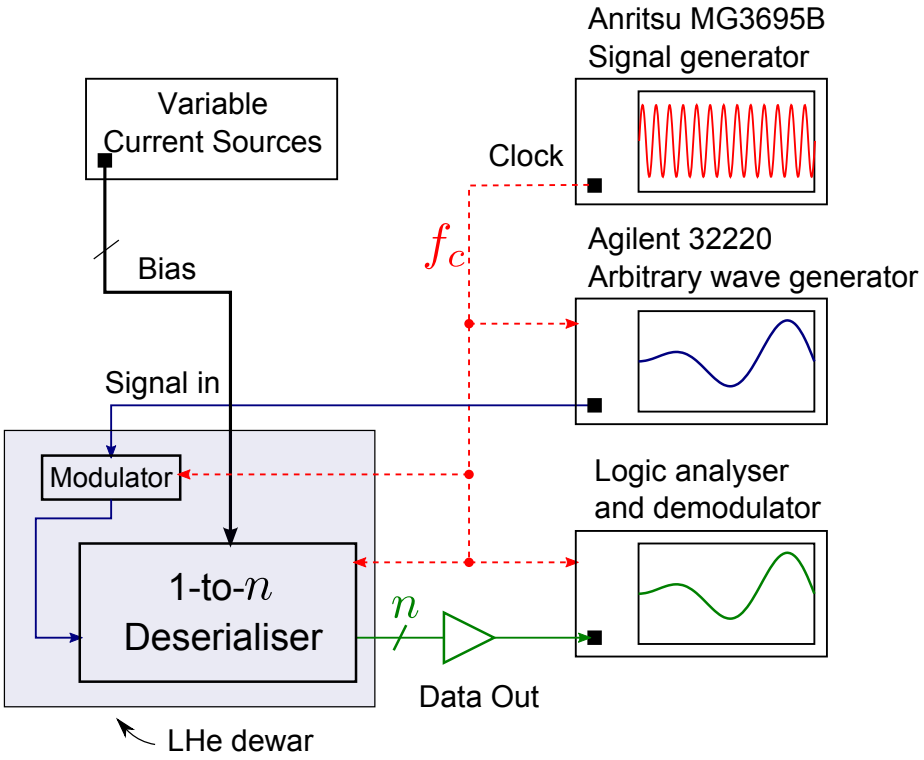

Figure 6.28: High-speed test setup for deserialisers.

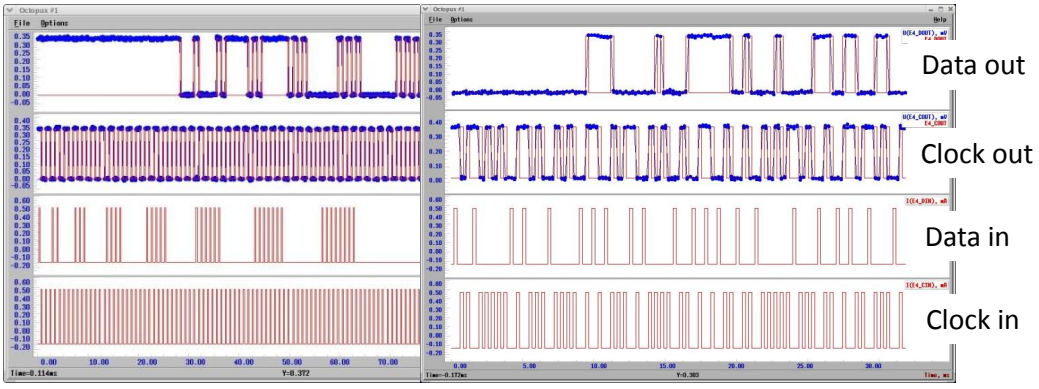

(a) (b)

Figure 6.29: Measured correct functionality of 16-bit eSFQ shift registers: simple pattern (a), randomly generated data and clock pattern (b). All 9 shift registers were fully operational.

Table 6.1: Experimentally determined functional bias margins for eSFQ shift registers across three chips (a comprehensive set of devices in different bias configurations) measured to establish functional correctness of the devices and attempt to identify desirable traits of the bias line layout.

Length	**Bias Line**		**Comment**	**Bias margins** [mA]	
	L_q	L_b		All pins	Single pin
32	Narrow	Medium		$13.9 \pm 27.5\%$	$15.4 \pm 14.5\%$
16	Narrow	Medium	Bias Tree	N/A	$6.71 \pm 35.5\%$
32	Wide	Medium		$14.0 \pm 29.8\%$	$13.7 \pm 24.5\%$
16	Narrow	Short		$8.38 \pm 21.8\%$	$8.2 \pm 24.6\%$
16	Narrow	Medium		$7.0 \pm 30.8\%$	$7.7 \pm 21.6\%$
16	Narrow	Long		$7.47 \pm 23.1\%$	$7.9 \pm 18.8\%$
32	Narrow	Medium	Flux bias	$13.0 \pm 27.0\%$	Not tested
16	Narrow	Medium	Flux bias	$7.6 \pm 13.0\%$	Not tested
16	Wide	Medium	Flux bias	$8.4 \pm 19.7\%$	Not tested

The measured bias margins roughly conform to expectations extrapolated from simulated results. Simulations relied on (small) 4-bit configurations to reduce computation time to design-friendly speeds and were performed in high-speed testbeds, not reflecting the actual devices under test or their periphery. The observed (rough) agreement with simulations confirmed that the designed structures are robust and scale well.

As Figure 6.30 indicates, the MeSR-based shift registers (with magnetic flux bias) functioned only when the magnetic bias is applied, which corresponds to our simulations. Bias margins do not seem dependent on the length of the shift register structure, although the dataset is too small to identify definite trends.

Simulations did not yield conclusive results for designing optimal bias line inductances and limiting inductors, although minimising the bias line inductance L_q seemed to be helpful in obtaining better bias-current distributions. The test structures, meant to comprehensively test the effect of these inductances on margins, did not show conclusive patterns either. We conclude from this that these inductances, at least when constrained to the parameter ranges investigated, do not significantly affect bias current distribution. This is a desirable result. Promising are the high bias margins measured for the 16-bit shift register biased from a single pin feeding a bias tree. This suggests that the off-chip biasing effort for eSFQ systems should be comparable with systems based on conventional RSFQ.

Of the high-speed structures, all tested structures passed low-speed testing (this is merely a confirmation of the functional tests, but does indicate that our Suite II chip designs are sound). The high-speed manual test setup is separate

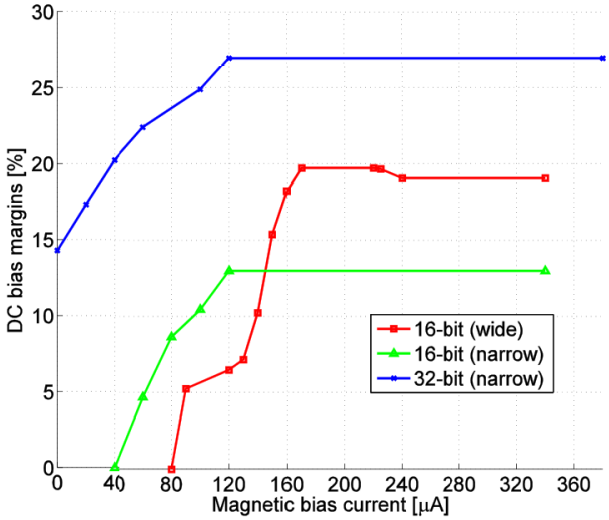

Figure 6.30: Observed bias margins for magnetically biased shift registers as a function of magnetic bias current.

from the Octopux-based test setup, in which all signals can be controlled programmatically in a closed system.

First, we tested a 32-bit eSR-based shift register with the first test setup described in Figure 6.26. The clock frequency f_c was varied from 1 GHz to 24 GHz, producing the output traces recorded in Figure 6.31. The differences between the offsets ($\Delta\tau_{\text{off}}$) are as expected from (6.5), confirming that the tested shift register indeed shifts by 32 bits (clock periods). Each tested shift register performed as expected here.

Next, the margins of several shift registers, eSR- and MeSR-based, were investigated by manually varying the bias currents and observing the output. Although many structures were manufactured, as described in detail in Table A.13, only so many were tested as deemed necessary to identify trends. This is common practice in superconducting circuit design, as testing is expensive, due to the cost of liquid helium and limited time on the specialised test equipment. These difficulties are generally more pronounced for high-speed tests than for functional tests.

Measured shift register margins at high speeds for the first test setup are recorded in Figure 6.32.

The measured margins are lower than found by 4-bit simulations, but the tested structures are longer than 4 bits, and thus this is expected. It is normal for measured bias margins to be lower than predicted by simulations, as many real effects are not considered adequately in simulations. These effects include fabrication tolerances, flux trapping, parasitic magnetic fields (for example due to bias currents), other parasitics, impurities, undersired side-effects/loads of the testing setup, temperature fluctuations, electromagnetic interference and many more. In fact, the purpose of analysis and optimisation at the circuit

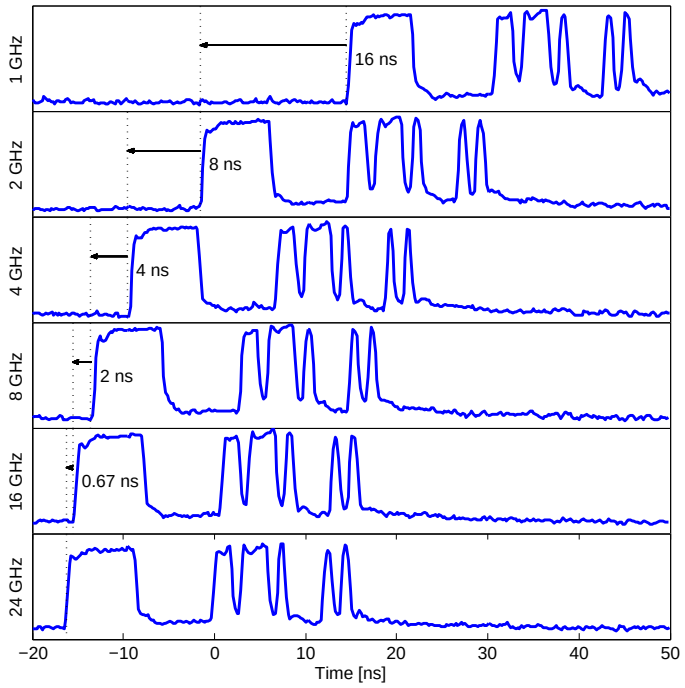

Figure 6.31: Amplified output of a 32-bit shift-register. The output pattern matches the input pattern, and the (correct) spacings of τ_{off} are shown for different clock frequencies.

level is to achieve circuits that are sufficiently resistant to these effects, not to predict bias margins.

In the light of all this, we are proud to see that the region of operation extends continuously over the investigated frequencies, with no significant narrowing at either edge (indicating that the devices under test should continue to perform beyond the investigated frequencies)!

Furthermore, we note that the height of the operational region for the 32-bit shift registers with a wide bias line is greater than for those with a narrow bias line. This is a trend we have observed in some simulations (though not conclusively). We do not see a similar trend among the 16-bit shift registers with corrective bias, which suggests that the corrective bias (applied magnetically) could have an effect on the bias current distribution. This should be explored further in future refinements of the concept.

For our second test setup involving average voltage measurement, we endeavoured to compare the operational region of `eSR-` and `MeSR` shift registers (as discussed in the list of Suite III considerations in Section 6.6.1). The two 32-bit structures on Chip `[...]` were subjected to the second test setup, described previously in Figure 6.27. Chosen clock frequencies f_c ranged from $1 - 20\,\text{GHz}$, whereas the modulation frequency was fixed at $f_m = 4\,\text{kHz}$ and the low-pass filter cut-off was at about 50 kHz. Deterioration of correct op-

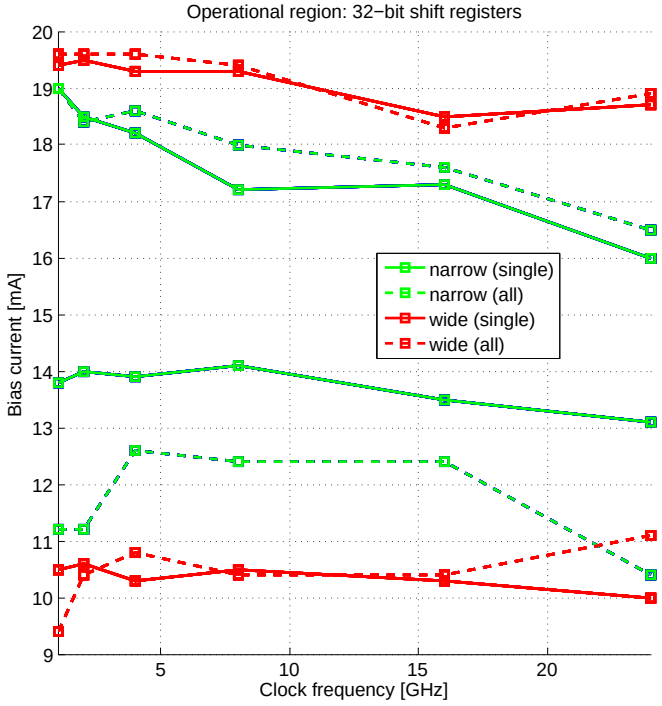

(a)

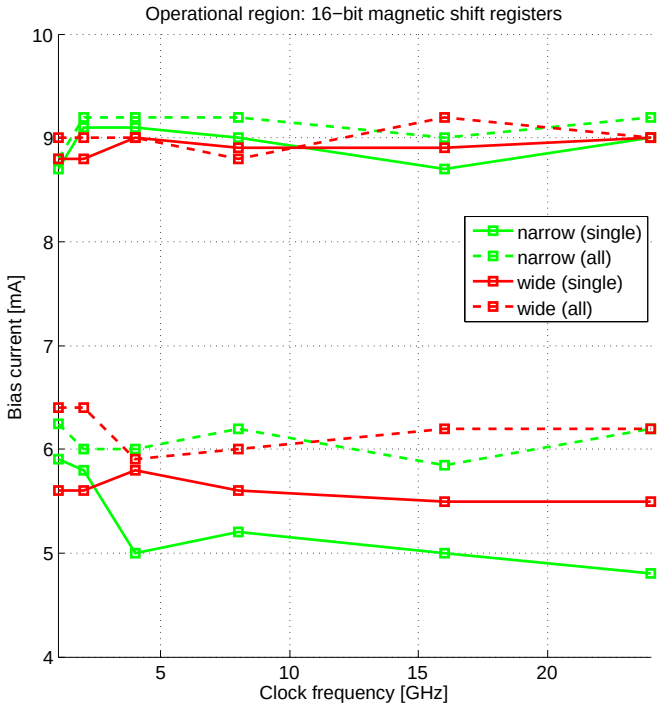

(b)

Figure 6.32: Observed operational range for eSFQ shift-registers without (eSR) (a) and with (MeSR) (b) a magnetic corrective bias. Upper and lower lines of the same style form boundaries of the region of operation.

eration was detected by a visible reduction in steepness of the envelope edge gradient. Note that this test is flawed on its own, as no shift can be reliably measured. Its results must be viewed in context of the previous test, during which correct shifts were established.

The range of bias currents for which correct operation was observed are visually reported in Figure 6.33.

Again, recorded margins compare well with simulated margins, although this agreement is far more pronounced in the case of the MeSR shift register. Again, the differences in operational region between single-pin biasing and all-pin biasing are not very pronounced in the MeSR case, confirming that the corrective magnetic bias enables a safer, more predictable design.

In terms of power dissipation, somewhat surprisingly, both devices perform equally well. To obtain a useful metric for comparison, we divide (6.3) by the employed clock frequency f_c and the number of bits n ($n = 32$ in this case), obtaining the dissipated energy per-bit and per-clock event,

$$E_b = I_b \Phi_0, \tag{6.8}$$

and insert for I_b the center point of the operational region. Using the worst-case center across all investigated frequencies, the figures reported in Table 6.2 are obtained.

Table 6.2: Bit energy E_b of 32-bit eSR- and MeSR-based shift registers for different bias configurations.

	Simulation	**Single pin**	**All pins**
eSR-based	1.01 aJ	0.81 aJ	0.91 aJ
MeSR-based	0.83 aJ	0.84 aJ	0.87 aJ

The bit energy of the MeSR-based shift register is in better agreement with simulations, whereas the bit energy of the eSR-based shift register is substantially lower than derived from simulations. We believe that deviations can be attributed to inconsistencies between simulated and experimental structures (in terms of n, for example), fabrication tolerances, experimental influence and noise.

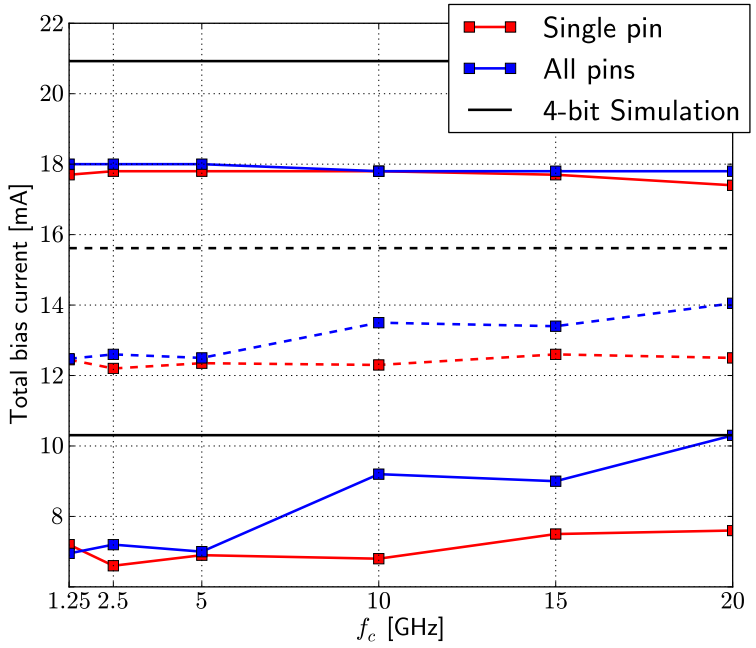

(a)

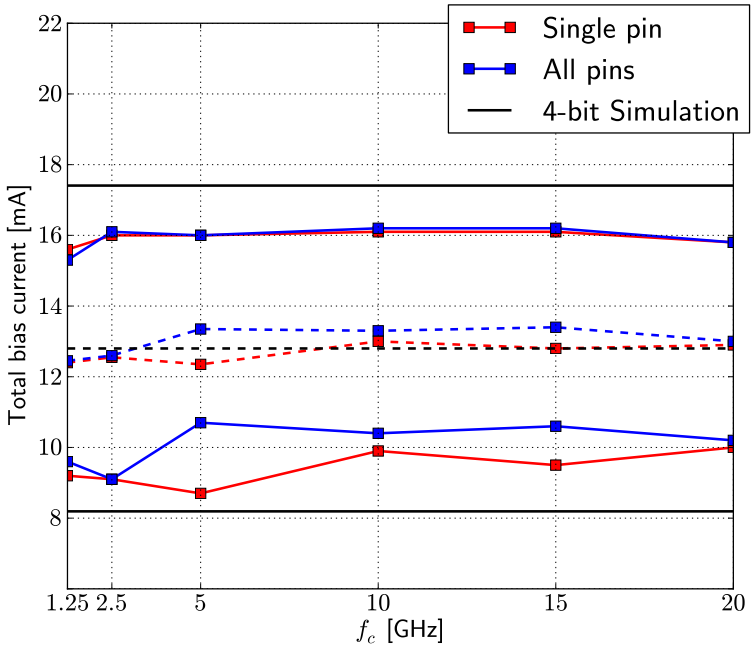

(b)

Figure 6.33: Calculated operational range for the second high-speed test setup, for eSFQ shift-registers without (eSR) (a) and with (MeSR) (b) a magnetic corrective bias. Upper and lower solid lines of the same style form boundaries of the region of operation, whereas dashed lines indicate the corresponding centers of the region.

Long shift register

Suite III included chips with a long (184-bit) shift register based on eSR. The choice of eSR is motivated by wider experimentally observed bias margins. The bias pad layout for this chip (see Figure 6.21) was conservative as well, totalling 23 pins, at 8 pins per bit. This facilitates maximum flexibility during the experiment.

Functional testing of the 184-bit shift register was not done in Octopux, but performed manually for a complex data pattern and a periodic clock. Correct operation was verified by inspection on an oscilloscope, with correct operation defined as the correct propagation of the pattern at the correct shift. An example of correct operation at MHz frequencies is shown in Figure 6.34. The shift of $184 \cdot \frac{1}{3.2\,\mathrm{MHz}} = 57.5\,\mu\mathrm{s}$ is correct.

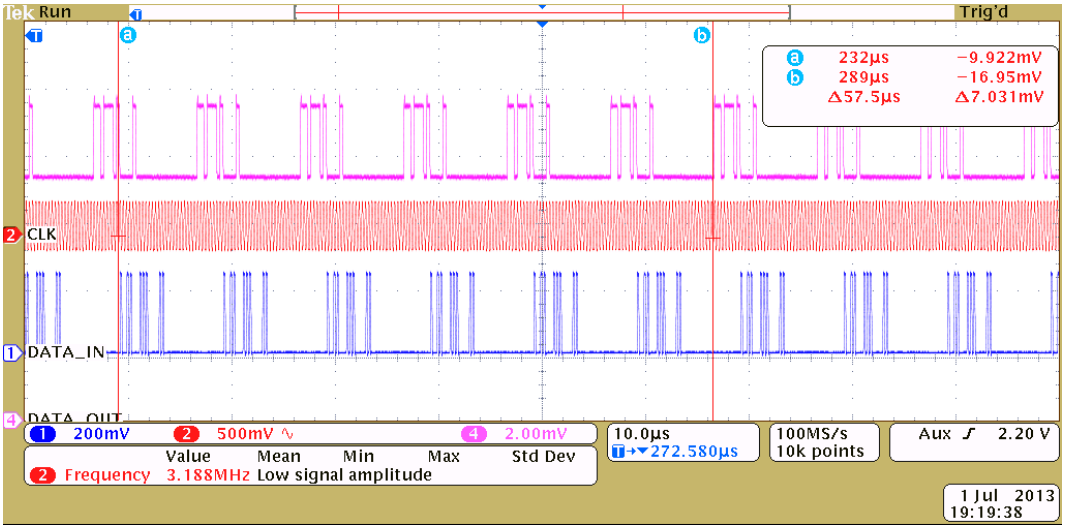

Figure 6.34: Oscilloscope traces for a functional 184-bit shift register test. The complex test pattern is reproduced with the correct delay. A 3.2 MHz clock and data signal are applied.

For verification at high-speed, the first high-speed shift register test (see Figure 6.26), meant to verify the correct shift, was performed. An additional, more comprehensive test meant to test complex, arbitrary data signals at high speeds was designed but could not be carried out due to lack of equipment at the time. The shift for two different clock frequencies is compared in Figure 6.35.

To exclude experimental equipment delays, we compared delays of all measured frequencies to a reference frequency chosen at $f_1 = 1\,\mathrm{GHz}$ (akin to the earlier high-speed shift test on short shift registers). The delay differences, measured in this way, are compared against the correct value predicted by

(6.5) in Figure 6.36. Clearly, the measured delays compare well to the theory, supporting correct operation.

Operational bias margins were measured for several different frequencies and bias configurations. Verification by visual inspection was performed at $f_c = 3.2\,\text{MHz}$. Verification by delay difference was employed for $f_c = 184\,\text{kHz}$ and $f_c = 24.3\,\text{GHz}$. The measured bias margins and associated power dissipation are reported in Table 6.3.

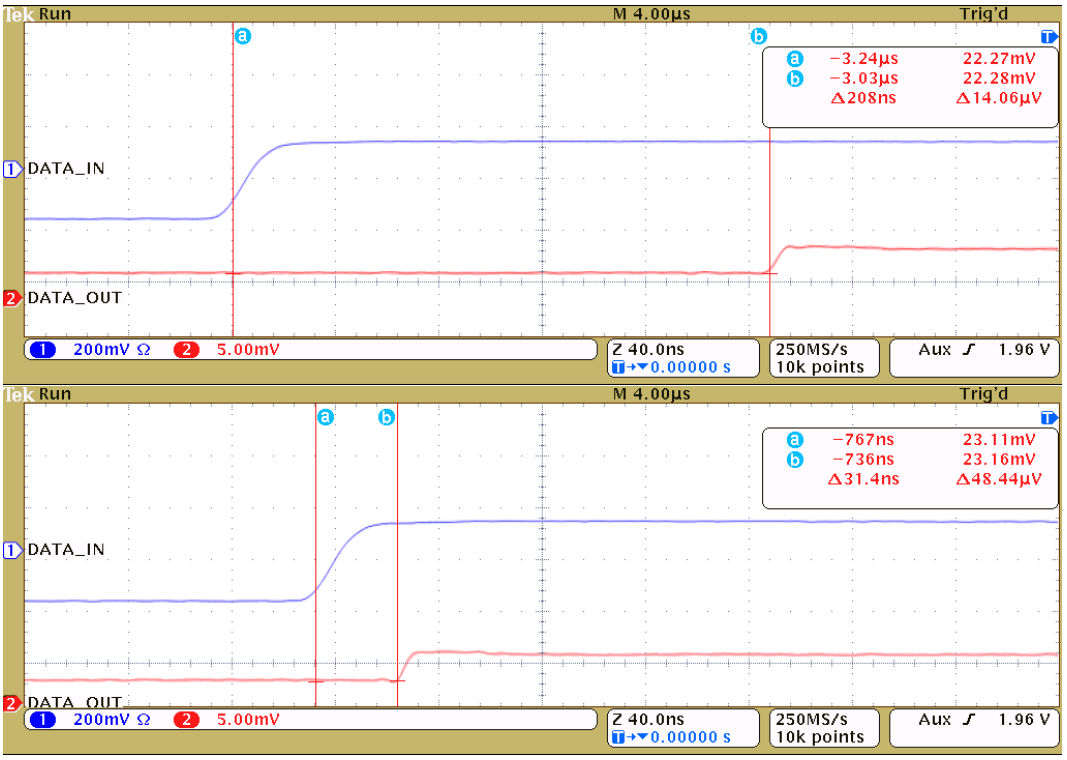

Figure 6.35: Oscilloscope traces for high-speed 184-bit shift register test. The upper set of traces corresponds to $f_c = 1\,\text{GHz}$, the lower to $f_c = 24.9\,\text{GHz}$ clock. In each case, the upper trace is the input data signal and the lower trace is the output data signal.

The measured margins for the 184-bit shift register are lower than for the 32-bit shift register. This is not surprising (and commonly reported in the literature), for at least two reasons related to fabrication tolerances:

1. The effect of fabrication tolerances on a single shift-register cell could be profound or negligible, depending on the process, but any effect is cumulative. Hence the probability of a repetition of n cells to work correctly for any given operating point is roughly P_c^n, where P_c is the probability of correct operation of a single cell.

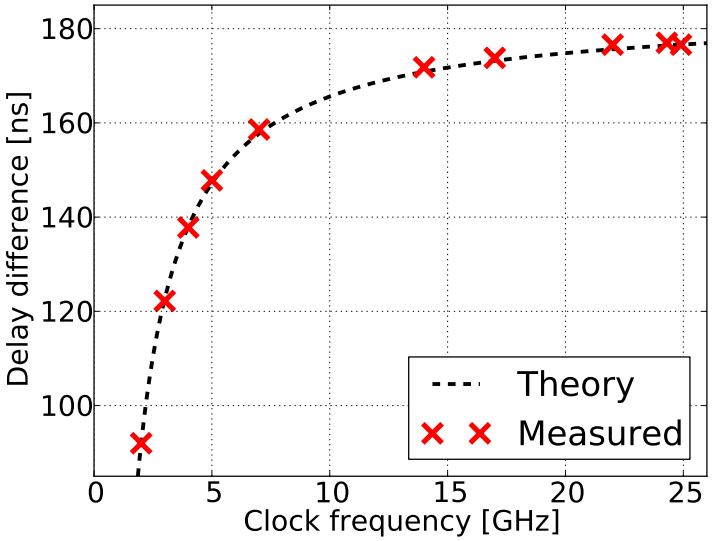

Figure 6.36: High-speed test results for 184-bit shift register. The ordinate refers to the delay difference between the investigated clock frequency and the reference frequency (1 GHz).

Table 6.3: Measured bias margins and associated power dissipation for the 184-bit shift register for several bias configurations.

f_c	**Bias-injection Configuration**		
	All pins[1]	**All pins**[1]	**3 pins**[2]
	6 varied equally	6 varied	3 varied equally
24.3 GHz[3]	97.6 ± 15.6 mA	*(not tested)*	105.1 ± 10.4 mA
	1.10 ± 0.18 aJ/bit	-	1.18 ± 0.12 aJ/bit
3.2 MHz[4]	100.6 ± 4.7 mA	98.5 ± 11.1 mA	no margin
	1.13 ± 0.05 aJ/bit	1.11 ± 0.12 aJ/bit	-
184 kHz[3]	96.5 ± 12.3 mA	*(not tested)*	97.1 ± 10.3 mA
	1.08 ± 0.14 aJ/bit	-	1.09 ± 0.12 aJ/bit

[1] All 23 pins received bias current, the current into 6 pins was varied to obtain the margins.

[2] Two pins at the edges and one pin in the middle of the shift register received current and were varied.

[3] Delay test, as in Fig. 6.35.

[4] Complex pattern test, as in Fig. 6.34.

2. A larger shift register occupies greater chip area, meaning that global fabrication tolerances could play an increasingly important role. By contrast, in small structures local tolerances are important and global tolerances are less likely to play a role.

The per-bit energy dissipation E_b of the 184-bit shift register is closer to the simulated value than for shorter shift registers (compare Table 6.2). Good agreement with simulations here is evidence for the correctness of our design method, as a progressive narrowing around the center of the operational region as circuit complexity increases suggests a good choice of nominal design value. Supporting evidence for correct operation lies in the consistency of the recorded margins across the frequencies investigated.

Bias configuration played a non-negligible role in observable bias margins. For complex patterns, no operation could be achieved when biasing from 3 pins only, and for no pattern could operation be achieved when biasing from 1 pin only. We explain this by noting that the bias line was laid-out very close to the active cells (this was not a design decision, simply a convenience measure). For this large shift register, biasing from few pins means placing large amounts of current near active Josephson technology. Although we attempted to shield the bias line and laid-out the cell in a superconducting cavity for this reason, later simulations by Fourie with new software [155] unavailable at design time showed that a significant electromagnetic field still coupled into the cell cavity, potentially supressing critical currents and interfering with cell operation. (Compelling visual evidence of this was exhibited by Fourie at a recent conference [156].)

We believe that there are several solutions to this problem. Better shielding should be achievable with processes currently under development and slated for reliable operation shortly, as wiring and biasing layers can be placed below the ground plane [**?**]. Also, we propose a means to better control bias current distribution in Section 6.7.

Deserialisers

All fabricated and tested deserialisers passed functional testing. An example snapshot of an exhaustive functional test performed with Octopux [152] is shown in Figure 6.37. As for the functional shift register tests performed with Octopux, both regular and random patterns were employed to verify correct operation. Margins recorded during functional testing with Octopux are reported in Table 6.4.

In general, bias margins are good and demonstrate scalability of the concept. Comparatively low bias margins were recorded for the 16-bit deserialiser, but these were not punishingly or even unexpectedly low. We attribute this observation to the relative complexity of the test structure and to an ill-designed monitor setup. The high fan-out of the deserialiser requires a large number

of SFQ-DC monitors as well as several DC-SFQ converters to apply the test patterns. Due to the high pin count of the test structure (resulting from the requirements of two bias lines and parallel output), all peripheral cells were biased from a single pin, which yielded low margins for the peripheral bias, potentially depressing the margins of the device under test.

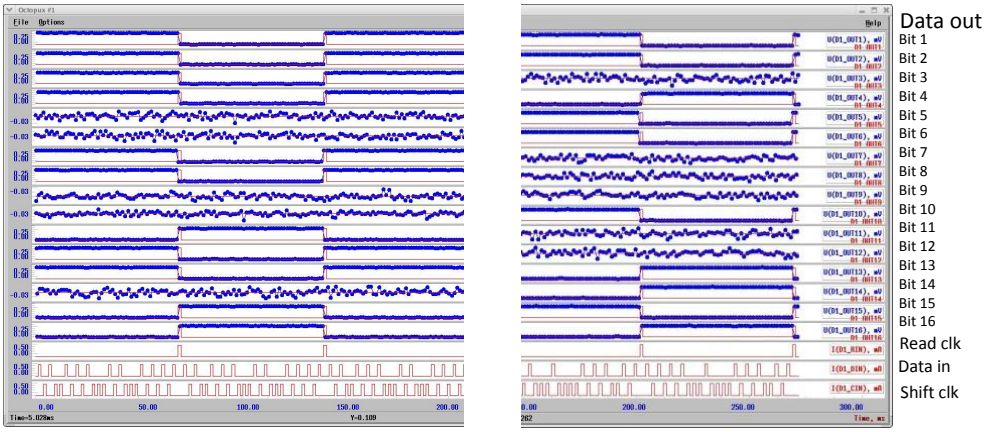

(a) (b)

Figure 6.37: Measured correct functionality of 16-bit deserialiser for different input patterns, (a) and (b). This circuit divides the input serial stream into 16 output parallel streams. Random data and shift clock pulses are used for testing. Note that measured voltages are scaled automatically, making some constant-level bits appear noisy.

The deserialiser structures under test worked without the magnetic flux bias applied (this is consistent with simulations). As demonstrated in Figure 6.38, applying the magnetic flux bias improves bias margins, doubling them in the case of the 16-bit deserialiser. For the 16-bit deserialiser, the clock and read bias margins exhibit comparable absolute values and dependence on the magnetic flux bias, which is consistent with simulations. The 8-bit structure does not mirror this symmetry, which may indicate higher susceptibility to the periphery.

Testing the deserialisers at high speed is currently ongoing. Since loose amplifiers (not part of other equipment) and high-speed oscilloscope channels are limited, and high-speed patterns cannot be applied reliably from off-chip, it was not possible to test a manufactured free-standing deserialiser (without the generating modulator) with our test setup at high speeds. We have tested it manually at low speeds, applying a unique pattern at kHz frequencies and displaying the output on an oscilloscope (see Figure 6.39a).

Table 6.4: Experimentally determined bias margins for eSFQ deserialisers across two chips (a comprehensive set of devices in different bias configurations) measured to establish functional correctness of the devices and attempt to identify desirable traits of the bias line layout.

Length	**Bias Line**		**Comment**	**Bias margins** [mA]
	L_q	L_b		All pins
16	Narrow	Medium	Flux bias	$6.8 \pm 14.4\%$ (Clock)
				$7.7 \pm 12.3\%$ (Read)
8	Narrow	Medium	Flux bias	$3.7 \pm 22.9\%$ (Clock)
				$3.3 \pm 31.9\%$ (Read)
4	Narrow	Medium	Flux bias	$1.4 \pm 47.4\%$ (Clock)
				$1.5 \pm 36.5\%$ (Read)

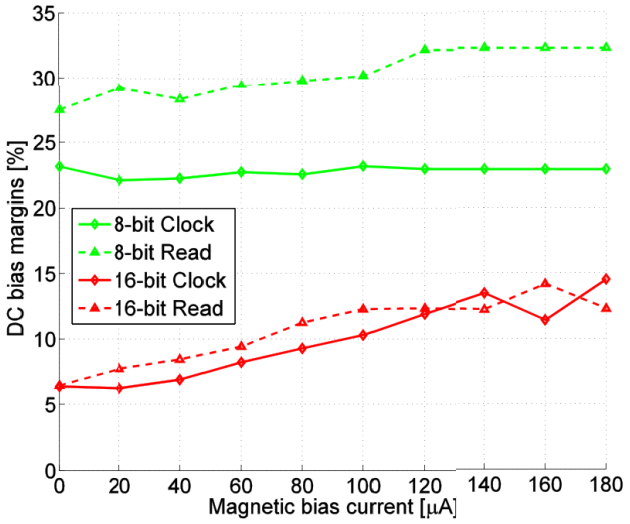

Figure 6.38: Observed bias margins for (magnetically biased) deserialisers as a function of magnetic bias current.

In the absence of an input signal, the employed modulator emits the repeated pattern `10101010...`. We confirmed that the deserialiser in series with the modulator worked correctly by applying a low-speed clock and displaying the output on an oscilloscope. As Figure 6.39b shows, the output is correct.

Unfortunately, we discovered a design error that precluded application of the magnetic corrective bias current to any manufactured deserialiser in the Suite II collection of high-speed test chips. As the deserialiser was optimised in the presence of this corrective bias current, the observed margins without it were somewhat narrower than expected. It was judged unlikely to work at high frequencies (without much effort and expense), thus we decided on a redesign instead. Redesigned deserialiser test chips were fabricated as part of Suite III. Unfortunately, tests of these redesigned chips have not yet been performed due to constraints on time and travel funds.

The demodulator setup does not work at low frequencies (this is so by design), which means that we could not employ it to demodulate our low-frequency signals. However, this would not have been of much value even if it were possible.

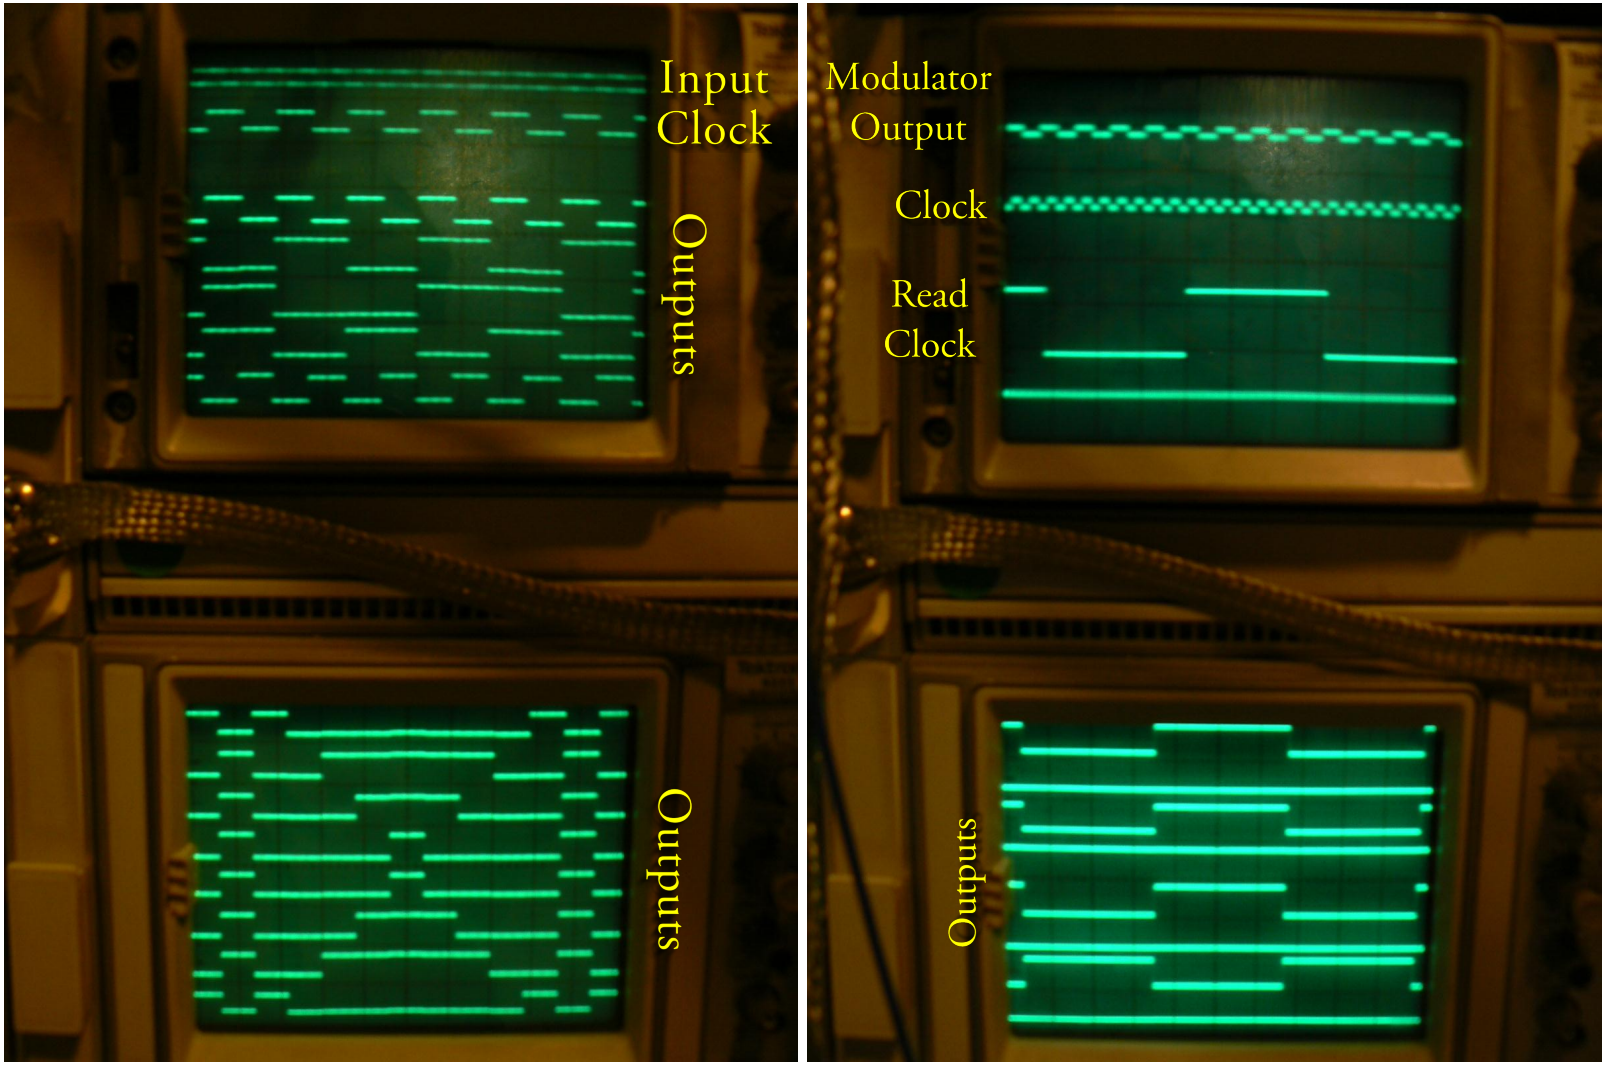

(a) (b)

Figure 6.39: Output of free-standing deserialiser (a) and deserialiser behind modulator (b).

6.6.4 Discussion

We are very happy with both the functional tests and the performed high-speed tests. Every structure that was tested in Octopux passed functional testing and had acceptable margins that promise reproducibility and scaling. The response to the magnetic flux bias was as predicted, which commends our design technique and simulations. Inconsistencies can, in our opinion, easily be attributed to experimental factors and fabrication tolerances outside our control.

The high-speed tests went well too, the limitations lay in the room temperature equipment rather than the superconducting chips. Applied patterns were reproduced correctly at clock speeds up to 24 GHz for shift registers and at low speeds for the deserialisers. We are confident that higher speeds, such as 20 GHz, for the deserialiser, are entirely possible, and lament that we did not have the time and equipment for more extensive testing before the cutoff time for this text.

It has been experimentally verified that our designed deserialiser is capable of reducing the data rate of the output from a superconducting flash ADC by the desired factor 16 (see Figure 6.9). Although we could not confirm this experimentally for high speeds, we believe that the volume of our high-speed tests of other structures (which went through the same design procedure) and the correct performance of the deserialiser at low speeds make the correct operation of our deserialiser at high speeds (20 GHz+) very likely.

The energy required for deserialisation alone should be no more than $\sim$ 1 aJ/deserialiser bit, which corresponds to a power dissipation of $P \sim 0.3\,\mu W$ per ADC output bit at $f_{clk} = 20\,GHz$. This is negligible in the power budget even within the cryocooled environment (the resistor ladder of the ADC alone will probably dissipate several mW, as explored in Section 5.6.1).

Perhaps most importantly, we have demonstrated, for the first time by simulation and experimentally, that eSFQ is a viable logic family. Its ultra-low-power dissipation, high speed, elegance, low IC footprint and biasing effort comparable to RSFQ, makes eSFQ an excellent candidate for further development with the eventual goal of high-performance computing and signal processing. There is no doubt that successes in this direction would have profound benefits for the SKA and beyond.

6.7 Towards broader application of eSFQ concepts

There are at least two aspects of eSFQ logic as described in this chapter that preclude it from universal application:

1. A bias-injection point at which a 2π increase occurs during every clock

period is required, but not all gates, particularly unclocked ones, have such a point.

2. The bias-current distribution achieved by eSFQ bias networks is not as predictable as for conventional RSFQ.

In this section we provide an overview of a possible approach to a solution of these problems.

6.7.1 The conventional RSFQ TFF

The RSFQ *toggle flip-flop*, or TFF, was first proposed by Likharev and Semenov in 1991 [9]. Since then it has become a popular circuit for technological demonstrations [19, 48] and finds many applications in practical circuits. The TFF is schematically depicted in Figure 6.40.

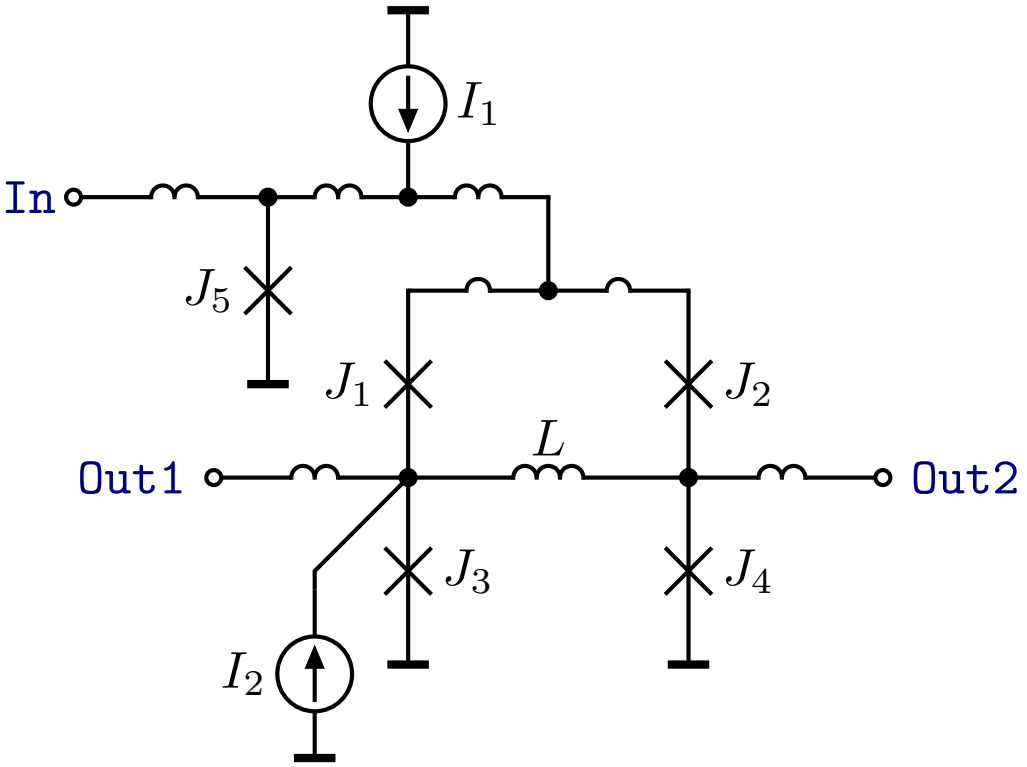

Figure 6.40: Schematic of the RSFQ TFF.

Parameters of the TFF are dimensioned so that it can have one of two states, which correspond to the direction of the steady-state current through storing inductor L.

Initially, the bias current I_2 ensures that the current flows, on the schematic, from right to left, which one might say corresponds to half an anti-fluxon deposited in the lower loop J_2, L, J_4. A pulse arriving at `In` causes J_5 to switch,

sending a current down the two legs (J_1,J_3) and (J_2,J_4). This current superimposes onto the current in J_3, causing J_3 to switch. Correspondingly, an SFQ pulse exits at `Out1`. The switching of J_3 also deposits a flux quantum into storing loop J_2, L, J_4, essentially switching the direction of the current through L. The storing loop now "stores" half a fluxon and J_4 is biased. Note that the 2π phase-jump of J_3 is equalised by a switch of J_2.

A subsequent SFQ pulse arriving at `In` again causes J_5 to switch and a current to flow down the two legs of the TFF. Now J_4 is biased and switches, causing an SFQ pulse to exit at `Out2`. Also, an anti-fluxon is deposited into the lower storing loop, resetting the TFF into the initial state. J_1 switches to equalise the 2π phase-jump across J_4.

A popular variant of this cell is the *asymmetric TFF*, which discards one of the outputs, resembling a digital frequency divider. Every second pulse produces an output, every other pulse is discarded. This TFF is "asymmetric" as it is generally optimised so that one of the outputs is not produced at all, yielding asymmetric parameters across the two legs.

6.7.2 Conversion to eSFQ

The ubiquity of the TFF makes it a desirable circuit for conversion to eSFQ. However, the TFF does not have a `Clock` terminal, whereas until now we have discussed only clocked gates for conversion in a synchronous system. A natural bias current injection point, at which a 2π phase increase occurs during every clock period, is thus not available.

While there are ballistic versions of the TFF [137], which do not require an explicit bias current, they do not exhibit high parameter margins and generally present a greater load to adjacent cells, making their arbitrary inclusion in standard circuits cumbersome.

The solution lies in a re-examination of the natural bias current injection points of previous eSFQ circuits. All bias currents are injected either atop junctions that switch during every clock period (that transport the clock), or decision-making pairs that are interrogated during every clock period. This idea of injecting atop a decision-making pair can be extended to a "decision-making tuple" (DMT), as depicted in Figure 6.41.

An interrogatory SFQ pulse propagates along the clock line until it reaches the DMT, down which it sends a current. Now, depending on how the DMT junctions J_1 through J_n are biased and dimensioned, exactly one of the junctions switches, equalising the 2π phase difference on the clock line and enabling the clock pulse to continue propagating along the clock line. In this case, the eSFQ bias current can be injected anywhere on the clock line, including atop the DMT.

This idea cannot be extended to DMTs of arbitrary height n, as a larger n brings with it a larger series inductance of the path of interrogation. Too large

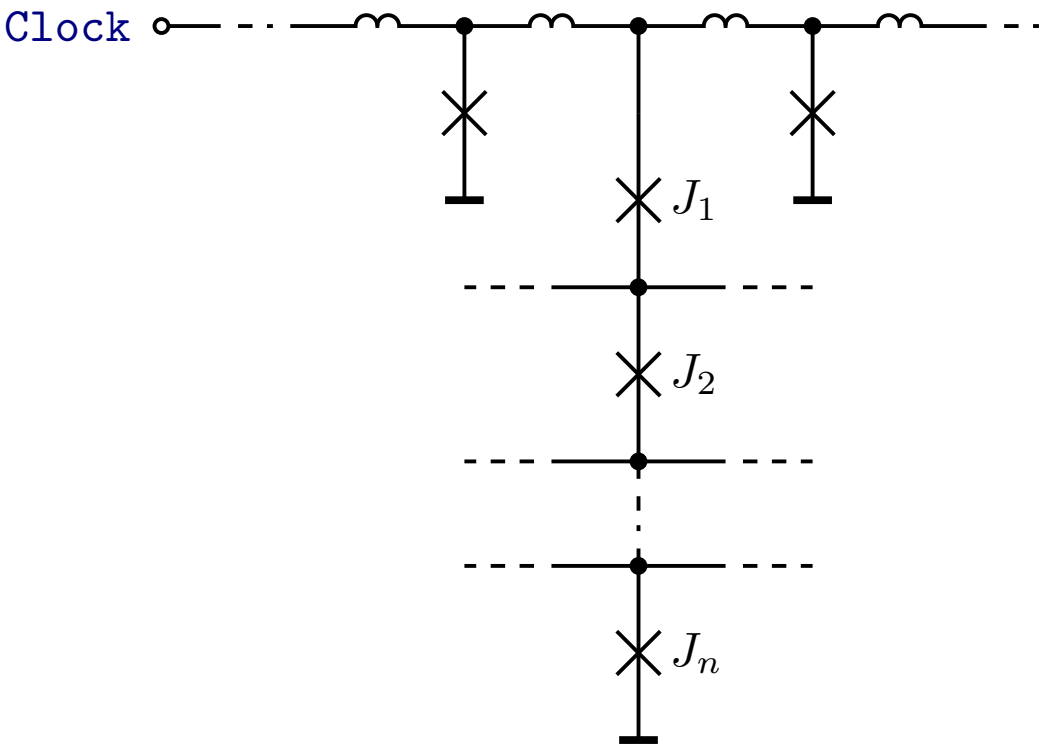

Figure 6.41: A decision-making tuple (DMT), interrogatable by a clock signal.

an inductance cannot be interrogated by an SFQ pulse, as the interrogation current would be too small to enable sufficient parameter margins.

For our purposes, we propose to employ a DMT of height $n = 3$, of which the lower two junctions are replaced by the legs of the TFF. Our concept is depicted in Figure 6.42. Dr Oleg Mukhanov of Hypres, Inc. is credited with the invention of this concept.

The TFF core is inserted into an eSFQ shift register derived from the two-junction RSFQ shift register demonstrated by Mukhanov in 1993 [138]. An eSFQ version of this shift register has recently been demonstrated by Kaplan and Vernik, but this work is as yet unpublished.

The function of `eTFF` and the previously discussed TFF are substantially the same, but the exact workings of the circuit need re-examination. There are essentially four important states, made up from the two possible states of the TFF core multiplied by the two states of "priming" inductor L_{da}. When inductor L_{da} is storing, a current flows down the legs of the TFF core. This current is not large enough to switch any junction in the core, but it "primes" the core for switching when the DMTs (J_5, J_1, J_3) and (J_5, J_2, J_4) are interrogated by the clock. The four states and the reaction of the TFF to an interrogating clock signal are summarised in Table 6.5.

Whether storing inductor L_{da} is storing upon a clock event thus determines whether the TFF is "evaluated" or not. Inductors L_{da} and L_{db} both make up

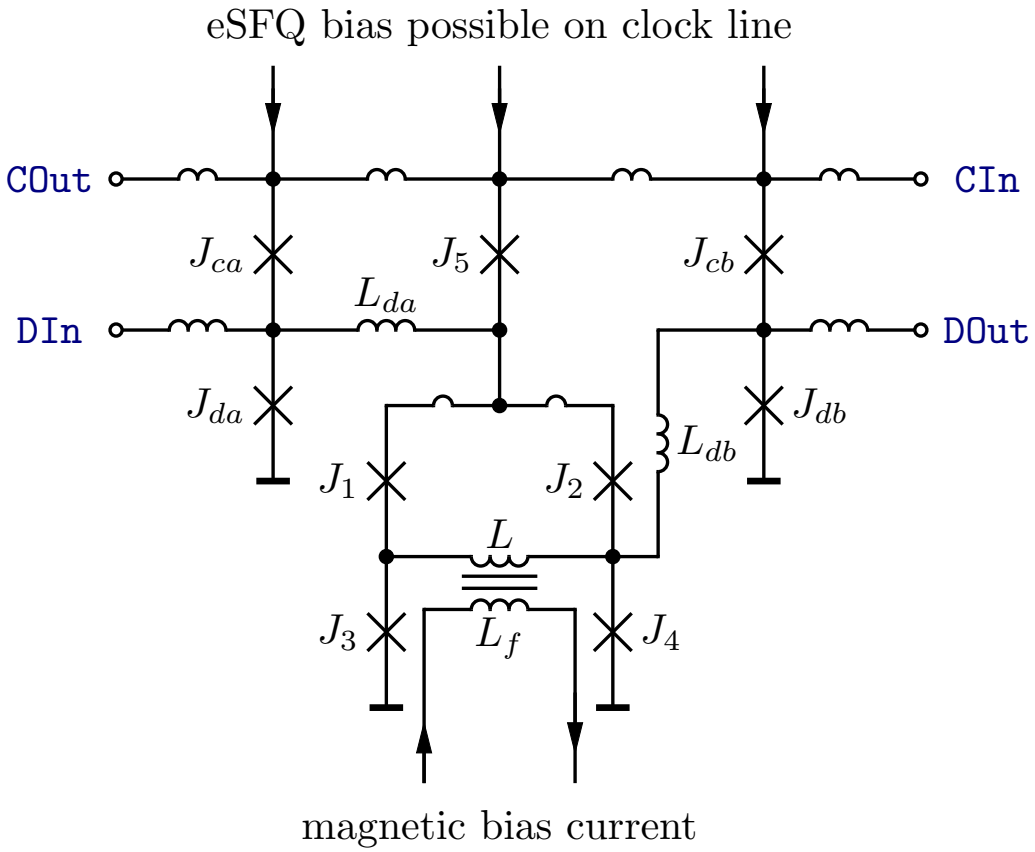

Figure 6.42: Schematic of cell `eTFF`, an eSFQ compatible version of the ubiquitous TFF cell.

Table 6.5: States of `eTFF` cell and response to interrogation.

State of L_{da}	**State of J_3, J_4**	**Response to interrogation**
not primed	J_3 biased	J_5 switches. No change in state.
not primed	J_4 biased	J_5 switches. No change in state.
primed	J_3 biased	J_3 and J_2 switch. J_4 becomes biased. L_{da} becomes non-storing (not primed).
primed	J_4 biased	J_4 and J_1 switch. L_{db} becomes storing, propagating a digital `1` along the shift register. J_3 becomes biased. L_{da} becomes non-storing (not primed).

the data path of the shift register, hence every second pulse arriving on the data line of the shift register is propagated, every other pulse is discarded (here by means of the termination resistor R_{term}). This corresponds to the operation of the asymmetric TFF gate described earlier.

In this way, the TFF is artificially synchronised through the addition of a clock line, but retains high parameter margins because it is possible to bias it.

6.7.3 Simulation of `eTFF` cells

The actual circuit that was designed and optimised is depicted in Figure 6.43. There are small topological differences between it and the conceptual schematic shown in Figure 6.42, but these are insubstantial and do not distract from the functionality described in the previous subsection.

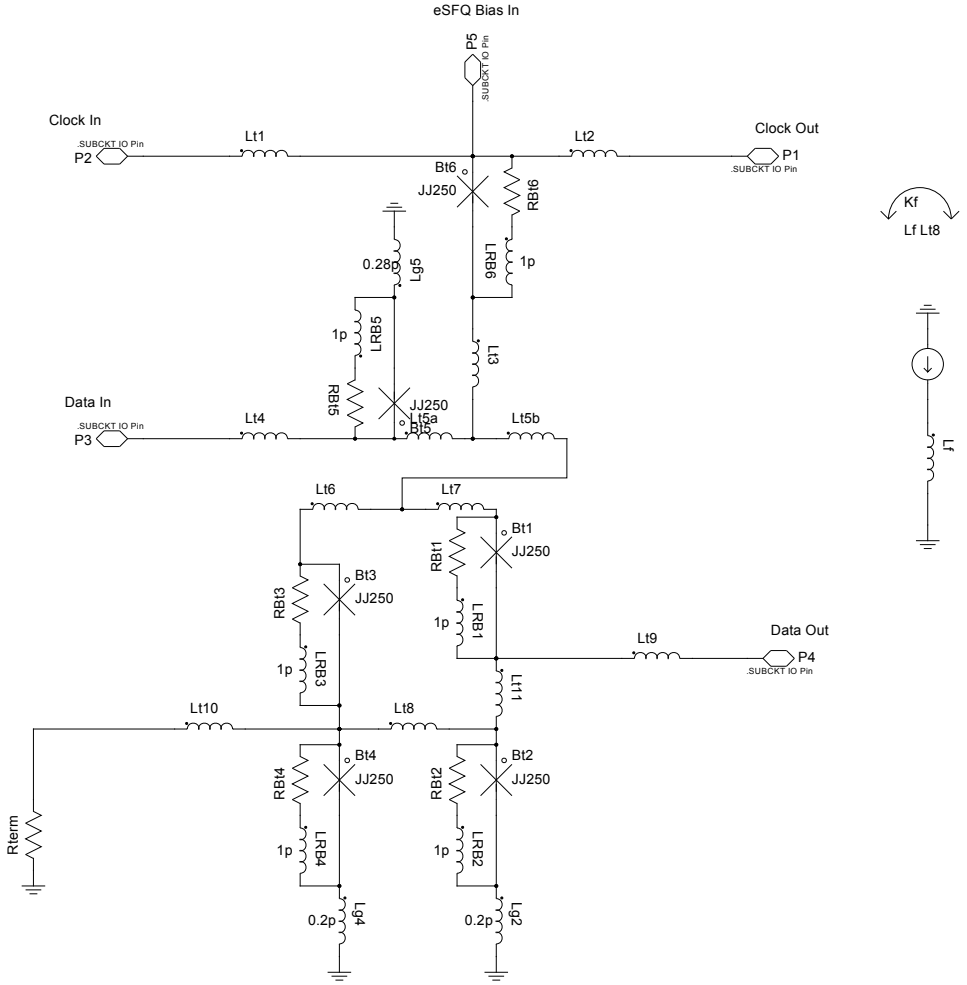

Figure 6.43: Schematic of TFF core of designed `eTFF` implementation. Details can be found in Section A.5.4.

A simulation of the designed eTFF sandwiched within an 8-bit shift register is recorded in Figure 6.44. Clearly, the incident data pulses are divided, with every second pulse transmitted by the eTFF.

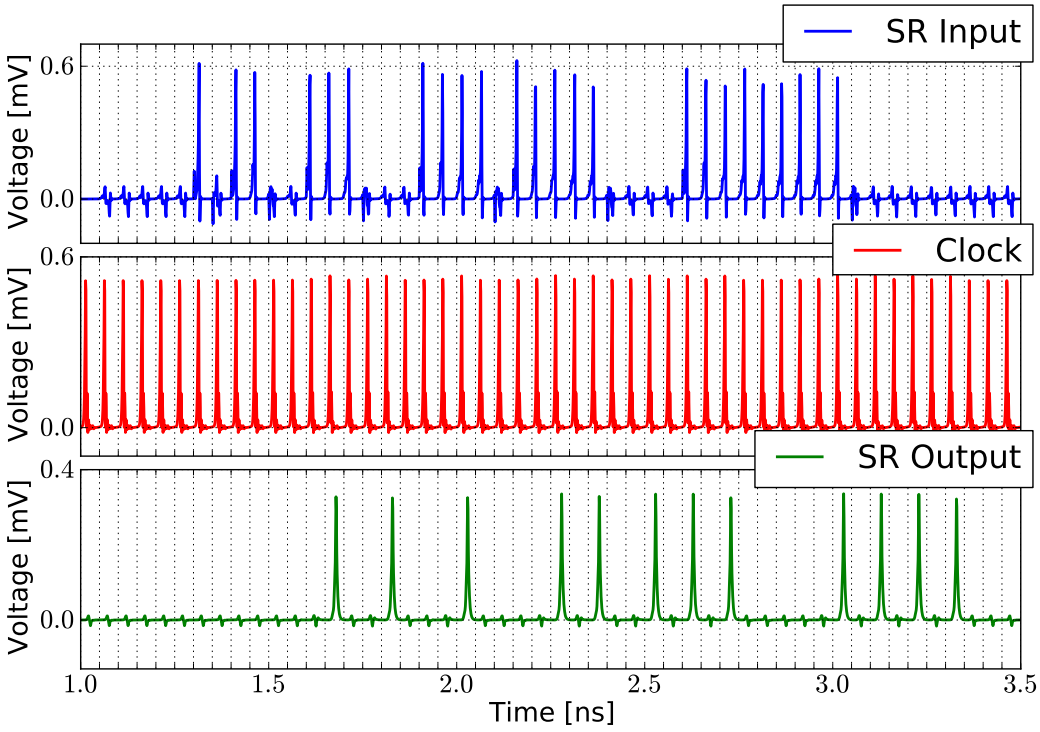

Figure 6.44: A 20 GHz simulation of the designed eTFF cell within an 8-bit shift register.

Employing the correct output from the simulation in Figure 6.44 as a verification condition, parameter margins for the eTFF cell are determined. The critical parameter margin is junction J_{t6} (J_5 from Figure 6.42), at $[-27\%, 20\%]$. This seems entirely feasible for fabrication.

6.7.4 Layout of eTFF test

The eTFF cell described in Figure 6.43 was laid out and, at the cut-off time for this text, was in fabrication. The layout of the cell is depicted in Figure 6.45.

The test layout, depicted in Figure 6.46, substantially mimics the circuit that produced the traces in Figure 6.44. An 8-bit shift-register flanks the eTFF cell under test. A comparator at the head of the shift register [138] is employed to inject data into the shift register, whereas the clock is applied from the tail of the shift register (realising the counterflow clocking scheme common in shift registers).

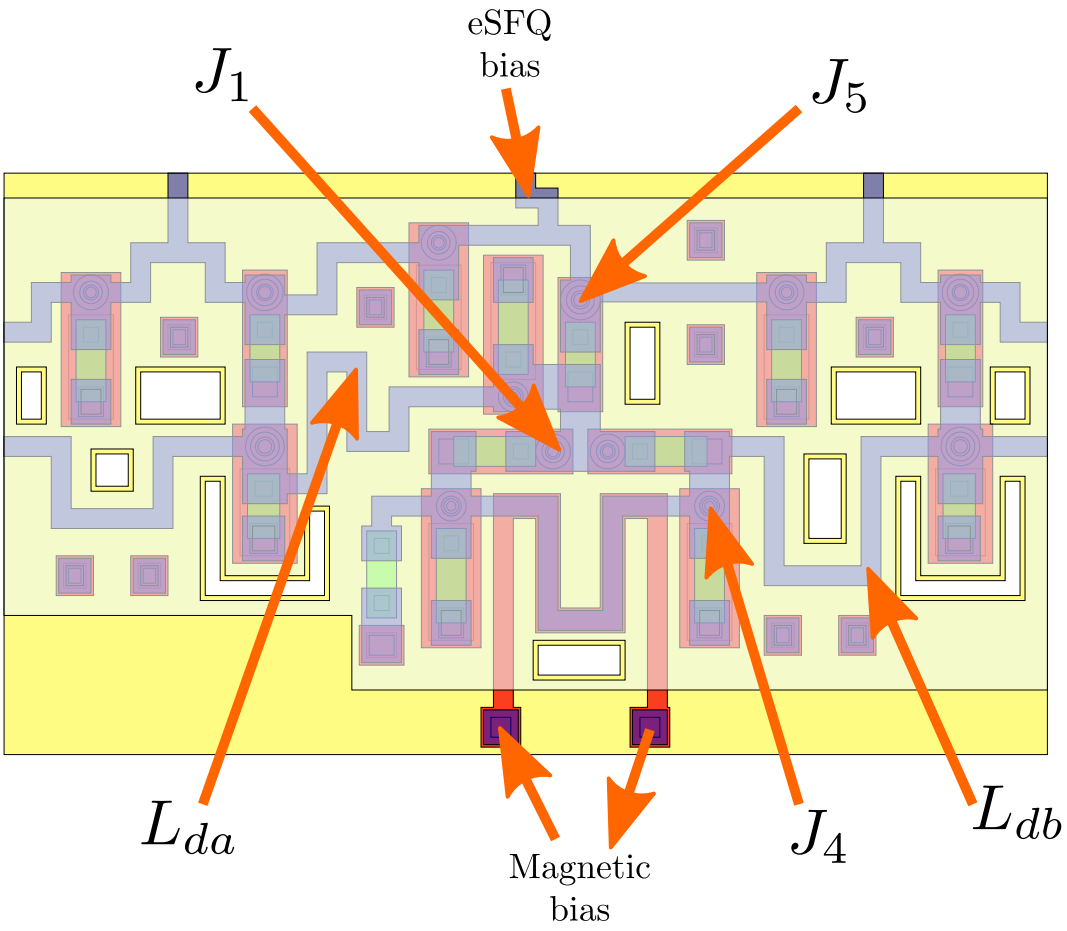

Figure 6.45: Realised physical implementation of eTFF cell in Hypres's 4.5 kA cm^{-1} process. Details, including extracted paramters, may be found in Section A.5.4.

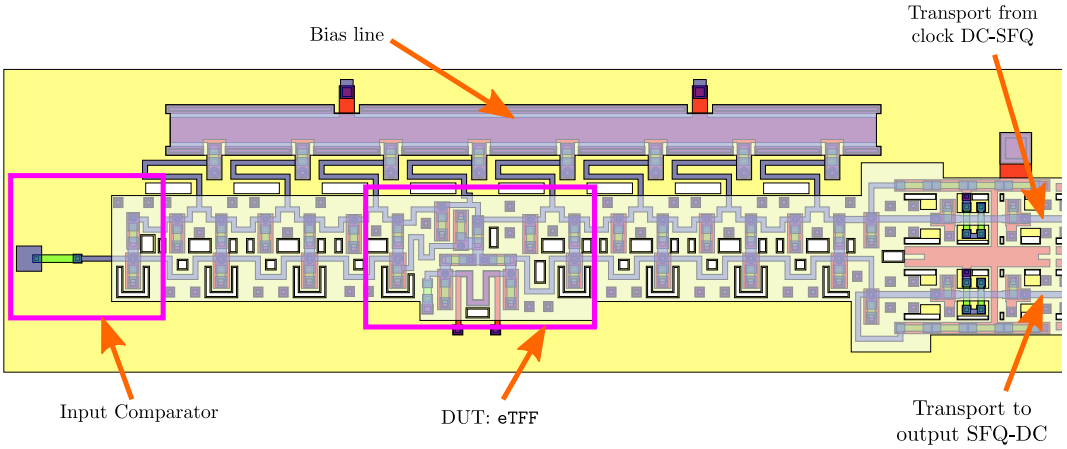

Figure 6.46: Realised physical implementation of eTFF test in Hypres's 4.5 kA cm^{-1} process.

Functional tests, for example using Octopux [152], can be performed by applying random data and clock patterns and evaluating the output. In this way, bias margins can be determined. This is the first step towards verification of the `eTFF` concept. At the time of writing, the test circuits are in fabrication.

6.8 Chapter summary and conclusion

First gates for a new ultra-low-power logic family were designed and demonstrated experimentally in this chapter. Measured margins were decent and high-speed operation was demonstrated as conclusively as the fabrication/test setup allowed. A 1-to-16 deserialiser appropriate for readout of an ADC for the SKA has been experimentally demonstrated. We summarise work reported and insights gleaned in this chapter as follows:

1. Although conventional RSFQ power dissipation is 3-4 orders of magnitude below that of conventional room temperature technologies, it is still high compared to the switching energy of the Josephson junction.

2. By eliminating the resistor bias network, static power dissipation falls to zero. The eSFQ logic family is a most elegant method for achieving this, avoiding the pitfalls and inconveniences of other methods examined here at the expense of redesigning certain logic gates.

3. Gates in eSFQ logic require the bias current introduced through the decision-making pair, which decreases parameter margins. We developed two solutions to this: overdamping the upper DMP junction or introducing a corrective flux bias. The former has the disadvantage of decreased operating frequency. The latter has the (in our opinion preferrable) disadvantage of mildly increased logistical requirements for operation. Acceptable margins can be achieved with both methods, at least for our demonstrators.

4. Different bias line designs have not resulted in a marked difference in bias current distribution. Experimental results lead to similar conclusions, hence we conclude that the bias line is not subject to non-trivial design constraints. This is a desirable result, however bias current distribution can be improved by more careful design of the bias line at the expense of robustness against fabrication tolerances.

5. Introducing a magnetic flux bias has shown a decrease in dynamic power dissipation in simulations. This is promising, but we have yet to back this up conclusively with experimental results. Our experiments show that the dynamic power dissipation of structures that employ a corrective magnetic flux bias is at least no worse than those that do not.

6. Deserialisers based on eSFQ principles have been designed and verified experimentally. They have been shown to work in an ADC readout configuration. As they will form a substantial component of an eventual ADC, the power savings compared to RSFQ are significant.

7. The eSFQ concept can be extended to nominally unclocked gates by redesigning them in a synchronised fashion. We have demonstrated the validity of this concept by designing an eSFQ TFF cell.

8. Even when including the cryocooling penalty, eSFQ can still remain 2-3 orders of magnitude below conventional (semiconducting) technologies in terms of power dissipation. This is a very bold claim and makes assumptions that will require validation as practical renditions of the technology are developed.

Conclusion: eSFQ is a viable logic family that achieves ultra-low-power dissipation in an elegant, practical manner. This enables ADC readout design without placing strain on the power budget in the cryocooler. Furthermore, even when taking the cryocooling penalty into account, it paves the way for future gate development in an effort to build low-power processing hardware useful for the SKA, achieving significant power savings. Ultra-low-power superconducting high-performance signal processing is a very fertile area of research that has generated much recent international attention [22, 116].

6.9 Chapter research output

1. An article is about to be published on IEEE Xplore (not peer-reviewed) in conjunction with a poster that was exhibited at the International Superconductor Electronics Conference '13 in Cambridge, MA, USA, on 8 July 2013.

   M.H. Volkmann, A. Sahu, A.V. Dotsenko and O.A. Mukhanov, "Operation of Practical eSFQ Circuits," *Proceedings of the Internation Superconductor Electronics Conference '13*, to be available on IEEE Xplore soon.

   Abstract:
   We have designed and tested energy-efficient single flux quantum (eSFQ) circuits for use in practical circuits. The first circuit is a 184-bit shift register for acquisition memory to store short transient events. The second circuit is a 16-bit deserializer integrated with an analog-to-digital converter modulator. The deserializer converts a high speed oversampling delta modulator 1-bit data stream into 16 parallel digital outputs

at 1/16 of input data rate. In this work, we report preliminary functional and high-speed operation of some of these eSFQ circuits including their margins of operation at different clock speeds, maximum operation frequencies, and power dissipation.

2. An article was published in the international journal *IEEE Transactions on Applied Superconductivity* in conjunction with an oral presentation at the Applied Superconductivity Conference '12 in Portland, OR, USA in October 2012.

   M.H. Volkmann, C.J. Fourie, A. Sahu, O.A. Mukhanov, "Experimental investigation of energy-efficient digital circuits based on eSFQ logic," *IEEE Transactions on Applied Superconductivity*, vol. 23, no. 3, 1301505, June 2013.

Abstract:
We present the design and experimental results of ultra-low-power digital circuits based on the recently introduced energy-efficient eSFQ logic. Similar to another low-power SFQ logic, ERSFQ, the eSFQ circuits make use of superconducting DC bias current dividers and thus avoid static power dissipation. As a result, per-gate power dissipation is reduced by two orders of magnitude and static power dissipation is zero. While ERSFQ circuits have already been demonstrated, here we report the design and implementation of eSFQ circuits. We design, simulate, fabricate and test a set of eSFQ circuits including shift registers and demultiplexers (deserialisers). The eSFQ circuits are fabricated using the HYPRES 4.5 kA/cm2 standard process. We integrate a low-pass analog-to-digital modulator with our eSFQ deserialiser and output voltage drivers to form an ultra-low-power analog-to-digital converter testable at high speed. In this paper, we demonstrate eSFQ design methods, address the differences from standard RSFQ design and demonstrate the viability and performance metrics of eSFQ circuits through simulations as well as functional and high-speed tests. Specifically, we present measured parameter margins and power vs. clock frequency dependencies and bit-error rates. Applications and scaling of the energy-efficient eSFQ logic for digital and mixed-signal circuits are discussed.
Acknowledgements:
This work is supported in part by US DoD contract W911NF-09-C-0036.

This work details efforts beyond the present author's first internship at Hypres, including work performed during his second internship in 2012.

The focus is on high-frequency testing of the developed circuits, as well as their integration with more complex devices. The conference oral presentation is scheduled for Tuesday 9 October 2012 at the Applied Superconductivity Conference in Portland, OR, USA. Online publication of the article is expected later in the year (with print publication to follow). The authors and their roles are the same as for Item 3, although the present author performed the high-frequency tests.

3. An article has been published in the prestigious international journal *Superconductor Science and Technology*:

M.H. Volkmann, C.J. Fourie, A. Sahu, O.A. Mukhanov, "Implementation of energy efficient single flux quantum (eSFQ) digital circuits with sub-aJ/bit operation," *Superconductor Science and Technology*, vol. 26, 015016, January 2013.

Abstract:

We report the first experimental demonstration of recently proposed energy-efficient single flux quantum logic, eSFQ. This logic represents the next generation of RSFQ logic eliminating dominant static power dissipation associated with a dc bias current distribution and providing over two orders of magnitude efficiency improvement over conventional RSFQ logic. We further demonstrate that the introduction of passive phase shifters allows the reduction of dynamic power dissipation by about 20%, reaching ~ 0.8 aJ per bit operation. Two types of demonstration eSFQ circuits, shift registers and demultiplexers (deserialisers), were implemented using the standard HYPRES 4.5 kA/cm^2 fabrication process. In this paper, we present eSFQ circuit design and demonstrate the viability and performance metrics of eSFQ circuits through simulations and experimental testing.

Acknowledgements:

The work is supported in part by US DoD contract W911NF-09-C-0036. The authors wish to thank S. Kaplan for many discussions on RSFQ and eSFQ operation and NioCAD Pty (Ltd.) for circuit design software and support, A. Kirichenko and A. Inamdar for the contribution of peripheral cells for the IC layouts, as well as I. Vernik for help and verification of circuit layouts. Special thanks go to the HYPRES fabrication team of D. Yohannes, J. Vivalda, R. Hunt, D. Donnelly, D. Amparo, who delivered the integrated circuits on time and within specifications. We also thank M. Manheimer and S. Holmes for attention to this work.

This article details much of the present author's work from his 2011 internship at Hypres, in Elmsford, NY, USA. Tests conducted beyond that internship are also included. The authors are the present author and his supervisor (Prof. Fourie), as well as staff at Hypres. Dr Mukhanov was instrumental in conceptualising the idea and providing guidance. Prof. Fourie, besides his supervisory role, assisted with verification of circuits. Mr Sahu performed the functional liquid-Helium tests.

Chapter 7

Conclusions

7.1 Summary of results

In previous chapters, we examined superconducting Flash ADCs based on the QOS comparator in detail. Furthermore, we developed first demonstrators of the eSFQ logic family and experimentally verified their operation. We summarise what we learned next.

7.1.1 ADCs for the SKA - Chapter 2

The high clock speeds of superconducting electronics allow us to digitise the entire band of interest of the SKA, $B \sim 10\,\mathrm{GHz}$, simultaneously when employing a superconducting flash ADC clocked at $f_{\mathrm{clk}} \sim 20\,\mathrm{GHz}$.

As great sensitivity is one of the foremost drivers of the SKA, and the SKA will be used to observe wideband events, the high instantaneous bandwidth achievable by the superconducting flash ADC is of great value.

Whereas flash ADCs have high bandwidth, literature shows that the achievable dynamic range of flash ADCs lags behind those of oversampling ADCs such as the sigma-delta ADC. As semiconductor sigma-delta ADCs are well-established and capable, and as superconducting sigma-delta ADCs cannot be tuned, we believe that a superconducting flash ADC would be of greater benefit to the SKA. We expect that the comparatively low dynamic range can be worked around by other methods (whereas no other methods exist, to our knowledge, to achieve an instantaneous bandwidth covering the band of interest of the SKA).

For this reason, we chose the superconducting flash ADC as potential frontend for the SKA.

7.1.2 ADC model - Chapter 4

Key to the superconducting Flash ADC is the multi-threshold comparator (MTC). Key to the MTC is the quasi-one-junction SQUID (QOS), the response

of which we analysed analytically in Chapter 4. We learned that the ideal response mainly depends on the circuit parameters I_{cq}, the critical current of the quantiser junction, and L_q, the quantiser inductance. The quantiser inductance L_q determines I_m, the quantiser period.

An n-bit flash ADC requires n comparators placed in parallel. A proportion of the input signal is applied to each comparator, and as n grows, so grows the ratio between the largest and smallest signal applied to a comparator.

We identified the decision-making pair (DMP) grey-zone, a result of thermal noise, which is significant even at 4.2 K, as a prime source of SNR-depression. The effect of the grey-zone becomes more pronounced as n increases. For example, for $n = 4$ and $I_{cq} = 125\,\mu\text{A}$, a DMP grey-zone of $I_g \sim 8\,\mu\text{A}$ depresses achievable SNR by 3 dB. For $n = 8$, however, the 3dB-depression point is reached at $I_g \sim 1\,\mu\text{A}$ already.

7.1.3 ADC circuits - Chapter 5

We attempted to design a superconducting MTC in the IPHT RSFQ1F process. We found that a lower bound is imposed on L_q by certain design rules (critical current of superconducting lines), whereas an upper bound is imposed on L_q by other design rules (minimum junction I_c and parasitic inductive load through minimum dimensions) and the Josephson inductance of the comparator. Fortunately we found a range of values for L_q that satisfies both constraints. We thus designed, laid out and extracted `MC1`, based on the QOS comparator.

Unfortunately the permissible grey-zone for QOS comparators is quite low. To overcome this limitation, we turned to the complementary QOS comparator. This circuit makes use of two (complementary) QOS quantisers, which are fed by a floating current source. Whereas this goes beyond our model of the QOS from Chapter 4, the knowledge gained about grey-zones and scaling remains applicable.

7.1.4 Experimental verification - Chapter 5

To investigate the effect of junction damping on the grey-zone of MTCs, we designed a low-speed cryogenic experiment, to be followed by a high-speed cryogenic experiment. We designed experiments that are analogous in SPICE and in the cryogenic environment to estimate the grey-zone associated with MTCs. Next we designed 5 MTCs (inclusive of the first MTC based on the QOS comparator). The SPICE results matched our expectations, but the experimental results of the CQOS comparators were not as good as expected. We have explained this sufficiently through a design-rule miscommunication and other experimental factors not modelled.

Nonetheless, the experimentally observed quantities show that even our first batch of comparators exceeds the requirements for a 4-bit superconducting

flash ADC, and further simulations show that employing a (readily available) $J_c = 4.5\,\mathrm{kA/cm^2}$ process would easily achieve the clock frequency constraints. We expect substantial improvement from our second experiment, those structure are currently in fabrication.

7.1.5 Readout - Chapter 6

Due to the data rates produced by superconducting circuits, special measures must be taken to make the data accessible to conventional technology (semiconductors). As raw single-bit data-links with conventional technologies run at several Gbps, we must reduce the data rate of the flash ADC output, which comes to at least 20 Gbps per bit (f_{clk}). Furthermore, currently available high-speed superconducting output drivers have a maximum frequency of ~ 2 Gbps, further underlining the need for a reduced data rate.

We opted for the deserialiser approach as a data-rate reduction measure, as it is straightforward and does not result in any data loss. We turned to the newly invented eSFQ logic family to design our readout circuitry. Gates based on eSFQ logic are ultra-low-power, without many of the disadvantages of other ultra-low-power SFQ logic families.

We successfully demonstrated working shift-registers and deserialisers implemented in eSFQ logic, the first demonstrated eSFQ logic circuits of any kind. These circuits were fabricated with Hypres's 4.5 $\mathrm{kA/cm^2}$ process and experimentally verified at low-speed and, to a large degree, at high-speed. Unfortunately fabrication problems again delayed some test structures beyond the cutoff time for this work, but since the experiments have already been designed and the previous suites of chips performed as expected, we expect that the last batch of tests will only confirm our expectations.

The energy and power requirements our ultra-low-power circuits is remarkable: sub-aJ/bit per operation was demonstrated.

Recent research also conducted within the scope of this work has shown that the concept of eSFQ logic is extendable to arbitrary logic gates (even unclocked gates), with experiments designed and currently in fabrication.

7.1.6 Software - Chapter 3

During our efforts, it became clear that the state of software available for superconducting design is problematic. Much software in use today is custom-written by students for a specific purpose and not well maintained. Other software is commercially available and of high quality, but not designed specifically for superconducting circuits. The price of such software products is usually high, which makes it difficult to justify their purchase for an unintended purpose. The only software product that is commercially developed specifically for superconducting circuit design is in its early stages and not yet suitable for serious designs.

For these reasons, we wrote our own software from scratch to facilitate design and analysis of our circuits. We developed tools that facilitate parametric modelling as well as high-performing, parallelised analysis tools (used extensively in work presented in Chapters 5 and 6). Furthermore, we have parallelised FastHenry, a popular backend for inductance extraction, an important link in the design toolchain.

7.2 Conclusions drawn

7.2.1 Advantages of superconductors in the SKA

We admit that a rigorous investigation of the advantages of superconductors for the SKA, especially as a comparison to likely conventional technologies, is outside the scope of this work. We are not qualified to make such predictions and believe that input from authorities in the field of radio astronomy (which we are not) is required to fully illuminate all facets of this discussion. We discuss the advantages of a superconducting frontend only on the sensitivity of the SKA.

What is clear is that superconducting flash ADCs can digitise the entire band of interest of the SKA instantaneously. We are not aware of a means of achieving this with conventional technologies at the feed, and we do not expect such a means to be available by 2024.

The downside of superconducting flash ADCs is that their dynamic range is limited. There are many reasons for this, the one we focused on in this work is the grey-zone of the multi-threshold comparators. In Chapter 5 we concluded that a minimal-grey-zone comparator can be fabricated with little SNR depression for a 4-bit ADC. In fact, more bits should be possible too.

In Chapter 2 we claimed that the improved bandwidth of flash ADCs should outrank their lower dynamic range. We revisit this issue now from the perspective of sensitivity, an important figure of merit of radio telescopes and an area in which the SKA is expected to outclass all existing radio telescopes.

The sensitivity of an interferometer for continuum observations σ is given by [157]

$$\sigma = \frac{\sqrt{2}k_B T_{\text{sys}}}{A_{\text{eff}}\eta_c\sqrt{N(N-1)B\tau_{\text{int}}}}, \tag{7.1}$$

where η_c is the efficiency of the correlator, N is the number of dishes, B is the observed bandwidth (instantaneous bandwidth, as we are discussing continuum observations), τ_{int} is the integration time devoted to the observation, T_{sys} is the system noise temperature, A_{eff} is the effective aperture area and k_B is Boltzmann's constant. Most of these parameters do not feature in our discussion, hence we do not explain them here in detail, but refer to Thompson et al instead [157]. Note that a lower sensitivity σ is better.

Since the sensitivity of a radio-telescope is proportional to $1/\sqrt{B}$, a superconducting flash ADC is almost certainly of great advantage for any survey, especially those where wideband events are observed (such as for Cradle of Life experiments). If we assume that we can achieve an instantaneous bandwidth of $B = 10\,\text{GHz}$ (which seems likely), and assume further that conventional technologies will not achieve instantaneous bandwidths greater than $B = 2.5\,\text{GHz}$ (which would require sampling at 5 GHz or greater, a tall order at cryogenic temperatures), then, all other parameters remaining equal, a superconducting flash ADC can improve the sensitivity by factor 2!

Of course, the low dynamic range of flash ADCs may also play a role. We know that the dynamic range (number of bits) of the digitised data determines the correlator efficiency η_c. Typical values are $\eta_c \approx 0.873$ for a 2-bit correlator, $\eta_c \approx 0.962$ for a 3-bit correlator and $\eta_c \approx 0.983$ for a 4-bit correlator [157]. We propose to achieve at least 4 bits, and seem to be able to achieve this. The grey-zone and other non-idealities may depress this figure, but we do not expect it to depress it beyond 6 dB. Hence, effectively at least 3 bits should be achieved.

Consequently, we assume at least $\eta_c = 0.962$, which decreases our achievable efficiency by a maximum of approximately 4%. Of course, if we compare this to our factor 2 sensitivity improvement due to greater bandwidth, we still have a 96% residual improvement. It seems the prediction we made in Chapter 2 has paid off in this respect.

Naturally, a conventional correlator will not be 100% efficient, exhibiting $\eta_c < 1$, making this debate largely academic. It is clear that, on the scales which we propose, an increase in bandwidth has a far greater positive effect on the sensitivity than a decrease in dynamic range (if applicable) has a negative effect.

7.2.2 Applicability of superconductors in the SKA

The circuits we described, designed, fabricated and measured in this work are manufactured as integrated circuits, with niobium the superconductor of choice. Niobium has a critical temperature of $T_c = 9.2\,\text{K}$, but the integrated circuits are usually designed to work at 4.2 K, the boiling point of helium.

The primary choice of niobium is usually justified by its useful material properties. A pure metal, it is easy to work with, deposit and etch, while remaining sturdy enough to keep its shape essentially indefinitely. Perhaps more importantly, it is homogeneous and isotropic in many respects. The same cannot be said for known high-temperature superconductors such as YBCO. The material properties of YBCO make it difficult to manufacture reproducible Josephson junctions, for example.

It is quite probable that the feeds of the SKA receivers will be cryocooled, as are those of KAT-7 and will be those of MeerKAT. The exact specifications of the cryogenic environment still need to be finalised, but a temperature below

20 K seems likely. Whether a temperature as low as 4.2 K will be desired or would be practical remains an open question. Many other superconductors exist below 20 K. niobium-nitride (NbN), for example, has $T_c \approx 16$ K. In fact, NbN processes exist and were once more popular than pure-niobium processes. Thus, NbN would be a real possibility if temperatures will not be low enough for pure niobium. A further recent contender is magnesium-diboride, MgB_2, which has $T_c \approx 39$ K and has enjoyed recent attention in superconductor electronics [158].

It is known that the technology required for the SKA does not exist today. Superconducting circuits do exist, and have proved valuable, but need some development to be viable as a product of which high uptime is expected. Issues that need to be addressed for the SKA's purposes are flux trapping (shielding), accurate threshold biasing, low-jitter clocking and, possibly, integration. Since technology development needs to occur in any case, and superconductors promise a greater reward than conventional technologies, we believe that a superconducting frontend is a suitable candidate for the SKA.

7.2.3 Recent competition

Recent progress in semiconductor electronics has facilitated the emergence of certain semiconductor ADCs that boast benchmarks comparable (in some respects) to the architecture investigated in this dissertation. Perhaps the most pertinent example is a 26 GS/s ADC with 2.9 effective bits produced by the Hittite Microwave Corporation [159].

While this is very impressive, there is one important aspect that rules out direct comparison of this ADC and our architecture. The reported power requirement of this ADC is greater than 4 W [159], exceeding the cryocooler heat lift of small-scale cryocoolers employed in superconducting electronics. Hence, to fully make use of the high sample rate, such a part would have to be placed outside the cryocooled receiver assembly, requiring also a very wideband signal line to penetrate said assembly (placing a significant additional heat load onto the cryocooler).

7.3 How to proceed

While the work we presented in this text is promising, it is only the beginning. Much more needs to be done before a superconducting frontend is realised. If a committed effort is made, we strongly believe that such a realisation and its deployment are feasible.

7.3.1 Research and collaboration

This work has already produced significant research output, underlining its calibre and actuality. Two international journal papers have been published, with an additional journal paper to be submitted soon. Three international conference posters and one oral presentation were reported in tandem, with more to be reported this year. International collaborators are co-authors on the majority of the papers, improving South Africa's image abroad. Further work on which the present author is co-author has been published in and will be submitted to peer-reviewed journals.

The applied superconductivity base in South Africa is small, essentially confined to the University of Stellenbosch for our purposes. Fabrication and experimental verification of designed circuits must currently be done overseas. The present author has completed two internships at the world's leading superconductor electronics design company in the USA (Hypres), where he has left a lasting impression and built strong foundations for future collaboration. He has also completed a short research visit to the IPHT in Jena, Germany, European stronghold of applied superconductivity, as well as an internship at the CiS Research Institute in Erfurt, Germany.

These and other international partners are an excellent resource on which to rely for further research along the lines of this work.

7.3.2 Experimental confirmation

Experimental confirmation of all of the measured test structures designed as part of this work was successful. Fabrication errors and delays have hampered experimental confirmation of some other test structures. This is unfortunate, but can be remedied.

Further confirmation of existing (designed) test structures should be completed before proceeding to the next design steps. The present author expects to complete further testing of structures (described in Sections 5.5 and 6.7) in late 2013 or early 2014.

7.3.3 Further analysis

In parallel, further analyses of the advantages and disadvantages of a superconducting frontend are a must. We have demonstrated in this work that a superconducting frontend for the SKA can be a reality and that it is likely to dramatically improve important figures of merit for the SKA.

Next, while waiting for further experimental results, we propose that a rigorous exploration of the advantages and disadvantages for the SKA of a flash ADC-based frontend is conducted, with the results summarised and published.

From there, choices should be made for n, I_m and f_c (determined by required receiver dynamic range, ADC sensitivity and instantaneous bandwidth

respectively) and a fully-integrated ADC laid-out, fabricated and tested. This is a tall order, as the testing especially, at least at high clock frequencies, will require custom (room temperature) equipment to be manufactured to process the immense data volume.

At this point, the applicability of the superconducting frontend to the SKA should be formally evaluated and a decision made about its inclusion in the SKA project. We are confident that this decision can be a positive one.

7.4 Last word

We have conceived, designed, analysed and experimentally verified key components of a potential superconducting frontend for the Square Kilometre Array. Harnessing the electromagnetic properties of superconductors, we have devised a relatively simple circuit that should shorten the receiver chain, quadruple instantaneous bandwidth, retain acceptable signal-to-noise ratio and, consequently, at least for continuum sampling, double the sensitivity of the Square Kilometre Array.

To boot, while devising a convenient, no-loss readout circuit, an integral part of the front-end, we have developed first demonstrators of a convenient ultra-low-power logic family, which were experimentally verified at up to 24 GHz (at which point our test equipment was the limiting factor). This ensures that our frontend will not add to the power budget of the cryocooled feed, but also paves the way towards a potential revolution in exascale supercomputing and an SKA backend with a significantly reduced power budget.

Bibliography

[1] D. van Delft and P. Kes, "The discovery of superconductivity," *Physics Today*, pp. 38–43, September 2010.

[2] U. S. Department of Energy - Basic Energy Sciences (public domain), http://www.ccas-web.org/superconductivity/#image1, last accessed on 29 September 2012.

[3] M.-L. Doan, "Meissner effect: levitation of a magnet above a superconductor," http://en.wikipedia.org/wiki/File:Meissner_effect_p1390048.jpg, October 2007, last accessed on 28 September 2012. Licensed under the Creative Commons Attribution-Share Alike 3.0 Unported license available at http://creativecommons.org/licenses/by-sa/3.0/deed.en.

[4] Wikimedia Commons. File:Atmospheric_electromagnetic_opacity.svg. Last accessed 6 June 2012. [Online]. Available: http://commons.wikimedia.org/wiki/File:Atmospheric_electromagnetic_opacity.svg

[5] P. E. Dewdney, P. J. Hall, R. T. Shilizzi, and T. J. L. W. Lazio, "The Square Kilometre Array," *Proceedings of the IEEE*, vol. 97, no. 8, pp. 1482 – 1496, 2009.

[6] EMSS Antennas, 18 Techno Avenue, Technopark, Stellenbosch, 7600, South Africa.

[7] SKA South Africa, Rosebank, Johannesburg, South Africa.

[8] C. J. Fourie, "A 10 GHz oversampling Delta modulating analogue-to-digital converter implented with hybrid superconducting digital logic," Master's thesis, University of Stellenbosch, March 2001.

[9] K. K. Likharev and V. K. Semenov, "RSFQ logic/memory family: A new Josephson-junction technology for sub-Terahertz-clock-frequency digital systems," *IEEE Transactions on Applied Superconductivity*, vol. 1, no. 1, pp. 3 – 28, Mar. 1991.

[10] M. Maruyama, H. Suzuki, Y. Hashimoto, and M. Hidaka, "A flash A/D converter using complementarily-combined SQUIDs," *IEEE Transactions on Applied Superconductivity*, vol. 19, no. 3, pp. 680 – 684, 2009.

[11] H. Suzuki, M. Oikawa, K. Nishii, K. Ishihara, K. Fukiwara, M. Maruyama, and M. Hidaka, "Design and demonstration of a 5-bit flash-type SFQ A/D converter integrated with error correction and interleaving circuits," *IEEE Transactions on Applied Superconductivity*, vol. 21, no. 3, p. 671, June 2011.

[12] Hypres Inc., *Niobium Integrated Circuit Fabrication Design Rules*, 175 Clearbrook Road, Elmsford, NY, Jan 2008.

[13] Department of Quantum Detection, IPHT Jena, *Rapid Single Flux Quantum (RSFQ) - Design Rules for Nb/Al2O3-Al/Nb-Process at IPHT Jena*, Jun. 2007.

[14] N. Pandor, "Statement by the South African Minister of Science and Technology: Square kilometre array organisation site decision," http://www.ska.ac.za/releases/20120525.php, May 2012, last accessed 28 August 2013.

[15] R. de Bruyn Ouboter, "Heike Kamerlingh Onnes's discovery of superconductivity," *Scientific American*, pp. 98 – 103, March 1997.

[16] U. S. Department of Energy - Office of Electricity Delivery and Energy Reliability, "Long Island HTS power cable project fact sheet," http://energy.gov/sites/prod/files/oeprod/DocumentsandMedia/LIPA_5_16_08.pdf, May 2008, last accessed on 28 September 2012.

[17] J. Clarke and A. I. Braginski, Eds., *The SQUID Handbook: Applications of SQUIDs and SQUID Systems*. Wiley-VCH, October 2006, vol. 2.

[18] W. Perold, "Complementary output switching logic: a superconducting voltage-state logic family operating at microwave frequencies," in *Communications and Signal Processing, 1998. COMSIG '98. Proceedings of the 1998 South African Symposium on*, September 1998, pp. 435 –440.

[19] W. Chen, A. V. Rylyakov, V. Patel, J. E. Lukens, and K. K. Likharev, "Rapid Single Flux Quantum T-flip flop operating up to 770 GHz," *IEEE Transactions on Applied Superconductivity*, vol. 9, no. 2, pp. 3212 –3215, 1999.

[20] D. Gupta, D. E. Kirichenko, V. V. Dotsenko, R. Miller, S. Sarwana, A. Talalaevskii, J. Delmas, R. J. Webber, S. Govorkov, A. F. Kirichenko, I. V. Vernik, and J. Tang, "Modular, multi-function digital-rf receiver systems," *IEEE Transactions on Applied Superconductivity*, vol. 21, no. 3, pp. 883–890, 2011.

[21] R. F. Service, "What it'll take to go exascale," *Science*, vol. 335, no. 6067, pp. 394–396, 2012.

[22] M. A. Manheimer and D. S. Holmes, "Energy-efficient superconducting supercomputing - power budgets and requirements," *Applied Superconductivity Conference (ASC'12)*, 2012, 2EE-01.

[23] F. Ortmann and R. Gerber, "SuperGreen computing," *Fluxonics Newsletter*, vol. 1, pp. 6–7, February 2012.

[24] H. Nyquist, "Certain topics in telegraph transmission theory," *Transactions of the A.I.E.E.*, pp. 617–644, February 1928.

[25] O. A. Mukhanov, "Energy-efficient single flux quantum technology," *IEEE Transactions on Applied Superconductivity*, vol. 21, no. 3, pp. 760 – 769, June 2011.

[26] ——, *In: 100 Years of Superconductivity*. Boca Raton: Taylor & Francis, 2011, ch. 7 - History of Superconductor Analog-to-Digital Converters, pp. 440–458.

[27] Last accessed on 28 September 2012. [Online]. Available: http://www.skatelescope.org/

[28] S. S. W. Group. (2011) The square kilometre array design reference mission. Last accessed 18 September 2012. [Online]. Available: http://www.skatelescope.org/uploaded/3517_DRM_v1.0.pdf

[29] "The technology — official SKA website," http://www.skatelescope.org/the-technology/ [Last accessed 18 September 2012].

[30] J. Jonas, "MeerKAT the south african array with composite dishes and wide-band single pixel feeds," *Proceedings of the IEEE*, vol. 97, no. 8, pp. 1522 –1530, aug. 2009.

[31] A. Chippendale and A. Schinckel, "ASKAP: Progress towards 36 parabolic reflectors with phased array feeds," in *Proc. XXXth URSI General Assembly and Scientific Symp*, 2011, pp. 1–4.

[32] J. Hessels and B. Stappers, "LOFAR: A powerful and flexible observatory for pulsars and fast transients," in *General Assembly and Scientific Symposium, 2011 XXXth URSI*, aug. 2011, pp. 1 –4.

[33] Last accessed 28 September 2012. [Online]. Available: http://public.ska.ac.za/kat-7

[34] G. Harp, R. Ackermann, Z. Nadler, S. Blair, M. Davis, M. Wright, J. Forster, D. DeBoer, W. Welch, S. Atkinson, D. Backer, P. Backus, W. Barott, A. Bauermeister, L. Blitz, D.-J. Bock, G. Bower, T. Bradford, C. Cheng, S. Croft, M. Dexter, J. Dreher, G. Engargiola, E. Fields, C. Heiles, T. Helfer, J. Jordan, S. Jorgensen, T. Kilsdonk, C. Gutierrez-Kraybill, G. Keating, C. Law, J. Lugten, D. MacMahon, P. McMahon, O. Milgrome, A. Siemion, K. Smolek, D. Thornton, T. Pierson, K. Randall, J. Ross, S. Shostak, J. Tarter, L. Urry, D. Werthimer, P. Williams, and D. Whysong, "Primary beam and dish surface characterization at the allen telescope array by radio holography," *IEEE Transactions on Antennas and Propagation*, vol. 59, no. 6, pp. 2004 –2021, June 2011.

[35] K. Murayama, Y. Maeda, Y. Hiratsuka, F. Imai, and K. Kang, "A compact HTS filter subsystem," *IEEE Transactions on Applied Superconductivity*, vol. 11, no. 1, p. 1400, March 2001.

[36] "The Large Hadron Collider," http://public.web.cern.ch/public/en/lhc/lhc-en.html, last accessed on 23 September 2012.

[37] D. Zhang, S. Dai, Z. Wang, J. Zhang, N. Song, X. Li, F. Zhang, X. Du, M. Qiu, Z. Zhang, W. Zhou, G. Zhang, L. Xiao, and L. Lin, "Research on stability of superconducting magnet for MRI," *IEEE Transactions on Applied Superconductivity*, vol. 21, no. 3, pp. 2100 –2103, June 2011.

[38] A. Karpov, D. Miller, F. Rice, J. A. Stern, B. Bumble, H. G. LeDuc, and J. Zmuidzinas, "Low noise 1.2 THz SIS mixer for Herschel radio observatory," *ISSTT Proceedings*, vol. 2, pp. 70–75, 2004.

[39] J. Mosher, P. Lewis, and R. Leahy, "Multiple dipole modeling and localization from spatio-temporal MEG data," *IEEE Transactions on Biomedical Engineering*, vol. 39, no. 6, pp. 541–557, june 1992.

[40] H. Rogalla and P. H. Kes, Eds., *100 Years of Superconductivity*, 1st ed. Taylor & Francis, Dec. 2011. [Online]. Available: http://www.worldcat.org/isbn/1439849463

[41] T. P. Orlando and K. A. Delin, *Foundations of Applied Superconductivity*. Addison-Wesley Publishing Company, 1991.

[42] A. M. Kadin, *Introduction to Superconducting Circuits*, 1st ed. Wiley-Interscience, 1999.

[43] J. Satchell, "Limitations on hts single flux quantum, logic," *IEEE Transactions on Applied Superconductivity*, vol. 9, no. 2, pp. 3841–3844, 1999.

[44] W. Meissner and R. Ochsenfeld, "Ein neuer Effekt bei Eintritt der Supraleitfaehigkeit," *Naturwissenschaften*, vol. 21, no. 44, pp. 787 – 788, 1933.

[45] M. Rudolph, J. Nagel, M. Meckbach, M. Kemmler, M. Siegel, K. Ilin, D. Koelle, and R. Kleiner, "Direct current superconducting quantum interferometers with asymmetric shunt resistors," *Applied Physics Letters*, vol. 101, no. 5, p. 052602, 2012.

[46] B. Josephson, "Possible new effects in superconductive tunnelling," *Physics Letters*, vol. 1, no. 7, pp. 251 – 253, 1962. [Online]. Available: http://www.sciencedirect.com/science/article/pii/0031916362913690

[47] B. D. Josephson, "The discovery of tunnelling supercurrents," *Rev. Mod. Phys.*, vol. 46, no. 2, pp. 251 – 254, April 1974.

[48] J. Kunert, O. Brandel, S. Linzen, O. Wetzstein, H. Toepfer, T. Ortlepp, and H.-G. Meyer, "Recent development in superconductor digital electronics technology at fluxonics foundry," *IEEE Transactions on Applied Superconductivity*, vol. 23, no. 4, p. 1101707, August 2013.

[49] D. Yohannes, S. Sarwana, S. K. Tolpygo, A. Sahu, Y. A. Polyakov, and V. K. Semenov, "Characterization of HYPRES's 4.5 kA/cm2 & 8 kA/cm2 Nb/AlOx/Nb fabrication processes," *IEEE Transactions on Applied Superconductivity*, vol. 15, no. 2, pp. 90–93, June 2005.

[50] J. G. Proakis and D. G. Manolakis, *Digital Signal Processing.* Pearson Education, Inc., 2007.

[51] C. E. Shannon, "Communication in the presence of noise," *Proceedings of the Institute of Radio Engineers*, vol. 37, no. 1, pp. 10–21, January 1949.

[52] H. H. Zappe, "Ultrasensitive analog-to-digital converter using Josephson functions," *IBM Tech. Disc. Bull.*, vol. 17, p. 3053, March 1975.

[53] R. E. Harris, C. A. Hamilton, and F. L. Lloyd, "Multiple-quantum interference superconducting analog-to-digital converter," *Applied Physics Letters*, vol. 35, p. 720, 1979.

[54] C. Hamilton and F. Lloyd, "Design limitations for superconducting A/D converters," *IEEE Transactions on Magnetics*, vol. 17, no. 6, pp. 3414 – 3419, nov 1981.

[55] L. Abelson and G. L. Kerber, "Superconductor integrated circuit fabrication technology," *Proceedings of the IEEE*, vol. 92, pp. 1517–1533, 2004.

[56] D. K. Brock, A. M. Kadin, A. F. Kirichenko, O. A. Mukhanov, S. Sarwana, J. A. Vivalda, W. Chen, and J. E. Lukens, "Retargeting RSFQ cells to a submicron fabrication process," *IEEE Transactions on Applied Superconductivity*, vol. 11, no. 1, pp. 369–372, 2001.

[57] S. K. Tolpygo, D. Yohannes, R. T. Hunt, J. A. Vivalda, D. Donnelly, D. Amparo, and A. F. Kirichenko, "Process development for superconducting integrated circuits with 80GHz clock frequency," *IEEE Transactions on Applied Superconductivity*, vol. 17, no. 2, pp. 946 – 951, June 2007.

[58] O. A. Mukhanov, D. Gupta, A. M. Kadin, and V. K. Semenov, "Superconductor analog-to-digital converters," *Proceedings of the IEEE*, vol. 92, pp. 1564–1584, 2004.

[59] S. Sarwana, D. Gupta, A. F. Kirichenko, T. Oku, C. Otani, H. Satu, and H. M. Shimizu, "High-sensitivity high-resolution dual-function signal and time digitzer," *Applied Physics Letters*, vol. 80, pp. 2023–2025, 2002.

[60] S. B. Kaplan, P. D. Bradley, D. K. Brock, D. Gaidarenko, and D. Gupta, "A superconductive flash digitizer with on-chip memory," *IEEE Transactions on Applied Superconductivity*, vol. 9, no. 2, pp. 3020–3025, 1999.

[61] H. Suzuki, M. Oikawa, K. Nishii, and M. Hidaka, "Investigation of a 5-bit flash-type SFQ A/D converter using 10 kA/cm2 niobium process and locally isolated grounds," *IEEE Transactions on Applied Superconductivity*, vol. 23, p. 1400805, 2013.

[62] J. A. Cherry and W. M. Snelgrove, *Continuous-Time Delta-Sigma Modulators for High-Speed A/D Conversion*. Kluwer Academic Publishers, 2000.

[63] J. Przybysz, D. Miller, E. Naviasky, and J. Kang, "Josephson sigma-delta modulator for high dynamic range A/D conversion," *IEEE Transactions on Applied Superconductivity*, vol. 3, no. 1, pp. 2732 –2735, mar 1993.

[64] J. F. Bulzacchelli, H.-S. Lee, J. A. Misewich, and M. B. Ketchen, "Superconducting bandpass $\Delta\Sigma$ modulator with 2.23-GHz center frequency and 42.6-GHz sampling rate," *IEEE Journal of Solid-State Circuits*, vol. 37, no. 12, December 2002.

[65] Q. P. Herr, D. L. Miller, A. A. Pesetski, and J. X. Przybysz, "A quantum-accurate two-loop data converter," *IEEE Transactions on Applied Superconductivity*, vol. 19, no. 3, pp. 676 – 679, June 2009.

[66] A. Kashlinsky and M. J. Rees, "Formation of Population III starts and pregalactic evolution," *Monthly Notices of the Royal Astronomical Society*, vol. 205, pp. 955–971, December 1983.

[67] "Submarine cable map 2012," http://submarine-cable-map-2012.telegeography.com/.

[68] S. Kaplan and H. Engseth, "A high-speed superconducting digital autocorrelator operating at 20.48GHz," *Supercond. Sci. Technol.*, vol. 20, pp. S310 – 314, 2007.

[69] WRCAD, "XIC and WRSpice," Whiteley Research Inc., 456 Flora Vista Avenue, Sunnyvale, CA 94086, USA, http://www.wrcad.com.

[70] M. M. Khapaev, *3D-MLSI The program for extraction of 3D inductances of multilayer superconductor circuits*, Moscow State University, September 2004, available at http://vm.cs.msu.ru/sotr/vmhap/3dmlsi/3dmlsm.pdf. Last accessed on 18 June 2012.

[71] P. Deitel and H. Deitel, *C How to Program.* Prentice Hall, 2012.

[72] Central limit theorem. Encyclopedia of Mathematics. [Online]. Available: http://www.encyclopediaofmath.org/index.php?title=Central_limit_theorem&oldid=26376

[73] J. Satchell, "Stochastic simulation of SFQ logic," *IEEE Transactions on Applied Superconductivity*, vol. 7, no. 2, pp. 3315–3318, 1997.

[74] S. Polonsky, "PSCAN'96: new software for simulation and optimization of complex RSFQ circuits," *IEEE Transactions on Applied Superconductivity*, vol. 7, no. 2, pp. 2685–2689, June 1997.

[75] T. Harnisch, J. Kunert, H. Toepfer, and H. Uhlmann, "Design centering methods for yield optimization of cryoelectronic circuits," *IEEE Transactions on Applied Superconductivity*, vol. 7, no. 2, pp. 3434 –3437, June 1997.

[76] F. G. Ortmann, A. van der Merwe, H. R. Gerber, and C. J. Fourie, "A comparison of multi-criteria evaluation methods for RSFQ circuit optimization," *IEEE Transactions on Applied Superconductivity*, vol. 21, no. 3, pp. 801–804, June 2011.

[77] D. Hodges and D. Pederson, "Invited: The here and now of computer-aided circuit design," in *Solid-State Circuits Conference. Digest of Technical Papers. 1974 IEEE International*, vol. XVII, February 1974, pp. 38 – 39.

[78] E. S. Fang and T. van Duzer, "A Josephson integrated circuit simulator (JSIM) for superconductive electronic applications," in *2nd ISEC*, no. 407-410, Tokyo, Japan, 1989.

[79] S. V. Polonsky, V. K. Semenov, and P. N. Shevchenko, "PSCAN: personal superconductor circuit analyser," *Supercond. Sci. Technol.*, vol. 4, no. 11, pp. 667–670, 1991.

[80] Institut für Photonische Technologien e.V., P.O. Box 100239, 07702 Jena, Germany, http://www.ipht-jena.de.

[81] "LAyout System for Individuals," http://lasihomesite.com/, last accessed 10 August 2012.

[82] "AutoHotKey," http://www.autohotkey.com.

[83] juspertor UG (haftungsbeschränkt), "Layouteditor," Hans-Durach-Str. 4, 82008 Unterhaching, Germany, http://www.layouteditor.net.

[84] C. J. Fourie, O. Wetzstein, T. Ortlepp, and J. Kunert, "Three-dimensional multi-terminal superconductive integrated circuit inductance extraction," *Supercond. Sci. Technol.*, vol. 24, 2011.

[85] P. I. Bunyk and S. V. Rylov, "Automated calculation of mutual inductance matrices of multilayer superconductor integrated circuits," *Ext. Abs. ISEC*, p. 62, 1993.

[86] M. Kamon, M. J. Tsuk, and J. K. White, "Fasthenry: a multipole-accelerated 3-D inductance extraction program," *IEEE Transactions on Microwave Theory and Techniques*, vol. 42, pp. 1750–1758, 1994.

[87] S. R. Whiteley, "FastHenry 3.0wr," available online at http://www.wrcad.com/.

[88] C. J. Fourie, "Calibration of inductance calculations to measurement data for superconductive integrated circuit processes," *IEEE Transactions on Applied Superconductivity*, vol. 23, p. 1301305, 2013.

[89] I. Haverkamp, O. Wetzstein, T. Ortlepp, R. Stolz, H.-G. Meyer, and H. Toepfer, "Optimization of a digital squid magnetometer in terms of noise and distortion," *Supercond. Sci. Technol.*, vol. 25, 2012, 065012.

[90] C. J. Fourie and M. H. Volkmann, "Oral contribution: Full-gate verification of SFQ circuit layouts with InductEx," in *European Conference on Applied Superconductivity*, 2013.

[91] NioCAD Team, *NioPulse User Manual*, 2nd ed., NioCAD, Pty (Ltd), Block D, De Wagenweg Office Park, Stellentia Avenue, Stellenbosch, 7600, South Africa, http://www.niocad.com, April 2012.

[92] Cadence Design Systems, Inc., 2655 Seely Avenue, San Jose, CA 95134, USA, http://www.cadence.com.

[93] NioCAD, Pty (Ltd), "NioPulse software suite," Block D, De Wagenweg Office Park, Stellentia Avenue, Stellenbosch, 7600, South Africa, http://www.niocad.com.

[94] Y. Saad and M. H. Schultz, "GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems," *SIAM J. Sci. Stat. Comput.*, vol. 7, pp. 856–869, 1986.

[95] OpenMP Architecture Review Board, "OpenMP application program interface version 3.0," May 2008. [Online]. Available: http://www.openmp.org/mp-documents/spec30.pdf

[96] "Boost C++ libraries," http://www.boost.org.

[97] N. Wirth, "What can we do about the unnecessary diversity of notation for syntactic definitions?" *Communications of the ACM*, vol. 20, no. 11, pp. 822–823, November 1977.

[98] P. Febvre, *MANUAL to handle JSIM and optimizer program AUTO5 form University of Technology Ilmenau*, University of Savoie, written in the framework of the European FLUXONICS initiative.

[99] H. Nyquist, "Thermal agitation of electric charge in conductors," *Phys. Rev.*, vol. 32, p. 110, 1928.

[100] J. Johnson, "Thermal agitation of electricity in conductors," *Phys. Rev.*, vol. 32, p. 97, 1928.

[101] I. Foster, *Designing and Building Parallel Programs*. Addison-Wesley Publishing Company, 1995, Section 1.4.4.

[102] K. E. Atkinson, *An Introduction to Numerical Analysis*, 2nd ed. John Wiley & Sons, 1989.

[103] E. Jones, T. Oliphant, P. Peterson *et al.*, "SciPy: Open source scientific tools for Python," 2001–. [Online]. Available: http://www.scipy.org

[104] H. Ko and T. van Duzer, "A new high-speed periodic-threshold comparator for use in a Josephson A/D converter," *IEEE Journal of Solid-State Circuits*, vol. 23, no. 4, pp. 1017 –1021, aug 1988.

[105] P. E. Black. (2011, December) Gray code. in Dictionary of Algorithms and Data Structures. U. S. National Institute of Standards and Technology. Last accessed 6 June 2012. [Online]. Available: http://www.nist.gov/dads/HTML/graycode.html

[106] B. P. Lathi, *Linear Systems and Signals*. Oxford University Press, July 2004.

[107] L. Schuchman, "Dither signals and their effect on quantization noise," *Communication Technology, IEEE Transactions on*, vol. 12, no. 4, pp. 162 –165, December 1964.

[108] J. Faulhaber, *Academia Algebrae - Darinnen die miraculosische Inventiones zu den höchsten Cossen weiters continuirt und profitiert werden.* Johann Remelins Kunst und Buchhändlers, Ulm., 1631.

[109] T. Filippov and M. Znosko, "Time characteristics of a Josephson-balanced comparator," *Supercond. Sci. Technol.*, vol. 12, pp. 776–778, 1999.

[110] G. van Rossum, "Python tutorial," Technical Report CS-R9526, May 1995, centrum voor Wiskunde en Informatica (CWI), Amsterdam.

[111] T. Ortlepp, S. Miyajima, H. Toepfer, and A. Fujimaki, "Josephson comparator with modified dynamic behavior for improved sensitivity," *Journal of Applied Physics*, vol. 111, p. 123901, 2012.

[112] T. Haddad, O. Wetzstein, S. Engert, H. Toepfer, and T. Ortlepp, "Investigation of the relationship between the gray zone and the clock frequency of a Josephson comparator," *Supercond. Sci. Technol.*, vol. 24, p. 095010, 2011.

[113] T. Ortlepp, M. H. Volkmann, and Y. Yamanashi, "Memory effect in balanced josephson comparators," in *Extended Abstracts of the International Superconducting Electronics Conference in Cambridge*, 2013.

[114] C. J. Fourie, "InductEx," http://www.sun.ac.za/inductex, last accessed on 28 September 2012.

[115] S. Yorozu, Y. Kameda, H. Terai, A. Fujimaki, T. Yamada, and S. Tahara, "A single flux quantum standard logic cell library," *Physica C - Superconductivity and its Applications*, vol. 378, pp. 1471–1474, 2002.

[116] IARPA, "Broad agency announcement: Cryogenic computing complexity (C3) program," https://www.fbo.gov/index?s=opportunity&mode=form&tab=core&id=55cefaec5a3dcac236383a90951e7f71, July 2013, last accessed 28 August 2013.

[117] A. Bozbey, S. Miyajima, T. Ortlepp, and A. Fujimaki, "Design and circuit analysis of quasi-one junction SQUID comparators for low temperature detector array read-out," *Journal of Superconductivity and Novel Magnetism*, September 2010.

[118] B. Ebert, "Noise-induced jitter and decision errors in Rapid Single Flux Quantum Electronics," Ph.D. dissertation, Technische Universität Ilmenau, 2010.

[119] S. Miyajima, T. Ortlepp, H. Toepfer, A. Bozbey, and A. Fujimaki, "Experimental demonstration and numerical analysis of microampere gray zone width with enhanced operating margin in shunted quasi-one

junction superconducting quantum interference device comparators," *Japanese Journal of Applied Physics*, vol. 52, p. 033101, 2013.

[120] Q. Herr and P. Bunyk, "Implementation and application of first-in first-out buffers," *IEEE Transactions on Applied Superconductivity*, vol. 13, no. 2, pp. 563 – 566, june 2003.

[121] J. A. Nelder and R. Mead, "A simplex method for function minimization," *The Computer Journal*, vol. 7, no. 4, pp. 308–313, 1965.

[122] Agilent Technologies, *Datasheet: 33500B Series Waveform Generators*, August 2012, http://cp.literature.agilent.com/litweb/pdf/5991-0692EN.pdf, Last accessed 28 August 2013.

[123] V. Ambegaokar and B. I. Halperin, "Voltage due to thermal noise in the D.C. Josephson effect," *Physical Review Letters*, vol. 22, p. 1364, 1969.

[124] XILINX, *Virtex-6 FPGA Data Sheet: DC and Switching Characteristics*, January 2012, http://www.xilinx.com/support/documentation/data_sheets/ds152.pdf, Last accessed on 1 September 2012.

[125] A. X. Widmer and P. A. Franaszek, "A DC-balanced, partitioned-block, 8b/10b transmission code," *IBM Journal of Research and Development*, vol. 27, p. 440, 1983.

[126] S. Borkar and A. Chien, "The future of microprocessors," *Commun. ACM*, vol. 5, no. 5, pp. 67–77, May 2011.

[127] Q. P. Herr, A. Y. Herr, O. T. Oberg, and A. G. Ioannidis, "Ultra-low-power superconductor logic," *Journal of Applied Physics*, vol. 109, no. 103903, May 2011.

[128] M. Tanaka, M. Ito, A. Kitayama, T. Kouketsu, and A. Fujimaki, "18-GHz, 4.0-aJ/bit operation of ultra-low-energy Rapid Single-Flux-Quantum shift registers," *Japanese Journal of Applied Physics*, vol. 51, 2012, 053102.

[129] T. Ortlepp, O. Wetzstein, S. Engert, J. Kunert, and H. Toepfer, "Reduced power consumption in superconducting electronics," *IEEE Transactions on Applied Superconductivity*, vol. 21, p. 770, 2011.

[130] T. Instruments, *Precision, Low Power Instrumentation Amplifiers*, October 1995, http://www.ti.com/lit/ds/sbos051b/sbos051b.pdf Last accessed on 28 September 2012.

[131] S. Sarwana, D. E. Kirichenko, V. V. Dotsenko, A. F. Kirichenko, S. B. Kaplan, and D. Gupta, "Multi-band digital-RF receiver," *IEEE Transactions on Applied Superconductivity*, vol. 21, no. 3, pp. 677–680, 2011.

[132] H. Suzuki, A. Inoue, T. Imamura, and S. Hasuo, "A Josephson driver to interface Josephson junctions to semiconductor transistors," *Dig. Tech. Papers, International Electron Devices Meeting*, pp. 290–293, 1988.

[133] T. Ortlepp, L. Zheng, S. R. Whiteley, and T. V. Duzer, "Design guidelines for Suzuki stacks as reliable high-speed Josephson voltage drivers," *Superconductor Science and Technology*, vol. 26, p. 035007, 2013.

[134] A. Rylyakov, "New design of single-bit all-digital RSFQ autocorrelator," *IEEE Transactions on Applied Superconductivity*, vol. 7, p. 2709, 1997.

[135] Y. Yamanashi, T. Nishigai, and N. Yoshikawa, "Study of LR-loading technique for low-power single flux quantum circuits," *IEEE Transactions on Applied Superconductivity*, vol. 17, p. 150, 2007.

[136] D. Kirichenko, S. Sarwana, and A. Kirichenko, "Zero static power dissipation biasing of RSFQ circuits," *IEEE Transactions on Applied Superconductivity*, vol. 21, p. 776, 2011.

[137] S. V. Polonsky, V. K. Semenov, P. Bunyk, A. F. Kirichenko, A. Kidiyarova-Shevchenko, O. A. Mukhanov, P. Shevchenko, D. Schneider, D. Y. Zinoviev, and K. K. Likharev, "New RSFQ circuits," *IEEE Transactions on Applied Superconductivity*, vol. 3, p. 2566, 1993.

[138] O. A. Mukhanov, "Rapid Single Flux Quantum (RSFQ) shift register family," *IEEE Transactions on Applied Superconductivity*, vol. 3, p. 3102, 1993.

[139] J. Majer, J. Butcher, and J. Mooij, "Simple phase bias for superconducting circuits," *Applied Physics Letters*, vol. 80, p. 3638, 2002.

[140] A. Ustinov and V. Kaplunenko, "Rapid single-flux quantum logic using π-shifters," *Journal of Applied Physics*, vol. 94, p. 5405, 2003.

[141] B. Dimov, D. Balashov, M. Khabipov, T. Ortlepp, F.-I. Buchholz, A. Zorin, J. Niemeyer, and F. Uhlmann, "Implementation of superconductive passive phase shifters in RSFQ digital circuits," *Supercond. Sci. Technol.*, vol. 21, 2008, 45007.

[142] O. Mielke, T. Ortlepp, J. Kunert, H.-G. Meyer, and H. Toepfer, "Controlled initialization of superconducting π-phase shifters and possible applications," *Supercond. Sci. Technol.*, vol. 23, 2010, 055003.

[143] S. Narayana and V. K. Semenov, "Power independent RSFQ logic cells," *Extended Abstracts of 11th International Superconducting Electronics Conference (ISEC'07)*, 2007, o-A01.

[144] O. Mielke, T. Ortlepp, P. Febvre, and F. Uhlmann, "Reduced probability of noise introduced malfunction in RSFQ circuits by implementing intrinsic π-phaseshifter," *IEEE Transactions on Applied Superconductivity*, vol. 19, p. 621, 2009.

[145] T. Ortlepp, O. Mielke, J. Kunert, and H. Toepfer, "Improved operation range of digital superconductive electronics by implementing passive phaseshifters," *Physica C*, vol. 470, p. 1955, 2010.

[146] A. F. Kirichenko, V. K. Semenov, Y. K. Kwong, and V. Nandakumar, "4-bit rapid single-flux-quantum decoder," *IEEE Transactions on Applied Superconductivity*, vol. 5, p. 2857, 1995.

[147] S. B. Kaplan and O. A. Mukhanov, "Operation of a high-speed superconducting demultiplexer using rapid single flux quantum (RSFQ) technology," *IEEE Transactions on Applied Superconductivity*, vol. 5, p. 2853, 1995.

[148] D. Y. Zinoviev and K. K. Likharev, "Feasibility study of RSFQ-based self-routing nonblocking digital switches," *IEEE Transactions on Applied Superconductivity*, vol. 7, p. 3155, 1997.

[149] S. V. Polonsky, V. K. Semenov, and A. F. Kirichenko, "Single flux quantum B flip-flop and its possible applications," *IEEE Transactions on Applied Superconductivity*, vol. 14, p. 9, 1994.

[150] M. Jeffery, T. Van Duzer, J. R. Kirtley, and M. B. Ketchen, "Magnetic imaging of moat-guarded superconducting electronic circuits," *Applied Physics Letters*, vol. 67, no. 12, pp. 1769 –1771, September 1995.

[151] D. Gupta, T. V. Filippov, A. F. Kirichenko, D. E. Kirichenko, I. V. Vernik, A. Sahu, S. Sarwana, P. Shevchenko, A. Talalaevskii, and O. A. Mukhanov, "Digital channelizing radio frequency receiver," *IEEE Transactions on Applied Superconductivity*, vol. 17, no. 2, pp. 430–437, June 2007.

[152] D. Y. Zinoviev and Y. A. Polyakov, "Octopux: an advanced automated setup for testing superconductor circuits," *IEEE Transactions on Applied Superconductivity*, vol. 7, p. 3240, 1997.

[153] R. P. Robertazzi and S. V. Rylov, "Synchronous flux quantizing analog-to-digital-converter," *IEEE Transactions on Applied Superconductivity*, vol. 3, no. 4, pp. 3114–3116, 1993.

[154] S. V. Rylov and R. P. Robertazzi, "Superconducting high-resolution A/D converter based on phase modulation and multichannel timing arbitration," *IEEE Transactions on Applied Superconductivity*, vol. 5, no. 2, pp. 2260–2263, 1995.

[155] S. M. Anton, I. A. B. Sognnaes, J. S. Birenbaum, S. R. O'Kelley, C. J. Fourie, and J. Clarke, "Mean square flux noise in SQUIDs and qubits: numerical calculations," *Superconductor Science and Technology*, vol. 26, p. 075022, 2013.

[156] C. J. Fourie, S. M. Anton, and J. Clarke, "Magnetic field calculations in the vicinity of superconductive circuit structures," in *Poster at the International Superconductor Electronics Conference '13*, 2013.

[157] A. R. Thompson, J. M. Moran, and G. W. S. Jr., *Interferometry and Synthesis in Radio Astronomy*. Wiley-VCH, 2001.

[158] T. H. Kim and J. S. Moodera, "Magnesium diboride superconductor thin film tunnel junctions for superconductive electronics," *Journal of Applied Physics*, vol. 100, p. 113904, 2006.

[159] Hittite Microwave Corporation, "Datasheet: 26 GS/s 3-bit analog-to-digital converter," http://www.hittite.com/content/documents/data_sheet/hmc5831lp9b.pdf, last accessed 28 August 2013.

Appendix A

Circuits and Layouts

A.1 Comparators

A.1.1 MC1 — Quasi-one junction SQUID-based multi-threshold comparator

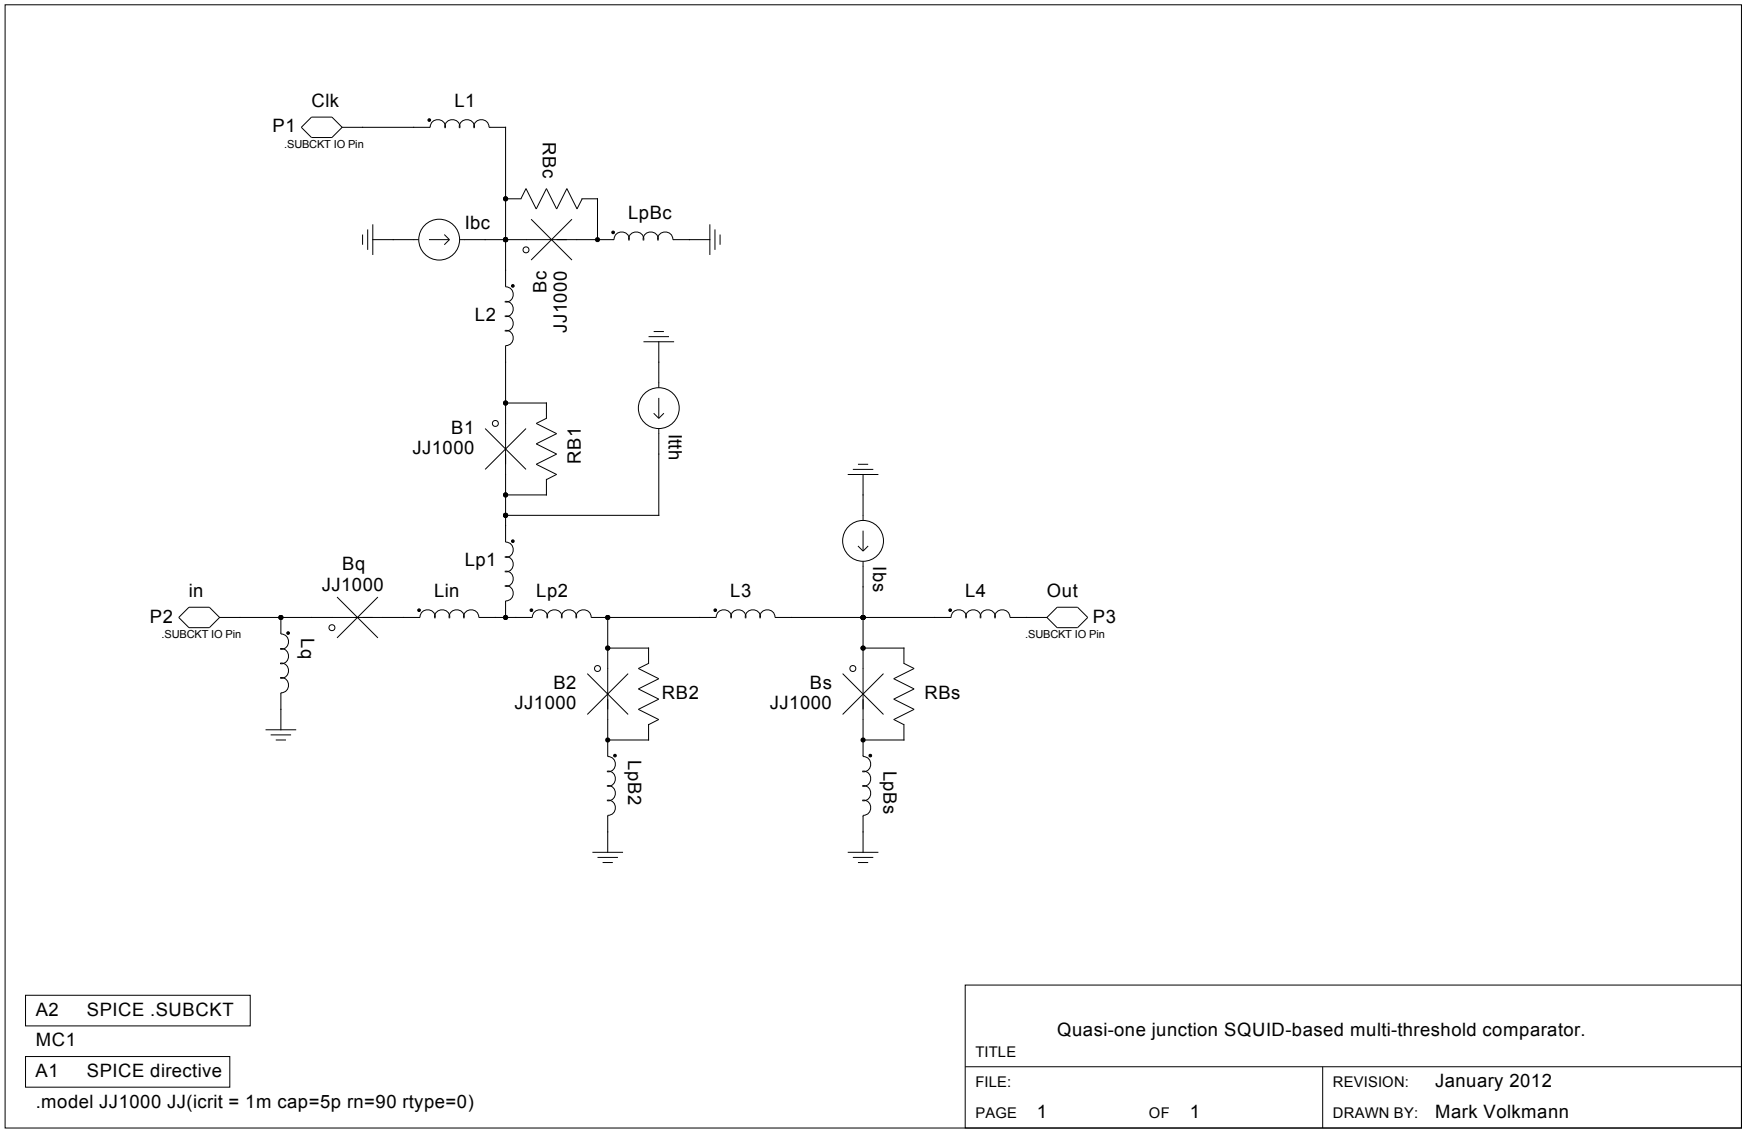

Figure A.1: Final circuit diagram of QOS-based multi-threshold comparator MC1 (Parameter values in Table A.1).

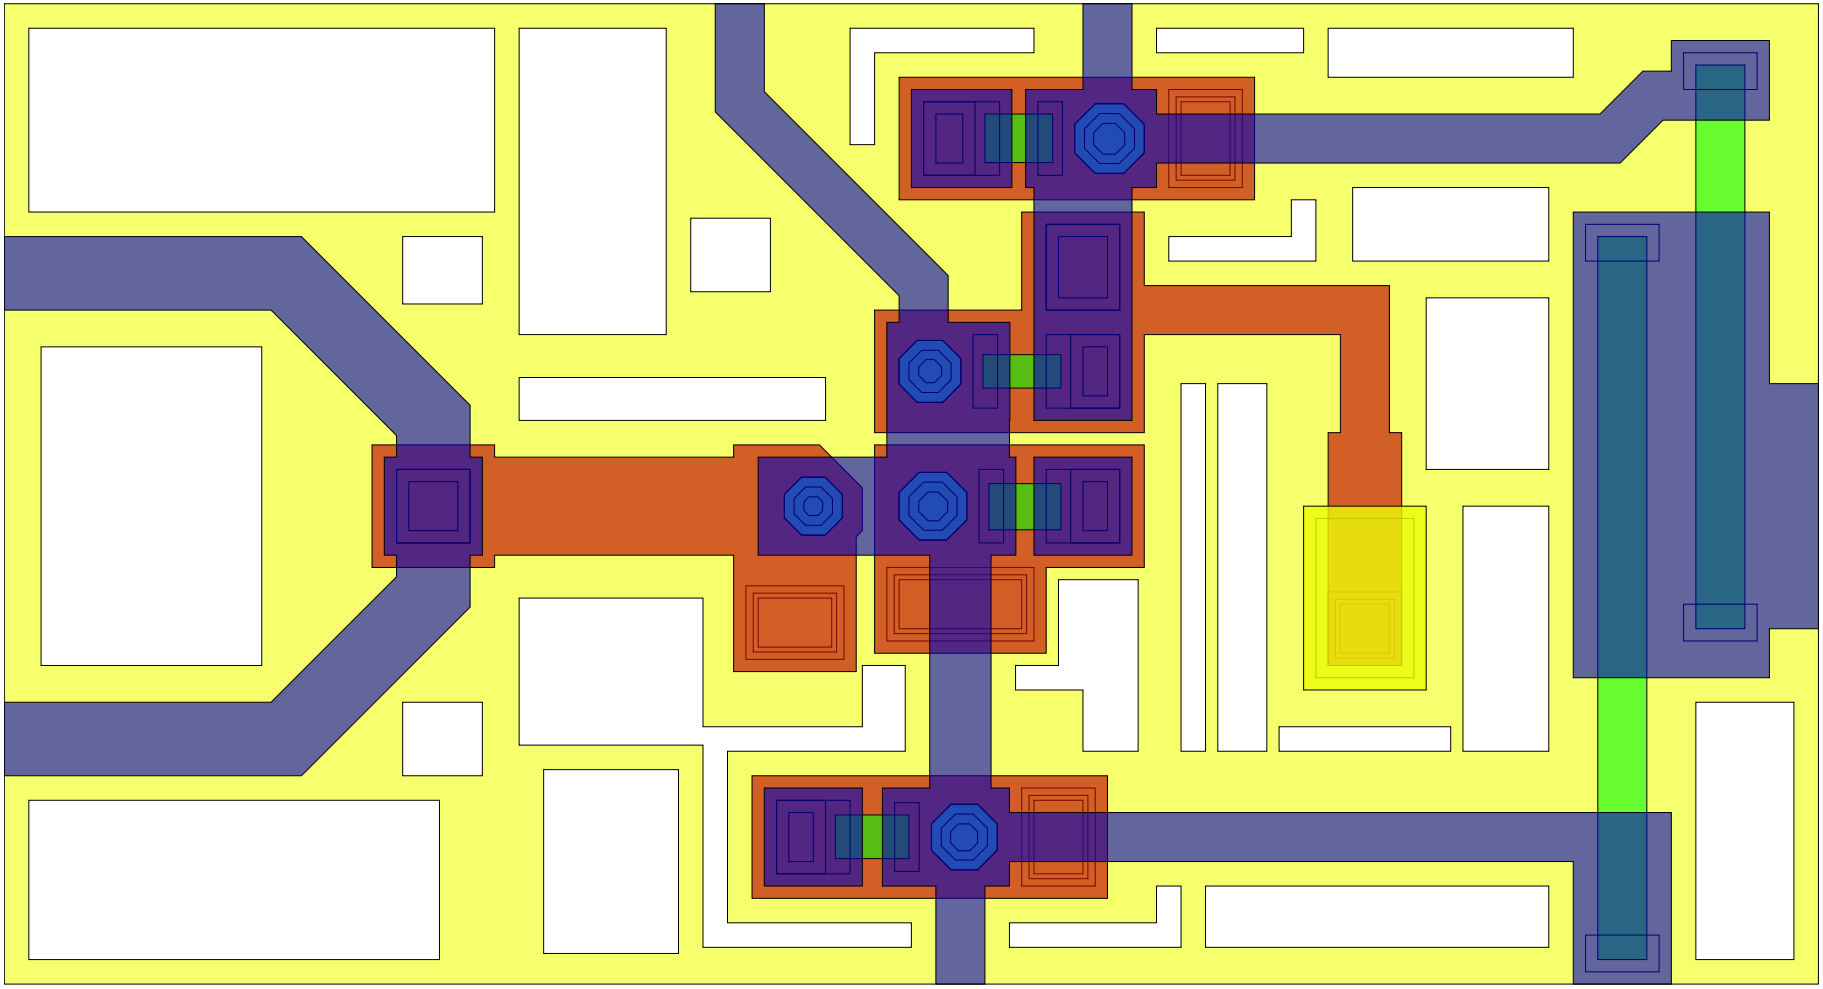

Figure A.2: Final layout of QOS-based multi-threshold comparator (MC1).

A.1.2 MC2 — Complementary quasi-one junction SQUID-based multi-threshold comparator

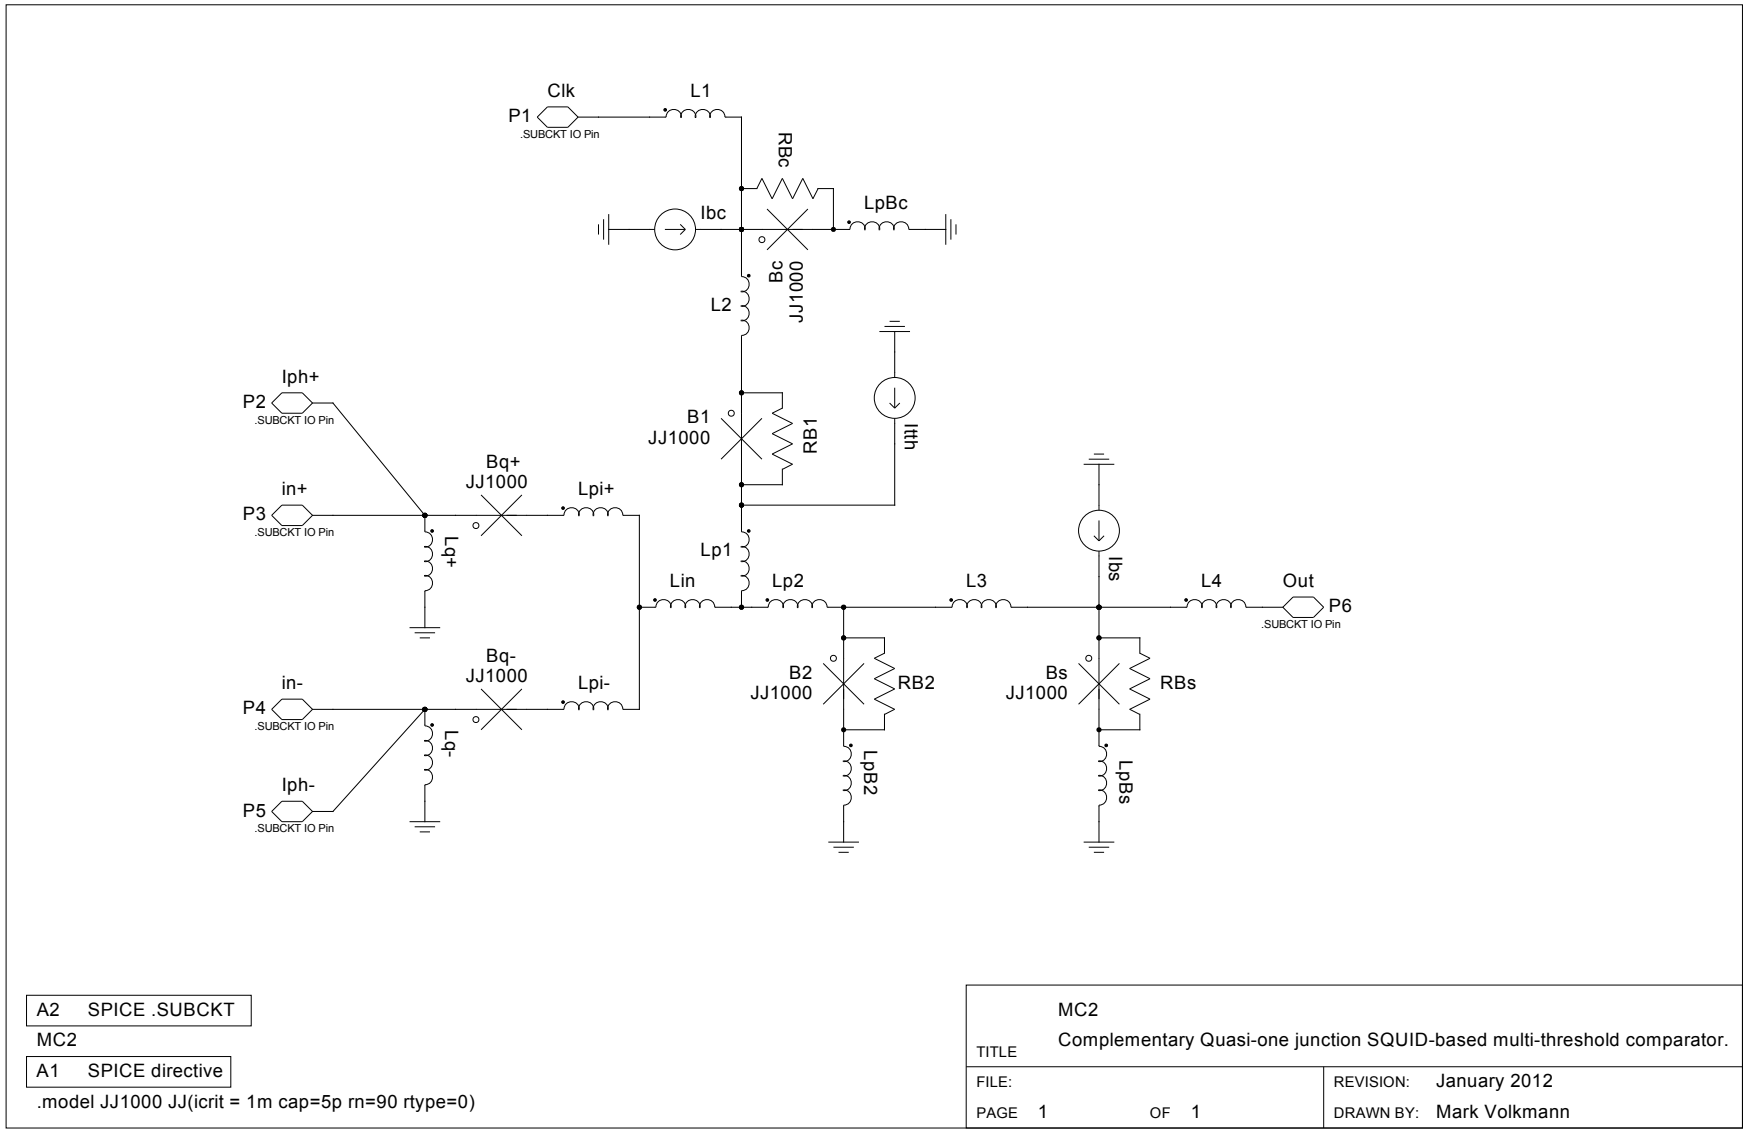

Figure A.3: Final circuit diagram of CQOS-based multi-threshold comparator MC2 (Parameter values in Table A.2).

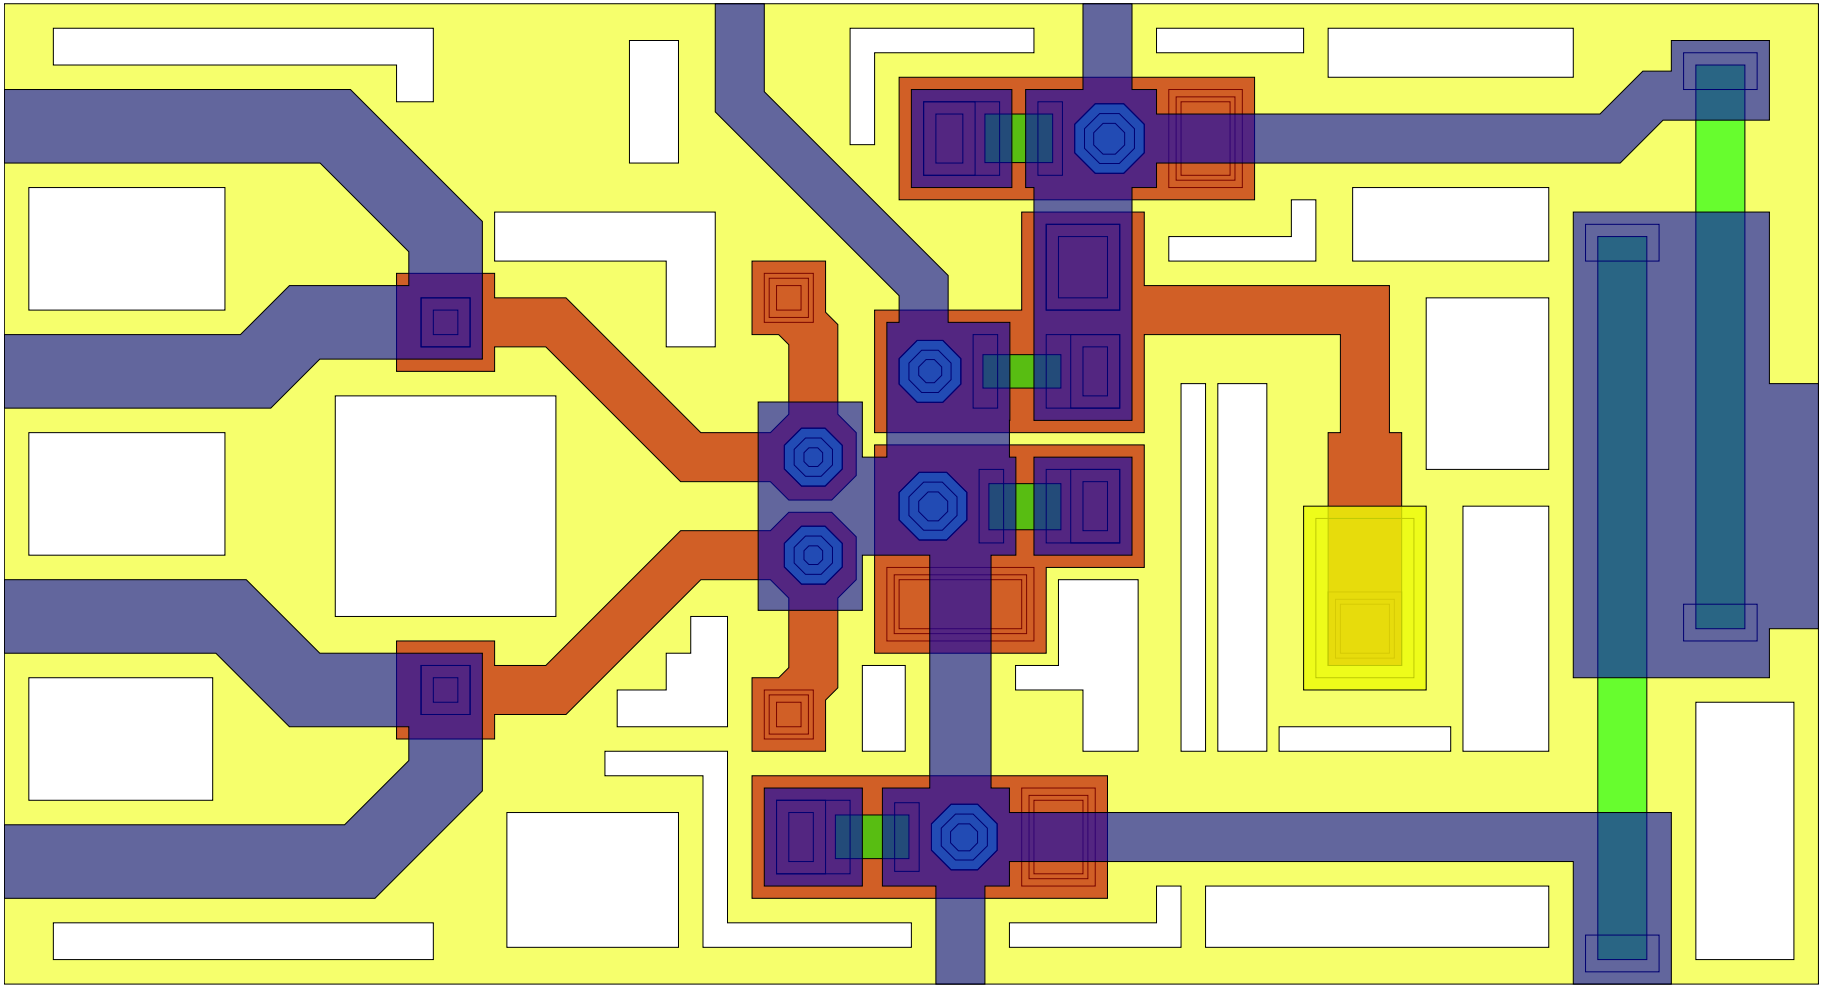

Figure A.4: Final layout of CQOS-based multi-threshold comparator MC2.

Table A.1: Extracted parameters of MC1.

Parameter	**Value**	**Parameter**	**Value**
L_1	1.93 pH	L_q	0.167 pH
L_2	1.90 pH	I_{cq}	125 µA
L_3	3.87 pH	β_{cq}	$\gg 1$
L_4	2.19 pH	I_{cB1}	175 µA
L_{in}	1.03 pH	I_{cB2}	275 µA
L_{p1}	0.969 pH	β_{cB1}	1
L_{p2}	0.221 pH	β_{cB2}	1
L_{pB2}	0.035 pH	I_{cBc}	325 µA
L_{pBc}	0.114 pH	I_{cBs}	250 µA
L_{pBs}	0.126 pH	β_{cBs}	1
I_{bc}	250 µA	β_{cBc}	1
I_{bs}	175 µA	I_{tth}	190 µA

Table A.2: Extracted parameters of MC2.

Parameter	**Value**	**Parameter**	**Value**
L_1	1.92 pH	L_q^+	1.31 pH
L_2	1.91 pH	L_q^-	1.31 pH
L_3	3.84 pH	I_{cq}^+	125 µA
L_4	2.19 pH	I_{cq}^-	125 µA
L_{in}	0.58 pH	β_{cq}^+	$\gg 1$
L_{p1}	0.921 pH	β_{cq}^-	$\gg 1$
L_{p2}	0.254 pH	I_{cB1}	175 µA
L_{pB2}	0.036 pH	I_{cB2}	275 µA
L_{pBc}	0.117 pH	β_{cB1}	1
L_{pBs}	0.127 pH	β_{cB2}	1
L_{pi}^+	0.466 pH	I_{cBc}	325 µA
L_{pi}^-	0.439 pH	I_{cBs}	250 µA
I_{bc}	250 µA	β_{cBc}	1
I_{bs}	175 µA	β_{cBs}	1
I_{ph}^+	0	I_{tth}	175 µA
I_{ph}^-	−380 µA		

A.1.3 MC3 — Complementary quasi-one junction SQUID-based multi-threshold comparator

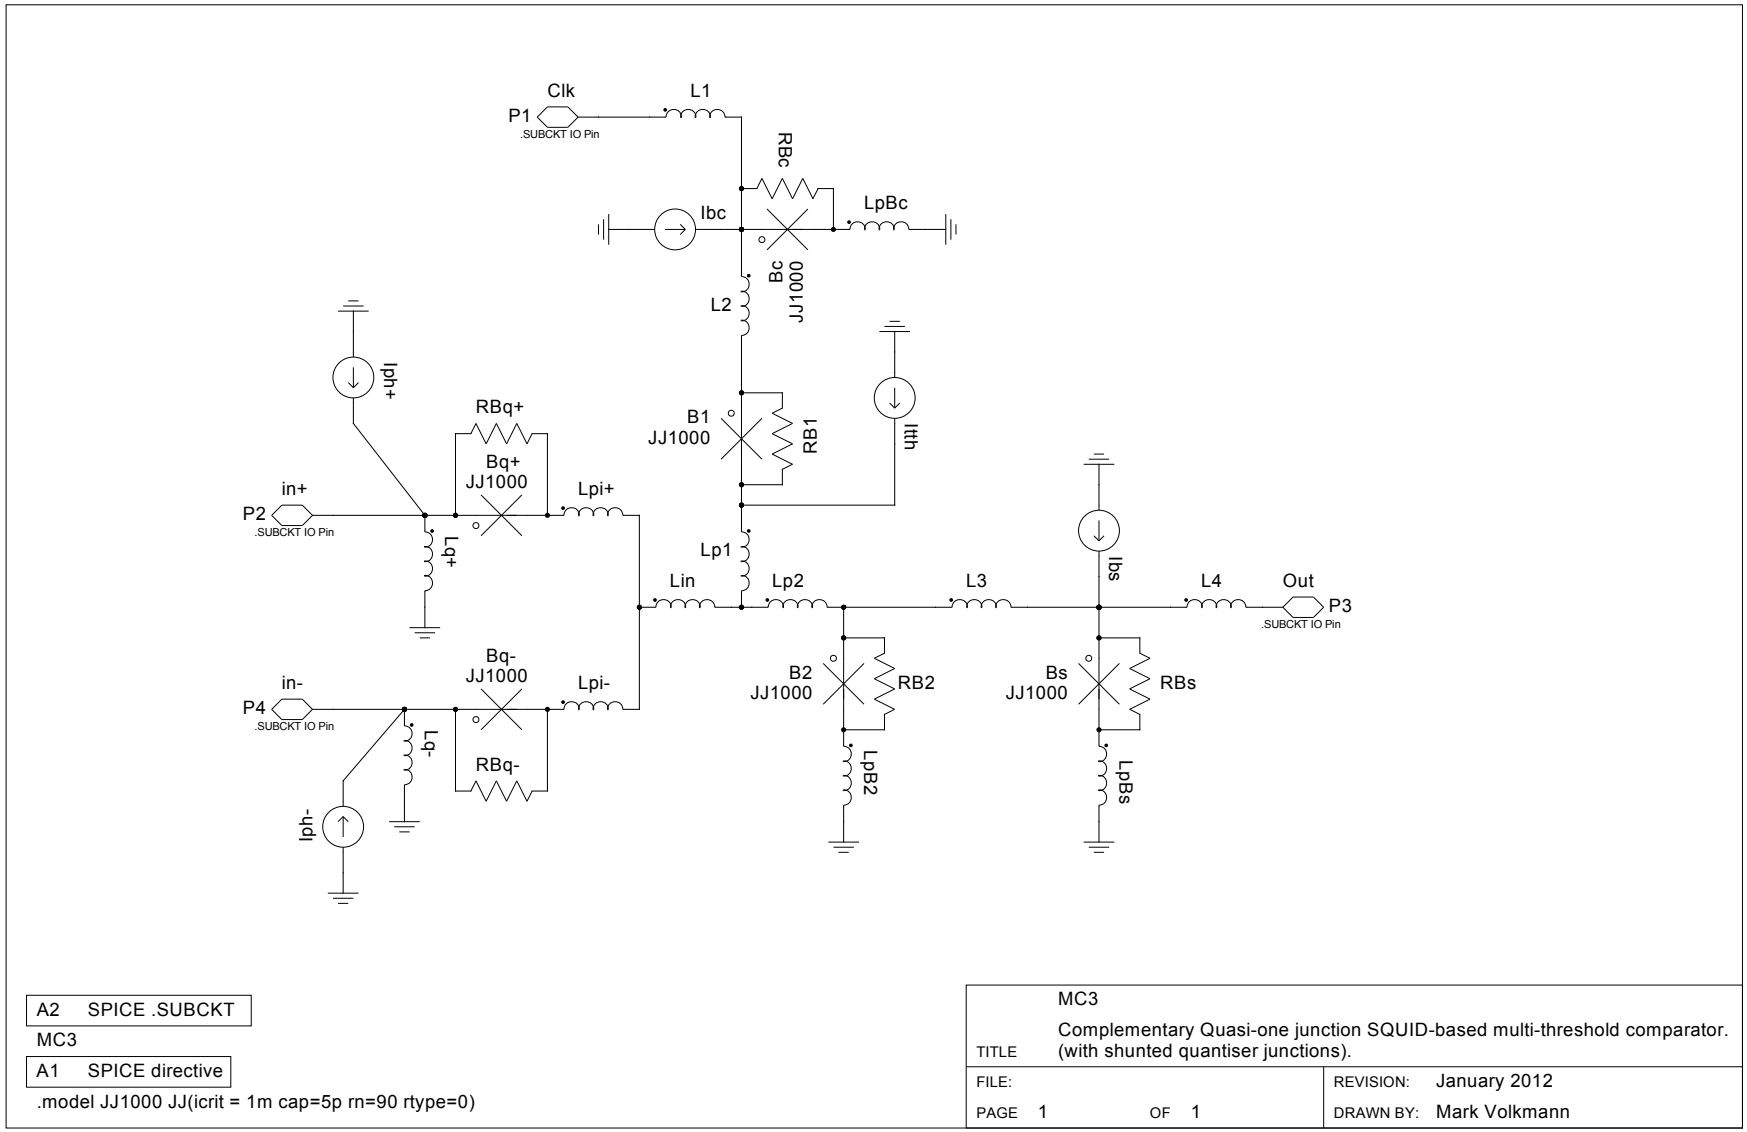

Figure A.5: Final circuit diagram of CQOS-based multi-threshold comparator MC3 (Parameter values in Table A.3).

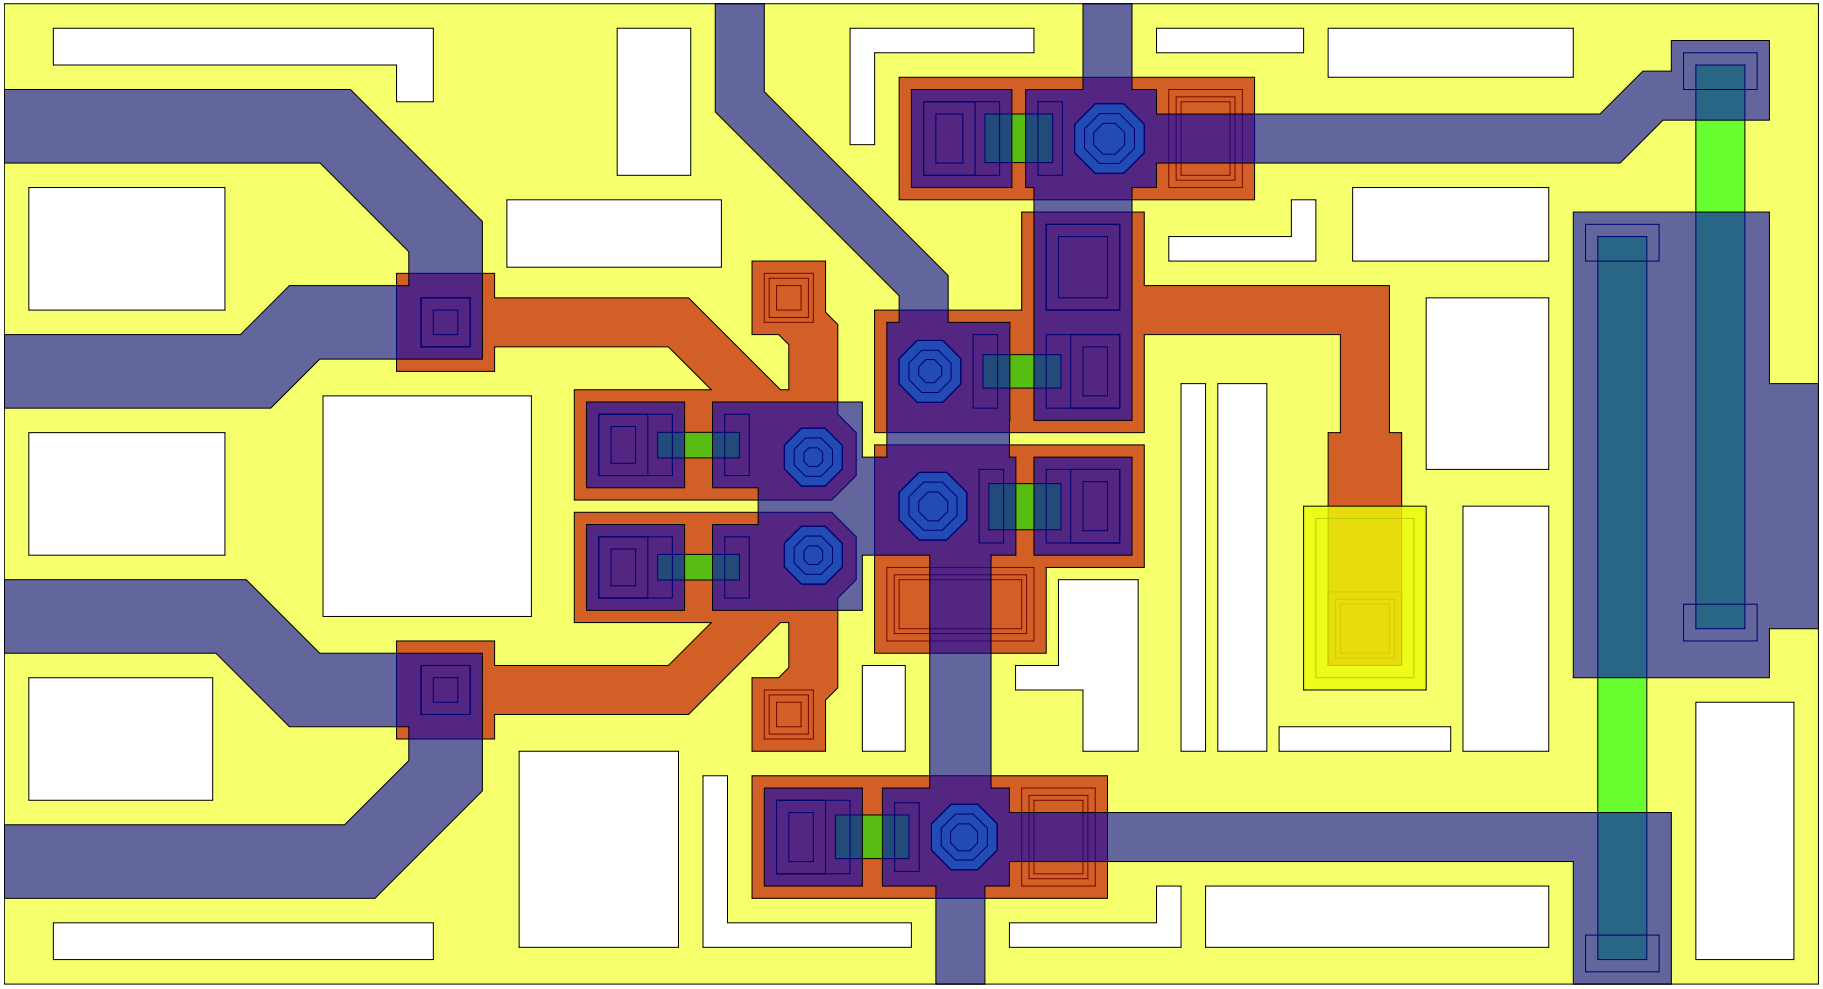

Figure A.6: Final layout of CQOS-based multi-threshold comparator MC3.

A.1.4 MC4 — Complementary quasi-one junction SQUID-based multi-threshold comparator

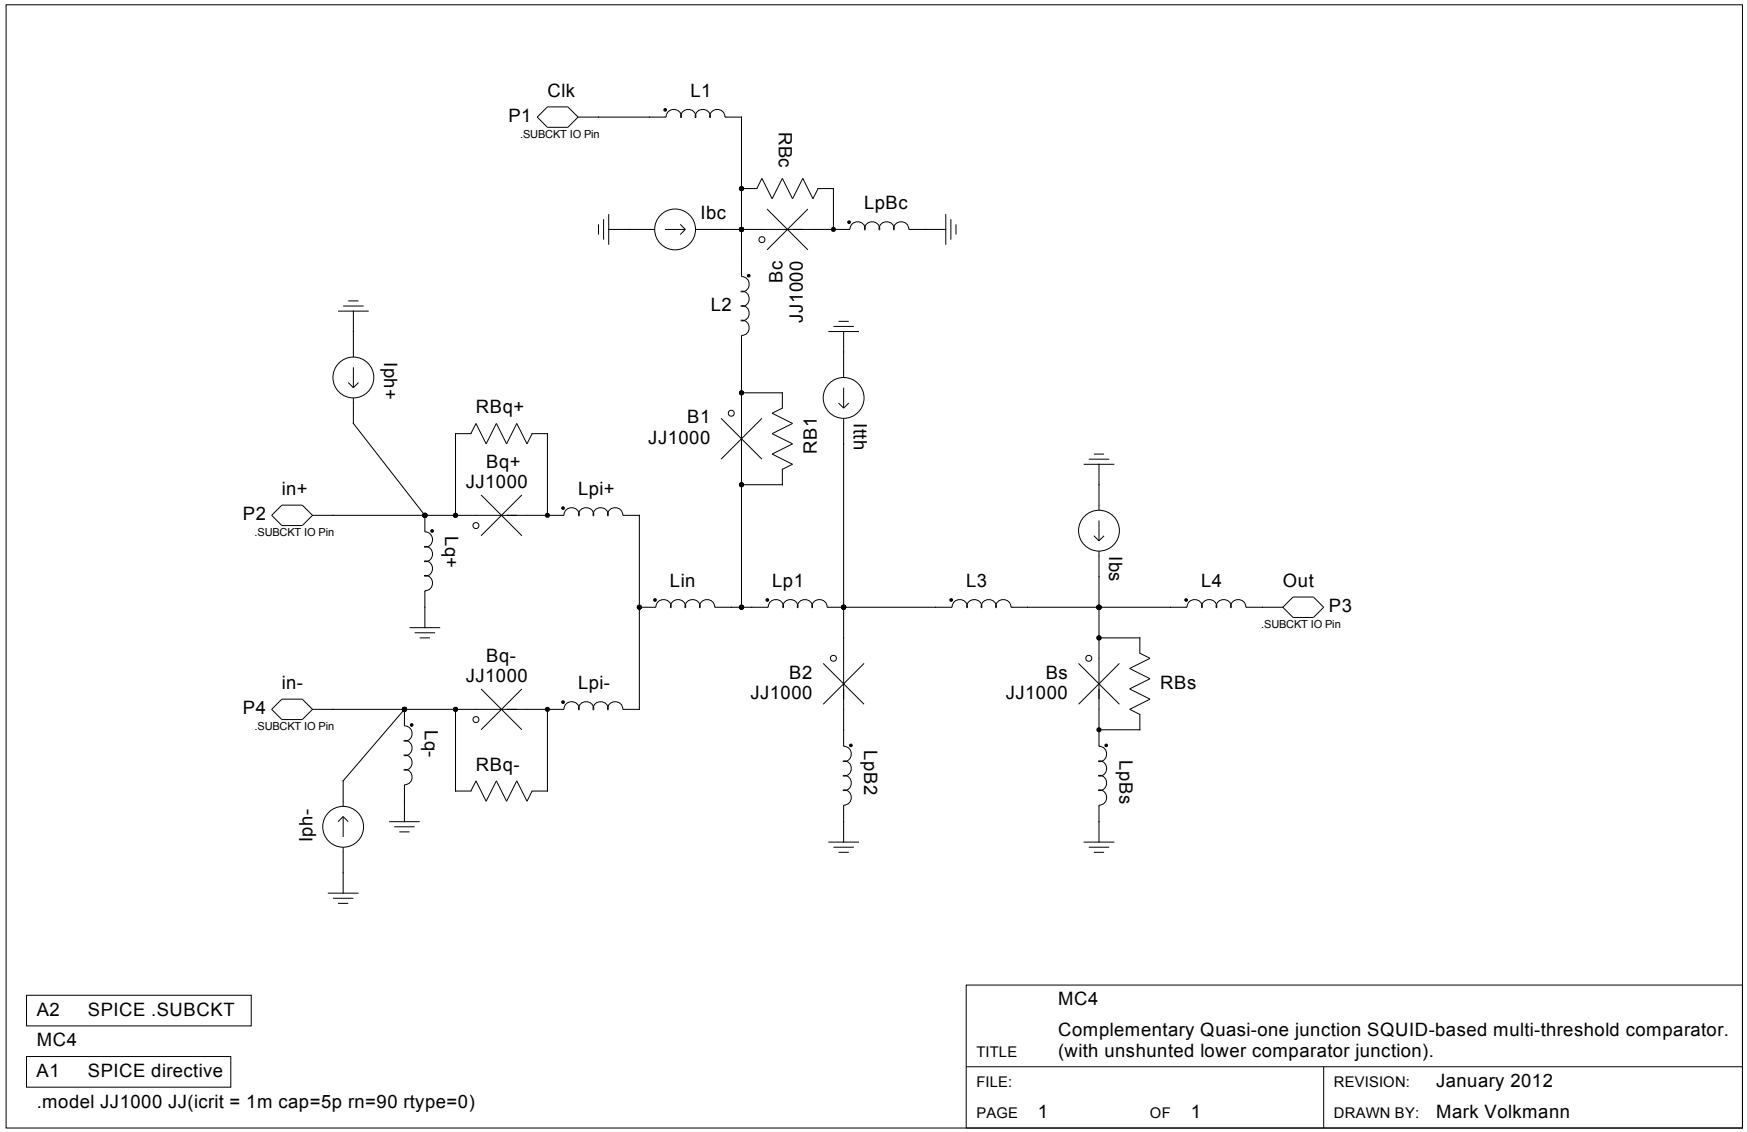

Figure A.7: Final circuit diagram of CQOS-based multi-threshold comparator MC4 (Parameter values in Table A.4).

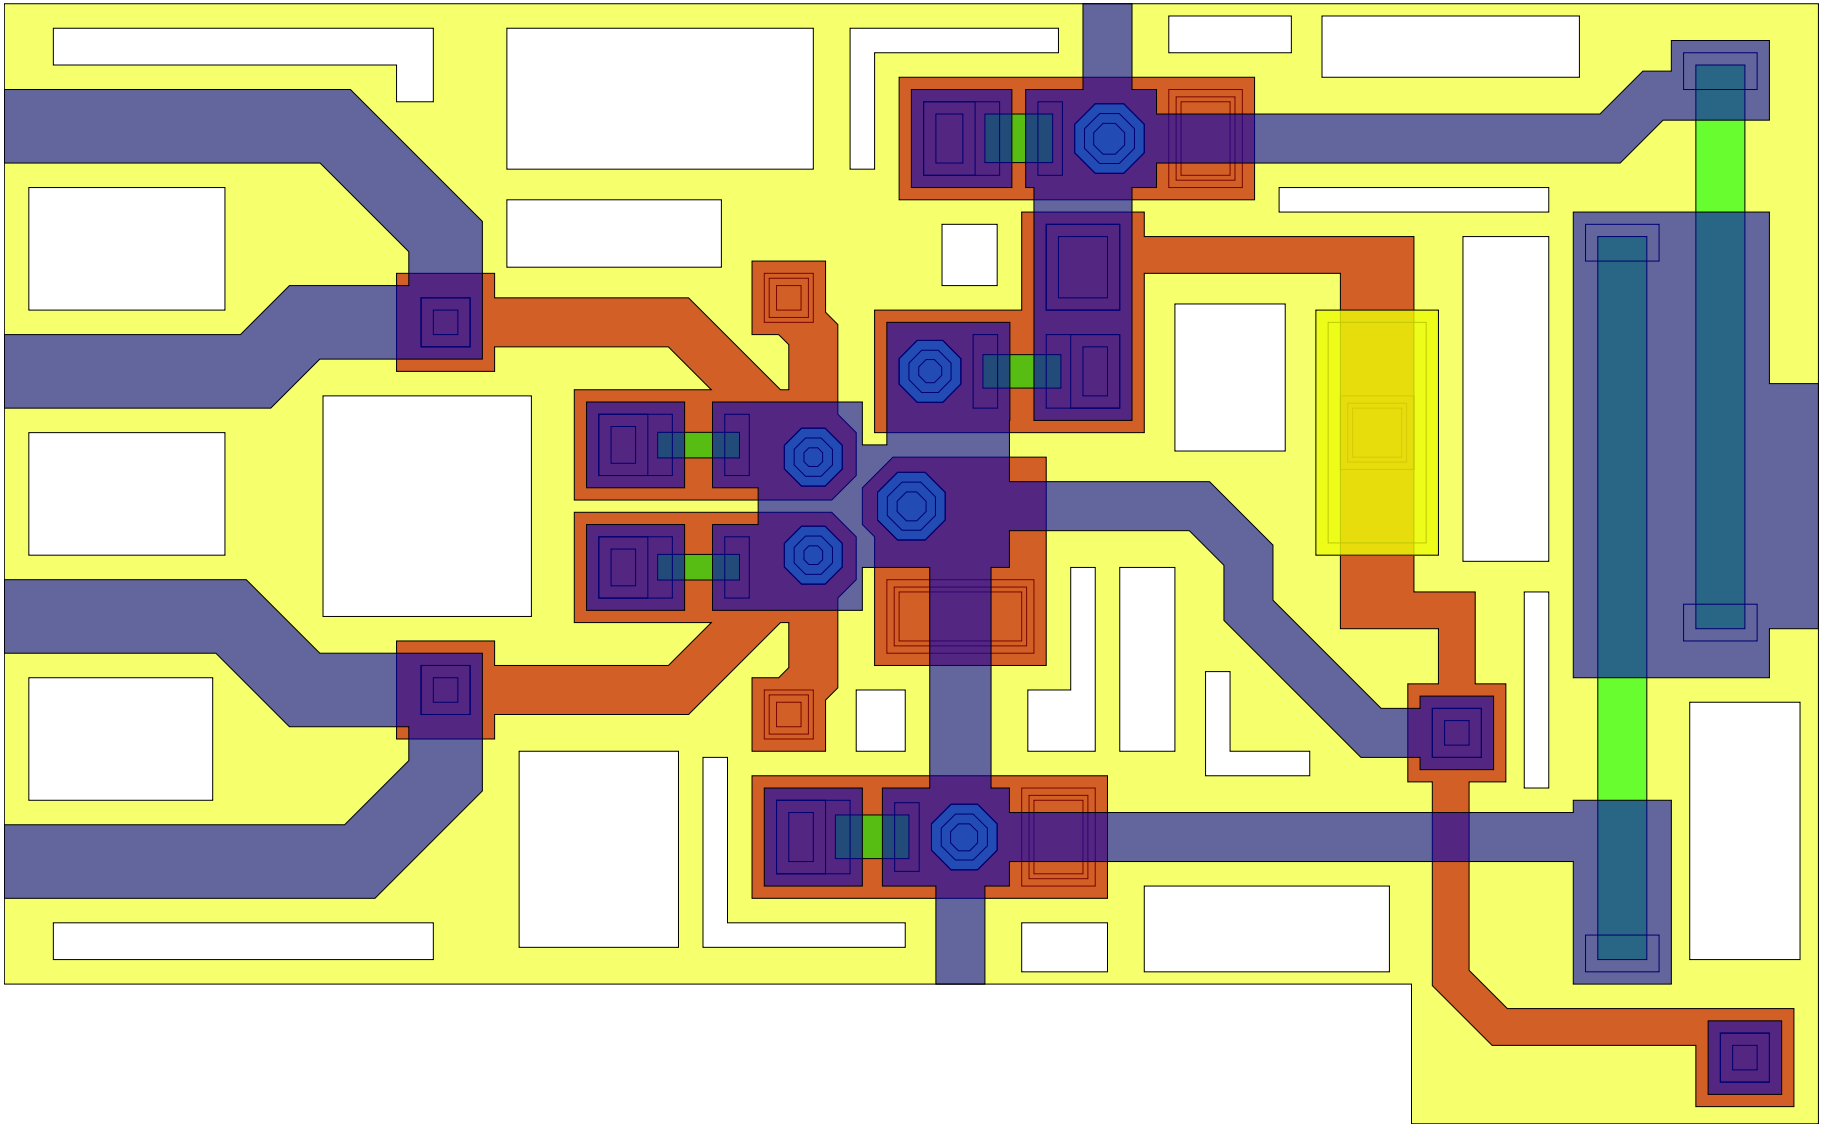

Figure A.8: Final layout of CQOS-based multi-threshold comparator MC4.

Table A.3: Extracted parameters of MC3.

Parameter	**Value**	**Parameter**	**Value**
L_1	1.92 pH	L_q^+	1.082 pH
L_2	1.91 pH	L_q^-	1.086 pH
L_3	3.84 pH	I_{cq}^+	125 µA
L_4	2.19 pH	I_{cq}^-	125 µA
L_{in}	0.59 pH	β_{cq}^+	1
L_{p1}	0.924 pH	β_{cq}^-	1
L_{p2}	0.252 pH	I_{cB1}	175 µA
L_{pB2}	0.035 pH	I_{cB2}	275 µA
L_{pBc}	0.116 pH	β_{cB1}	1
L_{pBs}	0.126 pH	β_{cB2}	1
L_{pi}^+	0.561 pH	I_{cBc}	325 µA
L_{pi}^-	0.561 pH	I_{cBs}	250 µA
I_{bc}	250 µA	β_{cBc}	1
I_{bs}	175 µA	β_{cBs}	1
I_{ph}^+	0	I_{tth}	195 µA
I_{ph}^-	−380 µA		

Table A.4: Extracted parameters of MC4.

Parameter	Value	Parameter	Value
L_1	1.94 pH	L_q^+	1.090 pH
L_2	2.85 pH	L_q^-	1.090 pH
L_3	3.50 pH	I_{cq}^+	125 µA
L_4	2.13 pH	I_{cq}^-	125 µA
L_{in}	0.302 pH	β_{cq}^+	1
L_{p1}	0.226 pH	β_{cq}^-	1
L_{pB2}	0.161 pH	I_{cB1}	175 µA
L_{pBc}	0.113 pH	I_{cB2}	275 µA
L_{pBs}	0.145 pH	β_{cB1}	1
L_{pi}^+	0.563 pH	β_{cB2}	$\gg 1$
L_{pi}^-	0.507 pH	I_{cBc}	325 µA
I_{bc}	250 µA	I_{cBs}	250 µA
I_{bs}	175 µA	β_{cBc}	1
I_{ph}^+	0	β_{cBs}	1
I_{ph}^-	−380 µA	I_{tth}	195 µA

A.1.5 MC5 — Complementary quasi-one junction SQUID-based multi-threshold comparator

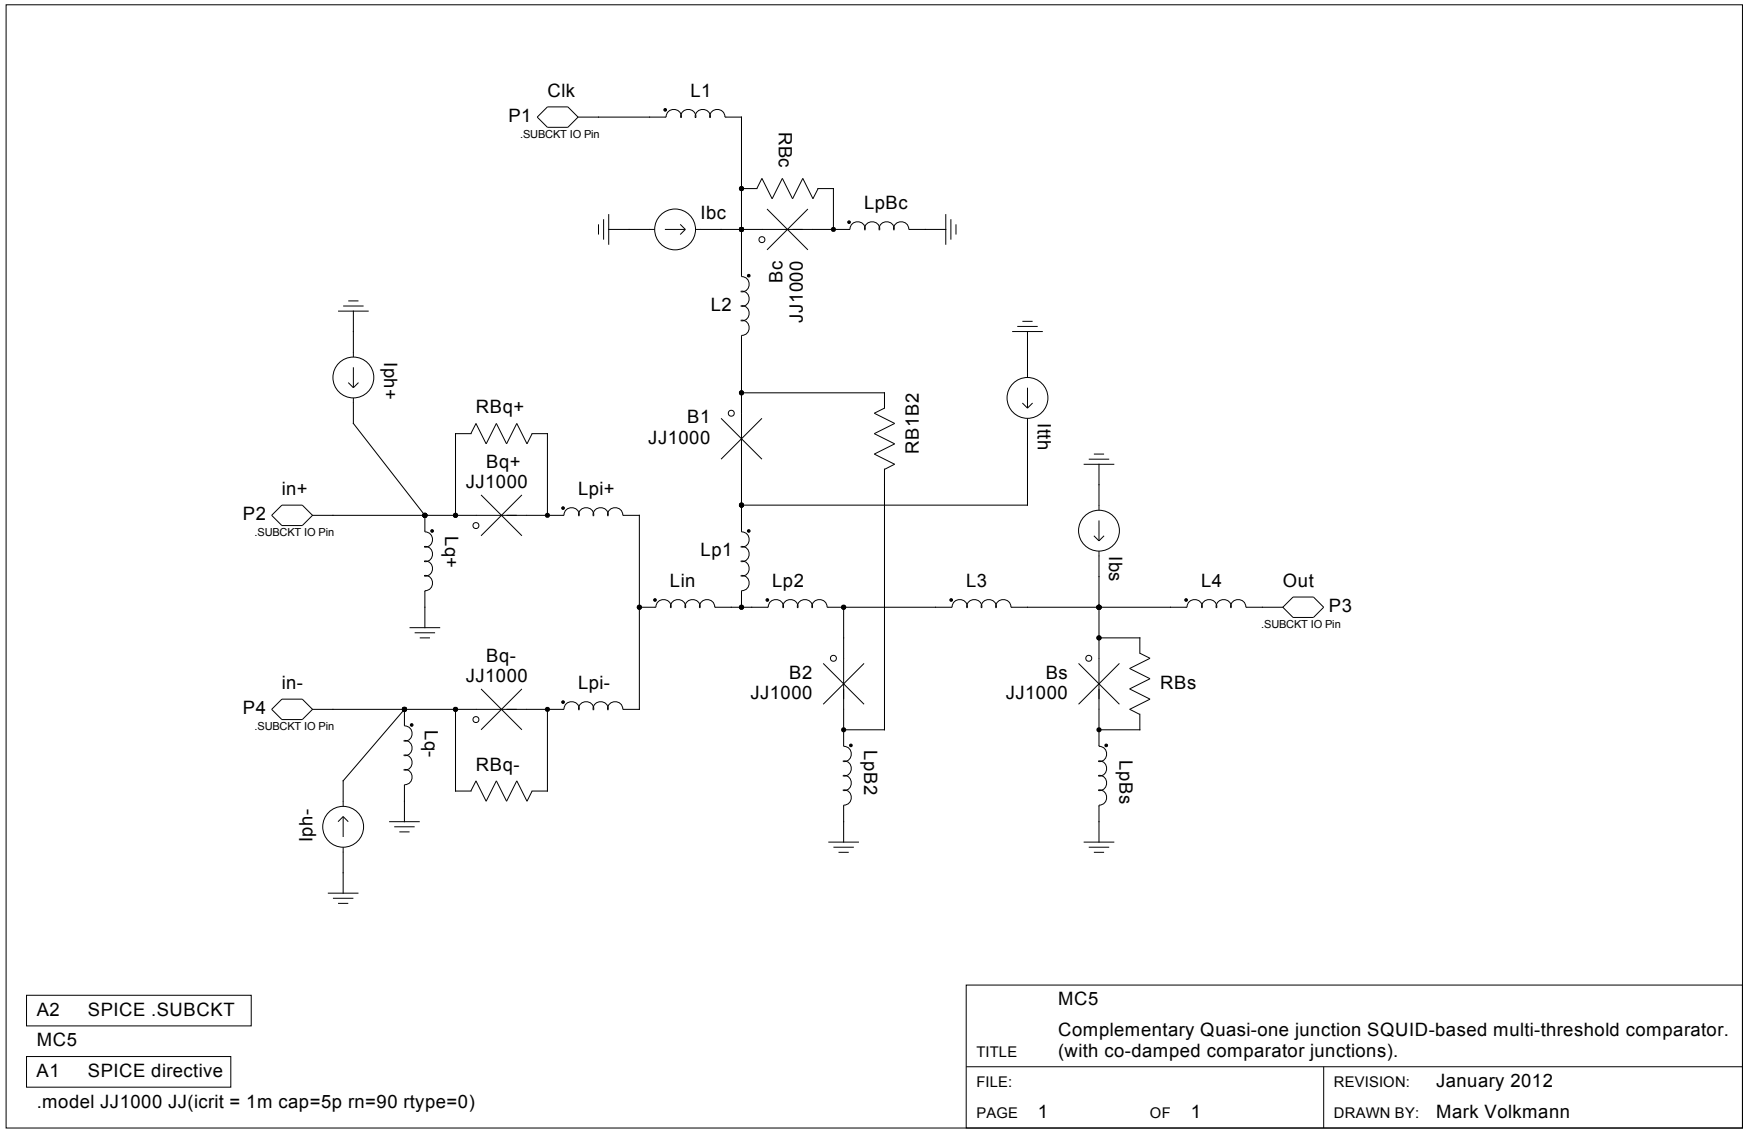

Figure A.9: Final circuit diagram of CQOS-based multi-threshold comparator MC5 (Parameter values in Table A.5).

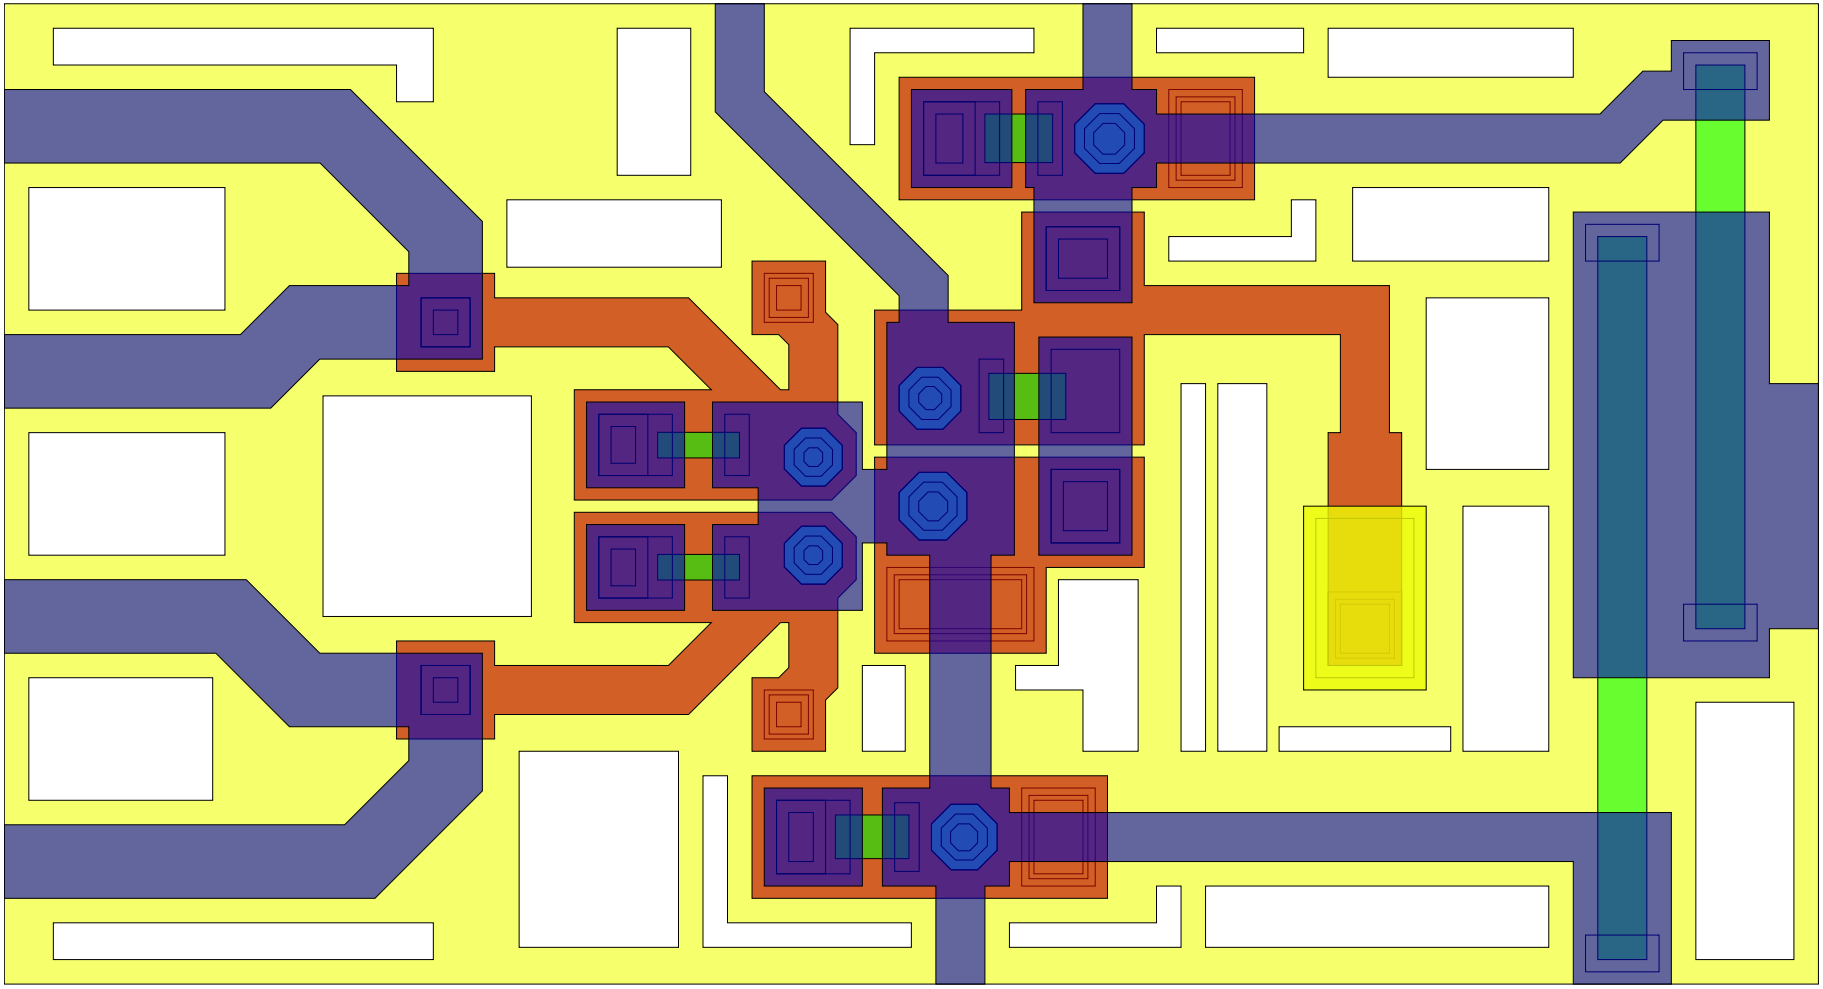

Figure A.10: Final layout of CQOS-based multi-threshold comparator MC5.

Table A.5: Extracted parameters of MC5.

Parameter	**Value**	**Parameter**	**Value**
L_1	1.92 pH	L_q^+	1.089 pH
L_2	1.90 pH	L_q^-	1.092 pH
L_3	3.83 pH	I_{cq}^+	125 µA
L_4	2.19 pH	I_{cq}^-	125 µA
L_{in}	0.773 pH	β_{cq}^+	1
L_{p1}	0.802 pH	β_{cq}^-	1
L_{p2}	0.246 pH	I_{cB1}	175 µA
L_{pB2}	0.033 pH	I_{cB2}	275 µA
L_{pBc}	0.117 pH	β_{cB1}	co-damped
L_{pBs}	0.126 pH	β_{cB2}	co-damped
L_{pi}^+	0.591 pH	I_{cBc}	325 µA
L_{pi}^-	0.563 pH	I_{cBs}	250 µA
I_{bc}	250 µA	β_{cBc}	1
I_{bs}	175 µA	β_{cBs}	1
I_{ph}^+	0	I_{tth}	195 µA
I_{ph}^-	−390 µA		

A.1.6 MX — Improved QOS comparator

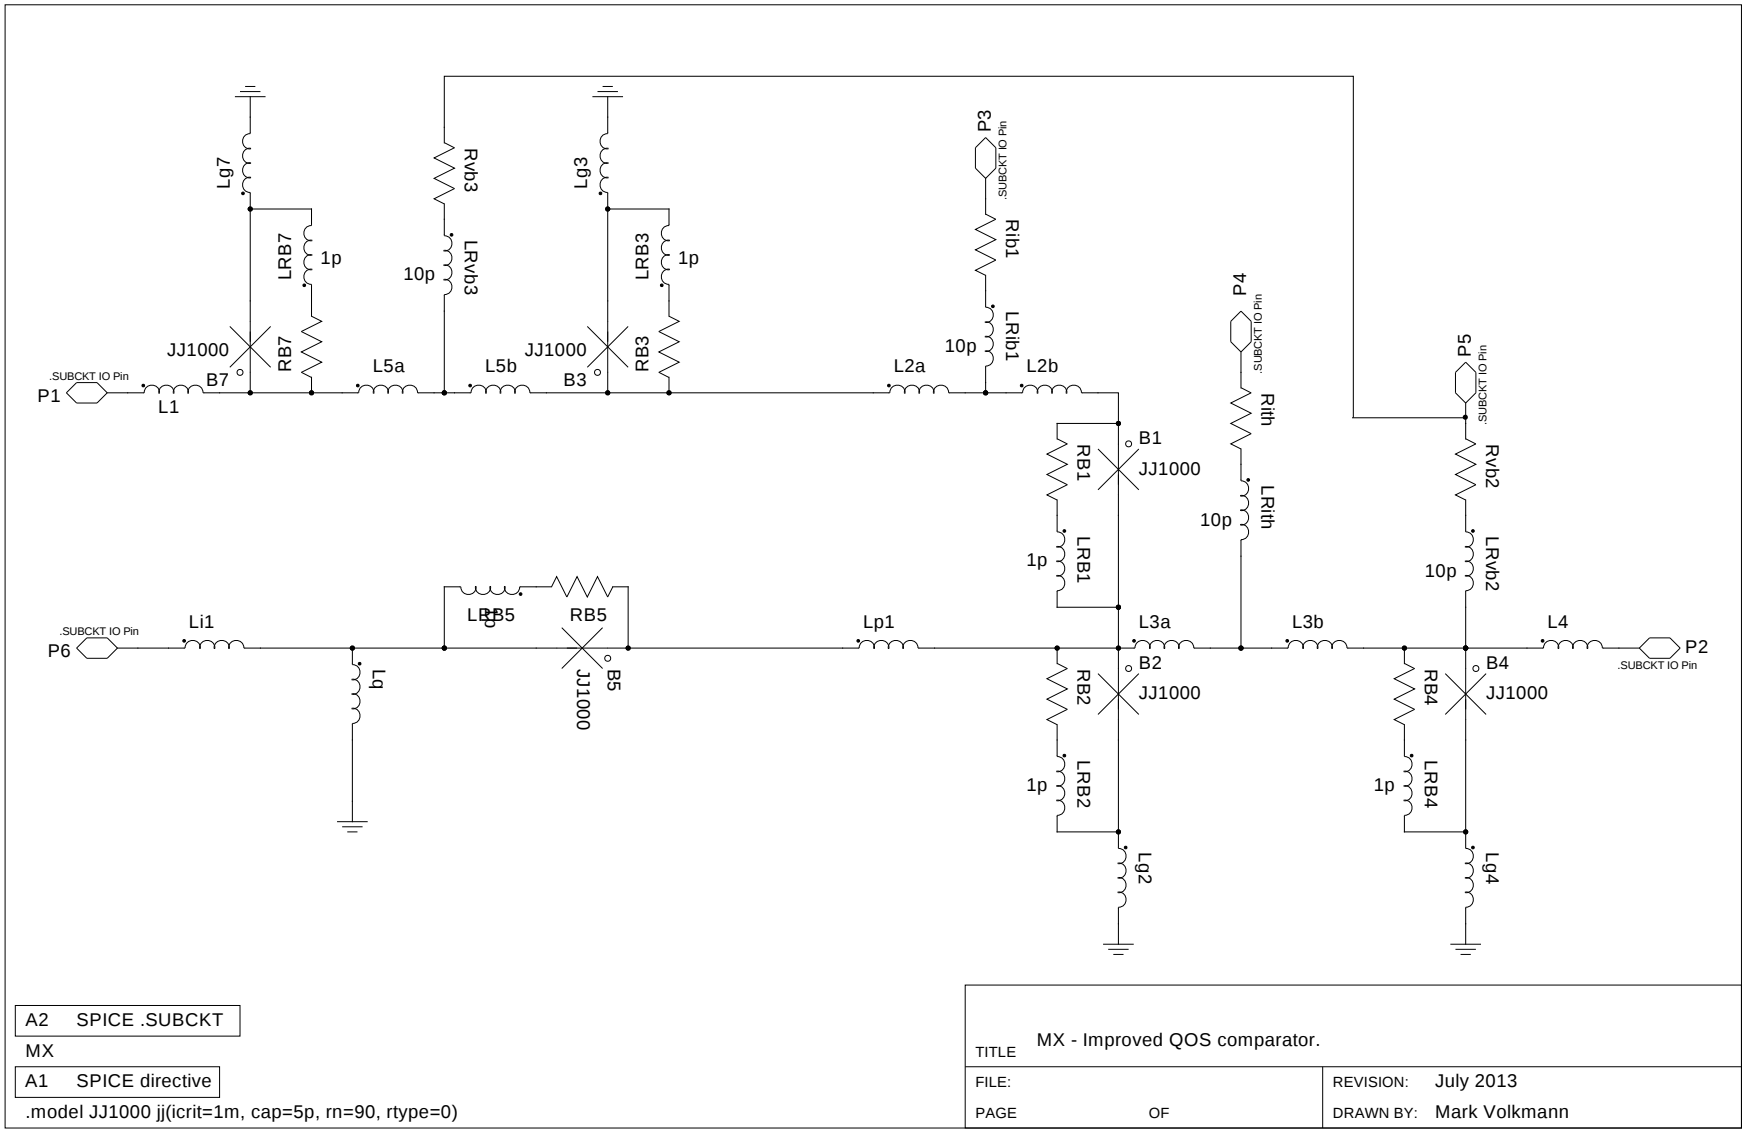

Figure A.11: Final circuit diagram of QOS-based multi-threshold comparator `MX` (Parameter values in Table A.6).

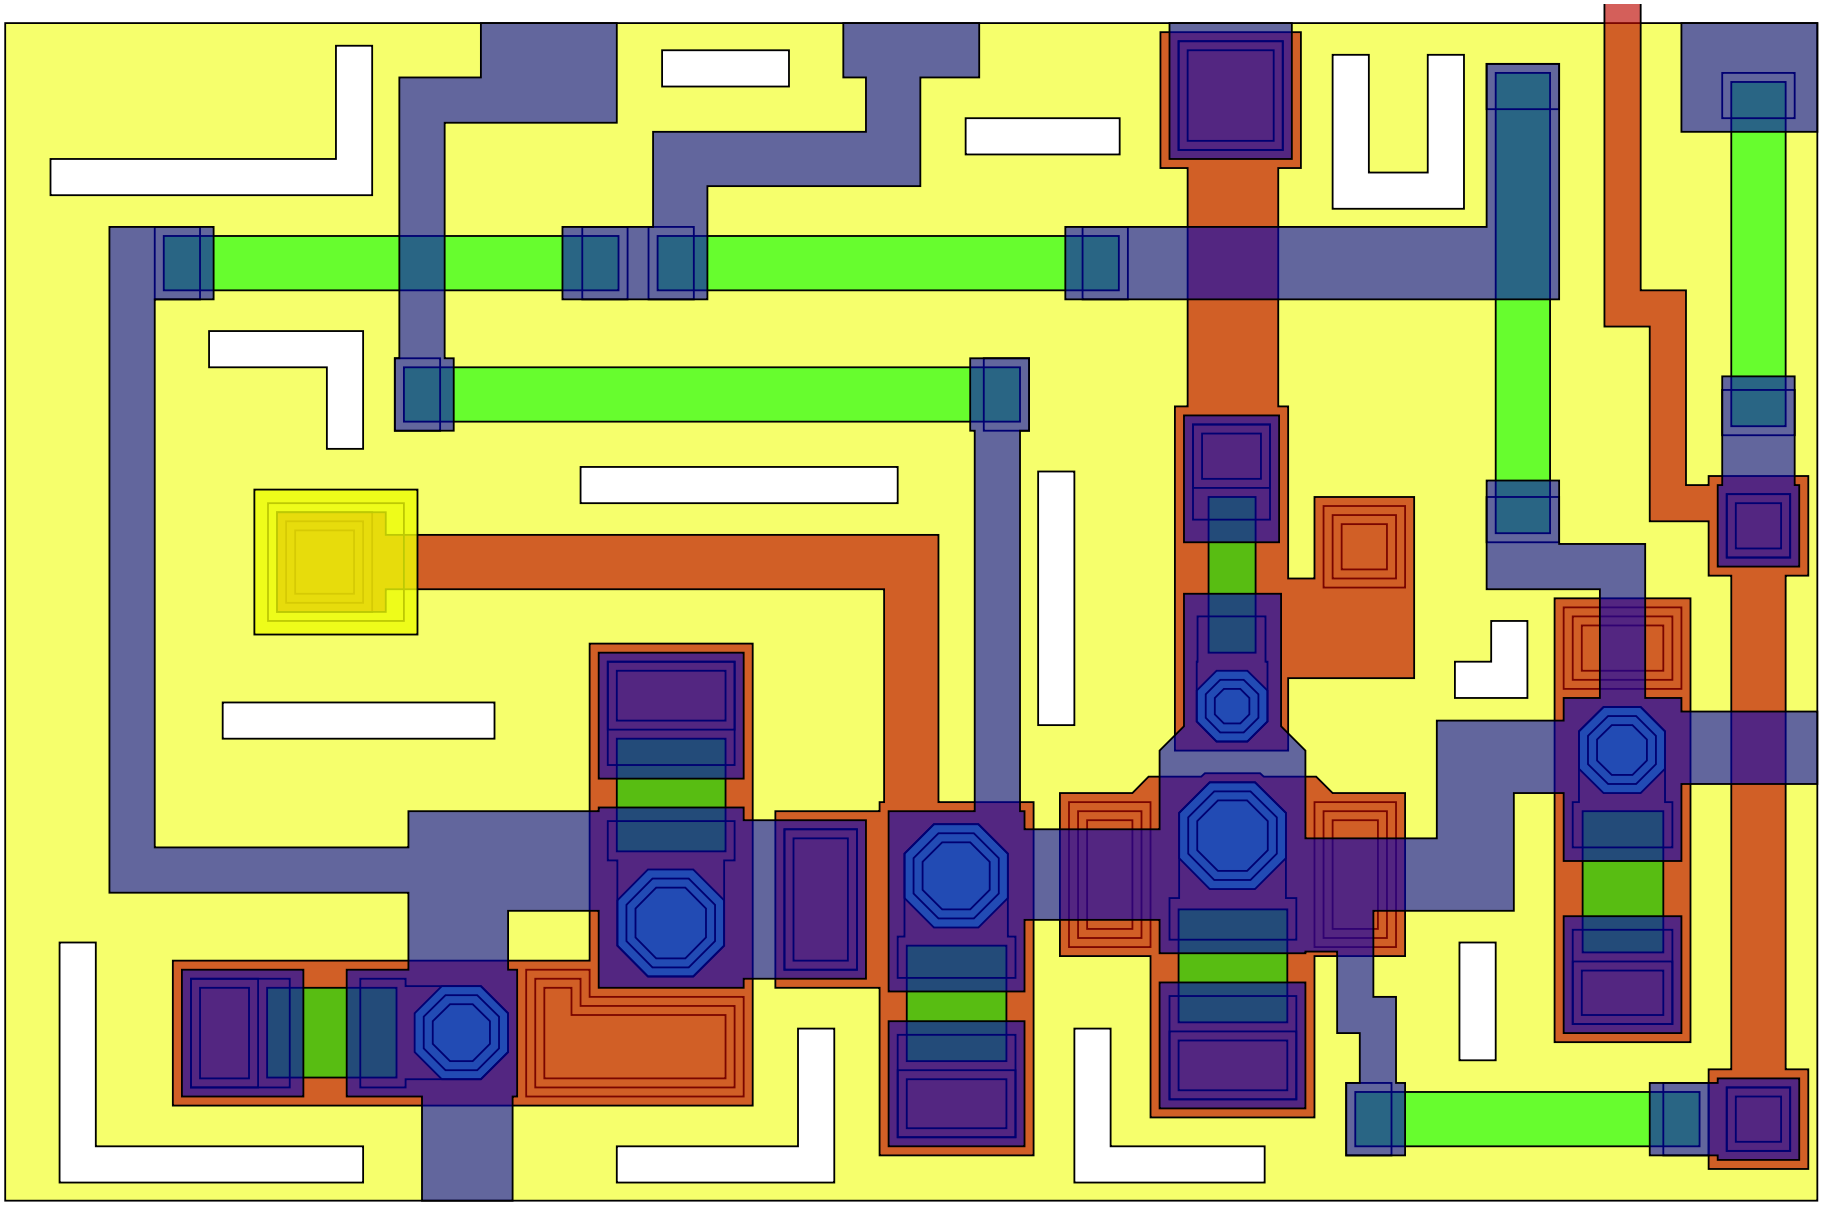

Figure A.12: Final layout of QOS-based multi-threshold comparator `MX` (Parameter values in Table A.6).

A.1.7 MCX — Improved CQOS comparator

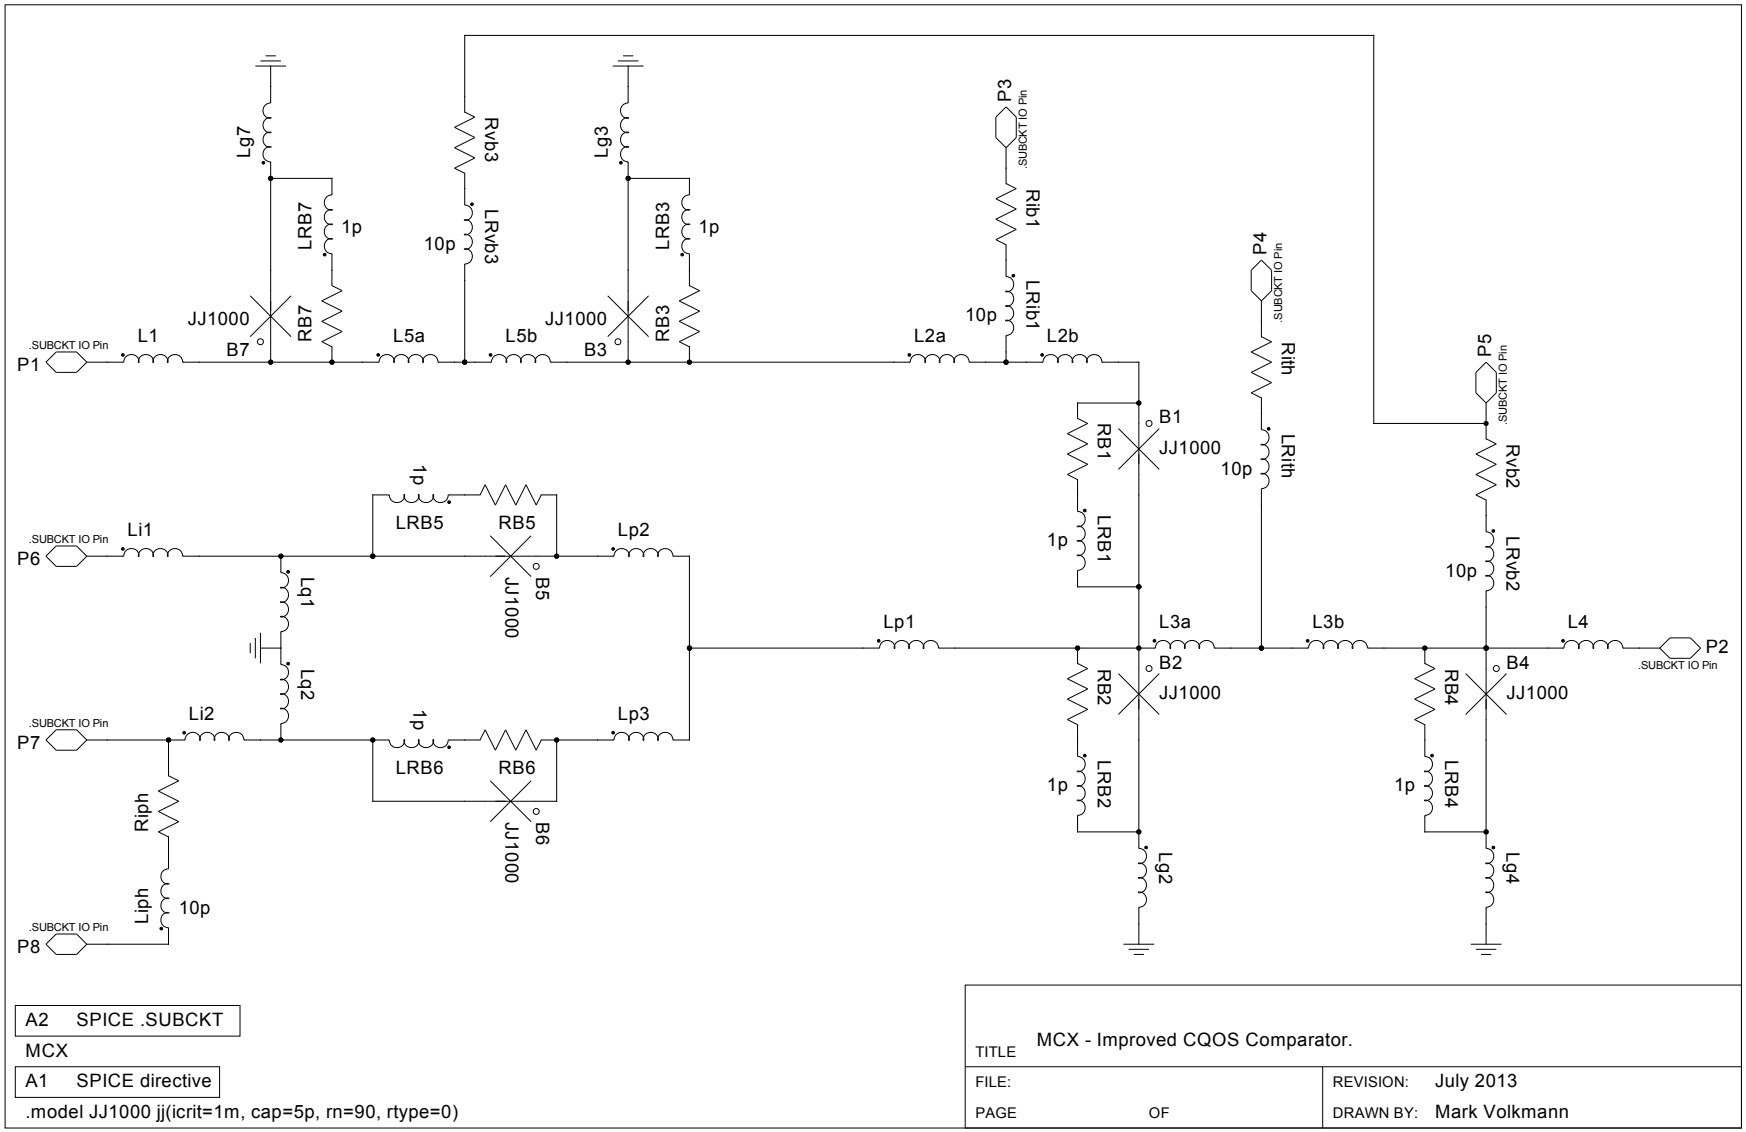

Figure A.13: Final circuit diagram of CQOS-based multi-threshold comparator `MCX` (Parameter values in Table A.7).

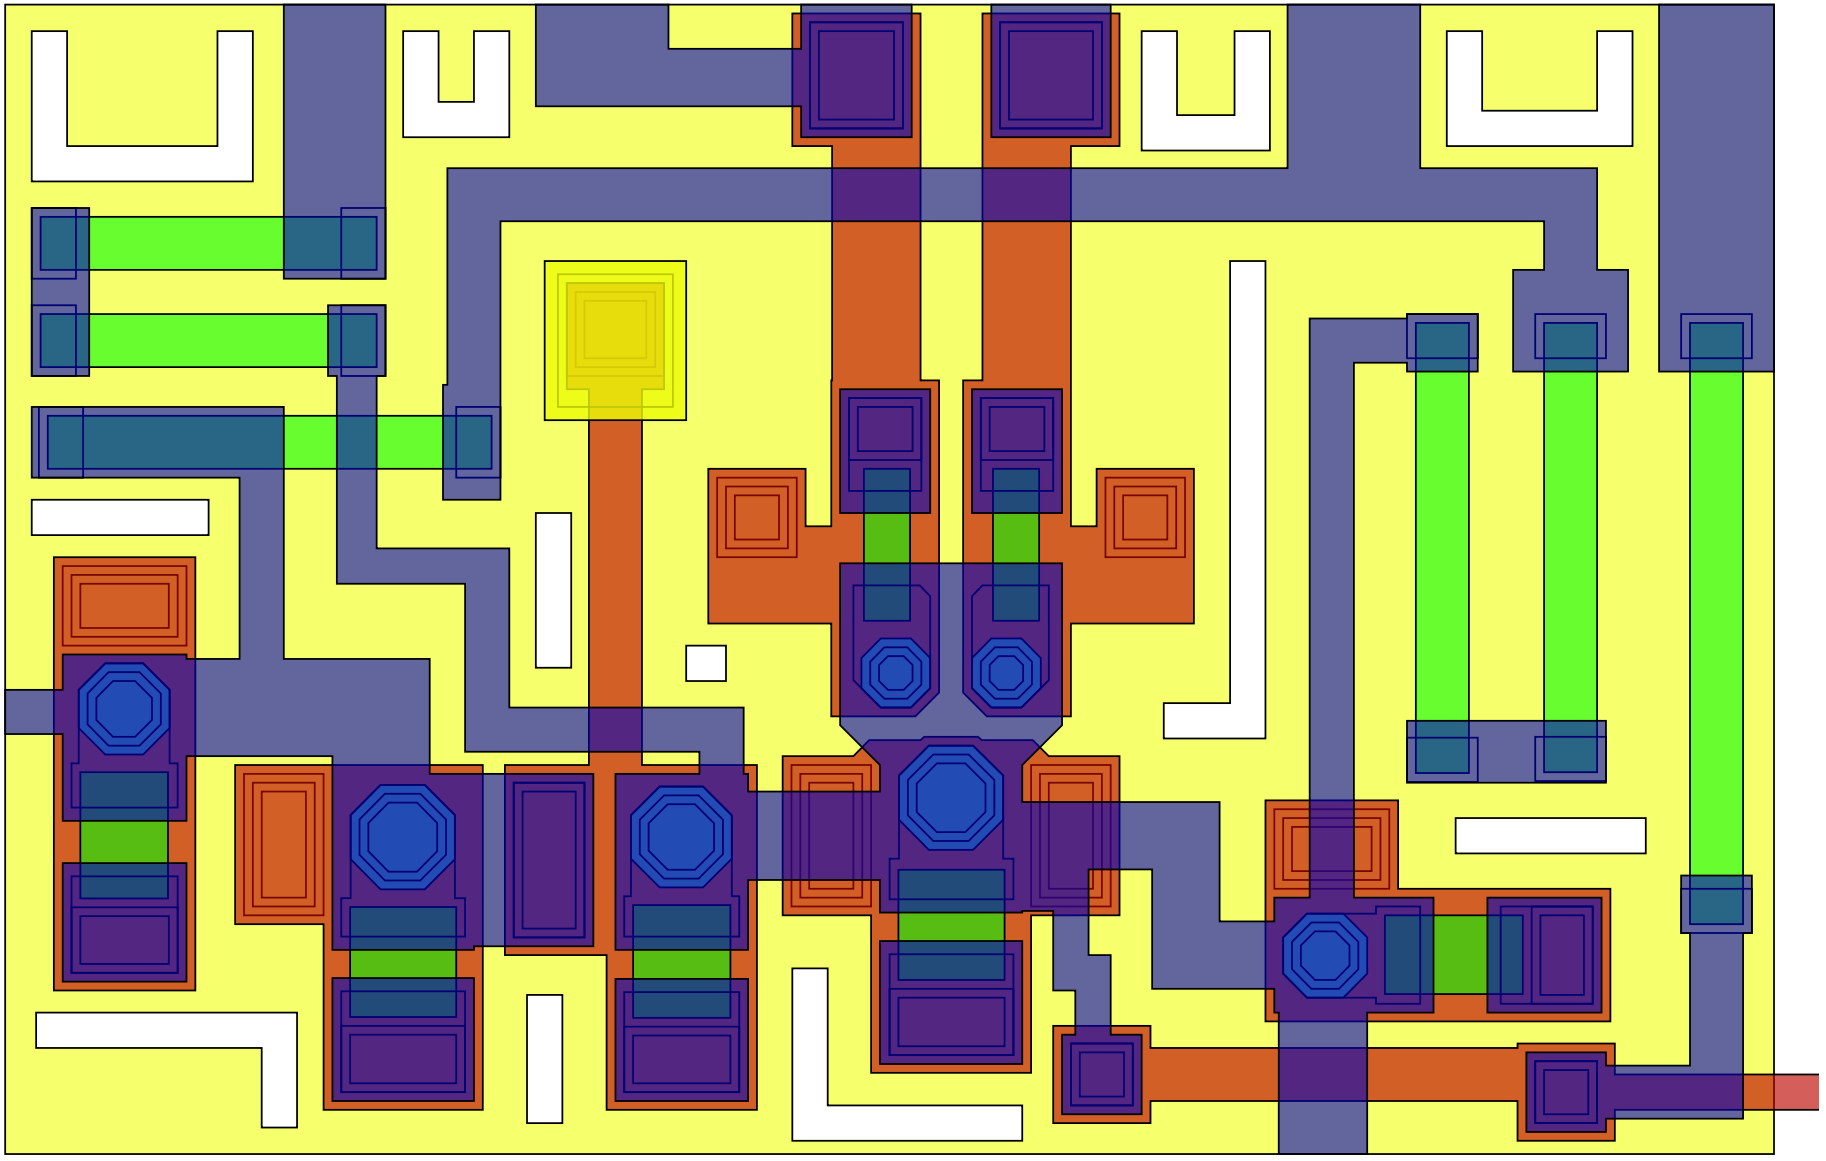

Figure A.14: Final layout of CQOS-based multi-threshold comparator `MCX` (Parameter values in Table A.7).

Table A.6: Extracted and nominal parameters of MX.

Parameter	Value	Parameter	Value
L_{2a}	1.36 pH	L_{2b}	1.66 pH
L_{3a}	0.44 pH	L_{3b}	3.60 pH
L_{5a}	0.95 pH	L_{5b}	1.55 pH
L_{p1}	0.83 pH	L_{g2}	0.20 pH
L_{g3}	0.21 pH	L_{g4}	0.22 pH
L_{g7}	0.19 pH	L_q	0.80 pH
L_1	1.44 pH	L_4	1.88 pH
R_{vb2}	14.3 Ω	I_{cB5}	125 µA
R_{vb3}	7.03 Ω	I_{cB7}	325 µA
I_{cB3}	500 µA	I_{cB4}	250 µA
I_{cB1}	450 µA	I_{cB2}	500 µA
β_{c*}	1		

Table A.7: Extracted and nominal parameters of MCX.

Parameter	Value	Parameter	Value
L_{2a}	1.34 pH	L_{2b}	1.64 pH
L_{3a}	0.33 pH	L_{3b}	3.81 pH
L_{5a}	0.97 pH	L_{5b}	1.53 pH
L_{p1}	0.51 pH	L_{g2}	0.21 pH
L_{g3}	0.23 pH	L_{g4}	0.27 pH
L_{g7}	0.19 pH	L_{q1}	0.79 pH
L_{q2}	0.79 pH	L_{p2}	0.44 pH
L_1	1.43 pH	L_4	2.01 pH
L_{p3}	0.44 pH	I_{cB5}	125 pH
R_{vb2}	14.3 Ω	I_{cB6}	125 µA
R_{vb3}	7.03 Ω	I_{cB7}	325 µA
I_{cB3}	500 µA	I_{cB4}	250 µA
I_{cB1}	450 µA	I_{cB2}	500 µA
β_{c*}	1		

A.2 Test-beds

A.2.1 MC-GREY — Grey comparator test-bed

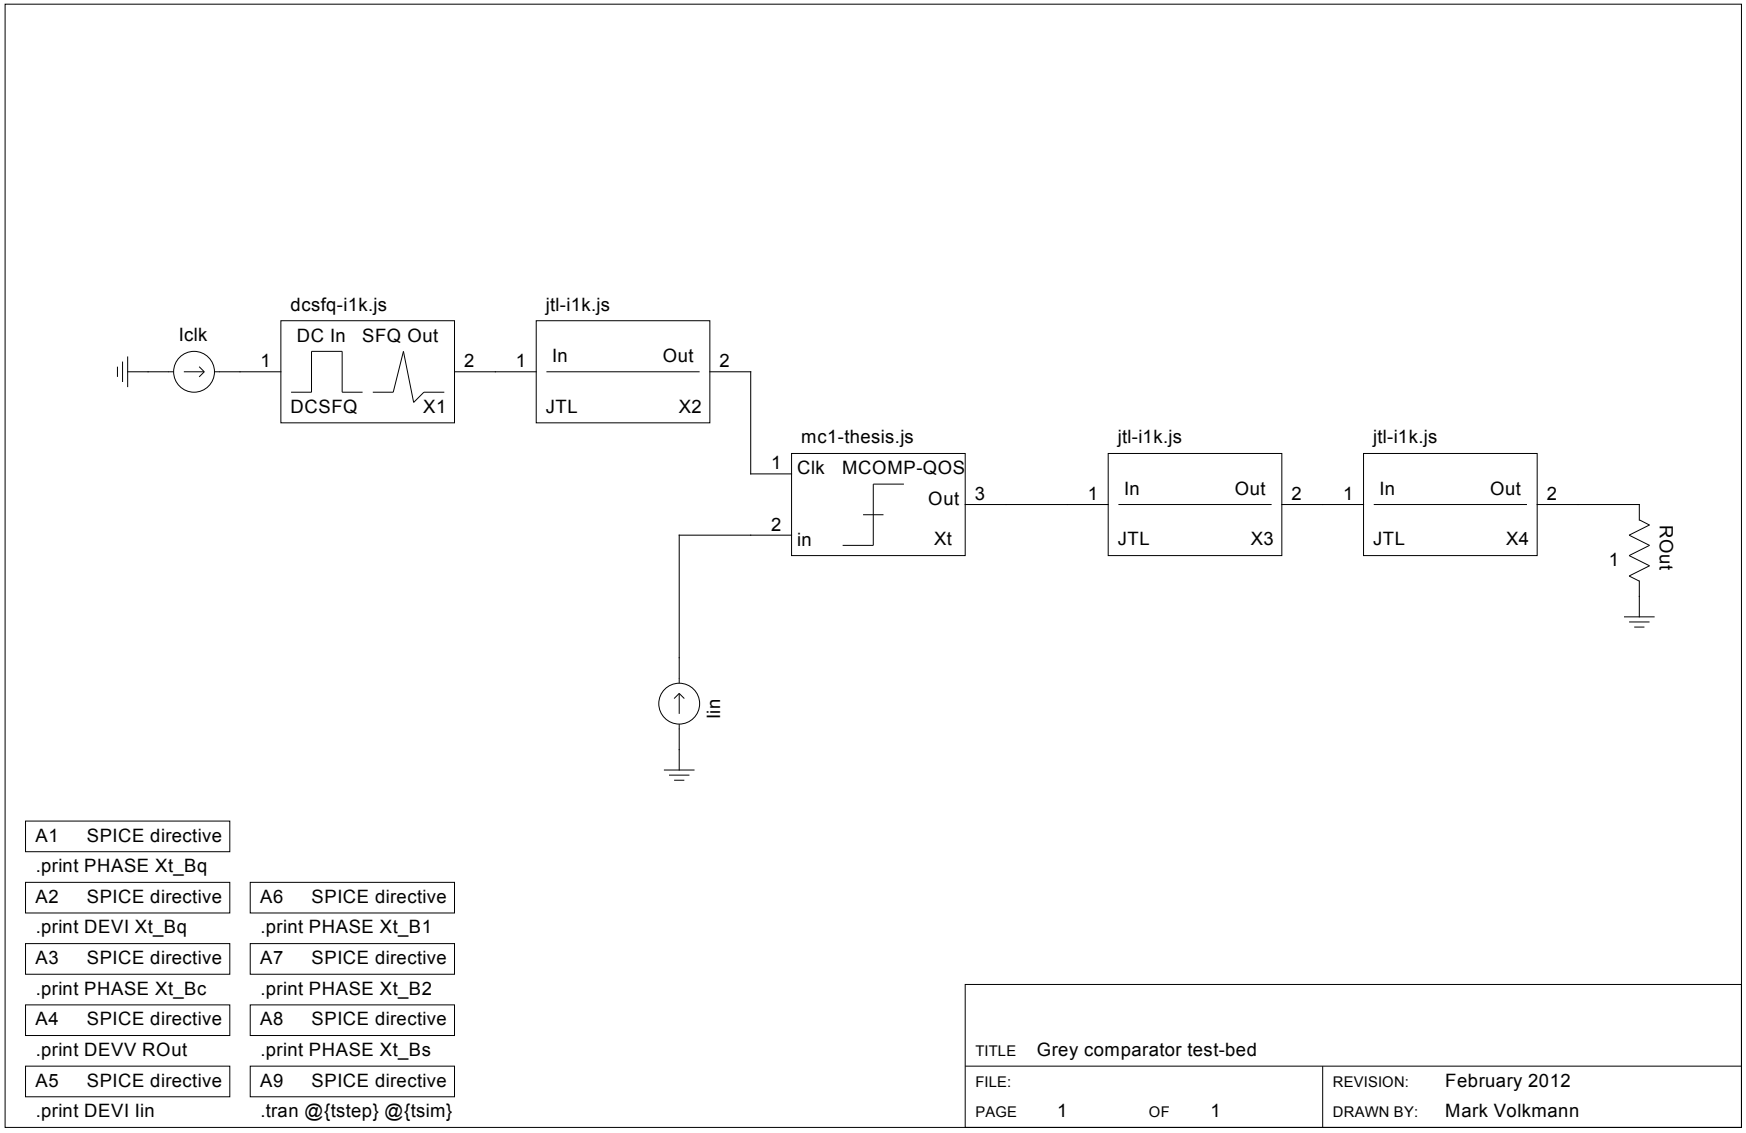

A.3 Comparator chip 1

A.3.1 mvv2-gz-inv — Test structures for grey-zone investigation.

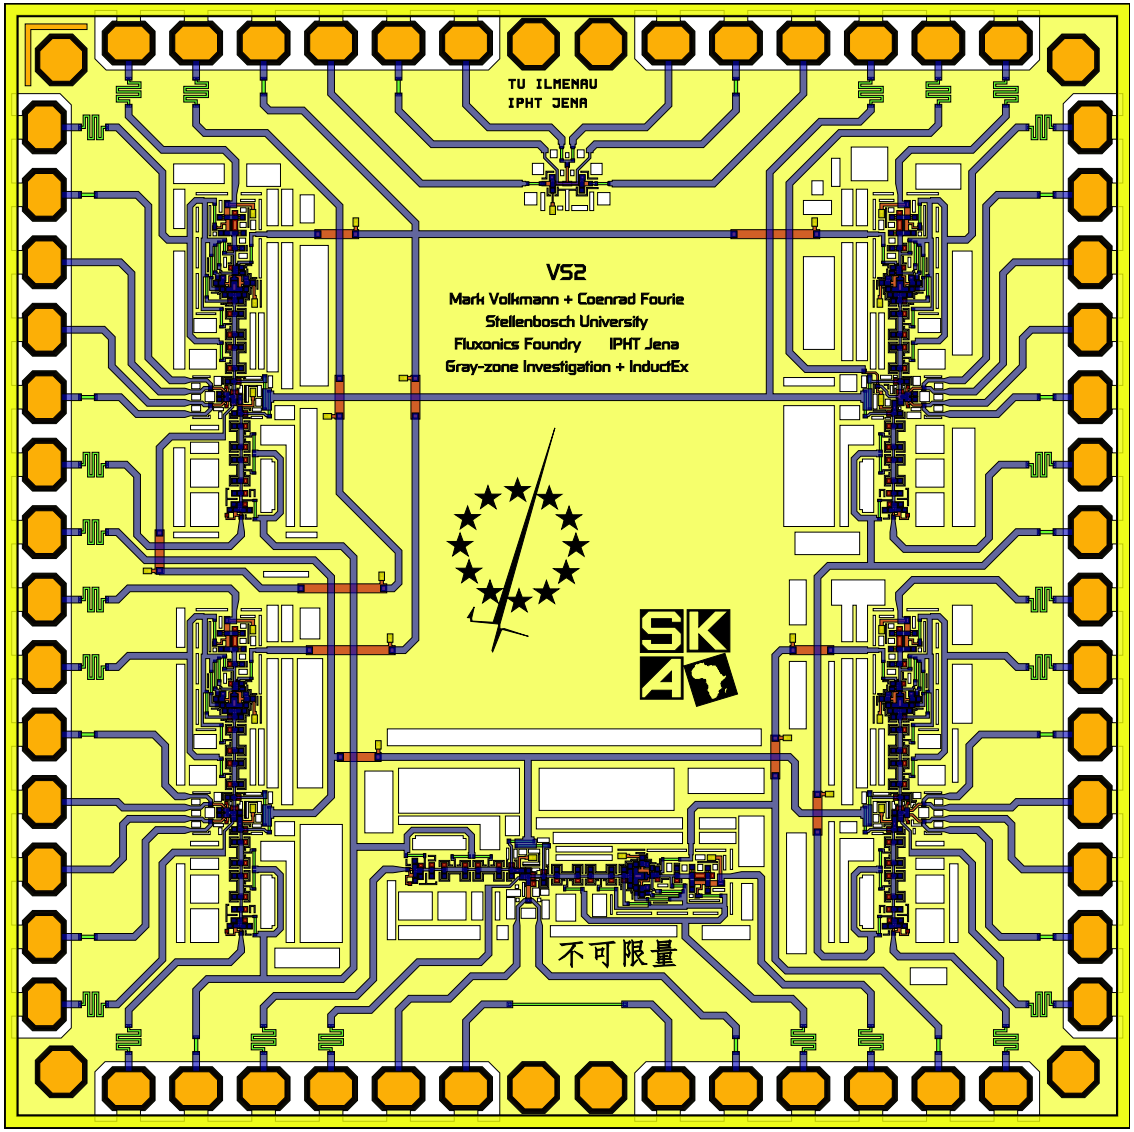

Figure A.16: Chip mvv2-gz-inv, to experimentally verify grey-zones of (complementary) QOS comparators.

A.4 Comparator chip 2

A.4.1 `mvw4-gz-c1-inv` — Test structures for investigation at high speed.

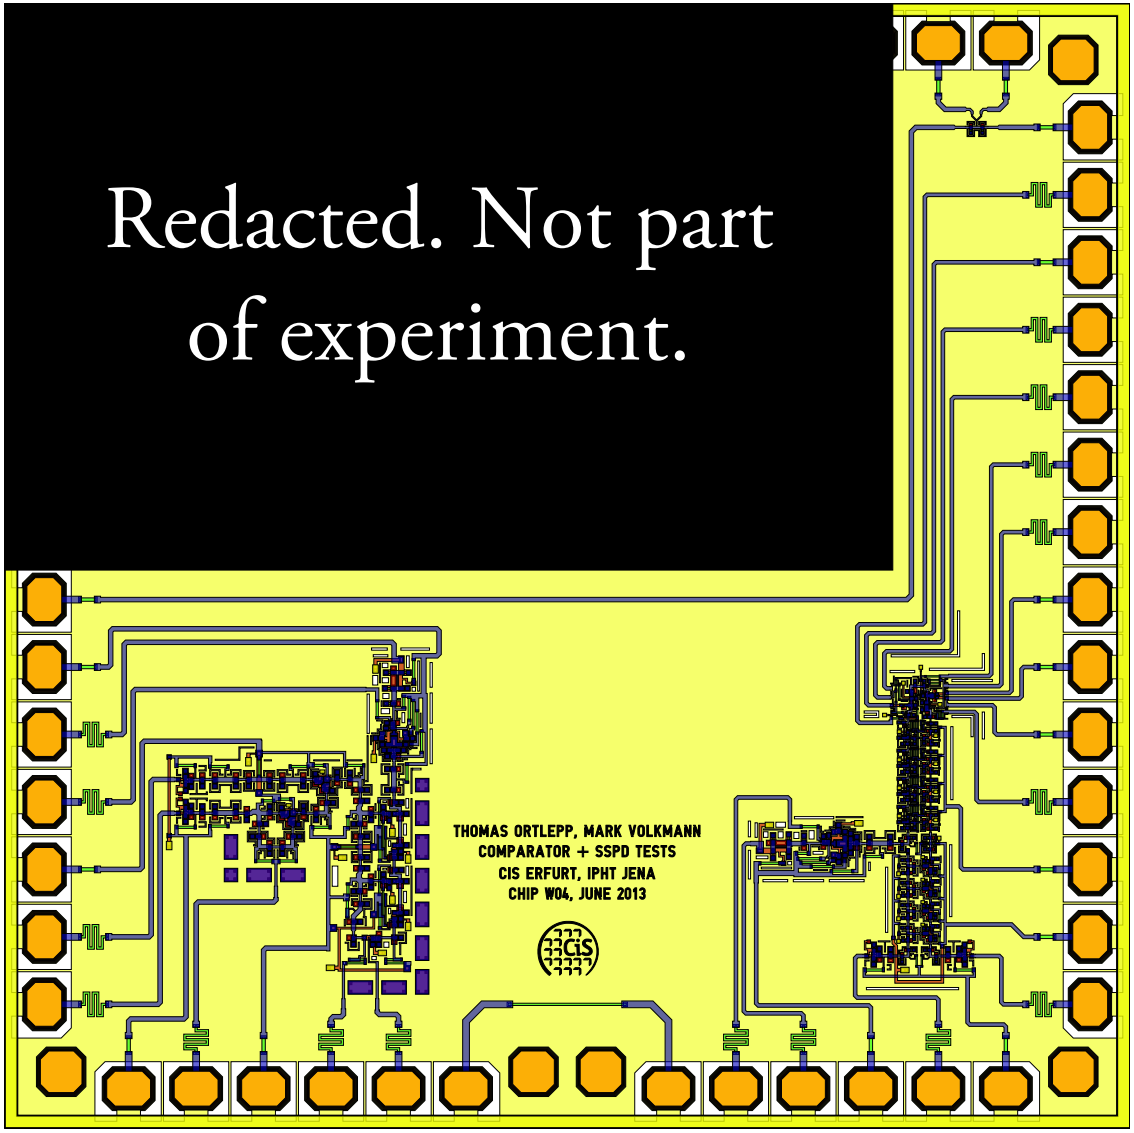

Figure A.17: Chip `mvw4-gz-c1-inv`, to experimentally verify grey-zones and correlation of (complementary) QOS comparators.

A.5 eSFQ cells

A.5.1 eSR — Shift register cell

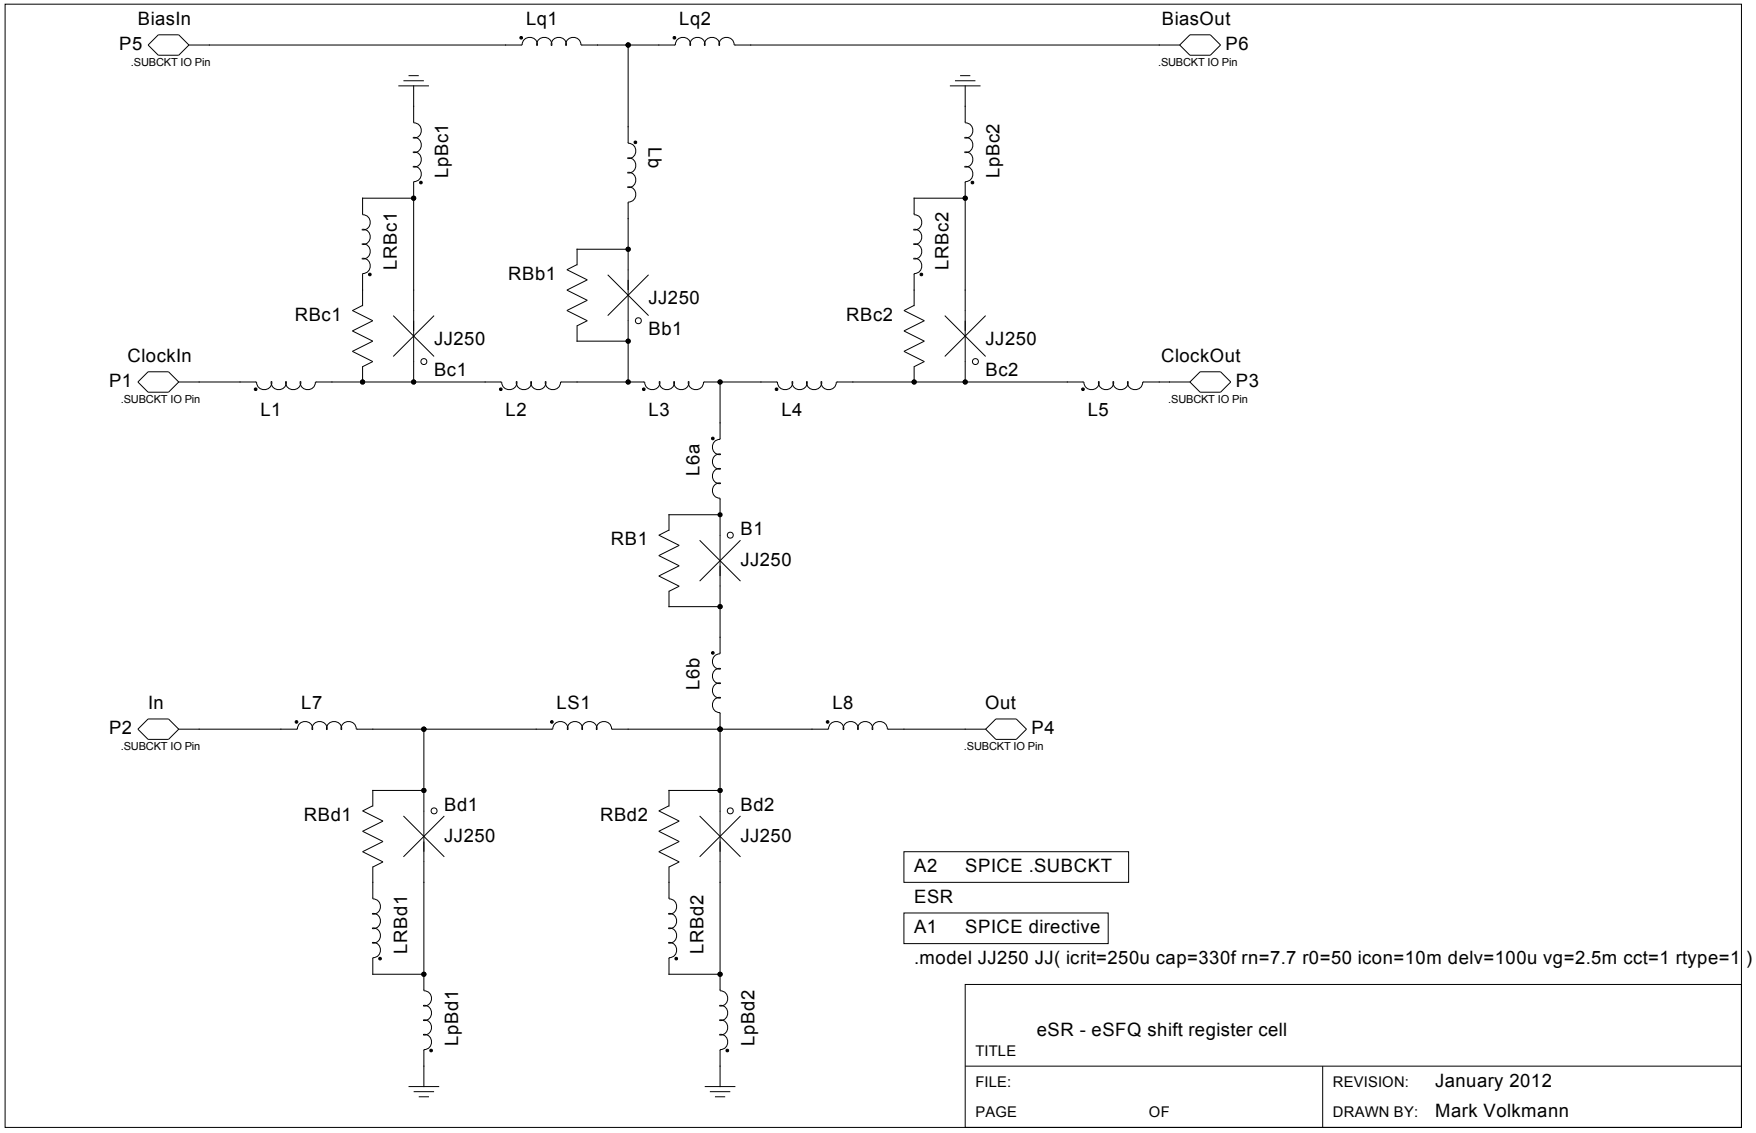

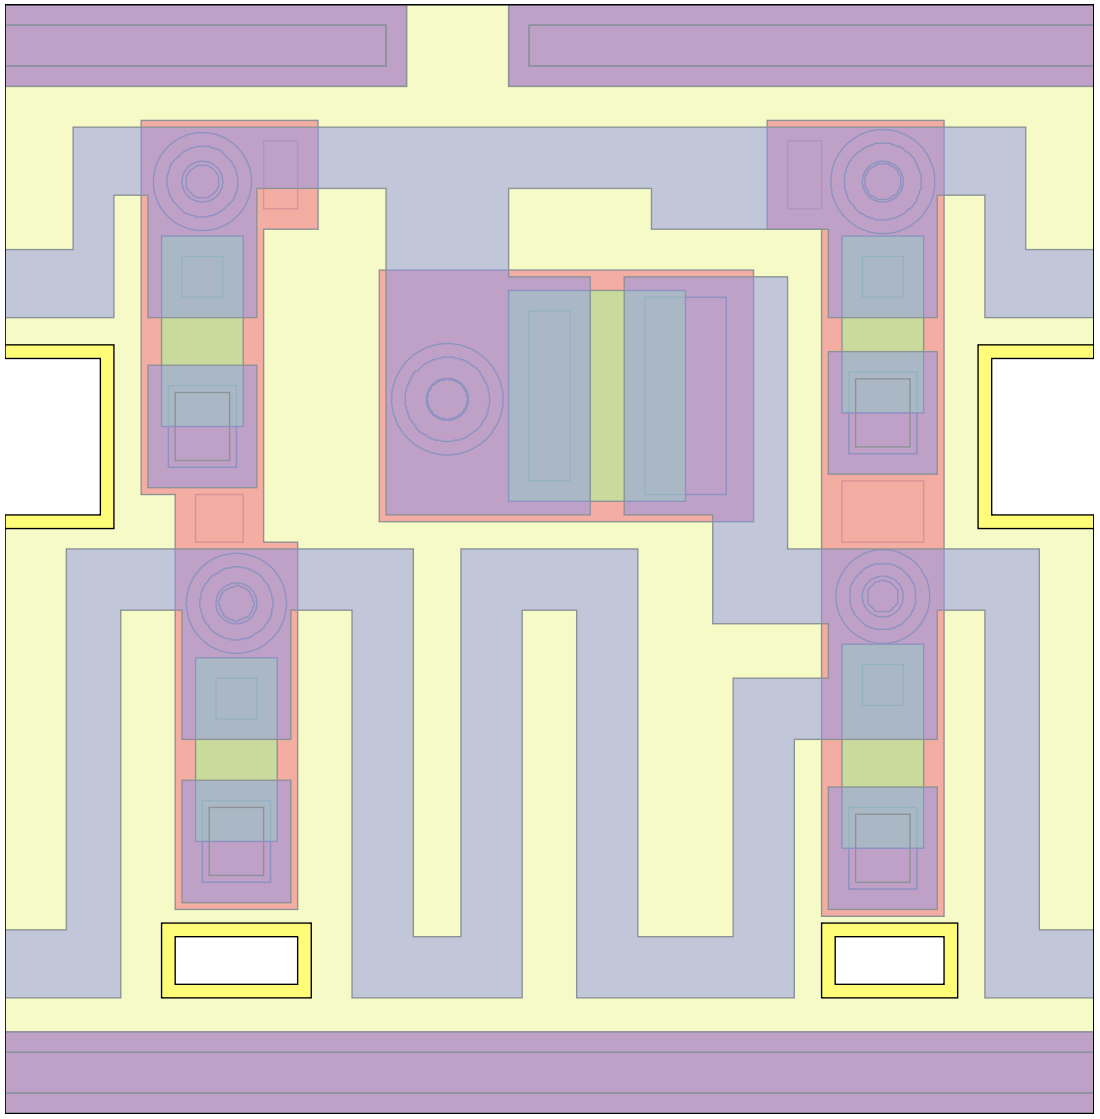

Figure A.19: Final layout of eSR shift register cell (Parameter values in Table A.8).

A.5.2 MeSR — Shift register cell with magnetic flux bias

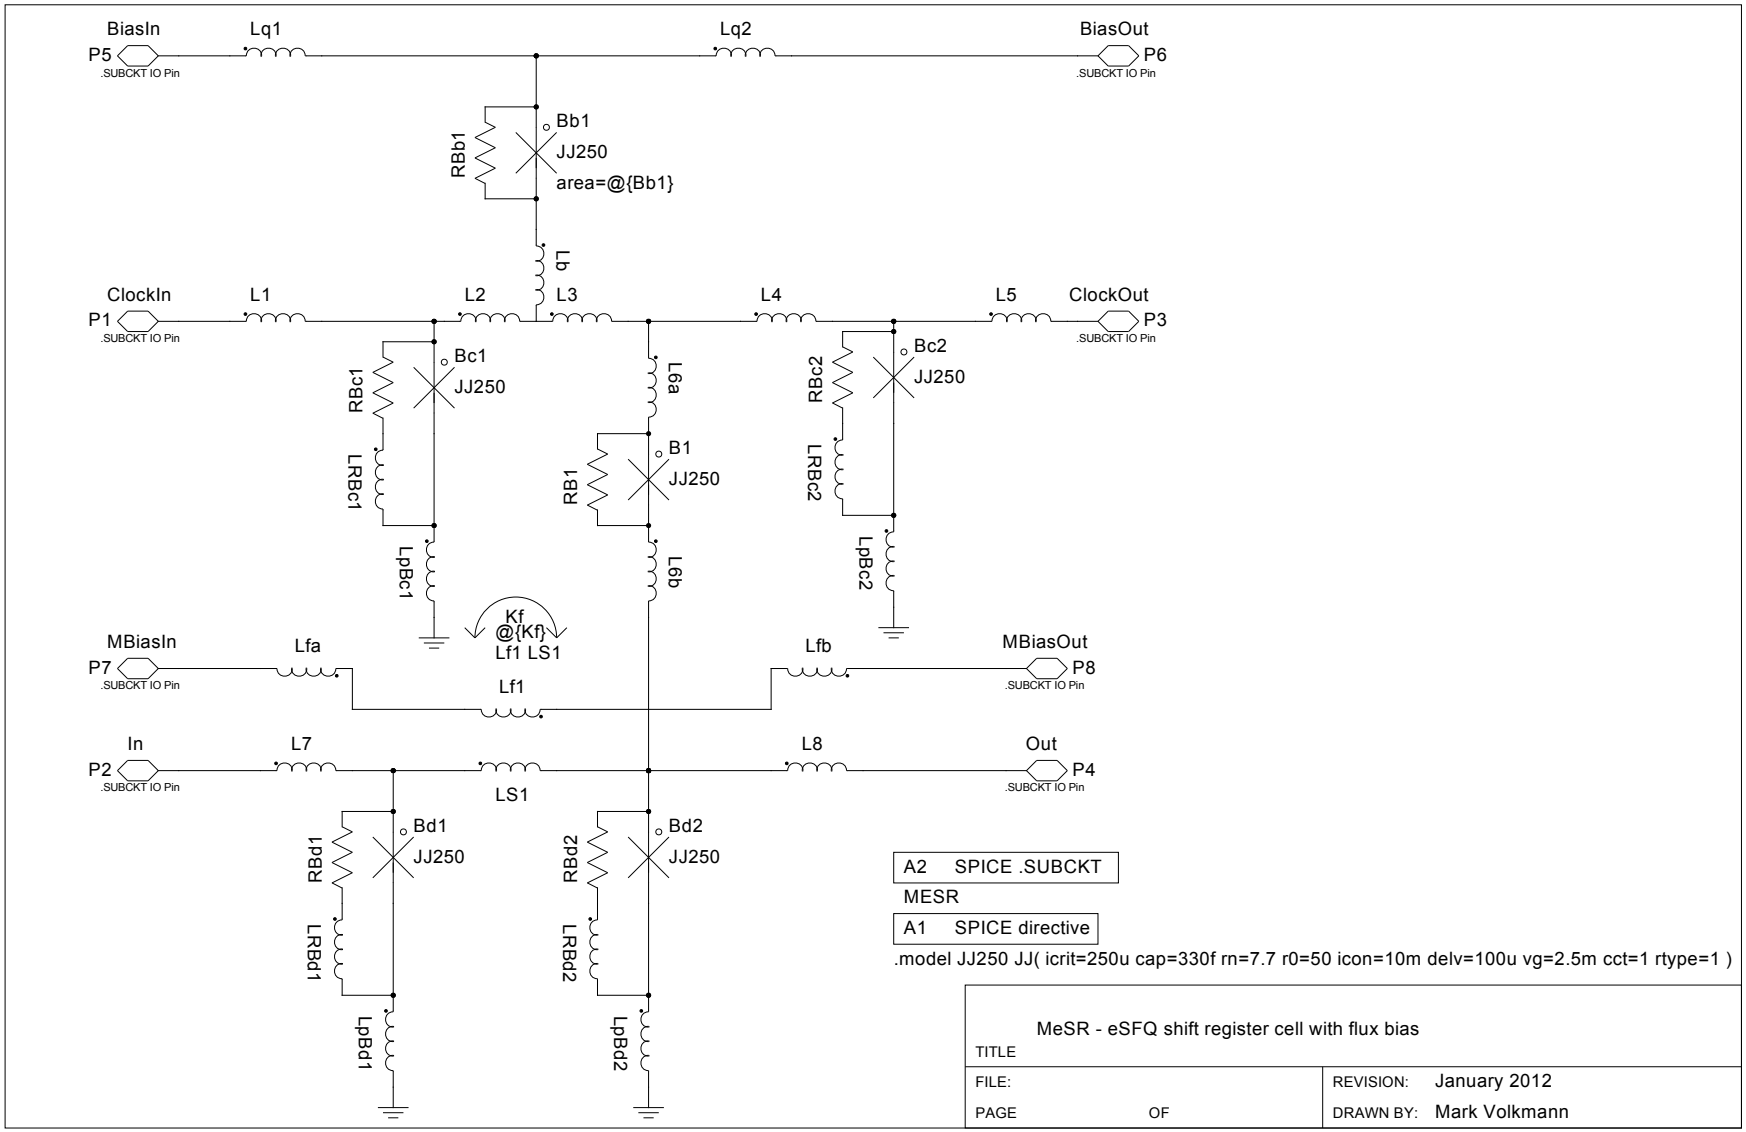

Figure A.20: Final circuit diagram of MeSR shift register cell (Parameter values in Table A.9).

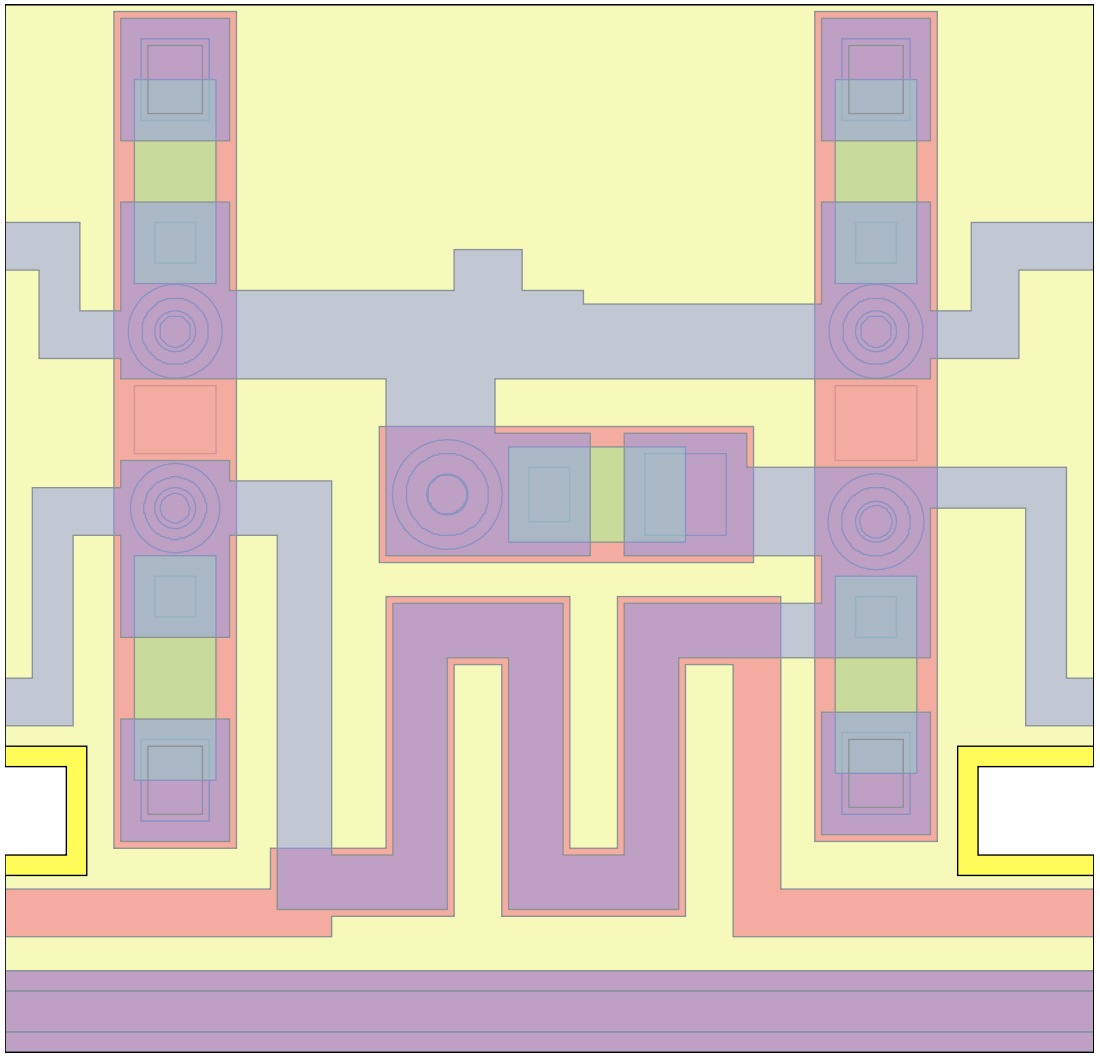

Figure A.21: Final layout of MeSR shift register cell (Parameter values in Table A.9).

Table A.8: Extracted (inductances) and nominal parameters of eSR.

Parameter	Value	Parameter	Value
L_1	2.19 pH	β_{cB1}	≈ 0.25
L_2	1.17 pH	β_{c*}	≈ 1
L_3	0.09 pH	L_{pBc1}	0.125 pH
L_4	1.91 pH	L_{pBc2}	0.125 pH
L_5	2.06 pH	L_{pBd1}	0.125 pH
L_6	2.1 pH	L_{pBd2}	0.182 pH
L_7	3.5 pH	L_{RB*}	~ 1 pH
L_8	3.4 pH	I_{cBb1}	575 µA
L_{S1}	10.8 pH	I_{cBd1}	188 µA
I_{cBc1}	213 µA	I_{cBd2}	225 µA
I_{cBc2}	250 µA	I_{cB1}	313 µA

Table A.9: Extracted (inductances and coupling factor) and nominal parameters of MeSR.

Parameter	Value	Parameter	Value
L_1	1.96 pH	β_{cB1}	≈ 0.25
L_2	1.02 pH	β_{c*}	≈ 1
L_3	0.16 pH	L_{pBc1}	0.114 pH
L_4	1.81 pH	L_{pBc2}	0.121 pH
L_5	2.16 pH	L_{pBd1}	0.409 pH
L_6	1.7 pH	L_{pBd2}	0.193 pH
L_7	2.84 pH	L_{RB*}	~ 1 pH
L_8	4.16 pH	I_{cBb1}	525 µA
L_{S1}	11.2 pH	I_{cBd1}	163 µA
I_{cBc1}	188 µA	I_{cBd2}	200 µA
I_{cBc2}	188 µA	I_{cB1}	288 µA
L_{f1}	11.5 pH	k_f	0.25

A.5.3 eDES — Deserialiser cell

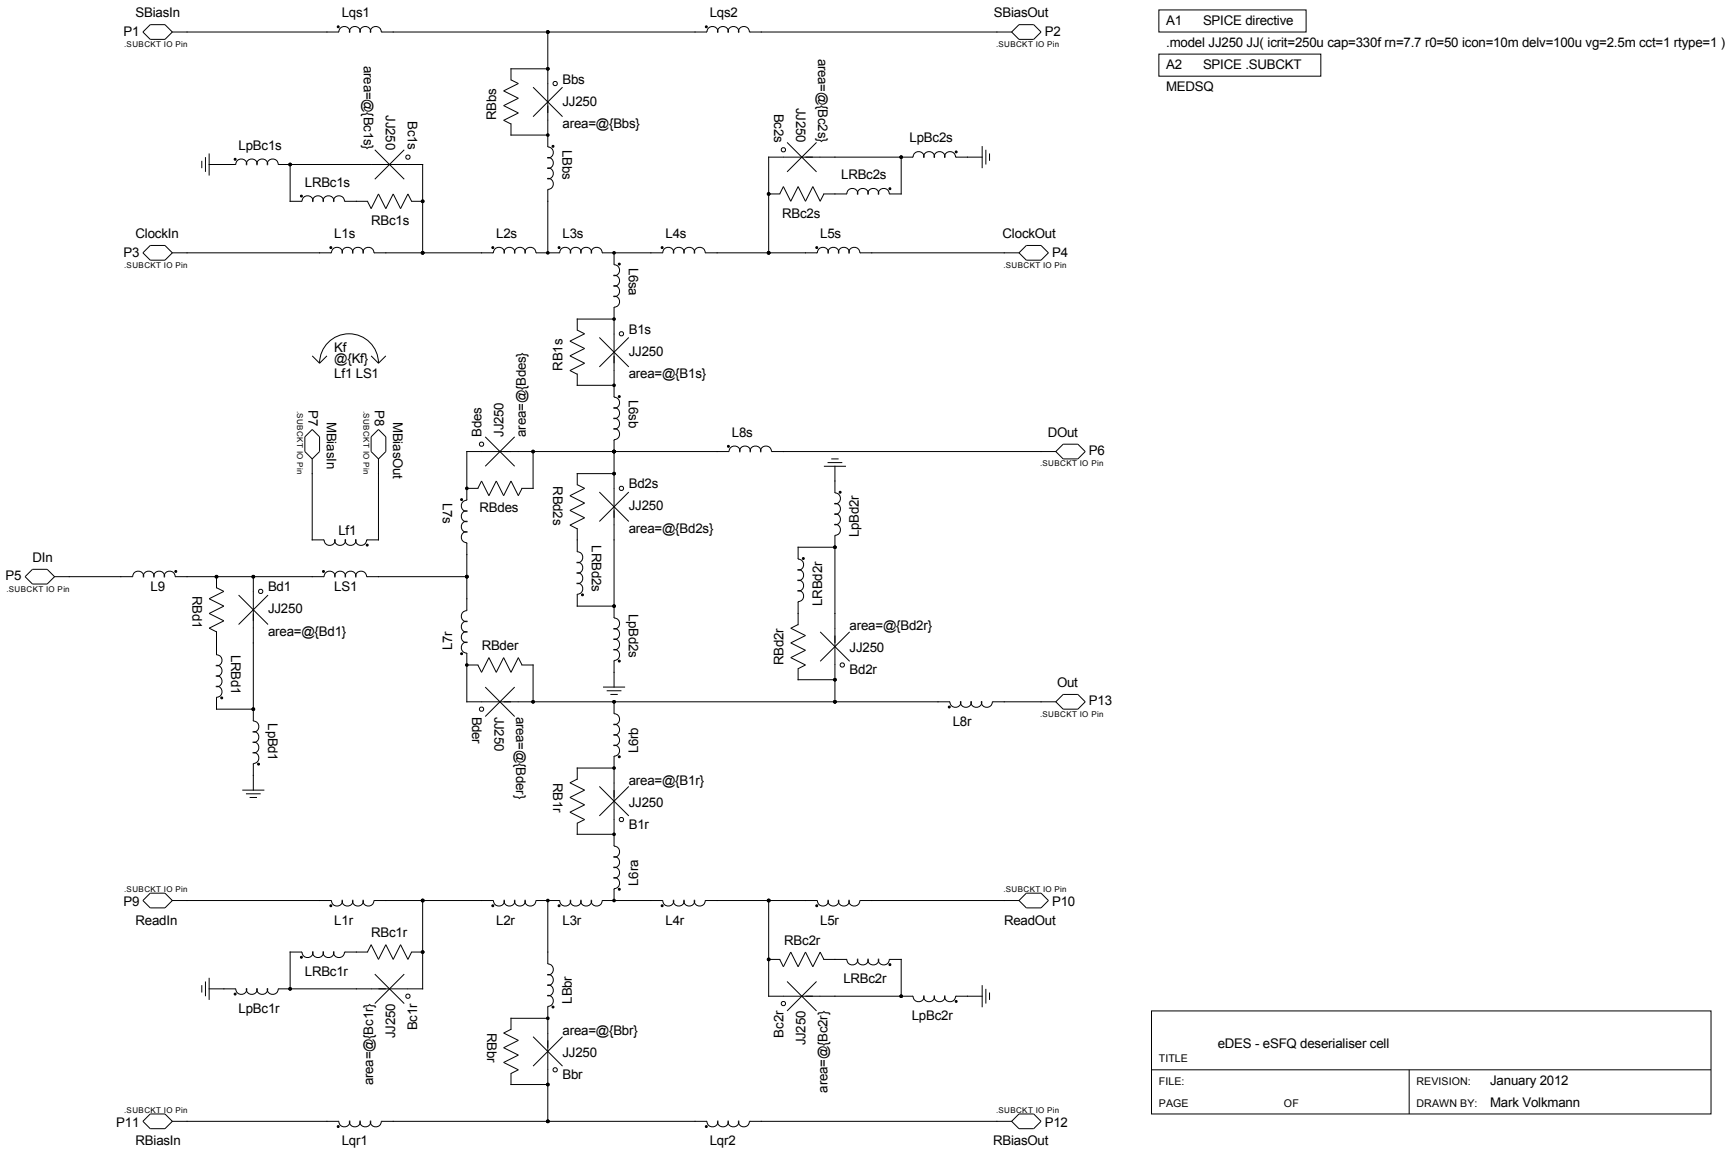

Figure A.22: Final circuit diagram of eDES deserialiser cell (Parameter values in Table A.10).

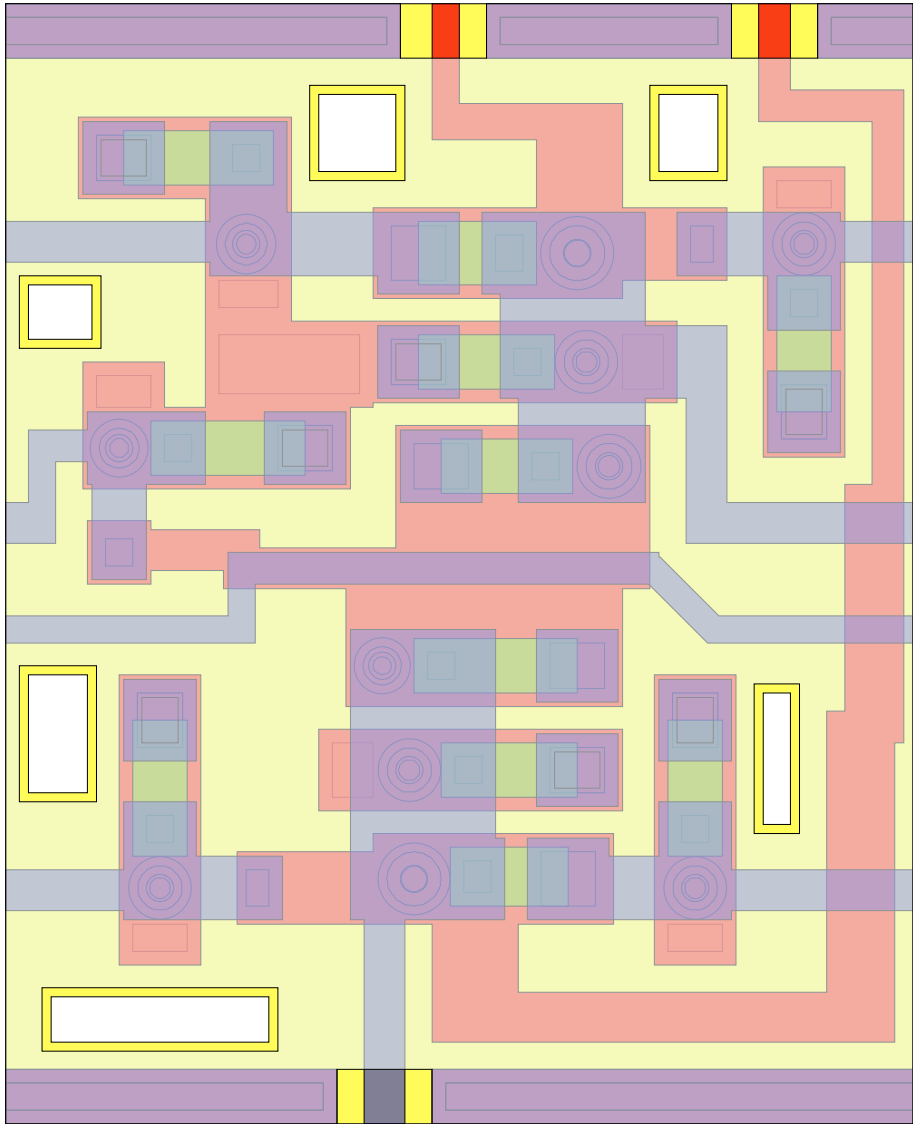

Figure A.23: Final layout of eDES deserialiser cell (Parameter values in Table A.10).

Table A.10: Extracted (inductances and coupling factor) and nominal parameters of eDES.

Parameter	Value	Parameter	Value
L_{1s}	1.25 pH	L_{1r}	2.52 pH
L_{2s}	1.42 pH	L_{2r}	1.38 pH
L_{3s}	0.05 pH	L_{3r}	0.22 pH
L_{4s}	1.95 pH	L_{4r}	1.65 pH
L_{5s}	2.89 pH	L_{5r}	1.73 pH
L_{6s}	0.20 pH	L_{6r}	0.2 pH
L_{7s}	0.9 pH	L_{7r}	0.55 pH
L_{8s}	4.2 pH	L_{8r}	3.4 pH
L_9	1.96 pH	L_{S1}	3.1 pH
L_{f1}	16.1 pH	K_f	≈ 0.24
I_{cBc1s}	163 µA	I_{cBd2s}	188 µA
I_{cBc2s}	188 µA	I_{cBd1}	150 µA
I_{cBc1r}	163 µA	I_{cBdes}	200 µA
I_{cBc2r}	188 µA	I_{cBder}	138 µA
I_{cBd2s}	188 µA	β_{c*}	≈ 1
I_{cB1s}	288 µA	I_{cB1r}	225 µA
I_{cBbs}	500 µA	I_{cBbr}	500 µA
L_{pBc1s}	0.153 pH	L_{pBc2s}	0.162 pH
L_{pBc1r}	0.125 pH	L_{pBc2r}	0.125 pH
L_{pBd1}	0.125 pH	L_{pBd2s}	0.281 pH
L_{pBd2r}	0.280 pH	L_{pB*}	~ 1 pH

A.5.4 eTFF — Toggle flip-flop cell

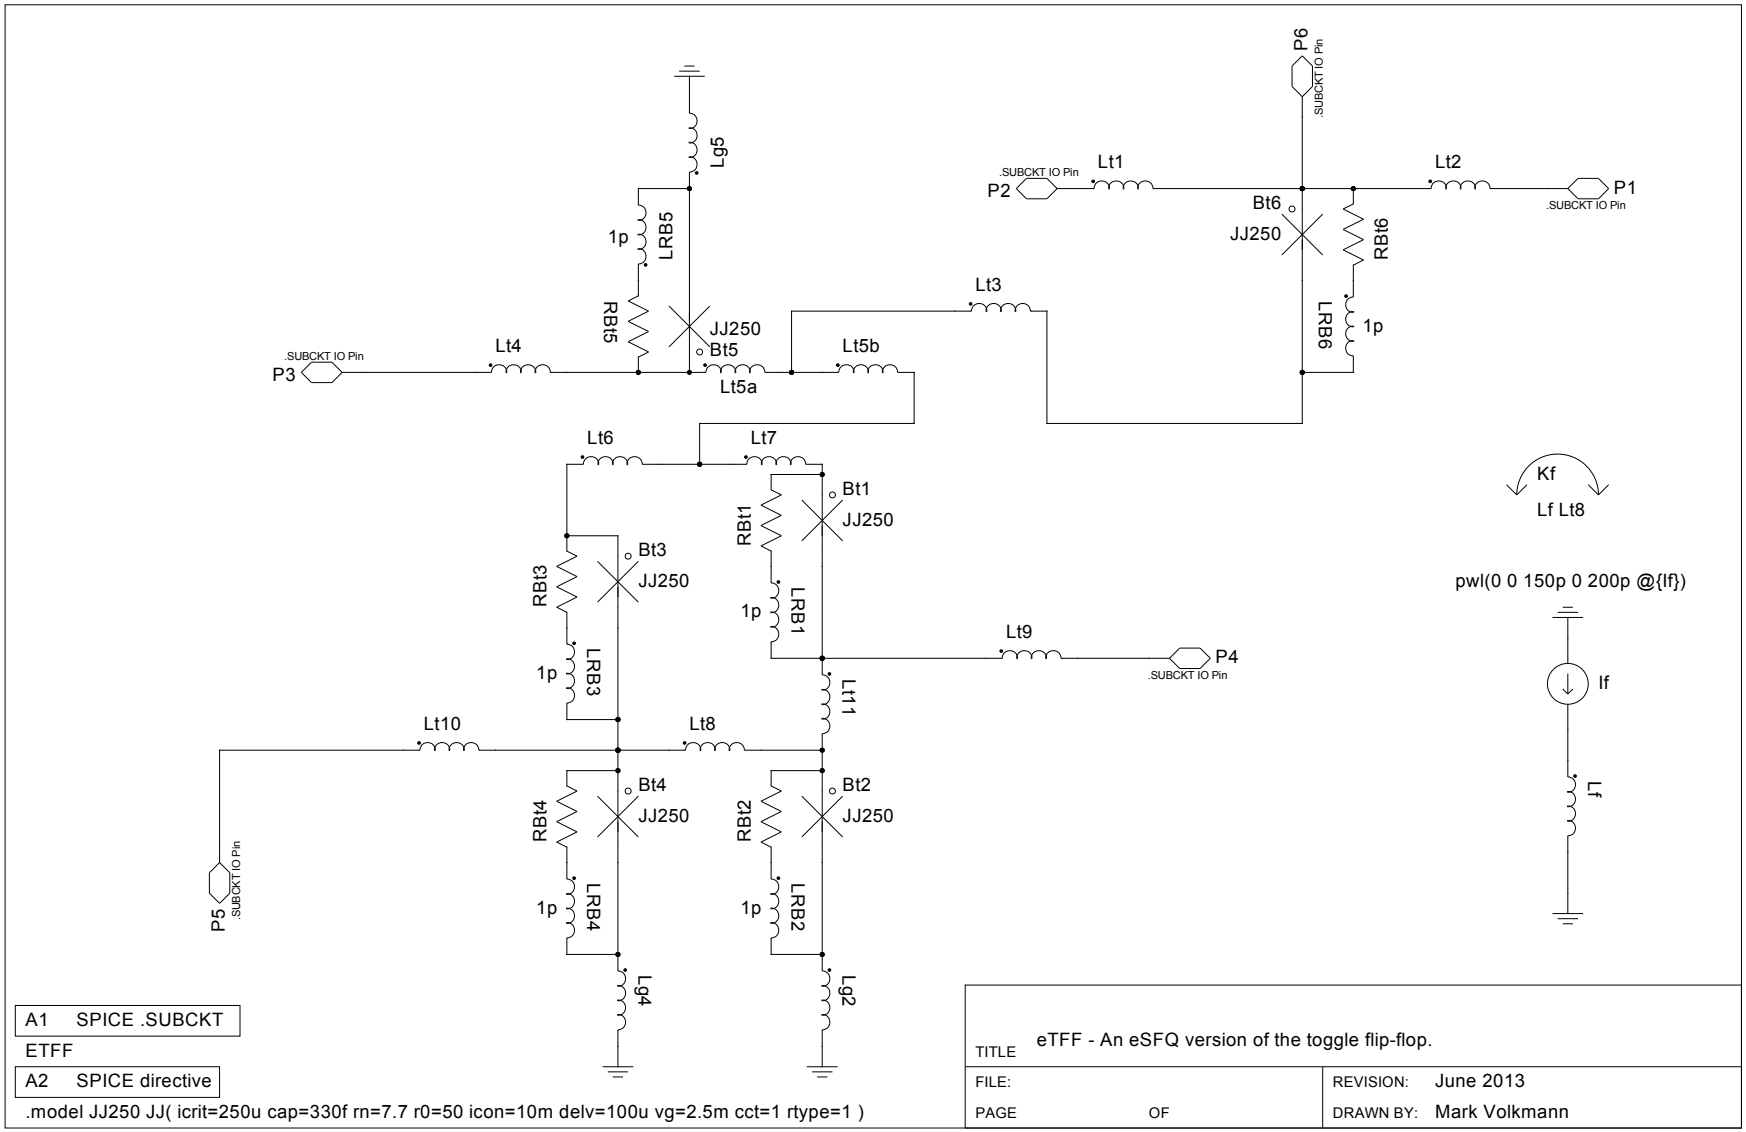

Figure A.24: Final circuit diagram of eTFF toggle flip-flop cell (Parameter values in Table A.11).

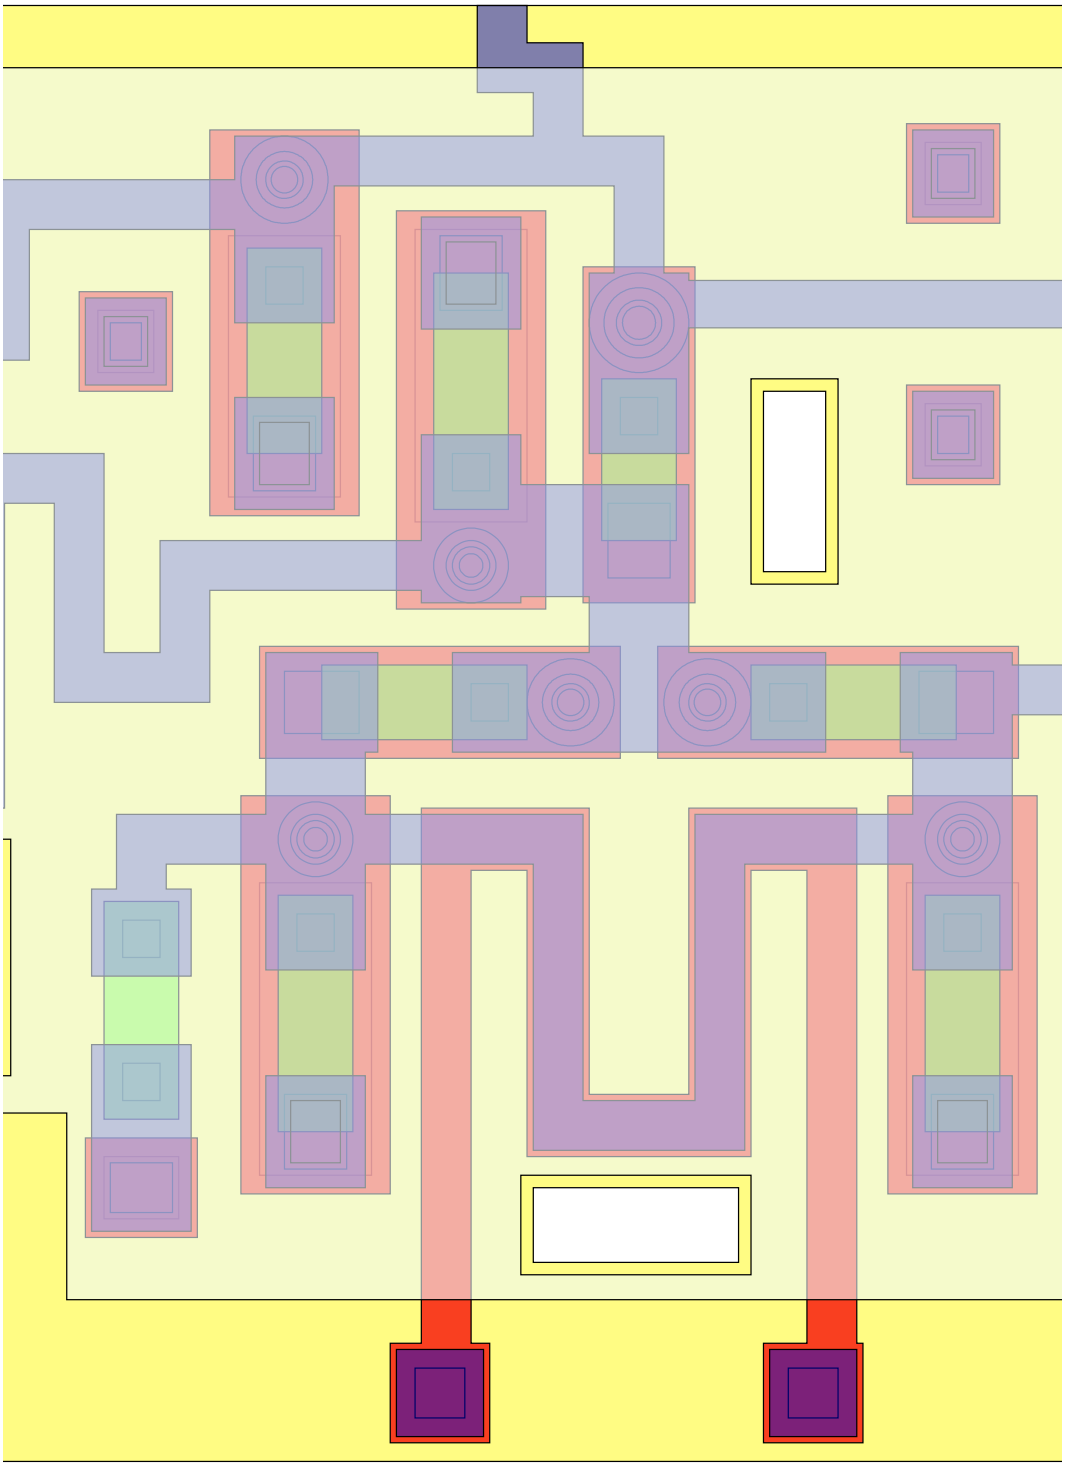

Figure A.25: Final layout of eTFF toggle flip-flop cell (Parameter values in Table A.11).

Table A.11: Extracted (inductances and coupling factor) and nominal parameters of `eTFF`.

Parameter	**Value**	**Parameter**	**Value**
L_{t1}	1.56 pH	L_{t6}	1.86 pH
L_{t2}	1.65 pH	L_{t7}	1.39 pH
L_{t3}	0.80 pH	L_{t8}	9.16 pH
L_{t4}	6.85 pH	L_{t9}	1.64 pH
L_{t5}	0.55 pH	L_{t10}	4.01 pH
L_{t11}	0.24 pH	L_{g4}	0.28 pH
L_{g2}	0.47 pH	L_{g5}	0.20 pH
I_{cBt1}	161 µA	I_{cBt2}	125 µA
I_{cBt3}	161 µA	I_{cBt4}	125 µA
I_{cBt5}	125 µA	I_{cBt6}	250 µA
L_f	18.1 pH	K_f	0.2
β_{c*}	1		

A.6 eSFQ chips — functional tests

A.6.1 Data transport chip batch summary

Table A.12: Structures developed to functionally investigate various eSFQ shift registers and deserialisers (Suite I).

Structure	Cell	Location[1]	Length	Limiting line L_b[2]	Bias line (L_q)	Bias[3]	Magnetic bias[4]
Investigate series bias-line width – `mv348_esr_lqvar`							
E1	eSR	South	32 bits	medium	narrow	8 points	No
E2	eSR	Centre	16 bits	medium	narrow	1-to-4 tree	No
E3	eSR	North	32 bits	medium	wide	8 points	No
Investigate limiting inductor size – `mv348_esr_lbvar`							
E4	eSR	South	16 bits	low	narrow	4 points	No
E5	eSR	East	16 bits	medium	narrow	4 points	No
E6	eSR	North	16 bits	high	narrow	4 points	No
Investigate magnetic current biasing – `mv348_esr_mag`							
M1	MeSR	South	32 bits	medium	narrow	8 points	Yes
M2	MeSR	East	16 bits	medium	narrow	4 points	Yes
M3	MeSR	North	16 bits	medium	wide	4 points	Yes
Investigate 16-bit eSFQ deserialiser – `mv348_edes16_mag`							
D1	eDES	Centre	16 bits	medium	narrow	8 points	Yes
Investigate 8-bit and 4-bit eSFQ deserialisers – `mv348_edes8a4_mag`							
D2	eDES	East	8 bits	medium	narrow	4 points	Yes
D3	eDES	West	4 bits	medium	narrow	2 points	Yes

[1] On the GDSII layout
[2] Inductor carrying current into cell from bias line
[3] Refers to different injection points on the same bias line
[4] In addition to the inductively injected bias

A.6.2 `mv348-esr-lqvar` — Structures for investigation of L_q size.

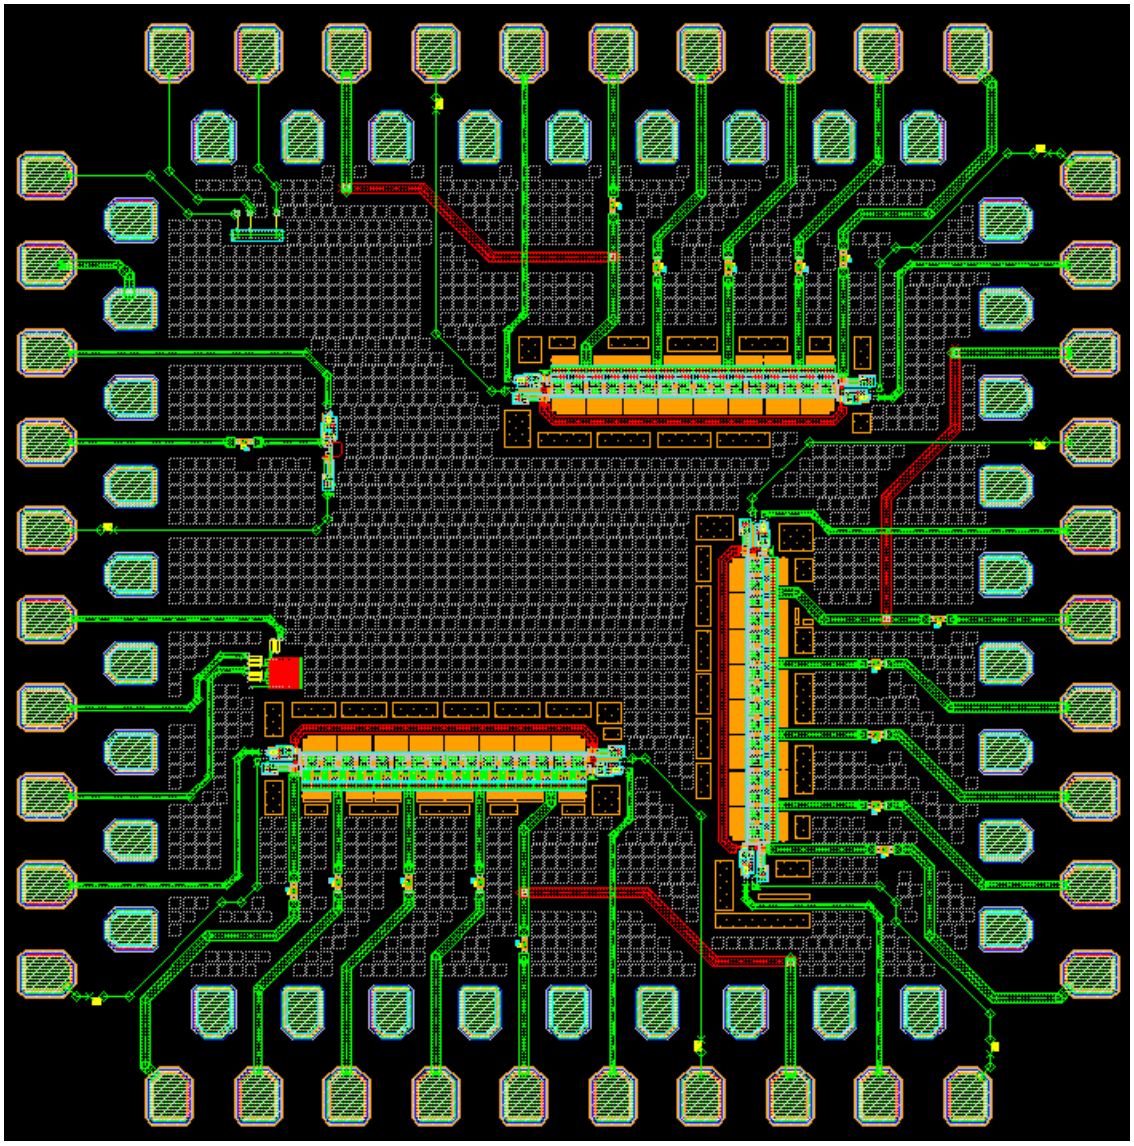

Figure A.26: Chip `mv348-esr-lqvar`, to experimentally verify eSR shift registers and investigate effect of L_b.

A.6.3 `mv348-esr-lbvar` — Structures for investigation of L_b size.

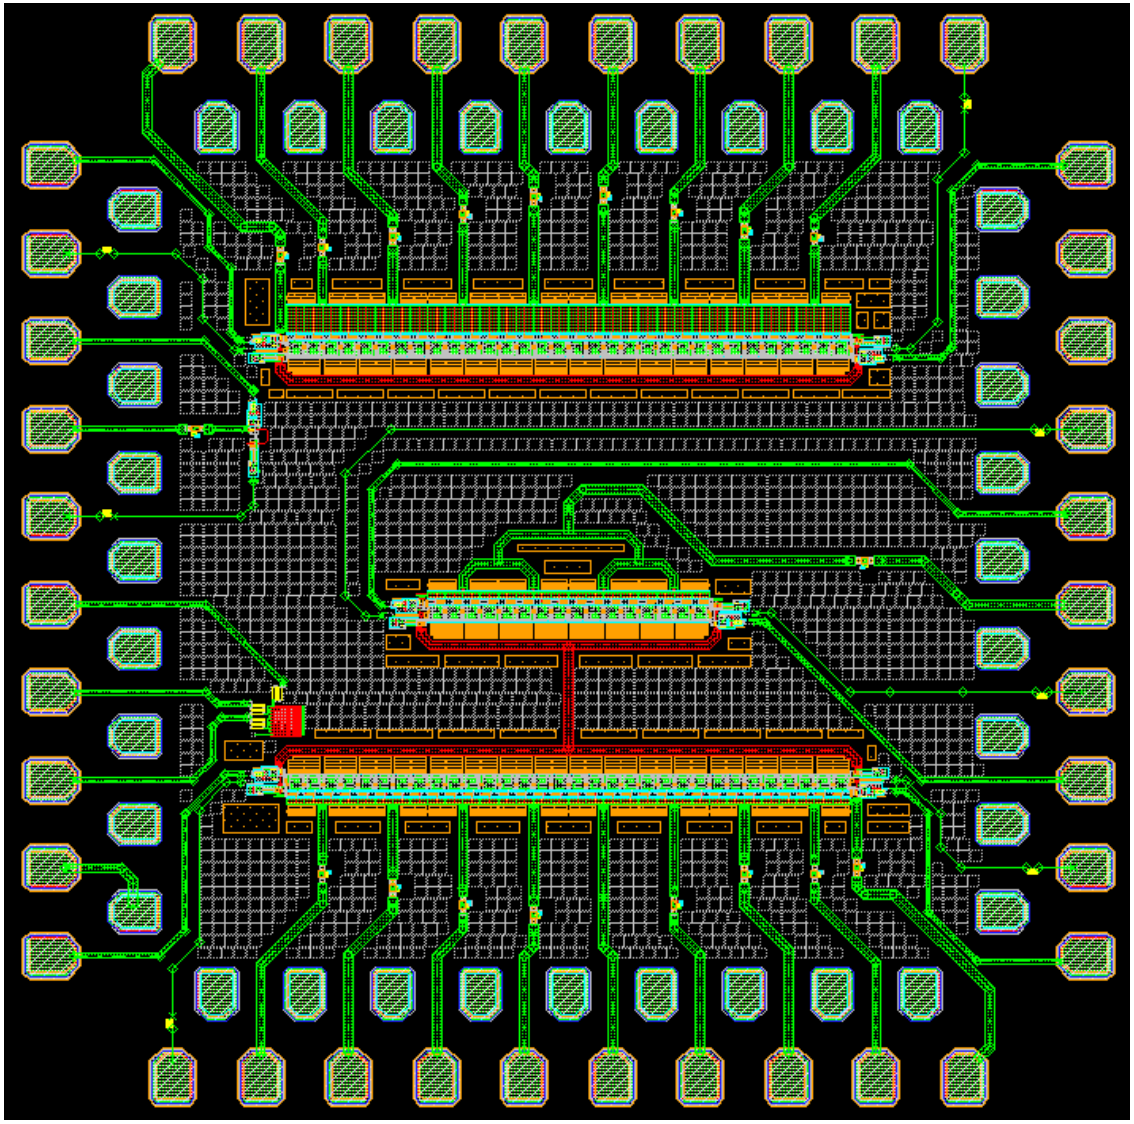

Figure A.27: Chip `mv348-esr-lbvar`, to experimentally verify eSR shift registers and investigate effect of L_b.

A.6.4 `mv348-esr-mag` — Structures for investigation of magnetically biased shift registers.

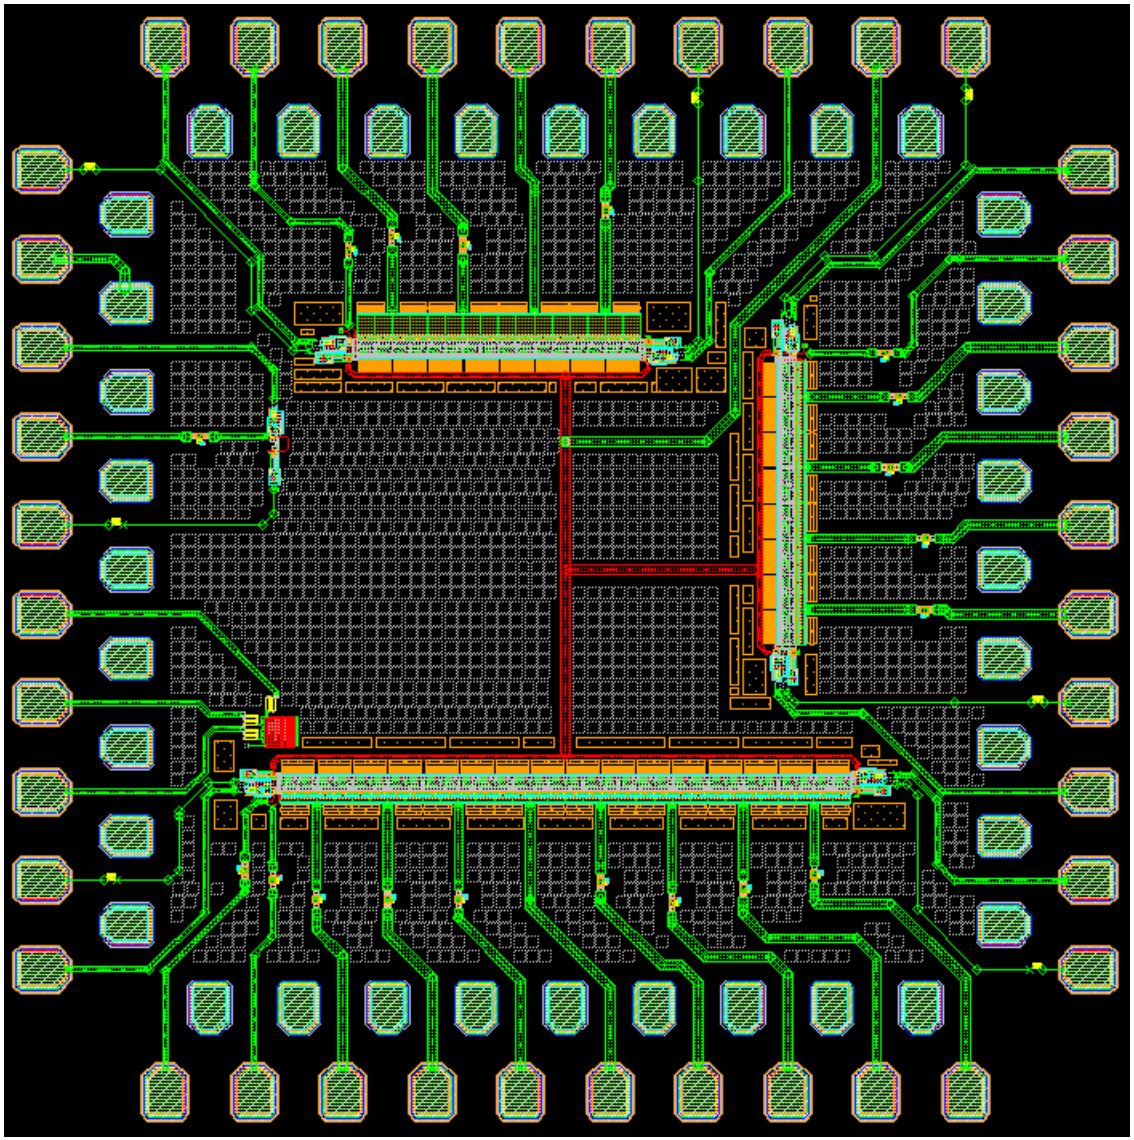

Figure A.28: Chip `mv348-esr-mag`, to experimentally verify MeSR shift registers and investigate effect of L_b.

A.6.5 `mv348-edes16-mag` — Structures for investigation of 16-bit deserialiser.

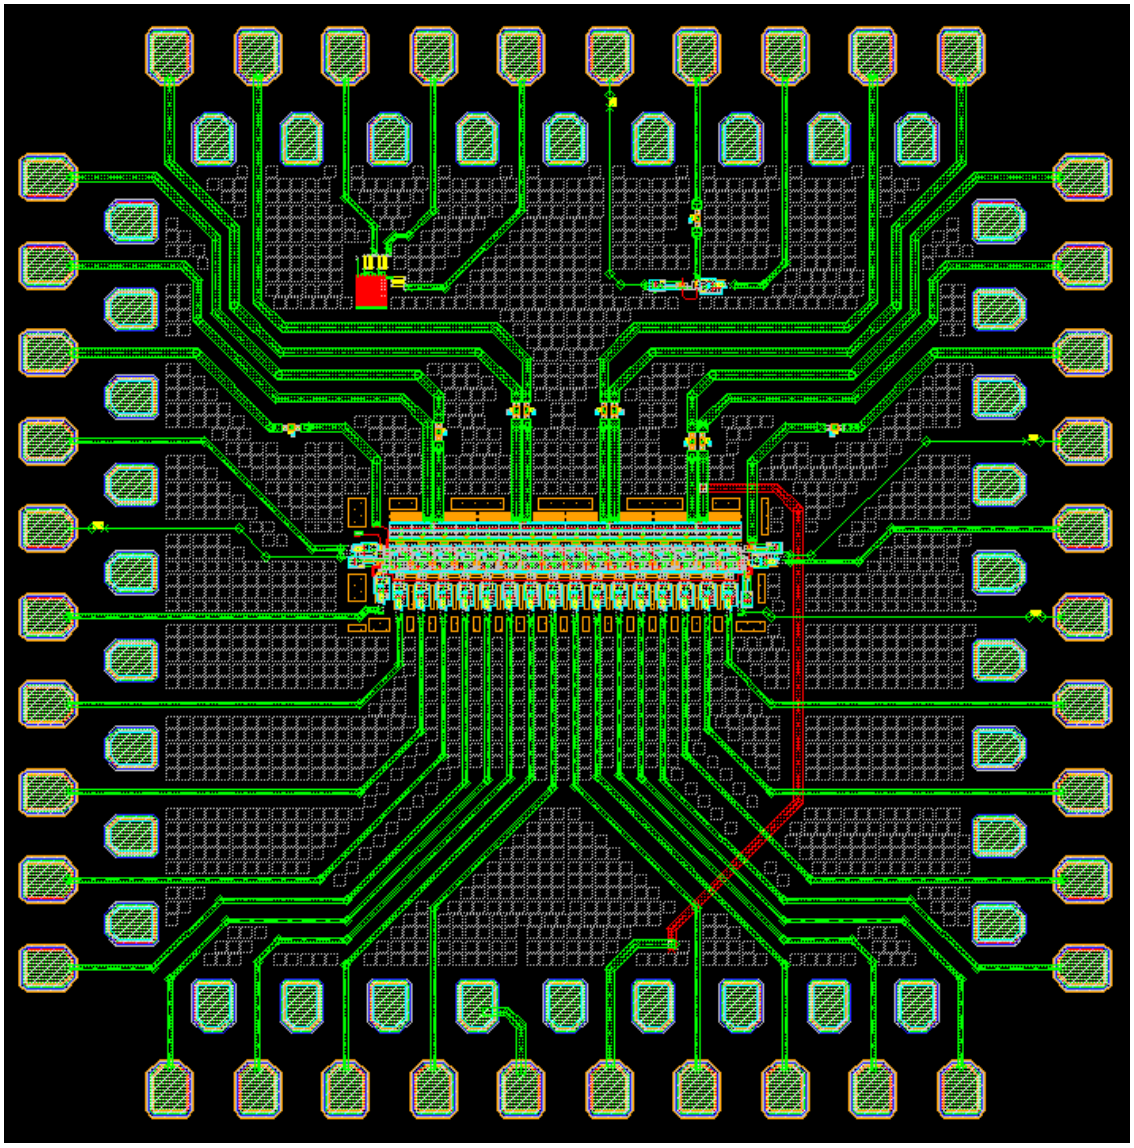

Figure A.29: Chip `mv348-edes-mag`, to experimentally verify eDES deserialiser.

A.6.6 `mv348-edes8a4-mag` — Structures for investigation of shorter deserialisers.

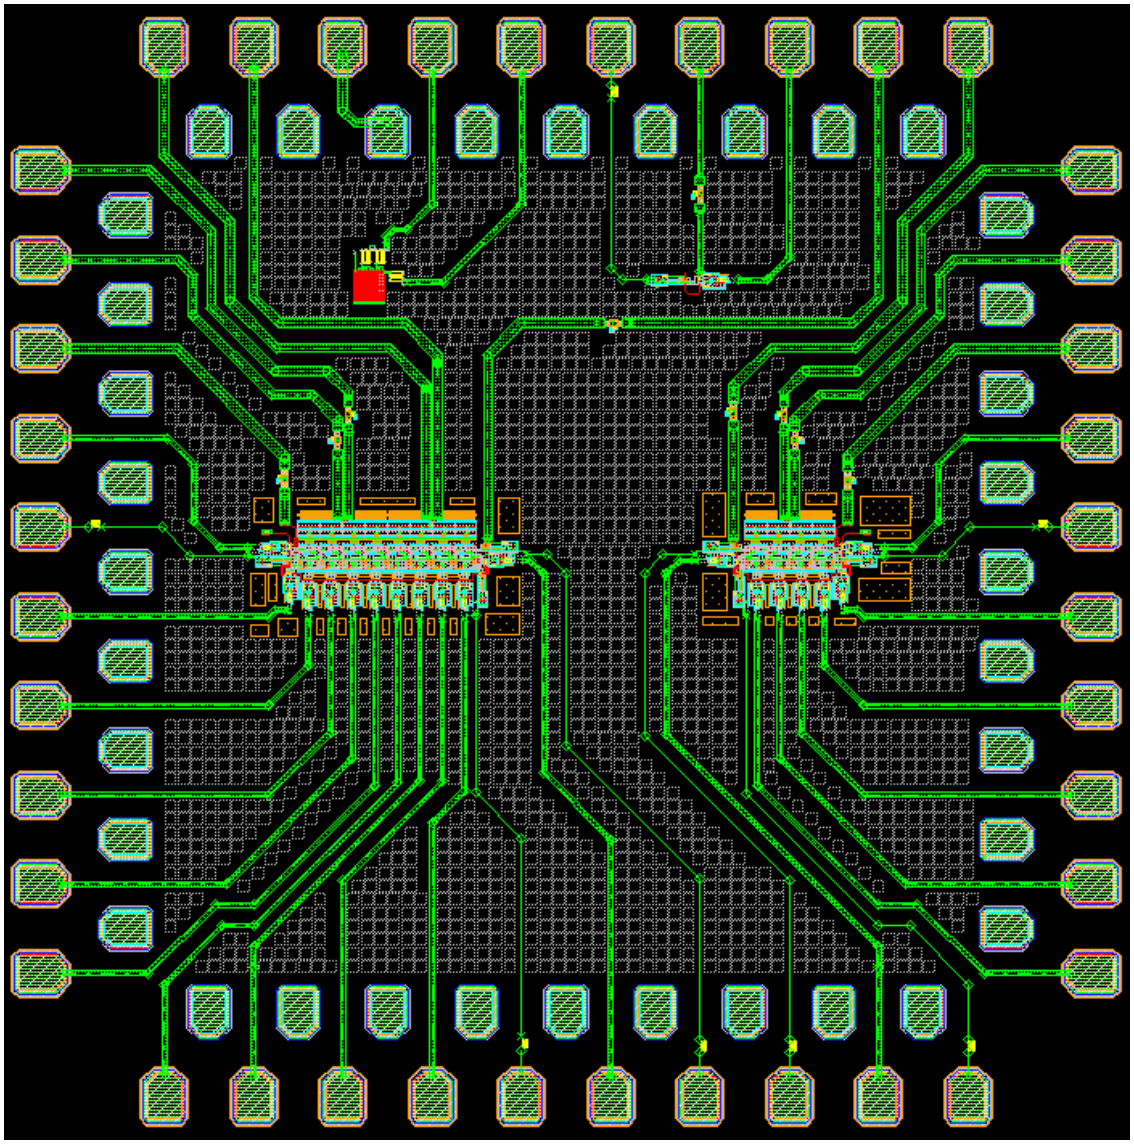

Figure A.30: Chip `mv348-edes8a4-mag`, to experimentally verify shorter eDES deserialisers.

A.6.7 `mv356-etff-8b` — Structures for investigation of eTFF cell

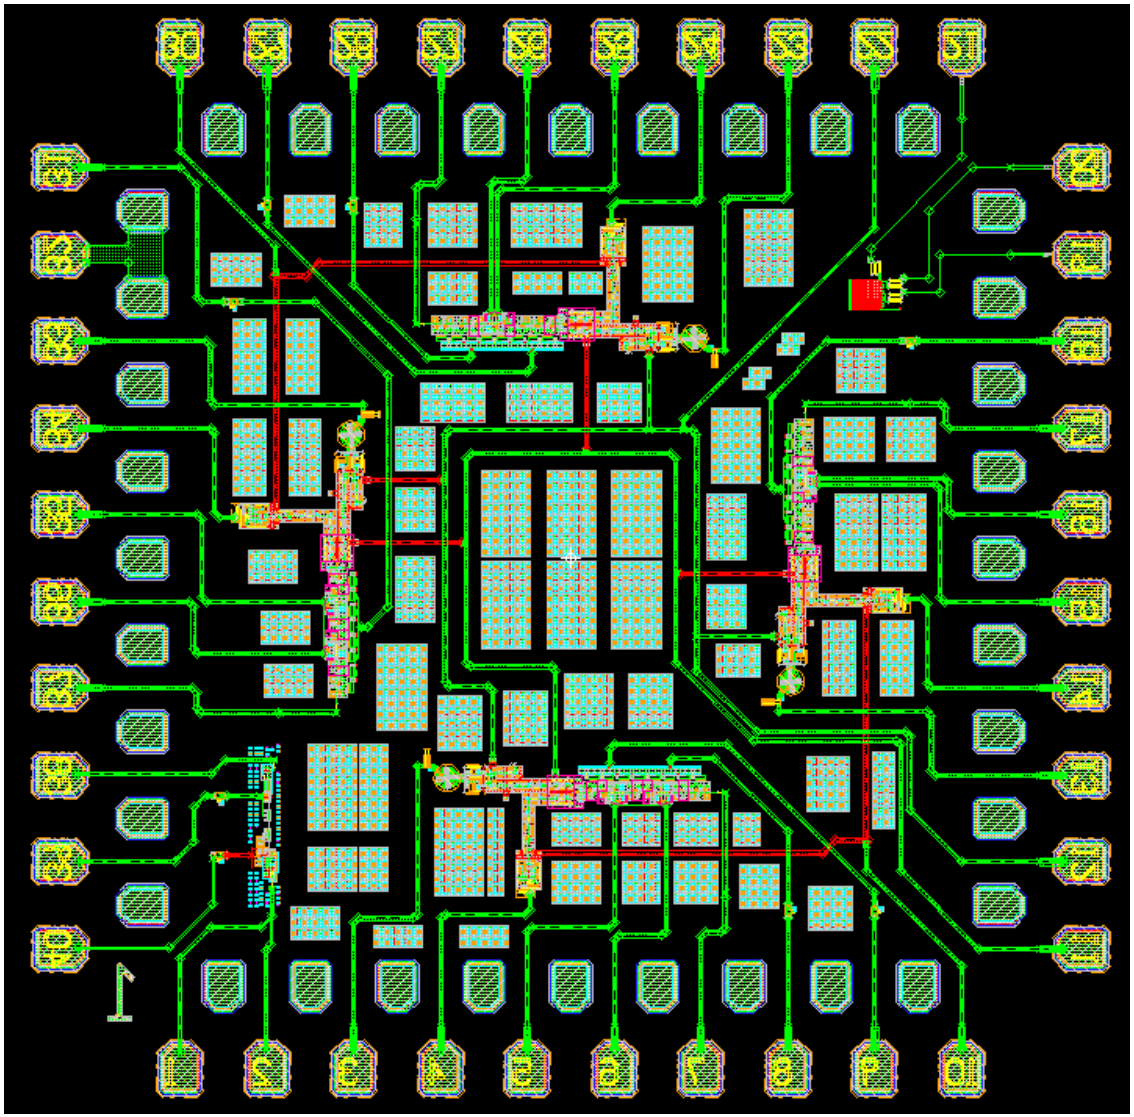

Figure A.31: Chip `mv348-etff-8b`, to experimentally verify functional operation of the `eTFF` cell, embedded in various configurations in 8-bit shift registers.

A.7 eSFQ chips — high-speed tests

A.7.1 Chip batch summary

Table A.13: Structures developed to investigate high-speed operation of various eSFQ shift registers and deserialisers (Suite II).

Structure	Cell	Location[1]	Length	Limiting line L_b[2]	Bias line (L_q)	Bias[3]	Magnetic bias[4]
Investigate series bias-line width – `mv350_esr_lqvar_hf`							
EH1	eSR	South	32 bits	medium	narrow	8 points	No
EH2	eSR	Centre	16 bits	medium	narrow	1-to-4 tree	No
EH3	eSR	North	32 bits	medium	wide	8 points	No
Investigate limiting inductor size – `mv350_esr_lbvar_hf`							
EH4	eSR	South	16 bits	low	narrow	4 points	No
EH5	eSR	East	16 bits	medium	narrow	4 points	No
EH6	eSR	North	16 bits	high	narrow	4 points	No
Investigate magnetic current biasing – `mv350_mag_hf`							
MH1	MeSR	South	32 bits	medium	narrow	8 points	Yes
MH2	MeSR	East	16 bits	medium	narrow	4 points	Yes
MH3	MeSR	North	16 bits	medium	wide	4 points	Yes
Investigate eSFQ deserialisers at high speeds – `mv350_pmd_edes16_hf`							
DH1	eDES + PMD	South	16 bits	medium	narrow	2×4 points	Yes
DH2	eDES	North	16 bits	medium	narrow	2×4 points	Yes

[1] On the GDSII layout
[2] Inductor carrying current into cell from bias line
[3] Refers to different injection points on the same bias line
[4] In addition to the inductively injected bias

A.7.2 `mv350_esr_lqvar_hf` — Structures for high-speed investigation of L_q size.

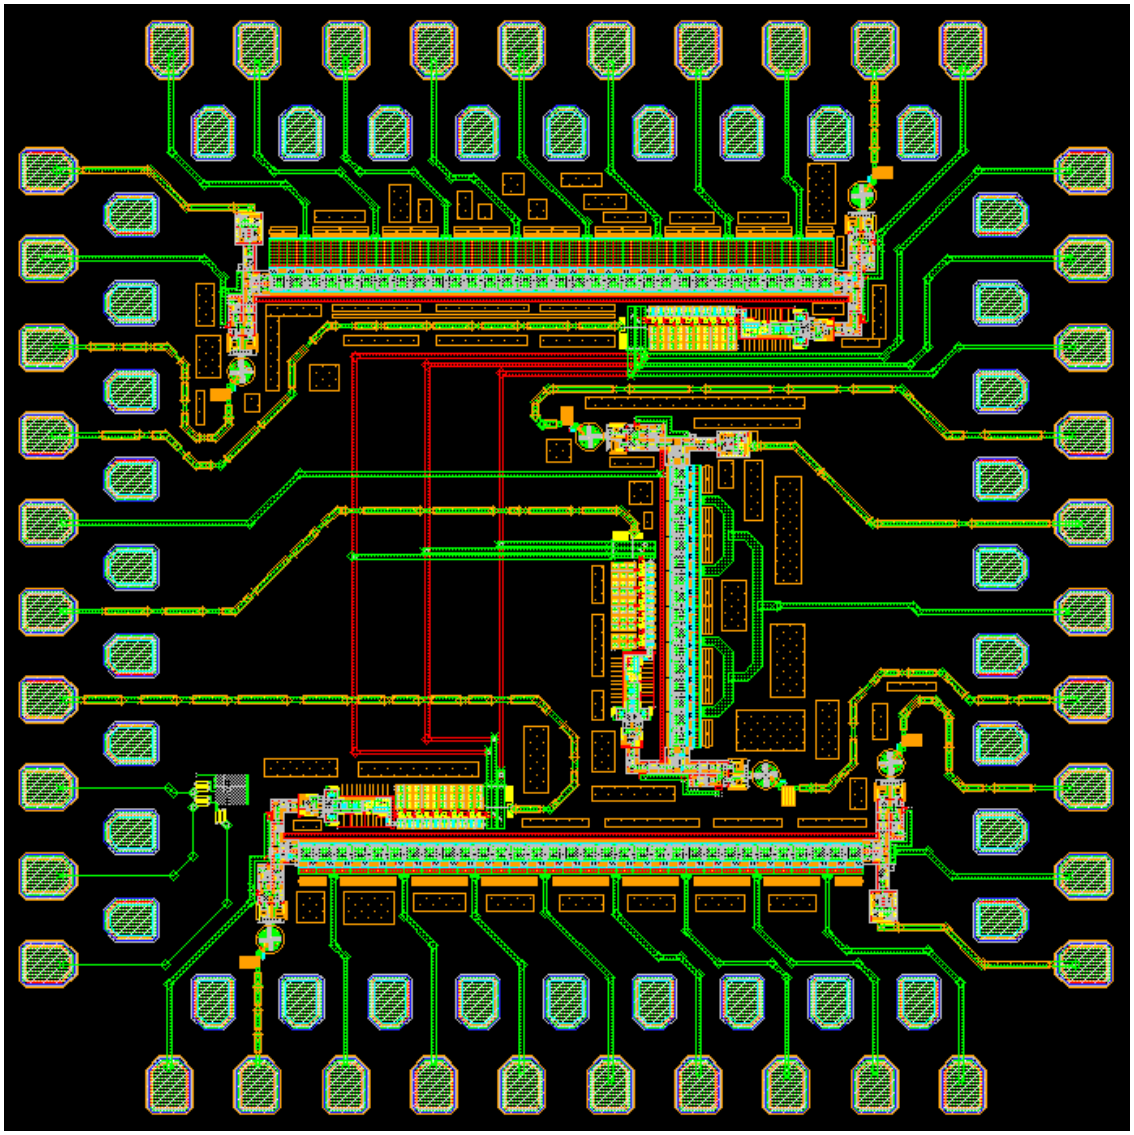

Figure A.32: Chip `mv350_esr_lqvar_hf`, to experimentally verify high-speed operation of eSR shift registers and investigate effect of L_b.

A.7.3 `mv350_esr_lbvar_hf` — Structures for high-speed investigation of L_b size.

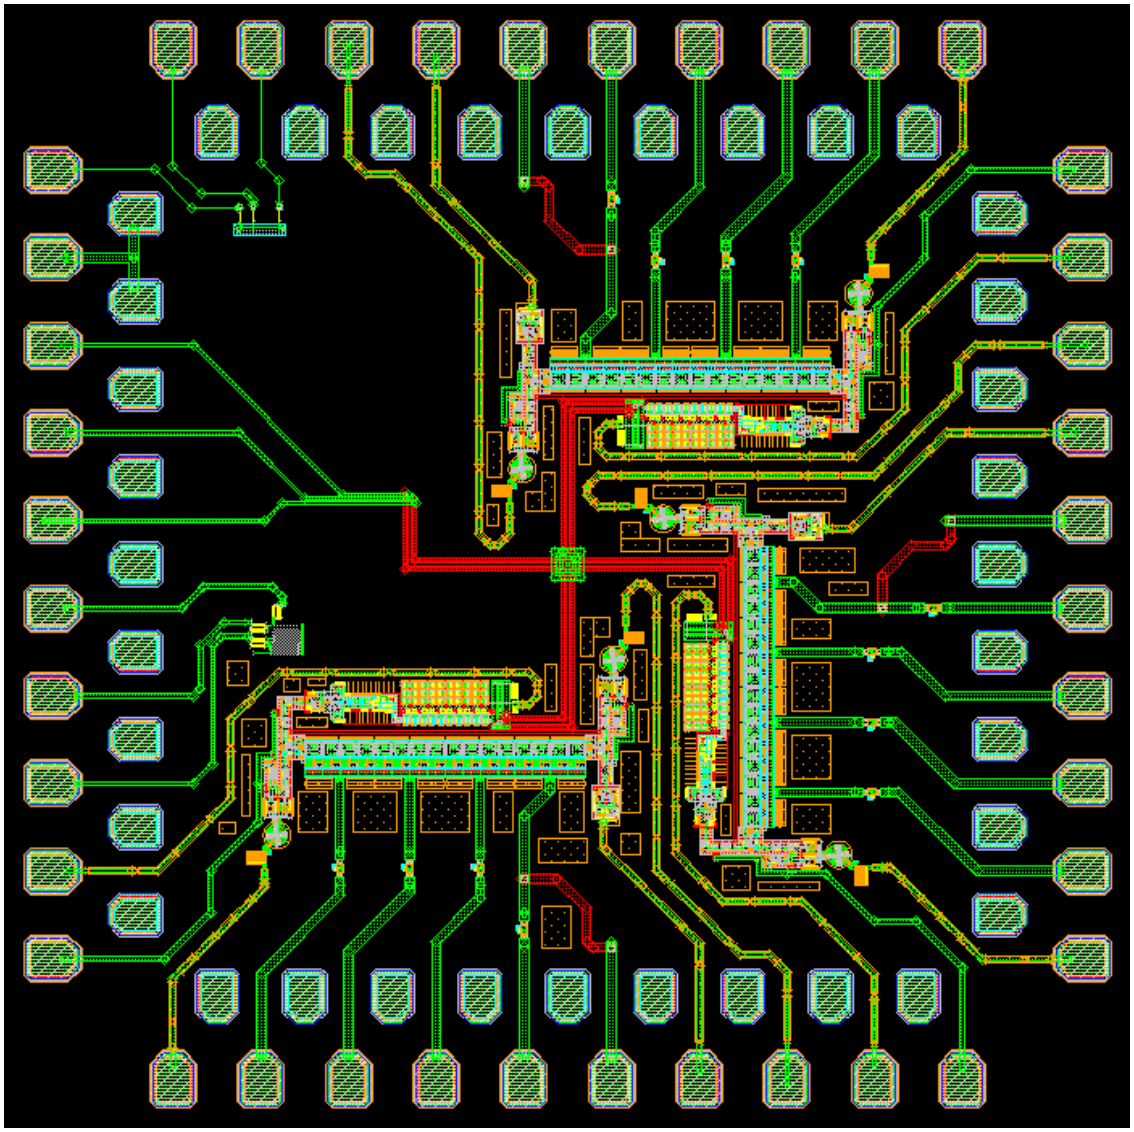

Figure A.33: Chip `mv350_esr_lbvar_hf`, to experimentally verify high-speed operation of eSR shift registers and investigate effect of L_b.

A.7.4 `mv350_mag_hf` — Structures for high-speed investigation of shift registers with corrective flux bias.

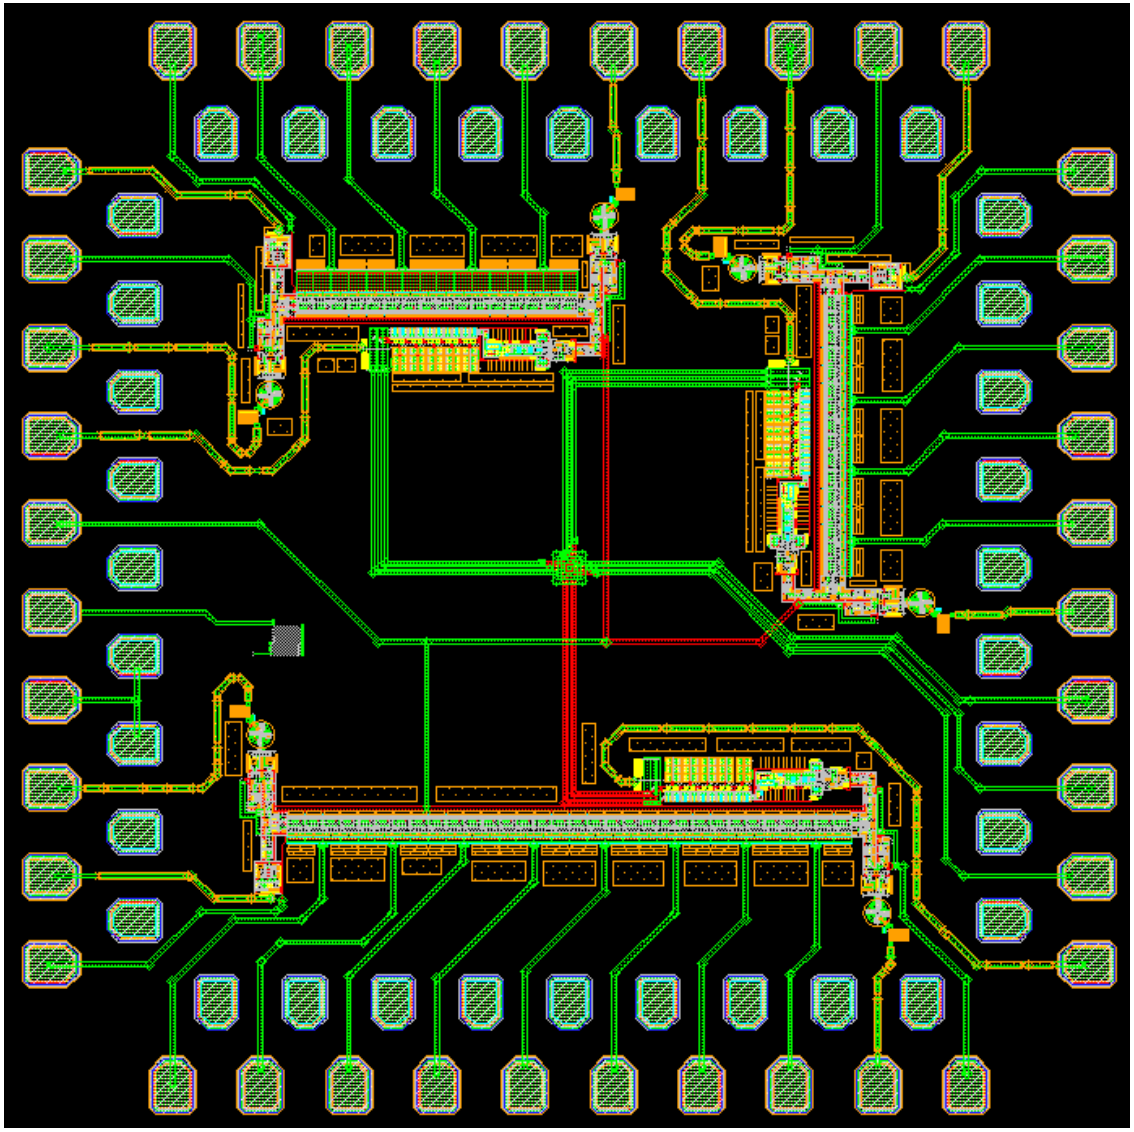

Figure A.34: Chip `mv350_mag_hf`, to experimentally verify high-speed operation of MeSR shift registers with corrective flux bias.

A.7.5 `mv350_pmd_edes16_hf` — Structures for high-speed investigation of deserialisers.

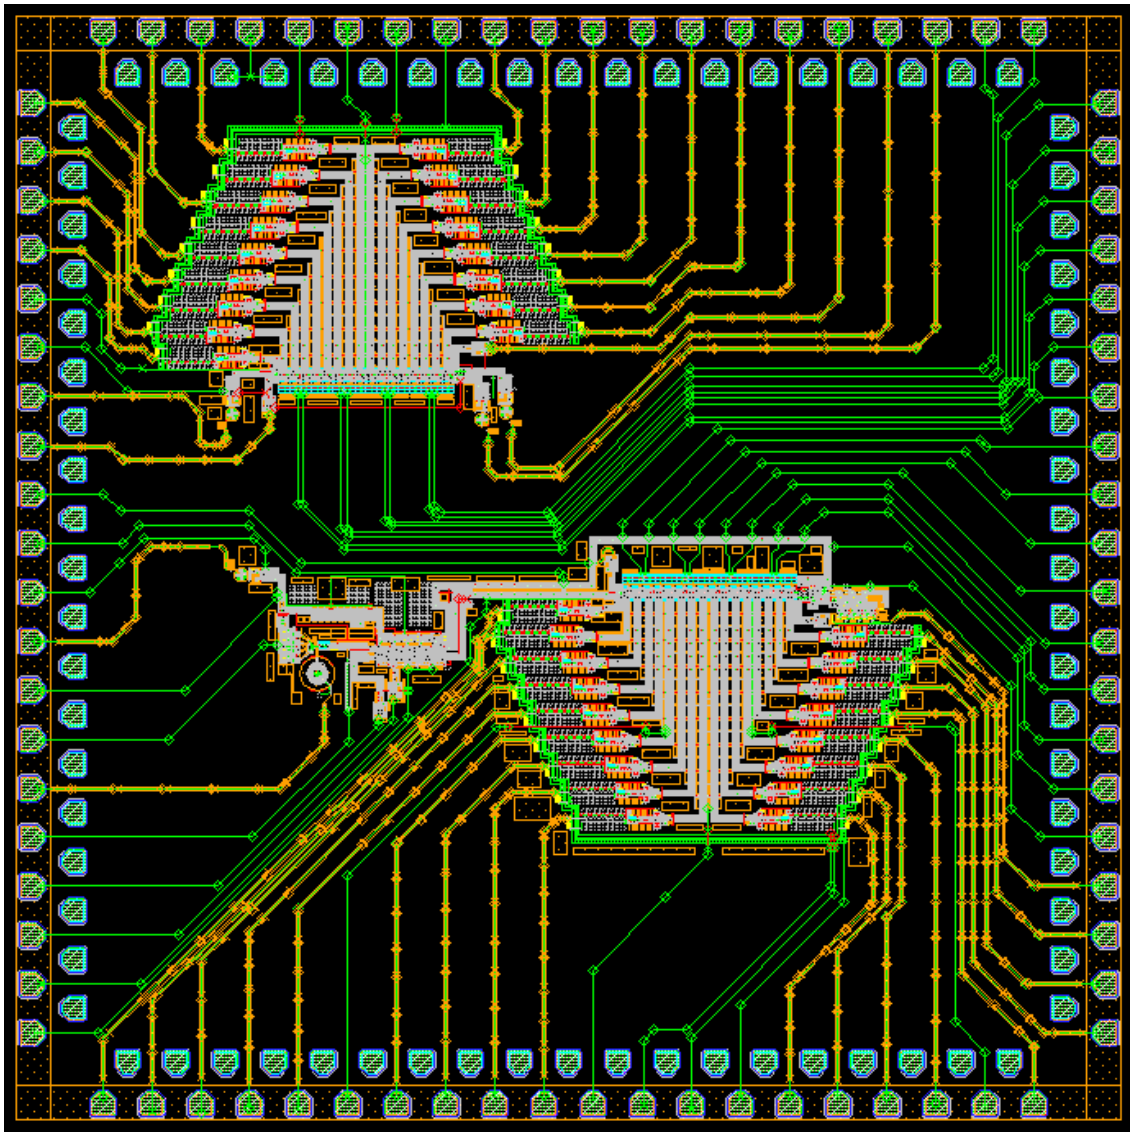

Figure A.35: Chip `mv350_pmd_edes16_hf`, to experimentally verify high-speed operation of eDES deserialisers.

Table A.14: Structures developed to investigate high-speed operation of various eSFQ shift registers and deserialisers (Suite III).

Structure	Cell	Location[1]	Length	Limiting line L_b[2]	Bias line (L_q)	Bias[3]	Magnetic bias[4]
Investigate power dissipation (comparison) – `mv352_magcomp_hf`							
EC1	eSR	North	32 bits	medium	narrow	8 points	No
MC1	MeSR	South	32 bits	medium	narrow	8 points	Yes
Investigate power dissipation (comparison), no clock divider – `mv352_magcomp_hf_nodiv`							
EC2	eSR	North	32 bits	medium	narrow	8 points	No
MC2	MeSR	South	32 bits	medium	narrow	8 points	Yes
Investigate longer structure – `mv352_esrlong_hf`							
EL1	MeSR	Centre	184 bits	medium	narrow	23 points	No
Investigate longer structure, no clock divider – `mv352_esrlong_nodiv_hf`							
EL2	MeSR	Centre	184 bits	medium	narrow	23 points	No
Investigate pattern reproduction, no clock divider – `mv352_esrlong_dsw_nodiv_hf`							
EL3	MeSR	Centre	184 bits	medium	narrow	23 points	No
Investigate eSFQ deserialisers at high speeds – `mv352_pmd_edes16_b_hf`							
DB1	eDES + PMD	South	16 bits	medium	narrow	2×4 points	Yes
DB2	eDES	North	16 bits	medium	narrow	2×4 points	Yes
Investigate eSFQ deserialisers at high speeds, refresh – `mv352_pmd_edes16_it2_hf`							
DI3	eDES + PMD	South	16 bits	medium	narrow	2×4 points	Yes
DI4	eDES	North	16 bits	medium	narrow	2×4 points	Yes

[1] On the GDSII layout
[2] Inductor carrying current into cell from bias line
[3] Refers to different injection points on the same bias line
[4] In addition to the inductively injected bias

A.7.6 mv352_magcomp_hf — Investigate power dissipation (comparison)

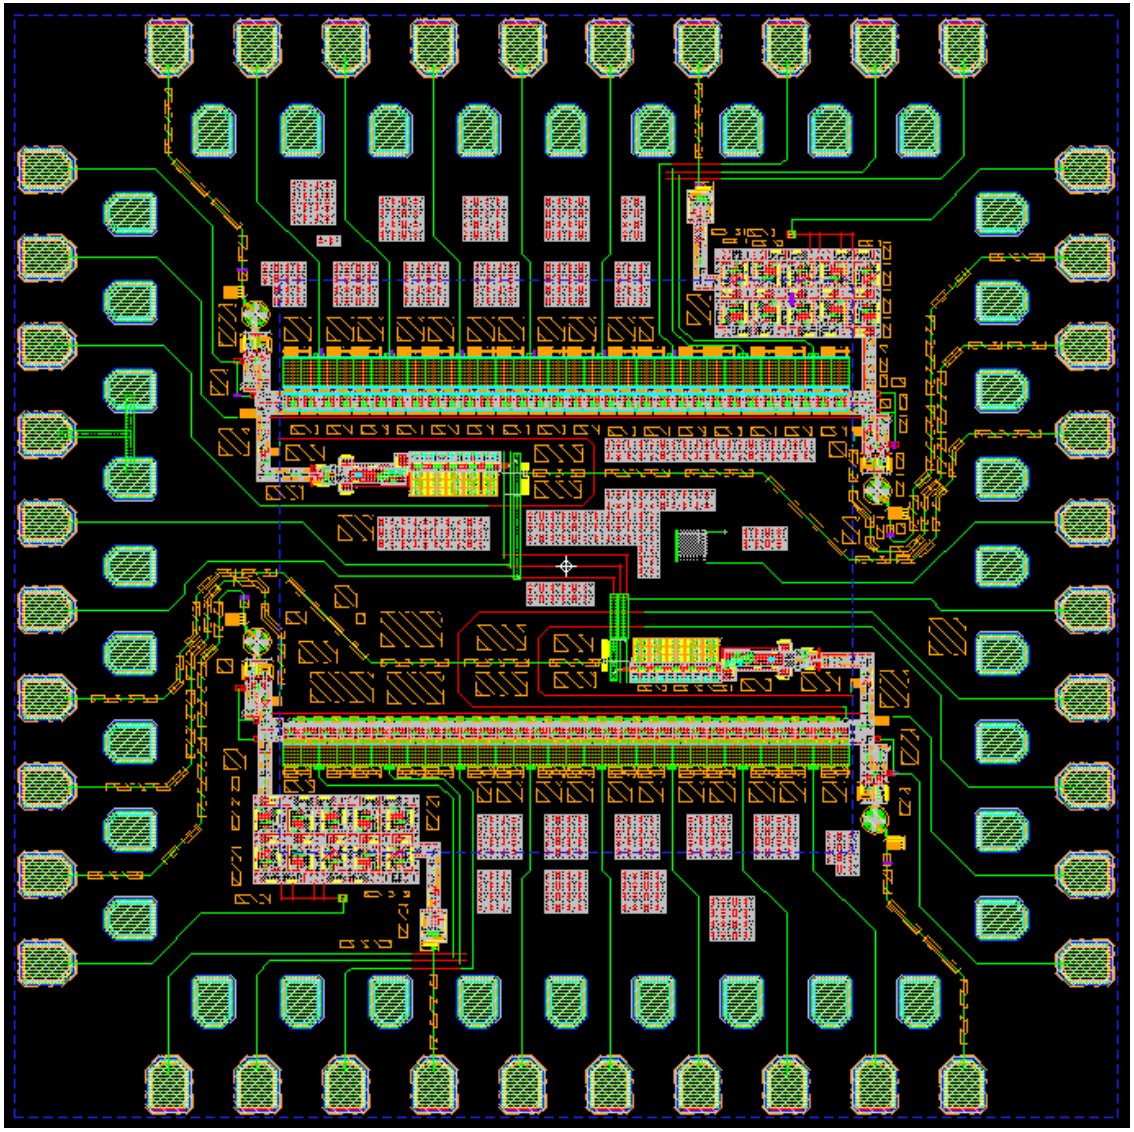

Figure A.36: Chip mv352_magcomp_hf, to experimentally confirm power dissipation of eSR and MeSR cells.

A.7.7 `mv352_magcomp_nodiv_hf` — Investigate power dissipation (comparison), no clock divider

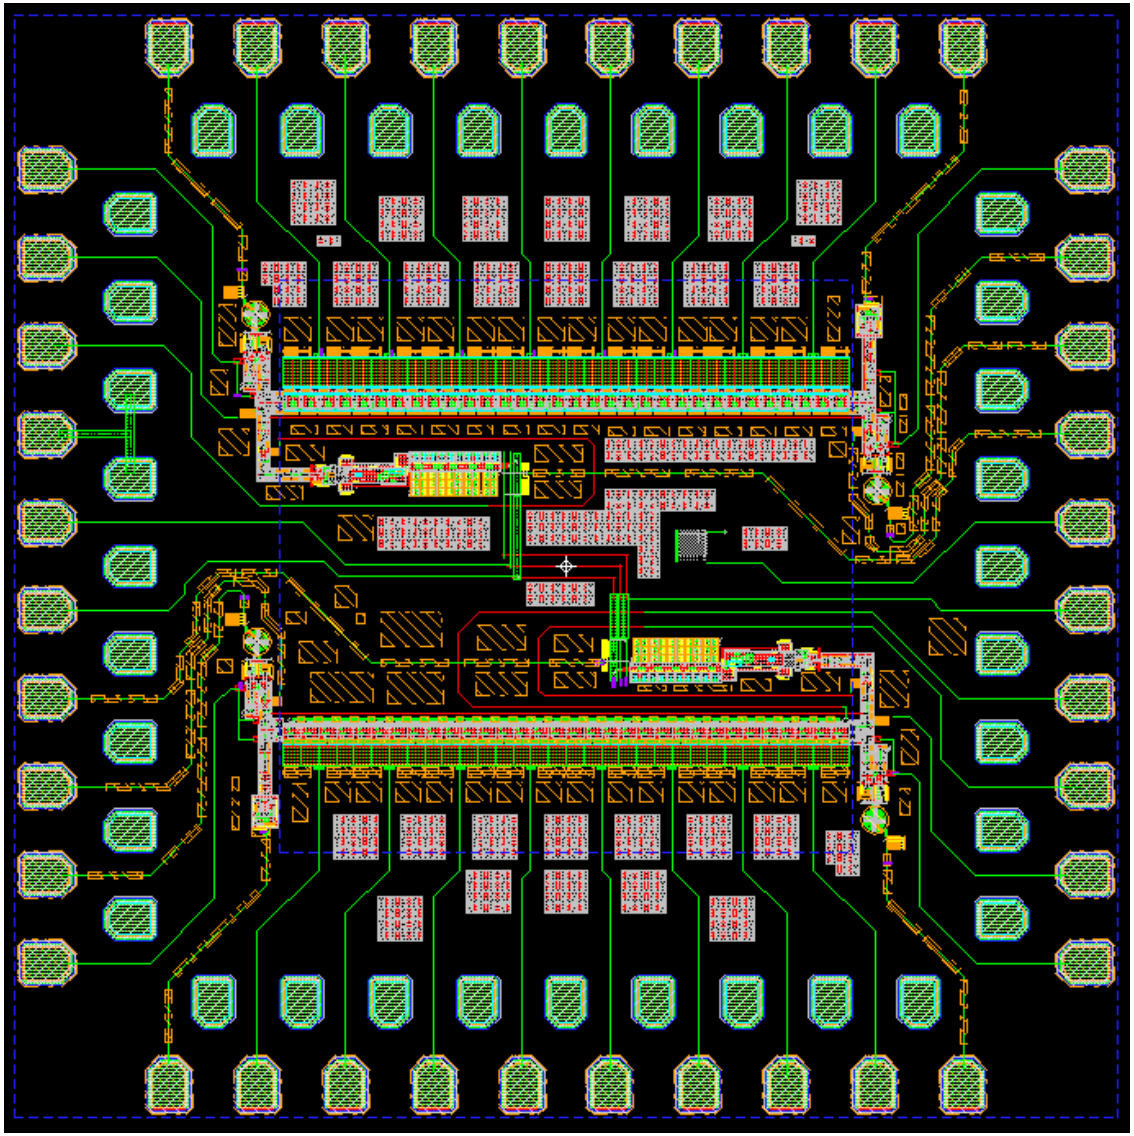

Figure A.37: Chip `mv352_magcomp_nodiv_hf`, to experimentally confirm power dissipation of eSR and MeSR cells, without a clock divider.

A.7.8 `mv352_esrlong_hf` — Investigate longer structure

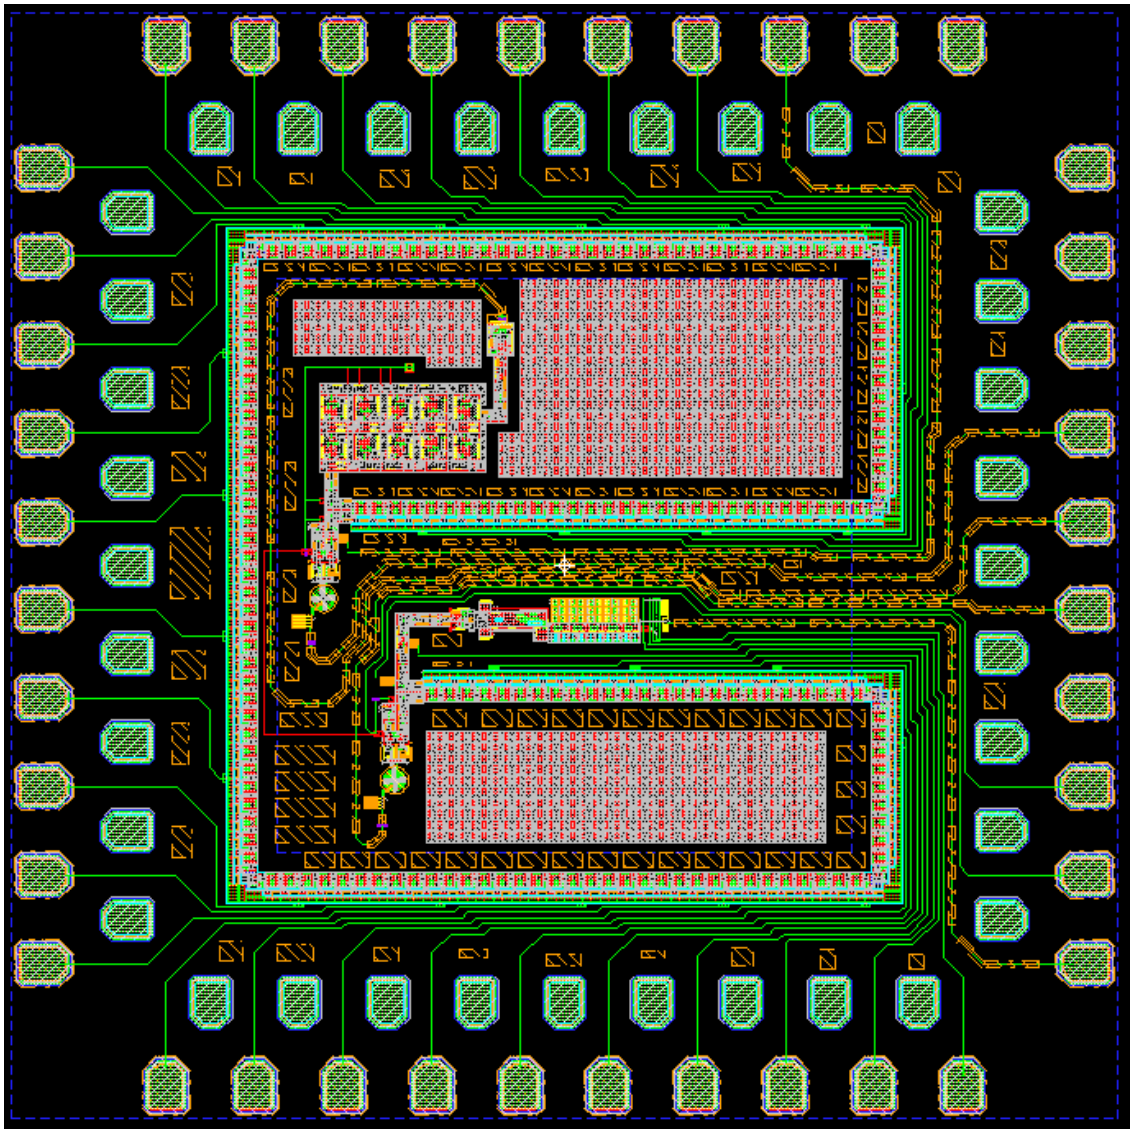

Figure A.38: Chip `mv352_esrlong_hf`, to experimentally confirm operation and power dissipation of longer shift register.

A.7.9 `mv352_esrlong_nodiv_hf` — Investigate longer structure, no clock divider

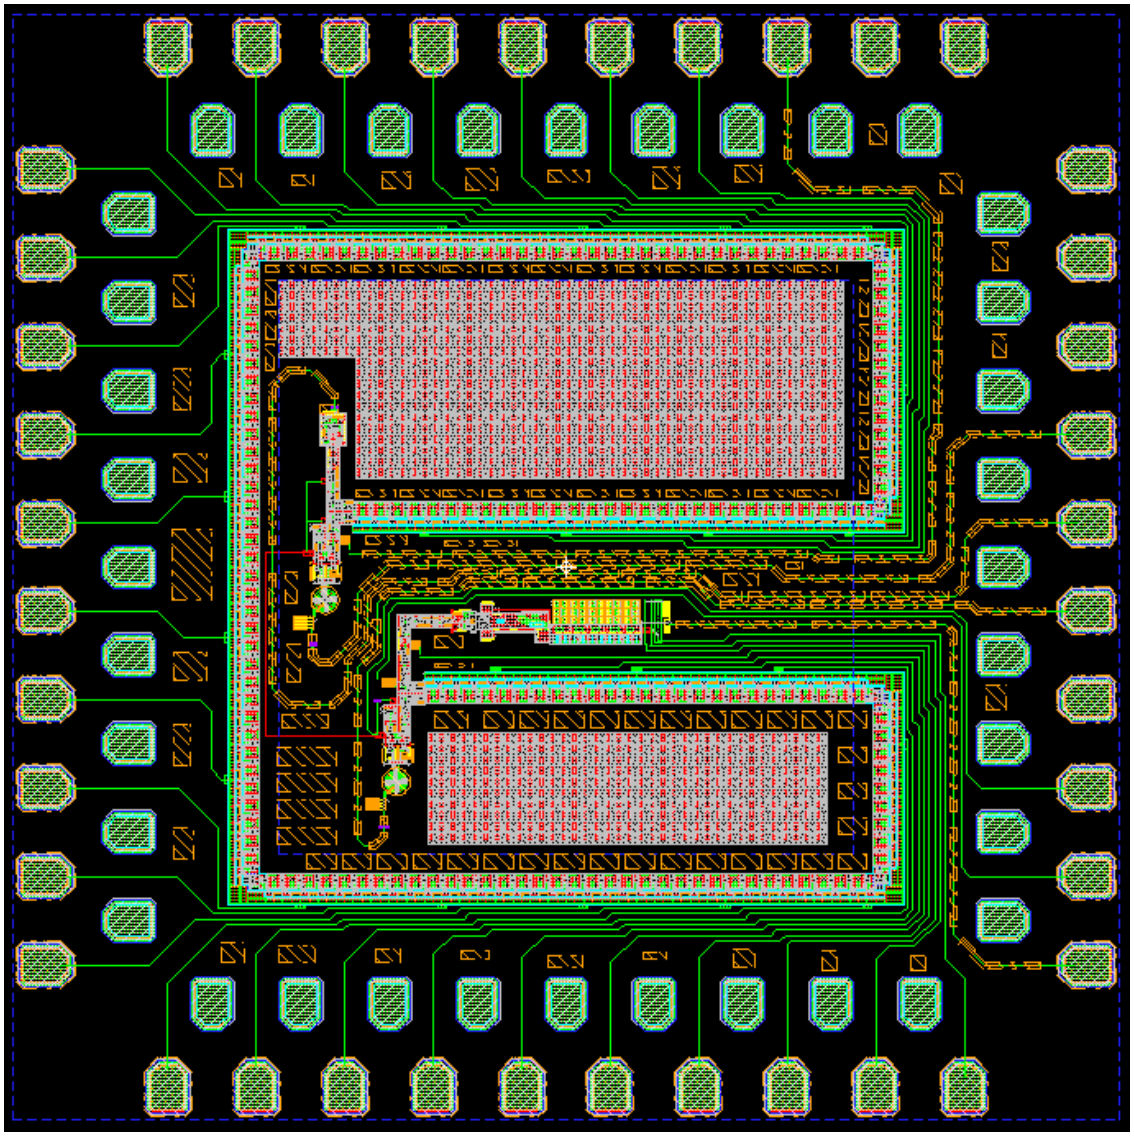

Figure A.39: Chip `mv352_esrlong_nodiv_hf`, to experimentally confirm operation and power dissipation of longer shift register, without a clock divider.

A.7.10 `mv352_esrlong_dsw_nodiv_hf` — Investigate pattern reproduction, no clock divider

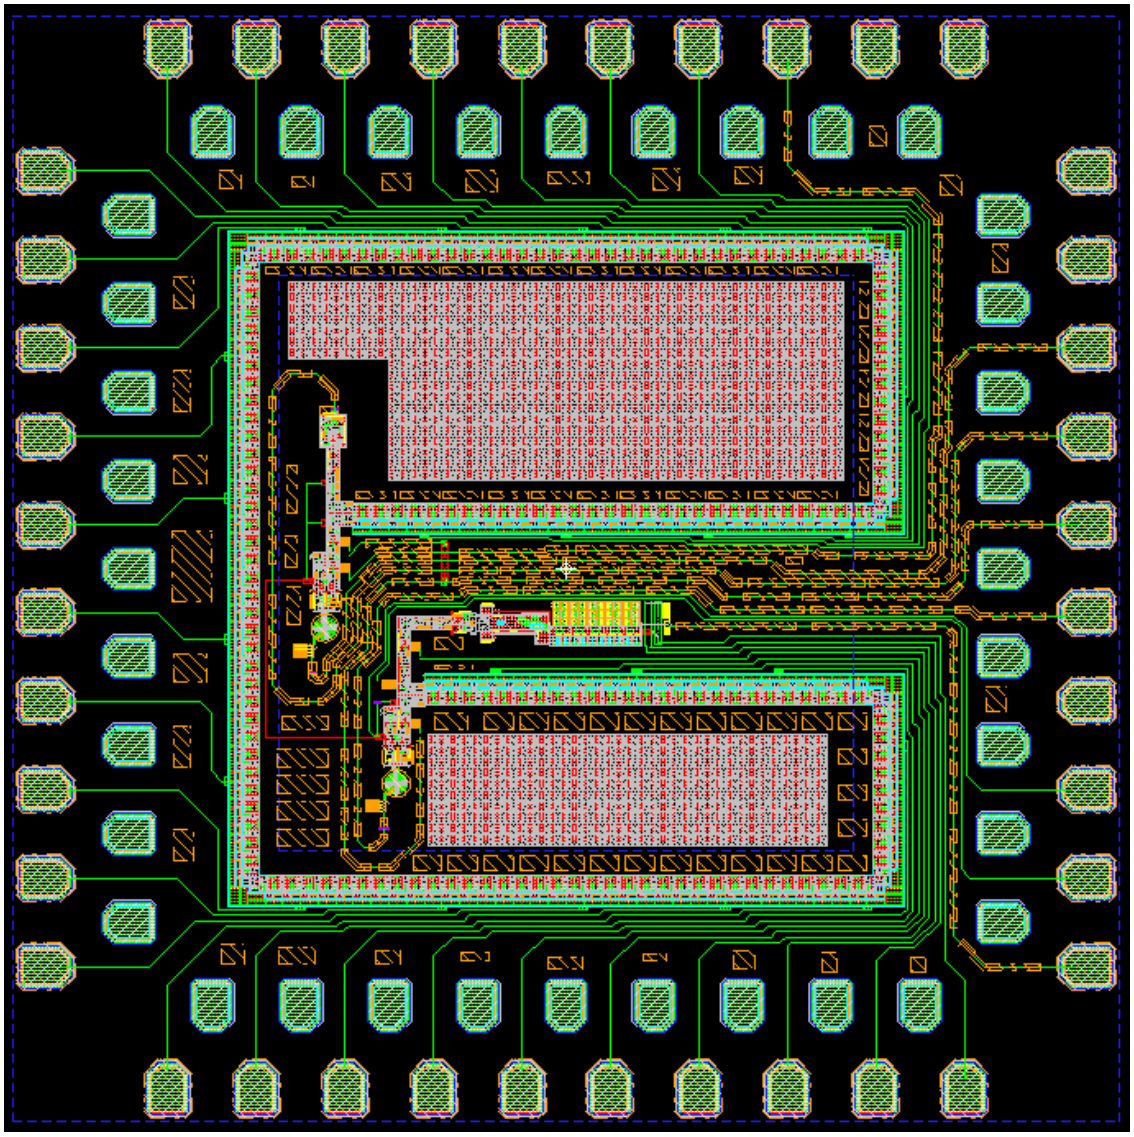

Figure A.40: Chip `mv352_esrlong_dsw_nodiv_hf`, to experimentally verify correct pattern reproduction of longer shift register at high speeds, without a clock divider.

A.7.11 mv352_pmd_edes16_b_hf — Investigate eSFQ deserialisers at high speeds

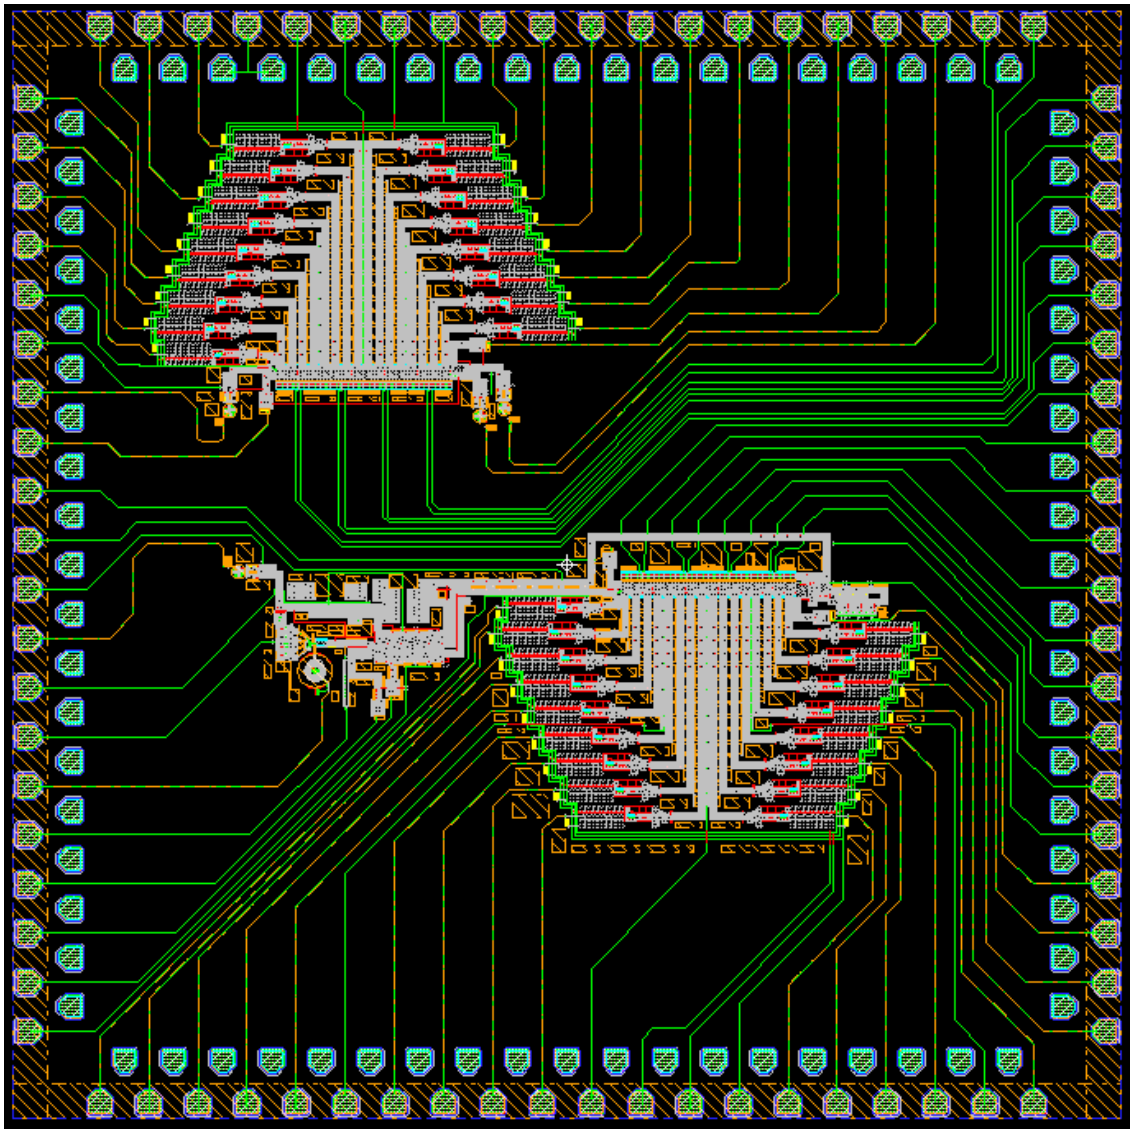

Figure A.41: Chip mv352_pmd_edes16_b_hf, to experimentally verify high-speed operation of eDES deserialisers, with correction.

A.7.12 `mv352_pmd_edes16_it2_hf` — Investigate eSFQ deserialisers at high speeds

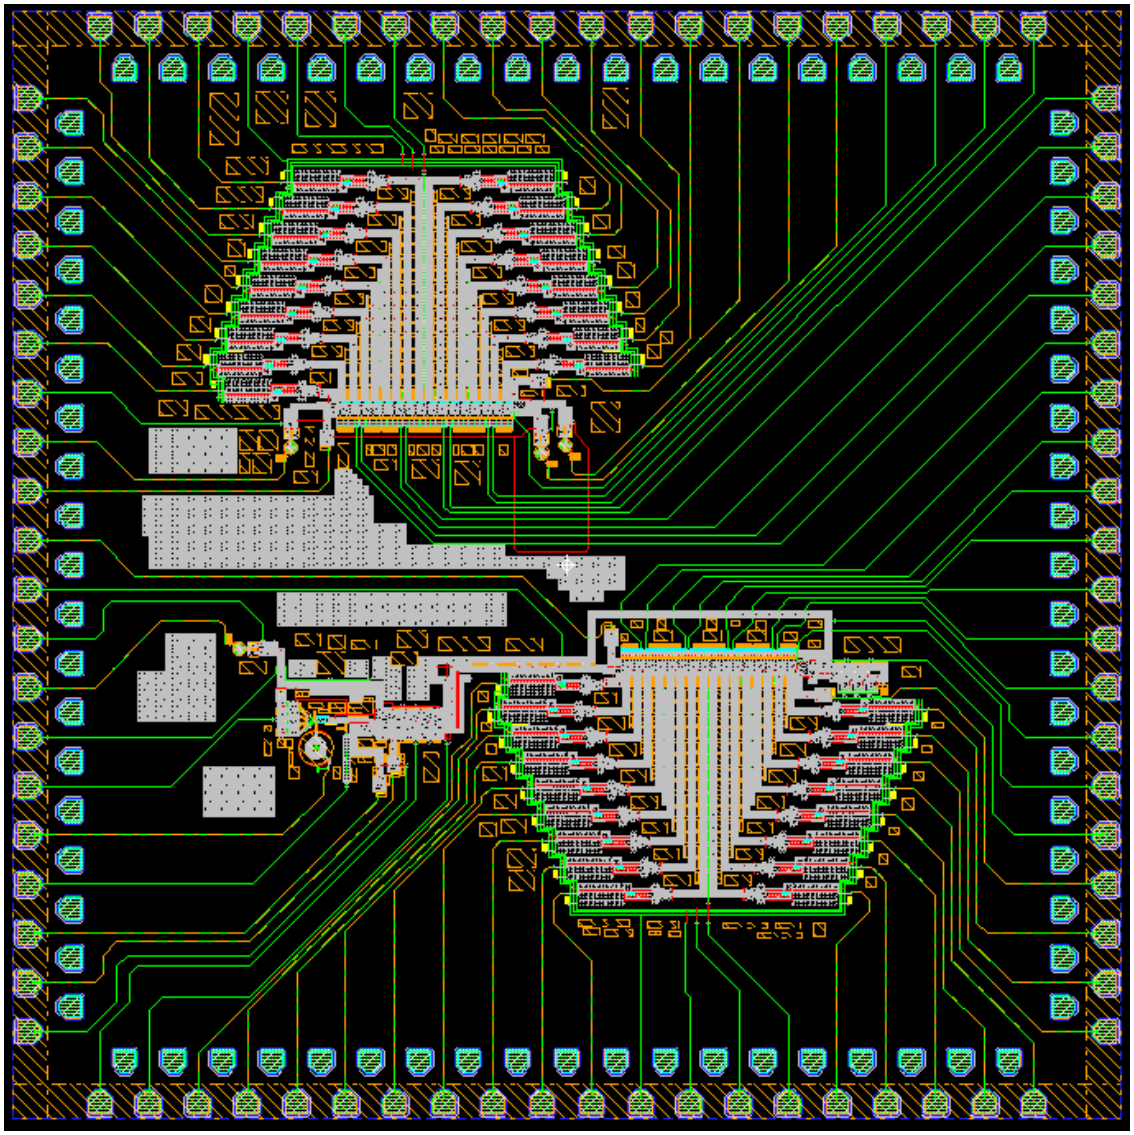

Figure A.42: Chip `mv352_pmd_edes16_it2_hf`, to experimentally verify high-speed operation of eDES deserialisers, with newer output drivers and better biasing.

Appendix B

Selected Code

B.1 Selected gEDA modifications (Scheme)

B.1.1 Modifications to support JSIM in backend `gnet-spice-sdb.scm`

Several modifications were made to add JSIM support to the gEDA framework. An extract of the Josephson junction translation code is shown in Listing B.1. Some of the code is retained from the original gEDA `spice-sdb` backend.

Listing B.1: Code for Josephson junction translation in `gnet-spice-sdb.scm`

```
1  ;
       ;-------------------------------------------------------------------------
2  ;;  write Josephson junction in wrspice format. Paul Bunyk, Sep 2,
       2005
3  ;;
4  ;;  Also enable write Josephson junction in JSIM format, in the
       presence
5  ;;  of the 'jsim' calling-flag. Mark Volkmann, Dec 8, 2011
6  ;
       ;-------------------------------------------------------------------------
7  (define spice-sdb:write-josephson-junction
8    (lambda (package port)
9
10     (debug-spew (string-append "Found␣Josephson␣junction.␣␣Refdes␣=␣"
           package "\n"))
11
12     ;; first write out refdes and attached nets
13     (spice-sdb:write-component-no-value package port)
14
15     ;; next, add a dummy node for JJ phase. Unlike in Xic netlister,
           give it
16     ;; a reasonable name, not a number, e.g., refdes.
17     ;; Omit this dummy node for JSIM.
18     (if (calling-flag? "jsim" (gnetlist:get-calling-flags))
19       (debug-spew "Omitting␣JJ␣phase␣node␣for␣JSIM.\n")
20       (display (string-append package "␣") port))
21
22     ;; next write JJ model name, if any.
```

```
23      (let ((model-name (gnetlist:get-package-attribute package "
            model-name")))
24          (if (not (string=? model-name "unknown"))
25                  (display (string-append model-name "␣" ) port))
26      )
27
28          (let ((value (gnetlist:get-package-attribute package "area")))
29                  (if (not (or (string=? value "unknown") (string=?
                        value "?")))
30                  (display (string-append "area=" value "␣" ) port)
31                  (if (spice-sdb:parameterise?) (display (string-append
                        "area=@{" package "}") port)))
32      )
33
34      ;; Next write out attribtes if they exist.  Use
35      ;; a list of attributes which can be attached to a junction.
36      ;;(let ((attrib-list (list "area")))
37      ;;  (spice-sdb:write-list-of-attributes package attrib-list port)
38      ;;          ;; write the off attribute separately
39          ;;      (display " " port))  ;; add additional space. . . .
40
41      (newline port)
42    )
43  )
```

B.1.2 Extract of modifications to support parametric modelling in `gnet-spice-sdb.scm`

An extract of the inductor translation code is reproduced in Listing B.2. Some of the code is retained from the original gEDA `spice-sdb` backend.

Listing B.2: Excerpt of code for inductor translation in `gnet-spice-sdb.scm`

```
1       ;; Next write inductor value, if any.  Note that if the
2       ;; component value is not assigned, then it will write "unknown",
3       ;; unless the ``param'' calling flag is present, in which case it
4       ;; will write ``@{<refdes>}'' to interface with the SU toolset.
5       (let ((value (gnetlist:get-package-attribute package "value")))
6                   (if (not (string=? value "unknown"))
7                   (display (string-append value "␣" ) port)
8                   (if (spice-sdb:parameterise?) (display (string-append
                        "@{" package "}") port)))
9       )
```

B.2 Selected signal-processing code (Python)

B.2.1 Program `adcn_sample`

Listing B.3: Implementation of `adcn_sample` program.

```
1   from numpy import * #array functions, such as sin, sqrt etc.
2   from numpy.random import randn, rand
3   from scipy.optimize import newton #for solving transcendental equation
4
5   phi0 = 2.06783376e-15
```

```
6
7  # Helper function for zero finding
8  def f1(i,ij,Ic,L):
9    return Ic*sin((2*pi*L/phi0)*(i - ij)) - ij
10
11 # Helper function for zero finding
12 def f2(i,Ic,L):
13   return newton(lambda x: f1(i+0.25*phi0/L, x, Ic, L), i)
14
15 # Perform quantisation function
16 def quant(xi,Ic,L):
17   ij = [f2(i,Ic,L) for i in xi]
18   return array(ij)
19
20 # Noisily discriminate signal
21 def ndisc(Iq,Ig):
22   IN = sqrt(2*pi)*Ig*randn(size(Iq))
23   return array(Iq + IN) >= 0    # Threshold assumed 0.
24
25 # Convert single bit vector from Gray code to single decimal
26 def Grayword2dec(word):
27   binarray = zeros(size(word),bool)
28   for i in range(len(word)):
29     if i == 0:
30       binarray[i] = word[i]
31     else:
32       binarray[i] = binarray[i-1]^word[i]
33   return int(''.join(['1' if b else '0' for b in binarray]), 2)
34
35 # Convert bit matrix from Gray code to decimal vector
36 def Gray2dec(bits):
37   N = len(bits[0])
38   bits = transpose(bits)
39   return array([Grayword2dec(bits[i]) for i in range(N)])
40
41 # Generate coded output decimal value
42 def coded(x,n,Ic,L,Ig):
43   bits = zeros((n,len(x)),bool)
44   for i in range(n):
45     dd = 2**(n-i)
46     bits[i] = ndisc(quant(x/dd,Ic,L),Ig)
47   return Gray2dec(bits)
48
49 # Convert output decimal value to floating-point value in input range
50 def decode(yhat,n,Delta):
51   ym = ((2**n)-1)/2.0
52   return (yhat - ym)*Delta
53
54 # Wring passed signal through noisy digitiser
55 def digitise(x,n,Ic,Ig,Delta):
56   L = phi0/(2*Delta)
57   xm = (((2**n))/2.0)*Delta
58   yhat = coded(x+xm,n,Ic,L,Ig)
59   y = decode(yhat,n,Delta)
60   et = y - x
61   return (y,x,et)
62
63 # Generate some random uniform samples
64 def makex(N,n,Delta):
65   return array(2**n*Delta*(rand(N)-0.5))
66
67 # Entry point to program
68 def adcn_sample(n,N,Ic,Ig,Delta):
```

```
69    x = makex(N,n,Delta)
70    return digitise(x,n,Ic,Ig,Delta)
```

B.2.2 Program snr_estimator

Listing B.4: Implementation of snr_estimator program.

```
1   import adcn_sample                          # Functions for sample
        generation
2   from scipy.special import erfinv            # The inverse error function
3   from numpy import *                         # Some useful functions
4   from joblib import Parallel, delayed        # Structures for
        parallelisation
5
6   Nthreads = 8     # Number of threads
7
8   # The quantile function, for confidence bounding
9   def quantile(halfalpha):
10    return abs(sqrt(2)*erfinv(2*halfalpha-1))
11
12  # Actual worker to be enveloped by thread
13  def worker(n,N,Ic,Ig,Delta):
14    (y,x,et) = adcn_sample.generate_samples(n,N,Ic,Ig,Delta)
15    Pe = sum((et*et))/N
16    Py = sum((y*y))/N
17    return 10*log(Py/Pe)/log(10)
18
19  # Infrastructure for parallelisation
20  def batch_job(Njobs,n,N,Ic,Ig,Delta):
21    return Parallel(n_jobs=Nthreads)(delayed(worker)(n,N,Ic,Ig,Delta)
          for i in range(Njobs))
22
23  # Entry point to program
24  #   theta = (n,Delta,Ic,Ig)
25  #   psi = (N,alpha,beta)
26  def snr_estimator(theta,psi):
27    (n,Delta,Ic,Ig) = theta
28    (N,alpha,beta) = psi
29    SNRdBs = []
30    epsilon = 10**12
31    muhat = 0
32    R = 0
33    z = quantile(alpha/2)
34    while R < 50 or epsilon > beta*muhat:    # While confidence not
          reached
35      SNRdBs += batch_job(4*Nthreads,n,N,Ic,Ig,Delta)
36      R = R+4*Nthreads
37      muhat = mean(SNRdBs)
38      sigmahat = sqrt(var(SNRdBs))
39      if R > 1:
40        epsilon = z*sigmahat/sqrt(R)
41      if R == 1 or R % (Nthreads*4) == 0:
42        print 'It␣{}:␣SNR␣=␣{},␣betahat␣=␣{}'.format(R,muhat,epsilon/
              muhat)
43    print "SNR␣=␣{}␣(epsilon␣=␣{},␣alpha␣=␣{},␣R␣=␣{})".format(muhat,
          epsilon, alpha, R)
44    return (muhat,epsilon,R)
45
46  # Return conservative estimate (lower bound of confidence interval)
47  def snr_flash_conservative(theta,psi):
```

```
48    (muhat,epsilon,R) = snr_flash(theta,psi)
49    return muhat - epsilon
```

B.2.3 Function qcomp_pg

Listing B.5: Implementation of qcomp_pg function which simulates noisy comparator operation.

```
1   import ibceval     # Switching probability curve functions
2
3   phi0 = 2.06783376e-15
4
5   # Simulate comparator defined by rhobar.
6   #  rhobar = (Ith1r,Ig1r,Ith2r,Ig2r,Ith1f,Ig1f,Ith2f,Ig2f,Im)
7   def qcomp_pq(xi,rhobar):
8     (Ith1r,Ig1r,Ith2r,Ig2r,Ith1f,Ig1f,Ith2f,Ig2f,Im) = rhobar
9     # Find thresholds to enable bin selection.
10    Ithr = (Ig2r*Ith1r + Ig1r*Ith2r)/(Ig1r+Ig2r)
11    Ithf = (Ig2f*Ith1f + Ig1f*Ith2f)/(Ig1f+Ig2f)
12    defi = 0.25*Im - Ithr
13    xd = fmod(xi + defi,Im)
14    xk = zeros(size(xd))
15    ik = zeros(size(xd),dtype=int)
16    out = zeros(size(xd),dtype=bool)
17    for k in range(len(xd)): # perform bin selection
18      if abs(xd[k] - Ithr) < abs(xd[k] - Ithf):
19        ik[k] = 1
20        xk[k] = xd[k] - defi
21      else:
22        ik[k] = 2
23        xk[k] = xd[k] - defi
24    for k in range(len(xd)): # employ relevant probability curve
25      if ik[k] == 1:
26        xtest = ibceval.sprob2([xk[k]],Ith1r,Ig1r,Ith2r,Ig2r)[0]
27      else:
28        xtest = ibceval.sprobfall2([xk[k]],Ith1f,Ig1f,Ith2f,Ig2f)[0]
29      out[k] = True if rand() < xtest else False # Noisy discrimination
30    return out
```