

Inductance in Superconductor Integrated Circuits

by

Coenrad Johann Fourie

*Dissertation presented for the degree of Doctorate in Engineering in the
Faculty of Engineering at Stellenbosch University*



Supervisor: Prof. Dirk I. L. de Villiers

March 2023

Declaration

By submitting this thesis 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.

March 2023

Copyright © 2023 Stellenbosch University

All rights reserved

Abstract

Inductance in Superconductor Integrated Circuits

C. J. Fourie

Dissertation: DEng

March 2023

This dissertation presents an overview of the research and publications of the candidate and his research group on the design of superconductor integrated circuits around inductance as a key circuit parameter, on the development and verification of inductance extraction tools for complex, three-dimensional integrated circuit models, and on the application of self- and mutual inductance extraction and magnetic field analysis to the improvement of superconductor circuit and system design. The research spans more than two decades, and culminates in the extraction of compact simulation models for the analysis of superconductor integrated circuits in the presence of trapped flux and external magnetic fields, which was not previously possible. The golden thread that ties all of his work together is inductance in superconductor integrated circuits.

Uittreksel

Induktansie in Supergeleier Geïntegreerde Stroombane

C. J. Fourie
Proefskrif: DIng
Maart 2023

Hierdie proefskrif bied 'n oorsig van die navorsing en publikasies van die kandidaat en sy navorsingsgroep op die ontwerp van supergeleier geïntegreerde stroombane rondom induktansie as 'n kernparameter, op die ontwikkeling en verifikasie van induktansie-onttrekkingsagteware vir komplekse, drie-dimensionele geïntegreerde stroombaanmodelle, en op die toepassing van self- en wedersydse induktansie-onttrekking en magnetiese veldanalises op die verbetering van supergeleier stroombaan- en stelselontwerp. Die navorsing strek oor meer as twee dekades, en lei tot die onttrekking van kompakte simulasiemodelle vir die analise van supergeleier geïntegreerde stroombane in die teenwoordigheid van vasgevangde magnetiese vloed en eksterne magnetiese velde. Hierdie vermoë was nie voorheen beskikbaar nie. Die goue draad wat al sy werk saambind is induktansie in supergeleier geïntegreerde stroombane.

Acknowledgements

Firstly, a big thanks to my family – my wife Dr Frana Fourie and our daughters Aletta and Lize – who put up with my never-ending work-at-home hours (more so during the COVID-19 epidemic) and my extended international research visits which I could only fractionally redeem by carting back home the odd *かわい* soft toy, Lego set, exotic stationary collection or affordable electronics. Your love and support gave my life substance and meaning.

As for my research career: anything meaningful that I have accomplished was in service of, or in collaboration with the community of physicists and engineers in applied superconductivity without whose support, encouragement and critique I would not have had the motivation to persist.

I would like to thank Professor Willem Perold who supervised my postgraduate studies and introduced me to the fine art of research with tangible outputs. Willie, as I could informally call him after our first Samuel Adams beer in the bar of a budget hotel in New Jersey, scraped together money to send me to international conferences that would shape my career. I am still striving to offer my own students the same opportunities.

More thanks go to the other professors who probably did not realise how much of an impact they had: Johannes Cloete as *The Engineer*, the respected lead figure of the Electronics and Electromagnetics group at Stellenbosch University's E&E Engineering Department, who forged a collegial environment between the professors and the postgraduate students with extended weekly late-morning coffee visits to Spice Café in Stellenbosch. Petrie Meyer, Keith Palmer and Johann de Swardt were part of that group that left indelible imprints on my training as engineer, researcher and academic.

My postgraduate students did a lot of the heavy lifting. I am thankful to them all, but have to thank a few by name: Mark Volkmann, who demanded a better InductEx and was quite happy to teach me how to code it; Kyle Jackman, who *made* InductEx better; Johannes Delpont, who created JoSIM and can quite likely compile an OS on an abacus with sufficient beads; Lieze Johnston (née Schindler) who spun a cell library from the wispiest of starting instructions; and Ruben van Staden, who dreams bigger than all the rest of us combined. Of these, I owe a debt of gratitude to the postdoc engineers, Kyle, Johannes and Lieze, who became a close group of friends through our frequent travels – especially so after one very memorable long-haul flight back home out of LAX – and who found great delight in exploring with me the culinary wonders that the world could offer us: sushi, shabu-shabu and sake in Japan; cheese, burn-it-yourself Crème Brûlée and *les vins des Côtes du Rhône* in France; single malt whisky (straight from the distillery) in Scotland; and Red Robin Whiskey River BBQ burgers all across the US. Without them, my efforts under ColdFlux could never have succeeded.

Lastly, I would like to thank all the international researchers and colleagues who helped shape my career and research. Their research exploits are published, the stuff of legend, or both. But it is the personal time they afforded me that immersed me in their cultures and

customs and immeasurably enriched my life experience. Pascal Febvre, who hiked with me across cities like Paris, Yokohama, Tampa, or middle-of-nowhere stretches of *veld* in South Africa, and who gate-crashed a private Bon Jovi concert with me in Charlotte, North Carolina, one night after ASC; Nobuyuki Yoshikawa, who taught me how to appreciate Japanese culture and cuisine (and 山崎!) and travelled with me by car, Shinkansen, local rail and foot all over Japan in search of temples, castles, mountains and the perfect ramen shop – whether it is tucked under the foot of Mount Fuji or the glittering hulk of Yokohama station; Thomas Ortlepp, who showed me the wonders of the *Autobahn* and in whose home I watched South Africa crush England in the 2007 Rugby World Cup final; Hannes Toepfer, who hiked with me through the Thuringian hills in search of blueberries and trout; Oleg Mukhanov, who invited me onto the ISEC international advisory board, inspired me to visit exotic places (when Oleg tells you about any place he has been to, you *want* to go there), and treated my family to memorable dinners on the Hudson; Vasili Semenov, who invited my family for Easter dinner at home, patiently let my small children feed his bread to wild geese on Long Island and unwittingly introduced them to Masha and the Bear; Massoud Pedram, who let me experience some of the finest cross-cultural dining in Los Angeles (almost always with a bottle of Oregon Pinot Noir); Denis Cr  t  , who taught me to dine like a Frenchman (even for lunch); Christopher Ayala and Olivia Chen, who showed my team and I around Japanese cities at night, and to whom we could show the wondrous night sky of the Southern Hemisphere; Igor Vernik, Alex Kirichenko and Timur Filippov, who all treated me to Russian hospitality in America; and Scott Holmes, SETA for SuperTools and its seedling, who immersed me in American history through interesting road trips, visits to museums and Revolutionary War fortifications, and the thorough exploration of Washington, DC, by foot.

Lastly, the research presented here is based upon work supported financially by the South African National Research Foundation, grant numbers 69006, 78789, 86021, 92426, 93586, 95237, 105859 and 120459, with additions developed under contracts W911NF14-C-0089 and W911NF-17-1-0120 with the Intelligence Advanced Research Projects Activity (IARPA)

Contents

Declaration	i
Abstract	ii
Uittreksel	iii
Acknowledgements	iv
Nomenclature	x
List of Figures	xviii
List of Tables	xix
1 Introduction	1
1.1 Background	1
1.2 A short profile	1
1.3 Contributions	3
1.4 Layout of the dissertation	4
2 Superconductor integrated circuits	5
2.1 Background	5
2.1.1 Flux quantization	6
2.1.2 The Josephson junction	6
2.1.3 The superconducting quantum interference device	10
2.2 Fabrication processes	12
2.2.1 Dimensions	12
2.2.2 Steps in IC fabrication	12
2.2.2.1 Wafer preparation	13
2.2.2.2 Deposition	13
2.2.2.3 Oxidation	13
2.2.2.4 Planarisation	13
2.2.2.5 Photolithography	13
2.2.2.6 Etching	13
2.2.2.7 Anodization	14
2.2.3 Monolayer fabrication processes	14
2.2.4 The main fabrication processes	14
2.2.4.1 Leibniz IPHT - FLUXONICS 1 kA cm^{-2}	15
2.2.4.2 Seeqc - #QC1000A 1 kA cm^{-2}	18
2.2.4.3 AIST - standard 2.5 kA cm^{-2} and advanced 10 kA cm^{-2}	19

2.2.4.4	MIT Lincoln Laboratory - SFQ5ee 10 kA cm ⁻²	20
2.3	Early superconductor digital circuits	21
2.4	The Rapid Single-Flux Quantum logic family	22
2.4.1	Origins of RSFQ	22
2.4.2	Data storage and transmission in RSFQ circuits	23
2.4.3	RSFQ circuit theory	25
2.4.3.1	Phase-based equations	25
2.4.3.2	Basic Josephson transmission line	27
2.4.3.3	Symmetrical Josephson transmission line	29
2.4.3.4	Basic storage element - the D Flip-Flop	30
2.4.3.5	Bias resistors	34
2.4.3.6	Design conclusion	35
2.5	Contributions to RSFQ	35
2.5.1	Logic cells	35
2.5.2	Special-purpose cells	36
2.5.2.1	DCRL	36
2.5.2.2	RSFQ-COSL output driver	36
2.5.3	A superconductor programmable gate array	37
2.5.4	Asynchronous logic: RSFQ-AT	43
2.5.5	Cell libraries	44
2.5.5.1	Routing architecture	45
2.5.5.2	Routing track block	47
2.5.5.3	ColdFlux RSFQ cell library	50
2.5.6	Microprocessors	53
2.6	Beyond RSFQ: ultra-low power logic	53
2.6.1	Low voltage bias	55
2.6.2	ERSFQ	55
2.6.3	eSFQ	57
2.6.3.1	eSFQ shift register	59
2.6.3.2	An eSFQ T flip-flop	60
2.6.4	Adiabatic Quantum Flux Parametron logic	63
2.7	Summary	66
3	Inductance calculation	67
3.1	Theory	67
3.1.1	Inductance of a conductor	67
3.1.2	Inductance in a superconductor circuit	67
3.2	Background	68
3.2.1	Superconductor integrated circuits	68
3.2.2	Parameter extraction	69
3.2.2.1	Resistance extraction	69
3.2.2.2	Capacitance extraction	69
3.2.2.3	Inductance extraction	69
3.2.3	Known inductance extraction tools	71
3.2.3.1	Normal conductors	71
3.2.3.2	Superconductors	71
3.3	InductEx: Three-dimensional inductance calculation	72
3.3.1	Early contributions	72

3.3.1.1	Meshing	75
3.3.1.2	Method of images	77
3.3.2	The first InductEx	79
3.3.3	Multiterminal netlists	81
3.3.4	Validation	83
3.3.4.1	Validation of inductance over holes	86
3.3.5	Fabrication-ready layout processing	88
3.3.5.1	Full-circuit layouts	88
3.3.5.2	Resistance	89
3.3.6	Solution speedup	91
3.3.6.1	FastHenry overview	92
3.3.6.2	FastHenry characterisation	93
3.3.6.3	Fast FastHenry (FFH)	95
3.3.7	A new engine: TetraHenry	97
3.3.7.1	Tetrahedral modelling	97
3.3.7.2	Volume Integral Equation	98
3.3.7.3	Discretization	98
3.3.7.4	Method of Moments	99
3.3.7.5	Volume Loop Basis Function	100
3.3.7.6	Electrostatic Analogy	102
3.3.7.7	Iterative Solver and Preconditioning	102
3.3.7.8	Results: Small Superconducting Structures	103
3.3.8	Sheet currents, triangular and hybrid meshes	104
3.3.8.1	Derivation of Surface Integral Equation	105
3.3.8.2	Discretization	106
3.3.8.3	Surface loop basis function	108
3.3.8.4	Hybrid Meshing	109
3.3.9	Coupling from flux trapped in holes	111
3.3.10	Coupling from external fields	113
3.3.11	Compact simulation models	115
3.3.11.1	Errors in extracted results	115
3.3.11.2	Fundamental cycles	116
3.4	Experimental verification	117
3.4.1	Measurement of inductance	117
3.4.2	Published inductance results	121
3.4.3	Calibration	122
3.4.3.1	Hypres fabrication processes	122
3.4.3.2	AIST processes	125
3.4.3.3	FLUXONICS process	125
3.4.3.4	MIT Lincoln Laboratory SFQ4ee and SFQ5ee processes	125
3.4.4	Mutual inductance in sub-micron structures	126
3.5	Conclusion on contributions	131
4	Tool chain	132
4.1	Background	132
4.2	Electrical simulation engine: JoSIM	133
4.2.1	History	133
4.2.1.1	Modified nodal analysis	134

4.2.2	JoSIM	134
4.2.2.1	Integration method	135
4.2.2.2	MNA component stamps	135
4.2.2.3	JoSIM application	137
4.3	Device level tools	139
4.3.1	Technology CAD	139
4.4	Cell level tools	142
4.4.1	Characterisation	142
4.4.2	Optimisation	144
4.4.3	Timing extraction and state verification	146
4.5	Chip level tools	148
4.5.1	Interconnect analysis	148
4.5.2	Synthesis, placement and routing	148
4.5.3	Static timing analysis	148
4.5.4	Layout-versus-schematic verification	148
4.6	Summary	149
5	Application	150
5.1	Ground planes and return currents	150
5.1.0.1	Single ground plane	150
5.1.0.2	Ground contacts	150
5.1.0.3	Multiple ground planes	151
5.2	Magnetic fields	154
5.3	Digital circuits	154
5.3.1	Coupling from bias lines	154
5.3.2	Inductive SFQ pulse transfer	156
5.4	Analogue devices: SQUID magnetometers	159
5.4.1	SQUID parameters of interest	159
5.4.2	Analysis of a planar direct-coupled SQUID: the M2700	160
5.4.2.1	Full inductance circuit model	161
5.4.2.2	Compact simulation model	164
5.4.2.3	Single inductance model	164
5.4.2.4	Feedback coil coupling	165
5.4.3	Flux noise	166
5.5	Flux trapping	166
5.5.1	Background to flux trapping	166
5.5.2	Modelling of flux trapping	167
5.5.3	Verification of flux trapping	168
6	Final conclusion	172
	References	173

Nomenclature

- AQFP** Adiabatic quantum flux parametron
- BDF** Backward differential formula
- CMOS** Complementary metal-oxide semiconductor
- CMP** Chemical mechanical polishing
- COSL** Complementary output switching logic
- CPR** Current-phase relation
- DC** Direct current
- DCRL** DC-resettable latch
- DDF** D flip-flop
- DMP** Decision-making pair
- DRO** Destructive readout register (also a DFF)
- DSNDO** Destructive-shift nondestructive read-out
- DUT** Device under test
- EDA** Electronic design automation
- ERSFQ** Energy-efficient rapid single flux quantum
- eSR** eSFQ shift register
- FDM** Finite difference method
- FEM** Finite element method
- FIFO** First-in-first-out
- FFH** Fast FastHenry
- FJTL** Feeding Josephson transmission line
- FMM** Fast multipole method
- FPGA** Field-programmable gate array
- GMRES** Generalised minimal residual method

HDL Hardware description language
HDPE High-density plasma etching
HFQ Half-flux-quantum
HTS High temperature superconductor
IC Integrated circuit
JTL Josephson transmission line
KCL Kirchhoff's current law
KVL Kirchhoff's voltage law
LHS Left hand side
LTS Low temperature superconductor
LUT Lookup table
LVS Layout-versus-schematic
MEMS Microelectromechanical system
MeSR Magnetic flux-biased eSFQ shift register
MNA Modified nodal analysis
MNPA Modified nodal phase analysis
MoM Method of moments
MPS Multipole setup
MQS Magnetoquasistatic
PCB Printed circuit board
PDN Power distribution network
PECVD Plasma-enhanced chemical vapour deposition
PTL Passive transmission line
QFP Quantum flux parametron
RAM Random access memory
RCA Ripple carry adder
RCSJ Resistively and capacitively shunted junction
RF Radio frequency
RHS Right hand side

RMSE Root mean square error

RSFQ Rapid single flux quantum

RSFQ-AT RSFQ asynchronous transmission

RWG Rao-Wilton-Glisson

SC Superconductor/superconducting

SCE Superconductor electronics

SI *Système International d'Unités* – International System of Units

SIS Superconductor-insulator-superconductor

SL Surface loop

SPGA Superconducting programmable gate array

SQUID Superconducting quantum interference device

STA Static timing analysis

SVD Singular value decomposition

SWG Schaubert-Wilton-Glisson

TCAD Technology computer-aided design

TFF T flip-flop

TSV Through-silicon via

TTH TetraHenry

VIE Volume integral equation

VJIE Volume electric current integral equation

VL Volume loop

List of Figures

2.1	A superconductor-insulator-superconductor Josephson junction and circuit symbol.	7
2.2	Scanning electron microscope image of the cross-section of an integrated circuit SIS Josephson junction from the MITLL SFQ4ee process.	8
2.3	Circuit schematic for the generalised RCSJ model.	8
2.4	Experimentally measured I-V curve of a niobium-aluminium-oxide-niobium planar SIS junction.	9
2.5	Circuit schematic for the generalised RCSJ model with an external shunt resistor.	10
2.6	A schematic of a two-junction SQUID with a symmetrical feed.	11
2.7	Circuit schematic of a dc SQUID with generalised Josephson junctions. . .	11
2.8	Microphotograph of SQUIDS in a monolayer process.	14
2.9	InductEx model of a monolayer SQUID.	15
2.10	The FLUXONICS process layer stack.	16
2.11	Cross-section of a shunted Josephson junction fabricated with the FLUXONICS process.	16
2.12	Cross-section of a model of a shunted Josephson junction.	17
2.13	The Seecq QC1000A process layer stack.	19
2.14	Cross-section of a model of a JTL in the Seecq QC1000 process.	19
2.15	The AIST ADP2 process layer stack.	20
2.16	The MIT Lincoln Laboratory SFQ5ee layer stack.	21
2.17	Electron microscope photograph of a manufactured COSL gate.	22
2.18	Experimental measurements of a COSL OR gate and a COSL NAND gate at 8 GHz.	22
2.19	RSFQ basic circuit blocks.	23
2.20	Basic RSFQ OR gate schematic.	24
2.21	Simulated transient response of RSFQ OR gate.	24
2.22	Current-phase relation of an inductor.	25
2.23	Current-phase relation of two inductors coupled with mutual inductance. .	26
2.24	A basic RSFQ Josephson transmission line.	27
2.25	Simulated phase response of a basic JTL chain.	29
2.26	A symmetrical RSFQ Josephson transmission line.	29
2.27	A basic RSFQ D flip-flop.	30
2.28	Mealy state diagram of the RSFQ DFF.	32
2.29	Simulated response of RSFQ DFF for set and reset inputs.	33
2.30	Margins of first-pass DFF design without a test for a set input in the “1” state.	33
2.31	Margins of first-pass DFF design when a set input in the “1” state is included. .	33
2.32	Optimum margins of DFF for all possible input combinations.	34

2.33	Schematic of a JTL showing resistive biasing and a common voltage rail. . .	35
2.34	Circuit schematic of an RSFQ DC-resettable latch.	36
2.35	Microphotograph of an RSFQ DC-resettable latch fabricated with the FLUX-ONICS process.	37
2.36	Circuit schematic for an RSFQ-to-COSL converter.	37
2.37	Simulated transient response of the RSFQ-COSL converter at a clock frequency of 10 GHz.	38
2.38	Microphotograph of RSFQ-to-COSL converter manufactured in Hypres 4.5 kA cm ⁻² process.	38
2.39	Schematic diagram of a general architecture for a symmetrical array FPGA.	39
2.40	Schematic of a 4-LUT SPGA with routing architecture and switches.	40
2.41	Layout of a 4-LUT SPGA with HUFFLE bipolar drivers on a 5 mm×5 mm chip for the Hypres 4.5 kA cm ⁻² process.	41
2.42	Connection diagram of a 4 track wide unidirectional Wilton switch block.	42
2.43	VPR place and route visualisation showing the use of configurable logic blocks for a combinational 8-bit ripple carry adder.	43
2.44	Schematic diagram of a generic two-input RSFQ-AT logic cell.	44
2.45	RSFQ-AT implementation of the half-adder and full-adder.	44
2.46	Simulated response of the RSFQ-AT full-adder.	45
2.47	Schematic representation of the row-based place and route strategy used for RSFQ circuits in ColdFlux.	46
2.48	Simplified illustration of a cross-section of the ColdFlux layout stack in the MIT-LL SFQ5ee process showing the assignment of passive transmission line layers.	48
2.49	Dimensions of the basic routing track block.	48
2.50	Cross-section of a three-dimensional simulation model for the M3-to-M1 stripline transition with an optimally filled via.	48
2.51	Three-dimensional rendering of an arbitrary 4 × 2 track block composition for the MIT-LL SFQ5ee process.	49
2.52	An RSFQ splitter cell layout that fits the routing block architecture.	50
2.53	Layout image of an RSFQ OR2T cell in the ColdFlux library.	53
2.54	System schematic of a basic microprocessor.	54
2.55	Conventional RSFQ biasing.	54
2.56	Exploitation of the current limiting properties of the Josephson junction to achieve desired bias current distribution in an ERSFQ circuit.	56
2.57	InductEx inductance extraction model for bias section of an ERSFQ circuit.	57
2.58	The eSFQ biasing principle.	58
2.59	Circuit schematic of an RSFQ D flip-flop and the eSFQ conversion of the D flip-flop.	58
2.60	Circuit schematic and simulated response of an eSFQ shift register cell – the eSR.	59
2.61	Circuit schematic and simulated response of an eSFQ shift register cell with magnetic flux bias – the MeSR.	60
2.62	Microphotographs of the eSR cell and the MeSR cell fabricated with the Hypres 4.5 kA cm ⁻² process.	61
2.63	InductEx model of the bias junction and inductor for the ESR, matching the maximum size investigated in 2011.	61
2.64	Modern InductEx model, with cuboid segments, of the eSR.	62

2.65	Microphotograph of an eSFQ TFF.	62
2.66	InductEx model for inductance extraction of the eSFQ TFF.	63
2.67	AQFP buffer cell schematic and simulation test circuit.	64
2.68	AQFP buffer cell layout drawing with InductEx ports for inductance extraction and flux trapping analysis.	65
2.69	JoSIM simulation result of AQFP buffer chain no fluxons in any moat and with three fluxons in moat FBL.	65
2.70	InductEx model of the AQFP buffer cell with current distribution generated by a positively oriented fluxon in moat FBL.	66
3.1	A straight line microstrip.	72
3.2	Dimensions of a straight line microstrip.	73
3.3	Traditional approximation of the effective path length for inductance estimation around a corner in a thin-film inductor.	74
3.4	Typical chip package section for analysis with FastHenry and individual self and mutual inductances resulting from FastHenry solution.	75
3.5	FastHenry segment shown with 3 height filaments, 5 width filaments and node for connection.	75
3.6	Common layout structures for inductance calculation.	76
3.7	Segmented line with interleaved cuboid segments.	76
3.8	Current distribution in the lowest filaments of a superconducting conductor and the highest filaments in a superconducting ground plane calculated with a cuboid mesh.	77
3.9	Inductance of superconducting microstrip calculated with cuboid mesh and normalized to the smallest solution.	78
3.10	Position of the reflection plane at λ_{eff} for a superconducting microstrip over ground.	79
3.11	Simplified top view of part of a Josephson transmission line with two junctions connected by an inductor, showing the cake-slicing segmentation process.	80
3.12	Microphotograph and InductEx model of a section of an RSFQ-to-COSL converter fabricated in the Hypres 1 kA cm ⁻² process.	81
3.13	A circuit with inductance, resistance and mutual inductance.	82
3.14	Segmented models and microphotographs of SQUID layouts in the IPHT RSFQ1D process.	85
3.15	Microphotograph of a reference SQUID and several SQUIDs with varying size ground plane holes underneath the loop inductor.	86
3.16	InductEx model of a SQUID meshed with cuboid segments and the extraction netlist with all excitation ports and inductors.	87
3.17	Measured and extracted inductance results for L_{loop} of test SQUIDs manufactured with the FLUXONICS process with loop inductors over ground plane holes.	87
3.18	Microphotographs of circuits with ground plane holes under inductors manufactured with the FLUXONICS process.	88
3.19	InductEx extraction model of an LR-biased RSFQ Toggle-flip-flop.	89
3.20	Schematics of the layer stack surrounding a resistive layer between two superconductive metal layers for popular fabrication processes in 2014.	90

3.21	Excerpts from the layer definition files for the layer stacks between the nearest metal layers above and below the resistive layer for popular fabrication processes in 2014.	91
3.22	A conductor excited by a voltage source, with discretised filaments connected to nodes, and modelled as a circuit.	92
3.23	InductEx calculation models for the FastHenry engine.	94
3.24	Breakdown of the time spent on the main solution steps in FastHenry for different extraction models.	95
3.25	Full-SWG basis functions in arbitrary body with piecewise constant electrical parameters.	99
3.26	Full-SWG basis function.	99
3.27	Close volume loop basis function and unclosed volume loop basis function.	100
3.28	Top view of tetrahedral mesh of rectangular conductor with two terminals.	101
3.29	Convergence rate of GMRES for the microstrip line example and the multilayer example.	103
3.30	Current density of a microstrip line above a ground layer.	104
3.31	Current density of a multilayer example with coupled structures.	104
3.32	Visualisation of the sheet current model for triangle T_m^+ , with projected triangles at heights h_m^0 and h_m^1	106
3.33	RWG basis function at material interface with different conductivities.	107
3.34	Closed and unclosed surface loop basis functions.	108
3.35	Hybrid loop basis function, consisting of both RWG and SWG basis functions.	109
3.36	Current density of a $50 \mu\text{m} \times 5 \mu\text{m}$ microstrip line (triangular meshing) with a via-interconnect (tetrahedral meshing).	110
3.37	The inductance of a microstrip line, with a via-interconnect, for different meshing techniques.	110
3.38	Definition of paths for every hole in an extraction model.	111
3.39	A circuit model for two holes coupled to a circuit with inductance, resistance and mutual inductance.	112
3.40	A circuit model for a moat and an external magnetic field coupled to a circuit with inductance, resistance and mutual inductance.	114
3.41	Graph and schematic of AQFP buffer cell with fundamental inductors.	117
3.42	Basic equivalent circuit of a dc SQUID.	118
3.43	Equivalent circuit of a dc SQUID used for inductance measurement on an integrated circuit.	118
3.44	InductEx model of an inductance measurement SQUID for the Hypres 4.5 kA cm^{-2} process with loop inductor in layer M2.	119
3.45	Measured SQUID voltage as a function of modulation current for a dc SQUID fabricated with the Hypres 4.5 kA cm^{-2} process with loop inductor in layer M2.	120
3.46	Equivalent circuit of a differential-arm dc SQUID used for inductance measurement on an integrated circuit.	120
3.47	Measured SQUID voltage as a function of control current for a dc SQUID fabricated with the MITLL SFQ5ee 10 kA cm^{-2} process, with coupled inductors in layer M5.	121
3.48	InductEx model of a differential-arm inductance measurement SQUID for the MITLL SFQ5ee 10 kA cm^{-2} process with L_1 and L_{ctrl} in layer M5.	122

3.49	Difference between InductEx calculations and average measurements of inductance in M1 microstrip over M0 ground for the Hypres 4.5 kA cm ⁻² process with nominal process parameters.	123
3.50	Difference between InductEx calculations and average measurements of inductance for the Hypres 4.5 kA cm ⁻² process with nominal process parameters and fixed segmentation size.	124
3.51	Difference between InductEx calculations and average measurements of inductance for the Hypres 4.5 kA cm ⁻² process with calibrated process parameters.	124
3.52	Rendered image of the 3D inductance model created by InductEx for MITLL SFQ4ee calibration structures.	127
3.53	Rendered image of the 3D inductance model created by InductEx for an MITLL SFQ4ee JTL layout with shadow casting.	130
4.1	Meshed 2D model from FLOOSS.	140
4.2	Boundary identification and extrusion of a FLOOXs-generated mesh with Silverlinings.	140
4.3	Current density across center strip edge of process modelled PTL.	141
4.4	Creation of a prismatic spline sweep to create a three-dimensional representation of a process-modelled object.	141
4.5	Side view of grounded, shunted Josephson junction from TCAD extraction.	142
4.6	Three-dimensional rendering of grounded, shunted Josephson junction from TCAD extraction.	142
4.7	A process-extracted edge-swept model of a full 100 μm × 70 μm OR2T cell with the M7 skyplane removed for visualisation purposes.	143
4.8	An InductEx model of a JTL with the M7 skyplane removed for visualisation purposes and current density due to the bias current.	143
4.9	Comparative results for a genetic and random optimization sequence starting with the same unoptimized COSL set-reset flip-flop.	145
4.10	Typical margin analysis plots for an unoptimised and optimised circuit.	146
4.11	Extracted delay of a JTL with nominal $I_C = 250 \mu\text{A}$ in the MIT Lincoln Laboratory SFQ5ee process as a function of applied bias voltage and characteristic voltage V_C	147
4.12	Mealy state diagram of an RSFQ XOR gate with three states and a non-functional RSFQ XOR gate with five states.	147
5.1	Mask layout of AND gate in CONNECT cell library with port definitions for InductEx modelling.	151
5.2	Modulus of current density in the main ground plane (M7) of an AND gate in the CONNECT cell library for the AIST ADP2 process.	152
5.3	Four layout patterns of two striplines with length L_m separated by a distance S_m between the centre lines.	152
5.4	A stripline in M6 connected to a stripline in M1, through holes in ground planes M2, M3, M4 and M5 of the MIT-LL SFQ5ee process	153
5.5	Simulated results for layout P1 to P4 as a function of stripline spacing.	153
5.6	Simulated operating field margins of an RSFQ circuit with DC-to-SFQ converter, two JTLs and an SFQ-to-DC converter as a function of normalized bias current.	155

5.7	Current distribution as calculated with InductEx for a bias line in M7 over a solid ground plane in M4 near an unshielded victim SQUID and a caged bias line in M5 near a victim SQUID shielded with an M7 skyplane.	156
5.8	Schematic circuit diagram of inductively coupled pulse transfer cell.	157
5.9	Microphotograph of the inductively coupled pulse transmission cell.	158
5.10	InductEx extraction model of complete inductively coupled pulse transmission cell.	158
5.11	A packaged M2700 YBCO SQUID from Star Cryoelectronics.	160
5.12	The M2700 SQUID from Star Cryoelectronics viewed close up.	161
5.13	Netlist of M2700 with feedback coil and external magnetic field current sources and equivalent inductances included for parameter extraction.	162
5.14	Close-up view of the InductEx segmented mesh for the M2700 extraction model, showing only the active and unused SQUID loops, the bias pins and the connection to the pickup loop.	162
5.15	JoSIM simulation netlist of the full M2700 SQUID model in a z -directed field.	163
5.16	JoSIM simulation output of for the full M2700 SQUID simulation model with flux modulation.	163
5.17	Schematic of a 2-junction SQUID with one equivalent inductance L_T and coupling from a z -directed field.	164
5.18	Mesh of M2700 SQUID and feedback coil generated with InductEx.	165
5.19	Circuit schematic of a flux trapping hole coupled to the inductors in a SQUID circuit.	168
5.20	Layout showing just the test SQUID and moat configurations for five flux linkage experiments.	169
5.21	Overlay of microscope photograph and InductEx simulation of the magnetic field created by a trapped fluxon for a test SQUID.	170
5.22	Critical current of SQUIDs in flux linkage experiments.	171

List of Tables

2.1	Effects of bias current to critical current ratio on JTL delay and margins.	29
2.2	Measured bias margins for fabricated ColdFlux RSFQ cells.	51
2.3	List of ColdFlux RSFQ library cells.	52
3.1	Measured and InductEx-extracted inductance results for test SQUIDs over 4 chips manufactured on one wafer in the IPHT RSFQ1D process.	84
3.2	Measured and InductEx-extracted mutual inductance results for test SQUIDs manufactured in the IPHT RSFQ1D process.	86
3.3	Measured and InductEx-extracted SQUID loop inductances over ground plane holes.	88
3.4	Calculation times for original FastHenry and Fast FastHenry (FFH) with different preconditioner options and processor core counts.	96
3.5	Performance comparison between TTH and FFH.	105
3.6	InductEx process parameters for Hypres mask aligner and wafer stepper processes.	125
3.7	InductEx modelling parameters for five MITLL SFQ4ee layer definition file sets.	126
3.8	Calibrated values of process (layer) parameters for five MITLL SFQ4ee layer definition file sets.	128
3.9	Results for five MITLL SFQ4ee layer definition file sets.	128
3.10	RMSE results for very narrow coupled structures in MITLL SFQ5ee process.	129
4.1	MNA component stamps for voltage method and trapezoidal integration	136
4.2	MNA component stamps for phase method and trapezoidal integration	138
4.3	MNPA component stamps for second order BDF integration	138
4.4	Comparison of electrical simulator execution speed	139
5.1	Measured and InductEx-extracted inductances of inductively coupled pulse transmission cell.	157

Chapter 1

Introduction

1.1 Background

Digital and analogue superconductor integrated circuits have found significant application in diverse areas of science and industry over the past four decades. From single-photon detectors for astronomy to digital radio systems for military applications, from high-Q filters for mobile communications to ultra-low power digital systems, and from very sensitive magnetometers to quantum electronics systems.

With the maturation of superconductor integrated circuits from experimental devices to complex commercial systems came the increased demand for design tools and high fidelity parameter extraction and verification capabilities.

The aim of this dissertation is to illustrate the contributions that my research group and I have made to the development of parameter extraction and layout verification tools for superconductor integrated circuits in particular, but with application to all integrated circuits. I also show how my group and I applied these tools to device and circuit design. I believe that this work has made a significant contribution to the international effort to advance superconductor integrated circuits and systems from laboratory experiments to industrial applications. My journey through more than two decades of research and development is described in the text.

1.2 A short profile

I started my tertiary education in 1995 when I enrolled for my BEng degree at the Department of Electronic Engineering at Stellenbosch University, fully intent on wrapping up a four-year degree before heading into the defence industry. Vacation work quickly soured me on industry, and a talented array of dynamic professors in electronics and electromagnetics inspired me to turn to research and a postgraduate career. Through chance I was assigned a final year undergraduate project on Rapid Single Flux Quantum superconductor circuits under Professor Willem Perold, who had just recently returned from a sabbatical at the University of California at Berkeley, where he worked under the world-renowned Professor Theodore van Duzer on superconductor digital circuits.

I stayed on for an MEng degree under Professor Perold, where I soon ran into practical difficulties with integrated circuit layout for a superconductor analogue-to-digital converter that required non-existing tools to solve. Professor Perold arranged for me to attend my first international conference, the IEEE Applied Superconductivity Conference, in Virginia Beach in 2000. There he introduced me to a number of high-profile international

researchers who, rather than to criticise my naïve efforts, chose to express confidence that I would make a success of my research project. Those words of encouragement nudged me towards research for the long haul. I subsequently stayed at Stellenbosch for a PhD under Professor Perold while joining the Department of Electrical and Electronic Engineering as a Junior Lecturer to start an academic career that has thus far spanned two decades.

After befriending Dr Thomas Ortlepp at the 2002 IEEE Applied Superconductivity Conference in Houston, Texas, where we both stayed in the cheapest accommodation in a University of Houston residence (to this very day the worst accommodation that I have had the displeasure of paying for), I accepted his invitation to attend the third RSFQ Design Workshop at the Ilmenau University of Technology, in Ilmenau, Germany, in 2005. I attended with my colleague, Retief Gerber, who also did his postgraduate studies under supervision of my mentor, Prof. Willem Perold. The workshop opened our eyes to work of other research groups in a way that no previous conference could, and we returned home with the idea to develop software modules that could improve the design capabilities and turn-around time of the RSFQ designers that we met at the workshop. On a side note: the workshop also introduced us to the concept of a European hotel breakfast, and the all-day sustenance that it could provide!

Retief Gerber and I, with Willem Perold, subsequently developed a plan to obtain funding for software development. After several failed pitches and rewrites, the South African National Research Foundation's Innovation Fund eventually provided seed money for a 30-month development phase that had to culminate in a product ready for commercialisation. We hired engineers and software developers and started NioCAD as a project through Stellenbosch University in January 2007.

I returned to Ilmenau for the fourth RSFQ Design Workshop in September 2007 and stayed on for a month on a research visit, where Thomas Ortlepp assigned me to RSFQ circuit design with the tools used at Ilmenau. It was my introduction to Paul Bunyk's inductance extraction tool, Lmeter [1], which was far more useful than the first version of InductEx. Lmeter had shortcomings, but it was faster than the FastHenry engine used by InductEx, and it could handle an RSFQ circuit netlist with multiple connected inductors and multiple terminals.

At the same time, acquaintances made at the RSFQ Design Workshop led to an invitation to partner with 14 European research institutions to bid for a European Commission Framework Programme 7 project. The proposal, titled "Shrink-Path of Ultra-Low Power Superconducting Electronics" or S-Pulse, was selected for funding. It was headed by Dr Hans-Georg Müller at the Institute für Photonische Technologien (IPHT) in Jena, Germany, and provided some travel money that allowed project partners to visit each other for a month or two over the next few years. Under this project, I was able to visit IPHT Jena several times to test superconductor integrated circuits that I had fabricated with both IPHT and Hypres. I carried the brittle, exposed dice from Hypres in small plastic containers in my carry-on luggage, and Olaf Wetzstein at IPHT patiently did the testing in liquid helium. The test results confirmed circuit operation and inductance result and contributed to my confidence in tool development.

After the NioCAD project had run its course, we had great-looking tools that lacked technical depth and customers. I returned my focus to research on inductance extraction because I needed better, faster methods for circuit analysis, and I believed that other researchers might need something more powerful than Lmeter too. I wrote InductEx so that it could be used with minimal setup, and then set out to test it in the wild. At the conclusion of S-Pulse, IPHT gave me access to experimentally measured results for

various SQUID layout with which InductEx could be validated. Oleg Mukhanov, then at Hypres, did the same with a vast set of measurements, and with Nobuyuki Yoshikawa's help I could also validate InductEx for the AIST processes.

After a month of intense development during a research visit to France, I gave InductEx to Romain Collot in Pascal Febvre's lab to help him validate his RSFQ layouts, to Naoki Takeuchi in Nobuyuki Yoshikawa's lab to help him squash the parasitic coupling in AQFP gate layouts, and soon after to Vasili Semenov who contacted me when his layouts ran over the support cliff of existing tools. Soon, results were reported at conferences with acknowledgements to InductEx, and more users started to line up. The tipping point was IARPA's C3 programme, when Hypres and IBM turned to InductEx to analyse layouts. I got a call from Massoud Pedram at USC in the early days of C3, who asked me to join a bid for an upcoming IARPA seedling programme on tool development for SCE integrated circuits. We got the project, delivered results, and managed to progress to the fully funded IARPA SuperTools project. The opportunities afforded me by SuperTools and the people I met – researchers, test and evaluation teams and programme management – allowed me to expand *and retain* my team, push my research into unknown territory, and expand my horizons.

Towards the end of SuperTools, I was awarded a B1 research rating after rigorous international peer review by the South African National Research Foundation – a fairly rare rating that signifies an accomplished researcher with considerable international recognition in their field.

1.3 Contributions

The vast majority of my contribution to the research field has revolved around the extraction of inductance from integrated circuit structures, hence the title of this dissertation. I started with inductance calculation simply because I wanted to verify the analytical approximations used for interconnect inductance – simple microstrip lines in single ground plane thin-film integrated circuit layouts. Once I could get a numerical solution with my own three-dimensional models, using the open source field solver FastHenry to find current distribution and inductance, I realised that I could make layouts that were significantly more complicated. As my group and I started to depend more on the numerical inductance calculations, I automated the methods as a programme that I called InductEx.

When I could extract inductance from multi-terminal systems, it seemed that all the research on inductance in superconductor circuits had been done, and I could return to circuit design. I could not imagine how much more was to come. Improvements in inductance extraction lead to wider applications, which demanded ever more improvements. This recently culminated in the answer to a question as old as superconductor electronics: “What does a trapped fluxon do to a circuit?”

My contributions are detailed in this dissertation, but to summarise it has always been about inductance:

- I contributed to the formalisation of RSFQ circuit design theory, specifically to calculate the range of inductance in circuit components.
- I contributed to numerical inductance extraction of complex, multi-layer three-dimensional integrated circuit layouts.

- I contributed to the understanding of the effect of ground plane currents, magnetic fields and trapped fluxons on superconductor integrated circuits – through their coupling to circuit inductors – and the improvement of layouts to counter detrimental effects.

1.4 Layout of the dissertation

This dissertation is structured by topic, so that it is not entirely chronological. It is presented in four main chapters.

Chapter 2 presents superconductor integrated circuits, which is how I was introduced to inductance extraction. Although much of my work is applicable to conventional integrated circuits, I have only ever designed superconductor integrated circuits. The work progresses from small logic cells to a full cell library.

Chapter 3 details my primary contribution to the research field: inductance extraction and magnetic field analysis of complicated superconductor integrated circuit layouts. It starts with my early attempts, born out of frustration, to just get a passable value for the inductance of an intra-gate connection – and builds up to a powerful commercial tool that can handle thousands of self- and mutual inductance calculations in chip-scale layouts.

Chapter 4 fills in the details around other tools that, together with parameter extraction tools (of which inductance is the main focus), complete a superconductor electronics design tool chain.

In Chapter 5, I show how the tools that my group and I developed are applied to real engineering problems, and highlight that my contribution is not purely academic, but has made meaningful impact in the applied superconductivity community.

Chapter 2

Superconductor integrated circuits

2.1 Background

The invention of the monolithic integrated circuit had a profound impact on the pace of technological progress, eventually leading to the dawn of the information age.

Although almost every integrated circuit manufactured in the world today is still a monolithic semiconductor machine, the capability to integrate thousands (and later millions and billions) of devices with inter-device wiring on a single crystal die of pure silicon have also enabled us to manufacture non-semiconductor integrated circuits.

The most mature non-semiconductor technology uses low temperature metallic superconductors. Such superconductor thin-film devices have been patterned as integrated circuits on silicon substrates for more than three decades. Cryogenic cooling requirements limited the application of early superconductor integrated circuits, but at their introduction in the mid to late eighties most superconductor digital circuits could outperform state-of-the-art semiconductor digital circuits by orders of magnitude in terms of switching speed (or clock speed). Superconductor circuit designers sold visions of a future filled with superconductor supercomputers to leverage generous government research grants, but the practical difficulties in turning a great circuit invention into a global industry hampered the growth of superconductor integrated circuit technology. From the early 1990's, the superconductor circuit community watched in dismay as the semiconductor industry powered ahead, charged by the advent of complementary metal-oxide semiconductor (CMOS) technology and relentless and mind-blowing progress in scaling of circuit feature size. With every semiconductor technology node the number of transistors on a semiconductor chip increased, their ingress into almost every device and aspect of human technology became more complete, the astronomical profits of what became a \$500 billion per year industry increased, and the dwindling support and funding for superconductor integrated circuits dragged heavier on morale and progress.

It has been argued several times that the lack of design tools for superconductor integrated circuits stymied the growth of the technology, especially when one considers just how instrumental the tools of companies such as Synopsys, Cadence and Mentor Graphics have been to the growth of the semiconductor industry. I have devoted the majority of my research career, and thus my life's work, to narrow some of the gaps between superconductor and semiconductor design tools. That is the focus of later chapters.

This chapter details my involvement in superconductor integrated circuit design, which I have mostly done either to verify and apply design tools, or at least with my group's EDA tools. Through circuit layout, the connection between the work done on circuit

design and my primary research focus has always been: integrated circuit inductance.

2.1.1 Flux quantization

A thorough treatment of superconductivity and superconductor devices can be found in literature. I touch on the basics here to define the parameters and equations used in the circuit design and inductance extraction discussions that follow later in the text.

One important aspect of a superconductor, derived from the macroscopic quantum model, is that flux through a hole in a superconductor is quantized [2], [3]:

$$\Phi = \frac{h}{2e}n, \quad (2.1)$$

where h is Planck's constant and e is the electron charge. From (2.1), the smallest non-zero flux enclosed by a superconducting loop is thus

$$\Phi_0 = \frac{h}{2e} = 2.0679 \times 10^{-15} \text{Wb}. \quad (2.2)$$

This quantum of flux is called a *fluxon*.

If a superconducting loop contains an odd number of π -phase shift Josephson junctions [4], [5], then the flux through the loop is

$$\Phi = (n + 1/2)\Phi_0. \quad (2.3)$$

The smallest captured flux in a π -phase shifted loop is thus $\pm\frac{1}{2}\Phi_0$, or a half flux quantum. However, the minimum *change* in captured flux is still one fluxon.

2.1.2 The Josephson junction

In 1962, Brian Josephson predicted the tunneling of a supercurrent I through a thin isolation barrier between two bulk superconductor regions [6], [7]. The Josephson effect, as it has come to be known, was demonstrated the next year.

Detailed background on superconductivity, Cooper electron pairs and the derivation of equations that govern the electrical behaviour of the Josephson junction can be found in literature. For this discussion, only the most relevant characteristics are presented.

A Josephson junction can be made in a number of ways. One is to use a weak link, where a superconducting line is narrowed to create a junction that limits a passing supercurrent to a small cross-section. The current density across the junction can then be made to exceed the critical current density (J_C) of the superconducting material.

Another technique is to use electron tunneling across a superconductor-insulator-superconductor (SIS) barrier to achieve the Josephson effect [3], provided that the barrier is sufficiently thin – in the order of tens of angstrom. A schematic diagram for an SIS junction is depicted in Figure 2.1. This type of Josephson junction is predominantly used in thin-film integrated circuits for almost all digital and most analogue applications due to high uniformity of critical current between different junctions on a chip.

The state of all the Cooper pairs in each superconducting “bank” of the junction terminals is described by the macroscopic wave function

$$\Psi(\mathbf{r}, t) = |\Psi(\mathbf{r}, t)| e^{j\theta(\mathbf{r}, t)}. \quad (2.4)$$




Figure 2.1: (a) A superconductor-insulator-superconductor Josephson junction and (b) the circuit symbol for the Josephson junction.

If the phase $\theta(\mathbf{r}, t)$ in each superconducting terminal of the junction is Θ_1 and Θ_2 respectively, the phase difference across the junction is

$$\phi = \Theta_1 - \Theta_2 - \frac{2\pi}{\Phi_0} \int_1^2 A_x d\mathbf{l}, \quad (2.5)$$

where A_x is the x -directed component of the magnetic vector potential of a magnetic field through the junction.

If there is no magnetic field through the junction, or it is negligible – as is assumed for digital circuit design and analysis – the phase difference reduces to $\phi = \Theta_1 - \Theta_2$. The supercurrent through the junction is then given by

$$I_S = I_C \sin \phi, \quad (2.6)$$

where I_C is the critical current of the Josephson junction. For a very detailed review of the current-phase relation (CPR) of any type of junction, see [8]. If the junction is small, which is again assumed for digital circuit design and analysis,

$$I_C = J_C A, \quad (2.7)$$

where J_C is the critical current density of the junction configuration – which is fixed for a wafer – and A is the contact area between the insulator and the superconductor terminal on either end.

The cross-section of a planar SIS junction manufactured with the MITLL SFQ4ee process is shown in Figure 2.2. The aluminium/aluminium-oxide tunnel barrier is less than 10 nm thick.

When the voltage over the junction exceeds the gap voltage

$$V_g = \frac{2\Delta}{e}, \quad (2.8)$$

where Δ is the energy required to add an unpaired electron or hole to the superconductor and e is the electron charge, Cooper pairs are broken and single or “normal” electrons called *quasiparticles* are formed. These quasiparticles behave differently to electrons in normal metals. The quasiparticle current increases linearly above V_g , so that the slope corresponds to a “normal resistance” R_n . The quasiparticle current I_n follows from Ohm’s law.

When the Josephson junction is at a temperature above 0 K but the voltage is below V_g , thermal motion of the charge carriers breaks some of the Cooper pairs to produce a non-zero density of quasiparticles that contribute to the normal current. The slope at which the normal current due to these quasiparticles increases with the voltage over the




Figure 2.2: Scanning electron microscope image of the cross-section of an integrated circuit SIS Josephson junction from the MITLL SFQ4ee process ([9], reproduced with permission). The tunnel barrier is a thin gray line between the JJ and M5 metals.

junction is fairly linear, and is thus approximated as an effective “subgap” resistance, R_{sg} . The linear approximations for R_n and R_{sg} are not valid near the gap voltage, and the discontinuity cannot be ignored in circuit simulators.

By the very nature of its construction a planar SIS junction has significant capacitance, so that a displacement current is present when voltage over the junction varies. A generalised model for the SIS Josephson junction thus contains supercurrent, normal and displacement current branches, as shown in Figure 2.3. This circuit model is called the resistively and capacitively shunted junction (RCSJ) model.




Figure 2.3: Circuit schematic for the generalised RCSJ model.

The total junction current is thus

$$I = I_C \sin \phi + \frac{V}{R} + C \frac{dV}{dt}. \quad (2.9)$$

It can be shown that the voltage-phase relation for a Josephson junction is

$$\frac{d\phi}{dt} = \frac{2\pi}{\Phi_0} V. \quad (2.10)$$

Equation (2.10) is also known as the second Josephson equation. Applied to (2.9), we have

$$I = I_C \sin \phi + \frac{1}{R} \frac{\Phi_0}{2\pi} \frac{d\phi}{dt} + C \frac{\Phi_0}{2\pi} \frac{d^2\phi}{dt^2}. \quad (2.11)$$

The current can be written as a dimensionless equation

$$\frac{I}{I_C} = \sin \phi + \frac{d\phi}{d\tau'} + \beta_C \frac{d^2\phi}{d(\tau')^2} \quad (2.12)$$

with

$$\tau_J = \frac{\Phi_0}{2\pi I_C R} \quad (2.13)$$

and

$$\tau_{RC} = RC \quad (2.14)$$

representing the time constants for the basic Josephson junction and the RC combination respectively. Also,

$$\tau' = \frac{t}{\tau_J}. \quad (2.15)$$

We also introduce the dimensionless Stewart-McCumber parameter as

$$\beta_C = \frac{\tau_{RC}}{\tau_J} \quad (2.16)$$

and thus

$$\beta_C = \frac{2\pi I_C R^2 C}{\Phi_0}. \quad (2.17)$$

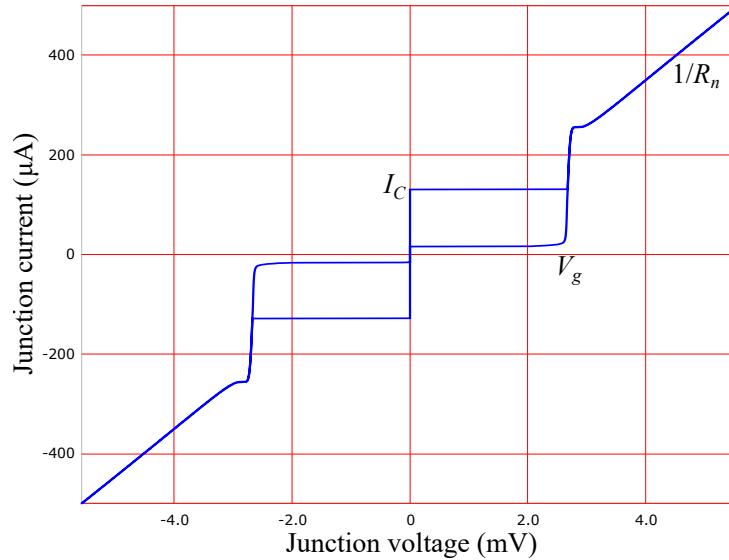


Figure 2.4: Experimentally measured I-V curve of a niobium-aluminium-oxide-niobium planar SIS junction.

The measured current-voltage (I-V) response of an unshunted Josephson junction is shown in Figure 2.4 when the junction current is swept up and down as a triangular signal. The junction was fabricated in the MITLL SFQ5ee process. It can be seen that the response is hysteretic. The critical current, gap voltage and normal resistance can be identified clearly.

In digital circuits, the hysteresis of the junction and the resistance discontinuity are ameliorated by the addition of a shunt resistance – which is external to the junction in most processes. If the shunt resistance R_s is much smaller than the normal resistance, we approximate the effective junction resistance as R_s . The equivalent circuit schematic for a shunted junction is shown in Figure 2.5. It is important to note that an external shunt resistor has an inductive component – typically in the low picohenry range – that adversely affects the dynamics of the junction.




Figure 2.5: Circuit schematic for the generalised RCSJ model with an external shunt resistor.

It is also important to note that the RCSJ model is only valid in the case of short junctions where $a_0 < \lambda_J$ so that the phase difference across the junction behaves like a point-like variable. Here, a_0 is the linear geometric dimension of the junction (the side length for a square junction, or the diameter for a round junction) and

$$\lambda_J = \sqrt{\frac{\Phi_0}{2\pi\mu_0 J_C (2\lambda_L + t_{ox})}} \quad (2.18)$$

is the Josephson penetration depth [10]. J_C is the current density from (2.7).

More information on the dynamics of Josephson junctions and the effects of the time constants are available in literature.

2.1.3 The superconducting quantum interference device

Two Josephson junctions can be combined in parallel to produce a device known as a superconducting quantum interference device (SQUID). A two-junction SQUID is also called a dc SQUID.

If the junctions have equal values for I_C , the total current entering the loop is

$$I_T = I_1 + I_2 = I_C \sin \phi_1 + I_C \sin \phi_2 \quad (2.19)$$

$$\Rightarrow I_T = 2I_C \cos\left(\frac{\phi_1 - \phi_2}{2}\right) \sin\left(\frac{\phi_1 + \phi_2}{2}\right) \quad (2.20)$$

Through contour integration of supercurrent densities and magnetic vector potentials around the contour C [2] it can be shown that

$$I_T = 2I_C \cos\left(\frac{\pi\Phi_a}{\Phi_0}\right) \sin\left(\phi_1 + \frac{\pi\Phi_a}{\Phi_0}\right), \quad (2.21)$$

where Φ is the magnetic flux in the SQUID loop enclosed by the contour C .




Figure 2.6: A schematic of a two-junction SQUID with a symmetrical feed. The integration path C for analysis is shown by the dotted line.




Figure 2.7: Circuit schematic of a dc SQUID with generalised Josephson junctions.

A practical SQUID – used as a magnetometer – produces a voltage related to the applied flux. Consider a dc SQUID with generalised Josephson junctions as shown in Figure 2.7, where junction capacitance is neglected for simplicity by assuming $\beta_C \ll 1$. A constant flux Φ_a threads the SQUID loop, of which the inductance is neglected for now.

If a voltage V exists over the SQUID, then 2.21 becomes:

$$I_T = 2I_C \cos\left(\frac{\pi\Phi_a}{\Phi_0}\right) \sin\left(\phi_1 + \frac{\pi\Phi_a}{\Phi_0}\right) + \left(\frac{1}{R_1} + \frac{1}{R_2}\right) V. \quad (2.22)$$

If we define a new phase

$$\phi = \phi_1 + \frac{\pi\Phi_a}{\Phi_0} \quad (2.23)$$

and keep the external flux Φ_a constant, then

$$\frac{d\phi}{dt} = \frac{d\phi_1}{dt} = \frac{2\pi}{\Phi_0}. \quad (2.24)$$

It can then be shown [2] that

$$I_T = I_{CT} \sin \phi + \frac{1}{R} \frac{\Phi_0}{2\pi} \frac{d\phi}{dt}, \quad (2.25)$$

where

$$I_{CT} = 2I_C \left| \cos\left(\frac{\pi\Phi_a}{\Phi_0}\right) \right|, \quad (2.26)$$

and R is the parallel combination of R_1 and R_2

$$R = \frac{R_1 R_2}{R_1 + R_2}. \quad (2.27)$$

The amplitude I_{CT} is periodic in the applied flux, with maximum at $\Phi_a = n\Phi_0$ and minimum at $\Phi_a = n\Phi_0/2$.

If a dc current I_T is greater than $2I_C$ at zero applied flux, then a time-dependent voltage V will develop over the SQUID. The dc value of V can be shown to be

$$\langle V \rangle = I_T R \sqrt{1 - \left[\frac{2I_C}{I_T} \cos \left(\frac{\pi \Phi_a}{\Phi_0} \right) \right]^2}. \quad (2.28)$$

It can now be seen that $\langle V \rangle$ is also periodic in applied flux, with a periodicity of exactly Φ_0 . This periodicity is exploited in the measurement of inductance in superconductor circuits, as is shown in Section 3.4.1.

2.2 Fabrication processes

The development of integrated circuit parameter extraction tools, which I detail later, requires a thorough understanding of the processes by which an integrated circuit is manufactured.

In my opinion, integrated circuits are the most complex objects ever created by humans. Modern semiconductor ICs can have a hundred billion components, all wired together on a single crystal silicon sliver of about $50 \mu\text{m}$ thick and a few square centimetres in area. A microchip can take hundreds of person-years to design with the best electronic design automation software available, requires fabrication in plants that can cost tens of billions of US dollars to build, and dissipates so much energy that it would evaporate in seconds under full load if not for some very good thermal engineering.

Superconductor ICs are decades behind semiconductor ICs in terms of complexity and scale, but the basic fabrication steps are very similar.

2.2.1 Dimensions

In the integrated circuit community, length is mostly given in centimetres. Despite using the International System of Units (SI) everywhere else, I use centimetres when referring to fabrication parameters to align with the most commonly provided numbers for IC fabrication processes.

2.2.2 Steps in IC fabrication

Circuit designers require a good understanding of the fabrication process steps of a target process to be able to design integrated circuits for that process.

Electronic design automation (EDA) tool development, especially as far as circuit and device parameter extraction from IC layouts is concerned, requires a more thorough understanding of the steps in the IC fabrication process to allow process-specific and generic process support.

Short of actual process engineering, the development of technology computer-aided design (TCAD) tools requires the most thorough understanding of fabrication steps. My Masters student Heinrich Herbst dove into this topic for his work in SC fabrication process modelling [11] under the ColdFlux project.

2.2.2.1 Wafer preparation

Most modern integrated circuits are made on silicon substrates. These substrates, in the form of silicon wafers, are cut from a single crystal ultra-pure silicon boule that is grown from a seed crystal with the Czochralski method. The wafer surfaces are then smoothed and polished in an electrochemical process.

2.2.2.2 Deposition

The metals in a superconductor IC, niobium as superconductor, aluminium for the Josephson junction barrier and molybdenum for resistors, are deposited through DC magnetron sputtering [9], [12].

The inter-layer dielectric SiO_2 (silicon dioxide) is deposited with plasma-enhanced chemical vapour deposition (PECVD).

2.2.2.3 Oxidation

Oxidation is defined as a reaction where electrons escape from an atom, molecule or ion. In a fabrication process, thermal oxidation is effected when a conducting layer is heated and exposed to pure oxygen to form an isolating barrier.

2.2.2.4 Planarisation

Planarisation is the process whereby features across the surface of a wafer are removed through polishing to smooth, flat surface. Planarisation ensures that the next layer to be deposited has an even foundation on which to rest and is not influenced by the topography of lower layers. It is essential to maintain low parameter spreads. In modern IC fabrication, planarisation is done through chemical mechanical polishing (CMP).

2.2.2.5 Photolithography

Photolithography is used to form masks over a wafer where etching is either allowed or blocked. Photoresist, a light-sensitive polymer, is spin-coated over a wafer and baked. The photoresist is then exposed to light, usually ultraviolet light, that is shone through glass-and-metal masks. Depending on the polarity of the photoresist, exposed areas either becomes soluble or not. After exposure, the soluble photoresist is washed off to expose underlying areas for etching. After etch completion, remaining photoresist is stripped off with a liquid resist stripper.

2.2.2.6 Etching

Etching is the process where exposed areas on a wafer, not protected by a hardened photoresist, are etched away by either a liquid etching agent (acid, or wet etch) or for smaller feature sizes by a dry etch technique. For dry etching, high-density plasma etching (HDPE) is mostly used. HDPE provides a combination of chemical etching from the reaction of the etching target film with the plasma and physical ion etching from directional bombardment of the wafer with ions. Ion bombardment also etches the photoresist, but the process is designed for excellent selectivity, so that etch rates in the target material are higher than in photoresist.

2.2.2.7 Anodization

Anodisation is an electrolytic process where a target metal is used as the anode during an oxidising process. The metal layer is thus coated with an oxidised barrier, which can either be used to create the isolating barrier of a Josephson junction or aid with isolation of metal layers

2.2.3 Monolayer fabrication processes

Superconductor integrated circuit fabrication processes can be as simple as a single layer – usually YBCO – superconductor deposited on a substrate before etching and junction patterning [13]. It should be noted that reliable patterning of Josephson junctions with tight control over I_C spreads at any desired location on the substrate is anything but simple, as is evident from the complexities of targeted weakening of superconducting regions through ion irradiation [14]. Digital superconductor circuits have been demonstrated in such monolayer high temperature superconductor (HTS) IC processes [15], although monolayer HTS processes are more popular for analogue devices such as SQUID magnetometers and gradiometers with step-edge [16]–[18] or bicrystal grain boundaries [19].




Figure 2.8: (a) Microphotograph of SQUIDs in a monolayer process before ion irradiation to create the Josephson junctions and (b) a close-up of one of the SQUIDs [20]

Modelling a monolayer process for inductance extraction is straightforward, as was shown in [20] with InductEx for the monolayer SQUIDs depicted in 2.8. An illustration of an inductance extraction model, with excitation ports defined, is shown in Figure 2.9.

2.2.4 The main fabrication processes

The most reliable Josephson junctions for narrow tolerance digital circuits are planar SIS junctions made with niobium as the superconductor electrodes and aluminium oxide as the barrier material (see the cross-sectional electron microscope photograph in Figure 2.2). Low temperature superconductor (LTS) processes that use such planar SIS junctions require at least two superconductor layers – one for each electrode of the junction. With both of the electrode layers also used for wiring, signal crossover is easily supported (unlike with monolayer processes). A third superconductor layer can then be used as a ground plane.

Each integrated circuit fabrication process has its own design rules, layer stack, fabrication sequence and material parameters. These may differ substantially between processes, so that an integrated circuit layout cannot simply be transferred between two processes




Figure 2.9: InductEx model of a monolayer SQUID.

without significant alterations. In general, layouts are not adapted between processes, but rather started anew for each process.

Detailed knowledge of the design processes is not only required for circuit layout, but also for the development of parameter extraction tools. Information about a design process should never be hard-coded into any EDA tool. For this reason, I developed InductEx to handle a generic fabrication process that could handle any or all of the fabrication steps and layer parameters of each of the processes that I had knowledge of, and designed a powerful and flexible input format with which to describe any of the known processes to InductEx. The most well-known and widely used of these processes are described below.

2.2.4.1 Leibniz IPHT - FLUXONICS 1 kA cm^{-2}

The 1 kA cm^{-2} niobium process [21] from what is today Leibniz IPHT, located in Jena, Germany, has only three superconductor layers – one for ground and two dual-purpose layers for junction electrodes and for wiring – and is the easiest to model in InductEx. The process, referred to as the FLUXONICS process, uses multiple anodisation and isolation mask layers which simplifies modelling support. The process does not use planarisation, so that layers are deformed in elevation due to the existence or absence of metal and isolation objects below. The FLUXONICS process layer stack is illustrated in Figure 2.10.

Some concepts and terminology can be defined and explained from the layer stack. In superconductor integrated circuit fabrication processes, it is standard practice to label all superconductive layers with a name starting with “M” for “metal”. The lowest layer, which is typically deposited directly on the wafer substrate, is usually named M0. Subsequent superconductor metal layers are incrementally named M1, M2, etc.

Isolation layers are typically labelled “I”, with a number corresponding to the nearest superconductor metal layer below. Sequential letters are used again after the isolation level number to denote different isolation layers between the same set of superconductor metal layers. In the FLUXONICS process, masks are used to pattern via holes on M0 and M1 that will be masked from anodisation. These layers are then I0A and I1A. The silicon dioxide isolation layers directly above each anodisation layer are then called I0B and I1B respectively.




Figure 2.10: The FLUXONICS process layer stack without the final isolation and pad layers.

Resistive metal layers are labelled “R” with a sequence number that corresponds to that of the nearest superconductor metal layer below.

It is clear that a via connection from the upper metal layer M2 to the lower metal layer M0 requires contact through six isolation layers and M1, so that a parameter extraction tool must be able to build a mesh that connects metals through all these layers, *and crucially, block a connection if an overlapping via pattern on any of the isolation layers is absent.*




Figure 2.11: Cross-section of a shunted Josephson junction fabricated with the FLUXONICS process. Electron microscope image provided by Leibniz IPHT.

An electron microscope image of the cross-section of a shunted Josephson junction with a nearby ground connection from M1 to M0 is shown in Figure 2.11. It is clear that etch stops at each isolation layer prevent etching through I2 from etching I1B and I1A below, so that a via from M2 to R1 may overlap the edges of the resistor. This keeps modelling for layout extraction simple, in that a connection from one metal layer to another can be found simply by calculating the intersection of all polygons on all isolation layers between the metal layers. Vias do not usually connect through a superconductor layer if the metal is absent. The resistive layer R1 is special, because a via stack through I1A, I1B and I1C must connect M1 to M2 *in the absence of an object on R1*. For generic parameter extraction modelling, which can be adapted to fit any process, I thus developed a layer bypass procedure that can be linked to any layer – such as R1 – so that all polygons on a layer mask are subtracted from the via stack intersections and the remaining intersections are the used to connect the surrounding metal layers in the mesh model.

An important part of extraction tool development is to support fabrication artefacts that have significant influence on extracted results, and for the sake of simplicity to ignore artefacts that have almost no influence.

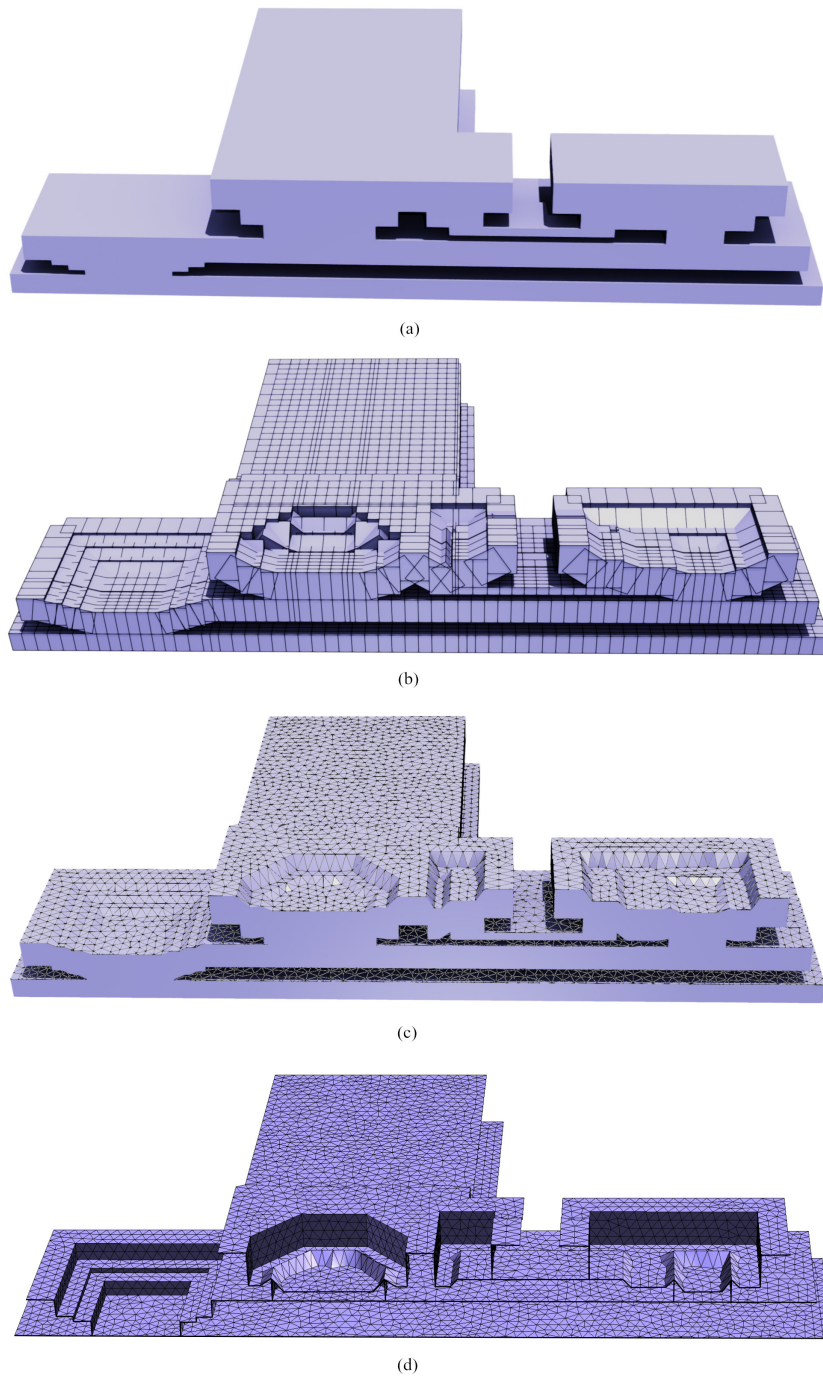


Figure 2.12: Cross-section of a model of a shunted Josephson junction similar to that in Figure 2.11 with (a) no elevation change and (b) a cuboid mesh, (c) a tetrahedral mesh and (d) a triangular mesh with elevation change.

In the FLUXONICS process, an anodisation layer consumes some of the thickness of the superconductor layer that becomes anodised. Typically, anodisation is only absent directly beneath a via contact. It is thus sufficient to ignore any layer thickness reduction due to anodisation during modelling, *and rather to model this by reducing the thickness of the metal layer in the process definition parameters passed to the extraction tool*. For InductEx, this means defining a slightly lower metal layer thickness in the Layer Definition File for the process.

The much more important artefact is elevation change in upper layers due to local etching or deposition on lower layers. This is shown with exaggeration in Figure 2.10, but clearly visible in the cross-section in Figure 2.11. A simplified model that *does not include* elevation change will significantly overestimate the inductance of a line in M2 over a ground plane in M0 where there is no M1 in between. This happens to be the most widely used inductance structure in the FLUXONICS process, and it is therefore essential to model elevation change. A model without elevation change is shown in Figure 2.12(a). It overestimates inductance in typical M2 inductors in real circuit layouts by as much as 10 % compared to the more representative models with elevation change shown in Figure 2.12(b) for cuboid segments, Figure 2.12(c) for tetrahedral segments and Figure 2.12(d) for triangular segments. For this reason, I developed InductEx to model elevation change with good precision.

2.2.4.2 Seeqc - #QC1000A 1 kA cm^{-2}

The #QC1000A 1 kA cm^{-2} process (hereafter referred to as the QC1000A process) from Seeqc, Inc. in Elmsford, New York, is a continuation of the fabrication process developed by Hypres, Inc. over most of the last thirty years. Hypres offered several variations on their fabrication process, of which the 4.5 kA cm^{-2} process [22] (which was a refinement of the earlier 1 kA cm^{-2} process) has arguably been the workhorse for superconductor IC fabrication in the Western hemisphere for more than two decades. I developed the very first version of InductEx specifically for the Hypres process in the early 2000's. The layer stack included four niobium layers, with M0 used mostly as the ground plane, Josephson junctions and shunt resistors fabricated between the wiring layers M1 and M2, and a third wiring layer M3 used mostly for skyplanes (shielding). The process was eventually extended to include a custom number of planarised layers below M0 [23].

When the fabrication process was transferred to Seeqc, Inc. a few years ago, where the focus is on qubits and energy efficient SFQ interface circuits, the critical current density was lowered to 1 kA cm^{-2} again.

The layer stack has essentially remained the same for many years, so that I only discuss the latest version of the Seeqc process here. The layer stack is illustrated in Figure 2.13

The differences between this process and the FLUXONICS process show that an extraction tool should be flexible enough to allow accurate modelling of integrated circuit structures even when construction of a layer stack varies significantly between processes. The key features that have to be supported in modelling are:

- Selective layer planarisation to support the planarisation in this process of M0 which sits above at least one conductive layer (MN1) in the layer stack.
- Via connection algorithms that can provide selective etch-stop when one mask layer is used for multiple isolation layers. In this process, I1 is fabricated in two steps to sandwich the resistive layer R2, and a single I1 via must connect conductive layers




Figure 2.13: The Seeqc QC1000A process layer stack.

M2 and M1 except where an object on R2 is in the way. In that case, the connection must be from M2 to R2 only.

- Via construction over the edge of a conductor must be possible without breaking the mesh. It is commonplace for designers using this process to save layout area and reduce inductance of shunt resistors by drawing I1 vias over the edge of an R2 resistor to short M2 to R2 and M1 with one via etch – this is depicted in the centre of the layer stack illustration in Figure 2.13.
- Flip-chip or multi-chip module scenarios need chip-to-chip or chip-to-carrier interconnects. The Seeqc QC1000A process supports bump bonds that are deposited directly on the top resistive pad layer R3.

In order to support these layout features in InductEx, I added layer arithmetic operators with which auxiliary layers can be created and inserted into the model layer stack. The resistive layer R2 is modelled correctly if an auxiliary lower layer for I1 is created from the difference between objects on I1 and R2, and slotted below layer R2 in the layer creation order. This is shown in Section 3.3.5.2 as layer I1BL in Figure 3.20.

The cross-section of an InductEx model for a JTL that uses MN1 for bias and internal inductors is shown in Figure 2.14.




Figure 2.14: Cross-section of a model of a JTL in the Seeqc QC1000 process.

2.2.4.3 AIST - standard 2.5 kA cm^{-2} and advanced 10 kA cm^{-2}

The National Institute of Advanced Industrial Science and Technology (AIST) advanced process (ADP2), located in Tsukuba, Japan, has nine niobium layers [24]. It evolved

from a fabrication process developed at what was known earlier as the International Superconductivity Technology Center (ISTEC) [25] to first have ten metal layers [26] and eventually only nine for higher yield [27]. It uses complemented caldera planarisation (described by [28] and [29]) in all but the upper two layers and introduced the concept of a buried power plane and multiple ground plane-encased passive transmission line layers below the ground plane. The ADP2 layer stack is illustrated in Figure 2.15.

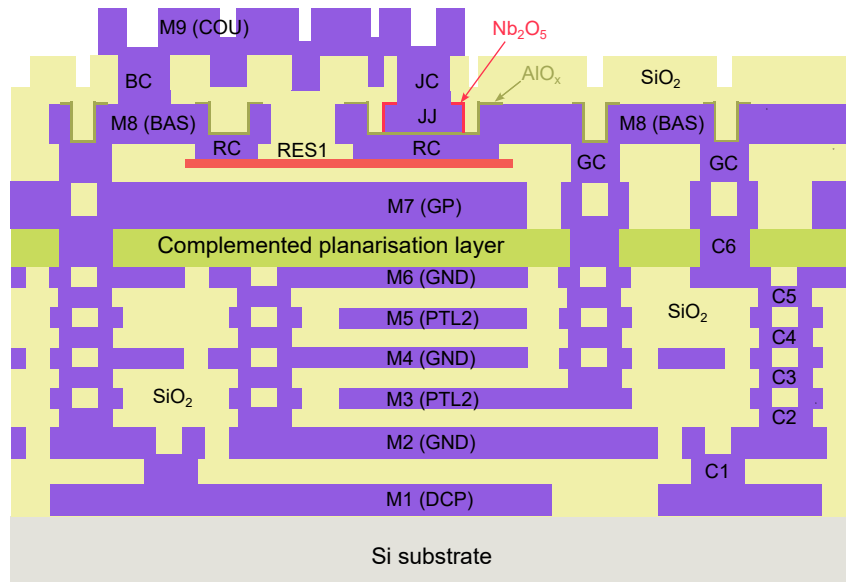


Figure 2.15: The AIST ADP2 process layer stack.

The AIST ADP2 process was the first process with partial planarisation for which I had to develop planarisation support during parameter extraction model construction. Other differences to the FLUXONICS and QC1000A processes that require model support are:

- A positive-mask ground plane, which differs from the negative-mask ground planes used in the FLUXONICS and Seeqc processes and for which default ground modelling involves the creation of a bounded ground “cast” from the inverse of the ground plane mask and subsequent cropping of the ground cast to reduce segment counts.
- A resistor layer between the main ground plane and the base electrode layer of the Josephson junctions – here between layers M7 (GP) and M8 (BAS).
- The use of multiple masks to etch vias through the same isolation layer. The ADP2 process uses GC to punch a via from M8 to M7 through the same isolation in which RC punches a via from M8 to RES1.

2.2.4.4 MIT Lincoln Laboratory - SFQ5ee 10 kA cm^{-2}

A fully planarised fabrication process for superconductor electronics has been developed at MIT Lincoln Laboratory in Lexington, Massachusetts. Early development of the process was driven by the IARPA Cryogenic Computing Complexity (C3) programme [30]. The process evolved through a number of process nodes, such as SFQ3ee, SFQ4ee, SFQ5ee, SFQ6ee and more [31], where the “ee” indicates that the process supports energy

efficient superconductor circuit design. Successive nodes feature more superconducting niobium layers and decreasing minimum linewidth. The SFQ5ee node of this process was designated as the process for which all electronic design automation tools and cell libraries had to be designed under the IARPA SuperTools programme.

A detailed description of the SFQ5ee process has been published [9], with the layer stack illustrated in [32]. A redacted illustration of the layer stack is shown in Figure 2.16.




Figure 2.16: The MIT Lincoln Laboratory SFQ5ee layer stack.

2.3 Early superconductor digital circuits

A project by IBM in the 1970's and early 1980's [33] was aimed at the development of superconductor digital computers. The project was abandoned in 1983 due to thermal cycling issues with the soft lead superconductors that degraded the parameters of Josephson junctions. It also used a latching logic design that could switch on in picoseconds, but required high power radio frequency (RF) signals at the desired clock frequency to reset latched cells.

The problem of degradation due to thermal cycling was solved through progressively better planar SIS Josephson junctions, with the advent of the Nb-Al oxide-Nb and Nb-Al oxide-Al-Nb junction [34], [35] leading to low parameter spread Josephson junctions that are stable with thermal cycling. These junctions made digital circuits based on Josephson junctions possible, and integrated circuits were soon demonstrated successfully [36].

The early niobium-based Josephson circuits were still based on voltage state operation. Japanese groups used four-junction logic (4JL) [37], with a 4-bit counter shown in simulation to have a projected delay time of 315 ps [38]. This was significantly faster than semiconductor integrated circuits at the time, but semiconductors would catch up within about a decade.

At Berkeley, flash analogue-to-digital converters based on one-junction and two-junction SQUIDs [39] were modified to act as voltage-state logic gates [40]. This evolved into Complementary Output Switching Logic (COSL) [41] (see Figure 2.17) that was demonstrated first between 5 and 10 GHz [42] (Figure 2.18) and would eventually reach 18 GHz.

However, Josephson junction-based voltage level switching logic just could not go significantly faster, and the concept, including COSL, was abandoned when single-flux-quantum logic was already an order of magnitude faster.

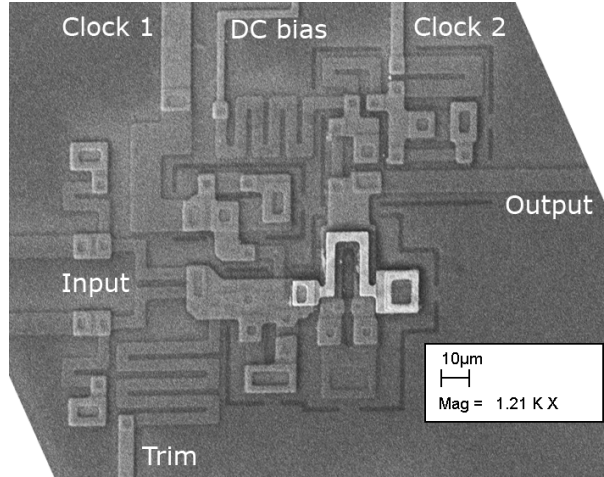


Figure 2.17: Electron microscope photograph of a manufactured COSL gate.




Figure 2.18: Experimental measurements of (left) a COSL OR gate and (right) a COSL NAND gate at 8 GHz. The output is the lower trace in each case. Oscilloscope screenshots with permission from Prof. Willem Perold, Stellenbosch University, South Africa.

2.4 The Rapid Single-Flux Quantum logic family

2.4.1 Origins of RSFQ

In 1985, Russian scientists Konstantin Likharev, Oleg Mukhanov and Vasili Semenov proposed a new way to do Josephson junction computing [43], where digital bits are not presented by dc voltage levels as in semiconductor digital circuits or earlier superconductor latching logic, but rather by the presence or absence of very short voltage pulses with quantised area:

$$\int V(t)dt = \Phi_0 \approx 2.07 \text{ mV}\cdot\text{ps}. \quad (2.29)$$

These pulses, integrated over time to the magnetic flux quantum, are called single-flux-quantum (SFQ) pulses, and all logic circuits that function with SFQ pulses are referred to as SFQ logic.

The first SFQ logic family relied on resistive interconnects between junctions, and was called Resistive SFQ (RSFQ). It suffered from narrow parameter margins, but was soon improved by the same team [44]. Resistive interconnects were replaced with more Josephson junctions, which broadened operating margins and increased switching speed. The logic family was then named Rapid Single-Flux-Quantum (RSFQ) and it dominated the superconductor digital logic arena for three decades – only recently making way for more energy efficient derivatives such as Energy-Efficient RSFQ [45] and ac-biased SFQ logic such as Adiabatic Quantum Flux Parametron (AQFP) [46].

RSFQ has been described in excellent detail by Likharev and Semenov [47], with extended logic cells detailed by Mukhanov [48]. Many an effort has been expended on documenting the design of RSFQ circuits, but the most comprehensive and legible is arguably the design guide compiled in the form of a PhD thesis by Dimov [49].

2.4.2 Data storage and transmission in RSFQ circuits

Most RSFQ circuits can be constructed from a combination of three basic subcircuit blocks: transmission blocks, storage elements [49] and decision-making pairs. The basic blocks are illustrated in Figure 2.19.




Figure 2.19: RSFQ basic circuit blocks.

A transfer block reproduces an incoming SFQ pulse and passes it to the next element. A storage element can store circulating current as a magnetic flux quantum in a loop. A decision making pair (DMP), when excited with an input SFQ pulse, switches one of two junctions depending on the current flowing through each junction.

As an example, the simplified circuit schematic of an RSFQ OR gate from [50] is shown in Figure 2.20. The simulated voltage response of the circuit to SFQ input voltages is shown in Figure 2.21(a), with a clock period marked. The absence of an SFQ pulse at the output during a clock period corresponds to a digital “0”, while an SFQ pulse at the output during a clock period corresponds to a digital “1”.

For circuit analysis tools and circuit design, it is more convenient to look at the circuit in terms of phase. The simulated phase over the three input junctions and the output junction is shown in Figure 2.21(b). Every junction switch results in a 2π phase change.




Figure 2.20: Basic RSFQ OR gate schematic.

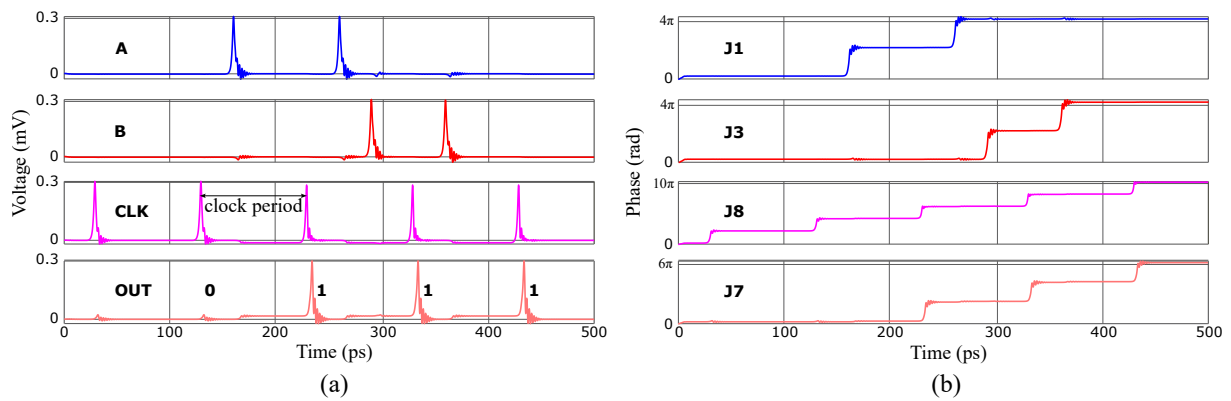


Figure 2.21: Simulated transient response of RSFQ OR gate. (a) Input and output voltages and (b) phase over Josephson junctions.

2.4.3 RSFQ circuit theory

Although many journal papers and doctoral dissertations examine RSFQ circuit design with fairly vague descriptions of phase difference, current flow and junction switching, a formalised circuit theory for RSFQ design – down to the level where it can be taught at the graduate level – did not exist in literature when Dr Lieze Schindler joined my research group. I tasked her with formalising the design of RSFQ circuits from phase equations, which she did during her PhD studies [51].

The design of an RSFQ circuit requires the selection of SQUID loops that will allow the desired number of states, the selection of where to apply inputs to achieve desired switching and state transition, and then the selection of appropriate circuit component values to create a functional circuit: junction critical currents, shunt resistance values, bias current values and *interconnect inductance*. Inductance is the last and most difficult to determine. It is also the most difficult to achieve in layout and to verify in layout extraction.

2.4.3.1 Phase-based equations

From the Josephson voltage-phase relation (2.10) and the relation of voltage over an inductor to current change through the inductor,

$$v = L \frac{di}{dt} \quad (2.30)$$

it follows that

$$L \frac{di}{dt} = \frac{\Phi_0}{2\pi} \frac{d\phi}{dt} \quad (2.31)$$

and thus that the current-phase relation of an inductor as shown in Figure 2.22 is

$$i_L = \left(\frac{\Phi_0}{2\pi} \right) \frac{\phi_1 - \phi_2}{L}. \quad (2.32)$$




Figure 2.22: Current-phase relation of an inductor.

Written differently, the phase developed over an inductor is related to the junction current as:

$$\phi_L = \frac{2\pi L}{\Phi_0} i. \quad (2.33)$$

Current is thus linearly related to phase over an inductor, as it is to voltage for a resistor.

We can show through similar derivation that for two inductors coupled with mutual inductance M , as shown in Figure 2.23,

$$\phi_1 = \frac{2\pi}{\Phi_0} (L_1 i_1 + M i_2), \quad (2.34)$$




Figure 2.23: Current-phase relation of two inductors coupled with mutual inductance.

and

$$\phi_2 = \frac{2\pi}{\Phi_0}(L_2 i_2 + M i_1). \quad (2.35)$$

The current-phase relation of a Josephson junction is

$$i_J = I_C \sin \phi. \quad (2.36)$$

The total equivalent inductance L_{Jt} of a Josephson junction as a function of the current through the junction can be calculated by setting the voltage over the junction equal to the derivative of the total flux in the equivalent inductance:

$$v_J = \frac{d[L_{Jt} i_J]}{dt}. \quad (2.37)$$

From (2.36) it follows that

$$\phi_J = \sin^{-1} \left(\frac{i_J}{I_C} \right). \quad (2.38)$$

Substitution of (2.38) and (2.10) into (2.37) yields

$$L_{Jt} i_J = \frac{\Phi_0}{2\pi} \sin^{-1} \left(\frac{i_J}{I_C} \right), \quad (2.39)$$

so that

$$L_{Jt} = \frac{\Phi_0 \sin^{-1} (i_J / I_C)}{i_J}. \quad (2.40)$$

Equations (2.32), (2.33), (2.36), (2.38) and (2.40) can be used for dc (operating point) circuit analysis of SFQ circuits in any state, when the dc voltage over superconducting branches is zero and no current flows through the resistive and capacitive branches of the RCSJ model.

The magnetic flux enclosed by a loop is

$$\Phi = LI. \quad (2.41)$$

For an arbitrary superconducting loop through branches with l inductors and j Josephson junctions, the loop flux is

$$\Phi = \sum_{m=1}^l L_m I_{L_m} + \sum_{n=1}^j L_{Jt_n} I_{J_n}. \quad (2.42)$$

Substitution of (2.33) and (2.39) into (2.42) yields

$$\Phi = \frac{\Phi_0}{2\pi} \left(\sum_{m=1}^l \phi_{L_m} + \sum_{n=1}^j \phi_{J_n} \right). \quad (2.43)$$

2.4.3.2 Basic Josephson transmission line

The most basic Josephson transmission line (JTL) is shown in Figure 2.24. Its function is to transmit an SFQ pulse from the input source to the output load.




Figure 2.24: A basic RSFQ Josephson transmission line.

In order to derive design equations, the input to the JTL is modelled as a phase source ϕ_{in} that has the same startup or dc value as the phase over every junction. The phase is ramped up by 2π when an SFQ pulse is introduced to the circuit.

Each Josephson junction is biased with a current I_B and connected to the next junction through a superconducting inductance L . The Josephson junction critical current I_C is selected first. A practical value that can be manufactured reliably with a target fabrication process is chosen so as not to be too small to be susceptible to noise errors, and not so large that the junction does not behave like a short junction anymore. For a typical planar niobium aluminium oxide junction process with $J_C \leq 1 \text{ kA cm}^{-2}$, $I_C = 250 \mu\text{A}$ is a standard choice. For a 1 kA cm^{-2} process, such as the FLUXONICS process, the side length of a square junction will then be $5 \mu\text{m}$, while λ_J (2.18) is about $10 \mu\text{m}$. For a 10 kA cm^{-2} process, such as the MIT-LL SFQ5ee process, the diameter of a circular junction will be around $1.8 \mu\text{m}$, while λ_J is about $4 \mu\text{m}$. In all of these processes, a $250 \mu\text{A}$ Josephson junction has dimensions that are no more than half of the limit for short junction behaviour.

Each Josephson junction is shunted with a resistor to damp oscillation. The conventional approach is to make the Stewart-McCumber parameter $\beta_C \approx 1$ [47], but it has since been shown that RSFQ gate speed can be adjusted by varying β_C in the range 1-4 without appreciable degradation in circuit operating margins [52]. For the IARPA SuperTools project, where switching speed was of primary importance, we chose to set $\beta_C = 2$. From (2.17), with $I_C = 250 \mu\text{A}$, critical current density $J_C = 100 \mu\text{A } \mu\text{m}^{-2}$ and $C = 70 \text{ fF } \mu\text{m}^{-2}$ for the MIT Lincoln Laboratory SFQ5ee process, $R_{shunt} = 3.88 \Omega$.

The next step is to select a bias current value I_B , as a fraction a of the critical current a junction, so that

$$I_B = aI_C. \quad (2.44)$$

If a is too close to 1 it can be shown that a particular junction becomes susceptible to erroneous switching when dynamic loads nearby skew the current distribution in a circuit, and the critical margin for the junction area becomes small. If a is too close to 0, the

junction switching time becomes long. In practice a is selected in the range from 0.6 to 0.8. For this discussion, $a = 0.7$ so that $I_B = 175 \mu\text{A}$.

The final design choice, before any switching dynamics are taken into consideration, is the interconnect inductance L . When a junction switches it undergoes a phase change that will equal 2π if the junction has the same current through it after switching. Consider the circuit in Figure 2.24, where an incoming SFQ pulse has switched J_{-2} and J_{-1} and is just switching J_0 . As the junction switches, the phase ϕ_{J_0} ramps up and causes a phase difference $(\phi_{J_0} - \phi_{J_1})$ to appear over the inductor between junctions J_0 and J_1 . This phase difference causes current flow through the inductor to add to the current through J_1 . L should be chosen to allow enough current through it to switch J_1 , which is already biased with $I_B = aI_C$. The peak inductor current should thus be larger than $(1 - a)I_C$, or

$$L \leq \frac{\Phi_0 (\phi_{J_0} - \phi_{J_1})_{peak}}{2\pi (1 - a)I_C}. \quad (2.45)$$

If an SFQ input successfully switches the cascade of junctions, then the peak phase difference $(\phi_{J_0} - \phi_{J_1})$ will not exceed 2π . However, increase in current through J_1 causes its phase to increase *while the phase over J_0 is still rising*, so that the peak phase difference is actually roughly between π and $\frac{4\pi}{3}$. Design should accommodate the lower value of π , so that

$$L_{max} = \frac{\Phi_0}{2(1 - a)I_C}. \quad (2.46)$$

For the circuit shown here, the maximum limit to L is thus 13.8 pH. In simulation, with the parameters of a junction from the MITLL SFQ5ee 10 kA/cm² process, the JTL starts to exhibit delayed switching at $L = 15$ pH and fails to switch completely at $L = 18$ pH. Equation (2.46) underestimates the maximum inductance for a functional circuit, which is safe.

There is no theoretical lower limit to the inductance L to sustain SFQ pulse transmission, but practical requirements impose limits. The distance between Josephson junctions results in non-zero inductance, while the need to prevent large bias current redistribution when a dynamic load changes state favours larger inductance.

In practice, the transmission inductance is thus usually selected as $L = \frac{\Phi_0}{2I_C}$, or 4.1 pH in this example.

For the MIT Lincoln Laboratory SFQ5ee process, the JoSIM model for a Josephson junction with unit area of 1 μm^2 and $I_C = 100 \mu\text{A}$ is defined as:

```
.model jj1 jj(rtype=1, vg=2.8mV, cap=0.07pF, r0=160, rn=16, icrit=0.1mA)
```

With this junction, scaled by an area parameter to 250 μA , the JTL chain is simulated. The phase response is shown in Figure 2.25. The delay time t_{delay} of a single stage is 1.67 ps.

For the same MITLL SFQ5ee junction model, the delay time and margins for I_B and I_C are shown in Table 2.1 as a function of a , the fraction of bias current to critical current. It is clear that a lower value of a increases switching delay, while a higher value of a lowers the margins on bias current and junction critical current. Critical margins are only tested to $\pm 90\%$, and the margins on L never fall below 90% when a is adjusted between 0.5 and 0.9.




Figure 2.25: Simulated phase response of a basic JTL chain.

Table 2.1: Effects of bias current to critical current ratio on JTL delay and margins.

Bias factor a	0.5	0.6	0.7	0.8	0.9
Delay time (ps)	2.4	1.97	1.67	1.45	1.25
I_B margins (%)	-53/+90	-61/+67	-67/+43	-71/+25	-74/+11
I_C margins (%)	-50/+47	-40/+63	-30/+79	-20/+90	-10/+90

2.4.3.3 Symmetrical Josephson transmission line

Since the introduction of a symmetrical JTL by Polonsky *et al.* [53], with two Josephson junctions and a single, centred bias, as shown in Figure 2.26, most JTL circuits have been implemented in this way. There is no obvious improvement in margins or stability when the symmetrical JTL is connected to a circuit with mismatched input or output phase, but the symmetrical JTL allows cell layout with a similar footprint area as other logic gates, and reduces the area consumed by the bias resistor by a factor of four, with only one $7.14\ \Omega$ resistor for every two junctions compared to two $14.28\ \Omega$ resistors in the basic JTL.




Figure 2.26: A symmetrical RSFQ Josephson transmission line.

The component values determined for the basic JTL in Section 2.4.3.2 are used, so that $I_{C1} = I_{C2} = 250\ \mu\text{A}$, $I_B = 2aI_C = 350\ \mu\text{A}$, and $L_1 = L_2 = L_3 = L_4 = \frac{1}{2} \frac{\Phi_0}{2I_C} = 2.05\ \text{pH}$.

This symmetrical JTL allows straight-forward connection between cells: the inductance L_4 in series with L_1 of the next cell adds to a total of $\frac{\Phi_0}{2I_C}$ H for pulse transmission. When all cells in a library are designed to match the symmetrical JTL, the input and output inductances are thus designed to be $\frac{\Phi_0}{4I_C}$ H.

One requirement of the IARPA SuperTools project was to make RSFQ cells technology-independent by using parameterised descriptions. The parameterised values for the symmetrical RSFQ JTL, for arbitrary values of β_C , I_C and a , and with junction capacitance C scaled for I_C are:

$$I_B = 2aI_C$$

$$R_{shunt} = \sqrt{\frac{\beta_C \Phi_0}{2\pi I_C C}}$$

$$L_1 = L_2 = L_3 = L_4 = \frac{\Phi_0}{4I_C}$$

2.4.3.4 Basic storage element - the D Flip-Flop

A basic storage element can be created if a SQUID loop is used to store a circulating current that can be read out at a later time. It is often named a destructive readout register (DRO) or a D flip-flop (DFF). The latter designation is used here.




Figure 2.27: A basic RSFQ D flip-flop.

A basic RSFQ DFF is shown in Figure 2.27. A SQUID formed by junctions J_2 and J_3 and inductor L_2 provides two states: a “zero” state where the magnetic flux stored in the loop is zero, and a “one” state where the magnetic flux stored in the loop is one fluxon.

In the startup or zero state, J_2 is biased at approximately aI_B . J_3 is unbiased and sinks the total current arising from phase differences $(\phi_2 - \phi_3)$, $(\phi_{reset} - \phi_3)$ and $(\phi_o - \phi_3)$, which is less than the current through J_2 . Before any junction has switched, all phases are within 2π rad from the reference phase (ground).

The magnetic flux enclosed by the SQUID loop, from (2.43), is

$$\Phi = \frac{\Phi_0}{2\pi} [-\phi_2 + (\phi_2 - \phi_3) + \phi_3] = 0.$$

The DFF is set to store a fluxon when an input SFQ pulse is applied at *set*. Inductor L_2 must be designed to limit the current arising from the phase difference $(\phi_2 - \phi_3)$ so that J_3 does not switch. Junction J_1 is added in series with the *set* input to create a decision making pair (DMP). If a second *set* input is applied while the DFF is in the set state (and ϕ_{set} increases by another 2π rad), J_1 must switch to provide a 2π rad phase increase that prevents increased static current through L_1 into the DFF.

For reading out the DFF, an SFQ pulse is applied at the *reset* input. Junction J_4 forms a DMP with J_3 , so that J_4 switches when *reset* is pulsed while the DFF is in the zero state, and J_3 switches when the DFF is in the one state. The phase at ϕ_3 only increases when J_3 switches; in which case the phase difference between ϕ_2 and ϕ_3 reduces to reset

the SQUID loop, and the resulting phase increase between ϕ_3 and ϕ_o sends current to J_o to produce an SFQ output.

For circuit design, we assume that the inputs are connected to circuits identical to the basic JTL so that the closest junctions to each input have critical current I_C and are biased at aI_C . J_o at the output is assumed to have the same parameters. The dc (startup) phase $\phi_{set} = \phi_{reset} = \sin^{-1} a$, while it is assumed that $\phi_o \approx \sin^{-1} a$.

Junction critical current is selected first. For compatibility with neighbouring circuits, I_{C1} to I_{C4} can be set to I_C . For this discussion, $I_C = 250 \mu\text{A}$. (A note: in practice, the DFF is often implemented with $I_{C1} = I_{C4} \approx 0.9I_C$ to improve operating margins, but successful design is possible when all the junctions have identical I_C).

As for the JTL, shunt resistors are found from (2.17) for a selected value of β_C . In keeping with the selections made for the JTL, the junctions are slightly underdamped with $\beta_C = 2$, so that $R_{shunt} = 3.88 \Omega$.

The bias current I_B is selected next. To match the phase ϕ_2 to ϕ_{set} , I_B is selected as aI_C . Some current is diverted through L_2 because of the phase drop to ϕ_3 , but since a value is not designed for L_2 yet it is assumed that $\phi_2 \approx \sin^{-1}(aI_C/I_C) = \sin^{-1} a$.

For the example here, $a = 0.7$, which yields $I_B = 175 \mu\text{A}$ and $\phi_2 = 0.7754$ rad.

Lastly, inductance values are designed. As discussed in Section 2.4.3.2, the input and output transmission inductors are selected to have $L = \frac{\Phi_0}{2I_C}$.

From Section 2.4.3.3, half of the transmission inductance at each input and output is assigned to the preceding or succeeding circuit. These inductors, labelled L_x in Figure 2.27, are thus defined as:

$$L_x = \frac{\Phi_0}{4I_C}. \quad (2.47)$$

For standard input and output Josephson junctions with $I_C = 250 \mu\text{A}$, the value of L_x is 2.07 pH.

Now we have

$$L_3 + L_x = \frac{\Phi_0}{2I_C}, \quad (2.48)$$

or

$$L_3 = \frac{\Phi_0}{4I_C} = 2.07 \text{ pH}.$$

It is a good assumption that the dc currents through junctions J_1 and J_4 are close to zero when the DFF is in state zero. The total equivalent inductance of each junction, from (2.40) and the small-angle approximation $\sin \theta \approx \theta$, is then

$$L_{Jt}|_{I_J \approx 0} = \frac{\Phi_0 \sin^{-1}(I_J/I_C)}{2\pi I_J} = \frac{\Phi_0 I_J/I_C}{2\pi I_J} = \frac{\Phi_0}{2\pi I_C}. \quad (2.49)$$

Thus it follows that

$$L_1 = L_4 \approx \frac{\Phi_0}{2I_C} - \frac{\Phi_0}{2\pi I_C} - L_x = 0.8 \text{ pH}.$$

All that is left to design is L_2 , the storage inductor that enables the two states in the SQUID loop.

In the set or “one” state, J_2 has switched one time more than J_3 to store one fluxon in the SQUID loop. Although the currents through J_2 and J_3 differ, the phase difference ($\phi_2 - \phi_3$) can be approximated very well as 2π .

For a generic solution, let the desired current ratio I_{J3}/I_C when the DFF is in the set or “one” state be b . The current in J_3 is:

$$I_{J3} = \frac{\Phi_0}{2\pi} \left[\frac{\phi_2 - \phi_3}{L_2} + \frac{\phi_{reset} - \phi_3}{L_4 + L_{Jt4} + L_x} + \frac{\phi_o - \phi_3}{L_3 + L_x} \right] \quad (2.50)$$

$$= \frac{\Phi_0}{2\pi} \left[\frac{\phi_2 - \phi_3}{L_2} + \frac{2I_C(\phi_{reset} - \phi_3)}{\Phi_0} + \frac{2I_C(\phi_o - \phi_3)}{\Phi_0} \right]. \quad (2.51)$$

With ϕ_{reset} and ϕ_o close to $\sin^{-1} a$:

$$I_{J3} = bI_C \approx \frac{\Phi_0}{2\pi} \left[\frac{2\pi}{L_2} + \frac{4I_C(\sin^{-1} a - \sin^{-1} b)}{\Phi_0} \right]. \quad (2.52)$$

Rewriting (2.52) in terms of L_2 yields:

$$L_2 \approx \frac{\Phi_0}{I_C \left[b + \left(\frac{2}{\pi} \right) (\sin^{-1} b - \sin^{-1} a) \right]}. \quad (2.53)$$

It is possible to select b over the range from around 0.5 to 1, but if $b \neq a$ then the DFF will sink or source current from or to the *reset* input and the circuit at the output in the set state. If b is chosen as 0.7, it follows that $L_2 \approx 11.8$ pH.

At this stage, the design has to be evaluated. An analysis of the loop flux signature of the circuit for an exhaustive combination of inputs, described in detail in [54], yields the Mealy state diagram shown in Figure 2.28. The DFF has two states, inputs at *reset* in the “0” state and *set* in the “1” state do not cause state change, and an output pulse is only produced when *reset* is applied in the “1” state. The DFF thus behaves as intended.




Figure 2.28: Mealy state diagram of the RSFQ DFF. The two states are “0” and “1”, lowercase labels represent inputs, and the uppercase label with a filled circle represents an SFQ output.

If the DFF is tested under normal conditions, where none or only one set input is applied between any two reset inputs (so that set is never applied in the “1” state), the simulated circuit response is shown in Figure 2.29. With this response used to verify operation, the margins are calculated as shown in Figure 2.30.

The critical margin is +28% on I_C of junction J_4 (listed as “B4” in the margin output due to the JoSIM label), which appears to be good enough that no optimisation is necessary.

However, when the simulation is altered to test a set input in the “1” state, the margins are calculated as shown in Figure 2.31, with the critical margin now +10% on I_C of junction J_1

It is tempting to run this circuit through an optimiser, but more efficient to look at the source of failure. If I_C for junction J_1 is increased to 280 μ A in simulation (slightly




Figure 2.29: Simulated response of RSFQ DFF for set and reset inputs.

```

B1   : 84 [ ***** ] ***** ] 90
B2   : 90 [ ***** ] ***** ] 60
B3   : 44 [ ***** ] ***** ] 32
B4   : 30 [ ***** ] ***** ] 28
IB2  : 79 [ ***** ] ***** ] 87
L1   : 90 [ ***** ] ***** ] 90
L2   : 72 [ ***** ] ***** ] 90
L3   : 90 [ ***** ] ***** ] 90
L4   : 90 [ ***** ] ***** ] 90
Critical margin: 28% ['B4+']
    
```

Figure 2.30: Margins of first-pass DFF design without a test for a set input in the “1” state.

```

B1   : 84 [ ***** ] *** ] 10
B2   : 11 [ ***** ] ***** ] 60
B3   : 44 [ ***** ] ***** ] 32
B4   : 30 [ ***** ] ***** ] 28
IB2  : 79 [ ***** ] ***** ] 13
L1   : 90 [ ***** ] ***** ] 90
L2   : 72 [ ***** ] ***** ] 14
L3   : 90 [ ***** ] ***** ] 90
L4   : 90 [ ***** ] ***** ] 90
Critical margin: 10% ['B1+']
    
```

Figure 2.31: Margins of first-pass DFF design when a set input in the “1” state is included.

higher than the margin of failure), a set input in the “1” state switches J_2 instead of J_1 . That is incorrect operation. From simulation, the current in J_2 in the “1” state is $50 \mu\text{A}$. This needs to be lowered to almost zero. Without adjusting the bias current, we can reduce L_2 to divert more current to J_3 in the “1” state. If L_1 is chosen as 8 pH , which is very close to $\frac{\Phi_0}{I_C}$, the current through J_2 reduces to $2 \mu\text{A}$ in the “1” state, and the critical margin improves to $+30\%$ (on I_C of J_4).

The DFF becomes even more robust if I_C for both junctions J_1 and J_4 is reduced by 10% , with a critical margin of $+43\%$ on I_C of J_2 as shown in Figure 2.32. At this point, *further optimisation by computer is entirely unnecessary.*

B1	:	66	[*****		*****]	51
B2	:	64	[*****		*****]	43
B3	:	70	[*****		*****]	46
B4	:	62	[*****		*****]	44
IB2	:	56	[*****		*****]	72
L1	:	90	[*****		*****]	90
L2	:	59	[*****		*****]	90
L3	:	90	[*****		*****]	90
L4	:	90	[*****		*****]	90
Critical margin: 43% ['B2']								

Figure 2.32: Optimum margins of DFF for all possible input combinations.

The DFF cell can now be adjusted for any required standard critical current with parametric equations for the component values. For arbitrary values of β_C , I_C and a , and with junction capacitance C scaled for I_C , these equations are:

$$\begin{aligned}
 I_{C(J2)} &= I_{C(J3)} = I_C \\
 I_{C(J1)} &= I_{C(J4)} = 0.9I_C \\
 I_B &= aI_C \\
 R_{shunt(J2)} &= R_{shunt(J3)} = \sqrt{\frac{\beta_C \Phi_0}{2\pi I_C C}} \\
 R_{shunt(J1)} &= R_{shunt(J4)} = \sqrt{\frac{\beta_C \Phi_0}{2\pi I_C C (0.9)^2}} \\
 L_1 = L_4 &= \frac{\Phi_0}{4I_C} - \frac{\Phi_0}{2\pi I_C} \\
 L_2 &= \frac{\Phi_0}{I_C} \\
 L_3 &= \frac{\Phi_0}{4I_C}
 \end{aligned}$$

The use of parametric equations to describe the circuit allows I_C to be adjusted if the cell library is retargeted to a new standard junction critical current. The DFF remains fully functional, with the critical margin above 40% , when I_C is changed between $50 \mu\text{A}$ and $500 \mu\text{A}$ – and beyond – although this represents the range of interest.

The same technique can be applied to other logic circuit such as AND, OR and NOT gates to derive parametric equations for all circuit elements.

2.4.3.5 Bias resistors

Although RSFQ circuit schematics often show current sources for the bias currents, implementation in an integrated circuit is done with resistors. This is shown for the

symmetrical JTL in Figure 2.33.

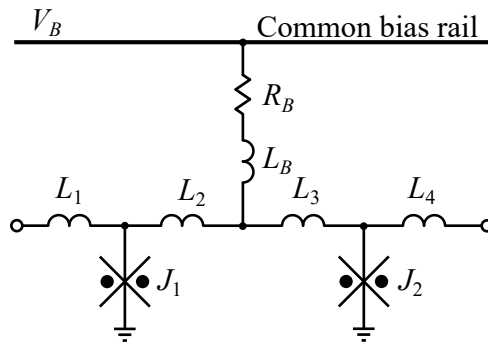


Figure 2.33: Schematic of a JTL showing resistive biasing and a common voltage rail.

The bias voltage V_B is selected to be much larger than the average voltage over the junctions generated by switching. In general [55], the bias voltage is selected as

$$V_B \approx 10I_C R_S, \quad (2.54)$$

where I_C is the critical current and R_S the shunt resistance of a typical junction.

For the FLUXONICS process, and with standard $I_C = 250 \mu\text{A}$ and $R_S = 1 \Omega$, (2.54) yields $V_B = 2.5 \text{ mV}$.

As shown earlier, the area of an SFQ pulse is 2.067 Vs . If the junction is switched at a frequency f , the average voltage over the junction is thus

$$V_J = \Phi_0 f. \quad (2.55)$$

At a frequency of 10 GHz , $V_J \approx 20 \mu\text{V}$. For a frequency of 100 GHz , $V_J \approx 0.2 \text{ mV}$, which is still much smaller than the selected value of V_B .

The bias resistor and any superconductor line segments that connect it between the common bias rail and the circuit have a combined inductance L_B . Although it is commonly ignored, it has a significant effect on circuit operation. This is shown in more detail in Section 2.6.1.

2.4.3.6 Design conclusion

More complex RSFQ logic cells can be designed with the technique described in this Section. Proper phase-based circuit analysis during the design phase has been shown to yield parameterised component values that allow for quick reassignment of designs to different fabrication processes or different standard Josephson junction critical currents.

The design technique also delivers nominal circuits with better margins than when components are chosen in an arbitrary fashion or when component values are varied in a circuit simulation until a circuit becomes functional, so that less optimisation time is required.

2.5 Contributions to RSFQ

2.5.1 Logic cells

I designed a set of basic RSFQ logic gates for my Masters degree [56] before developing phase-based circuit analysis theory or before I had access to tools for layout verification. It

was this exercise that propelled me into design tool development and inductance extraction research.

My postgraduate student Dr Rodwell Bakolo designed an RSFQ cell library [57] with the tools that were developed for the NioCAD project and InductEx. It is described in detail in his Masters thesis [58]. One important step forward was the introduction of single-clock NAND, NOR and XNOR gates [59].

2.5.2 Special-purpose cells

2.5.2.1 DCRL

As a cell required for programmable logic, I developed a DC-resettable latch (DCRL) [60]. The DCRL is a non-destructive readout set-reset cell that can only be reset when a DC current, coupled magnetically from an isolated reset line, is ramped up beyond the reset threshold. The reset current can be threaded past any number of DCRL cells, so that the entire switch fabric of a circuit such as a superconducting programmable gate array (SPGA) can be reset with a single signal.

An improved layout with a circuit test [61] was done for the IPHT FLUXONICS process after I demonstrated that InductEx is accurate for the calculation of inductance over ground plane holes. The circuit schematic of the DCRL is shown in Figure 2.34.

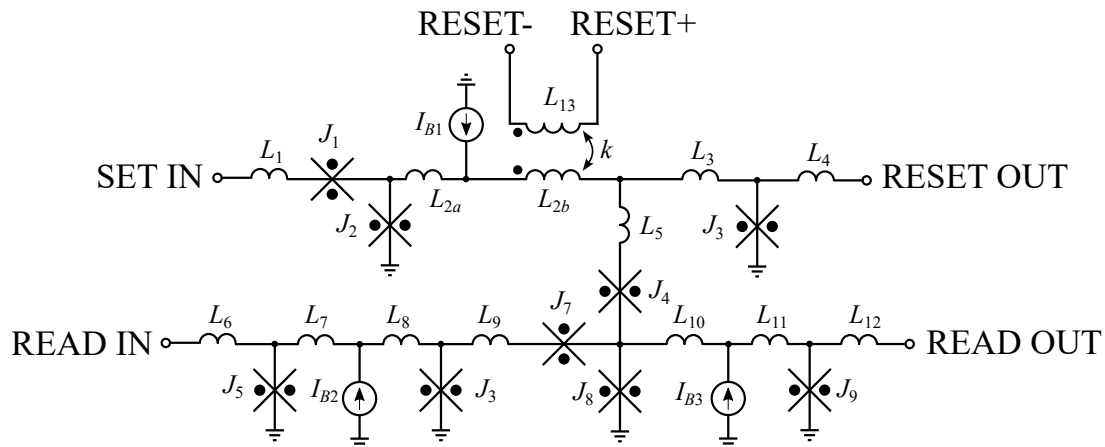


Figure 2.34: Circuit schematic of an RSFQ DC-resettable latch.

The circuit was laid out for the FLUXONICS process, and a ground plane hole was used to improve the mutual inductance between the reset line and the circuit. The fabricated circuit is shown in Figure 2.35. Test results are only available for the inductance of L_{2b} and the mutual inductance to the reset line, but not for the logic functionality of the circuit.

2.5.2.2 RSFQ-COSL output driver

My PhD supervisor, Professor Willem Perold, was with the group of Professor Ted van Duzer at UC Berkeley when they developed the COSL family of ac biased voltage-state superconductor logic. He had students working on COSL circuits, and I contributed logic cells such as a set-reset flip-flop [62].

The cell that stood out, though, was an RSFQ-to-COSL converter that would allow easier readout of RSFQ outputs to room-temperature electronics. [62]. The circuit




Figure 2.35: Microphotograph of an RSFQ DC-resettable latch fabricated with the FLUXONICS process.

schematic is shown in Figure 2.36, simulation results in Figure 2.37 and a microscope photograph of a the manufactured circuit (here with the RSFQ and COSL bias inputs separated) in Figure 2.38.

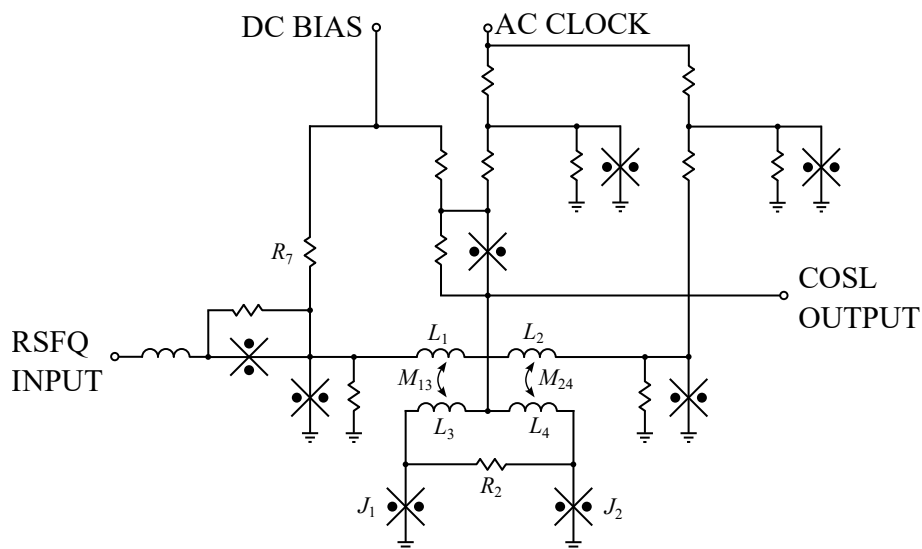


Figure 2.36: Circuit schematic for an RSFQ-to-COSL converter.

The layout of this circuit exposed the limits of techniques used for integrated circuit inductance estimation when the coupling structure for inductors L_1 to L_3 and L_2 to L_4 was laid out.

2.5.3 A superconductor programmable gate array

The field-programmable gate array (FPGA) [63] is an incredibly powerful and popular integrated circuit for low-volume niche digital applications, and is indispensable in systems that require rapid reconfiguration.

In a superconductor electronics system, where the design and fabrication turnaround time is long, and where replacement of an integrated circuit in a cryogenic environment

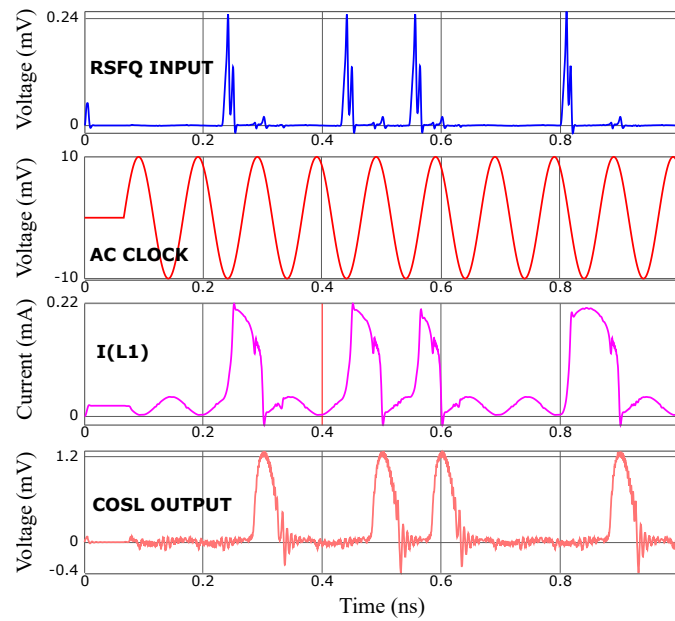


Figure 2.37: Simulated transient response of the RSFQ-COSL converter at a clock frequency of 10 GHz.




Figure 2.38: Microphotograph of RSFQ-to-COSL converter manufactured in Hypres 4.5 kA cm^{-2} process.

can be a slow and cumbersome process, the possibility to reconfigure a circuit while it is cryocooled would ease system design and upkeep. Reprogrammable circuits such as FPGAs would be ideal for this.

The idea of a superconducting programmable gate array (SPGA) originated from my supervisor, Prof Willem Perold and his then final year undergraduate student (and my contemporary) Peter Gross. A functional SPGA would make repurposing of an SCE IC much easier, as circuit functionality can be changed programmatically while the circuit is at cryogenic temperatures. Perold and Gross focused on application, while circuit implementation fell to me as one of the goals of my PhD research.

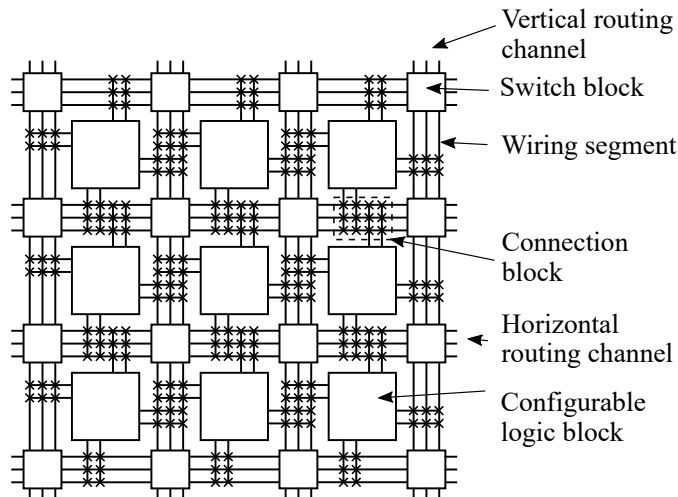


Figure 2.39: Schematic diagram of a general architecture for a symmetrical array FPGA.

The general architecture of a mesh or island style FPGA is shown in Figure 2.39. The configurable logic blocks consist of reprogrammable lookup tables (LUTs), connection blocks that have reprogrammable switches that connect the input and output lines of a logic block to the routing tracks, and switch blocks have reprogrammable switches that allow connection of tracks from horizontal and vertical routing channels.

Some design selections are required, such as:

- The number of inputs K to the LUT in a configurable logic block and the size 2^K of the LUT.
- The number of tracks in a routing channel.
- The number and pattern of switches in a switch block. If every vertical track can connect to every horizontal track, the switch block size is vast and most switches will be superfluous in a routed solution, while sparse switch patterns limit routability [64].

I chose to implement the entire SPGA fabric in RSFQ [65], but had to design some circuit blocks to make it possible. One of these was the DCRL [60], which formed the programmable bits of every LUT as well as the programmable contacts in connection and switch blocks. Another was a functional bipolar current circuit, the hybrid unlatching flip-flop logic element (HUFFLE) [66]–[68] that was used to switch a bipolar current through a large loop inductor that could couple magnetically to tens of cells.

Routing tracks were unidirectional, and composed of Josephson transmission lines. A matrix programming strategy was devised in order to programme the SPGA. Two

concatenated shift registers were loaded in series, with one holding the row programming data pattern and the other the column access selection that would set the current. Switches could be set per column when a HUFFLE for a specific column was switched to energise switches coupled magnetically to the HUFFLE loop inductor. The entire SPGA could be reset – with all switches “open” and lookup tables cleared – through the application of a single dc current pulse that threaded all the DCRL circuits.

The SPGA circuit blocks were demonstrated through the implementation of a programmable frequency divider [65], as the cells were too large to allow a functional SPGA to be implemented on a single chip.

My Masters student Hein van Heerden finally put together a full circuit implementation that included a very limited set of logic blocks and programming fabric [69] and used the same building blocks and programming matrix developed for my PhD. The implementation of the programming frame and LUT decoder is discussed in detail in [69] and [70]. The layout was the first ever to use InductEx for intra- and inter-gate inductance extraction – an absolute necessity due to the complexity of the coupling structures in the HUFFLE and the braided multi-line coupling to the LUT input decoders. The use of InductEx also made it possible to reduce the cell sizes sufficiently to allow four configurable logic blocks and the routing architecture and switch blocks with the complete programming fabric to fit on a single 5 mm×5 mm die for the Hypres 4.5 kA cm⁻² process. The SPGA schematic is shown in Figure 2.40 and the chip layout is shown in Figure 2.41.

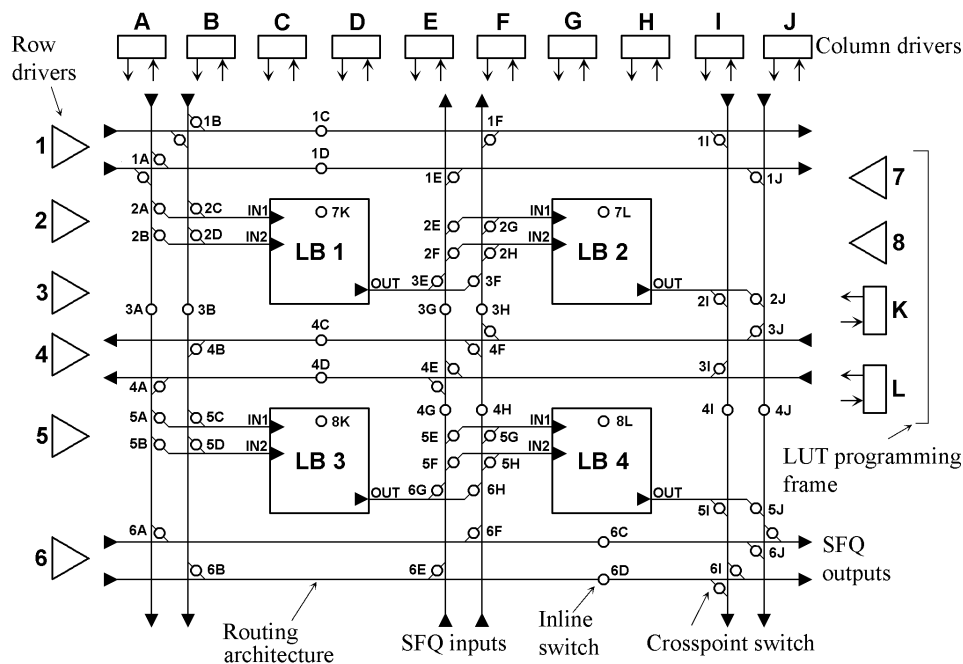


Figure 2.40: Schematic of a 4-LUT SPGA with routing architecture and switches. Programming fabric is omitted.

The final design used 4250 junctions and required a bias current of 560 mA. Of the layout area, 5% was devoted to the LUT programming frame, 10% to the 4 logic blocks, 20% to the 10-column by 6-row programming frame, and 65% to the switch matrices and routing architecture. Routing used the same metal layers as logic cells, so that a more complex SPGA would only become possible when fabrication technology improved to allow passive transmission line routing over (or under) the logic circuits.

Even though the SPGA was limited and dominated by the switching fabric, it has




Figure 2.41: Layout of a 4-LUT SPGA with HUFFLE bipolar drivers on a $5\text{ mm} \times 5\text{ mm}$ chip for the Hypres 4.5 kA cm^{-2} process.

served as a reference design for improved superconductor FPGAs [71]–[73].

A more comprehensive design was done by my PhD student Calvin Maree [74]. Calvin found that the HUFFLE-based programming fabric could be replaced by a much more efficient serial programming structure that used a Destructive-Shift Nondestructive Read-Out (DSNDO) cell as a switch and for LUT memory. Calvin formalised the programming by using the academic tool Versatile Place and Route (VPR) [75] as part of the Verilog-To-Routing (VTR) [76] open source FPGA EDA tool flow to map logic to the SPGA.

The addition of a bypassable output latch provides support for both combinational and sequential circuits, while decoding of the address for LUT read-out is done with either an RSFQ demultiplexer at the input of a LUT or an RSFQ multiplexer at the output of the LUT.

Calvin highlighted other issues too. Due to the clocked nature of RSFQ circuits, reliably levelling the outputs from different logic blocks requires either a data-driven self-timed (dual-rail) clock strategy [77] – which doubles the already area-intensive routing and switching fabric – or a dual clock approach with a fast clock and a slow clock. Both options were investigated.

Calvin also showed that a dense switch block, where every input can connect to every output, is unnecessarily expensive because most switch blocks are sparsely connected. The Wilton switch block [78], adapted for unidirectional signal transmission and illustrated for a width of four channels in Figure 2.42, provides good routing flexibility and efficient switch sparsity. It was thus used for SPGA design.




Figure 2.42: Connection diagram of a 4 track wide unidirectional Wilton switch block.

As an example of the use of configurable logic blocks and routing resources in an SPGA a purely combinational ripple carry adder (RCA) was mapped to a 7×7 -tile dual-rail SPGA by using ODIN [79], ABC [80] and VPR [75], [81]. The VPR placement and routing results are shown in Figure 2.43. The SPGA architecture uses 3-input LUTs and 4-channel wide routing with Wilton switch blocks as depicted in Figure 2.42. The 8-bit input vectors a and b , the 8-bit output vector sum and the carry-in and carry-out signals are mapped to the pins on the periphery.

Demonstration of the programming and successful simulation of the RCA in an SPGA architecture showed that small but useful reprogrammable circuits could be implemented on an SPGA, although the low density of existing SFQ logic cell layouts and fabrication processes did not allow for any of the MCNC20 benchmark circuits [82] to fit on a $5 \text{ mm} \times 5 \text{ mm}$ die with the technology available by 2019.

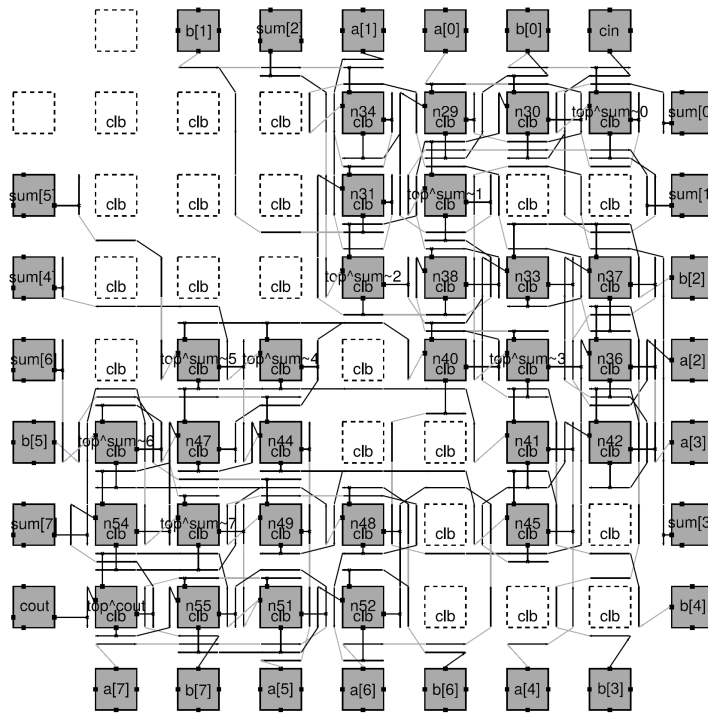


Figure 2.43: VPR place and route visualisation showing the use of configurable logic blocks for a combinational 8-bit ripple carry adder.

2.5.4 Asynchronous logic: RSFQ-AT

One of the drawbacks of standard RSFQ is that the logic cells need to be clocked individually. Despite the extremely fast switching time of RSFQ logic, system throughput is constrained when a logic cell has to wait for a clock pulse to produce and pass an output to the next logic cell. In general, where all logic cells are clocked at the same time, this limits data propagation to one logic level per clock cycle. An RSFQ microprocessor which has several logic levels in the instruction decoder and arithmetic logic unit would thus require multiple (tens or more) of clock cycles to complete one instruction cycle.

There are ways to improve the efficiency, such as designing a circuit for wave-pipelining [83]–[85] to improve throughput. However, fully asynchronous circuits would make system implementation much easier.

Retief Gerber, who I co-supervised with Prof. Willem Perold, introduced a concept of asynchronous RSFQ circuits that have timing elements and a clock line that accompanies each SFQ output and input [86]. This asynchronous family was named RSFQ asynchronous transmission (RSFQ-AT), and a logic cell would produce a clock signal for each output as soon as all the input clocks have arrived. A schematic diagram for a two-input RSFQ-AT logic cell is shown in Figure 2.44.

RSFQ-AT differs from dual-rail asynchronous RSFQ logic [87] where every signal is accompanied on a separate wire by its complement. In theory, RSFQ-AT can be laid out with a smaller footprint than a dual-rail asynchronous cell.

Retief selected the Muller C-element [88] for clock synchronisation, and demonstrated the concept on a half-adder and a full-adder [86]. The RSFQ-AT adder circuits are shown in Figure 2.44, where the AND, XOR and OR gate symbols represent RSFQ-AT AND, XOR and NOT gates respectively.

The simulated voltage response of the RSFQ-AT full-adder to the inputs a , b and




Figure 2.44: Schematic diagram of a generic two-input RSFQ-AT logic cell.




Figure 2.45: RSFQ-AT implementation of the (a) half-adder and (b) full-adder.

carry in (ci) and their respective clock signals, with outputs and the outputs *sum* (s) and *carry out* (co), together with the clock for each signal, is shown in Figure 2.46. The full-adder has three logic levels, and would thus require three clock cycles to compute co and two clock cycles to compute s with standard RSFQ logic.

A review of asynchronous SFQ technologies by my PhD student Dr Louis Müller [89] found that RSFQ-AT yielded reasonable speed with compact layout for smaller systems, but that the clock-follow-data requirement could be problematic in large systems (although lock-step passive transmission line routing would negate most problems). Delay-insensitive RSFQ [90] was found to be faster, more compact and more robust than the other asynchronous methods at the time.

2.5.5 Cell libraries

My involvement with cell library design started from a tool design perspective, when Retief Gerber and I worked on ways to define cells for easy migration between fabrication processes [91]. Some of the ideas around technology portable layout are still in use today for the layout synthesis tool SPiRA that my group developed under the IARPA ColdFlux project.

RSFQ NOR, XNOR and NAND cells were developed by Dr Rodwell Bakolo [57] while I was his PhD supervisor. Rodwell optimised the cells to work with both the FLUXONICS 1 kA cm^{-2} and the Hypres 4.5 kA cm^{-2} processes, and created a layout for each cell in both fabrication technologies. The difference in junction and feature size resulted in layout sizes of $300 \mu\text{m} \times 300 \mu\text{m}$ for the NAND and NOR gates in the FLUXONICS process compared to $200 \mu\text{m} \times 200 \mu\text{m}$ for the same gates in the Hypres process. For the XNOR gate, layout size was $450 \mu\text{m} \times 300 \mu\text{m}$ for FLUXONICS and $300 \mu\text{m} \times 200 \mu\text{m}$ for Hypres. Although large, it fitted with the general FLUXONICS library at the time where a JTL cost $150 \mu\text{m} \times 150 \mu\text{m}$ and a DFF cost $300 \mu\text{m} \times 150 \mu\text{m}$.

Due to the lack of a testing budget, no results were obtained for Rodwell's cell library, but the design and layout experience stayed with my group until it was needed for a much




Figure 2.46: Simulated response of the RSFQ-AT full-adder for the Hypres 4.5 kA cm^{-2} process of 2004.

more complete cell library.

Under the IARPA SuperTools programme's ColdFlux project [92], our team was obligated to deliver a complete RSFQ cell library that could:

- be designed with the tool chain developed under ColdFlux,
- be extracted and verified with the ColdFlux tool chain, and
- could be used for timing extraction, gate-level synthesis, placement and routing of RSFQ-based superconductor integrated circuits with the ColdFlux tool chain.

The ColdFlux requirements meant in practice that a complete cell library could not be designed until decisions about how cells would be placed and routed were in place. The planarisation of the MIT Lincoln Laboratory SFQ5ee and SFQ6ee processes, and the design rule limitations imposed by maximum and minimum layer fill requirements, meant that the cell library had to be designed with care. The project provided valuable insight into how cells had to be designed not just for optimum margins and maximum speed, but also for design rule limitations, flux trapping robustness, track-based placement and automated routing.

2.5.5.1 Routing architecture

In order to allow automated place and route tools to work, cell layouts need to conform to certain rules derived from the limitations and requirements of the fabrication process. Under ColdFlux, we opted for row-based place and route [93], which is how most semiconductor logic circuits are placed and routed. This is illustrated in Figure 2.47.

Under the ColdFlux place-and-route strategy, logic cell rows are alternated with clock distribution rows which consist mainly of pulse splitter cells. Each cell has a width and height that is an exact integer multiple of a *smallest track block*. Each track block must




Figure 2.47: Schematic representation of the row-based place and route strategy used for RSFQ circuits in ColdFlux.

allow one signal interconnect line to pass through it. Cell height is set to a fixed value for all cells in a library version, while width is adapted to fit all the required components into a cell. Rows can be abutted directly against the bias lines top and bottom, but may be made wider during place-and-route to find a routing solution in a densely connected system.

The distance between cells is unknown until after placement, but can be anything from a few cell widths (tens to hundreds of micrometres) to about twice the side length of a chip (ten to fifty millimetres). We thus opted for dedicated passive transmission line (PTL) signal interconnects between all cells.

For superconductor logic, signal interconnects with PTL have been demonstrated before [94] and have been used with great effect to realise complex SFQ circuits such as microprocessors [95]. One advantage of PTL is that it is designed for a given characteristic impedance to which every cell's signal output drivers and input receivers are matched. With the same characteristic impedance for every PTL, and the same layer thicknesses for all routing layers, the width of every PTL is the same. Since the PTL interconnects must be able to route "over" a cell, but all layers above the main ground plane M4 are used for circuit electronics, routing is done below the ground plane. We assigned the electronics and the routing resources of the ColdFlux layout stack in the MIT-LL SFQ5ee process as illustrated in Figure 2.48.

PTL conductors require good ground planes, regular stitches between multiple ground planes (for stripline) and vias with ground pins for layer transition. Layout is made more complicated by the minimum and maximum fill requirements of a fabrication process such as the MITLL SFQ5ee process. The minimum and maximum fill limits apply locally (for every $200\ \mu\text{m} \times 200\ \mu\text{m}$ block as well as over the entire chip area. It is thus not possible to cast a mostly solid ground plane for the transmission lines (or even for the logic circuits).

The very successful Japanese CONNECT cell library uses a $30\ \mu\text{m} \times 30\ \mu\text{m}$ [96] track block that allows PTL to cross horizontally and vertically. It uses $\beta_C = 2$ for all logic cells, and has a DC power plane on the lowest metal layer. The basic cell tile provides for two routing layers in each axial direction underneath a cell. Keeping in mind the lessons I learned from analysing the CONNECT track block for ground return current effects [97], and incorporating the design rule limitations of a planarised process, I designed a track block solution that would draw on the strengths of the CONNECT track block and work for the ColdFlux cell library [98].

2.5.5.2 Routing track block

The ColdFlux track block is shown in Figure 2.49. It was fixed at a size of $10\ \mu\text{m} \times 10\ \mu\text{m}$, with holes in the lower ground plane M0 to serve as flux trapping moats and to meet the maximum layer fill requirement. The track block can be tiled over the entire active chip area. Stripline PTL with width $4.5\ \mu\text{m}$ fits inside the track block with sufficient clearance to the ground contacts in the four corners to meet minimum distance separation requirements on every layer. The PTL has characteristic impedance $Z_o = 5.35\ \Omega$ [98]. PTL can be routed in layer M3 with M4 and M2 as ground, or in layer M1 with M2 and M0 as ground. These are called PTL1 and PTL2 respectively for routing purposes.

A rendering from an InductEx model in Figure 2.50 shows how the PTL conductors fit into a series of track blocks strung together.

I created a set of build-and-fill rules [98] that, when applied by a layout synthesis tool after place-and-route, builds the track routing, ground fill and electronics layers as illustrated in Figure 2.51.




Figure 2.48: Simplified illustration of a cross-section of the ColdFlux layout stack in the MIT-LL SFQ5ee process showing the assignment of passive transmission line layers.




Figure 2.49: Dimensions of the basic routing track block. All dimensions are in μm .




Figure 2.50: Cross-section of a three-dimensional simulation model for the M3-to-M1 stripline transition with an optimally filled via. The vertical dimension has been scaled up for clarity. The current density profile when the input on M3 and the output on M1 are excited is shown, as calculated with InductEx.

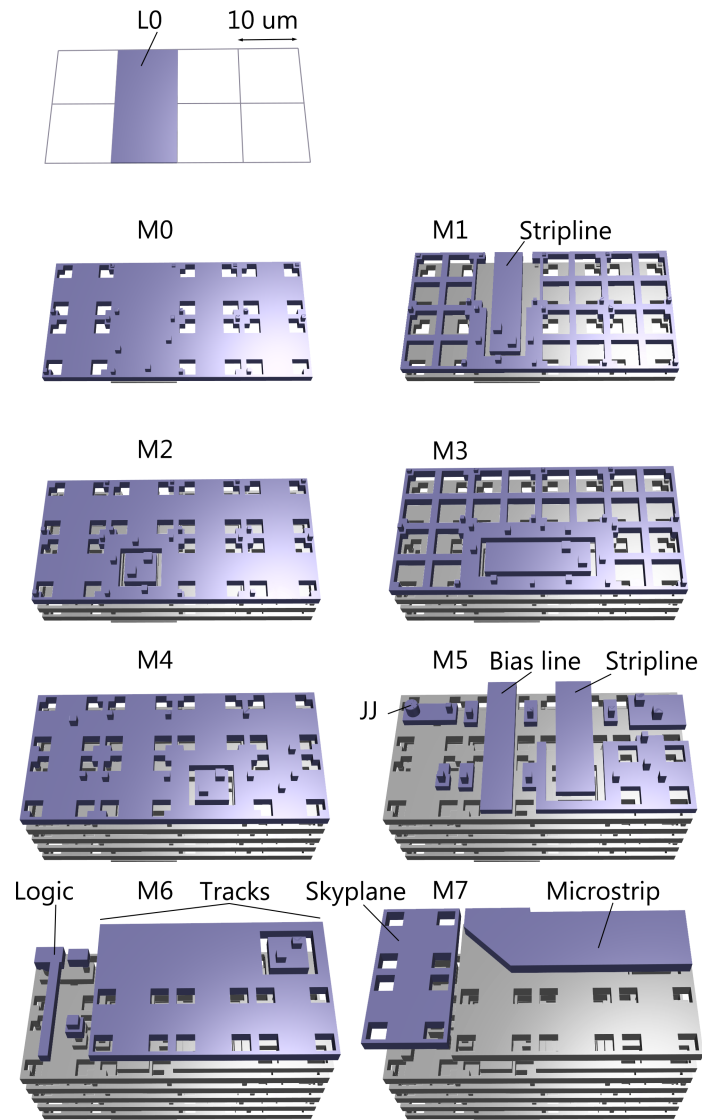


Figure 2.51: Three-dimensional rendering of an arbitrary 4×2 track block composition for the MIT-LL SFQ5ee process. The vertical dimension has been scaled up by a factor of eight for clarity.

2.5.5.3 ColdFlux RSFQ cell library

As part of the SuperTools ColdFlux project [92], Lieze Schindler designed and implemented an RSFQ cell for her PhD [50], [99]. The cell library was designed for the MIT Lincoln Laboratory SFQ5ee process. An example of an RSFQ pulse splitter with integrated PTL drivers and receivers, laid out for the track block architecture to $40\ \mu\text{m} \times 50\ \mu\text{m}$ in size, is shown in Figure 2.52.

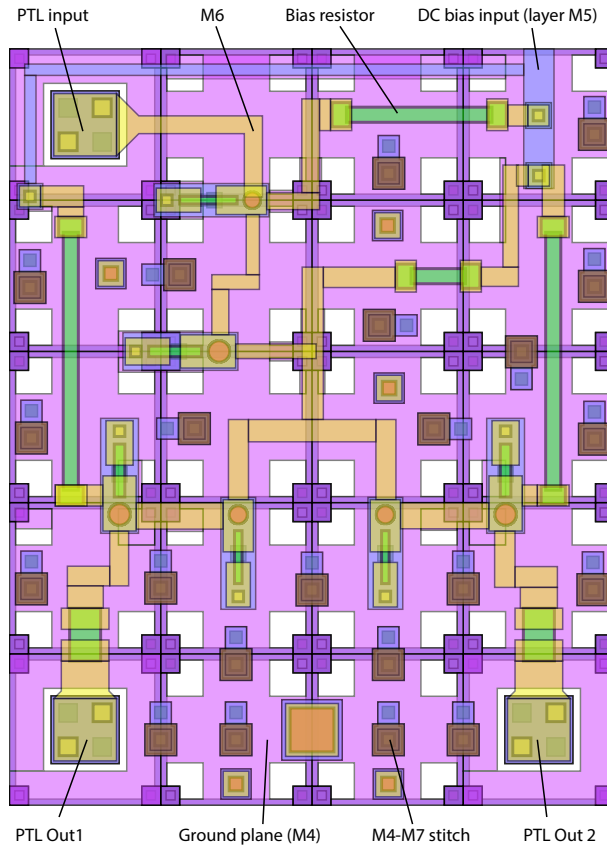


Figure 2.52: An RSFQ splitter cell layout that fits the routing block architecture. The layout is $40\ \mu\text{m} \times 50\ \mu\text{m}$ in size.

All cells were analysed for margins and yield roll-off, and optimised with an optimiser developed for the ColdFlux project. Many of the cells were fabricated by MIT Lincoln Laboratory under SuperTools as part of ColdFlux-allocated chip runs and tested with a liquid helium immersion probe by NIST in Boulder, Colorado. Gates worked successfully at low frequency (in the kilohertz range). We lacked the infrastructure to add high frequency (Gigahertz range) testbenches and interconnects on the chips, and the test harness used by NIST did not support microwave feed-in. The bias margins for successful operation at low frequency are shown in Table 2.2.

The RSFQ cells delivered under ColdFlux are listed in Table 2.3. The list includes cells without PTL input receivers or output drivers, with inputs and outputs placed for direct interconnect through standard interface inductance when cells are abutted.

For our initial RSFQ cell library, Lieze laid out conservatively with large distances between components and by including bias lines in separate track blocks *inside* each cell to minimise coupling from the bias lines to circuit elements. After the successful cell tests, my PhD student Ms Tessa Hall optimised the cell library for smaller layouts. We elected to provide bias inputs both at the top and bottom of each cell (where previously the bias

Table 2.2: Measured bias margins for fabricated ColdFlux RSFQ cells.

Cell	Measured bias current operating margins (mA)					
	Run 1	Run 2	Run 3	Run 4	Run 5	Run 6
DFFT	4.82-5.99	4.90-5.90	5.00-5.99	4.76-6.05	4.72-5.99	4.80-6.05
NDROT	6.50-7.09	6.25-7.09	6.29-7.09	6.14-6.60	6.14-7.22	6.18-7.18
NOTT	4.79-5.55	4.70-5.50	4.77-5.45			
OR2T	5.66-6.90	5.50-6.87	5.53-6.93			
XORT	6.04-7.75	6.37-7.75	6.04-7.67	6.11-7.72	6.55-7.10	6.05-7.14

inputs were only fed from the top). Tessa removed the shielded bias lines inside each cell and compressed the component spacing while leaving the moats intact. The result is on average a 50% reduction in cell size, with some cells reduced by up as much as 67%. The layout of the OR2T cell is shown in Figure 2.53 for both library versions as an example.

Table 2.3: List of ColdFlux RSFQ library cells.

Interfacing Cells	
DCSFQ	DC to SFQ pulse converter
DCSFQ-PTLTX	DC to SFQ pulse converter with PTL transmitter
SFQDC	SFQ pulse to DC converter
PTLRX-SFQDC	SFQ pulse to DC converter with PTL receiver
Interconnects	
JTL	Josephson transmission line
JTLT	PTL connection Josephson transmission line
SPLIT	Splitter
SPLITT	PTL connection splitter
MERGE	Merger
MERGET	PTL connection merger
PTLTX	PTL transmitter
PTLRX	PTL receiver
Always0 (sync)	Synchronous always zero
Always0 (async)	Asynchronous always zero
Always0T (sync)	PTL connection synchronous always zero
Always0T (async)	PTL connection asynchronous always zero
Buffers	
DFF	D flip-flop
DFFT	PTL connection D flip-flop
NDRO	Non-destructive readout
NDROT	PTL connection non-destructive readout
BUFF	Buffer for clock balancing
BUFFT	PTL connection buffer for clock balancing
Logic Cells	
AND2	2-Input AND
AND2T	PTL connection 2-input AND
OR2	2-Input OR
OR2T	PTL connection 2-input OR
XOR	Exclusive OR
XORT	PTL connection Exclusive OR
NOT	Inverter
NOTT	PTL connection inverter
XNOR	Exclusive OR with inverter
XNORT	PTL connection exclusive OR with inverter

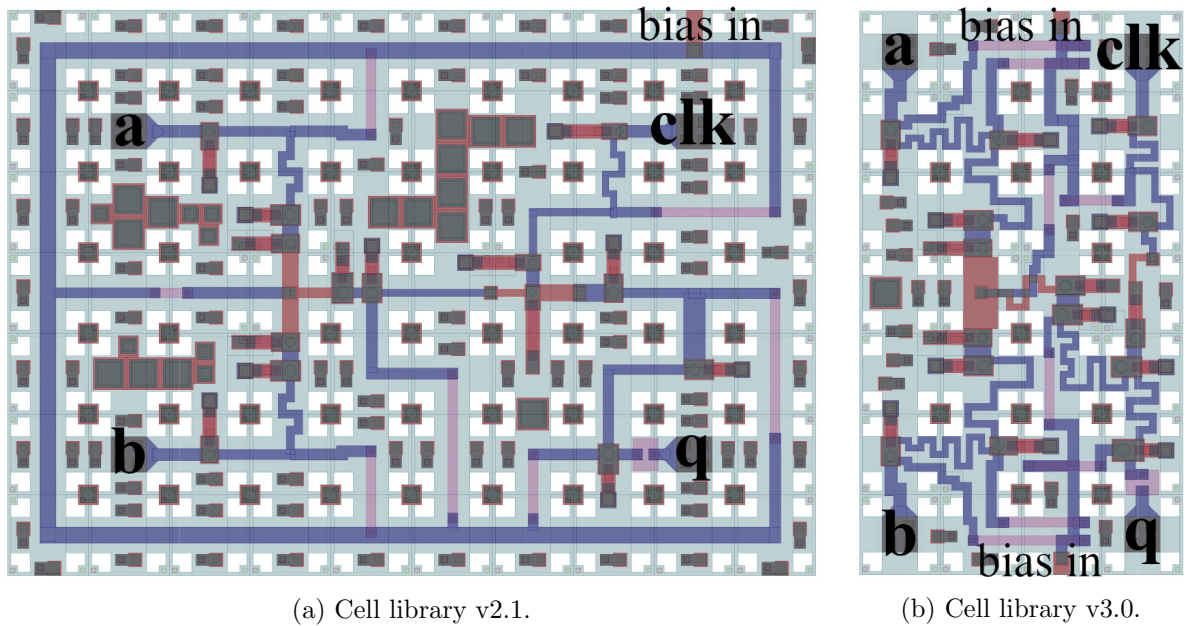


Figure 2.53: Layout image of an RSFQ OR2T cell in the ColdFlux library. Some layers have been omitted for clarity. The v3.0 cell is 60% smaller than the v2.1 cell.

2.5.6 Microprocessors

My PhD student Dr Louis Müller demonstrated that RSFQ-AT could be used at a system level by designing a microprocessor that utilised RSFQ-AT in combination with conventional RSFQ cells. The microprocessor was designed [100] according to the system schematic shown in Figure 2.54.

The hybrid RSFQ/RSFQ-AT microprocessor used a 4-bit address word, a 9-bit instruction word, 4-bit data word, a pipeline with fetch and decode/execute, a programme counter, instruction register and accumulator. A 4×16 -bit operand memory and a 9×16 -bit instruction memory were used. The operations *Load*, *Store*, *Add*, *Subtract*, *Jump*, *Conditional jump*, *AND*, *OR*, *XOR*, *Push*, *Pop* and *Stop* were supported.

The RSFQ/RSFQ-AT microprocessor required 5300 Josephson junctions, but was only simulated in Verilog. Funding constraints, and limitations to fabrication technology at the time meant that the microprocessor was never fabricated. However, the design process strengthened the impression on me that a full suite of design tools was required to enable fast, efficient and reliable design of any complex SFQ system.

2.6 Beyond RSFQ: ultra-low power logic

RSFQ was developed during a time when integration density was very low and the research focus was on ultimate speed with little concern about power consumption.

However, as presented in great detail by Mukhanov in his seminal paper on energy efficiency of SFQ technologies [55], the static power dissipation of resistor-biased RSFQ can be two orders of magnitude larger than the dynamic power dissipation at full switching speed, which limits the efficiency of very large scale RSFQ systems.

Figure 2.55 shows the resistive biasing of a conventional RSFQ circuit. As shown earlier, the voltage over a junction integrates to Φ_0 over time for every switching event

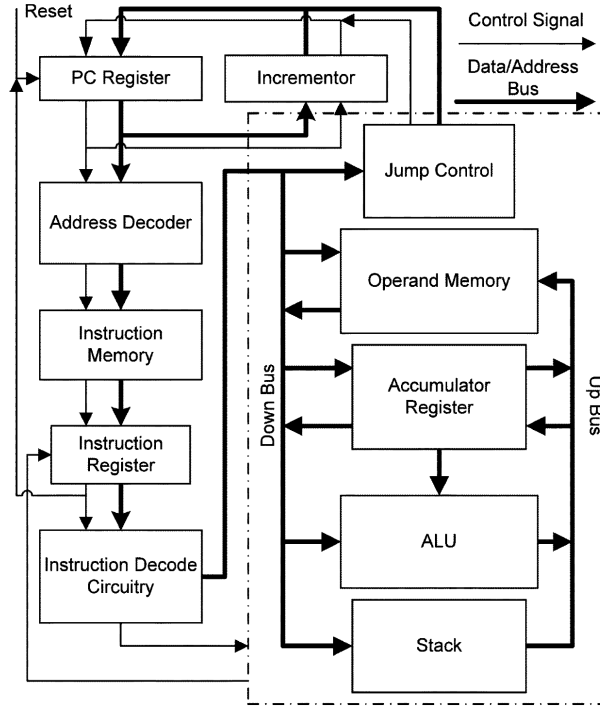


Figure 2.54: System schematic of a basic microprocessor.

that increases phase by 2π . If every biased junction in a cell switches at a switching frequency f , the average voltage of the cell is

$$v_i = \Phi_0 f. \quad (2.56)$$

The dynamic power dissipation in RSFQ circuits – used when switching Josephson junctions – is thus

$$P_D = I_B \Phi_0 f, \quad (2.57)$$

where I_B is the total bias current of the RSFQ circuit and f is the clock frequency.

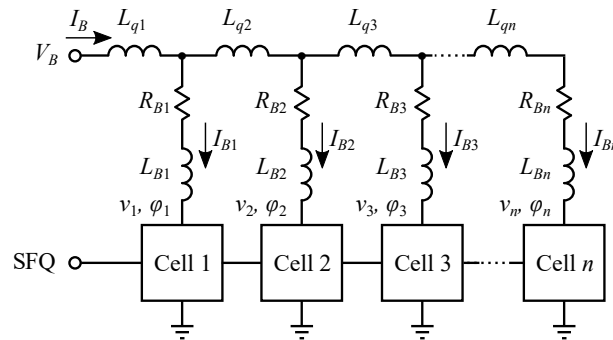


Figure 2.55: Conventional RSFQ biasing.

The static power dissipation in an RSFQ circuit is

$$P_S = \sum_{i=1}^n I_{B_i}^2 R_{B_i}. \quad (2.58)$$

If $V_B \gg v_i$, then

$$P_S \approx V_B I_B. \quad (2.59)$$

The ratio of static to dynamic power dissipation is then

$$\frac{P_S}{P_D} \approx \frac{V_B}{\Phi_0 f}. \quad (2.60)$$

For typical bias voltages of a few millivolt – we commonly use $V_B \approx 2.5$ mV, from (2.54) as discussed in Section 2.4.3.5 – and clock frequencies in the range of 10 GHz to 100 GHz, this ratio ranges from one to two orders of magnitude. At $f = 100$ GHz, $P_S \approx 12P_D$.

2.6.1 Low voltage bias

One way to reduce the static power consumption P_S of an RSFQ circuit is to reduce the bias voltage V_B [101]. Such *LR-biasing* has been investigated experimentally, with limitations characterised [102], [103]. The limitations can also be observed in simulation.

Static power dissipation decreases linearly as the bias voltage V_B is lowered. The value of the bias resistor R_B also decreases, so that the inductance of the bias branch, L_B , has to be increased (see Figure 2.33 and Figure 2.55 for symbol definitions). As discussed in [102], the time constant

$$\tau = \frac{L_B}{R_B} \quad (2.61)$$

must be maintained in the range

$$\Delta t \ll \tau \ll T, \quad (2.62)$$

where Δt is the SFQ pulse width for a given junction and T is the clock or switching period of the RSFQ cell.

If the LR time constant τ is too low – and thus L_B too small – the operating margins of a cell are reduced significantly as bias current reduces when a junction switches. If L_B is too large, the maximum switching frequency of the cell is reduced [55], [102]. This has been demonstrated experimentally [104], where bias voltage as low as 20 μ V was demonstrated at a maximum operating frequency of 4.7 GHz.

It is possible to reduce the static and dynamic power dissipation with half-flux-quantum (HFQ) SFQ circuits that use a combination of conventional (0-shifted) and π phase-shifted Josephson junctions [105], [106], and ultimately combine such HFQ circuits with low-voltage LR biasing, as was demonstrated recently [107].

My contribution to LR biasing was in the determination of the inductance L_B to satisfy (2.62) when ground plane holes are used to increase inductance and decrease layout size [108]. An LR-biased Toggle flip-flop is shown in Figure 3.18(b) with the inductance extraction model in Figure 3.19 in Section 3.3.4.1.

2.6.2 ERSFQ

The obvious method to eliminate static power dissipation is to get rid of the resistive bias network entirely.

One solution, called energy-efficient RSFQ (ERSFQ) [45], [55], replaces every bias resistor with a shunted ($\beta_C \leq 1$) Josephson junction, as illustrated in Figure 2.56. The critical current of each bias junction J_{Bi} is set to the value of the required bias current I_{Bi} . The overdamped junction then acts as a very effective current limiter [55].

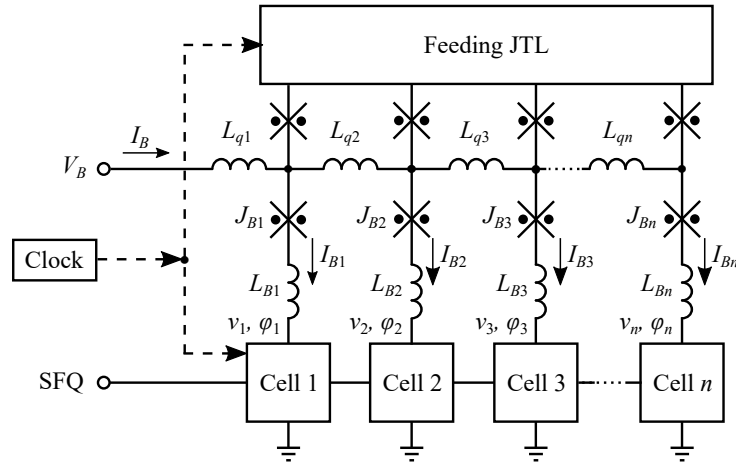


Figure 2.56: Exploitation of the current limiting properties of the Josephson junction to achieve desired bias current distribution in an ERSFQ circuit.

When a total bias current I_B equal to the sum of all bias junction critical currents is injected into the bias network, each bias junction is forced to transmit its critical current and the circuit is biased as intended. There is then zero voltage drop across the bias junctions, so that static power dissipation $P_S = 0$.

However, when biased gates are active (and switching), the voltage at the bias injection points (v_1, v_2 , etc.) is non-zero at $v_i = \Phi_0 f$, where i is the cell number and f is the switching frequency of that cell. For simplicity here, it is assumed that all cells switch at the clock frequency $f = f_{\text{clk}}$. To maintain the correct region of operation of the current-limiting junctions, the voltage on the bias line must at least be equal to that of the injection points. This is achieved by connecting a *feeding Josephson transmission line* (FJTL) to the bias line, as illustrated in Figure 2.56. If the FJTL is clocked at the same frequency f as the main circuit, the voltage on the bias line is maintained at $v \approx \Phi_0 f_{\text{clk}}$. When a cell is inactive the voltage difference between the bias line and the cell's bias current injection point is absorbed through continuous switching of the bias junction.

In practice, ERSFQ circuit margins are better if the FJTL is overpumped (clocked faster than the main circuit) [109]. It is also shown that the operating margins extend into the region where the bias limiting junctions switch continuously and a voltage of more than $100 \mu\text{V}$ is maintained on the bias line. In this case, the static power dissipation is not zero, but still significantly lower than that of standard RSFQ. A detailed analysis of the role and behaviour of the FJTL is shown in [71]

Even though average current transmitted by a limiting junction equals its critical current, the instantaneous current through the junction changes by $\Delta I_{Bi} = \Phi_0 / L_{Bi}$. This instantaneous change in current must be lower than the critical current margin of a cell to ensure operation, so that a high inductance L_{Bi} is required. If ΔI_{Bi} is limited to 5%, then for a bias current of around $200 \mu\text{A}$, the bias inductance must exceed 200 pH . For circuits with $100 \mu\text{A}$ bias current, a relatively enormous 400 pH bias inductance is required.

It is difficult to achieve high inductance in compact integrated circuit layouts, and my contribution to ERSFQ has been the analysis of coiled and high kinetic inductance layer

structures used to increase this bias inductance as much as possible. These contributions have mainly been through private communication, and results cannot be shared here, with the exception of a rendered inductance extraction model shown in Figure 2.57. The model, for an ERSFQ circuit fabricated with the MITLL SFQ4ee process, shows how coils are used under the main ground plane to increase the inductance L_B of bias structures. More recently, high kinetic inductance layers of NbN have been used to increase inductance [110], but calculation with Inductex is no less important.




Figure 2.57: InductEx inductance extraction model for bias section of an ERSFQ circuit with four coils underneath the main ground plane, which has been rendered transparent. The circuit above the ground plane is omitted for clarity

2.6.3 eSFQ

The size of the FJTL layouts and the large bias inductors – as well as small but non-zero static power dissipation – of ERSFQ circuits leave room for improvement. One contender is eSFQ, a circuit technology that performs synchronous phase compensation [55] and obviates the need for the FJTL or large bias inductors.

With eSFQ, the DMP inherent to every clocked RSFQ cell, which undergoes a 2π phase shift during every clock cycle, is connected to the bias line. This is illustrated in Figure 2.58.

Most RSFQ cells can be converted to eSFQ. The RSFQ DFF, of which the circuit schematic is shown in Figure 2.59(a), can be converted to eSFQ by moving the bias injection point to just above the DMP at the clock input as shown in Figure 2.59(b). With every bias injection point subjected to a 2π phase increase for every clock cycle, no phase difference builds up between bias injection points, and consequently there is no unwanted redistribution of bias currents.

Naturally, eSFQ produces its own set of peculiar difficulties. The standard RSFQ T flip-flop (TFF) [47] is not clocked and thus does not have a suitable DMP to serve as the bias injection point. For such a cell, a supply-free version [53] is required that transmits an SFQ pulse ballistically if it is sandwiched between properly biased eSFQ cells [111]. Similar supply-free ballistic transmission cells are required for unlocked pulse transmission cells such as pulse splitters and mergers [55].

While still studying towards his Masters degree, my student Dr Mark Volkmann was invited to visit Hypres in 2011 under the tutelage of Dr Oleg Mukhanov. Mark spent a few months with Hypres as a visiting researcher, where under Dr Mukhanov's supervision and with the help of a talented team at Hypres he designed the first practical eSFQ circuits.

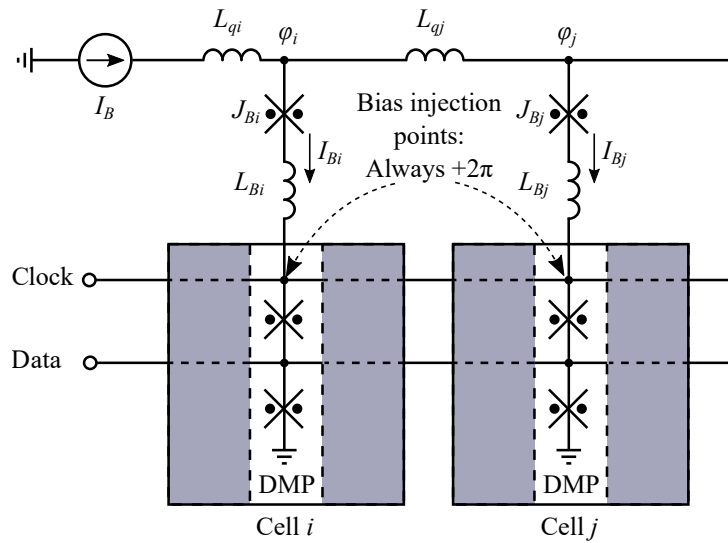


Figure 2.58: The eSFQ biasing principle. When all bias terminals are connected to the clock net, each bias injection point experiences the same $+2\pi$ phase shift during each clock period.

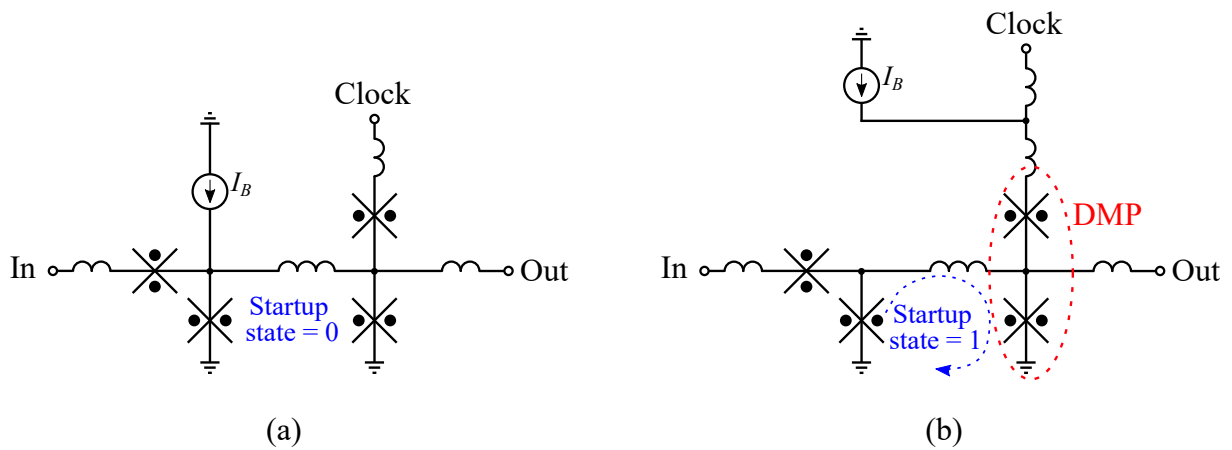


Figure 2.59: Circuit schematic of (a) an RSFQ D flip-flop and (b) the eSFQ conversion of the D flip-flop.

Mark designed an eSFQ shift register cell, called the eSR (for eSFQ shift register) and improved the design with a magnetic flux bias that resulted in the magnetic flux-biased eSR (MeSR) cell [112]. Mark also designed a deserialiser based on the shift-and-dump architecture [113]. All of these circuits were tested successfully at low and high frequency [114], but I will focus on the eSR and MeSR as my contribution to the analysis of these cells led to vast improvements in inductance calculation speed that is detailed in Section 3.3.6.

2.6.3.1 eSFQ shift register

The circuit schematic of an eSFQ shift register (eSR) is shown in Figure 2.60(a). Unlike a normal RSFQ shift register cell, the eSR is biased through the clock DMP. After bias current ramp-up at switch-on, junction J_2 is bias and J_1 is not. This corresponds to a logical “1” stored in the cell. After the first clock, the cell is cleared and current is redistributed to J_1 . This corresponds to a logical “0” in the cell. A clock pulse at C_{in} always switches J_5 , one of the two junctions J_2 and J_4 that comprise the DMP, and J_3 to maintain phase balance.

Simulation results are shown in Figure 2.60(b). It was found that J_3 has low margins, due to bias current distributing between both J_2 and J_1 when the eSR is in the set state. This can be improved by adding a corrective bias through magnetic coupling [112].

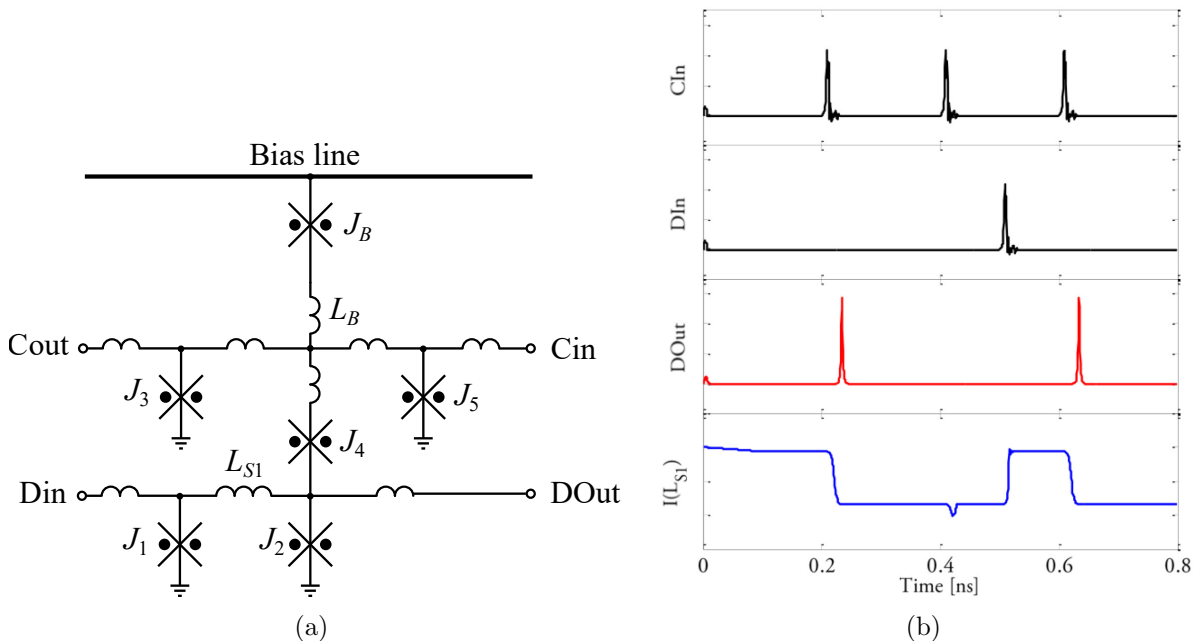


Figure 2.60: (a) Circuit schematic and (b) simulated response of an eSFQ shift register cell – the eSR.

When the eSR is altered to include a magnetic flux bias line that is coupled through its inductance L_F to the storage inductor L_{S1} , a magnetically introduced corrective flux bias can be applied to counteract leakage current from J_2 to J_1 and significantly improve the margins of the circuit. The circuit schematic of this magnetic flux biased eSR, the MeSR, is shown in Figure 2.61(a).

The flux bias has the added advantage that, as it is ramped up, it switches junction J_2 to set the MeSR to the “0” state to match conventional RSFQ shift register behaviour. The simulated response is shown in Figure 2.61(b).

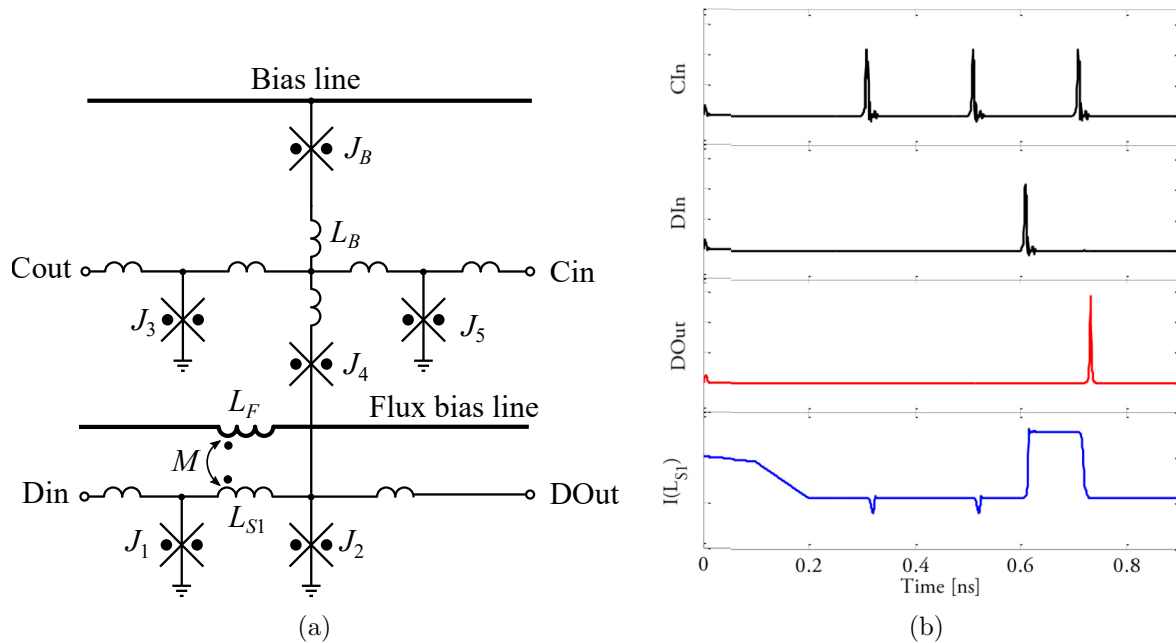


Figure 2.61: (a) Circuit schematic and (b) simulated response of an eSFQ shift register cell with magnetic flux bias – the MeSR.

The eSFQ shift registers were manufactured, as shown in Figure 2.62, and tested successfully [111].

Inductance extraction for the eSR and MeSR designs used InductEx as it existed at the start of 2011. As is explained in Chapter 3, the tool was still experimental and limited. Mark and I could only extract inductance of subsections of the cells due to the presence of a skyplane in M3 that resulted in a high segment count. One such a section, that for the bias limiting junction and the inductor L_B , is depicted in Figure 2.63. Full-cell inductance analysis, as depicted in Figure 2.64, was not yet possible. The inconvenience of having to break a layout down into smaller parts for analysis spurred me to re-engineer InductEx for support of larger models. Very soon after, circumstances would demand that very capability.

2.6.3.2 An eSFQ T flip-flop

Mark returned to Hypres for another research visit in 2012 to work with Dr Igor Vernik on the development of an eSFQ T flip-flop (TFF) proposed by Dr Oleg Mukhanov, as well as new eSFQ shift registers. The TFF schematic and simulated response are shown in [115].

These newer eSFQ cells demanded inductance extraction for layouts where several inductors *were laid out over holes in the ground plane* to increase inductance. To account for all the ground and skyplane currents, a layout simply could not be reduced to several small, separate extractions. Full-cell inductance extraction was absolutely necessary. By mid-2012, I had improved the model building and meshing capabilities of InductEx to handle the entire TFF layout and produce an extraction model that would fit inside the memory limitations of the computers that we had then.

The eSFQ TFF layout was fabricated and tested successfully [115], and is shown in Figure 2.65. The inductance extraction model that I constructed with InductEx to match the eSFQ TFF layout is shown in Figure 2.66.

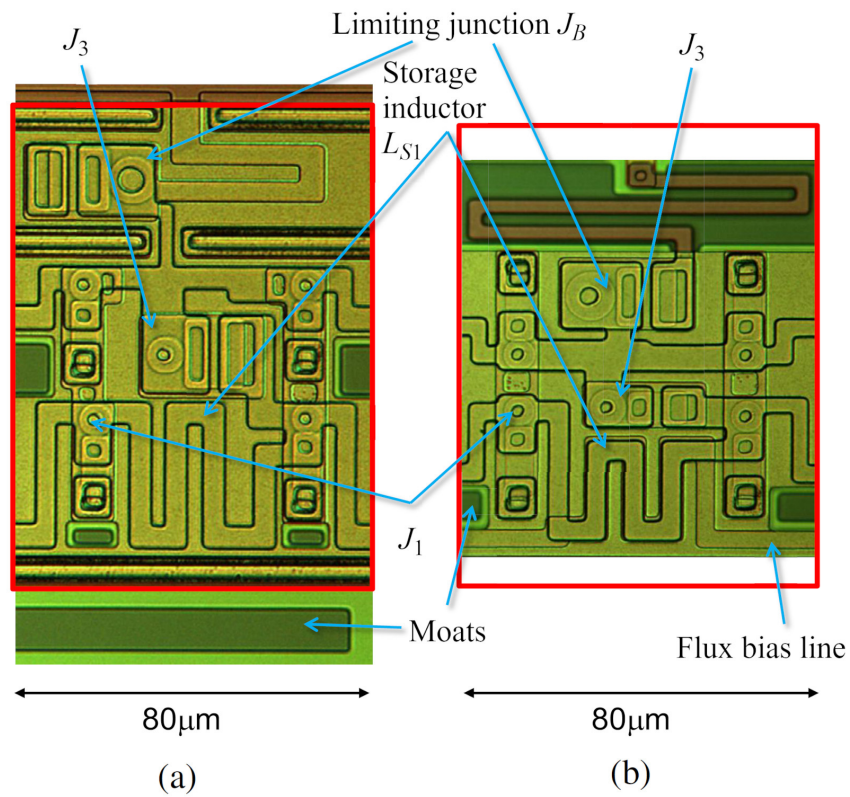


Figure 2.62: Microphotographs of (a) the eSR cell and (b) the MeSR cell [111] fabricated with the Hypres 4.5 kA cm^{-2} process.




Figure 2.63: InductEx model of the bias junction and inductor for the ESR, matching the maximum size investigated in 2011. The skyplane M3 has been omitted for clarity.




Figure 2.64: Modern InductEx model, with cuboid segments, of the eSR shown in Figure 2.62(a). The skyplane M3 has been omitted for clarity.

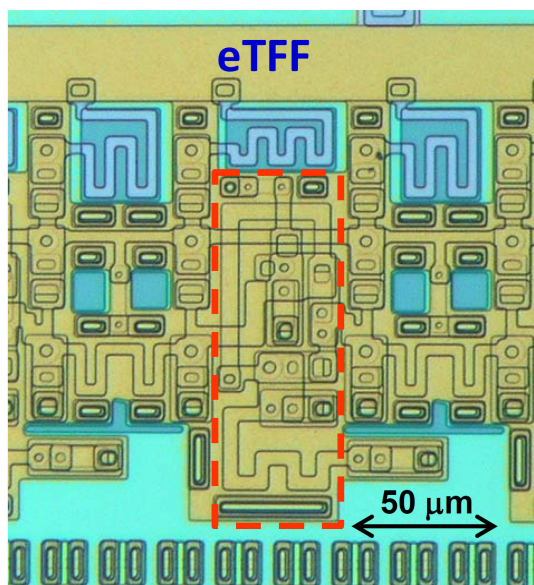


Figure 2.65: Microphotograph of an eSFQ TFF [115] fabricated with the Hypres 4.5 kA cm^{-2} process.




Figure 2.66: InductEx model for inductance extraction of the eSFQ TFF shown in Figure 2.65. The full model is shown on the left, with the skyplane in M3 omitted on the right for clarity.

What makes the model remarkable is that, by mid-2012, with 17 ports, 21 inductors and more than 50 000 segments – and especially due to the segment-heavy skyplane and abundance of sky-to-ground curtains through vias and metal strips in M1 and M2 – it represented the most complex circuit ever to be extracted as a single model. *Extraction took almost a day*, and prompted me to invest research effort into speeding up the calculation tools, as discussed in Section 3.3.6.

2.6.4 Adiabatic Quantum Flux Parametron logic

It was predicted that the bit energy of a logic gate can be reduced to the order of $k_B T$, where k_B is the Boltzmann constant and T is temperature, by adiabatically varying the shape of the potential of the gate from a single to a double well [116], [117]. The parametric quantron [118] and the quantum flux parametron (QFP) [119], [120] exploited this technique, with both using ac power sources for operation.

Due to high gain, high speed and robust operation, Prof. Nobuyuki Yoshikawa and his students at Yokohama National University in Japan investigated the QFP for its energy consumption when operated in the adiabatic mode [46]. Thus an exciting new logic family – adiabatic quantum flux parametron (AQFP) – was created and shown to be around two orders of magnitude more energy efficient than other low-energy SFQ logic families [121] when operated at around 10 GHz with unshunted junctions.

An AQFP circuit operates around a two-junction SQUID as shown in Figure 2.67(a). The junctions J_1 and J_2 , which can be resistively shunted but are left resistively unshunted for best energy efficiency [121], are shunted by a centre inductor L_q . An ac excitation (clock) through L_x , combined with a bipolar input signal current through L_{in} , stores a fluxon in either the left or right half of the SQUID bisected by L_q and sets the output current direction through L_{out} .

Dr Naoki Takeuchi, then a graduate student in the laboratory of Prof. Nobuyuki Yoshikawa at Yokohama National University, designed and tested the first AQFP circuits [46]. The first layouts employed straight microstrip lines for inductor layout and magnetic coupling control, but these were area inefficient. Smaller layouts that included any asymmetry would not function, or have very low margins. After meeting each other

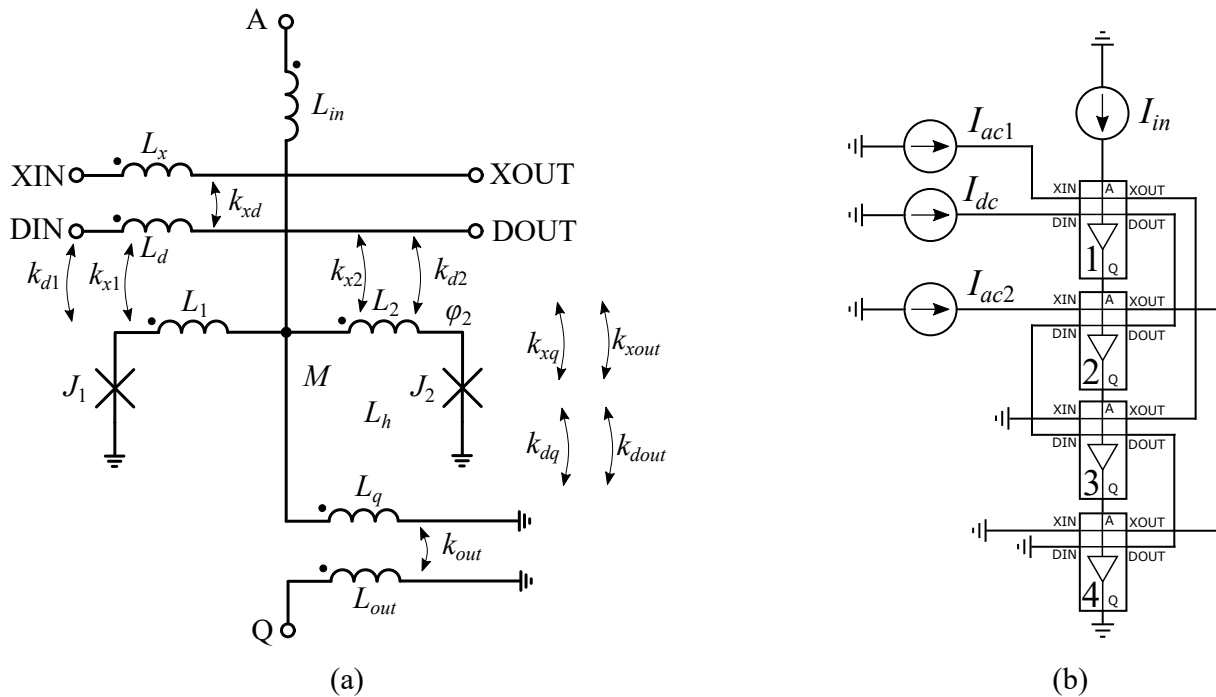


Figure 2.67: AQFP buffer cell (a) schematic and (b) simulation test circuit.

at a conference, I helped Dr Takeuchi build inductance extraction models with InductEx that could find all the parasitic (unwanted) coupling – such as that between L_x and L_q to name but one. Electrical simulations with the extracted parasitic inductances duplicated observed failures, and it was clear that layouts had to be extracted and improved with care. I gave advice on smaller symmetrical layouts. The resulting layout of an AQFP buffer cell, with symmetrical coils to cancel out parasitic coupling, is shown in Figure 2.68 with InductEx port labels for extraction included. Subsequent AQFP designs still rely on InductEx to calculate the self and mutual inductances, and to verify that remaining parasitic coupling is small enough not to affect circuit operating margins [122].

My contribution to the improvement of AQFP also included the calculation and verification of impedance of the microwave ac clock lines to minimise reflection [123]. Most recently, I applied InductEx and compact model extraction to find the coupling from fluxons in moats to AQFP cells – specifically to understand the role of trapped flux in the reduction of AQFP operating margins [124].

I have found that the inductor L_q in an AQFP gate is quite sensitive to coupling from a fluxon in a moat near that inductor, especially if the coupling from the moat is asymmetric to the standard dual-coil layout L_q and L_{out} . Even a single fluxon in one of the moats FBL or FBR in Figure 2.68 reduces operating margins significantly, while three fluxons in one of these moats renders the circuit inoperable (see the simulated results in Figure 2.69). A rendering of current density calculated by InductEx when a fluxon is trapped in moat FBL is shown in Figure 2.70, where it is clear that coupling to the dual-coil structure (which contains L_q and L_{out}) results in asymmetric current flow that unbalances the two-junction SQUID.

Engineering more robust AQFP layouts is an ongoing research effort where I continue to make a contribution to the field.

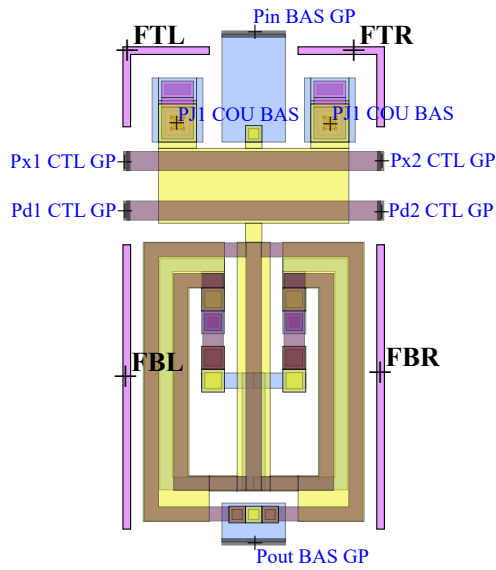


Figure 2.68: AQFP buffer cell layout drawing with InductEx ports for inductance extraction and flux trapping analysis.

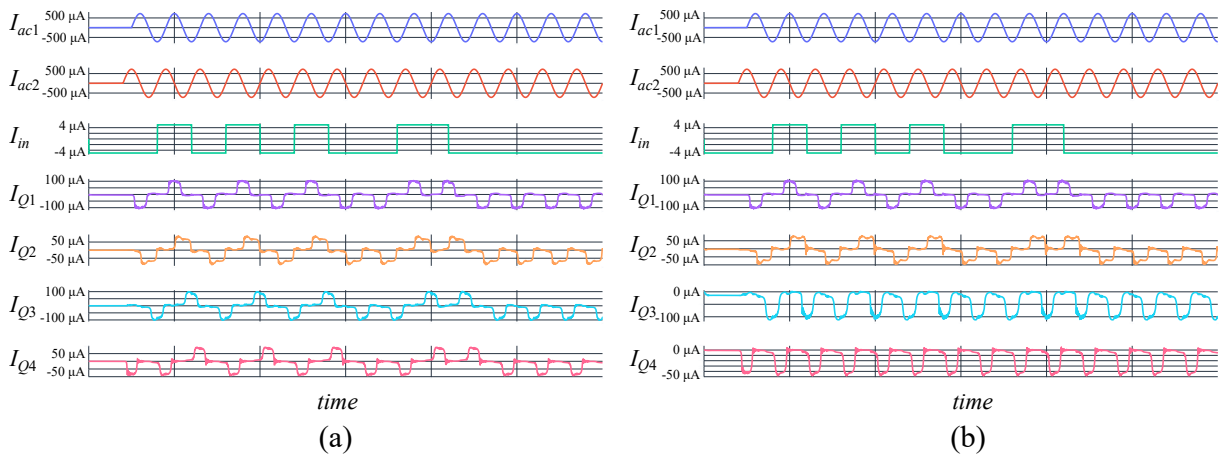


Figure 2.69: JoSIM simulation result of AQFP buffer chain from Figure 2.67(b) with (a) no fluxons in any moat – showing correct operation – and (b) three fluxons in moat FBL (see Figure 2.68) of buffer 2 – showing failure to switch buffers 3 and 4.




Figure 2.70: InductEx model of the AQFP buffer cell with current distribution generated by a positively oriented fluxon in moat FBL.

2.7 Summary

Although the bulk of my research career has been focused on the development of powerful software for design automation and verification in support of the superconductor integrated circuit design community, as will be presented in the following chapters, I have made contributions to circuit theory, as well as the design and verification of basic cells and cell libraries.

Through all of this work, the golden thread that strings everything together has been *inductance*.

Chapter 3

Inductance calculation

3.1 Theory

3.1.1 Inductance of a conductor

Inductance is a property of an electrical conductor that causes it to resist a change in electrical current flowing through it. Current in the conductor creates a magnetic field, of which at least a part cuts through the surface of a closed circuit formed by the conductor and its current return path. A time-varying current results in a time-varying magnetic field which, from Faraday's law of induction, induces an electromotive force in the circuit to oppose the change of current. If the current is increasing, the resulting voltage over the conductor is positive at the node where current enters the conductor, leading to (2.30) – the equation that relates voltage over a conductor to current change through the conductor.

From (2.30), inductance of a conductor is the constant L that relates voltage to the change in current. The unit of inductance is henry, named after the American scientist Joseph Henry (1797–1878) who discovered electromagnetic induction around the same time as Michael Faraday in the 1830's.

3.1.2 Inductance in a superconductor circuit

For an excellent and concise discussion on the background to self- and mutual inductance in superconductor circuits, see [125]. Briefly summarised, the self-inductance L of a superconducting circuit loop is defined through the total energy of the loop U and the current I in the loop as

$$U = U_m + U_k = \frac{1}{2}L_m I^2 + \frac{1}{2}L_k I^2 = \frac{1}{2}L I^2. \quad (3.1)$$

Here, U_m is the energy of the magnetic field around (and inside) the conductor and U_k is the kinetic energy of the current carriers inside the superconductor. The total inductance L is the sum of the magnetic inductance L_m and the kinetic inductance L_k . The magnetic inductance derives from the geometry of the loop (and thus the field cutting the loop surface) and is also called the geometrical inductance. The kinetic inductance is purely a function of material properties of the superconductor and the cross-sectional area and length of the superconductor, so that it *behaves like an imaginary resistance*.

The kinetic inductance of a section of superconductor with length l , constant cross-sectional area A , and uniform current distribution is

$$L_k = \frac{\mu_0 \lambda l}{A}, \quad (3.2)$$

where μ_0 is the permeability of free space ($\mu_0 = 4\pi 10^{-7}$ H/m) and λ is the London penetration depth [126]. The penetration depth is defined as

$$\lambda = \frac{m}{2n_s e^2}, \quad (3.3)$$

with m , e and n_s the electron mass, electron charge and Cooper pair number density respectively.

3.2 Background

Of the primary passive electrical components in an integrated circuit, namely resistance, capacitance and inductance, inductance has been the most neglected component. This is mostly due to its negligible effect inside semiconductor logic gates, and on the small effects of inductance on gate-to-gate interconnects when compared to the significant influence of resistance and distributed capacitance of gate-to-gate interconnects on circuit power consumption and switching speed of semiconductor digital ICs.

Even though it has long been neglected in digital semiconductor integrated circuits, inductance – which is considered parasitic – influences on-chip interconnect delay [127] and timing analysis [128], [129], cross-talk noise of on-chip interconnects [130], interconnect delay of signal nets in the presence of power supply grid noise [131] and noise analysis [132].

Recently, more attention has been given to inductors on semiconductor ICs.

For miniaturised power converters, microfabricated racetrack inductors of a few nanohenry [133] to a few microhenry [134] have been demonstrated.

Large-size on-chip inductors are also now used for transformers in microelectromechanical systems (MEMS) [135] where inductance values are more than 100 μ H.

Inductance of on-chip power distribution networks (PDNs) has also been considered [136]–[138].

Solenoid inductors in three-dimensional radio frequency chips, patterned out of thin films on the chips and through-silicon via (TSV) contacts [139], find application in energy storage and RF filtering.

However, such single inductors are still described and designed with analytical equations.

3.2.1 Superconductor integrated circuits

In semiconductor integrated circuits, in the words of Ismail [132]:

“... the industry applies a three-step design process for integrated circuits when handling inductance. First, employ design methodologies and techniques to reduce the inductance effects in the design. Second, use the well-developed RC-based design tools to optimize and verify the circuit. Third, wish nothing will go wrong.”

Such a strategy would not work for superconductor integrated circuits, where the design process must explicitly account for inductance. The inductance of gate interconnects and intra-gate connections determine circuit functionality by regulating fluxon

storage or transmission [47] and influence bias current distribution and circuit operating margins. Mutual inductance also affects data transmission between magnetically coupled (galvanically isolated) circuits [61], [140]–[144], and is of great significance in ac-biased circuits such as Adiabatic Quantum Flux Parametron (AQFP) logic. The range of inductance that yields successful circuit operation, such as fluxon storage as opposed to transmission, can be very narrow, and accurate calculation is therefore very important.

3.2.2 Parameter extraction

The three passive components of importance in an electrical circuit are resistors, capacitors and inductors.

3.2.2.1 Resistance extraction

Resistance is the simplest of the passive components to extract. Thin-film resistance extraction in integrated circuits requires only geometric (length, width, height) and material properties (bulk or sheet resistivity) of the conductor [145] to allow extraction. The simplest method uses heuristic extraction that finds a single resistance directly from layout artwork [146] for a given sheet resistivity. More accurate numerical extraction methods that also allow multi-terminal resistance calculation range from a nonuniform rectangular grid method to solve the node admittance matrix of integrated resistors [147], to a finite element method (FEM) solver that partitions layout artwork into contact and body regions and solves the admittance matrix when boundary conditions are enforced (external currents are applied at contact terminal boundaries and electric potential is constant at a contact terminal boundary) [148] and even to a boundary element method that allows arbitrary shapes and requires simpler meshes than equivalent FEM and finite difference method (FDM) calculators [149].

3.2.2.2 Capacitance extraction

Capacitance extraction is more complex and requires the addition of geometric properties (dimensions and distance or spacing) of the immediate neighbouring conductors and the dielectric constant of the space between conductors [145]. Capacitive effects are short-ranged [150], but complex three-dimensional extraction with boundary-element methods leads to dense linear systems that require accelerated iterative techniques [151] or sparsification [152] to become tractable or solve in reasonable time. For full-chip capacitance extraction, pattern matching and interpolation from look-up tables are generally used [150]. Irrespective of calculation times, modelling layouts for resistance and capacitance extraction is straight-forward.

3.2.2.3 Inductance extraction

Inductance is the most complex of the three passive components to extract. In the absence of magnetic materials, mutual and self-inductance are functions of system geometry [153]. Inductance describes magnetic flux generated by current flowing in a loop, with an additional kinetic term for superconductors [154] arising from the kinetic energy of the superelectrons. Inductance extraction in an integrated circuit thus requires modelling of a conductor and the current return path, which is frequency-dependent for normal conductors [155]. However, the current return path is not always known *a priori*, which

complicates modelling and calculations. In semiconductor integrated circuits, inductance paths can be minimised through careful design [156], thereby simplifying calculation of the loop inductance. At high frequencies, current return paths can be assumed to be the nearest power or ground lines [157]–[159], which often reduces models to two-dimensional complexity [158]. The method of return-limited inductances [150], [160] formalises this approach to create sparse inductance networks for full-chip extraction, although it has been shown to underestimate inductance [161] as feature size scales down.

Numerical calculation of inductance when return paths are not known *a priori* is made possible by partial inductances [162], for which flux linkage, and thus partial inductance, is calculated through loops between the conductor segments and infinity. Total loop inductance is then the sum of partial self and mutual inductances for all conductor segments that comprise the loop [159], [162], [163]. However, magnetic coupling acts over a long range and leads to a dense system [145], [150]. Furthermore, partial inductances are defined by loop areas to infinity, which makes partial inductance matrices dense and their solution computationally expensive. Even with calculation cost reduction techniques, such as the multipole-accelerated method used by FastHenry [164], matrix sparsification through the use of equipotential shells [165] and the K-method [166] and its extensions [167], [168], full-chip inductance extraction with such methods remains intractable. Full-chip extraction therefore requires simplified models, such as return-limited inductance [150], [160], two-dimensional transmission line methods [157], closed-form analytical models [163] or analytical formulae benchmarked against FastHenry solutions for a set of geometries [169].

One inductance modelling technique for on-chip inductances [170] used analytical formulae for self and coupling inductance of straight wire combinations with wire cross section down to around 1 μm . Although faster than FastHenry, the method has an average error of 10 % to FastHenry.

In superconductor integrated circuits, and more specifically the single flux quantum (SFQ) logic families to which most superconductor digital circuits belong today [47], gate interconnects are short and purely inductive unless ballistic SFQ pulse transfer is used [94], [171], [172]. For ballistic transfer, passive transmission line (PTL) interconnects are impedance matched, and accurate analytical methods exist to calculate the characteristic impedance of microstrip line [173] PTLs.

Superconductor ICs manufactured in the low critical temperature (low- T_c) processes from Seeqc (formerly operated by Hypres [22]), FLUXONICS at Leibniz-IPHT [21], AIST [25], [26], [174], MIT Lincoln Laboratory SFQ4ee and SFQ5ee [9], SFQ6ee [31] and similar foundries account for most superconductor ICs produced today. In these ICs, all ground plane and wiring layers are superconductive, with the exception of resistive layers in which Josephson junction damping resistors and current bias resistors are formed. At frequencies far below the energy gap frequency of a superconductor (approximately 700 GHz for niobium), the penetration depth (λ) is frequency independent. This means that current flows near the surface with the same distribution at all frequencies and that surface impedance is almost completely inductive [3], so that current return paths attempt to minimize loop inductance (and thus loop area). Current return paths are therefore in ground or shield layers closest to conductors, so that all inductance of circuits such as logic gates can normally be calculated by considering only the layout structures inside the circuit boundary. Interconnects can be interrupted at circuit boundaries for calculation, because return current in real environments will enter through the ground or shield layers directly below or above these interconnects. Exceptions exist, such as

monolayer high critical temperature (high-T_c) circuits which have no ground plane [20], but partial inductance extraction [162] (leading to partial element equivalent circuit or PEEC models [175]) can be shown to be sufficient to handle these.

In general, therefore, SC ICs do not require full-chip inductance extraction; only per-gate or per-circuit extraction. Provided that circuit size is reasonable, solutions with numerical methods could thus be made tractable. The capability to calculate the multi-terminal networks of self and mutual inductance in such circuits accurately is also required, which makes the use of numerical methods crucial. This is the purpose for which InductEx was developed [176].

3.2.3 Known inductance extraction tools

3.2.3.1 Normal conductors

For inductance extraction of normal conductors, FastHenry [164] was the gold standard when I started my research career. Several papers presented specific algorithms that were claimed to be from tens to hundreds of times faster [177], [178] than FastHenry, but I could only find one *published software tool* that claimed superior performance over FastHenry and was available for download: Inductwise [167]. Inductwise used a reluctance method and the authors showed results within 1 % of FastHenry below 10 GHz but worsening rapidly above that frequency. Speedup of 10 to 26 times was claimed. Unfortunately, the download link for Inductwise has been dead for many years.

3.2.3.2 Superconductors

Tools based on numerical methods to calculate inductance in superconductor circuits have been available for a long time.

INDEX [179] was presented in the early 1990's. It divided a complex two-dimensional net into rectangles for which individual inductances were calculated from analytical techniques derived from numerical simulations and curve-fitting. INDEX could then extract netlists from the results. However, I could not find any example sets for, or any evidence of widespread use of INDEX.

The superconductive integrated circuit design community, and especially designers of RSFQ circuits, have used Lmeter [1] with great success from the 1990's. Lmeter (sometimes written L-meter) is a multi-terminal self and mutual inductance network extraction programme, is ideal for per-gate or per-circuit inductance extraction, and used to be considered fast and reliable for conventional RSFQ logic [47]. Lmeter as a gate-level inductance extraction tool has found application in circuit designs that range from SFQ cells [180] to analogue-to-digital converters [181], logic blocks [182] and even full SFQ processors [183], and has also been integrated into the Cadence design environment [184] through a scripting interface with the SKILL language. However, some extraction models with very narrow inductive lines, ground plane features or the absence of a ground plane are either intractable with Lmeter, or solve with unacceptable accuracy.

The other popular and powerful superconductive inductance extraction utility, 3D-MLSI [185]–[188], handles multilayer structures with narrow lines and ground plane features with very good accuracy, but due to difficulty with modelling vias and building logic cell type layout models, it mostly found application in SQUID analysis [189]–[191] and for SQUID-based gradiometers [19].

The limitations of Lmeter and 3D-MLSI when complex layouts (with isolated ground planes, mutual coupling over ground plane holes, multiple shield layers, etc.) are extracted have caused some researchers to turn to FastHenry [164] with superconductivity support [192] (where the London equations are included in the imaginary part of the complex conductivity) to solve specific inductance extraction problems [142], [193]. However, direct use of FastHenry is cumbersome when models are created manually, solutions have large inaccuracy when current flow in short and wide conductors is not adequately modelled, and multi-terminal networks cannot be solved readily.

FastHenry uses rectangular segments – cuboids – for geometry discretisation, which makes segmentation less elegant and more complicated than the triangular meshes supported by Lmeter and 3D-MLSI. An orthogonal segmentation method was proposed earlier [194] and implemented in the first version of the inductance calculation utility InductEx [195], as is detailed later in this Chapter. An improved implementation of InductEx, capable of calculating multi-terminal self and mutual inductance extraction for complex SFQ layouts, was presented and demonstrated to be sufficiently accurate for advanced SFQ cell design [176]. InductEx, using FastHenry as the magnetoquasistatic 3D field solver, was made more efficient by changing FastHenry to calculate port currents instead of an inductance matrix, and letting InductEx calculate the multi-terminal inductance network through the solution of an overdetermined system of linear equations formed from branch currents and circuit meshes.

3.3 InductEx: Three-dimensional inductance calculation

This section only focuses on the contributions of my research group and I.

3.3.1 Early contributions

My contributions to inductance extraction started purely because I had to design the layout of inductors in superconductor circuit layouts and immediately ran into issues.




Figure 3.1: A straight line microstrip.

For a straight microstrip line, which comprises a ground plane and a conductor separated by some vertical distance as depicted in Figure 3.1, the analytical solution to inductance was derived by Chang [173]:

$$L = \frac{\mu_0}{WK} \left\{ h + \lambda_1 \left[\coth \left(\frac{t_1}{\lambda_1} \right) + \frac{2p^{1/2}}{r_b} \operatorname{csch} \left(\frac{t_1}{\lambda_1} \right) \right] + \lambda_2 \coth \left(\frac{t_2}{\lambda_2} \right) \right\}, \quad (3.4)$$

with the parameters:

$$\beta = 1 + \frac{t_1}{h}, \quad (3.5)$$

$$p = 2\beta^2 - 1 + [(2\beta^2 - 1)^2 - 1]^{1/2}, \quad (3.6)$$

$$\eta = p^{1/2} \left\{ \frac{\pi W}{2h} + \frac{p+1}{2p^{1/2}} \left[1 + \ln \left(\frac{4}{p-1} \right) \right] - 2 \tanh^{-1} p^{1/2} \right\}, \quad (3.7)$$

$$\Delta = \max(\eta, p), \quad (3.8)$$

$$r_{bo} = \eta + \frac{p+1}{2} \ln \Delta, \quad (3.9)$$

$$r_b = r_{bo} - [(r_{bo} - 1)(r_{bo} - p)]^{1/2} + (p+1) \tanh^{-1} \left(\frac{r_{bo} - p}{r_{bo} - 1} \right)^{1/2} - 2p^{1/2} \tanh^{-1} \left(\frac{r_{bo} - p}{p(r_{bo} - 1)} \right)^{1/2} + \frac{\pi W}{2h} p^{1/2}, \text{ for } 5 > \frac{W}{h} \gtrsim 1, \quad (3.10)$$

$$r_b = r_{bo}, \text{ for } \frac{W}{h} \geq 5, \quad (3.11)$$

$$r_a = \exp \left[-1 - \frac{\pi W}{2h} - \frac{p+1}{p^{1/2}} \tanh^{-1}(p^{-1/2}) - \ln \left(\frac{p-1}{4p} \right) \right] \quad (3.12)$$

and

$$K = \frac{h}{W} \frac{2}{\pi} \ln \frac{2r_b}{r_a}. \quad (3.13)$$

Here K is the fringe field factor that is a function of W , h and t_1 , t_1 is the thickness of the microstrip line, t_2 is the thickness of the ground plane, h is the distance between the top of the ground plane and the bottom of the microstrip line, and W is the width of the microstrip line as illustrated in Figure 3.2.




Figure 3.2: Dimensions of a straight line microstrip.

In the absence of numerical tools for inductance calculation, some designers calculated the per-square inductance of a microstrip line with (3.4) and then just counted squares over the length of an inductor, where for a straight-line inductor the number of squares equal the length divided by the width. Difficulties arise when an effective path length has to be obtained for a layout with a corner such as in Figure 3.3, when microstrip width steps up or down, or when one microstrip connects to another through a via.




Figure 3.3: Traditional approximation of the effective path length for inductance estimation around a corner in a thin-film inductor.

For the effective path length, an effective corner inductance would be calculated from the per-square inductance L_{sq} of a line with given dimensions (thickness, distance to ground plane, width) and added to an inductive path. The value of the corner inductance used to be a matter of debate, and ranged from $L_{sq}/\sqrt{2}$ to $0.56L_{sq}$ between research groups. One numerical method put it close to $0.5L_{sq}$ [196]. In addition to these inconsistencies, any corner with different arm widths would result in a different value for the effective inductance.

It was obvious that a proper numerical inductance calculator was necessary to validate a real circuit layout, but I could not find a reliable inductance calculator that was suitable to the geometry of my layouts. Ironically, if I had known then how to install and use Paul Bunyk's Lmeter [1], I would probably not have given a second thought to an inductance calculator, and would have gone on to live a very different life.

But I knew nothing at all about Lmeter, only that FastHenry [164] was a well-known inductance calculation tool and that it had by then already been adapted for superconductivity by Stephen Whiteley. Pressed for time to finish the layouts for my PhD, I set out to make FastHenry work with my layouts.

The immediate difficulty with choosing FastHenry as an inductance calculator for compact integrated circuit layout is that it was never designed to handle such structures. FastHenry was developed to find the inductance of printed circuit board (PCB) tracks and chip packaging leads, all of which can be modelled as long, slender conductors with uniaxial current flow. Half of the leads of a typical chip package to which FastHenry is suited for inductance calculation are shown in Figure 3.4(a). For calculation purposes, unity amplitude voltage sources are connected across the terminals of each lead, and FastHenry yields an inductance matrix that describes a set of inductors with mutual inductance as depicted in Figure 3.4(b).

As long as a circuit could be modelled as a set of isolated, coupled inductors, FastHenry could be applied to the extraction of inductance. However, FastHenry is formulated to model uniform current flow along the length of every mesh segment, with segments connected electrically through nodes at the centre of a segment's entry/exit face (see Figure 3.5). Current distribution in a superconducting line is concentrated on the outer surfaces, but penetrates the conductor with decaying amplitude determined by the London penetration depth (λ). FastHenry supports the subdivision of a segment into filaments




Figure 3.4: (a) Typical chip package section for analysis with FastHenry and (b) individual self and mutual inductances resulting from FastHenry solution.

across the width and height, as shown in Figure 3.5. Each filament then carries uniform current, but larger current density in the outer filaments better approximates the distribution in a superconducting line that is thicker or wider than the penetration depth.

The biggest drawback to a typical FastHenry model is that the interconnection of two segments results in a bad approximation of current distribution around corners, through vias or near tee connections with short arms – *the very geometries that heavily populate typical digital logic circuit layouts*. Examples of such structures are depicted in Figure 3.6. An earlier publication by Guan *et al.* [194] demonstrated how interleaving short segments in the three axial directions could handle current flow in each axial direction, although their work was aimed at building lookup tables for inductance structures.




Figure 3.5: FastHenry segment shown with 3 height filaments, 5 width filaments and node for connection.

I developed a similar but more structured technique to building interleaved meshes for the full layout of several structures used in RSFQ integrated circuit layout [197]. Most of these structures included Josephson junction layouts on either end, all were fully parameterised to adapt the calculations for different layout dimensions, and all layouts that my team and I did until 2003 used a concatenation of these structures to determine the inductance values in a logic circuit layout. All structures are meshed with short cuboid segments that are interleaved in the x and y directions for planar structures, with the addition of z -directed interleaved segments for vias between layers. The interleaving is illustrated in Figure 3.7.

3.3.1.1 Meshing

The uniform current distribution in a FastHenry mesh segment is a bad approximation of actual current flow in a superconductor where current is concentrated near the conductor surface and at the outer edges of a conductor over ground, as shown in Figure 3.8. This leads to significant overestimation of inductance. Therefore, just as current distribution

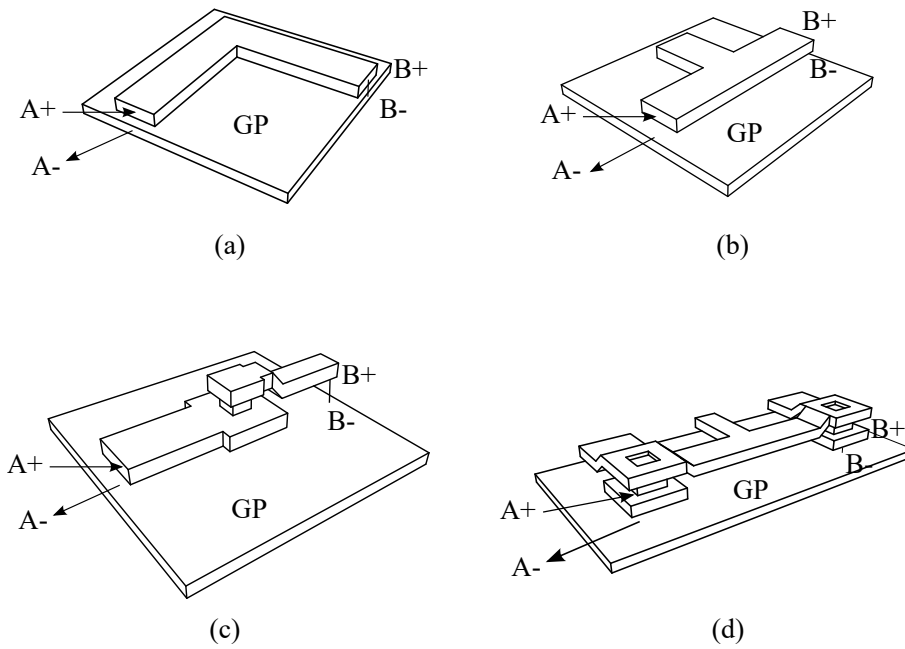


Figure 3.6: Common layout structures for inductance calculation: (a) cornered microstrip, (b) microstrip with tee-in, (c) via-connected microstrips and (d) microstrip connecting two Josephson junctions with shunt resistor connects and optional dc tee-in. The port at B is always shorted between the positive and negative terminals, while a unity amplitude voltage source excites the port at A to determine the inductance of the structure with the current return path in the ground plane.




Figure 3.7: Segmented line (a) with interleaved cuboid segments and (b) graphical representation with segment widths shrunk to one third of their actual values for visualisation purposes.

modelling is improved in the x and y directions of an in-plane conductor with smaller mesh elements, it can be improved in the z direction by slicing thin-film conductors into thinner slivers in the z direction. In FastHenry, the option to declare filaments over the height of a conductor provides that capability.




Figure 3.8: Current distribution in the lowest filaments of a superconducting conductor and the highest filaments in a superconducting ground plane calculated with a cuboid mesh. The line is $4\ \mu\text{m}$ wide, λ is $90\ \text{nm}$ and thickness is irrelevant. The graphs show the current distribution, from top to bottom, when 70 homogeneous $0.1\ \mu\text{m}$ segments, 8 homogeneous $1\ \mu\text{m}$ segments and $1\ \mu\text{m}$ segments with $0.1\ \mu\text{m}$ lambda edges (10 segments in total) are used.

The graphs in Figure 3.9 [176] show how decreasing segment size and increasing number of filaments (which translate to decreased thickness of slivers in the z direction) reduces the calculated inductance of a superconducting microstrip until it approaches an asymptote when all segments have a largest dimension in any axial direction that does not exceed the London penetration depth, λ .

Ideally, then, any structure for inductance extraction would have to be meshed with maximum segment sizes equal to λ to ensure accurate results. However, for typical thin-film niobium integrated circuits, $\lambda \approx 90\ \text{nm}$, which is smaller than the typical layer thickness of $200\ \text{nm}$ to $300\ \text{nm}$ and the typical line width of $1\ \mu\text{m}$ to $10\ \mu\text{m}$. From an engineering perspective, such fine meshing is thus prohibitively expensive for typical circuit layouts and limits the maximum size of extraction models to small logic gates.

Calculation time for typical interleaved meshed structures with FastHenry scales as $O(n^3)$. Just halving the segment size in the in-plane direction x and y could increase segment count by a factor of four, and calculation time by a factor of about sixty. Memory use would also increase, so that large, finely meshed structures would simply overrun the resources of a computer at any useful model size.

My aim was thus to develop meshing and modelling methods that could maintain desired accuracy at much larger segmentation size, and thus to provide extraction tools that can handle layouts that are significantly larger than any other method can reliably handle.

3.3.1.2 Method of images

Superconductor integrated circuit layouts make heavy use of ground planes. In earlier processes, a lower superconducting metal ground plane, usually named M0, would be

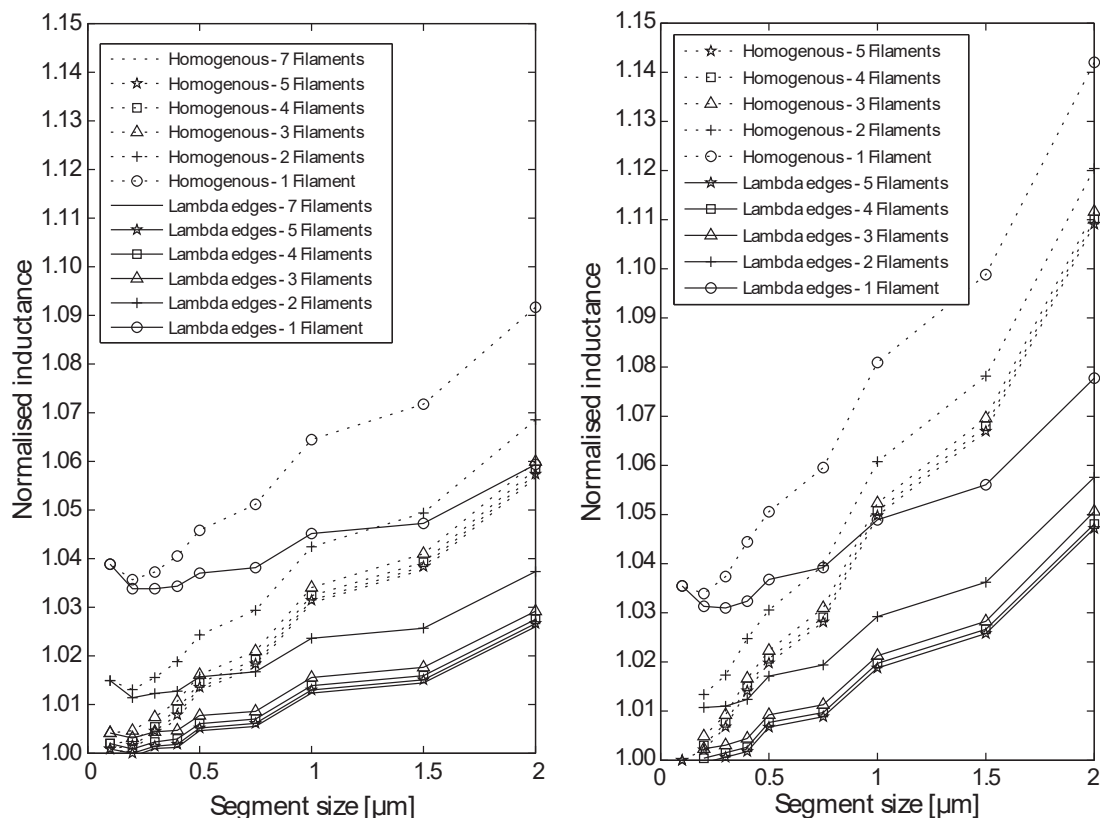


Figure 3.9: Inductance of superconducting structures calculated with cuboid mesh and normalized to the smallest solution for (a) a microstrip line $4\ \mu\text{m}$ wide and $10\ \mu\text{m}$ long and (b) a microstrip line $4\ \mu\text{m}$ wide with a corner and arms extending $5\ \mu\text{m}$ on each side of the bend. Both lines are $0.3\ \mu\text{m}$ thick and separated by $0.35\ \mu\text{m}$ from a $0.1\ \mu\text{m}$ thick ground plane extending $2\ \mu\text{m}$ beyond the line dimensions. London penetration depth for all superconductors is $90\ \text{nm}$. For homogenous models, all segments have the same dimensions. For “lambda edge” models, the width of all segments at the edges of every conductor are fixed to the London penetration depth.

filled underneath a circuit. The Josephson junctions are then patterned between two superconducting layers that form the base and counter electrodes, usually named M1 and M2 respectively.

When I started my work on RSFQ circuit layouts and inductance calculation, the solution of any densely interleaved mesh with FastHenry was excessively expensive. A typical mesh with around 1000 mesh elements would take a few minutes to solve on a personal computer, while meshes with 5000 to 10 000 mesh elements would take hours to solve. Models beyond that complexity were intractable.

Since most of the mesh elements in an early RSFQ layout were devoted to the ground plane, I investigated a “method of images” technique to replace the ground plane with a reflection plane [198] situated at a depth equal to the effective penetration depth below the ground plane surface [196]:

$$\lambda_{eff} = \lambda \coth\left(\frac{d}{\lambda}\right), \quad (3.14)$$

where λ is the bulk London penetration depth and d is the thickness of the superconducting film – here the thickness of the ground plane.




Figure 3.10: Position of the reflection plane at λ_{eff} for a superconducting microstrip over ground.

The method of images was efficient, and allowed us to extract inductance from larger circuit structures than was possible with full-cell ground plane modelling, but it had significant disadvantages:

- Neglect of the kinetic inductance component of the ground plane.
- No support for moats or other holes in the ground plane; or for closely cropped ground planes.
- No support for multiple ground planes, or sky planes used for stripline conductors or cell shielding.

Finally, when I formalised model building in the first version of InductEx, ground planes could be cropped automatically to within a specified distance from every conductor with very minimal degradation of accuracy, so that meshes with ground planes included became more efficient than meshes for the method of images, and the latter was discontinued.

3.3.2 The first InductEx

As our superconductor digital circuit layouts evolved, it became evident that it was cumbersome to adapt and expand the set of mesh models described in [65]. I thus

developed a software tool that would read in an arbitrary layout from a mask layout file, convert it to a three-dimensional model according to a user-defined fabrication process description, mesh the model and run the calculation through FastHenry. Meshing had to be completely automated, with support for mask-to-wafer offset (the change in layout size between the mask drawing and the manufactured integrated circuit). I devised a “cake-slicing” method, illustrated in Figure 3.11 [195] that allowed automated generation of interleaved cuboid segments.

This was the first version of InductEx [195]. A model generated with this version of InductEx is shown in Figure 3.12.

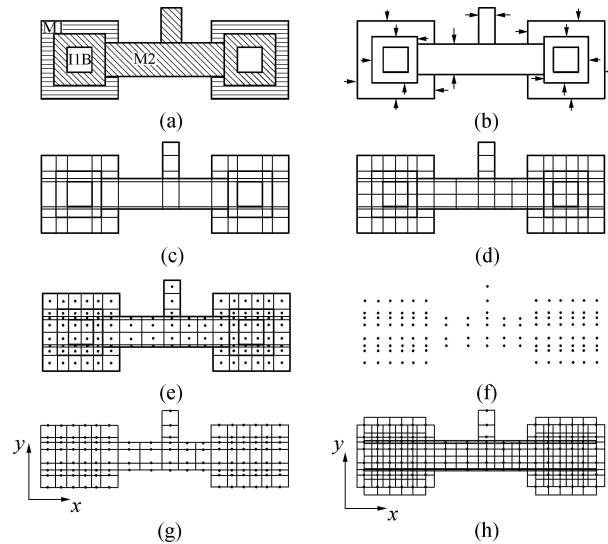


Figure 3.11: Simplified top view of part of a Josephson transmission line with two junctions connected by an inductor, showing the cake-slicing segmentation process. (a) Layout, (b) mask-to-wafer offsets applied, (c) all blocks sliced along block boundaries, (d) sub-slicing to limit maximum segment size, (e) segment centres identified, (f) nodes defined, (g) y -elements placed and (h) x -elements filled in. Vertical (z) segments are filled at junctions and metal-to-metal vias (not shown here).

In order to make InductEx usable as a tool, I developed it to read layouts from the industry standard GDSII binary stream file format. There were severe limitations, though. Layouts had to be flattened – containing no hierarchy or sub-cells. Only rectangular objects were accepted, and very strict text labels were used for port definition. The first version of InductEx could also only handle isolated inductors and coupling, but not netlists with multiple inductors connected to the same nodes. It also only worked under Windows.

I applied InductEx to the layout of RSFQ circuits [65] for the Hypres 1 kA cm^{-2} process as it existed circa 2003, as well as to an investigation into the expected variation of inductance due to process variations [195]. The results showed that process tolerances up to 10 % on isolation layer thickness would not produce more than 3 % to 6 % variation in the inductance of the typical layout structures that were used for RSFQ circuit layout at the time. I concluded that, if the on-chip inductance of a layout structure was close to the design value, process variations would not unduly degrade the operation of inductance-sensitive circuits.

At the time I could not obtain published results on any measurements of superconductor microstrip for integrated circuit processes against which to verify InductEx




Figure 3.12: (a) Microphotograph of a section of an RSFQ-to-COSL converter fabricated in the Hypres 1 kA cm^{-2} process showing a pair of coupled inductors and (b) the meshed inductance extraction model with reflection plane and images generated by the first version of InductEx.

calculations. All verification was thus done against analytical calculations with Chang's equation [173], which InductEx could easily match within a few percent when segment size was about one fifth to one tenth of line width. I had to trust that three-dimensional calculations would be of similar accuracy.

InductEx version 1 could handle models with 10 000 segments with relative ease, although calculations could run for an hour or more. All modules were compiled for 32-bit processors and thus had a memory limit of about 2 GB.

3.3.3 Multiterminal netlists

One of my responsibilities under the NioCAD project was to guide the development of inductance extraction capabilities. I developed an in-house proof-of-concept (informally numbered version 2 of InductEx) that used a port-to-port excitation scheme and the solution of simultaneous linear equations to find the inductance values of a multi-terminal netlist of interconnected inductors. It could solve all the self inductances in an RSFQ circuit, and was used to successfully design layouts for RSFQ cells in both the IPHT and Hypres processes. However, it could not handle mutual inductance, all layouts had to be reduced to rectangular blocks, and large circuits such as an RSFQ SFQ-DC converter had to be broken into two or more subsections to allow the models to fit in memory. The RSFQ SFQ-DC converter took around 20 hours to extract.

In July 2010, Dr Mark Volkmann, then a final year undergraduate electronic engineering student at Stellenbosch University, did his final year project under my supervision. He was tasked with designing a 20 GHz First-In-First-Out (FIFO) shift register, and his circuit layouts were so complex that InductEx could not extract a FIFO cell layout. I raced against time to develop a better extraction scheme and to speed up the FastHenry engine to allow Mark to finish his project. By October 2010, the port-to-port calculation scheme had been discarded, and InductEx now used a port excitation method whereby one port at a time was excited while all other ports were zeroed, and all port currents measured.

Mutual inductance calculation was also now possible. I made the segmentation algorithm more efficient to reduce segment count, and soon Mark could extract a full FIFO cell layout in about 90 minutes. I made further improvements by altering the preconditioner setup for FastHenry, and shaved 70 % off the calculation time for the largest circuits, so that a full FIFO cell could be extracted in under 20 minutes. By this time, NioCAD was no longer a project within Stellenbosch University, but a spin-off company to which I contributed research under contract. Under the commercial restrictions imposed by NioCAD (Pty) Ltd, this version of InductEx was limited to in-house use only. I numbered it version 4.0.




Figure 3.13: A circuit with inductance, resistance and mutual inductance.

Consider the circuit shown in Figure 3.13. During the magnetoquasistatic (MQS) solution of field equations for a given structure with a number of voltage ports, the current distribution in every segment of the structure is obtained. With only one port excited per solution, the voltage over every other port is zero and these ports are thus short circuits. The current through every port is found from the MQS solution, and through iteration the current through any circuit branch can be found as long as there are a sufficient number of ports.

With branch currents and port voltages known, equations can be composed to find the unknown branch inductance and resistance values, and any mutual inductance values.

From Kirchhoff's voltage law (KVL) it is known that the algebraic sum off all voltage drops around any loop in a circuit must equal zero. The first step towards building KVL equation is to find loops (or cycles in graph theory) that cover all the circuit components. There are four possible loops in the circuit in Figure 3.13, but it is possible to reduce the computation cost of very large circuit solutions by using only a fundamental set of cycles [199]. In this circuit, one set of fundamental cycles is shown. It contains three loops, marked as Loop1, Loop2 and Loop3.

The KVL equations for the three loops are:

$$\begin{aligned} V_1 &= I_1 j\omega L_1 - I_2 (R_2 + j\omega L_2) + I_2 j\omega M_{12} - I_1 j\omega M_{12} + I_3 j\omega M_{13} - I_3 j\omega M_{23} \\ V_1 - V_2 &= I_1 j\omega L_1 + I_2 j\omega M_{12} + I_3 j\omega M_{13} \\ V_3 &= I_3 (R_3 + j\omega L_3) + I_1 j\omega M_{13} + I_2 j\omega M_{23} \end{aligned}$$

Even if the voltages are all real-valued (typically 1 $\angle 0^\circ$ V), the presence of reactive components leads to branch currents with imaginary components. In order to manage the real and imaginary components of voltage and current, while keeping the values of resistance, inductance and mutual inductance strictly real, the KVL equations can be

expanded so that real and imaginary values of each loop voltage has a separate equation, and real and imaginary currents are handled separately:

$$\begin{aligned}
 V_{1r} &= -I_{1i}\omega L_1 - I_{2r}R_2 + I_{2i}\omega L_2 - I_{2i}\omega M_{12} + I_{1i}\omega M_{12} - I_{3i}\omega M_{13} + I_{3i}\omega M_{23} \\
 V_{1i} &= I_{1r}\omega L_1 - I_{2i}R_2 - I_{2r}\omega L_2 + I_{2r}\omega M_{12} - I_{1r}\omega M_{12} + I_{3r}\omega M_{13} - I_{3r}\omega M_{23} \\
 V_{1r} - V_{2r} &= -I_{1i}\omega L_1 - I_{2i}\omega M_{12} - I_{3i}\omega M_{13} \\
 V_{1i} - V_{2i} &= I_{1r}\omega L_1 + I_{2r}\omega M_{12} + I_{3r}\omega M_{13} \\
 V_{3r} &= I_{3r}R_3 - I_{3i}\omega L_3 - I_{1i}\omega M_{13} - I_{2i}\omega M_{23} \\
 V_{3i} &= I_{3i}R_3 + I_{3r}\omega L_3 + I_{1r}\omega M_{13} + I_{2r}\omega M_{23}
 \end{aligned}$$

where V_{1r} and V_{1i} denote the real and imaginary components of V_1 respectively and the notation applies to all voltages and currents. With real and imaginary components handled separately, the KVL equations can be written in matrix form as:

$$\begin{bmatrix} V_{1r} \\ V_{1i} \\ V_{1r} - V_{2r} \\ V_{1i} - V_{2i} \\ V_{3r} \\ V_{3i} \end{bmatrix} = \begin{bmatrix} -I_{1i} & -I_{2r} & I_{2i} & 0 & 0 & (-I_{2i} + I_{1i}) & -I_{3i} & I_{3i} \\ I_{1r} & -I_{2i} & -I_{2r} & 0 & 0 & (I_{2r} - I_{1r}) & I_{3r} & -I_{3r} \\ -I_{1i} & 0 & 0 & 0 & 0 & -I_{2i} & -I_{3i} & 0 \\ I_{1r} & 0 & 0 & 0 & 0 & I_{2r} & I_{3r} & 0 \\ 0 & 0 & 0 & I_{3r} & -I_{3i} & 0 & -I_{1i} & -I_{2i} \\ 0 & 0 & 0 & I_{3i} & I_{3r} & 0 & I_{1r} & I_{2r} \end{bmatrix} \begin{bmatrix} \omega L_1 \\ R_2 \\ \omega L_2 \\ R_3 \\ \omega L_3 \\ \omega M_{12} \\ \omega M_{13} \\ \omega M_{23} \end{bmatrix}$$

This system of linear equations,

$$\mathbf{b} = \mathbf{A}\mathbf{x}, \tag{3.15}$$

with \mathbf{b} as the vector of known loop voltages and \mathbf{A} as the matrix of known branch currents, can be solved for the vector of unknown component values, \mathbf{x} . Singular value decomposition (SVD) is used, and frequency is divided out to yield inductance and mutual inductance values.

With InductEx, the rows in \mathbf{b} and \mathbf{A} are duplicated for each voltage port as it is excited while the other ports are zeroed, for a total of $2N^2$ rows for a system with N ports. The computational cost of the SVD when all voltage loops in a circuits with many branches in the netlist are used can be significant (minutes), but the use of fundamental cycles keeps N low and results in SVD solutions that complete within milliseconds.

3.3.4 Validation

Under the S-Pulse project, I had the opportunity to fabricate integrated circuits with SQUIDs test structures from which inductance could be measured. I used InductEx to extract inductance from the SQUID layout structures, and had the circuits fabricated and measured at IPHT with the assistance of Dr Jürgen Kunert and Dr Olaf Wetzstein. The measurement of inductance in superconductor circuits is discussed in Section 3.4.1.

Some of the SQUIDs are shown in Figure 3.14. The InductEx models were constructed with a maximum segment size of 2.5 μm , and model size was managed by modelling vias

Table 3.1: Measured and InductEx-extracted inductance results for test SQUIDs over 4 chips manufactured on one wafer in the IPHT RSFQ1D process.

Chip	Inductor in layer	Measured (pH)	Uncalibrated (pH)	Calibrated (pH)	Error-calibrated to measured (%)
1	M2 ^a	9.83	9.20	9.89	+0.7
1	M1	5.61	5.72	5.62	+0.3
1	M1-M2	6.92	6.74	6.92	+0.0
2	M2 ^a	9.95	9.20	9.89	-0.6
2	M2 ^b	11.2	10.5	11.2	+1.7
2	M1 ^c	5.68	5.88	5.79	+1.9
3	M2	20.5	19.0	20.5	-0.1
3	M2 over hole ^d	17.4	19.6	20.2	+16
3	M1 over hole ^e	20.8	19.8	19.8	-5.1
4	M2 ^a	9.73	9.20	9.89	+1.7

^a Duplicate structure on different chips.

^b Structure shown in Figure 3.14(g) with SQUID loop inductor in M2 and control line in M1.

^c Structure with SQUID loop inductor in M1 inductively coupled to control line in M2.

^d SQUID with loop inductor in M2 that spans a hole in the ground plane.

^e SQUID with loop inductor in M1 that spans a hole in the ground plane.

as direct electrical connections between different layers. Lambda-width edge segments were used to improve current distribution modelling.

The calculation results were the first where InductEx could be compared directly to experimental measurements for actual three-dimensional circuit structures with vias and lines that cross over holes (as opposed to basic stripline structures). The calculation results were generally within 5 % of measured results, except for the lines over holes. The results were very encouraging, and validated the meshing and port excitation strategies that I employed in InductEx.

Given the known calculation errors resulting from segments that are larger than the London penetration depth, as shown in Figure 3.9, and the prohibitive cost (in terms of computing resources) of meshing structures the size of RSFQ logic gates, I investigated the possibility to increase the accuracy of calculations at a given segment size well in excess of the penetration depth. It turned out that by altering the penetration depth λ for the metal layers M1 and M2 from the process-specified 90 nm to $\lambda_{M1} = 83$ nm and $\lambda_{M2} = 140$ nm, the error between calculation and measurement for four calibration structures on one wafer could be reduced to a root mean square error (RMSE) of 0.45 %. The calibrated layer parameters were then applied to the calculation of the structures measured on a second wafer. The results are shown in Table 3.1 and Table 3.2.

It is clear that InductEx with cuboid segments for the FastHenry engine provided good calculation results for default layer parameters, and could be made more accurate with parameter calibration. At the time I could not explain the large error between calculation and measurement for the inductors over holes, and I was motivated to keep doing research on meshing, calculation methods and calibration.

It would become obvious later that calibration to structures that did not include holes was partly responsible for the error. Inadequate modelling of the ground plane – the current return path around the outside of a hole was not modelled – made up the rest. As I have witnessed countless times, inadequate or inaccurate modelling is the primary

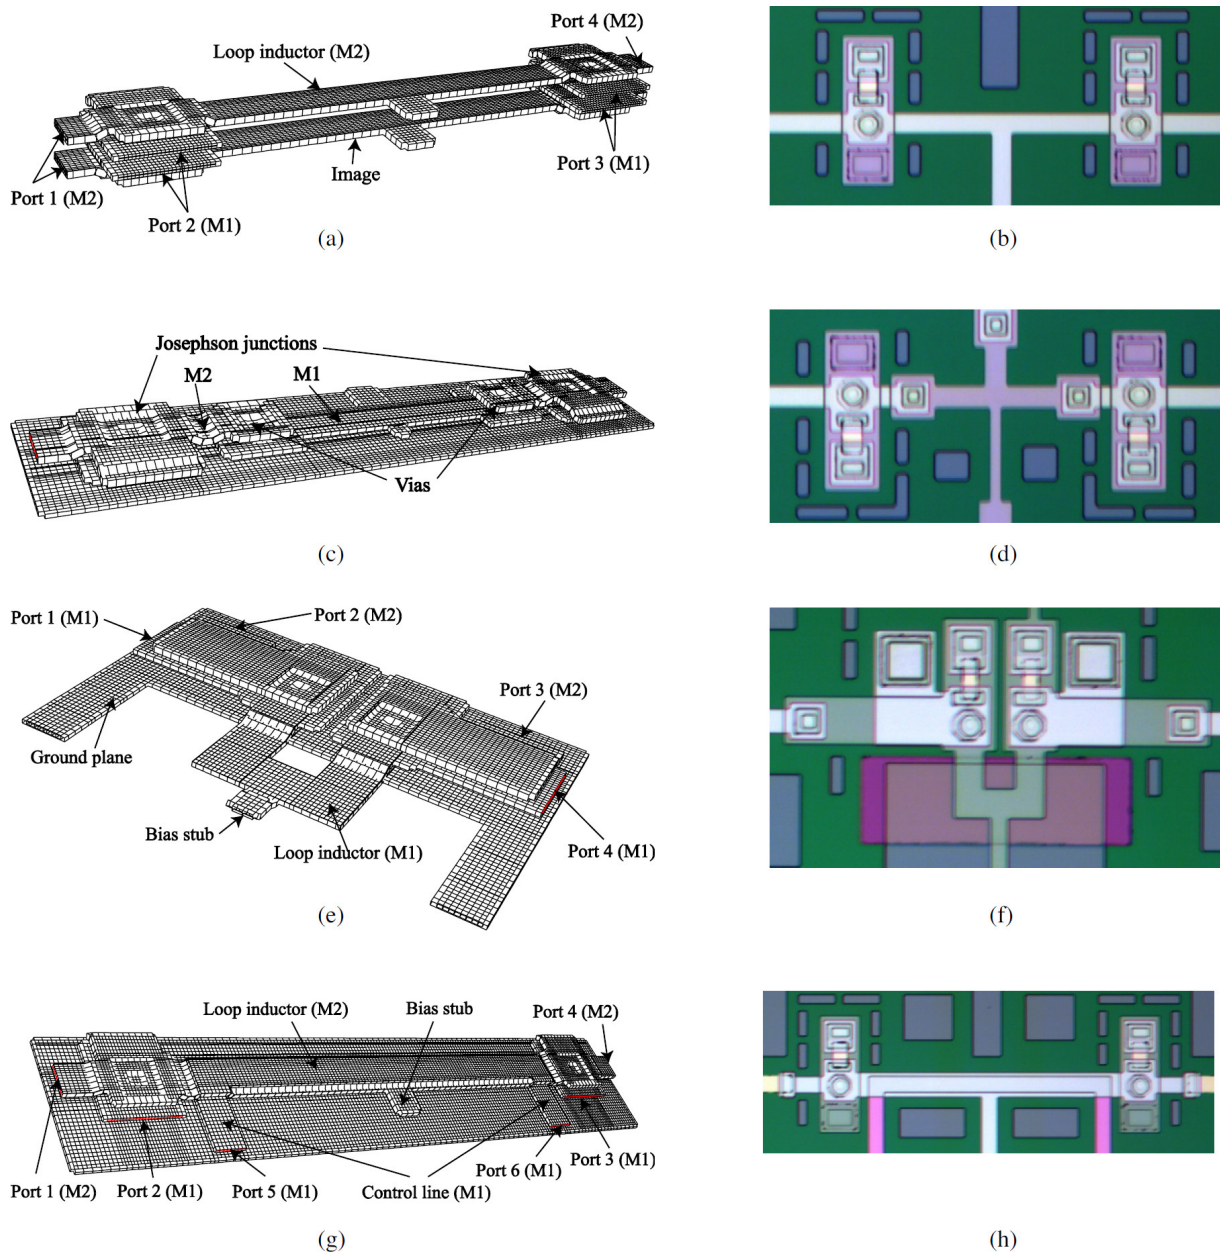


Figure 3.14: Segmented models and microphotographs of SQUID layouts in the IPHT RSFQ1D process. (a) Model of SQUID with $10\ \mu\text{m}$ wide loop inductor in metal layer M2, and image (using a reflection plane) and (b) M2 SQUID microphotograph. (c) Model of SQUID with vias and a $10\ \mu\text{m}$ wide loop inductor transitioning between layers M2 and M1, with ground plane included, and (d) VIA SQUID microphotograph. (e) Model of SQUID with inductor in M1 ($12.5\ \mu\text{m}$ wide) looping over a hole in the ground plane to form an enclosed hole of $12.5\ \mu\text{m} \times 10\ \mu\text{m}$ and (f) microphotograph of SQUID with ground plane hole. (g) Model of SQUID with $15\ \mu\text{m}$ wide loop inductor in M2 and $10\ \mu\text{m}$ wide control line passing between loop inductor and ground plane in M1 and (h) microphotograph of SQUID with coupled control line. For image clarity, vertical dimensions are enlarged five times.

Table 3.2: Measured and InductEx-extracted mutual inductance results for test SQUIDs manufactured in the IPHT RSFQ1D process.

Chip	Mutual inductance combination	Measured (pH)	Uncalibrated (pH)	Calibrated (pH)	Error-calibrated to measured (%)
2	M2-M1 ^a	3.84	3.98	4.00	+4.2
2	M1-M2 ^b	4.36	4.53	4.48	+2.3
4	M1-M2 ^b	4.32	4.53	4.48	+3.6

^a Structure shown in Figure 3.14(g) with SQUID inductor in M2 and coupled control line in M1.

^b Structure with SQUID loop inductor in M1 inductively coupled to control line in M2.

cause of calculation error.

3.3.4.1 Validation of inductance over holes

After the observation that inductance calculation results for structures over holes differed more from measurements than inductors over ground planes, and before the impact of incomplete modelling was understood, a targeted investigation of inductance over ground plane holes was launched. The dominant inductance extraction tool at the time, Lmeter, could not handle inductance over holes, and I undertook to develop and demonstrate proven capability in this to stimulate interest in my research outputs and InductEx in particular.

Between 2012 and 2013 I collaborated with Prof. Hannes Toepfer at Ilmenau University of Technology and a fabrication and test team at IPHT to design, fabricate and measure SQUID structures in the IPHT RSFQ1D process. I designed and analysed the SQUIDs with varying hole size shown in Figure 3.15, while Prof. Toepfer and Dr Olaf Wetzstein at IPHT made available results on a SQUID with a π -phaseshifter [200] with a ground plane hole underneath a conductor structure and on LR-bias inductors used for LR-biased RSFQ circuits.

The results were presented in [201]. For the test SQUIDs with varying ground plane hole size, some microphotographs from the manufactured chip are shown in Figure 3.15. An InductEx calculation model and the extraction netlist, which is the same for all the test SQUIDs, are shown in Figure 3.16. Experimental measurements provided the value for L_{loop} for every SQUID.




Figure 3.15: Microphotograph of, from left to right, a reference SQUID and several SQUIDs with varying size ground plane holes underneath the loop inductor. The SQUIDs were manufactured with the IPHT RSFQ1D (FLUXONICS) process. The largest hole depicted furthest to the right has dimensions $D = 100 \mu\text{m}$ and $W = 20 \mu\text{m}$.

The measured and calculated results are shown in Figure 3.17, and show excellent agreement. Here, of course, the ground plane was modelled to close around the hole and provide adequate current return paths.

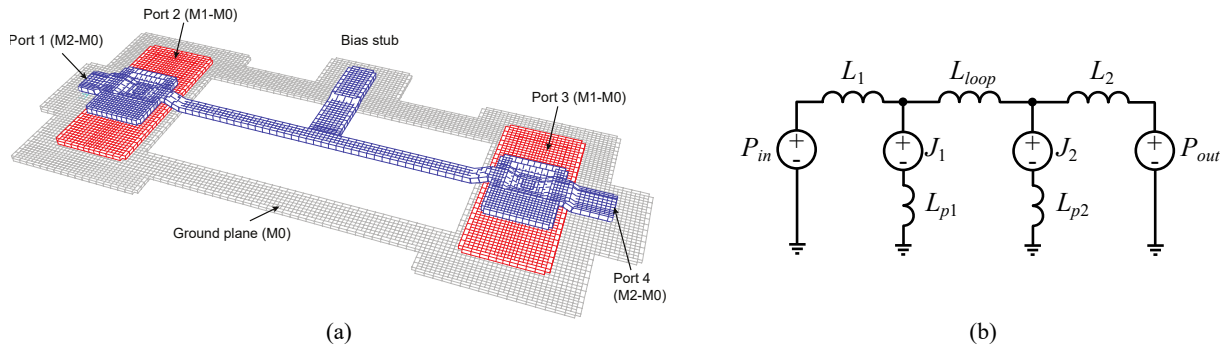


Figure 3.16: (a) InductEx model of the SQUID to the far right in Figure 3.15 meshed with cuboid segments and (b) extraction netlist with all excitation ports and inductors.




Figure 3.17: Measured and extracted inductance results for L_{loop} of test SQUIDs manufactured with the FLUXONICS process with loop inductors over ground plane holes. Line width $s = 5 \mu\text{m}$. All holes have length $D = 100 \mu\text{m}$, while width W on either side of the line is varied.

Table 3.3: Measured and InductEx-extracted SQUID loop inductances over ground plane holes.

	π -phaseshifter	LR-bias inductor
Calculated (pH)	4.19	79.6
Measured (pH)	4.10	79.5

As a further investigation, the inductance of two other ground plane hole structures, a π -phaseshifter and high-inductance bias line for LR-biasing, were calculated and compared with measurements. The layout structures are shown in Figure 3.18(a) and Figure 3.18(c) respectively. The results are listed in Table 3.3, and are very accurate. The improvement over earlier experiments lay once again in proper modelling of the ground plane around the holes to allow enough width to model the return current flow with sufficient accuracy.

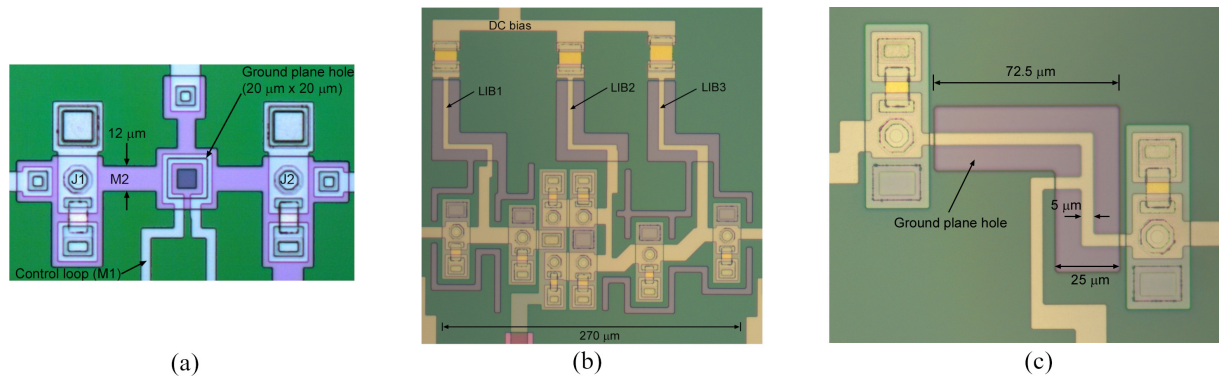


Figure 3.18: Microphotographs of circuits with ground plane holes under inductors manufactured with the FLUXONICS process. (a) A SQUID containing a π -phaseshifter, (b) a Toggle-flip-flop RSFQ circuit with LR-bias inductors and (c) a SQUID containing just the LR-bias inductor for measurement.

As a result of this work, it was no longer necessary to build test structures just to determine the inductance of a structure such as an LR-bias line. The InductEx model shown in Figure 3.19, with adequate ground plane modelling around every hole, allowed analysis of even the LR-bias inductors. Incidentally, this InductEx model showed that connection of the flux trap moats to the LR-bias holes raises the inductance of the LR-bias inductors significantly (by about 25 % for L_{IB1} and L_{IB3}) due to the longer return path for bias current around the resulting extended holes. The increased inductance was a bonus for LR-biasing, but what we did not know then (and could not model, calculate and analyse until several years later) was that coupling from both the bias current – returning around the moat-extended ground plane holes – and any fluxons trapped in these holes to the RSFQ cell’s internal inductors will severely distort quiescent current distribution and gate circuit operating margins.

3.3.5 Fabrication-ready layout processing

3.3.5.1 Full-circuit layouts

The evolution of InductEx and its supporting tools was driven by both technical requirements, such as capability and support for materials, processes and device geometries, and by user requirements. The demands from technical and user requirements increased in parallel as tool use increased.

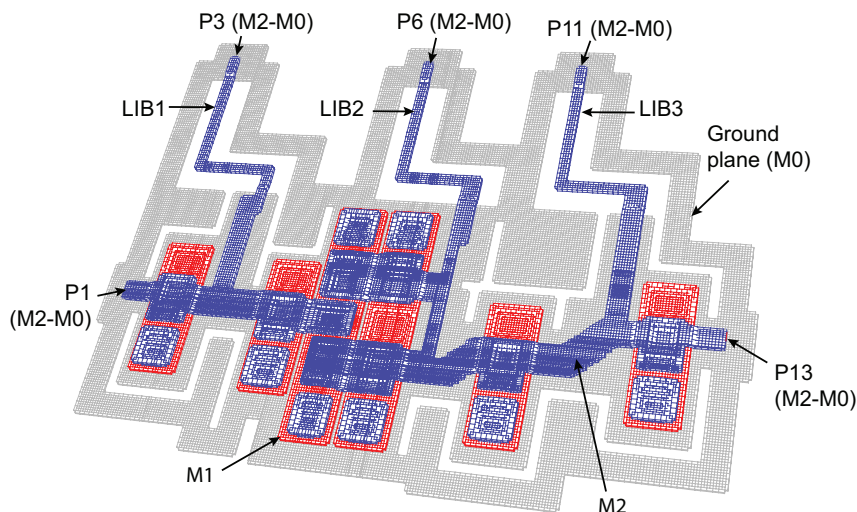


Figure 3.19: InductEx extraction model of an LR-biased RSFQ Toggle-flip-flop shown in Figure 3.18

Early use of InductEx required layouts to be flattened to remove hierarchy, and limited port definitions to very strict layout limits. While it reduced the development time, it cost considerable effort for layout engineers to generate layout snippets from full-chip layouts, flatten these, delete unwanted geometry (in order to reduce computational cost), and crucially *to reproduce any edits made during inductance matching on the extraction layout exactly in the original full-chip artwork*.

In order to promote InductEx from a tool only suited to academic laboratories to one fully functional in commercial and industrial environments, some improvements were made. Many of these improvements were specifically made to allow extraction from fabrication-ready full-chip layouts. The most important of these are:

- Full support for an unlimited hierarchy depth in GDS layouts. This required processing of structures with translation, rotation and array placements.
- Designation of a *top* structure during calculation which can differ from the top structure in the GDS file. This allows extraction of a specific cell or circuit layout (with its subcomponents of lower order in the layout hierarchy) from a larger layout, and is routinely used to target a specific cell or module in a full chip layout.
- Model reduction operators and designation objects that can be placed on non-fabrication layers, where these can remain in the tape-out layouts while instructing InductEx to remove or simplify layout objects that are not of interest during extraction.

3.3.5.2 Resistance

Originally, InductEx was developed purely to calculate inductance. Objects on resistive layers were not modelled, although the profile of such objects was considered in the topography of a layer stack to adjust the height of unplanarised layers above.

As soon as support for fabrication-ready layouts was added, it became evident that *the inductance of shunt resistor branches* was also of interest to circuit designers during layout validation. If the full layout structure of a resistor was modelled, then the resistance

could also be calculated (as a parasitic component to the branch inductance, from the perspective of InductEx). Support for resistance required changes to three-dimensional model building and an overhaul of the linear algebra required to solve the unknown *impedances* in a system with inductive and resistive branch components [202].

For the original field solver, FastHenry, resistivity had to be specified through a conductivity parameter *sigma*, which is the inverse of the bulk resistivity ρ and can be calculated from the per-square resistance and layer thickness of a resistive layer as

$$\sigma = \frac{1}{\rho} = \frac{1}{R_{sheet}d} \quad (3.16)$$

in siemens per unit length. R_{sheet} is the sheet resistance in ohm per square (Ω/\square) and d is the resistive layer thickness. For the resistive layer R1 in the FLUXONICS 1 kA cm^{-2} process, with $R_{sheet} = 1 \Omega/\square$ and $d = 80 \text{ nm}$, $\sigma = 12.5 \text{ S } \mu\text{m}^{-1}$.




Figure 3.20: Schematics of the layer stack surrounding a resistive layer between two superconductive metal layers for popular fabrication processes in 2014. (a) FLUXONICS 1 kA cm^{-2} , (b) Hypres 4.5 kA cm^{-2} and (c) AIST STP/ADP2 $2.5/10 \text{ kA cm}^{-2}$ processes.

The section of the layer stack surrounding the resistive layer in three processes as these existed in 2014 is shown in Figure 3.20. Until then, I simplified model building by assuming that a via between superconductor layers could cut all isolation and directly connect the nearest metal layers above and below. InductEx simply deleted vias where they overlapped resistors.

As is evident from Figure 3.20, different processes have very different approaches to the layer stack and mask sets. Different via masks can be used to cut into the same isolation, or certain vias are etch-stopped before passing through all the isolation layers, while other vias can etch-stop against a resistive layer or *bypass it entirely* to connect to the next metal layer below (thereby etching through another isolation layer). This forced a rethink of how the layer stack is programmed (through the LDF file) into InductEx. Details were published in [202]. In short, a generic layer description process was developed that can utilise non-fabrication masks to create copies of isolation layers and insert these anywhere in the layer stack, layer objects can be added to or subtracted from object sets on multiple other layers or from multiple other layers, and vias can be programmed to

bypass metal layers selectively where the via does not overlap a metal layer object. The end result is a programmable layer stack that is so capable that it is still able to support the latest evolution of all the fabrication processes to which InductEx is applied.

<pre> \$Layer Name = I1A Thickness = 0.07 Mask = -1 Filmtype = I IDensity = 1e-5 \$End \$Layer Name = I2A Thickness = 0.15 Mask = -1 Filmtype = I \$End \$Layer Name = R1 Thickness = 0.08 Sigma = 12.5 Mask = 1 Filmtype = R ViaBypass = TRUE \$End \$Layer Name = I2B Thickness = 0.15 Mask = -1 Filmtype = I \$End </pre>	<pre> \$Layer Number = 59 Name = I1BL Thickness = 0.1 Mask = -1 Filmtype = I LayerADD = 3 LayerSUB = 9 \$End \$Layer Number = 9 Name = R2 Thickness = 0.07 Sigma = 6.803 Mask = 1 Filmtype = R ViaBypass = TRUE \$End \$Layer Number = 3 Name = I1B Thickness = 0.1 Mask = -1 Filmtype = I \$End </pre>	<pre> \$Layer Number = 2 Name = GC Thickness = 0.15 Mask = -1 Filmtype = I PlanarModel = 1 \$End \$Layer Number = 3 Name = RES Thickness = 0.035 Sigma = 11.91 Mask = 1 Filmtype = R ViaBypass = TRUE \$End \$Layer Number = 9 Name = RC Thickness = 0.15 Mask = -1 Filmtype = I LayerADD = 2 \$End </pre>
(a)	(b)	(c)

Figure 3.21: Excerpts from the layer definition files for the layer stacks between the nearest metal layers above and below the resistive layer for popular fabrication processes in 2014. (a) FLUXONICS 1 kA cm^{-2} , (b) Hypres 4.5 kA cm^{-2} and (c) AIST ADP2 10 kA cm^{-2} process. (For the AIST STP 2.5 kA cm^{-2} process, the thickness of layer RES is $0.08 \mu\text{m}$ and $\sigma = 10.417 \text{ S } \mu\text{m}^{-1}$). Some layer parameters are omitted for brevity. The PlanarModel parameter in (c) selects the planarisation method for the AIST process.

3.3.6 Solution speedup

The original FastHenry [164] MQS solver for the calculation of inductance in three-dimensional structures, adapted for superconductivity and altered to support multi-terminal inductance extraction with InductEx, remained popular by 2013. It was originally developed for printed circuit board layouts with slender conductors. When applied to densely discretised IC structures, it became slow. It was accepted that a model with fewer than 10 000 mesh elements could solve in tens of minutes to an hour, that a model with 50 000 mesh elements could require more than a day to solve, and that 100 000 mesh elements was a theoretical upper limit for mesh element count. I remember students waiting 2 weeks for an extraction result on two coupled coils. By the time Dr Mark Volkmann visited Hypres to help design eSFQ shift register and TFF cells [115] (see Section 2.6.3) I was regularly analysing circuits such as the eSFQ TFF cell that took more than a day to extract.

Clearly, FastHenry was inefficient at higher mesh element counts, but simply limiting the number of elements in a model was not an option. Larger circuits with more ground

planes would eventually need millions of mesh elements. I thus set out to characterise the inefficiencies and try to find ways to increase the solution speed of FastHenry while also investigating the development of a new MQS solver. This was the first task for Dr Kyle Jackman when he started his Masters degree in 2014.

3.3.6.1 FastHenry overview

FastHenry combines three components: mesh analysis, an iterative solver known as the generalised minimal residual method (GMRES) [203] and the Fast Multipole Method (FMM) [204] from which the name was derived. Calculation models for FastHenry are discretised into cuboid filaments (also called segments) in which uniformly distributed currents and voltages are assumed to be sinusoidal and at steady state. Filaments are connected together through nodes, so that a discretised structure with connected filaments can be represented as an equivalent circuit with branches for the filaments, as shown in Figure 3.22 [164]. Now

$$\mathbf{Z}\mathbf{I}_b = (\mathbf{R} + j\omega\mathbf{L})\mathbf{I}_b = \mathbf{V}_b, \quad (3.17)$$

where \mathbf{I}_b and \mathbf{V}_b are vectors for the current and voltage phasors of each branch and \mathbf{Z} is the complex impedance matrix. Here, \mathbf{R} is the diagonal matrix of dc resistances and \mathbf{L} is the matrix of self-inductances of each filament on the diagonal and partial inductances between all filaments everywhere else.

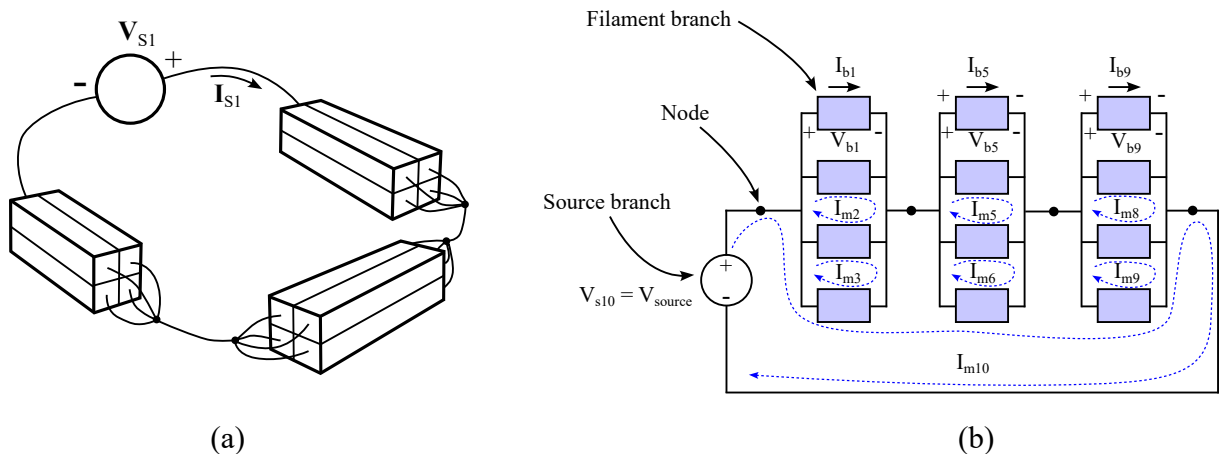


Figure 3.22: A conductor excited by a voltage source, (a) with discretised filaments connected to nodes, and (b) modelled as a circuit.

A nodal analysis formulation and current conservation at each node can be used to construct linear equations, the solution of which yields node voltages and branch currents. Solution by direct factorisation is prohibitively slow, and iterative methods converge slowly for the nodal analysis formulation [164]. FastHenry thus uses a mesh analysis method.

In mesh analysis, as illustrated in Figure 3.22(b), a mesh is any loop of branches that does not enclose any other branches. Here, the currents flowing around any mesh in the network are the unknowns and are represented by vector \mathbf{I}_m .

An m by n mesh matrix, \mathbf{M} , where m is the number of meshes and n is the number of filament branches, describes the position and orientation of the filaments in every branch. Entries are 0, 1 or -1. The mesh matrix is mostly empty – there are typically only a few filaments per mesh – and is therefore assumed to be sparse.

From Kirchhoff's voltage law, the sum of branch voltages around every mesh is zero, or

$$\mathbf{M}\mathbf{V}_b = \mathbf{V}_s, \quad (3.18)$$

where \mathbf{V}_s is the vector of source branch voltages.

The vector of mesh currents is related to the branch currents and mesh matrix as

$$\mathbf{M}^T\mathbf{I}_m = \mathbf{I}_b. \quad (3.19)$$

Combining (3.17), (3.18) and (3.19) yields

$$\mathbf{M}\mathbf{Z}\mathbf{M}^T\mathbf{I}_m = \mathbf{V}_s. \quad (3.20)$$

The system of linear equations in (3.20) is intractable with Gaussian elimination when there are thousands of filaments, let alone millions of filaments as in large circuit models. An iterative method, GMRES, is thus used. The costliest step during each GMRES iteration is the computation of the matrix-vector product $(\mathbf{M}\mathbf{Z}\mathbf{M}^T)\mathbf{I}_m^k$, where \mathbf{I}_m^k is the basis vector for the Krylov subspace computed at the k^{th} iteration [203]. This requires $O(m^2)$ operations. FastHenry reduces the matrix-vector product to $O(m)$ operations by using the Fast Multipole Method [204] to form an approximation to the matrix-vector product whenever needed, without ever computing $(\mathbf{M}\mathbf{Z}\mathbf{M}^T)$ explicitly.

The convergence rate of the GMRES iterative method is then improved by using a preconditioned system [205]. The equation

$$\mathbf{M}\mathbf{Z}\mathbf{M}^T\mathbf{P}\mathbf{x} = \mathbf{V}_s \quad (3.21)$$

has the same solution as (3.20) for some square matrix \mathbf{P} with the same dimension as $\mathbf{M}\mathbf{Z}\mathbf{M}^T$ if we set $\mathbf{I}_m = \mathbf{P}\mathbf{x}$. The matrix \mathbf{P} is a preconditioner. A good preconditioner is one that is quick to form and to apply, with a significant reduction in GMRES iterations. Naturally, if the time that it takes to compute the preconditioner is longer than the time saved in GMRES iterations, then the preconditioner is of no use.

Many preconditioning techniques exist, and most aim to make \mathbf{P} as close as possible to \mathbf{A}^{-1} . FastHenry supports sparsified preconditioners, one where the matrix \mathbf{L} of partial inductances in (3.17) is sparsified by dropping all the mutual inductances outside of the cubes formed during FMM; the other by using only the diagonal of \mathbf{L} . These are referred to as ‘‘Cube’’ and ‘‘DiagL’’ respectively.

3.3.6.2 FastHenry characterisation

With the calculation steps identified, we set out to characterise the typical time spent on each when representative superconductor integrated circuit models are solved. I selected four practical examples with increasing complexity:

1. A set of coupled coils with 2 ports and a slender line geometry for a digital SQUID magnetometer [206] from the FLUXONICS process shown in Figure 3.23(a), with 7635 filaments.
2. An 8-port AQFP cell with a single ground plane [207] from the AIST HSTP process shown in Figure 3.23(b), with 23 090 filaments.

3. A 21-port RSFQ toggle flip-flop (TFF) [108] with a single ground plane from the FLUXONICS process shown in Figure 3.23(c), with 37 274 filaments.
4. A 17-port eSFQ TFF (eTFF) [115] from the Hypres 4.5 kA cm⁻² process shown in Figure 3.23(d), with 47 341 filaments.

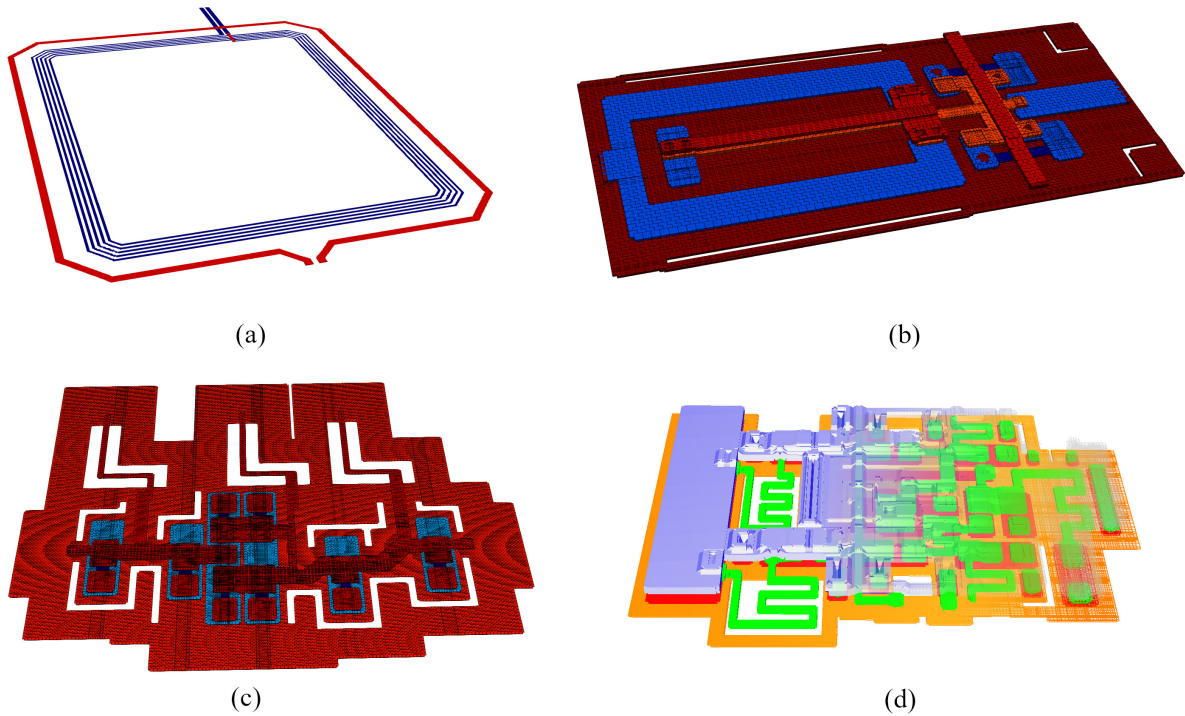


Figure 3.23: InductEx calculation models for the FastHenry engine. (a) Two coils, (b) an AQFP cell, (c) an RSFQ TFF and (d) an eSFQ TFF.

The results for the four example models are shown as pie charts in Figure 3.24. On all but the most simple structure, the computation of the preconditioner dominated (taking 98.9 % of the run time of more than a day on the eSFQ TFF cell).

Clearly, the preconditioner is not efficient for the models with interleaved mesh elements that are used for superconductor integrated circuit structure modelling. The default preconditioner, a cube-block sparsified-L method was shown earlier to outperform all the other preconditioners supported by FastHenry on industrial models with around 3000 filaments [208], although it was speculated that with more than an order of magnitude more non-zero elements than the diagonal-of-L sparsified preconditioner, cube-block might be prohibitively expensive for larger problems.

I thus tried the diagonal-of-L preconditioner, and found that, even though it converged slightly slower than the cube-block preconditioner, its calculation was much faster, thus significantly reducing total computation time. This is also shown in Figure 3.24.

Simply bypassing the preconditioner is not efficient either. In the larger models, as shown in Table 3.4, GMRES solution takes so many iterations without a preconditioner that it is generally not more efficient than solutions with the diagonal-of-L preconditioner. At this stage, the decision was made to abandon FastHenry entirely and to build a new MQS solver that is based on FastHenry, but with all the solution steps optimised. We eventually called it Fast FastHenry (FFH).




Figure 3.24: Breakdown of the time spent on the main solution steps in FastHenry for different extraction models when a cube-block (Cube) and diagonal-of-L (Diag) preconditioner is used. The steps are multipole setup (MPS), preconditioner calculation (Precon), GMRES and all other minor steps (Other).

3.3.6.3 Fast FastHenry (FFH)

When Dr Kyle Jackman started his postgraduate studies in my group, his task was to develop FFH. He started from the open source code base of FastHenry, but used modern C libraries and coding techniques to reduce memory use. FFH improved on FastHenry in three main areas: multipole setup (MPS), construction of the preconditioner and GMRES.

MPS involves the calculation of multipole to local expansion operators and the near-part matrices [209], which are used in the FMM. The dominant cost of MPS is the construction of the near-part matrices which store the near-field interactions between filaments. The entire circuit is divided into cubes. The near-field interactions within the finest cubes are calculated independently, which allows for easy parallelisation with negligible thread management overheads in FFH. Cubes are grouped together to ensure even load balance between threads, and the MPS then speeds up when more processing cores are used.

The preconditioner \mathbf{P} is sparse and it is therefore only necessary to store the non-zero values. FastHenry uses linked lists that are slow when values are added or modified. For FFH, the construction of the \mathbf{P} uses routines from the CXSparse library [210] which uses the compressed column format for storing sparse matrices. These routines reduce run time and memory usage compared to linked lists.

For the LU decomposition [211] of \mathbf{P}^{-1} , the SuperLU_MT library [212] is used in FFH. SuperLU_MT implements an asynchronous parallel supernodal algorithm for sparse Gaussian elimination [213] that vastly outperforms the LU decomposition algorithms in FastHenry. Combined, the routines from the CXSparse and SuperLU_MT libraries reduce the construction time of the preconditioner by 50 to 130 times in our test examples. We

Table 3.4: Calculation times for original FastHenry and Fast FastHenry (FFH) with different preconditioner options and processor core counts.

Layout model	FastHenry Cube	FastHenry DiagL	FastHenry no preconditioner	FFH DiagL 1 core	FFH DiagL 2 cores	FFH DiagL 4 cores
Coupled coils	9 s	8 s	35 s	4.7 s	3.1 s	2.4 s
AQFP cell	1 692 s	548 s	764 s	60 s	37 s	23 s
RSFQ TFF	47 632 s	6 627 s	14 457 s	273 s	181 s	131 s
eSFQ TFF	106 706 s	14 138 s	12 069 s	399 s	255 s	162 s

found that construction time for the Cube preconditioner is on average 7 times longer than that of the DiagL preconditioner with FFH on typical calculation models, so that even though Cube provides faster GMRES conversion it delivers an overall lower speed gain. FFH thus uses the DiagL preconditioner by default.

The dominant computation cost of GMRES is the matrix-vector product. FMM uses more than 90% of this time. FastHenry implements the FMM through an electrostatic analogy by integrating the vector potential across each filament [164]. The vector potential is decomposed into its x , y , and z components; each component considered a scalar electrostatic potential. In FFH, instead of evaluating the FMM separately for each dimension, a separate set of updating vectors that includes the real and imaginary parts is created for each dimension. Updating vectors are used for storing the results of each FMM stage and require negligible memory. This modification delivers a speed increase of nearly 4 times when computing the matrix-vector product. Furthermore, duplicating the updating vectors and assigning a set to each thread, several matrix-vector products (one for each GMRES) can be computed in parallel with negligible memory increase per additional thread. Finally, typical multiport extraction models use many excitation ports. GMRES is executed once for every port, so that most gain is obtained for multiport calculations when each GMRES is executed in a separate thread for every processor core available on a computer.

These improvements made FFH vastly superior to FastHenry in terms of calculation speed, while still delivering the same calculation results. A comparison of execution times is shown in Table 3.4.

It is hard to overstate the staggering improvement in calculation time achieved with FFH during the course of 2014. The extraction of inductance for the eSFQ TFF reduced from *more than a day to less than 3 minutes*. (Note: as of writing, InductEx with triangular segments solves the eSFQ TFF extraction *in 10 seconds!*)

With 64-bit support, improved memory management and better segmentation algorithms, the limit on inductance extraction model size of about 50 000 filaments, previously hemmed in by 2 GB of RAM or days of computing time, expanded vastly. At the time of writing we regularly extract inductance from models with millions of filaments, with computer memory the only limitation. A system with 128 GB of RAM can now easily handle a model with 5 million filaments and complete extraction in tens of minutes to a few hours. For typical logic gates with around 100 000 filaments, solution takes only a few seconds.

3.3.7 A new engine: TetraHenry

Although the cuboid meshes used with FastHenry and FFH work well for gate-level inductance extraction, it soon became evident that really large models – now supported with the enhancements of FFH – were being built for chip-scale extractions: SQUID layouts, ground plane return current modelling, and multi-gate layouts in processes with multiple ground layers where cuboid meshes were very ineffective. One important limitation to cuboid meshing is that a filament width has to be perpetuated along over the entire dimension of a structure in that direction. Alignment of via segments to lower and upper layers also enforce the same mesh widths over multiple layers, which often leaves ground planes with high filament density perpetuated to the ends of the plane. Another limitation is that complex curvatures are poorly modelled, especially narrow line spiral inductors or gradiometer loops.

In an effort to support more general layouts with higher fidelity, I tasked Dr Kyle Jackman with investigating the development of a solution engine that would process tetrahedral meshes for his Masters degree. We expected that such an engine would perform better on non-Manhattan geometries, but we never imagined *how much better* it would perform in terms of speed and solution accuracy for all layouts.

Kyle returned with a proof-of-concept ten days later. It was slow but very promising, and we redirected the majority of our development focus to implement it as the new engine TetraHenry [209], [214], which we often abbreviate as TTH. TetraHenry and the applications that it opened was so successful that Kyle's Masters was upgraded to a PhD which he completed in 2017. The theory behind and development of TetraHenry is presented in detail in Kyle's PhD dissertation [215], and very briefly summarised here.

3.3.7.1 Tetrahedral modelling

TetraHenry was first developed to handle meshes with tetrahedral volume elements that would allow multidimensional current flow in complex superconducting structures.

The volume electric current integral equation (VJIE), used in [216]–[219], was chosen as the most suitable method for modeling superconducting currents. The VJIE formulation is similar to the volume integral equation (VIE) used in FastHenry and FFH, and requires fewer iterations when using an iterative method [217].

Dr Jackman derived the VJIE formulation for superconducting structures, starting with the MQS Maxwell's equations and assuming sinusoidal steady-state. Volume Loop (VL) basis functions [220], a combination of Schaubert-Wilton-Glisson (SWG) functions [221], are used to discretise the VJIE. The Method of Moments (MoM) [222] is used to construct a linear system of equations from the VJIE, which is solved using the GMRES iterative method [203] just as with FastHenry and FFH. Similarly, the matrix-vector product in the GMRES is accelerated using the FMM [204]. Algorithmic improvements and parallelisation methods developed in [209] for GMRES and the FMM were modified and implemented in TTH.

Tetrahedral meshing is a non-trivial operation. Where I could develop efficient algorithms to do cake-slice cuboid meshing for FastHenry and FFH [195], tetrahedral meshing of complex structures is best done by a third party finite element mesh generator. After evaluating a few candidates, we settled on Gmsh [223].

A proof-of-concept is one thing, but a functional commercial grade tool is another beast entirely. We needed years to redesign and implement modelling methods in InductEx to handle generic layer interconnects and excitation port definitions for tetrahedral meshing.

3.3.7.2 Volume Integral Equation

The VJIE formulation is derived by starting with Maxwell's equations and assuming sinusoidal steady-state, as discussed in [164]. The displacement current is assumed negligible, i.e. MQS approximation, since the conductivity is large within the conductors. The VJIE can be obtained as follows [215]:

$$\frac{\mathbf{J}(\mathbf{r})}{\sigma(\mathbf{r})} + \frac{j\omega\mu}{4\pi} \int_{V'} \frac{\mathbf{J}(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} dv' = -\nabla\phi(\mathbf{r}), \quad (3.22)$$

In (3.22), $\mathbf{J}(\mathbf{r})$ and $\phi(\mathbf{r})$ are respectively the volume current and scalar potential. The conductivity $\sigma(\mathbf{r})$ can vary within the conductor, while the permeability is considered constant everywhere: $\mu(\mathbf{r}) = \mu_0$. Support for superconductivity is added to (3.22) by replacing $\sigma(\mathbf{r})$ with the complex conductivity, $k(\mathbf{r})$, using London equations and the two-fluid model [2]:

$$k(\mathbf{r}) = \tilde{\sigma}_0(\mathbf{r}) + \frac{1}{j\omega\mu\lambda(\mathbf{r})^2}. \quad (3.23)$$

The meshing engine does not account for the London penetration depth, but the non-uniform current density can be modeled by increasing the number of meshing layers near the surface of the superconductor.

3.3.7.3 Discretization

In order to model current flow in piecewise homogenous objects, the Full-SWG basis function [221] is used to expand $\mathbf{J}(\mathbf{r})$. Figure 3.25 shows an arbitrary body with piecewise constant electrical parameters, discretized using Full-SWG functions. Figure 3.26 shows the definition of the Full-SWG basis function. The two tetrahedrons, T_n^+ and T_n^- , are associated with the n th face in the discretized region. The position vectors $\rho_{\mathbf{n}}^+$ and $\rho_{\mathbf{n}}^-$ represent points in T_n^+ and T_n^- , respectively. The vector $\rho_{\mathbf{n}}^+$ is defined with respect to (thus from) the free vertex in T_n^+ . The vector $\rho_{\mathbf{n}}^-$ is defined towards the free vertex in tetrahedron T_n^- [221]. The sign of the two tetrahedrons depends on the choice of the direction of current flow.

To simplify the problem, the entire volume is first assumed to be a homogeneous dielectric body, preventing surface charges accumulation. The Full-SWG function can then be used within the entire volume:

$$\mathbf{f}_{\mathbf{n}}(\mathbf{r}) = \begin{cases} \frac{1}{3|v_n^+|} \rho_{\mathbf{n}}^+(\mathbf{r}), & \text{if } \mathbf{r} \in T_n^+ \\ \frac{1}{3|v_n^-|} \rho_{\mathbf{n}}^-(\mathbf{r}), & \text{if } \mathbf{r} \in T_n^- \\ 0, & \text{otherwise} \end{cases}, \quad (3.24)$$

where $|v_n^\pm|$ represents the volume of tetrahedron T_n^\pm . This function differs from the basis functions used in [221] and [219], which uses the area of the face to normalize $\mathbf{f}_{\mathbf{n}}(\mathbf{r})$. Using Full-SWG functions, the volume electric current density, $\mathbf{J}(\mathbf{r})$, can be expanded as follow:

$$\mathbf{J}(\mathbf{r}) = \sum_{n=1}^N i_n \mathbf{f}_{\mathbf{n}}(\mathbf{r}), \quad (3.25)$$

where N is the number of faces that make up the entire volume and i_n is the *branch* current through the n th face.




Figure 3.25: Full-SWG basis functions in arbitrary body with piecewise constant electrical parameters.




Figure 3.26: Full-SWG basis function.

3.3.7.4 Method of Moments

Following the Method of Moments and using (3.24) as testing functions, (3.22) can be converted to a system of N independent equations [215]:

$$ZI_{branch} = (R + j\omega L)I_{branch} = V_{branch}, \quad (3.26)$$

where I_{branch} and V_{branch} are vectors containing N branch currents and voltages, respectively. The matrices R and L represent the real and imaginary part of Z . The elements of matrices R and L at index (m, n) are computed as:

$$R_{m,n} = \int_{v_m} \frac{1}{k(\mathbf{r})} \mathbf{f}_m(\mathbf{r}) \cdot \mathbf{f}_n(\mathbf{r}) dv, \quad (3.27)$$

$$L_{m,n} = \frac{\mu}{4\pi} \int_{v_m} \int_{v_n} \frac{\mathbf{f}_m(\mathbf{r}) \cdot \mathbf{f}_n(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} dv' dv. \quad (3.28)$$

The element at index m of vector V_{branch} is computed as:

$$(V_{branch})_m = - \int_{v_m} \mathbf{f}_m(\mathbf{r}) \cdot \nabla \phi(\mathbf{r}) dv, \quad (3.29)$$

The volumes v_m and v_n represent the volumes of tetrahedrons $T_m^+ + T_m^-$ and $T_n^+ + T_n^-$, respectively. The weighting function, $\mathbf{f}_m(\mathbf{r})$, are defined as Full-SWG functions, as given in (3.24).

3.3.7.5 Volume Loop Basis Function

The divergence of the electric flux within the conductor should be zero, but the divergence of the basis function is non-zero [221]:

$$\nabla \cdot \mathbf{f}_n(\mathbf{r}) \begin{cases} \frac{1}{3|v_n^+|}, & \text{if } \mathbf{r} \in T_n^+ \\ \frac{1}{3|v_n^-|}, & \text{if } \mathbf{r} \in T_n^- \\ 0, & \text{otherwise} \end{cases} \quad (3.30)$$

Several schemes have been developed to ensure the divergence free condition, such as basis reduction [224] and volume loop (VL) basis set [225]. The basis reduction scheme reduces the number of unknowns, but the matrix equation has a large condition number; making it difficult to solve with an iterative method [219]. The VL basis function, associated with each edge, also ensures the divergence free condition. Figure 3.27 shows the closed and unclosed VL basis function.




Figure 3.27: (a) Close volume loop basis function and (b) unclosed volume loop basis function.

The volume loop basis function around the edge m can be defined as a combination of SWG functions [225]:

$$\mathbf{o}_m(\mathbf{r}) = \sum_{n=1}^N M_{m,n} \mathbf{f}_n(\mathbf{r}), \quad (3.31)$$

where $M_{m,n}$ is the value at index (m, n) of matrix M and can be either 0 or ± 1 , depending on the direction of the loop and the SWG function. The volume electric current density, $\mathbf{J}(\mathbf{r})$, can now be expanded in terms of VL basis functions:

$$\mathbf{J}(\mathbf{r}) = \sum_{m=1}^M i_m \mathbf{o}_m(\mathbf{r}) = \sum_{m=1}^M i_m \left\{ \sum_{n=1}^N M_{m,n} \mathbf{f}_n(\mathbf{r}) \right\}, \quad (3.32)$$

where i_m is defined as the *mesh* current circulating around loop m , which is a combination of several *branch* currents, i_n :

$$i_m = \sum_{n=1}^N M_{m,n} i_n. \quad (3.33)$$

The *mesh* currents are stored within the vector, I_{mesh} , and can be computed from the current vector, I_{branch} :

$$I_{mesh} = M I_{branch}. \quad (3.34)$$

Each row of M represents a single VL basis function. The column index of M determines which currents from I_{branch} , i.e. SWG functions $\mathbf{f}_n(\mathbf{r})$, form part of the VL basis function.

Equation (3.26) can now be transformed as follows:

$$MZI_{branch} = MV_{branch}, \quad (3.35)$$

and replacing I_{branch} with I_{mesh} ,

$$(MZM^T)I_{mesh} = V_{mesh}. \quad (3.36)$$

The vector, V_{mesh} , contains the voltages across each VL basis function and is defined as,

$$V_{mesh} = MV_{branch}. \quad (3.37)$$

It is shown in [225] that the values of the vector V_{mesh} will become zero for closed VL basis functions and will be equal to the voltage difference across the ends of unclosed VL basis functions:

$$(V_{mesh})_m = \begin{cases} 0, & \text{for closed loop } m \\ \phi(\xi)|_{\xi \in A_a} - \phi(\xi)|_{\xi \in A_b}, & \text{for unclosed loop } m \end{cases}. \quad (3.38)$$

The functions $\phi(\xi)|_{\xi \in A_a}$ and $\phi(\xi)|_{\xi \in A_b}$ represent the constant voltage potential across the two faces at the ends of an unclosed loop, with area A_a and A_b , respectively.

Figure 3.28 illustrates the setup of the VL basis functions within a rectangular conductor, with a voltage source connected to two terminals. The points represent edges viewed from above and the circles represent closed VL basis functions. Since displacement current is assumed negligible, current will not flow across the boundary and SWG basis functions are not required for boundary faces. However, the faces connected to the terminals require SWG basis functions, since current can flow across these faces. Closed VL basis functions around the terminal edges ensure that the terminal faces are shorted electrically. An unclosed VL basis function is defined between the two terminals and represents the voltage difference between the two terminals, as shown in (3.38).




Figure 3.28: Top view of tetrahedral mesh of rectangular conductor with two terminals.

3.3.7.6 Electrostatic Analogy

As demonstrated in [164], the FMM can be used to evaluate the matrix-vector product, ZI_{branch} , without explicitly forming Z . The matrix-vector product, ZI_{branch} , can be separated into a real and imaginary part:

$$ZI_{branch} = RI_{branch} + j\omega LI_{branch}. \quad (3.39)$$

The evaluation of RI_{branch} is not computationally expensive, since R is a sparse matrix. However, L is a dense matrix and LI_{branch} is computationally expensive to compute directly. Using the electrostatic analogy, it is possible to compute LI_{branch} by evaluating the electrostatic potential, produced by the surrounding charges, at each tetrahedron [164]. Each entry of the matrix-vector product, LI_{branch} , can be evaluated as follows:

$$(LI_{branch})_m = \sum_{n=1}^N \left(\frac{\mu}{4\pi} \int_{T_m^\pm} \int_{T_n^\pm} \frac{\mathbf{f}_m(\mathbf{r}) \cdot \mathbf{f}_n(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} dv' dv, \right) i_n, \quad (3.40)$$

where i_n is the *branch* current through face n , i.e. the coefficient in (3.25). Note that the integration in (3.40) is performed over the tetrahedrons T_m^\pm and T_n^\pm , and not over the volumes v_m and v_n , as given in (3.28). Equation (3.40) can also be written in terms of the magnetic vector potential, $\mathbf{A}(\mathbf{r})$,

$$(LI_{branch})_m = \int_{T_m^\pm} \mathbf{f}_m(\mathbf{r}) \cdot \mathbf{A}(\mathbf{r}) dv \quad (3.41)$$

where

$$\begin{aligned} \mathbf{A}(\mathbf{r}) &= \frac{\mu}{4\pi} \sum_{n=1}^N \left(\int_{T_n^\pm} \frac{\mathbf{f}_n(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} dv' \right) i_n \\ &= \frac{\mu}{4\pi} \sum_{n=1}^N \left(\int_{T_n^\pm} \frac{\rho_n^\pm}{|\mathbf{r} - \mathbf{r}'|} dv' \right) \frac{i_n}{3|v_n^\pm|}. \end{aligned} \quad (3.42)$$

This decomposition shows that $(LI_{branch})_m$ can be evaluated by integrating the magnetic vector potential, $\mathbf{A}(\mathbf{r})$, over each tetrahedron. The vector potential can be decomposed into its x -, y -, and z -components. Each component can be considered a scalar electrostatic potential generated by a collection of charges [164]:

$$\psi_p(\mathbf{r}) = \frac{\mu}{4\pi} \sum_{n=1}^N \left(\int_{T_n^\pm} \frac{(\rho_n^\pm)_p}{|\mathbf{r} - \mathbf{r}'|} dv' \right) \frac{i_n}{3|v_n^\pm|}, \quad (3.43)$$

where $p \in \{1, 2, 3\}$ and the scalar potential, $\psi_p(\mathbf{r})$, denotes the p th component of $\mathbf{A}(\mathbf{r})$. The product $(i_n/3|v_n^\pm|) (\rho_n^\pm)_p$ can be interpreted as the charge density within T_n^\pm .

Equation (3.43) can easily be accelerated using the FMM, since it involves the evaluation of electrostatic potential at tetrahedron m due to accumulative effect of n charges. Using the FMM, the matrix-vector product, LI_{branch} , can be computed in $O(m)$ operations [164].

3.3.7.7 Iterative Solver and Preconditioning

Equation (3.36) is solved iteratively using the GMRES method [203]. The matrix-vector product is the costliest step of the GMRES and is accelerated using the FMM, as discussed

in section 3.3.7.6. The convergence rate of the iterative method is reduced using a right preconditioned linear system [205]:

$$(MZM^T)Px' = V_{mesh}, \quad (3.44)$$

where P is the preconditioning matrix. The aim is to make P^{-1} as close as possible to MZM^T . Instead of sparsifying MZM^T , a better preconditioner is formed when only the matrix Z is sparsified. This approach is less computationally expensive and has shown to be effective [164]. The preconditioning matrix is then formed using ILU factorization [213]:

$$P^{-1} = M(Z_{sparse})M^T \approx LU. \quad (3.45)$$

The matrices L and U are the lower and upper triangular matrices, respectively. Two sparse formats for Z_{sparse} are evaluated: using the diagonal values of Z , referred to as *Diagonal-L*, or using the non-zero pattern of R , referred to as *Pattern-R*. Figure 3.29a shows the convergence of the GMRES for the superconducting microstrip line example in Figure 3.30, when applying no preconditioner, the *Diagonal-L* and the *Pattern-R* preconditioners. Figure 3.29b shows the convergence of the multi-layer example in Figure 3.31. From Figure 3.29a and 3.29b it is evident that *Diagonal-L* preconditioning accelerates the convergence of the GMRES method, compared to the linear system with no preconditioning. Constructing the *Pattern-R* preconditioner is more computationally expensive compared to the *Diagonal-L* preconditioner, but it delivers much faster convergence and reduces overall calculation time.




Figure 3.29: Convergence rate of GMRES for (a) the microstrip line example and (b) the multi-layer example.

3.3.7.8 Results: Small Superconducting Structures

To evaluate the efficiency and performance of TTH, several test structures were simulated with TTH and the results were compared to Fast FastHenry (FFH) [209]. Figure 3.30 and 3.31 show the current density calculated with TTH for a microstrip line and a multilayer structure, respectively. The geometry in Figure 3.31 was generated using InductEx. Table 3.5 shows the performance of TTH compared to FFH. The unknowns represent the number of SWG functions in TTH and filaments in FFH. Extracted values correspond

with both FFH and the method used in [196] with less than 1% error. The running time and memory usage of TTH are lower than FFH; however, the number of SWG functions necessary for the same level of accuracy can be higher compared to FFH. For example, the structure in Figure 3.31 does not model the London penetration depth accurately and, therefore, requires a lower discretization size compared to the structure in Figure 3.30. In order to lower the number of unknowns, i.e. increase the discretization size, the number of layers near the edges must be increased.




Figure 3.30: Current density of a $5\ \mu\text{m} \times 50\ \mu\text{m}$ microstrip line (thickness = 220 nm and penetration depth = 137 nm) 177.5 nm above ground layer (overhang = $6\ \mu\text{m}$, thickness = 300 nm, and penetration depth = 86 nm). Note: segment size and height division is for illustration purposes only.




Figure 3.31: Current density of a multilayer example with coupled structures. Penetration depth is 90 nm and thicknesses are respectively 200 nm, 250 nm and 350 nm for top, middle and ground layers. Ground overhang is $5\ \mu\text{m}$.

3.3.8 Sheet currents, triangular and hybrid meshes

Tetrahedral meshes are non-uniform, which removes the strict meshing requirements that cake slicing imposes on cuboid meshes where the width of a narrow segment must be perpetuated along all segments on the same axis. From my earliest attempts to analyse chip-scale structures it was clear that cuboid meshes were ill-suited to large ground

Table 3.5: Performance comparison between TTH and FFH. Bench-marked performed on a Intel Core i7-3612QM @2.1 GHz, running Windows 8.1.

Layout model	Unknowns	Inductance	CPU Time	Memory
Strip line (TTH)	121104	4.421 pH	84 s	1.51 GB
Strip line (FFH)	119824	4.426 pH	87 s	2.92 GB
Multilayer (TTH)	108404	1.471 pH	67 s	1.45 GB
Multilayer (FFH)	105655	1.461 pH	144 s	2.71 GB

planes with small features. When we introduced tetrahedral meshing, chip-scale modelling became tractable, but it was soon obvious that most tetrahedra are used to model vast expanses of thin film conductors where triangles with sheet currents can be even more efficient.

In almost every application of superconductor integrated circuit analysis, the circuits are constructed with thin superconducting films. If the thickness of the superconducting films are on the same order as the London penetration depth, the three-dimensional volume current density can be restricted to two dimensions. This is also known as the sheet current model, which has proven to be efficient for simulating the current density in multilayer superconductor films [185], [194], [226]–[228].

Meshing thin superconducting films with two dimensional triangular elements, instead of tetrahedral elements, significantly reduces the number of unknowns. Modelling the current density inside a cuboid conductor requires at least six tetrahedral elements, whereas the same cuboid can be modelled with only two triangular elements. Furthermore, each tetrahedron requires four SWG basis functions, one for each face, whereas a triangle requires only three Rao-Wilton-Glisson (RWG) basis functions [229], one for each edge. Theoretically, the number of unknowns can be reduced by a factor of $\frac{6 \times 4}{2 \times 3} = 4$, if triangular meshing is used instead of tetrahedral meshing. However, triangular meshing is limited to sheet current models; therefore, only practical for simulating thin superconducting films with finite thickness. The thickness of these films can be modelled with the special Green's functions defined in [185], [226], [227].

3.3.8.1 Derivation of Surface Integral Equation

The following can be assumed for a large class of thin film superconductor circuits [194]:

$$t_m \ll l \quad \text{and} \quad \lambda_m \sim t_m, \quad (3.46)$$

where t_m and λ_m are respectively the thickness and penetration depth of film m . The value l is the size of the circuit in the x, y -plane. If it is assumed that $J_z(r) = 0$, the volume current density, $\mathbf{J}(\mathbf{r})$, can be reduced to a sheet current density in the x, y -plane [226],

$$\mathbf{J}_m^s(\mathbf{r}) = \int_{h_m^0}^{h_m^1} \mathbf{J}(\mathbf{r}) dz, \quad (3.47)$$

where h_m^0 and h_m^1 is respectively the bottom and top z -coordinates of layer m . Taking the average of the current density over the height of film layer m ,

$$\mathbf{J}_m^s(\mathbf{r}) = t_m \mathbf{J}(\mathbf{r}), \quad (3.48)$$

the integral equation, (3.22), can be written in terms of sheet currents,

$$\frac{\mathbf{J}_{\mathbf{m}}^{\mathbf{s}}(\mathbf{r})}{t_m k(\mathbf{r})} + \frac{j\omega\mu}{t_m 4\pi} \int_{S'_m} \mathbf{J}_{\mathbf{m}}^{\mathbf{s}}(\mathbf{r}') G_0(\mathbf{r}, \mathbf{r}') ds' = -\nabla\phi(\mathbf{r}), \quad (3.49)$$

where $\mathbf{J}_{\mathbf{m}}^{\mathbf{s}}(\mathbf{r})$ is the sheet current and t_m is the thickness of film layer m . In the case of thin film superconducting materials with thickness,

$$t_m \ll \lambda_m, \quad (3.50)$$

the product $t_m k(\mathbf{r})$ effectively replaces the penetration depth, λ , with the perpendicular penetration depth [226], [230]:

$$\lambda_{\perp} = \frac{\lambda_m^2}{t_m}. \quad (3.51)$$

For single-layer problems, the free-space Green's function $G_0(\mathbf{r}, \mathbf{r}')$ can be used:

$$G_0(\mathbf{r}, \mathbf{r}') = \frac{1}{|\mathbf{r} - \mathbf{r}'|}. \quad (3.52)$$

However, for multi-layered films with finite thickness, the current density above and below the films have to be taken into account [226]. Figure 3.32 demonstrates the top and bottom surfaces of layer m , with a normal vector pointing in the z -direction. The integration is done over the two projected triangles parallel to T_m^+ , at heights $z = h_m^0$ and $z = h_m^1$. The Green's function for the interacting films m and n can be calculated as,

$$G_{m,n}(\mathbf{r}, \mathbf{r}') = \frac{1}{4} \sum_{k=0}^1 \sum_{l=0}^1 \left\{ \left\| \left(\mathbf{r} + \mathbf{n}_m \frac{t_m}{2} (-1)^k \right) - \left(\mathbf{r}' + \mathbf{n}_n \frac{t_n}{2} (-1)^l \right) \right\| \right\}^{-1}, \quad (3.53)$$

where \mathbf{n}_m and \mathbf{n}_n are the unit normal vectors of layers m and n , respectively. The Green's function in (3.53) is similar to the one used in [226] for finite thickness films.




Figure 3.32: Visualisation of the sheet current model for triangle T_m^+ , with projected triangles at heights h_m^0 and h_m^1 .

3.3.8.2 Discretization

As discussed in Section 3.3.7.3, a system of linear equations can be obtained from the integral equation, (3.49), using the Method of Moments [222]. The thin superconducting film is discretized using triangular elements, instead of tetrahedral elements. It is assumed that the electrical parameters in each triangle are constant.

The integral equation in (3.49) is discretized using the RWG basis function [229]. Figure 3.32 shows the definition of the RWG basis function. The two triangles, T_n^+ and T_n^- , are associated with the n th edge of the discretized region. The position vectors, ρ_n^+

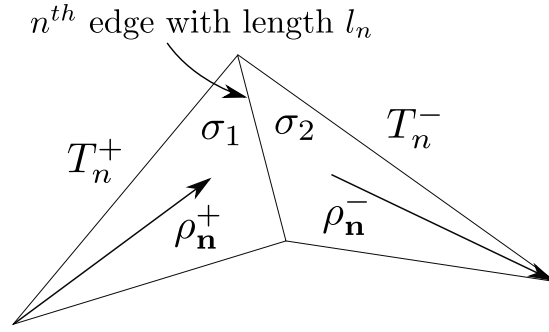


Figure 3.33: RWG basis function at material interface with different conductivities.

and ρ_n^- , represent points in T_n^+ and T_n^- , respectively. In triangle T_n^+ , the positive position vector, ρ_n^+ , is defined with respect to (thus from) the free vertex. The negative position vector, ρ_n^- in triangle T_n^- , is defined towards the free vertex [229]. The signs of the two triangles depend on the direction of current flow through edge n .

To simplify the problem, the entire problem domain is assumed to be a piecewise homogeneous dielectric body. Although the RWG basis function is defined for infinitely thin triangles, it is assumed to have finite thickness. The RWG function is defined as:

$$\mathbf{f}_n^s(\mathbf{r}) = \begin{cases} \frac{1}{2|a_n^+|} \rho_n^+(\mathbf{r}), & \text{if } \mathbf{r} \in T_n^+ \\ \frac{1}{2|a_n^-|} \rho_n^-(\mathbf{r}), & \text{if } \mathbf{r} \in T_n^- \\ 0, & \text{otherwise,} \end{cases} \quad (3.54)$$

where $|a_n^\pm|$ is the area of T_n^\pm . This function differs from the basis functions used in [229], which uses the length of the face to normalize $\mathbf{f}_n^s(\mathbf{r})$. Using the RWG function, the sheet current density, $\mathbf{J}^s(\mathbf{r})$, can be expanded as follows:

$$\mathbf{J}^s(\mathbf{r}) = \sum_{n=1}^N i_n \mathbf{f}_n^s(\mathbf{r}), \quad (3.55)$$

where N is the number of edges that make up the entire surface domain and i_n is the *branch* current through the n th edge.

The integral equation in (3.49) can be solved with the Method of Moments, as described in Section 3.3.7.4. Using the RWG function as weighting functions, $\mathbf{f}_m^s(\mathbf{r})$, a system of N linear equations can be obtained:

$$Z^s I_{branch} = V_{branch}. \quad (3.56)$$

The sheet impedance matrix Z^s can be decomposed into its real and imaginary parts:

$$Z^s = R^s + j\omega L^s. \quad (3.57)$$

where R^s and L^s are respectively the sheet resistance and inductance matrices. The entries of the sheet resistance matrix are computed as follow:

$$R_{m,n}^s = \int_{s_m} \frac{1}{t_m k_m} \mathbf{f}_m^s(\mathbf{r}) \cdot \mathbf{f}_n^s(\mathbf{r}) ds, \quad (3.58)$$

and the entries of the sheet inductance matrix:

$$L_{m,n}^s = \frac{\mu}{4\pi} \int_{s_m} \int_{s_n} \frac{1}{t_m t_n} \mathbf{f}_m^s(\mathbf{r}) \cdot \mathbf{f}_n^s(\mathbf{r}') G_{m,n}(\mathbf{r}, \mathbf{r}') ds' ds, \quad (3.59)$$

where t_m and t_n are the thicknesses of surfaces s_m and s_n . The value k_m is calculated from (3.23),

$$k_m = k(\mathbf{r}), \quad r \in s_m. \quad (3.60)$$

The values $R_{m,n}$ and $L_{m,n}$ correspond to the RWG basis functions m and n , respectively. The surfaces s_m and s_n represent the surfaces of the RWG-basis functions, which are a combination of $(T_m^+ + T_m^-)$ and $(T_n^+ + T_n^-)$, respectively. The voltage over each edge is stored in the vector, V_{branch} , and can be computed as follow:

$$(V_{branch})_m = - \int_{s_m} \mathbf{f}_m^s(\mathbf{r}) \cdot \nabla \phi(\mathbf{r}) \, ds \quad (3.61)$$

3.3.8.3 Surface loop basis function

Similar to the SWG basis function, discussed in Section 3.3.7.5, the divergence of the RWG basis function is also non-zero [229]. In order to ensure the divergence free condition, a surface loop (SL) basis function is used to discretize the integral equation in (3.49). This SL basis function is similar to the VL basis function described in Section 3.3.7.5. Figures 3.34a and 3.34b illustrate how closed and unclosed SL basis functions are constructed around nodes (vertices).




Figure 3.34: (a) Closed surface loop basis function. (b) Unclosed surface loop basis function with two boundary edges of lengths l_a and l_b .

Following the same approach described in Section 3.3.7.5, the SL basis function can be defined as a combination of RWG functions around node m :

$$\mathbf{o}_m^s(\mathbf{r}) = \sum_{n=1}^N M_{m,n} \mathbf{f}_n^s(\mathbf{r}). \quad (3.62)$$

The value of $M_{m,n}$ is either 0 or ± 1 , depending on the direction of the RWG function n in loop m . The value N is the total number of edges on the surface domain. Using the SL basis function, the sheet current can be expanded as follows:

$$\mathbf{J}^s(\mathbf{r}) = \sum_{m=1}^M i_m \mathbf{o}_m^s(\mathbf{r}) = \sum_{m=1}^M i_m \left\{ \sum_{n=1}^N M_{m,n} \mathbf{f}_n^s(\mathbf{r}) \right\}, \quad (3.63)$$

where i_m is defined as the *mesh* current circulating around node m . Once again, the MoM is used to obtain a matrix equation, see Section 3.3.7.5, using SL basis functions:

$$(MZ^s M^T) I_{mesh} = V_{mesh}. \quad (3.64)$$

It can be shown that the values of V_{mesh} will become zero for closed SL basis functions and will be equal to the voltage difference across the ends of an unclosed SL basis function:

$$V_m = \begin{cases} 0, & \text{for closed loop } m \\ \phi(\xi)|_{\xi \in l_a} - \phi(\xi)|_{\xi \in l_b}, & \text{for unclosed loop } m \end{cases} \quad (3.65)$$

The functions, $\phi(\xi)|_{\xi \in l_a}$ and $\phi(\xi)|_{\xi \in l_b}$, represent the constant voltage potential across the two edges at the ends of an unclosed SL basis function, with lengths l_a and l_b , respectively.

3.3.8.4 Hybrid Meshing

In order to provide the efficiency of triangular segments while retaining the volume discretization of tetrahedra where vias or thick-film structures co-exist with thin-film conductors, Dr Jackman devised a hybrid meshing capability that was implemented in TTH. I subsequently added layer-specific and layout object-specific selection of meshing method to InductEx so that a model can use a combination of triangular and tetrahedral segments.

Hybrid meshing can be used to improve calculation speed, by modeling thin superconductor films with triangles and complex via-interconnects with tetrahedrons.

To use triangles and tetrahedrons simultaneously, both the volume loop (3.31) and surface loop (3.62) basis functions are implemented. Hybrid loop basis functions are used at the interface that connects triangles with tetrahedrons, as shown in Figure 3.35. This hybrid loop basis function consists of both SWG and RWG functions, depending on the type of element (triangle or tetrahedron) in the loop. The single integral equations, (3.27) and (3.58), remain the same; whereas the double integral equations, (3.28) and (3.59), are a combination of triangular and tetrahedral elements.




Figure 3.35: Hybrid loop basis function, consisting of both RWG and SWG basis functions.

An example of a hybrid mesh is shown in Figure 3.36. The microstrip line is connected to the ground layer through a via-interconnect. The dimensions are the same as the microstrip in Figure 3.30. Triangular meshing is used for both the microstrip line and the ground layer; whereas the via is meshed with tetrahedrons. The inductance, as a function of the number of height layers, is shown in Figure 3.37. Two types of hybrid loop function are evaluated: the sheet current of each triangle enters a single face in the tetrahedral mesh or multiple faces in the tetrahedral mesh. If each surface is connected to a single tetrahedral face, the inductance of the hybrid mesh is higher (1.6% error), compared to the tetrahedral method. This is expected, since the area through which the current can flow is reduced. If each surface is connected to multiple tetrahedral faces, the area of the interface is increased and the extracted inductance matches the tetrahedral method with less than 0.5% error, as can be seen in Figure 3.37.




Figure 3.36: Current density of a $50 \mu\text{m} \times 5 \mu\text{m}$ microstrip line (triangular meshing) with a via-interconnect (tetrahedral meshing).




Figure 3.37: The inductance of a microstrip line, with a via-interconnect, for different meshing techniques.

3.3.9 Coupling from flux trapped in holes

With a new, fast engine that supports cuboid, tetrahedral and triangular segments, or any hybrid combination of these for the calculation of current distribution and the extraction of integrated circuit parameters, my attention turned to the remaining questions in superconductor integrated circuit layout verification in the electromagnetic environment:

- What is the effect of a trapped fluxon on circuit operation?
- How does an external magnetic field couple to (and affect) circuit operation?

For the analysis of flux trapped in holes, or moats, TetraHenry required support for the inclusion of holes as ports. Dr Kyle Jackman developed and implemented this functionality as part of his PhD dissertation [215].

Although flux trapping holes can exist in any axial direction (flux can also be trapped in the hole formed when two superconducting vias connect two superconducting planes), the analysis of flux trapping is usually confined to moats in the ground and sky planes – thus in the plane of the IC surface. Even so, the analysis method is generic and can be applied to any hole.

A path is defined through each hole under investigation, where each path is closed *outside* the furthest boundaries of the mesh as shown in Figure 3.38. Every VL or SL basis function in the mesh that encloses a hole is identified by finding if it closes through the surface of the closed hole path. The VL or SL basis functions for a hole then form a cycle that describes the hole inductance in series with a voltage excitation port.

I automated the setup in InductEx so that users only have to mark a hole with a label on the layout, after which InductEx handles path creation, port excitation and inductance extraction.




Figure 3.38: Definition of paths for every hole in an extraction model.

Once the incorporation of hole ports in TetraHenry was completed, it became possible to integrate holes into extraction models.

In order to add the effect of flux in holes to a circuit simulation, every hole is treated as an inductor coupled to every other inductor – including other holes – in a circuit layout. During the MQS solution of current distribution, hole ports are excited with $1 \angle 0^\circ$ V just as any other voltage port in the system. For inductance extraction, this means that the holes can be added to the KVL equations of the full system (as discussed in

Section 3.3.3) by adding one cycle or loop for every hole, and appropriately expanding the current matrix \mathbf{A} and unknown component value vector \mathbf{x} .

As an example, consider the circuit in Figure 3.39 with a circuit and two holes (or moats).

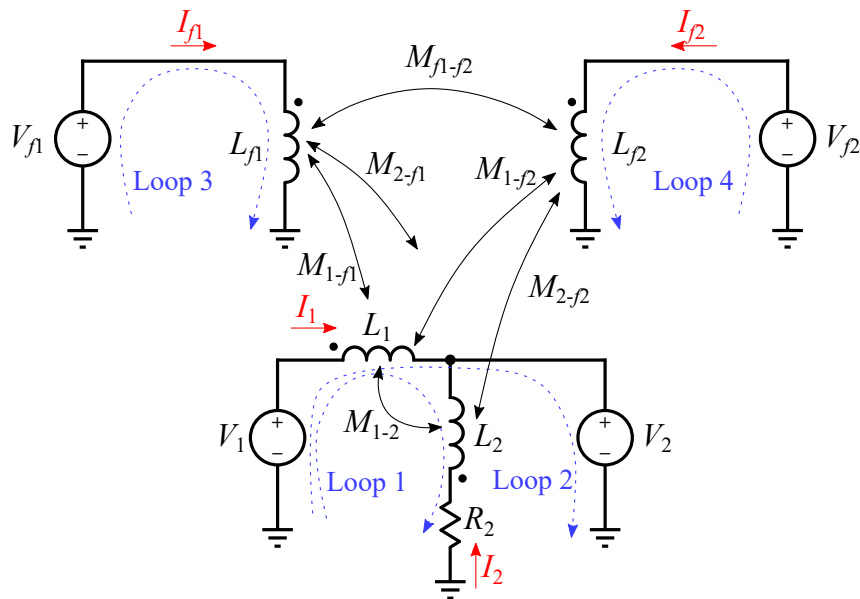


Figure 3.39: A circuit model for two holes coupled to a circuit with inductance, resistance and mutual inductance.

The KVL equations are written in matrix form, with:

$$\mathbf{b} = \begin{bmatrix} V_{1r} \\ V_{1i} \\ V_{1r} - V_{2r} \\ V_{1i} - V_{2i} \\ V_{f1r} \\ V_{f1i} \\ V_{f2r} \\ V_{f2i} \end{bmatrix}$$

$$\mathbf{A} = \begin{bmatrix} -I_{1i} & -I_{2r} & I_{2i} & 0 & 0 & (-I_{2i} + I_{1i}) & -I_{f1i} & -I_{f2i} & I_{f1i} & I_{f2i} & 0 \\ I_{1r} & -I_{2i} & -I_{2r} & 0 & 0 & (I_{2r} - I_{1r}) & I_{f1r} & I_{f2r} & -I_{f1r} & -I_{f2r} & 0 \\ -I_{1i} & 0 & 0 & 0 & 0 & -I_{2i} & -I_{f1i} & -I_{f2i} & 0 & 0 & 0 \\ I_{1r} & 0 & 0 & 0 & 0 & I_{2r} & I_{f1r} & I_{f2r} & 0 & 0 & 0 \\ 0 & 0 & 0 & -I_{f1i} & 0 & 0 & -I_{1i} & 0 & -I_{2i} & 0 & -I_{f2i} \\ 0 & 0 & 0 & I_{f1r} & 0 & 0 & I_{1r} & 0 & I_{2r} & 0 & I_{f2r} \\ 0 & 0 & 0 & 0 & -I_{f2i} & 0 & 0 & -I_{1i} & 0 & -I_{2i} & -I_{f1i} \\ 0 & 0 & 0 & 0 & I_{f2r} & 0 & 0 & I_{1r} & 0 & I_{2r} & I_{f1r} \end{bmatrix}$$

$$\mathbf{x} = \begin{bmatrix} \omega L_1 \\ R_2 \\ \omega L_2 \\ \omega L_{f1} \\ \omega L_{f2} \\ \omega M_{1-2} \\ \omega M_{1-f1} \\ \omega M_{1-f2} \\ \omega M_{2-f1} \\ \omega M_{2-f2} \\ \omega M_{f1-f2} \end{bmatrix}$$

The solution of \mathbf{x} is done exactly as in Section 3.3.3.

3.3.10 Coupling from external fields

As it became evident that InductEx could be used for more than just inductance extraction, my group and I began to investigate the effects of external magnetic fields on circuit operation. An external field induces current in circuit loops – even at dc for superconducting loops. All we had to do was to find a way to model and analyse this coupling.

Initially, we used large coils in three dimensions for which I built in modelling support in InductEx. These coils were handled just as other circuit inductors, and the coupling to designated circuit inductors could then be calculated with InductEx. In an electrical simulation, the coils would then be driven with an appropriately-sized current source to apply a fairly uniform field to a circuit under test. The methodology did not differ much from how magnetic fields were applied in laboratory test cases. My PhD student, Dr. Rodwell Bakolo, investigated circuits and magnetic field operating margins in this way [231], [232] and developed on-chip shielding techniques.

For more accurate modelling, and especially for compact simulation model extraction, Dr Kyle Jackman added the ability to excite uniform vector magnetic fields in FFH and TetraHenry. This allowed me to add automatic external magnetic field analysis to any circuit model extracted with InductEx, and then to calculate the coupling from these fields to all branches in a circuit model. Artificial coils were no longer needed.

For the MQS solution, a constant field vector with a defined magnitude and direction is applied and the current density in every mesh element calculated when all voltage ports are zeroed. For simulation purposes in JoSIM or a similarly capable SPICE engine, the external magnetic field is modelled as a current source driving a field inductance. We can select the current source amplitude to have any ratio of current in ampere to flux density in tesla. For simplicity, the amplitude is chosen as 1 A T^{-1} and the field inductance is set to 1 H. This inductance is then coupled with mutual inductance to every other inductor in the target circuit – including all flux trapping sites.

A circuit schematic that shows the inclusion of one external magnetic field vector modelled as a current source driving a field inductor is shown in Figure 3.39.

The current source prohibits currents in the circuit from inducing current in the magnetic field loop during simulation.




Figure 3.40: A circuit model for a moat and an external magnetic field coupled to a circuit with inductance, resistance and mutual inductance.

Since the field inductance L_{field} is known, and the voltage over L_{field} is solely determined by the field current I_{field} , calculation of the field coupling mutual inductances (such as $M_{1-field}$ and $M_{2-field}$) cannot be done with the same set of equations with which circuit inductances, resistances and mutual inductances between circuit elements are calculated.

Rather, a second set of equations is required after the circuit component values are calculated as discussed in Section 3.3.3 and Section 3.3.9. Thus, after solution of L_1 , L_2 , R_2 , L_{f1} , M_{1-2} , M_{1-f1} and M_{2-f1} , we turn to an MQS solution of the circuit currents I_1 , I_2 and I_{f1} when an external field with a defined flux density (1 T) is applied while all voltage sources are zeroed.

The KVL equations for three voltage loops are now:

$$\begin{aligned}
 0 &= I_1 j\omega L_1 - I_2 (R_2 + j\omega L_2) + I_2 j\omega M_{12} - I_1 j\omega M_{12} + I_{f1} j\omega M_{1-f1} \\
 &\quad - I_{f1} j\omega M_{2-f1} + I_{field} j\omega M_{1-field} - I_{field} j\omega M_{2-field} \\
 0 &= I_1 j\omega L_1 + I_2 j\omega M_{12} + I_{f1} j\omega M_{1-f1} + I_{field} j\omega M_{1-field} \\
 0 &= I_{f1} j\omega L_{f1} + I_1 j\omega M_{1-f1} + I_2 j\omega M_{2-f1} + I_{field} j\omega M_{f1-field}
 \end{aligned}$$

We already selected $I_{field} = 1$ A, so that the KVL equations can be separated into known values on the left hand side and unknown values on the right hand side as:

$$\begin{aligned}
 I_1 j\omega L_1 - I_2 (R_2 + j\omega L_2) + I_2 j\omega M_{12} - I_1 j\omega M_{12} \\
 + I_{f1} j\omega M_{1-f1} - I_{f1} j\omega M_{2-f1} &= -j\omega M_{1-field} + j\omega M_{2-field} \\
 I_1 j\omega L_1 + I_2 j\omega M_{12} + I_{f1} j\omega M_{1-f1} &= -j\omega M_{1-field} \\
 I_{f1} j\omega L_{f1} + I_1 j\omega M_{1-f1} + I_2 j\omega M_{2-f1} &= -j\omega M_{f1-field}
 \end{aligned}$$

The right hand side is purely imaginary, so that

$$\begin{aligned}
 I_{1r}\omega L_1 - I_{2i}R_2 - I_{2r}\omega L_2 + I_{2r}\omega M_{12} - I_{1r}\omega M_{12} \\
 + I_{f1r}\omega M_{1-f1} - I_{f1r}\omega M_{2-f1} &= -\omega M_{1-field} + \omega M_{2-field} \\
 I_{1r}\omega L_1 + I_{2r}\omega M_{12} + I_{f1r}\omega M_{1-f1} &= -\omega M_{1-field} \\
 I_{f1r}\omega L_{f1} + I_{1r}\omega M_{1-f1} + I_{2r}\omega M_{2-f1} &= -\omega M_{f1-field}
 \end{aligned}$$

It is thus algorithmically easy to solve the field coupling values. A vector \mathbf{b} is constructed with the imaginary voltage around every loop divided by frequency, excluding the effect of field coupling, as the row entries:

$$\mathbf{b} = \begin{bmatrix} I_{1r}L_1 - I_{2i}\frac{R_2}{\omega} - I_{2r}L_2 + I_{2r}M_{12} - I_{1r}M_{12} + I_{f1r}M_{1-f1} - I_{f1r}M_{2-f1} \\ I_{1r}L_1 + I_{2r}M_{12} + I_{f1r}M_{1-f1} \\ I_{f1r}L_{f1} + I_{1r}M_{1-f1} + I_{2r}M_{2-f1} \end{bmatrix}$$

The vector of unknown values contains only the field couplings:

$$\mathbf{x} = \begin{bmatrix} M_{1-field} \\ M_{2-field} \\ M_{f1-field} \end{bmatrix}$$

The \mathbf{A} matrix has a signed 1 if a field coupling acts on in a specific voltage loop – meaning the inductor on which the field coupling acts is in the loop. The sign has the opposite polarity as that of the inductor in the loop. For this example:

$$\mathbf{A} = \begin{bmatrix} -1 & 1 & 0 \\ -1 & 0 & 0 \\ 0 & 0 & -1 \end{bmatrix}$$

For complete treatment of an external magnetic field, the field is represented by the orthogonal components in the x , y and z directions and solved for each. The resulting extracted circuit thus has three extra inductors, L_{fieldx} , $L_{fielddy}$ and $L_{fielddz}$, each with inductance 1 H and each driven by a current source that represents the magnetic flux density in the respective axial direction at the circuit.

3.3.11 Compact simulation models

The culmination of all my research to date on inductance and parameter extraction methods and tools is the development of compact simulation models.

3.3.11.1 Errors in extracted results

Errors in calculated inductance values are introduced by perturbations to \mathbf{A} and \mathbf{b} in (3.15) (see Section 3.3.3) that arise from computer precision, mismatch between the inductors in a circuit netlist and the layout model, and from neglected coupling.

The use of double precision floating point numbers internally by InductEx and the TetraHenry engine keeps errors from computer precision small, so that these can be neglected.

The relative error

$$\|\mathbf{Ax} - \mathbf{b}\|/\|\mathbf{b}\| \quad (3.66)$$

can be calculated after solution of \mathbf{x} and evaluated to determine if the solution is reliable. InductEx calculates this relative error automatically and prints it as a percentage. There is no clear threshold as to what constitutes a sufficiently precise solution, but in general a worst-case error of more than 5 % is indicative of mismatch between the inductive branches in the circuit netlist and the layout, or considerable magnetic coupling in the circuit layout that has not been accounted for by the circuit netlist.

Mismatch between the inductive branches of a circuit schematic drawn up by a circuit designer and the actual physical layout can be difficult to identify, but an experienced circuit modeller can almost always construct a schematic that fits a layout well. A mismatched model results in a large condition number of \mathbf{A} . The condition number provides the ratio between the largest and smallest singular values of \mathbf{A} , and gives an indication of the largest error in precision. For every factor of ten increase in the condition number, one digit of precision in the solution is compromised. In general, the solution from a system with a condition number that exceeds 1×10^6 should be distrusted and the schematic circuit netlist should be inspected.

In cases where there are not enough ports to allow all branch impedances to be extracted, \mathbf{A} becomes rank deficient. In this case the minimum norm solution is often mathematically correct but physically impossible, so that negative inductances or coupling factors larger than 1 can be obtained. In general the solution from a rank deficient system should be disregarded, and the circuit model revisited to provide more ports.

The most difficult error to correct arises from the neglect of coupling between inductors. It is very easy for a circuit designer to unintentionally neglect a mutual inductance in a circuit netlist and thereby introduce significant errors into calculated inductance results. Circuits that have multiple mutual inductors in a tight layout, such as typical AQFP cells, are especially vulnerable to errors introduced by neglected coupling. It may seem intuitive to add all possible coupling to the circuit schematic (so that every inductor is coupled to every other inductor), but in a circuit netlist with more than one inductor connected galvanically to the same subnet the result is almost always an underdetermined system. I showed this for an AQFP buffer circuit [233] (see Figures 2.67 and 2.68), where there are 21 mutual inductances if coupling between every inductor pair is included in the circuit netlist. The system is rank deficient and the minimum norm SVD solution returns some negative inductances that makes calculation of the mutual inductances linked to those inductors impossible.

3.3.11.2 Fundamental cycles

For a circuit netlist to be an exact representation of the superconductor layout, the number of inductors should be equal to the number of *fundamental cycles* in the netlist graph. This is the *fundamental inductor set*, which represents the exact number of inductors required to extract a full inductance matrix, containing all the mutual inductances.

The fundamental cycles can be obtained by first constructing a spanning tree of the netlist graph and then identifying the chords of the spanning tree [234]. The chords are the edges (branches) of the graph that do not form part of the *spanning tree* [234]. Each fundamental cycle consists of a chord together with the path in the spanning tree connecting the endpoints of the chord. There are exactly $m - n + c$ fundamental cycles, where m is the number of edges, n is the number of vertices, and c is the

number of connected components. The fundamental cycles are linearly independent from the remaining cycles, because each fundamental cycle contains a unique chord. The fundamental inductor set can therefore be obtained by placing inductors only in the chord branches of the circuit netlist. This ensures that each fundamental cycle contains only one inductor.

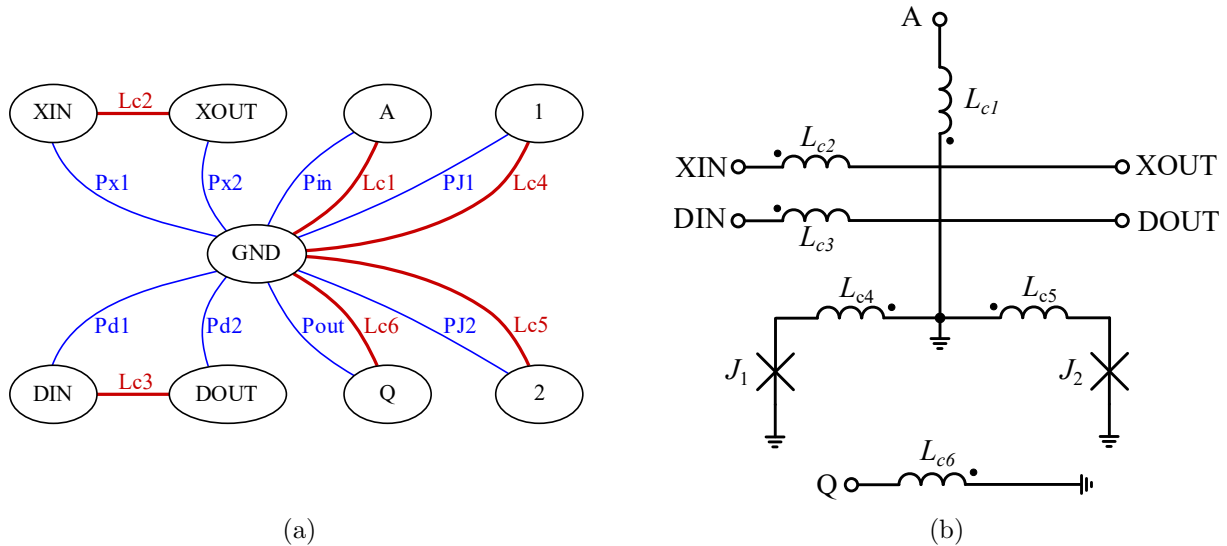


Figure 3.41: (a) Graph of AQFP buffer cell with red edges representing the *fundamental inductors*. (b) Schematic of the AQFP buffer cell with the *fundamental inductor set* (L_{c1} to L_{c6}).

Figure 3.41(a) shows the graph of the AQFP buffer cell from Figure 2.67 with the fundamental inductor set and Figure 3.41(b) shows the corresponding schematic.

Each fundamental cycle must contain at least one unique excitation port in order to calculate the mutual inductance between all the fundamental inductors. This can be accomplished by placing an excitation port in series with each chord. This ensures enough linear equations to calculate all the mutual inductances between all the fundamental inductors (L_{c1} to L_{c6} in Figure 3.41b).

The fundamental inductor set is mostly not intuitive. In the compact model here, for instance, there is no obvious inductor between L_{c4} and L_{c5} and ground to couple to the output inductor L_{c6} . The output stage coupling is fully modelled by the coupling between L_{c6} and *all the other inductors*. It is difficult to design and adjust the layout from the compact model, which is why compact model extraction is the final step before the simulation netlist is committed to a library.

3.4 Experimental verification

In order to verify the accuracy of inductance extraction, results were compared to experiments.

3.4.1 Measurement of inductance

When a dc SQUID, of which a basic circuit schematic is depicted in Figure 3.42, is biased with a current that exceeds the combined critical current of its two junctions, the SQUID

operates in the voltage mode. The average of the time-varying voltage developed over the SQUID, $\langle V \rangle$, from (2.28), is periodic with applied flux and has a periodicity of exactly Φ_0 .




Figure 3.42: Basic equivalent circuit of a dc SQUID.

It can be shown [235] that the phase difference over each junction, ϕ_1 and ϕ_2 , are related by

$$\phi_2 - \phi_1 = \frac{2\pi}{\Phi_0}(\Phi_a + LI_{circ}), \quad (3.67)$$

where L is the total loop inductance ($L_1 + L_2$) that includes both the magnetic and kinetic components of inductance, and Φ_a is the externally applied magnetic flux with flux density B through the effective SQUID loop area A_{eff} so that

$$\Phi_a = BA_{eff}. \quad (3.68)$$




Figure 3.43: Equivalent circuit of a dc SQUID used for inductance measurement on an integrated circuit. Modulation and control current directions are chosen to align with the circulating current in Figure 3.42.

Self and mutual inductance can be measured [236] when a dc SQUID is manufactured as shown in Figure 3.43. If modulation current I_{mod} is applied to the modulation input and extracted at the modulation output, a control current fed through the control pins, and the loop contains inductances L_{p1} and L_{p2} that are not in the path of the modulation current, then (3.67) expands to

$$\phi_2 - \phi_1 = \frac{2\pi}{\Phi_0}[\Phi_a + (L_1 + L_2 + L_{p1} + L_{p2})I_{circ} + (L_1 + L_2)I_{mod} + (M_1 + M_2)I_{ctrl}]. \quad (3.69)$$

Even if the critical currents of junctions J_1 and J_2 differ, and irrespective of the values of L_{p1} and L_{p2} , the circulating current I_{circ} is exactly the same at every 2π increment of $(\phi_2 - \phi_1)$ when a stable dc bias current I_b is applied. If the externally applied magnetic field Φ_a is kept constant – easily guaranteed with magnetic shielding – then a change in $(\phi_2 - \phi_1)$ of 2π rad is solely due to the inductances L_1 and L_2 , the mutual inductances M_1 and M_2 , and the change in modulation and control currents, so that

$$n\Phi_0 = (L_1 + L_2)\delta I_{mod} + (M_1 + M_2)\delta I_{ctrl}, \quad (3.70)$$

where n is the total number of voltage periods observed when either I_{mod} or I_{ctrl} is swept.

Self inductance is measured by changing I_{mod} while I_{ctrl} is kept constant (zero in practice), and mutual inductance is measured when I_{ctrl} is changed while I_{mod} is kept constant. Clearly, the control line can be omitted for experiments where only the self inductance is of interest.

A dc SQUID was manufactured with the Hypres 4.5 kA cm^{-2} process, with loop inductor in layer M2 and without a control line. A rendering of the three-dimensional extraction model generated by InductEx is shown in Figure 3.44. Cuboid segments are used for the FastHenry engine. With the process parameters provided in the Hypres design rules, InductEx calculates the series inductance of the junction arms, L_{p1} and L_{p2} as 0.14 pH each, and the total loop inductance (corresponding to $L_1 + L_2$ in Figure 3.43) as 11.1 pH .




Figure 3.44: InductEx model of an inductance measurement SQUID for the Hypres 4.5 kA cm^{-2} process with loop inductor in layer M2.

A measurement in liquid helium (performed with the assistance of Dr Olaf Wetzstein at IPHT in 2010), with modulation current swept from 0 to 1 mA at low frequency, is shown in Figure 3.45. The oscilloscope only recorded average voltage. The distance between the first and the sixth voltage peak is $895 \mu\text{A}$, so that from (3.70) for five periods we find:

$$L_{loop} = L_1 + L_2 = \frac{5\Phi_0}{895 \times 10^{-6}} = 11.6 \text{ pH}.$$

A more efficient design that uses a dc SQUID with two inductive branches to ground, and of which the modulation is a function of the difference between the inductance of each arm, has been presented [32]. A schematic diagram of such a differential-arm dc SQUID is shown in Figure 3.46. Such SQUID test structures were recently used to do an in-depth analysis of mutual inductance in SC integrated circuit structures with line widths down to $0.25 \mu\text{m}$ [125].

This design reuses the bias, modulation and control lines between SQUIDs and reduces the chip pads required per test cell – thereby allowing many more inductance tests to be




Figure 3.45: Measured SQUID voltage as a function of modulation current for a dc SQUID fabricated with the Hypres 4.5 kA cm^{-2} process with loop inductor in layer M2.




Figure 3.46: Equivalent circuit of a differential-arm dc SQUID used for inductance measurement on an integrated circuit.

included on a chip. In 2021, in collaboration with a research team from Synopsys under the IARPA SuperTools programme, Dr Kyle Jackman in my research group designed several differential-arm SQUID test structures to analyse self and mutual inductance in multiple layer configurations for the MITLL SFQ5ee process. Our test chip used six differential-arm SQUIDs strung together for every set of bias, modulation and control lines.

Voltage is still periodic with a periodicity of Φ_0 . In a constant magnetic field with a constant bias current, where n voltage periods are observed as a function of δI_{mod} and δI_{ctrl} , we have

$$n\Phi_0 = (L_1) \frac{\delta I_{mod}}{2m} + M\delta I_{ctrl}, \quad (3.71)$$

where m is the number of SQUIDs strung to the same modulation line, and the factor $1/2$ handles the equal division of modulation current between the left and right arms of the SQUID.

A measurement result of a small mutual inductance is shown in Figure 3.47. The zero voltage regions occur when the SQUID exits the voltage mode due to the low bias point. It does not affect the periodicity. This measurement was done for us in 2022 at NIST in Boulder, Colorado, by Dr Adam Sirois and Dr Manuel Castellanos-Beltran under the IARPA SuperTools programme.




Figure 3.47: Measured SQUID voltage as a function of control current for a dc SQUID fabricated with the MITLL SFQ5ee 10 kA cm^{-2} process, with coupled inductors in layer M5.

The period can be read from the graph as 13.25 mA, although we automate calculation with a discrete Fourier transform for large test sets. From (3.71), with $n = 1$, $\delta I_{mod} = 0$ and $\delta I_{ctrl} = 13.25 \times 10^{-3} \text{ A}$, the mutual inductance $M = 0.156 \text{ pH}$.

An InductEx model for the circuit tested above is shown in Figure 3.48. It is meshed with triangles. In order to reduce clutter on the image, edge segments and shadow casting patterns are omitted from the mesh shown here. With edge segments, shadow casting, a global mesh size of $1 \mu\text{m}$ (double the dimension of the line widths) and calibration to 48 test results for different inductance structures, the InductEx model yields a calculated result of $M = 0.158 \text{ pH}$.

3.4.2 Published inductance results

One step in the verification of InductEx is the comparison of extracted results to published experimental measurements. Some of these published results are referenced here.

Measured inductance results for the MIT Lincoln Laboratory 8-layer 250 nm [32] and SFQ5ee and SC1 processes [237] have been published, with reference to InductEx




Figure 3.48: InductEx model of a differential-arm inductance measurement SQUID for the MITLL SFQ5ee 10 kA cm^{-2} process with L_1 and L_{ctrl} in layer M5.

calculations. MIT Lincoln Laboratory also published measured results for high kinetic inductance structures [238] and detailed results for mutual inductance [125].

For stacked vias in a multilayer process, InductEx was applied to the AIST ADP process and compared to measured results [239].

3.4.3 Calibration

3.4.3.1 Hypres fabrication processes

As shown throughout this text, numerical inductance calculation is never completely accurate, but approaches analytical results for arbitrary small segment size – provided that the three-dimensional model and port placement are a good representation of the inductance structure.

However, small segment size in a 3D solver is very expensive in terms of calculation time and system memory. From an engineering perspective, numerical calculation parameters should thus be chosen to balance calculation time and accuracy over the entire range of feature sizes in a typical layout extraction problem.

From the earliest comparison between calculated and measured results on the FLUX-ONICS process, discussed in Section 3.3.4, it was evident that a close match between experiment and calculation could be obtained by selection of a modelling parameter – such as segment size – for calculation speed, and by subsequent adjustment of process parameters – specifically isolation layer thickness and superconductor London penetration depth – to compensate for offsets in the calculation results. The offsets are to be expected when current distribution in large segments cannot match the spatial distribution in a real structure. However, process engineers and managers pushed back against the reporting of any InductEx-specific process parameters that differed from the actual process (e.g. a London penetration depth for a niobium thin-film layer differing significantly from 90 nm) because it would seem to imply that the process was somehow not good.

At Hypres, Dr Oleg Mukhanov held a refreshingly different and very much engineering view: “artificial” process parameters in the InductEx layer definition file were entirely acceptable as a way to decrease the error in calculation results at a given segmentation size. Dr Mukhanov organised generous access to years of inductance test results from multiple chips over 22 wafers manufactured with mask aligner photolithography, and 5

wafers manufactured with wafer stepper photolithography for the Hypres 4.5 kA cm^{-2} process [22], [240].

The calibration procedure for the Hypres process, which is detailed in [201], started with an investigation into what would be the largest acceptable segment size (for lowest computing resource cost) that could be used with InductEx without introducing unacceptable error. I was free to define the limits of acceptability, and I thus pegged it as the largest segment size for which the average error between InductEx calculations and measured results at every available line width stayed flat over the range of line widths. With access to a vast trove of data points, I picked layer M1 microstrip over M0 ground and ran InductEx for models covering the entire measured range from $2.5 \mu\text{m}$ to $20 \mu\text{m}$ line width for the mask aligner data set, and from $0.8 \mu\text{m}$ to $10 \mu\text{m}$ line width for the wafer stepper data set. The results are shown in Figure 3.49.

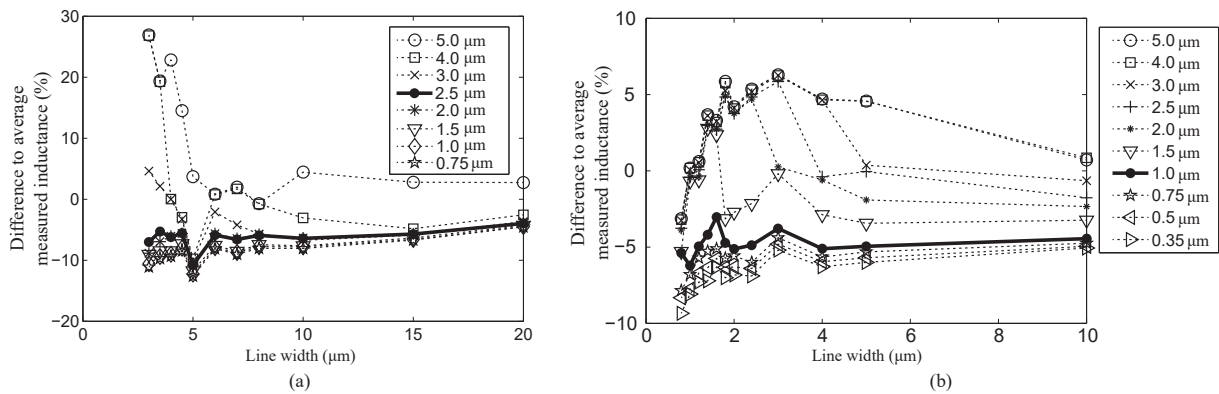


Figure 3.49: Difference between InductEx calculations and average measurements of inductance in M1 microstrip over M0 ground for the Hypres process with nominal process parameters (a) for 22 wafers manufactured with mask aligner photolithography and (b) for 5 wafers manufactured with wafer stepper photolithography. Each data trace represents a fixed maximum segmentation size, and the bold trace in each graph represents the largest segmentation size for which the error remains flat.

For both data sets, the largest segmentation size that yielded a flat error over line width (shown in bold in Figure 3.49) was equal to the smallest line width. It was immediately clear that a segment size larger than the smallest line width would lead to rapidly increasing error at lower line widths. I thus selected the minimum demonstrated line width for each wafer processing method – $2.5 \mu\text{m}$ for the mask aligner and $1.0 \mu\text{m}$ for the wafer stepper – as the segment size used for InductEx calibration. Inductance extraction of every structure was done with the selected segmentation size and nominal process parameters. The results are shown in Figure 3.50. A root mean square error (RMSE) of 7.6 % was obtained for the mask aligner data set, while the RMSE for the wafer stepper data set was 6.94 %.

A method was designed to calculate and plot the effect of varying each process parameter on every width for every microstrip and stripline combination. From these plots, parameters could be adjusted to nudge the results closer to measurement for every conductor-ground combination *over all line widths*. The calibrated process parameters derived from these adjustments are shown in Table 3.6 [201] when segment size is $2.5 \mu\text{m}$ for mask aligner photolithography and $1 \mu\text{m}$ for wafer stepper photolithography. The errors after calibration are plotted in Figure 3.51.

The RMSE after calibration, over all structures, is 2.25% or below, which is remarkably

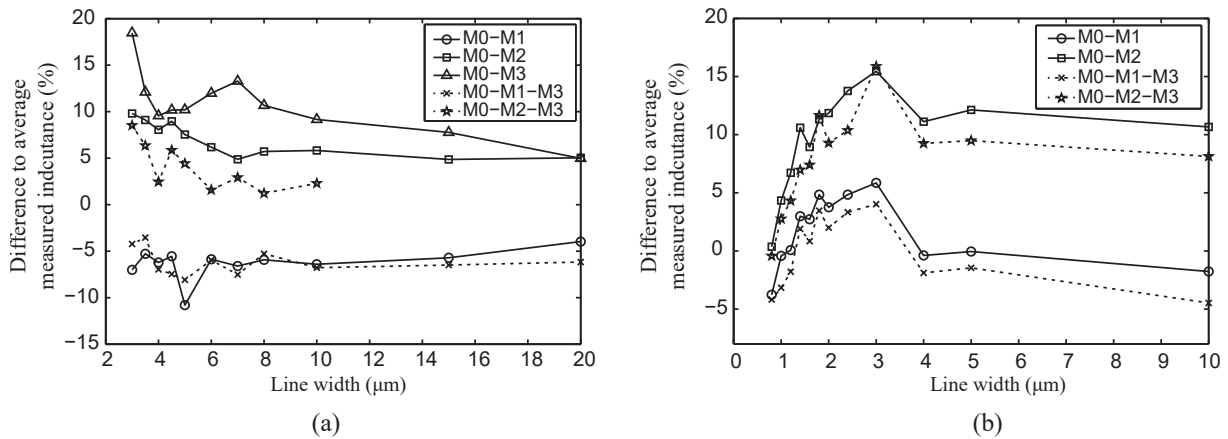


Figure 3.50: Difference between InductEx calculations and average measurements of inductance for the Hypres 4.5 kA cm^{-2} process with nominal process parameters (a) for mask aligner photolithography and a fixed segmentation size of $2.5 \mu\text{m}$ (RMSE = 7.6 %) and (b) for wafer stepper photolithography and a fixed segmentation size of $1.0 \mu\text{m}$ (RMSE = 6.94 %). Layer combinations are “ground-conductor” for microstrip, and “ground-conductor-ground” for stripline.

good considering the rough estimates on inductance used before the introduction of InductEx.

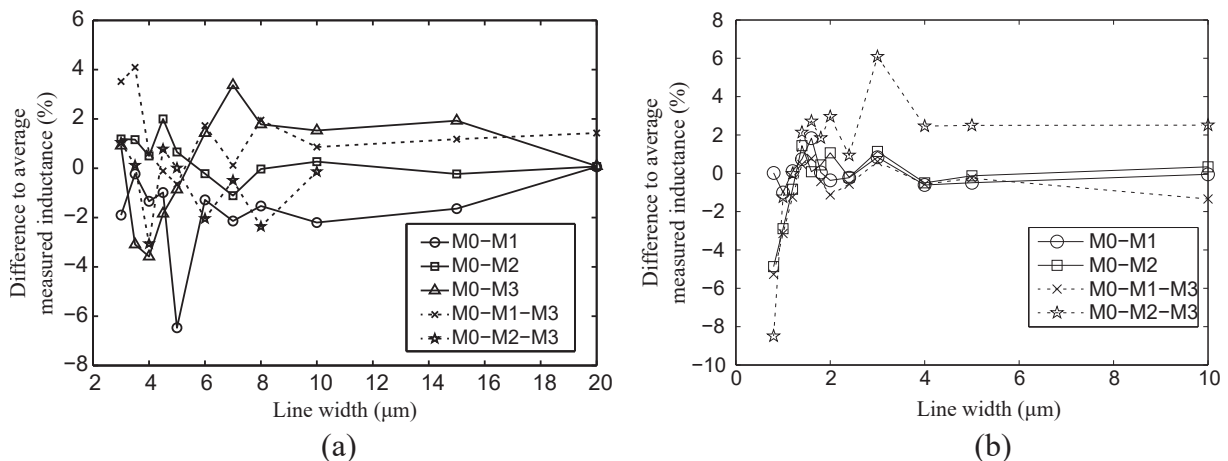


Figure 3.51: Difference between InductEx calculations and average measurements of inductance for the Hypres 4.5 kA cm^{-2} process with calibrated process parameters (a) for mask aligner photolithography with $2.5 \mu\text{m}$ segments and an RMSE of 1.87 % and (b) for wafer stepper photolithography with $1.0 \mu\text{m}$ segments and an RMSE of 2.25 %). Layer combinations are “ground-conductor” for microstrip, and “ground-conductor-ground” for stripline.

The calibrated parameter sets for the Hypres process were used by my group and other Hypres customers for several years, and it provided a blueprint for all subsequent process calibrations. The method is still used, although, at the time of writing, it is already possible to calibrate processes to the same RMSE with double or more the segmentation size through the use of triangular or tetrahedral meshes with current distribution enhancement methods such as shadow casting and edge meshing.

Table 3.6: InductEx process parameters for Hypres mask aligner and wafer stepper processes.

Description	Mask aligner		Wafer stepper	
	Nominal	Calibrated 2.5 μm	Nominal	Calibrated 1.0 μm
λ_{M0} (nm)	90	78	90	90
λ_{M1} (nm)	90	112	90	99
λ_{M2} (nm)	90	97.5	90	60
λ_{M4} (nm)	90	110	90	90
Bias M1 (μm)	0	0.015	0	0
Bias M2 (μm)	-0.2	-0.1	-0.1	-0.15
Bias M3 (μm)	-0.4	-0.05	-0.3	-0.3
Thickness I0 (nm)	150	150	150	200
Thickness I1 (nm)	200	190	200	140
Thickness I2 (nm)	500	495	500	500

3.4.3.2 AIST processes

Following the successful development of calibrated InductEx layer definition files for the FLUXONICS and Hypres processes, a similar calibration was done for the AIST 10 kA cm⁻² ADP2 process [26].

Before calibration, calculation results for nominal process parameters were compared to measurements for several pulse transfer circuits [144]. The results were published [241]. With the exception of an error of more than 10 % for one of the COU layer structures over a ground plane hole, all results were within 5 % of measurements.

In follow-up experiments, I designed structures for inductance calibration that would, for the first time, include inductors that threaded one or more ground planes [242]. Calibration was done similarly to that discussed in Section 3.4.3.1 and [201]. The final calibrated RMSE between extracted and measured results, over all layers through all ground plane threading combinations, was 2.24%, which is almost exactly the same as that obtained for the Hypres process in Section 3.4.3.1.

3.4.3.3 FLUXONICS process

As part of the measurement of inductance over ground plane holes [108] for the FLUXONICS process, I calculated the RMSE between measurements and extraction for all the calibration SQUIDS, and showed an RMSE of 2.2%. Again, the accuracy is very close to that for the Hypres and AIST processes at the time.

3.4.3.4 MIT Lincoln Laboratory SFQ4ee and SFQ5ee processes

As part of precursor work before the IARPA SuperTools project, I collaborated with Hypres under the IARPA C3 project to compile appropriate calibration sets with which InductEx extraction results with cuboid meshes could be matched closely to all measured structures [243] (we were still developing tetrahedral meshing at that time). A test set was designed to test the inductance of all layers on the MITLL SFQ4ee process between different sets of ground and sky planes. A rendering of the InductEx extraction model for two such structures is shown in Figure 3.52.

This experiment was the first in which the experimental structures were designed to test all layer combinations for layouts conforming to typical gate dimensions, and with uniform input/output sections to eliminate differences between structures due to modeling

assumptions. During the calibration procedure, process parameters – here only the layer penetration depths – were recursively adjusted as fitting parameters for the calculation models until the RMSE between all calculation results and experimental measurements reached a minimum.

This calibration highlighted the difference between modelling parameters and process parameters. Process parameters represent the physical properties of the process, such as layer thickness and penetration depth. Model parameters only influence the horizontal segment size applied to numerical models, and the number of filaments for cuboid segments that make up the height of a conductor. For the MITLL SFQ4ee process, the test structures used striplines or microstrip with width $1.2\ \mu\text{m}$. I showed earlier that good calibration is possible if the maximum segment size equals the minimum line width or smaller [201], so that the the maximum segment size was selected as $1\ \mu\text{m}$ for a nominal set (Set 2 in Table 3.7). The modeling parameters have a significant effect on calculation speed and memory use, so that alternate calibration sets were created for faster calculation (Sets 3, 4, and 5 in Table 3.7) where the height filament count is fixed to 1 and maximum segment size is stepped through $1\ \mu\text{m}$, $1.5\ \mu\text{m}$, and $2.5\ \mu\text{m}$. Finally, a high-fidelity set (Set 1) was created with a maximum segment size of $0.5\ \mu\text{m}$, which would be accurate for all line widths down to about $0.5\ \mu\text{m}$.

The calibrated process parameters are listed Table 3.8. Artificial values derived from calibration are in bold. None of the values differ significantly from the actual process parameters, except for the penetration depth of M6 in Set 5. This indicates that a maximum segment size of $2.5\ \mu\text{m}$ is already too coarse and this set should not be used for reliable calculations.

Table 3.7: InductEx modelling parameters for five MITLL SFQ4ee layer definition file sets.

Modelling parameter	Set 1	Set 2	Set 3	Set 4	Set 5
Segment size μm	0.5	1	1	1.5	2.5
M0-M4 height filaments	2	2	1	1	1
M5 height filaments	1	1	1	1	1
M6-M7 height filaments	2	2	1	1	1

The RMSE results for the calibrated layer parameters are listed in Table 3.9 for the self inductances, the mutual inductances, and for all self and mutual inductances together. The average segment count and solution time (on a dual-core Intel i7 mobile processor clocked at 2.9 GHz) per structure are also listed. The results are very good, with a smallest RMSE of just 0.9% for self-inductance with Sets 1, 2, and 3. These were, at the time, the best calibration results ever reported for InductEx, and was attributed to both the process quality and the rigorous design of the test structures to have the same junction, holes, and ground-to-sky via layouts.

While Set 3 was sufficiently accurate for the calculation of self inductances in gate layouts with the MITLL SFQ4ee and SFQ5ee process nodes, Set 2 was more reliable where mutual inductances were required.

3.4.4 Mutual inductance in sub-micron structures

Early validation of the accuracy of InductEx calculation models was done for integrated circuits where inductors had line widths of $5\ \mu\text{m}$ to $10\ \mu\text{m}$, which is significantly larger




Figure 3.52: Rendered image of the 3D inductance model created by InductEx for MITLL SFQ4ee calibration structures. (a) Most of the SQUID loop inductor in M4 and a lower ground plane in M3, and (b) most of the SQUID loop inductor in M3, a coupled line in M2 and a lower ground plane in M1. Both test structures have a sky plane in M7, which is peeled away for clarity in the rendered images. The dimensions of the structures are $60\ \mu\text{m} \times 22\ \mu\text{m}$

Table 3.8: Calibrated values of process (layer) parameters for five MITLL SFQ4ee layer definition file sets.

Layer parameter	Set 1	Set 2	Set 3	Set 4	Set 5
M0-M4 thickness (nm)	200	200	200	200	200
M5 thickness (nm)	135	135	135	135	135
M6-M7 thickness (nm)	200	200	200	200	200
I0-I3 thickness (nm)	200	200	200	200	200
I4 thickness (nm)	250	250	250	250	250
I4 thickness (nm)	270	270	270	270	270
I4 thickness (nm)	200	200	200	200	200
M0 λ (nm)	92	90	80	84	84
M1 λ (nm)	98	90	87	86	86
M2 λ (nm)	95	90	84	84	90
M3 λ (nm)	95	90	87	86	86
M4 λ (nm)	95	90	87	86	86
M5 λ (nm)	103	96	93	93	91
M6 λ (nm)	82	79	75	71	30
M7 λ (nm)	105	90	87	86	90

Table 3.9: Results for five MITLL SFQ4ee layer definition file sets.

Result	Set 1	Set 2	Set 3	Set 4	Set 5
Average number of segments	108 200	66 500	35 540	33 840	32 750
Average solution time (s)	351	201	90	78	77
RMSE of all self inductance	0.9%	0.9%	0.9%	1.6%	1.9%
RMSE of all mutual inductance	0.9%	1.4%	4.0%	4.0%	9.0%
RMSE of all self and mutual inductance	0.9%	1.1%	2.3%	2.5%	5.1%

than the conductor and isolation layer thickness of around 200 nm to 500 nm, and where a single ground plane was used.

Until around 2015, we could only guess how well accuracy would hold when device size shrinks, because fabrication processes did not support sub-micron feature sizes. Since the field equations do not change, I assumed that shrinking down the mesh size with feature size would be sufficient. I was only partly right. In the latest MITLL processes, inductor layouts with line width down to $0.25\ \mu\text{m}$ is possible, so that line width is then almost the same as its thickness and the isolation to the nearest ground plane. Inductors can also be threaded through multiple ground planes. In simple stripline configurations, self inductance scaled remarkably well, but mutual inductance results were observed to drift.

By 2022, as triangular segments replaced older cuboid segments for large inductance extraction models – those with millions of segments to model large swathes of layout over all the metal layers of the MITLL SFQ5ee process – I had increased the default segment size to $2\ \mu\text{m}$ and calibrated it so that layouts with line widths down to $1\ \mu\text{m}$ would be handled adequately. However, experimental measurements on weak coupling between very narrow stripline layouts were shown [125] in 2022 to cause significant overestimation of the mutual inductance by InductEx.

The cause of this overestimation comes from modelling. When lines are much narrower than the segment size – in this case down to $0.25\ \mu\text{m}$, or eight times narrower than the segment size – the ground plane segments are far bigger than the line width. The return current is then modelled inaccurately beneath lines. Where the ground plane segments overlap both lines, strong coupling is then created artificially. The InductEx model for such a circuit is shown in Figure 3.47.

The easy solution is to decrease maximum segment size to $0.25\ \mu\text{m}$, but the resource cost for large layouts is prohibitive. A more elegant solution is to cast “shadows” from every object to the nearest ground planes above and below, and to create mesh elements that have edges on the shadow boundaries, as is shown in Figure 3.53. Furthermore, the addition of narrow segments with width equal to the penetration depth around the outside of every conductors models the edge current distribution much better, and improves the accuracy of mutual inductance extraction.

I added these methods to InductEx, with shadow casting shown in Figure 3.53.

The RMSE results between measurement and calculation are shown in Table 3.10. The mutual inductance differs between about 30% of the self inductance for half of the structures, where coupling is between overlapping striplines on different layers, and 9% of the self inductance where coupling is between adjacent lines as shown in Figure 3.47. The table includes RMSE results when mutual inductance is normalise to self inductance.

It is clear that shadow casting and edge slicing bring self inductance and normalised mutual inductance within an RMSE of 3%. Crucially, this is for lines with width down to $0.25\ \mu\text{m}$, while segment size is $2\ \mu\text{m}$ – a significant result.

Table 3.10: RMSE results for very narrow coupled structures in MITLL SFQ5ee process.

Meshing method	RMSE of L	RMSE of M	RMSE of M normalised to L
Normal mesh	7.61%	29.5%	3.46%
Shadow casting	4.89%	8.19%	3.19%
Shadow casting and edge slicing	2.69%	4.66 %	1.47%

In the end, these mutual inductance experiments confirm that errors in extracted inductance arise almost solely from modelling.




Figure 3.53: Rendered image of the 3D inductance model created by InductEx for an MITLL SFQ4ee JTL layout. The top image shows the model without the skyplane in M7, with shadow casting to the ground plane M4 visible. The bottom image shows the model with the M7 skyplane and the shadows cast on it included.

3.5 Conclusion on contributions

My group's contributions to integrated circuit inductance calculation has been far-reaching.

At the start of my research career, inductance in superconductor integrated circuits was estimated from analytical approximations, or calculated for limited geometries with quasi-2D numerical methods. Today, InductEx is the inductance extraction utility of choice from academia to research laboratories and large military-industrial companies. When superconductor IC fabrication facilities publish inductance data on their processes, InductEx is now often used as the numerical reference [32], [237].

Chapter 4

Tool chain

4.1 Background

The design of an integrated circuit starts from a concept that can range from a few devices such as detectors in an analogue system to a complex behavioral description such as a digital microprocessor.

The end goal of the design stage is to deliver photomask layouts to a fabrication facility (or “fab”) in a process called “tape-out”.

In order to proceed from a design concept to tape-out, reliable and capable tool chains are required that can handle the implementation and verification of integrated circuit designs with thousands to millions of components.

After a decade of research and development in superconductor electronic design automation (EDA) tools, which included the NioCAD project – funded from 2007 to 2009 by the South African National Research Foundation’s Innovation Fund, and afterwards until its termination in 2012 by the Industrial Development Corporation – I published an audit of the status of superconductor electronic circuit design tools [244] with my then PhD student Mark Volkmann in 2013. We unpacked the available tools and user behaviour at the time, and concluded that, with the exception of a few efforts such as NioCAD, there had been little progress since the previous comprehensive EDA status assessment in 1999 [245]. We predicted that EDA tools would only really start to develop when the complexity of superconductor integrated circuits increased, and that it would most likely require an international open-source effort to develop and maintain a useful set of SCE EDA tools.

We were right on both counts.

In 2014, IARPA commissioned the Cryogenic Computing Complexity (C3) programme to develop complex superconductor microprocessors and memory. That programme, together with vast improvements in fabrication from MIT Lincoln Laboratory, exposed the dire need for better EDA tools. In 2015, IARPA started a seedling programme on SCE EDA tools in which I participated in collaboration with Prof. Massoud Pedram at the University of Southern California. In 2016, building on the results of the seedling programme, the IARPA SuperTools programme was announced. The end goal of SuperTools was to have commercial and open-source tool chains for SCE integrated circuit design, and it has advanced the state-of-the-art significantly in just five years.

My PhD student Dr Nicasio Maguu Muchuka had just published an open-source tool overview [246], while my then Masters student Dr Johannes Delpont and I had significant results from the seedling project. We also had results from the IARPA C3 programme

and could put together a formidable development plan for SuperTools. We joined Prof. Massoud Pedram at the University of South California again, broadened the project to include Professors Peeter Beerel, Murali Annavaram, Shahin Nazarian and Sandeep Gupta at USC, Prof. Mark Law at the University of Florida, Prof. Pascal Febvre at the University of Savoie-Mont Blanc, Prof. Yanzhi Wang at Northeastern University and Prof. Nobuyuki Yoshikawa at Yokohama National University.

I published an updated roadmap in 2018 [247] and led a team effort to publish a comprehensive overview of the ColdFlux tool chain [92] (then in development) under SuperTools. I have also been invited to deliver review or overview talks on CAD tool development for SFQ circuits, such as [248] and [249].

Since these publications came out, the software tool chain has advanced even further. My contributions to the tool chain are discussed below.

4.2 Electrical simulation engine: JoSIM

4.2.1 History

Electrical simulators are used to verify the transient behaviour of SCE circuits at the device and logic gate level for digital circuits, and are indispensable during the design process.

For conventional semiconductor integrated circuits, SPICE (Simulation Program with Integrated Circuit Emphasis) is almost universally used for electrical circuit simulation. SPICE was created at UC Berkeley as a class project in 1969-1970. It was released as open source software and evolved into a powerful and essential tool for simulation of integrated circuit operation and performance. Today there are many SPICE simulators, all based on the genetics of the original SPICE but with different implementations of numerical methods or code.

SPICE uses three main numerical methods: *Newton iteration* to find the solution of circuits with nonlinear elements, *sparse matrix methods* to fit enormous matrices into computer memory and solve LU decomposition in finite time, and *implicit integration* to integrate the differential equations that arise from reactive circuit components. Differences in the implementation of the numerical methods between SPICE simulators can (and do) cause different simulation results – something that any circuit designer should be aware of.

Standard SPICE engines lack support for the Josephson junction. The first circuit simulator specifically for SCE circuits was COMPASS [250], which initially supported the resistively shunted junction model. It was soon extended [251] to include the Werthamer microscopic tunneling model [252], but was never widely adopted.

A more widely used circuit simulator, PSCAN, was introduced in 1991. PSCAN uses a modified nodal phase method and supports the microscopic tunneling model as one option for Josephson junction simulation. PSCAN used dimensionless units for circuit elements, which makes the circuit schematics slightly different to those of a standard SPICE engine. Furthermore, PSCAN only supported inductive coupling in two-inductor transformers and could thus not model inductors coupled to multiple other inductors – an essential requirement for the analysis of circuits such as AQFP gates. This shortcoming was fixed when PSCAN was rewritten in Python and released as open source software PSCAN2 [253]. Although PSCAN2 is a powerful simulation engine that runs more than

an order of magnitude faster than PSCAN, it is not widely used, probably because of the lack of user manuals or example sets.

The most popular electrical simulators for SCE electronics have intrinsic support for the resistively and capacitively shunted junction (RCSJ) model of the Josephson junction [254], [255]. The RCSJ model, which balances accuracy and computing time, is sometimes referred to as the Stewart-McCumber model. JSPICE3 [256] had an intrinsic RCSJ model as a three-terminal device, with the third terminal used to probe the junction phase. The direct successor of JSPICE3, WRspice [192], is a very powerful voltage-based simulator with support for semiconductor devices. It supports the RCSJ Josephson junction model with a standard piecewise-linear model, an analytic exponentially derived approximation, and a fifth-order polynomial expansion model for quasiparticle resistance. Different critical current models are also supported, so that the magnetic field coupling can be modelled. WRspice was a commercial simulator, but was released as open source software when its creator, Dr Stephen Whiteley, joined Synopsys under the SuperTools programme.

JSIM [257] is a lightweight voltage-based simulator that was widely used until very recently for both analogue and digital SCE simulation. It was designed to operate on systems without large random access memory, which limits efficiency. JSIM only supports passive circuit elements and the Josephson junction (through to the RCSJ model) with a piecewise-linear quasiparticle resistance. It is limited to transient analysis. A modified version that includes limited thermal noise analysis support was released as JSIM_n [258]. A script running under Linux converts a noiseless simulation deck into a deck at a specified temperature with noise current sources added to all linear resistors, so that noise in the Josephson junction resistances is not supported.

4.2.1.1 Modified nodal analysis

A circuit simulator needs to find a set of linear equations of the form $\mathbf{Ax} = \mathbf{b}$ that can be solved simultaneously at a fixed point (static solution) or at any time step in a transient solution. The modified nodal analysis (MNA) method [259] is used in SPICE simulators to find this set of equations.

In MNA, the Kirchhoff current law (KCL) is used to do nodal analysis to find the voltages at each node. Branch conductances are stamped into the \mathbf{A} matrix, unknown node voltages are placed in the variable vector \mathbf{x} , and the sum of currents at each node (usually zero) is written into right hand side (RHS) vector \mathbf{b} . Where branch current is independent of node voltage, such as for voltage sources, or where node voltage is a function of branch current, such as for inductors, branch currents are added as unknowns to \mathbf{x} . Independent current and voltage sources contribute to \mathbf{b} .

The matrix \mathbf{A} on the left hand side (LHS) is square, and the contributions of every component are stamped into \mathbf{A} for the nodes to which it is connected and its independent or controlling branch currents.

At every time step, unknowns are computed as $\mathbf{x} = \mathbf{A}^{-1}\mathbf{b}$.

4.2.2 JoSIM

The Josephson simulator (JoSIM) [260], [261] was conceived under the ColdFlux project as a simulation engine that would exploit modern coding methods for improved speed and larger circuit support (with an initial aim of one million circuit components) than existing

superconductor circuit simulators. The work formed the main focus of the doctoral studies of my postgraduate student, Dr Johannes Delpont.

At its core, JoSIM is set apart from other simulators by the provision for two analysis modes for the solution of linear circuit equations: a traditional modified nodal voltage analysis mode and a modified nodal phase analysis mode. The two modes require different MNA stamps.

4.2.2.1 Integration method

Reactive components such as inductors and capacitors require an integration method to determine the current or voltage at every time step. The most basic of these methods is the first order backward Euler method, which interpolates the current value based on the previous value. This method suffers loss of accuracy unless time steps are sufficiently small, and is thus computationally expensive. SPICE simulators therefore use a second order method such as trapezoidal integration or backward differential formula (BDF, or Gear) [262] integration.

Trapezoidal integration is faster than Gear and is considered to be more accurate in standard SPICE simulators, but is known to cause numerical ringing. Gear integration dampens all ringing: numerical and physical, so that it could suppress physical ringing and result in incorrect simulation results if time steps are not sufficiently small.

Both integration methods were implemented in JoSIM.

The trapezoidal integration method is defined as

$$\left(\frac{dy}{dt}\right)_n = \frac{2}{h_n}(y_n - y_{n-1}) - \left(\frac{dy}{dt}\right)_{n-1}, \quad (4.1)$$

where n is the iteration count and h_n the current time step in the transient analysis. This integration method is suitably accurate for circuit simulation purposes, but a rapid change in y , such as during a 2π phase switch of a Josephson junction, tends to produce spikes in the derivative (voltage for a Josephson junction). We also observed excessive ringing with the trapezoidal method.

The Gear, or second order BDF method, is expressed as

$$\left(\frac{dy}{dt}\right)_n = \frac{3}{2h} \left[x_n - \frac{4}{3}x_{n-1} + \frac{1}{3}x_{n-2} \right], \quad (4.2)$$

where n is still the iteration count and h is the time step of the simulation. This method requires the results of the two previous time steps.

4.2.2.2 MNA component stamps

Each circuit component has an MNA stamp that dictates entries into the matrix \mathbf{A} and the vector \mathbf{b} .

The inductor is shown as an example of how a component stamp is created. The voltage over an inductor is

$$v = L \frac{di}{dt}. \quad (4.3)$$

If we apply the equation in (4.1) to (4.3) we obtain an equation for the inductor voltage that is dependent on the previous time step current and voltage:

$$V_n - \frac{2L}{h_n}I_n = -\frac{2L}{h_n}I_{n-1} - V_{n-1}, \quad (4.4)$$

where (4.4) can be written in general matrix form as

$$\begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & -1 \\ 1 & -1 & -\frac{2L}{h_n} \end{bmatrix} \begin{bmatrix} V^+ \\ V^- \\ I \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ -\frac{2L}{h_n}I_{n-1} - V_{n-1} \end{bmatrix}.$$

The stamps for a resistor, inductor, capacitor and voltage source can be seen in Table 4.1.

JoSIM uses the RCSJ model, for which the MNA stamp is identical to that used in JSIM [257]. It relies on a second order guess of the phase for the next time step.

$$\begin{bmatrix} \frac{1}{R} + \frac{2C}{h_n} & -\frac{1}{R} - \frac{2C}{h_n} & 0 \\ -\frac{1}{R} - \frac{2C}{h_n} & \frac{1}{R} + \frac{2C}{h_n} & 0 \\ -\frac{h_n}{2} \frac{2e}{\hbar} & \frac{h_n}{2} \frac{2e}{\hbar} & 1 \end{bmatrix} \begin{bmatrix} V^+ \\ V^- \\ \phi \end{bmatrix} = \begin{bmatrix} I_s \\ -I_s \\ \phi_{n-1} + \frac{h_n}{2} \frac{2e}{\hbar} V_{n-1} \end{bmatrix}$$

Here, the phase node ϕ is a virtual node not connected physically in the circuit and e and \hbar are the electron charge and the reduced Planck's constant respectively.

The current value on the RHS is defined as

$$I_s = -I_c \sin \phi^0 + \frac{2C}{h_n} V_{n-1} + C \dot{V}_{n-1}, \quad (4.5)$$

with the phase guess

$$\phi_n^0 = \phi_{n-1} + \frac{h_n}{2} \frac{2e}{\hbar} (V_{n-1} + v_n^0) \quad (4.6)$$

and the voltage guess

$$v_n^0 = V_{n-1} + h_n \dot{V}_{n-1} \quad (4.7)$$

This method of using the phase guess relies on a voltage guess and subsequently information about the previous two values of the junction voltage as well as their derivatives. When the voltage guess has a large magnitude, the phase guess becomes extremely large and can lead to simulation instability. This is mitigated at the start of simulation by pegging the phase guess to zero for the first few simulation time steps.

Table 4.1: MNA component stamps for voltage method and trapezoidal integration

	LHS	RHS
R	$\begin{bmatrix} \frac{1}{R} & -\frac{1}{R} \\ -\frac{1}{R} & \frac{1}{R} \end{bmatrix} \begin{bmatrix} V^+ \\ V^- \end{bmatrix}$	$\begin{bmatrix} 0 \\ 0 \end{bmatrix}$
C	$\begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & -1 \\ 1 & -1 & \frac{h_n}{2C} \end{bmatrix} \begin{bmatrix} V^+ \\ V^- \\ I_C \end{bmatrix}$	$\begin{bmatrix} 0 \\ 0 \\ -\frac{h_n}{2C}I_{n-1} - V_{n-1} \end{bmatrix}$
V	$\begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & -1 \\ 1 & -1 & 0 \end{bmatrix} \begin{bmatrix} V^+ \\ V^- \\ I_V \end{bmatrix}$	$\begin{bmatrix} 0 \\ 0 \\ V_n \end{bmatrix}$

The current through a Josephson junction is related to phase, so that the direct calculation of the phase is more practical. It therefore becomes sensible to perform the entire transient analysis in phase since each component affects the phase of the entire circuit.

Calculation of phase is done through the Josephson voltage-phase relation

$$v = \frac{\Phi_0}{2\pi} \frac{d\phi}{dt}. \quad (4.8)$$

This relation can be substituted into every voltage dependent equation and expanded using the trapezoidal rule, so that modified nodal phase analysis (MNPA) stamps can be created for each component.

The Josephson junction MNPA stamp sees the role of the voltage and phase swapped. This means that the phase is now the connected node which we want to calculate, and the voltage becomes a virtual non-connected node which is required only for calculation purposes.

$$\begin{bmatrix} 0 & 0 & \frac{1}{R} + \frac{2C}{h_n} \\ 0 & 0 & -\frac{1}{R} - \frac{2C}{h_n} \\ 1 & -1 & -\frac{h_n}{2} \frac{2e}{\hbar} \end{bmatrix} \begin{bmatrix} \phi^+ \\ \phi^- \\ V \end{bmatrix} = \begin{bmatrix} I_s \\ -I_s \\ \phi_{n-1} + \frac{h_n}{2} \frac{2e}{\hbar} V_{n-1} \end{bmatrix}$$

The phase for the next time step remains the same second order guess as in (4.6) which utilises a voltage guess as in (4.7).

Direct calculation of the phase allows the addition of DC external magnetic fields through mutual coupling with all the inductors in the circuit. This is a rather important feature in low temperature superconductivity due to the high susceptibility to external fields which is not trivial using voltage-based methods.

The MNPA stamps for other components are shown in Table 4.2 for comparison.

With second order BDF integration, the MNA and MNPA stamps are altered to accommodate the integration. The MNPA stamps of some components are shown in Table 4.3 for comparison.

4.2.2.3 JoSIM application

The original goal with the development of JoSIM was to have an electrical simulation engine that could handle the size of ColdFlux circuits *in the phase domain*. For fast LU decomposition, JoSIM uses KLU [263] and is written in modern C++.

By 2018, Dr Johannes Delpont had compared JoSIM to the available engines WRSpice and JSIM for speed and simulation size capabilities [261]. Several examples were simulated, each with a time step of 0.25 ps with the maximum time step set to the same for a total of 1000 ps. The examples were executed on a system with an Intel Core i5 and 8GB RAM running macOS Mojave. The results of these simulations are shown in Table 4.4. In small examples simulators are quite closely matched however as the size of simulation grows JoSIM starts to gain ground.

Towards the end of ColdFlux, a large SFQ clock distribution network was simulated with JoSIM. The simulation contained 23 592 967 circuit components, of which 3 670 017 were Josephson junctions. A representative simulation length to test the entire system required 64 GM of RAM and 5700 s to complete. JoSIM thus reached the ColdFlux requirements of more than 10 000 000 electrical components in a circuit that can be simulated with reasonable resources.

Table 4.2: MNA component stamps for phase method and trapezoidal integration

	LHS		RHS
R	$\begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & -1 \\ 1 & -1 & -\frac{\pi h_n R}{\Phi_0} \end{bmatrix}$	$\begin{bmatrix} \phi^+ \\ \phi^- \\ I_R \end{bmatrix}$	$\begin{bmatrix} 0 \\ 0 \\ \frac{\pi h_n R}{\Phi_0} I_{n-1} + \phi_{n-1} \end{bmatrix}$
L	$\begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & -1 \\ 1 & -1 & -\frac{2\pi}{\Phi_0} L \end{bmatrix}$	$\begin{bmatrix} \phi^+ \\ \phi^- \\ I_L \end{bmatrix}$	$\begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}$
C	$\begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & -1 \\ 1 & -1 & -\frac{\pi h_n^2}{2C\Phi_0} \end{bmatrix}$	$\begin{bmatrix} \phi^+ \\ \phi^- \\ I_C \end{bmatrix}$	$\begin{bmatrix} 0 \\ 0 \\ \frac{\pi h_n^2}{2C\Phi_0} I_{n-1} + \phi_{n-1} + h_n \dot{\phi}_{n-1} \end{bmatrix}$
V	$\begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & -1 \\ 1 & -1 & 0 \end{bmatrix}$	$\begin{bmatrix} \phi^+ \\ \phi^- \\ I_V \end{bmatrix}$	$\begin{bmatrix} 0 \\ 0 \\ \frac{\pi h_n}{\Phi_0} (V_n + V_{n-1}) + \phi_{n-1} \end{bmatrix}$

Table 4.3: MNPA component stamps for second order BDF integration

	LHS		RHS
R	$\begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & -1 \\ 1 & -1 & -\frac{2\pi R}{\Phi_0} \frac{2h_n}{3} \end{bmatrix}$	$\begin{bmatrix} \phi^+ \\ \phi^- \\ I_R \end{bmatrix}$	$\begin{bmatrix} 0 \\ 0 \\ \frac{4}{3}\phi_{n-1} - \frac{1}{3}\phi_{n-2} \end{bmatrix}$
L	$\begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & -1 \\ 1 & -1 & -\frac{2\pi}{\Phi_0} L \end{bmatrix}$	$\begin{bmatrix} \phi^+ \\ \phi^- \\ I_L \end{bmatrix}$	$\begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}$
C	$\begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & -1 \\ 1 & -1 & -\frac{4h_n^2}{9C} \frac{2\pi}{\Phi_0} \end{bmatrix}$	$\begin{bmatrix} \phi^+ \\ \phi^- \\ I_C \end{bmatrix}$	$\begin{bmatrix} 0 \\ 0 \\ \frac{8}{3}\phi_{n-1} - \frac{22}{9}\phi_{n-2} + \frac{8}{9}\phi_{n-3} - \frac{1}{9}\phi_{n-4} \end{bmatrix}$
V	$\begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & -1 \\ 1 & -1 & 0 \end{bmatrix}$	$\begin{bmatrix} \phi^+ \\ \phi^- \\ I_V \end{bmatrix}$	$\begin{bmatrix} 0 \\ 0 \\ \frac{4}{3}\phi_{n-1} - \frac{1}{3}\phi_{n-2} + \frac{2\pi}{\Phi_0} \frac{2h}{3} V_n \end{bmatrix}$
P	$\begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & -1 \\ 1 & -1 & 0 \end{bmatrix}$	$\begin{bmatrix} \phi^+ \\ \phi^- \\ I_\phi \end{bmatrix}$	$\begin{bmatrix} 0 \\ 0 \\ \phi_n \end{bmatrix}$

Table 4.4: Comparison of electrical simulator execution speed

Simulation		Execution Times(s)		
Circuit Description	JJ Count	JSIM	WRspice	JoSIM
Basic JTL	2	0.057	0.107	0.134
400 simulation I-V curve	400	60.5	56	54.87
4-bit KSA	2095	48.500	DNF	23.490
General Partial Products	3904	93	64	20.9
3000 JTL string	6006	276	159	91.89
4000 JTL string	8006	>3 600 000	232.7	130.87
5000 JTL string	10 006	DNF	263.8	169.81

*DNF: Did not finish. Non-convergence, time step too small

At the time of writing, simulations with more than 50 000 000 components and 7 000 000 Josephson junctions can execute successfully with about 80 GB of RAM use.

JoSIM is now the primary electrical simulator for my research group, but also a wider audience outside of SuperTools. It is so widely used today that the online user manual has even been translated into Japanese.

4.3 Device level tools

Device level tools are the closest to the physical materials from which the devices on an integrated circuit are constructed.

4.3.1 Technology CAD

As SC circuit structure and device dimensions shrink deeper below 1 μm , device design will depend increasingly on powerful technology computer-aided design (TCAD) tools that provide reliable process simulation capabilities. Process simulation for device design is already used extensively in semiconductor circuits. In SC circuits, TCAD tools would aid with process simulation for the design of small, high current density self-shunted Josephson junctions or memory elements that use magnetic devices or ferromagnetic junctions.

Prior to the SuperTools programme, there was no research into or tools for TCAD design in superconductor ICs. Through the ColdFlux project under SuperTools I met Professor Mark Law, who has had a successful career in TCAD research and development for semiconductor ICs. His group's Florida object-oriented device, process and reliability simulator (FLOOXs) has formed the basis of many a commercial process simulator. They adapted FLOOXs for superconductor IC fabrication processes [264] under the ColdFlux project.

Heinrich Herbst completed his Masters degree under my supervision on a tool called Katana that takes slices through a GDS layout file, determines the boundaries of objects on every layer in the two-dimensional cross-section formed by each slice, compiles a list of process instructions for FLOOXs and executes it, and then reads the two-dimensional mesh generated by FLOOXs [11]. An example of such a mesh obtained from FLOOXs is shown in Figure 4.1. The process-generated mesh was obtained for objects on three

niobium layers, separated by silicon dioxide isolation layers. The top layer is silicon dioxide, and the objects on the three niobium layers are separated in the centre by $0.2\ \mu\text{m}$ from other objects on the same layers. The curved profile of anisotropic etching is visible on the edges of the niobium layers, while silicon dioxide flows successfully into the etched regions.




Figure 4.1: Meshed 2D model from FLOOSS.

A tool called Silverlinings then finds the boundaries between different materials, removes the FLOOSS mesh, extrudes the boundaries into the third dimension to create three-dimensional objects (see Figure 4.2), and creates a new, more efficient mesh with Gmsh.




Figure 4.2: Boundary identification and extrusion of a FLOOSS-generated mesh with Silverlinings.

We used such TCAD-generated models of PTL layouts to evaluate the difference in electrical properties compared to when a simple rectangular cross-section for each layout element is used [265], [266] and found that all properties varied by less than 1% between the simple and TCAD-extracted models for a standard $5\ \Omega$ ColdFlux PTL, with the exception of peak current density. As shown in Figure 4.3, current crowds at the lower corner of the PTL line edge, and the peak current density is about 10% higher than that calculated for a rectangular model with 90 degree corners.

We concluded that process modelling is important for the extraction of Josephson junction device parameters, especially for very small junctions, but that there is no advantage to applying TCAD modelling for cell- and chip-level inductance extraction.

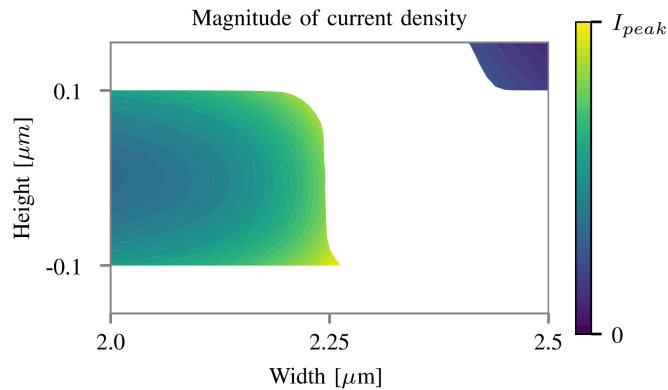


Figure 4.3: Current density across center strip edge of process modelled PTL.

To test this, and to show integration with InductEx, Heinrich developed a tool that would read a full cell layout, determine the edge profile of every layer from FLOOXs and turn it into a spline, and then build a meshed model where the edge profile is added as a prismatic spline sweep around the outer (or inner) edges of every object on a layer. The process is depicted in Figure 4.4.




Figure 4.4: Creation of a prismatic spline sweep to create a three-dimensional representation of a process-modelled object.

Although the automated model generator is not a full three-dimensional process simulation, but rather an application of two-dimensional process-modelled edge profiles layout objects, this is the only way to obtain a three-dimensional mesh for a structure of the size and complexity of a typical SFQ cell. The computing resources needed by FLOOXs makes full three-dimensional process modelling of anything more than a Josephson junction (device level) intractable.

Heinrich tested the tool on the layout of a full Josephson junction with shunt resistor and vias, as shown in Figure 4.5 and Figure 4.6. There is deliberate underetch of the metal layers to accentuate the edge profiles at this scale. In this layout, the via to the junction is smaller than the junction area. The profile gives insight into what the manufactured device would look like.

To test the capabilities of tool, the full layout of an OR2T cell from the ColdFlux RSFQ cell library (see Section 2.5.5.3), including all structures below the ground plane associated




Figure 4.5: Side view of grounded, shunted Josephson junction from TCAD extraction. Metal underetch is deliberate to accentuate the edge profile.




Figure 4.6: Three-dimensional rendering of grounded, shunted Josephson junction from TCAD extraction.

with the track block fabric and all fill structures, was modelled with process-determined edges. The full model, rendered with a ray tracing engine, is shown in Figure 4.7. The layout, at $100\ \mu\text{m} \times 70\ \mu\text{m}$, represents the largest full mesh model that we could create under the ColdFlux project.

To show integration with InductEx, ports were added to a similar swept-edge model created for a JTL and the mesh was analysed with InductEx as shown in Figure 4.8. No obvious difference in inductance could be observed between the model with process-extracted edges and a standard InductEx model with rectangular cross-sections for all objects, so that we concluded that the rectangular object modelling used by InductEx for both cuboid and tetrahedral meshes is still sufficient for superconductor integrated circuit modelling down to the minimum line widths of the MITLL SFQ5ee process.

4.4 Cell level tools

Cell level tools enable the SCE IC designer to design, characterise, optimise, layout and verify logic cells.

4.4.1 Characterisation

My first exposure to the analysis of SC circuit layouts was when my Masters degree supervisor, Prof. Willem Perold, asked me to help him model the effects of process-related tolerances on circuit parameters for COSL cells in the Hypres $1\ \text{kA cm}^{-2}$ process. We used actual layout dimensions to convert tolerances to variations in circuit component values [267], and found that simulations showed different – mostly lower – circuit yield (percentage of circuits that work as intended in a population with statistical variations




Figure 4.7: A process-extracted edge-swept model of a full $100\ \mu\text{m} \times 70\ \mu\text{m}$ OR2T cell with the M7 skyplane removed for visualisation purposes.




Figure 4.8: An InductEx model of a JTL with the M7 skyplane removed for visualisation purposes and current density due to the bias current.

applied) than circuits with blanket variations on all the component values as in other Monte Carlo analysis methods [268], [269].

I extended this work to a more thorough Monte Carlo analysis as part of my PhD, also under the supervision of Prof. Perold, and created a comprehensive modelling method for SC circuits [270]. I used this to great effect for analysing circuit yield, but the computational cost compared to the more intuitive and cheaper method of margin analysis [271], [272] meant that Monte Carlo analysis never gained widespread use. Monte Carlo analysis, despite its usefulness to verify that a circuit will very likely *always* work for given process tolerances, has two main drawbacks: it cannot tell us to which component a circuit is most sensitive, and it makes circuit optimisation very time-intensive and difficult once the yield is close to 100%.

Even though margin analysis checks only one component at a time (and thus disregards the effects that a variation in one component can have on the sensitivity of the circuit to another component), it still gives a very good indication of which components a circuit is most sensitive too. Margin analysis, with a binary search over every component value from the nominal to the upper and lower margins, is fast to compute and gives a result in terms of bias current margins that can easily be checked during experimental measurements.

I developed a margin analysis tool for my PhD for in-house use, but it was never released. The NioCAD project also included a margin analysis tool, but the first tool to be made open was written by Dr Mark Volkmann for his PhD [112]. It was simply called “Analyse”, and has since been superceded by “Optimum” developed by Dr Johannes Delpont under ColdFlux.

4.4.2 Optimisation

Once a working nominal circuit has been obtained for a new cell design, either from circuit equations or through manual manipulation of parameters on a chosen circuit configuration – a suboptimal technique that is disconcertingly widely practised – a margin analysis is performed to determine the narrowest (critical) margin. It is generally accepted that an SFQ circuit is robust if the critical margin is around 30%. Since we choose I_B to be roughly $0.7I_C$ (see Section 2.4.3), the margins on bias currents and junction critical currents are not likely to exceed 30% by much; no matter how good the circuit is.

Even 20% margins are generally acceptable on more complex cells. However, if the critical margin is below about 10%, and many margins significantly below 20%, circuit optimisation is required.

The number of components in a typical SFQ cell that need optimisation can be several tens, which results in such a large search space that only heuristic methods tend to be efficient. An early optimiser, MALT, was based on the method of inscribed hyperspheres [272]. It was not easily accessible, and did not gain wide use. COWBoy [273], another heuristic algorithm for margin optimisation, was integrated into PSCAN and is still in use in conjunction with PSCAN. However, it is not available as an optimisation tool when simulation engines such as JoSIM, JSIM and WRSpice are used.

Without access to PSCAN at the start of my research career, I investigated another heuristic method that was gaining popularity at the time: genetic algorithms. I developed an in-house tool to apply genetic algorithms to optimisation [274], [275], and used it to optimise the yield of all the RSFQ circuits developed during my Masters and PhD studies [56], [65].

Conventional genetic algorithms operate on problems that have been reduced to binary

strings [276] through a decoding function that maps the phenotype space (real-world parameters) to the genotype space [277]. These strings, called chromosomes or genomes, are able to reproduce, pair for crossover and undergo random bit mutation. The probability of reproduction is determined by circuit fitness. Pairing and crossover allows the exchange of genetic information, while mutation provides a way of introducing random jitter into solutions to prevent convergence on a single local optimum. Real-valued parameters can also be represented by binary substrings in the chromosome [276], but this limits the range of the solution. Since the parameters subject to optimization are the real-valued element values (resistance, inductance and critical current), I opted to map them directly to a genome of real values. This is sometimes referred to as an evolution strategy [277]. A starting population is drawn at random from the nominal circuit by spreading all the component values with a random distribution, and evaluating the yield of each circuit, and assigning a fitness value (such as the yield). A new population is then created according to a procreation probability based on fitness of the parents (circuits with higher fitness are more likely to be drawn as parents for the new population). New individuals (circuits) are then created by random crossover between the genomes of the parents and random mutations.

The genetic optimiser was compared against a random method, where a population is created by random variation of all components in the *best circuit of the previous generation*. The results in Figure 4.9 show that the genetic optimiser works, stalls when yield approaches 100%. Yield, found from expensive Monte Carlo analyses, turns out not to be a good fitness function. The tool was thus abandoned.

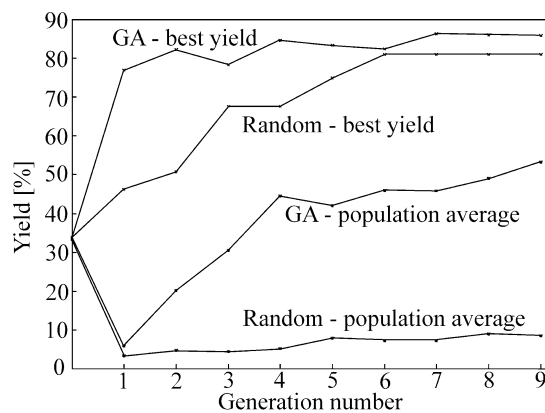


Figure 4.9: Comparative results for a genetic (GA) and random optimization sequence starting with the same unoptimized COSL set-reset flip-flop.

Margin analysis, which is computationally much more efficient than yield analysis, was used in conjunction with Monte Carlo analysis in an optimisation tool popular in Japan: SCOPE [278], whereas xopt, a centre-of-gravity method developed in Germany [279] was based purely on margin analysis. Figure 4.10 illustrates how margins typically appear before and after optimisation.

Under the NioCAD project, a metaheuristic optimiser was investigated [280], but it was found to have limited applicability to RSFQ circuits.

Under the ColdFlux project, my PhD student Dr Paul le Roux wrote an optimiser as part of the JoSIM-Tools package that uses differential evolution [281]. This optimiser was used for all ColdFlux cell library optimisation.




Figure 4.10: Typical margin analysis plots for (a) an unoptimised and (b) an optimised circuit.

4.4.3 Timing extraction and state verification

Under the ColdFlux project it was necessary to obtain hardware description language (HDL) models for every circuit in the cell library. These HDL models had to include state-dependent timing information so that functional verification of synthesised logic circuits could be performed.

I expanded earlier work by my PhD student Dr Louis Müller on automated state machine and timing characterisation [282] to develop a complete tool, TimEx [54]. With automatic cycle detection code similar to what I used in InductEx, TimEx finds every superconducting loop in a given superconductor circuit (called the *device under test*, or DUT) and then analyses the total flux around each loop at steady state after any input was changed. The flux must always sum to an integer number of Φ_0 – or fluxons – at steady state. The array of fluxon counts is then used as a signature for the state, so that TimEx can find the total number of states and identify which state is reached after an input causes a switching event.

TimEx is used for the automatic construction of a HDL models that contain all states, all input-dependent state transitions and all state-dependent critical times between any inputs, as well as state-dependent delay times between any input and output combination of an SFQ circuit. TimEx was verified against measured and published delay time measurements. The graph in Figure 4.11 shows the input-to-output delay time of a JTL as extracted with TimEx when the Stewart-McCumber parameter β_C is varied. Here, the characteristic voltage V_C is a normalised parameter used with PSCAN and defined [283] so that the Stewart-McCumber parameter is

$$\beta_C = \sqrt{V_C}, \quad (4.9)$$

The JTL uses junctions with nominal $I_C = 250 \mu\text{A}$ and is biased at $0.7I_C$ when the bias voltage is 2.5 mV. The characteristic voltage V_C is varied from 0.5 ($\beta_C = 0.25$) to 1.0 ($\beta_C = 1$). The extracted timing values follow the same dependence as that of measured results [283].

TimEx was released as an open-source deliverable under ColdFlux as a way to extract timing models and to automatically create HDL files. However, the main application was unforeseen: gate-level circuit designers started to use TimEx – with its exhaustive search of all combinations of inputs in every state as it maps out the Mealy state diagram of a circuit – to verify correct operation and identify any erroneous state or state transitions. An example is shown in Figure 4.12, where the TimEx-extracted Mealy finite state machine diagrams are shown for two RSFQ XOR gates, of which one functions correctly and has three states, and the other has two more erroneous states and is thus not a true XOR gate.




Figure 4.11: Extracted delay of a JTL with nominal $I_C = 250 \mu\text{A}$ in the MIT Lincoln Laboratory SFQ5ee process as a function of applied bias voltage and characteristic voltage V_C .




Figure 4.12: Mealy state diagram of (a) an RSFQ XOR gate with three states and (b) a nonfunctional RSFQ XOR gate with five states of which states 3 and 4 are illegal for XOR operation. A filled black circle indicates an SFQ output pulse at the output designated in capital letters.

4.5 Chip level tools

Chip level tools complete the tool chain for SCE integrated circuit design, but is mostly far removed from the physical level of component design and layout verification where most of my research has been focused. I thus give only a very brief overview of such tools

4.5.1 Interconnect analysis

In order to connect logic gates at any reasonable distance, passive transmission line (PTL) routing is used; mostly as microstrip or stripline.

My Phd student, Dr Paul le Roux, developed modelling methods for superconducting passive transmission lines [284]. We also investigated the effects of process-related deformation of transmission line structures on PTL characteristics [266], and developed SPICE model extraction techniques to represent PTL interconnects [285].

We have also investigated matching of PTL to circuits for improved performance [286].

4.5.2 Synthesis, placement and routing

With cells designed for row-based placement and routing [93], my research partner Prof. Massoud Pedram at the University of South California and his colleagues and students designed and built a comprehensive suite of tools for synthesis, placement and routing of combinational and sequential RSFQ systems called qPALACE [287].

In order to test our cell libraries, my Masters students Jude de Villiers and Edrich Verburg developed ViPeR, a tool suite for synthesis, placement and routing of combinational RSFQ circuits [288], [289]. ViPeR is more limited and simpler than qPALACE – the latter supports synchronous circuits and dual clock schemes – but has been used to synthesize and validate combinational circuits as large as a 32-bit ripple carry adder.

4.5.3 Static timing analysis

Static timing analysis (STA) is a technique that is used to provide an estimation of the expected timing (and power) of a digital circuit without the requirement for simulation. The number of timing paths from any input to any output increases as an exponential function with respect to the number of logic gates in the circuit. Therefore, it is impractical to perform a full-chip simulation at the electrical level. Timing information about a circuit is a crucial part of the standard cell-based design flow, and my then Masters student Dr Johannes Delpont developed an STA tool, SuperSTA [290], for the IARPA seedling project that eventually lead to SuperTools.

The STA tool could handle pre-placed STA and post-place-and-route STA, and relied on accurate timing models extracted with TimEx [54]. It has since been superseded by the qPALACE STA tool qSTA [291] developed by the group of Prof. Massoud Pedram under ColdFlux.

4.5.4 Layout-versus-schematic verification

Layout-versus-schematic (LVS) verification is used to determine if a layout matches an original circuit schematic. It is a crucial verification step, but has long been overlooked in SCE design tools. An implementation using Cadence DIVA was reported in 1997 [292], and groups with access to Cadence use a similar LVS implementation. My Masters

student Rebecca Roberts developed a limited LVS method specifically for SCE circuits in 2014 [293], but this has not yet been released as a standalone tool.

By the time SuperTools started, it was still not clear how to develop a proper, user-friendly LVS tool for SFQ circuit layouts. Recently, Dr Kyle Jackman developed a very good solution, which falls under the InductEx tool chain as InductEx-LVS. We have not published details yet.

4.6 Summary

Although my primary research focus has been the development of powerful inductance extraction and verification tools, superconductor electronics integrated circuit design requires a wider range of tools. My team and I contributed to such tools to help complete the design tool chain from circuit conception, to integrated circuit layout, verification and sign-off. Many of these tools also enable further analysis and verification of circuit layouts in conjunction with the inductance extraction tools. The most notable examples of such tools to which I contributed are: JoSIM, which made it possible to simulate trapped flux in circuits after inductance extraction; TCAD tools, which enable device analysis and parameter extraction from more realistic meshed models; state machine extraction tools that use flux signatures in inductive loops with or without Josephson junctions; and LVS tools that take inductance into account to verify full-chip layouts.

Chapter 5

Application

Although most of my research career has been devoted to the development of methods and tools to design and verify SCE integrated circuit components, with a strong focus on inductance and the interaction of magnetic fields with superconductor circuit structures, I have also applied these tools and methods to specific applications. These range from complex circuit design (detailed in Chapter 2) to the analysis of the effects of ground plane currents or the extraction of device parameters for analogue circuits.

I conclude with an overview of some of these applications.

5.1 Ground planes and return currents

5.1.0.1 Single ground plane

It has been demonstrated that SFQ circuits are sensitive to the dc bias return currents in the ground plane of a fabricated circuit, and that these currents in the ground plane can reduce circuit operating margins [294].

With the capabilities afforded us by InductEx, my team and I investigated the influence of the ground plane in superconductor circuits on the performance of RSFQ cells [295]. We demonstrated that we could derive simulation models that showed the effect of ground contact placement – in this case the extraction point for bias current from a single ground plane circuit layout – on circuit operational margins. This laid the groundwork for our subsequent development of compact simulation models to handle flux trapping and external magnetic fields.

5.1.0.2 Ground contacts

Before the availability of InductEx and its powerful current distribution calculation in three-dimensional structures, circuit designers rarely considered the effects of ground return currents on circuit performance.

The CONNECT cell library [27] uses a $30\ \mu\text{m} \times 30\ \mu\text{m}$ track block that allows easy tiling, and supports two PTL lines in the east-west and two PTL lines in the north-south direction. The track block was carefully designed to minimise coupling from the dc power layer [296], [297], which is placed at the very bottom of the layer stack. Due to the spacing of the PTLs, bias currents are fed up from the dc power plane through bias pillars on the corner of each track block. Research on effective moat strategies resulted in an optimum moat configuration [298] that had only one disadvantage: ground connections between the multiple ground plane layers were placed far from the dc bias pillars.

For an investigation into coupling from ground plane return currents to circuit components [97], [299] I analysed the RSFQ AND gate from the CONNECT cell library [27]. The AND-gate layout is shown in Figure 5.1, and it consists of four track blocks tiled together in 2×2 format with the central bias pillar removed.




Figure 5.1: Mask layout of AND gate in CONNECT cell library with port definitions for InductEx modelling. The dimensions are $60 \mu\text{m} \times 60 \mu\text{m}$.

As depicted in Figure 5.2, I used InductEx to calculate the bias return current distribution *in the ground plane* and showed that the placement of ground contacts is suboptimal and that it creates significant coupling between the dc bias return current and other circuit components. The simple solution of adding minimum-sized ground contacts next to the bias pillars reduces current flowing in the ground plane away from bias lines and thus also reduces coupling to circuit components. I could formalise layout requirements for minimum interference from bias return currents that was applied to the design of the ColdFlux track block architecture (see Section 2.5.5.1).

5.1.0.3 Multiple ground planes

Analytical, sheet inductance and 2D inductance estimation techniques break down completely when inductors between multiple ground planes are analysed. In [300] we showed that, as soon as any part of a circuit is threaded through a ground plane in a layer stack with multiple ground planes, the *position of the ground pillars* that seam the different ground planes together affect inductance significantly.

As an experiment, Kyle Jackman and I designed for stripline configurations or layout patterns for the MIT Lincoln Laboratory SFQ5ee process, as shown in Figure 5.3. For each pattern, the second metal was fixed on layer M1, while the first was stepped between M2 and M6 while all the intermediate layers were filled in as ground planes as shown in Figure 5.4. Ground connections between the upper, lower and intermediate ground planes were varied for the layout patterns. From P1 to P4, the ground connections are brought successively closer to the signal vias to reduce the size of the ground return current loops and the stray coupling of flux between the two signal lines.

The results in Figure 5.5 clearly show how the mutual inductance between the two striplines varies by almost three orders of magnitude between P1 and P4. We showed




Figure 5.2: Modulus of current density in the main ground plane (M7) of (a) an AND gate in the CONNECT cell library for the AIST ADP2 process using the standard device structure when the cell is biased with 2.5 mV, and (b) the same AND gate with the same bias, but with the ground contact stacks placed around the dc bias pillars.





Figure 5.3: Four layout patterns of two striplines with length L_m separated by a distance S_m between the centre lines.



(a) Visualisation of stripline through ground planes. (b) Ground return current through *ground pillars*

Figure 5.4: A stripline in M6 connected to a stripline in M1, through holes in ground planes M2, M3, M4 and M5 of the MIT-LL SFQ5e process. Ground planes are connected with two ground pillars next to the striplines. Current distribution was calculated with InductEx. All layers were scaled vertically for visualization purposes.




Figure 5.5: Simulated results for layout P1 to P4 as a function of stripline spacing.

that, when signal (or bias) lines traverse ground planes, there will be significant coupling between the lines unless ground vias – or stitches – directly surround all signal vias. The layout cost has to be absorbed, and we applied this ground connection strategy to the ColdFlux track block [98].

5.2 Magnetic fields

The calculation of inductances in a layout structure requires solution of the current distribution throughout the meshed structure. With the Biot-Savart law it is then possible to calculate the magnetic field strength or flux density at any given point in space.

I implemented this method in a scripting language, but it was slow. A significantly faster, multipole accelerated method, was then developed by Dr Kyle Jackman and implemented directly in FFH and TetraHenry [301]. I added methods to InductEx with which to define volumes or planes over which to calculate fields and developed the required mesh control methods to do so efficiently.

We have applied magnetic field analysis widely, but initially used it to evaluate the influence of bias current lines – especially trunk lines that approach 100 mA of current – on nearby SFQ circuit structures [302]. This type of analysis enables us to find the best shielding techniques and closest acceptable layout distance between trunk bias lines and SFQ circuits in row-based place-and-route layouts.

An early application of InductEx to the analysis of magnetic fields on circuits involved modelling of influence of external magnetic fields on the behaviour of RSFQ circuits [303]. Initially, magnetic field influence was analysed through the construction of external coils, all generated by InductEx, as part of an extraction model. Extracted circuits were then simulated with an electrical simulator to determine the effect of the coil-modelled field on circuit operating margins [304].

After the inclusion of magnetic field excitation directly into FFH and TetraHenry, the artificial coil constructs could be replaced with direct excitation of a magnetic field in all three axial directions. I showed an application where the external operating margins of a circuit were analysed as a function of external magnetic field strength in each axial direction [305]. The circuit analysed was modelled from the same circuit that was tested by Collot [306] and contained an RSFQ DC-to-SFQ converter, two JTLs and an SFQ-to-DC converter. The margin simulation results from the InductEx-extracted model are shown in Figure 5.6. These compare well to the measured results in [306], which shows that circuits can be analysed for response to a magnetic environment through the use of extracted models.

5.3 Digital circuits

The first application of InductEx was to digital circuit layout verification, and application is so widespread that anything approaching a comprehensive discussion would consume volumes. I focus on a small few other applications only.

5.3.1 Coupling from bias lines

Unwanted coupling between conductors, whether it is between clock and signal lines in AQFP, or between bias lines and flux storage loops in SFQ circuits, can create headaches

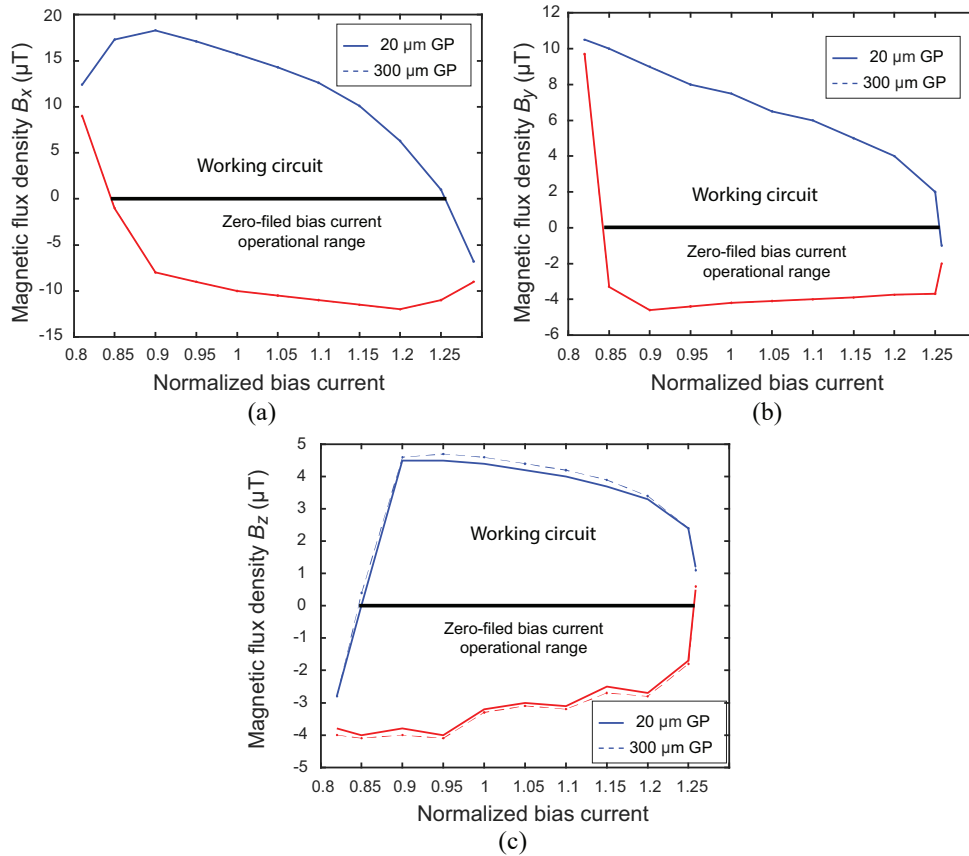


Figure 5.6: Simulated operating field margins of an RSFQ circuit with DC-to-SFQ converter, two JTLs and an SFQ-to-DC converter as a function of normalized bias current for (a) the x-directed field, (b) the y-directed field and (c) the z-directed field. Solid lines and dashed lines show margins for the models with a small and large ground planes that extend 20 μm and 300 μm beyond the circuit layout respectively.

for circuit designers. All coupling can be extracted with a standard multiterminal inductance extraction using InductEx, but one of the applications of InductEx is to specifically drive a control structure and to calculate the induced current in a victim loop.

For the design of ColdFlux RSFQ cell library, I analysed the track block layout for how well a sensitive layout, such as a SQUID loop, is isolated from a large current-carrying dc bias line [98]. The model and simulation result shown in Figure 5.7 shows the current density everywhere in a structure laid out for the MIT Lincoln Laboratory SFQ5ee process with the ColdFlux track block architecture, when a bias current line is placed at minimum separation of $10\ \mu\text{m}$ – in the closest track parallel to a SQUID loop – and excited with 100 mA of current. The results show that, for a bias line in M5 shielded with a sky plane layer in M6, and with the victim SQUID shielded by a skyplane in M7, the bias line can carry 345 mA before the current in the SQUID loop changes by 1%. With a limit of 100 mA imposed by the ColdFlux place and route rules, the track block layout thus allows bias lines to be routed in the closest tracks near SFQ layouts without fear of bias current-induced circuit failure.




Figure 5.7: Current distribution as calculated with InductEx for a bias line in M7 over a solid ground plane in M4 near an unshielded victim SQUID (on the left) and a caged bias line in M5 near a victim SQUID shielded with an M7 skyplane (on the right). The current density scale is in A/m^2 . Both bias lines are excited with a current of 100 mA.

5.3.2 Inductive SFQ pulse transfer

SFQ circuits are low power devices where switching energy is defined by the the magnetic flux quantum multiplied by the current – typically the bias current of a switching junction – associated with the flux quantum in a circuit. Switching energy is thus approximately $2 \times 10^{-19}\ \text{J}$ to $4 \times 10^{-19}\ \text{J}$ [55] for RSFQ circuits with junctions biased between $100\ \mu\text{A}$ and $200\ \mu\text{A}$. However, even in implementations where static power dissipation in dc biased logic such as RSFQ and its derivative families approaches zero, there still remains dc bias currents. For a typical RSFQ or related logic gate, the bias current is in the order of 1 mA. For a million-gate circuit, this translates to 1 kA of bias current into the chip.

The energy cost to regulating kiloamperes of stable dc current into what is essentially a short-circuited load limits the maximum efficiency of large scale systems. Added to

Table 5.1: Measured and InductEx-extracted inductances of inductively coupled pulse transmission cell.

	Self inductance (pH)	Mutual inductance (pH)
Calculated	13.4	10.1
Measured	14.0	9.93

that, the magnitude of the dc current creates enormous magnetic fields in an environment where the electronics are extremely sensitive to such fields.

For these reasons, methods to reuse or recycle current in SFQ circuits have often been investigated.

One solution is the use of serial biasing – the reuse of bias current for blocks of circuits with equal bias current requirements – where each block has a floating ground from which all bias return current is picked up and funneled to the bias input of the next circuit block [140], [141], [143]. The signal inputs and outputs between different blocks must be galvanically isolated, so that inductively coupled interconnects over edges of the ground plane islands for every block are used.

With InductEx demonstrated to be accurate for inductors over holes [108], I used it to design an inductively coupled transmission (TX) cell with an unsymmetrical transformer and a ground plane hole on the receiver side to improve the mutual inductance. The circuit schematic is shown in Figure 5.8 and the component values are presented in [61].

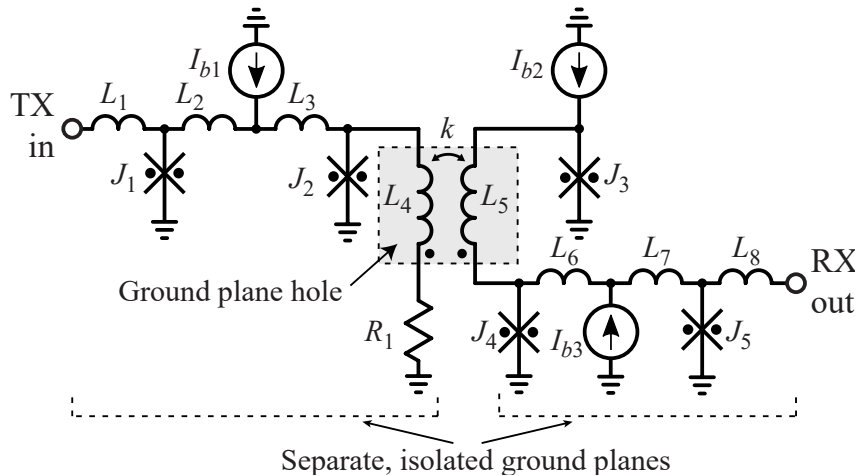


Figure 5.8: Schematic circuit diagram of inductively coupled pulse transfer cell.

From circuit simulation, the critical margin was predicted as approximately $\pm 10\%$ for I_{b2} . The circuit was fabricated in the FLUXONICS process and tested by Dr Olaf Wetzstein at Leibniz IPHT. A microphotograph is shown in Figure 5.9. The InductEx extraction model, using cuboid segments for the FastHenry engine, is shown in Figure 5.10. The cuboid mesh model used 21 000 filaments and solved in 10 minutes in 2012.

Successful pulse transmission was measured in liquid helium for low frequency tests (below 1 MHz), and the bias margins were measured as $\pm 5\%$.

The chip also included a test structure to verify the inductance and mutual inductance in the presence of isolated ground planes and a hole, and experimentally measured values compared to results extracted with InductEx are shown in Table 5.1. The measurements are the average from eight chips over two wafers.

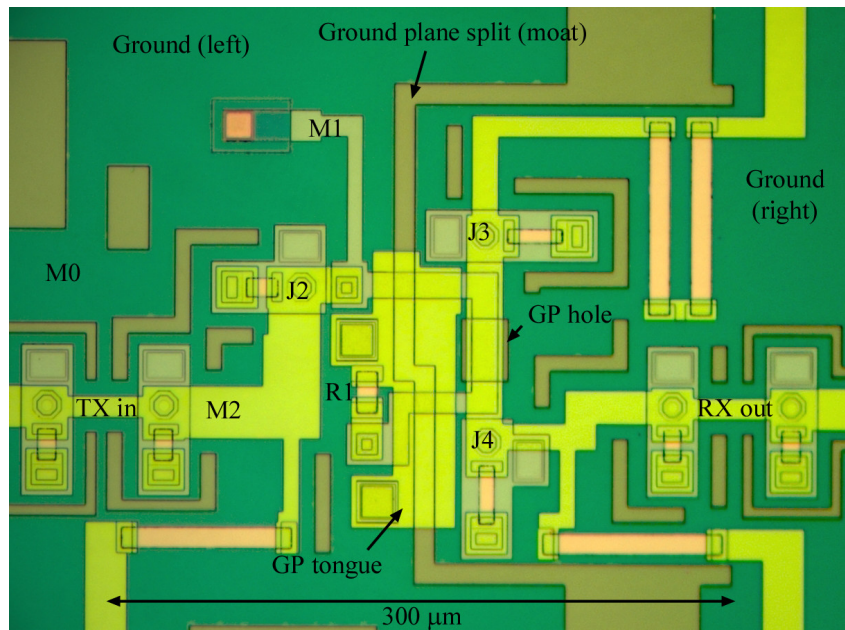


Figure 5.9: Microphotograph of the inductively coupled pulse transmission cell.

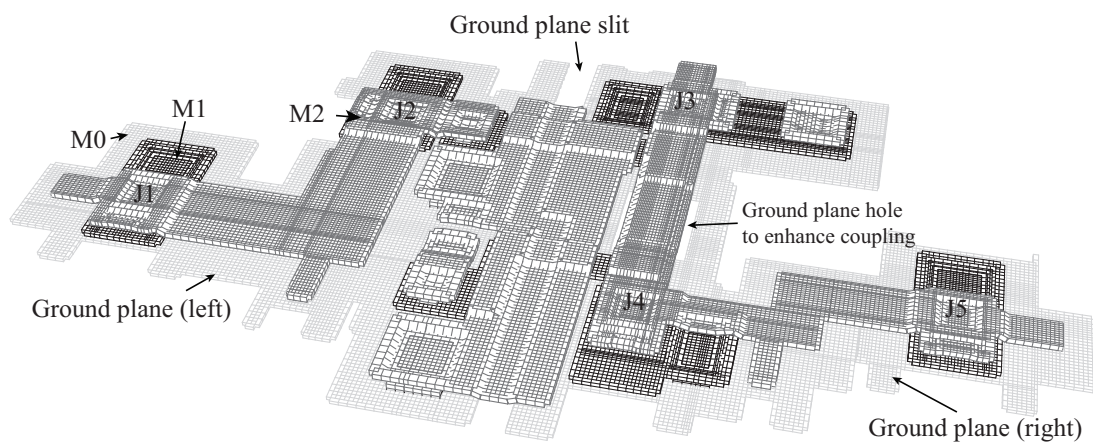


Figure 5.10: InductEx extraction model of complete inductively coupled pulse transmission cell.

Other successful results have been demonstrated [142], [144], [307], and a recent demonstration of 16 stacks of serially biased shift registers by Semenov and Polyakov relied on InductEx for design of the floating couplers [308].

5.4 Analogue devices: SQUID magnetometers

Magnetometry is a key application for dc SQUIDS [235]. For good magnetic field resolution, a large pickup area is required. However, some limitations apply. Foremost is the inductance associated with a pickup loop: larger loop area invariably means larger loop inductance.

It can be shown [2] that, for thermal noise not to degrade sensitivity of a SQUID, the magnetic energy stored in the SQUID loop should be larger than the thermal energy. Thus

$$L < \frac{\Phi_0^2}{k_B T}, \quad (5.1)$$

where k_B is the Boltzmann constant and T is temperature in kelvin. For a high temperature SQUID operating at 77 K, the loop inductance should thus be significantly smaller than 4 nH – or practically smaller than 1 nH. There is our first constraint.

Furthermore the modulation or screening parameter, β_L , defined as

$$\beta_L = \frac{2LI_C}{\Phi_0}, \quad (5.2)$$

is mostly designed to be close to 1 for optimal energy resolution [309].

Junction critical current should be large enough so that junction coupling energy is larger than thermal energy [309], which limits the lowest critical current at 77 K to around 10 μ A. For a dc SQUID with 10 μ A critical current for each junction, loop inductance L should thus be about 100 pH. This limits the loop area – and thus the pickup area – of the SQUID. The obvious solution for improved magnetic field resolution is to couple a much larger pickup coil to the SQUID loop in such a way that the total SQUID loop inductance does not cause β_L to deviate much from 1.

Moreover, if the SQUID has a root-mean-square magnetic flux noise $S_\Phi^{1/2}(f)$, then the rms magnetic field resolution of a magnetometer with a (presumed noiseless) flux transformer improves to $S_B^{1/2} = S_\Phi^{1/2}(f)/A_{eff}$ where A_{eff} is the effective area of the magnetometer [309], [310].

The effective area can be orders of magnitude larger than the SQUID loop area. It can also be expressed in terms of field-to-flux conversion efficiency

$$B_\Phi = \frac{\Phi_0}{A_{eff}}. \quad (5.3)$$

Technically, B_Φ is the magnetic flux density sensitivity (flux sensitivity for short), although it is commonly referred to as the (magnetic) field sensitivity. For simplicity it is referred to as field sensitivity here.

5.4.1 SQUID parameters of interest

For SQUID design the effective area is of interest. Traditionally, SQUID designs stick to well-defined geometries for which analytical formulae have been presented for effective

area [309] in terms of SQUID loop effective area and the inductance of the pickup loop – both of which are also derived from formulae for simple, well-defined geometries – and the coupling factor between the pickup and loop inductance. Most of these formulae have been derived empirically, such as that for the inductance of a square washer [311], [312]. The drawback of such empirical methods is that the formulae only hold within constricted parameter ranges, and that design of arbitrary or any complex shapes for the pickup loop is not practical.

My contribution to SQUID magnetometry, through my work on InductEx, is the tremendous simplification of the calculation of SQUID loop inductance, pickup loop inductance, effective area and thus field sensitivity. This is demonstrated in the next section.

5.4.2 Analysis of a planar direct-coupled SQUID: the M2700

A high temperature SQUID, an M2700 made by Star Cryoelectronics from a YBCO monolayer deposited on a bicrystal substrate, is depicted in Fig. 5.11 and Fig. 5.12. It is a smaller version of the M1000 SQUID described in [309].




Figure 5.11: A packaged M2700 YBCO SQUID from Star Cryoelectronics with heater, feedback coil and output transformer included. The pickup loop is visible as a dark rectangle in the centre of the disc on the left.

The package dimensions are obtained from the device data sheet and physical measurements with both vernier calipers and calibrated microscope images. For modelling, the feedback coil outer dimension is set to 3.96 mm, the length to 2.0 mm, and the depth of the coil below the SQUID die plane as 2.46 mm. The wire diameter is 76 μm . The coil has 175 turns, which is modelled with 10 up-down wrapping layers along the length of the coil for InductEx extractions.

The coil is generated with ObjectBuilder – an InductEx auxiliary tool developed specifically for such structures.

For calculation of the field sensitivity, three modelling techniques are shown. Equivalence demonstrates that any technique can be used.




Figure 5.12: The M2700 SQUID from Star Cryoelectronics viewed close up. The feedback coil is packaged below the square washer, and is not visible.

5.4.2.1 Full inductance circuit model

The M2700 SQUID magnetometer has a rectangular SQUID loop directly coupled to a larger pickup coil implemented as a square washer. Due to the presence of a second SQUID – which can be bonded in if the first does not function as expected – the bias feed line is slightly offset from the active SQUID. A full simulation model can be constructed by hand, and the equivalent InductEx extraction model can then be sketched as shown in Fig. 5.13. Ports J_1 and J_2 represent the Josephson junctions, L_s is the inductance of the small SQUID loop, and L_p is the inductance of the pickup loop. The inductance L_x models the offset location for bias current removal, which is in the centre of two SQUIDs of which only one is wired into the circuit. The circuit netlist is defined as depicted here, except that the external field source and inductance are not included, because InductEx adds that. A section of the InductEx extraction model is shown in Figure 5.14 to better illustrate the bias pin offset, the SQUID loop and the unused SQUID.

It is assumed that at low frequency (for magnetic field measurement), and definitely at dc, there is no coupling to the resistive bias lines. For the InductEx extraction model, the bias branch is thus flagged as resistive and it is subsequently omitted when coupling from the external field to the circuit is calculated by InductEx.

With the full circuit netlist, it is not possible to include mutual inductance between the pickup loop inductor and the SQUID loop inductor during inductance extraction, as this leads to a rank deficient branch current matrix with no good solution. This is not a problem, though, as these inductors are electrically connected in the circuit model, so that magnetic coupling between the two is absorbed into the individual self-inductances.

The InductEx extraction results are back-annotated to a simulation netlist, shown in Figure 5.15. The x - and y -directed field components have been omitted here for brevity. An external field is applied in the z direction (perpendicular to the plane of the pickup loop) in simulation by the field current source I_{FieldZ} , with a current-field conversion

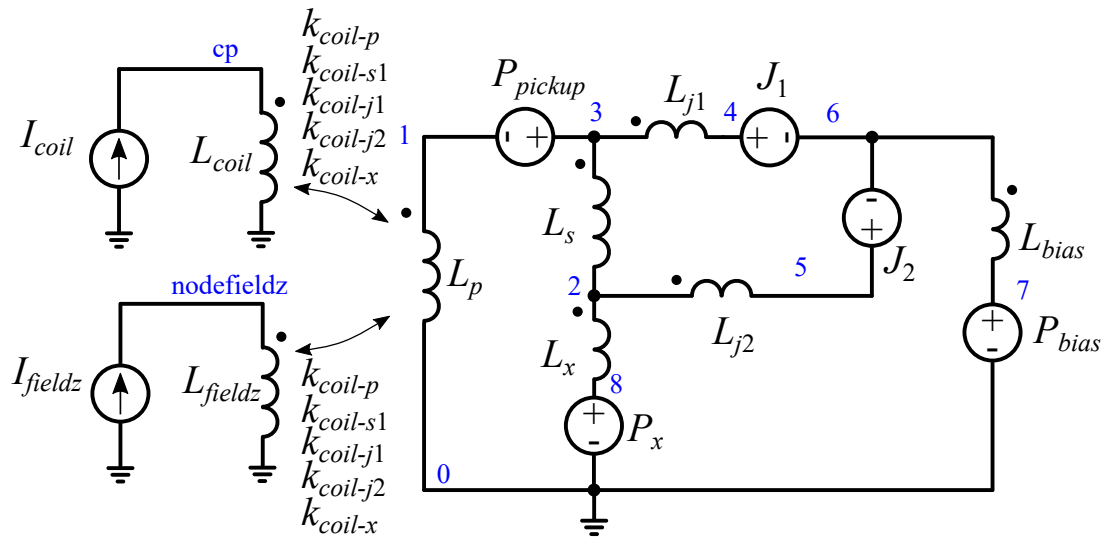


Figure 5.13: Netlist of M2700 with feedback coil and external magnetic field current sources and equivalent inductances included for parameter extraction.

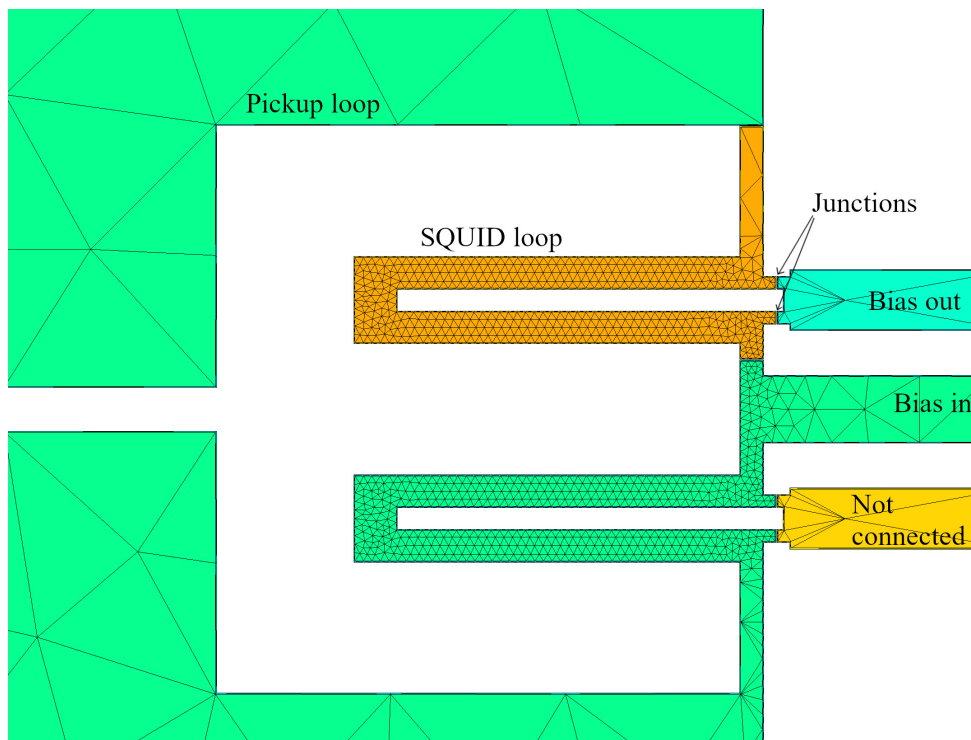


Figure 5.14: Close-up view of the InductEx segmented mesh for the M2700 extraction model, showing only the active and unused SQUID loops, the bias pins and the connection to the pickup loop.

factor of 1 A/T. The simulation sweeps the external field from 0 to 100 nT in 1 ns. The simulated voltage over the SQUID is shown in Figure 5.16

```

* Back-annotated simulation file written by InductEx
* External magnetic field coupling (z-directed):
IFieldz 0 NodeFieldz pwl(0 0 1000p 100n)
LFieldz NodeFieldz 0 1
KFZ0 LFIELDZ LJ2 0.123051103591173
KFZ1 LFIELDZ LS1 0.0757035879327776
KFZ2 LFIELDZ LJ1 0.249062061513407
KFZ3 LFIELDZ LX 0.0700479428665593
KFZ4 LFIELDZ LP -0.0345772692980707
* SQUID simulation model for M2700
B1 4 6 jj1 area=1
B2 5 6 jj1 area=1
LJ1 3 4 1.354E-011
LJ2 2 5 6.231E-012
Lx 2 0 1.827E-011
Rbias 6 7 1
IBias 0 7 pwl(0 0 5p 25u)
LS 3 2 6.42E-011
LP 3 0 1.94E-009
.model jj1 jj(rtype=0, vg=100uV, cap=85fF, rn=14ohm, icrit=0.01mA)
.tran 0.25p 1000p 0 0.25p
.plot v(6)
.plot i(lfieldz)
.END

```

Figure 5.15: JoSIM simulation netlist of the full M2700 SQUID model in a z -directed field.




Figure 5.16: JoSIM simulation output of for the full M2700 SQUID simulation model with flux modulation. Flux density in the z direction is swept from 0 to 100 nT in 1 ns.

From the graph in Figure 5.16 the change in flux density for one period is read as 25.9 nT, so that the field sensitivity is $25.9 \text{ nT}/\Phi_0$. This is slightly better than the value of $33 \text{ nT}/\Phi_0$ obtained from the SQUID's calibration document (a lower value is better, as it is more sensitive to magnetic flux density). The field sensitivity is sensitive to the London penetration depth (λ was assumed as 240 nm here) and is very sensitive to the area enclosed by (and thus the inductance of) the SQUID loop inductor L_s , so that the difference can be ascribed to both the uncertainty in penetration depth and the uncertainty in the dimensions of the SQUID loop.

The effective area is thus $A_{eff} = \frac{2.067e-15}{25.9e-9} = 7.98 \times 10^{-8} \text{ m}^2$.

5.4.2.2 Compact simulation model

A compact simulation model can be extracted with the InductEx compact model extraction tool. The resulting circuit has three inductive branches: L_s , L_p and L_{bias} , and includes the coupling between these inductors as well as the coupling from the external field. When simulated with JoSIM for an external field sweep, the measured field sensitivity is 25.0 nT.

5.4.2.3 Single inductance model

While the full and compact models are required to handle the asymmetry of the SQUID (due to the offset bias pins), the field sensitivity can be found without the use of an electrical simulation if the SQUID loop is reduced to a single equivalent inductance L_T that combines L_s , L_p and the parasitic inductances into one.




Figure 5.17: Schematic of a 2-junction SQUID with one equivalent inductance L_T and coupling from a z -directed field.

In this case, the InductEx analysis for the SQUID layout in an external magnetic field delivers a mutual inductance between each axial field direction and the single SQUID loop inductance. InductEx was developed to find the mutual inductance between the magnetic field in every axial direction and every inductor in a layout. The mutual inductance is calculated for a model where the external magnetic field is a current source with amplitude of 1 ampere per tesla, and the field inductance (coupled to the inductors in the circuit) is 1 H.

Phase change over the SQUID loop inductor due to the external field is thus

$$\delta(\phi_1 - \phi_2) = \frac{2\pi}{\Phi_0} M \delta I_{field}. \quad (5.4)$$

The field current required to let $\delta(\phi_1 - \phi_2) = 2\pi$, which is equivalent to increasing the flux through the SQUID loop by Φ_0 , is then

$$\delta I_{field} = \frac{\Phi_0}{M}. \quad (5.5)$$

Since δI_{field} is equivalent to the magnetic field sensitivity for any applied field direction, it follows that

$$B_{\Phi} = \frac{\Phi_0}{M} \quad (5.6)$$

and that the mutual inductance calculated for such a model with InductEx for is equal to the effective area A_{eff} of the SQUID.

The mutual inductance between the total SQUID inductance L_T and the equivalent inductance selected for the z -directed field is calculated through an easy setup with InductEx as 7.99×10^{-8} H, which equals the effective area found with the full inductance model above. The field sensitivity is 25.9 nT.

The simplification of the calculation of effective area of a SQUID from the actual layout dimensions, rather than from analytical guesswork, is one of my contributions to the research field.




Figure 5.18: Mesh of M2700 SQUID and feedback coil generated with InductEx.

5.4.2.4 Feedback coil coupling

For the feedback coil with 10 wraps, the coil inductance is extracted as $51.6 \mu\text{H}$ with a coil resistance of 7.3Ω if the bulk conductivity of the coil copper is assumed as 58.7 S m^{-1} . The field coupling is obtained from simulation of the SQUID in JoSIM. The SQUID is biased with $25 \mu\text{A}$ and the voltage period is measured when current through the feedback coil is swept up. One period of voltage modulation equals one fluxon coupled through

the SQUID loop, so that we can calculate the feedback coil coupling from the simulation plots as before. For this SQUID, the feedback coil coupling is found as $6.1 \mu\text{A}/\Phi_0$.

The feedback coil coupling is very sensitive to the position of the coil relative to the SQUID pickup loop. For the H225 SQUID, the calibrated feedback coupling is $13.2 \mu\text{A}/\Phi_0$. Minor changes in the vertical position or alignment of the coil (which we can do easily with InductEx) yields feedback coupling changes that fit with the calibrated measurements.

In summary, the full inductance mesh method gives results that are close to the calibrated characteristics of the SQUID and the coil, especially considering that many dimensions and material parameters are obtained from best guesses.

5.4.3 Flux noise

In 2013, Dr Steven Anton in the group of Prof. John Clarke at UC Berkeley investigated a numerical method to calculate mean square flux noise in SQUIDs and qubits. I calculated current distribution in a nonuniform mesh for his method with InductEx, and the results were published [313].

5.5 Flux trapping

One of the most consequential developments, if not *the* most, of my group's research and development of inductance extraction tools and phase-based simulators is the capability to analyse circuit behaviour in the presence of trapped flux. At the start of SuperTools, this was a requirement that we knew we had to meet, but that we did not know how to do, or even if we could do it at all. The successful demonstration and verification of flux trapping analysis as part of the InductEx tool chain represents the culmination of almost a decade of research effort, and is a fitting conclusion for this dissertation.

5.5.1 Background to flux trapping

It is known that magnetic flux frozen into a superconductor circuit structure during cooldown causes deviation in operating margins and in the worst case leads to circuit failure. It is one of the more serious problems associated with high density superconductor integrated circuits.

Flux trapping in narrow films was already observed and reported in 1982 [314]. The *moat* as a flux trapping mitigation device was presented in 1983 [315] for SQUIDs that relied on small holes in the ground plane to improve or engineer inductance. If the energy of a vortex occupying a moat is lower than that of the energy of a vortex occupying a nearby hole or pinning location, a fluxon is more likely to trap in the moat.

There are two ways for a trapped fluxon to affect a circuit [316]:

- If a vortex is frozen into a Josephson junction (penetrating both electrodes), the effective area and thus its critical current is reduced. The junction malfunctions if a fluxon is frozen into only one electrode.
- If a vortex is frozen into moat or any location outside of a Josephson junction, magnetic coupling to sensitive circuit structures and the current induced by that coupling can alter circuit behaviour.

It was believed [316] and shown experimentally since then that deviation in expected circuit behaviour due to magnetic coupling from fluxons in moats is much more common than that of fluxons frozen into Josephson junction electrodes. In 2007, no analysis tools existed that could allow circuit designers to model the coupling from trapped fluxons to a circuit layout, although 3D-MLSI [185] could at least calculate the current distribution from a trapped fluxon. It was postulated that such coupling could be completely eliminated or at least dramatically reduced.

For every 1 μT of magnetic flux density vertical to the surface of a superconductor chip at cooldown, five fluxons are trapped per $100\ \mu\text{m} \times 100\ \mu\text{m}$ circuit area. Circuit design thus includes mitigation strategies, such as moats in ground planes. If the number of fluxons per unit area exceeds the moat count, Pearl vortex formation is very likely. This can be easily be avoided by having a sufficient number of moats in a given area for specified maximum field tolerance.

The more important design question – one that cannot simply be waved away by adding more moats – is if a circuit will be fully operational when a fluxon is captured in any (or all) moats and if the moat placement or flux mitigation strategy can be engineered to ensure that the circuit is operational for any flux trapping combination.

SQUID microscopy has been used to investigate flux trapping in moats [317], but the work by Jeffery *et al.* used unbiased circuits, with no information available on circuit operation. They concluded that moats were effective at trapping fluxons, and that long (continuous) moats were better at higher flux density than shorter moats. However, these experiments were performed with moats that were more than $100\ \mu\text{m}$ apart.

Experiments on working circuits were conducted and published by Robertazzi *et al.* at Hypres in 1997 [318]. The conclusion was that very large moats, of about $3000\ \mu\text{m} \times 50\ \mu\text{m}$ outside all gates showed the largest operating margins, although the best flux density performance was obtained for many small holes. Circuits failed at a maximum flux density of 1 μT for all the moat strategies.

A thorough investigation of moat efficacy was published by Naryana *et al.* [319], with earlier measurements described in [316]. The results showed higher operating field margins when moats were long and closely spaced, and could find no obvious influence from moat shape. With moats spaced by between $20\ \mu\text{m}$ and $30\ \mu\text{m}$, circuits operated up to a cooldown flux density of 2 μT .

A simulation technique that uses InductEx to find the coupling between a moat and a circuit structure was shown by Yamanashi *et al.* [320]. The method approximated a hole in the ground plane as a loop around the ground plane with width equal to the penetration depth. One big drawback was that the rest of the ground plane had to be omitted, while another was that the conversion of moat inductance to current induced per fluxon was not trivial. Still, the results were reasonable for microstrip inductors.

5.5.2 Modelling of flux trapping

With the inductance of a hole and the coupling to every circuit inductance and other hole inductances calculated with InductEx as described in Section 3.3.9, a circuit model for simulation with a phase-based simulator such as JoSIM can be constructed as shown in Figure 5.19.

For electrical simulation, flux is applied to each hole as a phase of $2\pi n$ rad, where n is the signed number of fluxons in the hole, and the sign represents polarity. This method models hole-to-hole and circuit-to-hole coupling correctly, and improves on the incomplete




Figure 5.19: Circuit schematic of a flux trapping hole coupled to the inductors in a SQUID circuit.

model where current sources excite holes for simulation engines that cannot handle phase (as I described in [305]).

5.5.3 Verification of flux trapping

The ability to add a flux port to an inductance extraction model was added to FFH and TetraHenry by Dr Kyle Jackman [321]. Under the ColdFlux project, Kyle and I designed a set of flux linkage experiments [322], [323] to test the software and modelling methods. The experiments were fabricated with the MIT Lincoln Laboratory SFQ5ee process and were used to evaluate the magnetic coupling of different moat configurations to superconducting microstrip lines. All self and mutual inductances were determined from measurements of the SQUID I_C vs flux modulation period, while SQUID critical currents were measured from the current-voltage curves for swept bias current. A JoSIM simulation model was constructed for each experiment and populated with element values extracted from InductEx. For each flux linkage experiment, an on-chip coil, with inductance extracted using InductEx, was used to apply controlled magnetic flux density to the experiment while it was cooled in a magnetically shielded environment. After each cooldown, I_C of the SQUID was measured. For each experiment, the simulated critical current was obtained from an I-V curve generated by JoSIM from the simulation model.

Five of these experiments are shown as layout schematics in Figure 5.20.

For each of the flux linkage experiments, a controlled field was applied through a on-chip coil that focuses on the SQUID loop inductor and the moats in the immediate vicinity. The experimental setup is explained in detail in [322], and an illustration is shown in Figure 5.21. Calculation of the magnetic fields from which the stream lines for this image was created, is done with TetraHenry and was made faster than an early direct Biot-Savart implementation through a multipole accelerated field calculation [301].

Experimental and simulation results are shown in Figure 5.22. The excellent agreement between measurements and simulations validate both the compact circuit model and the parameter extraction results, which can thus be used with confidence. The results are analysed in detail in [323]. Briefly, we confirmed with experimental results what we predicted from simulation: moats closer to an inductor couple more strongly to the

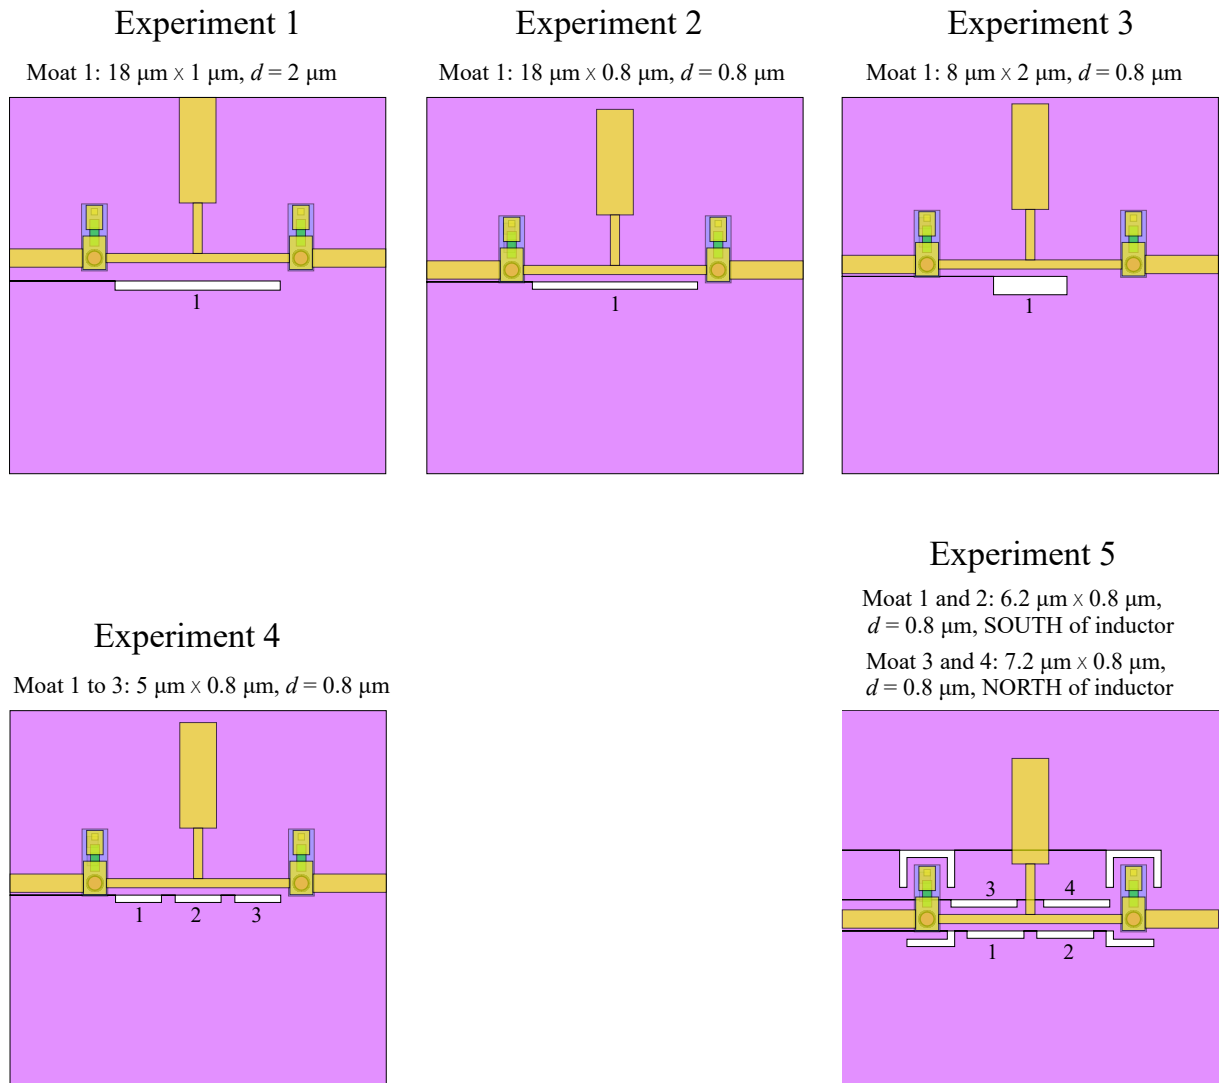


Figure 5.20: Layout showing just the test SQUID and moat configurations for five flux linkage experiments.




Figure 5.21: Overlay of microscope photograph and InductEx simulation of the magnetic field created by a trapped fluxon for a test SQUID.

inductor and changes the current distribution in the SQUID lead to changes in the critical current of between 5% and 10% for the closest, longest moats; moats fill up roughly evenly if they are in close proximity to each other, and are unlikely to take two fluxons when a nearby moat remains empty; and Pearl vortex formation in the ground plane is more likely once nearby moats all contain one fluxon.

From the flux linkage experiments and moat analysis with simulation, we were able to formulate moat design guidelines:

- There should be more moats in a given circuit area than expected fluxons for the target maximum magnetic flux density after shielding. Double trapping in any moat is unlikely in a uniform field before Pearl vortices appear.
- Moats should not run the entire length of a critical inductor.
- Moats perpendicular to an inductor have lower coupling to the inductor than those in parallel to an inductor.
- Moats should be evenly staggered on both sides of an inductor.
- Moats matched on ground and sky planes below and above a circuit layout should be used wherever possible, as this significantly reduces fluxon coupling to circuit inductors.

We have recently applied the analysis of circuit behaviour to trapped flux to sensitive circuits. One such application was the improvement of AQFP layouts [124].

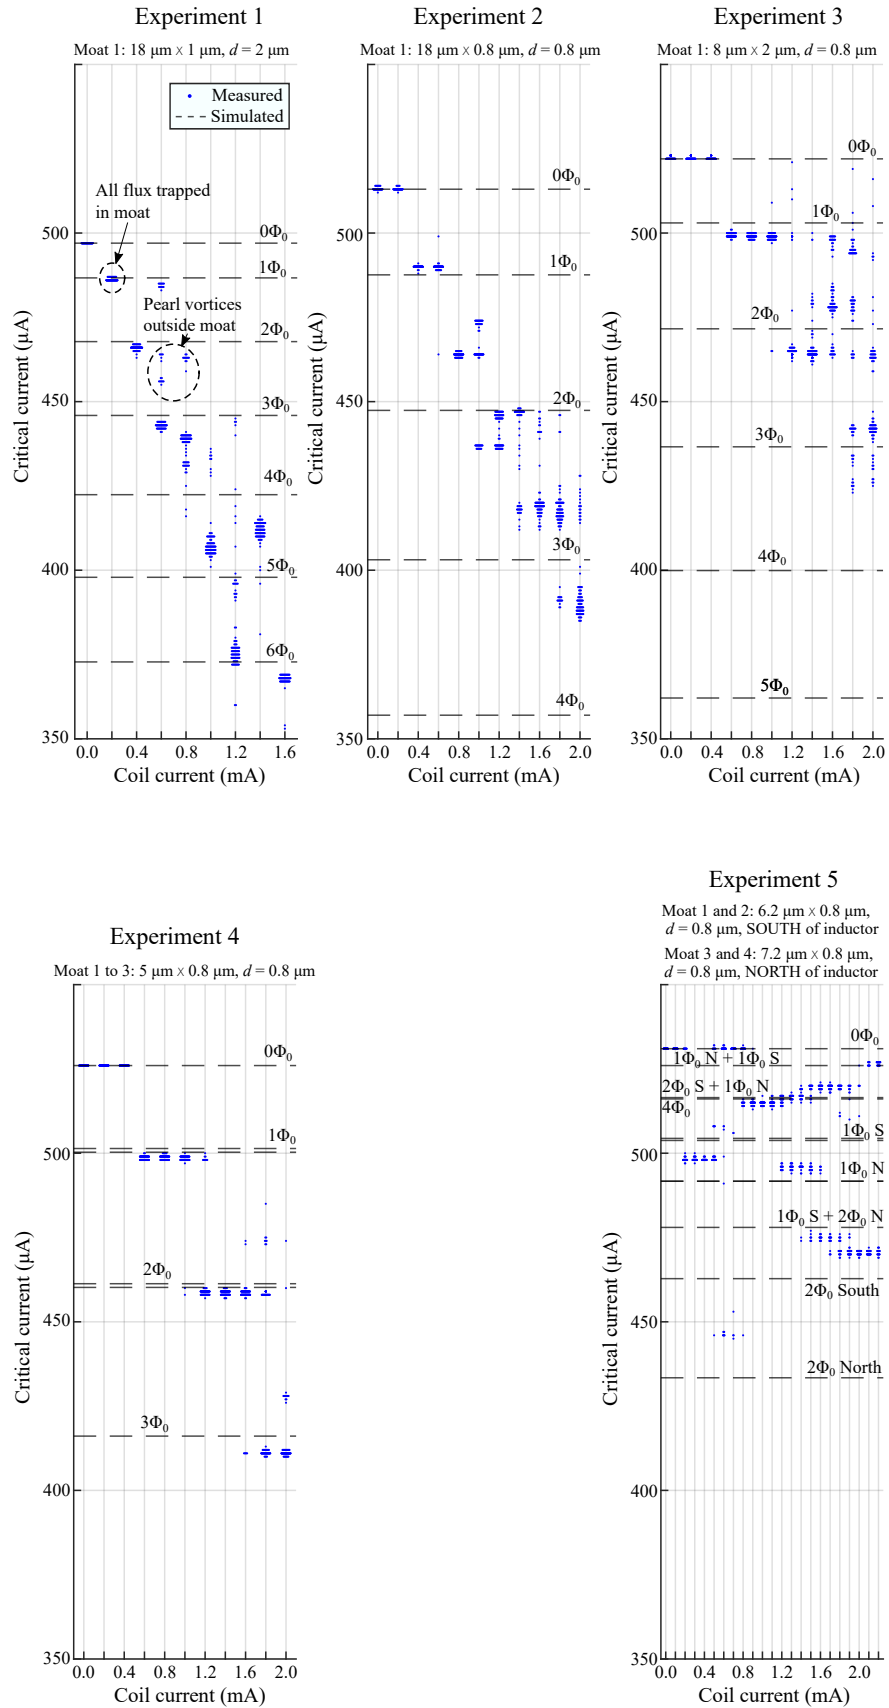


Figure 5.22: Critical current of SQUIDs in flux linkage experiments. All simulated trapped flux used flux return path around the edge of the chip.

Chapter 6

Final conclusion

This dissertation covers the main contribution of my research career, which started in 1999 in the Department of Electrical and Electronic Engineering at Stellenbosch University, to the field of applied superconductivity.

From RSFQ circuit design with a lack of tools for layout verification, I progressed to the research on and development of electronic design automation software. Inductance extraction was the key capability that could link many of the tools together. My research group and I developed and verified inductance extraction tools for complex, three-dimensional integrated circuit models, and applied self- and mutual inductance extraction and magnetic field analysis to the improvement of superconductor circuit and system design.

The research culminated in methods with which to extract compact simulation models for the analysis of superconductor integrated circuits in the presence of trapped flux and external magnetic fields. This provides much-needed design and verification capabilities to physicists and engineers in applied superconductivity as applications shift towards superconductor devices and circuits for quantum computing systems.

References

- [1] P. Bunyk and S. Rylov, “Automated calculation of mutual inductance matrices of multilayer superconductor integrated circuits,” *Proc. Ext. Abstracts 4th Int. Supercond. Electron. Conf. (ISEC’93), Boulder, CO*, p. 62, 1993.
- [2] T. P. Orlando and K. A. Delin, *Foundations of applied superconductivity*. Reading, MA: Addison-Wesley, 1991.
- [3] T. Van Duzer and C. W. Turner, *Superconductive devices and circuits*. Upper Saddle River, NJ: Prentice-Hall, 1999.
- [4] C. C. Tsuei, J. R. Kirtley, C. C. Chi, L. S. Yu-Jahnes, A. Gupta, T. Shaw, J. Z. Sun, and M. B. Ketchen, “Pairing symmetry and flux quantization in a tricrystal superconducting ring of $\text{YBa}_2\text{Cu}_3\text{O}_{7-\delta}$,” *Physical Review Letters*, vol. 73, pp. 593–596, 4 Jul. 1994.
- [5] C. Tsuei, J. Kirtley, M. Rupp, J. Sun, C. Chi, A. Gupta, S. Y.-J. Lock, and M. Ketchen, “Half-integer flux quantum effect in tricrystal cuprate superconductors,” *Physica C: Superconductivity*, vol. 263, no. 1, pp. 232–237, 1996.
- [6] B. Josephson, “Possible new effects in superconductive tunnelling,” *Physics Letters*, vol. 1, no. 7, pp. 251–253, 1962.
- [7] B. D. Josephson, “The discovery of tunnelling supercurrents,” *Reviews of Modern Physics*, vol. 46, pp. 251–254, 2 Apr. 1974.
- [8] A. A. Golubov, M. Y. Kupriyanov, and E. Il’ichev, “The current-phase relation in Josephson junctions,” *Reviews of Modern Physics*, vol. 76, pp. 411–469, Apr. 2004.
- [9] S. K. Tolpygo, V. Bolkhovsky, T. J. Weir, A. Wynn, D. E. Oates, L. M. Johnson, and M. A. Gouker, “Advanced fabrication processes for superconducting very large-scale integrated circuits,” *IEEE Transactions on Applied Superconductivity*, vol. 26, no. 3, Apr. 2016, Art. no. 1100110.
- [10] M. Schmelz, “Development of a high sensitive receiver system for transient electromagnetics,” Ph.D. dissertation, University of Twente, 2013.
- [11] H. F. Herbst, “Gate-level superconductor integrated circuit fabrication process modelling for improved layout extraction,” Masters thesis, Stellenbosch University, 2021.
- [12] S. K. Tolpygo, V. Bolkhovsky, T. J. Weir, L. M. Johnson, M. A. Gouker, and W. D. Oliver, “Fabrication process and properties of fully-planarized deep-submicron Nb/Al–AlOx/Nb Josephson junctions for VLSI circuits,” *IEEE Transactions on Applied Superconductivity*, vol. 25, no. 3, Jun. 2015, Art. no. 1101312.
- [13] S. A. Cybart, S. M. Anton, S. M. Wu, J. Clarke, and R. C. Dynes, “Very large scale integration of nanopatterned $\text{YBa}_2\text{Cu}_3\text{O}_{7-\delta}$ Josephson junctions in a two-dimensional array,” *Nano Letters*, vol. 9, no. 10, pp. 3581–3585, 2009.

- [14] N. Bergeal, J. Lesueur, M. Sirena, G. Faini, M. Aprili, J. P. Contour, and B. Leridon, "Using ion irradiation to make high-Tc Josephson junctions," *Journal of Applied Physics*, vol. 102, no. 8, 2007, Art. no. 083903.
- [15] H. Toepfer, T. Ortlepp, H. F. Uhlmann, D. Cassel, and M. Siegel, "Design of HTS RSFQ circuits," *Physica C: Superconductivity*, vol. 392-396, pp. 1420–1425, 2003.
- [16] V. Schultze, R. Stolz, R. Ijsselsteijn, V. Zakosarenko, L. Fritzsche, F. Thrum, E. Il'ichev, and H.-G. Meyer, "Integrated SQUID gradiometers for measurement in disturbed environments," *IEEE Transactions on Applied Superconductivity*, vol. 7, no. 2, pp. 3473–3476, 1997.
- [17] S. K. H. Lam, R. Cantor, J. Lazar, K. E. Leslie, J. Du, S. T. Keenan, and C. P. Foley, "Low-noise single-layer $\text{YBa}_2\text{Cu}_3\text{O}_{7-\delta}$ dc superconducting quantum interference devices magnetometers based on step-edge junctions," *Journal of Applied Physics*, vol. 113, no. 12, 2013, Art. no. 123905.
- [18] L. L. Kaczmarek, R. Ijsselsteijn, V. Zakosarenko, A. Chwala, H.-G. Meyer, M. Meyer, and R. Stolz, "Advanced HTS DC SQUIDS with step-edge Josephson junctions for geophysical applications," *IEEE Transactions on Applied Superconductivity*, vol. 28, no. 7, 2018, Art. no. 1601805.
- [19] A. Steppke, C. Becker, V. Grosse, L. Dörrer, F. Schmidl, P. Seidel, M. Djupmyr, and J. Albrecht, "Planar high-Tc superconducting quantum interference device gradiometer for simultaneous measurements of two magnetic field gradients," *Applied Physics Letters*, vol. 92, no. 12, 2008, Art. no. 122504.
- [20] T. Wolf, N. Bergeal, J. Lesueur, C. J. Fourie, G. Faini, C. Ulysse, and P. Febvre, "YBCO Josephson junctions and striplines for RSFQ circuits made by ion irradiation," *IEEE Transactions on Applied Superconductivity*, vol. 23, no. 2, Apr. 2013, Art. no. 1101205.
- [21] S. Anders, M. Blamire, F.-I. Buchholz, D.-G. Cr  t  , R. Cristiano, P. Febvre, L. Fritzsche, A. Herr, E. Il'ichev, J. Kohlmann, J. Kunert, H.-G. Meyer, J. Niemeyer, T. Ortlepp, H. Rogalla, T. Schurig, M. Siegel, R. Stolz, E. Tarte, H. ter Brake, H. Toepfer, J.-C. Villegier, A. Zagorskin, and A. Zorin, "European roadmap on superconductive electronics – status and perspectives," *Physica C: Superconductivity*, vol. 470, no. 23, pp. 2079–2126, 2010.
- [22] D. Yohannes, S. Sarwana, S. Tolpygo, A. Sahu, Y. Polyakov, and V. Semenov, "Characterization of HYPRES' 4.5 kA/cm² & 8 kA/cm² Nb/AlOx/Nb fabrication processes," *IEEE Transactions on Applied Superconductivity*, vol. 15, no. 2, pp. 90–93, Jun. 2005.
- [23] D. T. Yohannes, R. T. Hunt, J. A. Vivalda, D. Amparo, A. Cohen, I. V. Vernik, and A. F. Kirichenko, "Planarized, extendible, multilayer fabrication process for superconducting electronics," *IEEE Transactions on Applied Superconductivity*, vol. 25, no. 3, Jun. 2015, Art. no. 1100405.
- [24] S. Nagasawa, K. Hinode, T. Satoh, M. Hidaka, H. Akaike, A. Fujimaki, N. Yoshikawa, K. Takagi, and N. Takagi, "Nb 9-layer fabrication process for superconducting large-scale SFQ circuits and its process evaluation," *IEICE Transactions on Electronics*, vol. E97-C, no. 3, pp. 132–140, Mar. 2014.

- [25] M. Hidaka, S. Nagasawa, T. Satoh, K. Hinode, and Y. Kitagawa, “Current status and future prospect of the Nb-based fabrication process for single flux quantum circuits,” *Superconductor Science and Technology*, vol. 19, no. 3, S138–S142, Feb. 2006.
- [26] S. Nagasawa, T. Satoh, K. Hinode, Y. Kitagawa, M. Hidaka, H. Akaike, A. Fujimaki, K. Takagi, N. Takagi, and N. Yoshikawa, “New Nb multi-layer fabrication process for large-scale SFQ circuits,” *Physica C: Superconductivity and its Applications*, vol. 469, pp. 1578–1584, 2009.
- [27] A. Fujimaki, M. Tanaka, R. Kasagi, K. Takagi, M. Okada, Y. Hayakawa, K. Takata, H. Akaike, N. Yoshikawa, S. Nagasawa, K. Takagi, and N. Takagi, “Large-scale integrated circuit design based on a Nb nine-layer structure for reconfigurable data-path processors,” *IEICE Transactions on Electronics*, vol. E97-C, no. 3, pp. 157–164, Mar. 2014.
- [28] S. Nagasawa, K. Hinode, M. Sugita, T. Satoh, H. Akaike, Y. Kitagawa, and M. Hidaka, “Planarized multi-layer fabrication technology for LTS large-scale SFQ circuits,” *Superconductor Science and Technology*, vol. 16, no. 12, pp. 1483–1486, Nov. 2003.
- [29] T. Satoh, K. Hinode, S. Nagasawa, Y. Kitagawa, M. Hidaka, N. Yoshikawa, H. Akaike, A. Fujimaki, K. Takagi, and N. Takagi, “Planarization process for fabricating multi-layer Nb integrated circuits incorporating top active layer,” *IEEE Transactions on Applied Superconductivity*, vol. 19, no. 3, pp. 167–170, 2009.
- [30] M. A. Manheimer, “Cryogenic computing complexity program: Phase 1 introduction,” *IEEE Transactions on Applied Superconductivity*, vol. 25, no. 3, Jun. 2015, Art. no. 1301704.
- [31] S. K. Tolpygo, V. Bolkhovsky, R. Rastogi, S. Zarr, A. L. Day, E. Golden, T. J. Weir, A. Wynn, and L. M. Johnson, “Advanced fabrication processes for superconductor electronics: Current status and new developments,” *IEEE Transactions on Applied Superconductivity*, vol. 29, no. 5, 2019, Art. no. 1102513.
- [32] S. K. Tolpygo, V. Bolkhovsky, T. J. Weir, C. J. Galbraith, L. M. Johnson, M. A. Gouker, and V. K. Semenov, “Inductance of circuit structures for MIT LL superconductor electronics fabrication process with 8 niobium layers,” *IEEE Transactions on Applied Superconductivity*, vol. 25, no. 3, 2015, Art. no. 1100905.
- [33] W. Anacker, “Josephson computer technology: An IBM research project,” *IBM Journal of Research and Development*, vol. 24, no. 2, pp. 107–112, 1980.
- [34] M. Gurvitch, M. Washington, H. Huggins, and J. Rowell, “Preparation and properties of Nb Josephson junctions with thin Al layers,” *IEEE Transactions on Magnetics*, vol. 19, no. 3, pp. 791–794, 1983.
- [35] M. Gurvitch, M. A. Washington, and H. A. Huggins, “High quality refractory josephson tunnel junctions utilizing thin aluminum layers,” *Applied Physics Letters*, vol. 42, no. 5, pp. 472–474, 1983.
- [36] Y. Tarutani, M. Hirano, and U. Kawabe, “Niobium-based integrated circuit technologies,” *Proceedings of the IEEE*, vol. 77, no. 8, pp. 1164–1176, 1989.
- [37] S. Kosaka, A. Shoji, M. Aoyagi, F. Shinoki, S. Tahara, H. Ohigashi, H. Nakagawa, S. Takada, and H. Hayakawa, “An integration of all refractory josephson logic LSI circuit,” *IEEE Transactions on Magnetics*, vol. 21, no. 2, pp. 102–109, 1985.

- [38] H. Nakagawa, I. Kurosawa, S. Takada, and H. Hayakawa, “Josephson 4-bit digital counter circuit made by Nb/Al-oxide/Nb junctions,” *IEEE Transactions on Magnetics*, vol. 23, no. 2, pp. 739–742, 1987.
- [39] E. Fang, D. Hebert, and T. Van Duzer, “A multi-gigahertz, Josephson flash A/D converter with a pipelined encoder using large-dynamic-range current-latch comparators,” *IEEE Transactions on Magnetics*, vol. 27, no. 2, pp. 2891–2894, 1991.
- [40] H. Luong, D. Hebert, and T. Van Duzer, “Fully parallel superconducting analog-to-digital converter,” *IEEE Transactions on Applied Superconductivity*, vol. 3, no. 1, pp. 2633–2636, 1993.
- [41] W. Perold, M. Jeffery, Z. Wang, and T. Van Duzer, “Complementary output switching logic—a new superconducting voltage-state logic family,” *IEEE Transactions on Applied Superconductivity*, vol. 6, no. 3, pp. 125–131, 1996.
- [42] M. Jeffery, W. Perold, and T. van Duzer, “Experimental demonstration of complementary output switching logic approaching 10 Gb/s clock frequencies,” *IEEE Transactions on Applied Superconductivity*, vol. 7, no. 2, pp. 2665–2668, 1997.
- [43] K. Likharev, O. Mukhanov, and V. Semenov, “Resistive single flux quantum logic for the Josephson-junction digital technology,” in *Proceedings of the Third International Conference on Superconducting Quantum Devices, Berlin (West), June 25-28, 1985*, H.-D. Hahlbohm and H. Lübbig, Eds. Berlin, Boston: De Gruyter, 1985, pp. 1103–1108.
- [44] O. Mukhanov, V. Semenov, and K. Likharev, “Ultimate performance of the RSFQ logic circuits,” *IEEE Transactions on Magnetics*, vol. 23, no. 2, pp. 759–762, Mar. 1987.
- [45] D. E. Kirichenko, S. Sarwana, and A. F. Kirichenko, “Zero static power dissipation biasing of RSFQ circuits,” *IEEE Transactions on Applied Superconductivity*, vol. 21, no. 3, pp. 776–779, Jun. 2011.
- [46] N. Takeuchi, D. Ozawa, Y. Yamanashi, and N. Yoshikawa, “An adiabatic quantum flux parametron as an ultra-low-power logic device,” *Superconductor Science and Technology*, vol. 26, no. 3, Jan. 2013, Art. no. 035010.
- [47] 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.
- [48] O. Mukhanov, S. Polonsky, and V. Semenov, “New elements of the RSFQ logic family,” *IEEE Transactions on Magnetics*, vol. 27, no. 2, pp. 2435–2438, 1991.
- [49] B. Dimov, “General restrictions and their possible solutions for the development of Ultra High-Speed integrated RSFQ digital circuits,” Ph.D. dissertation, Technische Universität Ilmenau, 2005.
- [50] L. Schindler, “The development and characterisation of a parameterised rsfq cell library for layout synthesis,” Ph.D. dissertation, Stellenbosch University, 2021.
- [51] L. Schindler and C. J. Fourie, “Application of phase-based circuit theory to RSFQ logic design,” *IEEE Transactions on Applied Superconductivity*, vol. 32, no. 3, Apr. 2022, Art. no. 1300512.

- [52] B. Dimov, M. Khabipov, D. Balashov, C. Brandt, F.-I. Buchholz, J. Niemeyer, and F. Uhlmann, "Tuning of the RSFQ gate speed by different Stewart-McCumber parameters of the Josephson junctions," *IEEE Transactions on Applied Superconductivity*, vol. 15, no. 2, pp. 284–287, 2005.
- [53] S. Polonsky, V. Semenov, P. Bunyk, A. Kirichenko, A. Kidiyarov-Shevchenko, O. Mukhanov, P. Shevchenko, D. Schneider, D. Zinoviev, and K. Likharev, "New RSFQ circuits," *IEEE Transactions on Applied Superconductivity*, vol. 3, no. 1, pp. 2566–2577, 1993.
- [54] C. J. Fourie, "Extraction of dc-biased sfq circuit verilog models," *IEEE Transactions on Applied Superconductivity*, vol. 28, no. 6, Sep. 2018, Art. no. 1300811.
- [55] O. A. Mukhanov, "Energy-efficient single flux quantum technology," *IEEE Transactions on Applied Superconductivity*, vol. 21, no. 3, pp. 760–769, Jun. 2011.
- [56] C. J. Fourie, "A 10 GHz oversampling delta modulating analogue-to-digital converter implemented with hybrid superconducting digital logic," Masters thesis, Stellenbosch University, 2001.
- [57] R. S. Bakolo and C. J. Fourie, "Development of a RSFQ cell library for the University of Stellenbosch," in *IEEE Africon '11*, 2011, pp. 1–5.
- [58] R. S. Bakolo, "Design and implementation of a RSFQ superconductive digital electronics cell library," Masters thesis, Stellenbosch University, 2011.
- [59] R. S. Bakolo and C. J. Fourie, "New implementation of RSFQ superconductive digital gates," *SAIEE Africa Research Journal*, vol. 104, no. 3, pp. 90–96, 2013.
- [60] C. J. Fourie and W. J. Perold, "An RSFQ DC-resettable latch for building memory and reprogrammable circuits," *IEEE Transactions on Applied Superconductivity*, vol. 15, no. 2, pp. 348–351, 2005.
- [61] C. J. Fourie, O. Wetzstein, J. Kunert, and H.-G. Meyer, "SFQ circuits with ground plane hole-assisted inductive coupling designed with InductEx," *IEEE Transactions on Applied Superconductivity*, vol. 23, no. 3, Jun. 2013, Art. no. 1300705.
- [62] C. Fourie and W. Perold, "A single-clock asynchronous input COSL set-reset flip-flop and SFQ to voltage state interface," *IEEE Transactions on Applied Superconductivity*, vol. 15, no. 2, pp. 263–266, 2005.
- [63] S. D. Brown, R. J. Francis, J. Rose, and Z. G. Vranesic, *Field-programmable gate arrays*. Kluwer Academic Publishers, 1992.
- [64] S. Brown, J. Rose, and Z. Vranesic, "A detailed router for field-programmable gate arrays," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 11, no. 5, pp. 620–628, 1992.
- [65] C. J. Fourie, "A tool kit for the design of superconducting programmable gate arrays," Ph.D. dissertation, Stellenbosch University, 2003.
- [66] A. Hebard, S. Pei, L. Dunkleberger, and T. Fulton, "A DC-powered Josephson flip-flop," *IEEE Transactions on Magnetics*, vol. 15, no. 1, pp. 408–411, 1979.
- [67] Y. Hatano, H. Nagaishi, S. Yano, K. Nakahara, H. Yamada, S. Kominami, and M. Hirano, "An all DC-powered Josephson logic circuit," *IEEE Journal of Solid-State Circuits*, vol. 26, no. 8, pp. 1123–1132, 1991.

- [68] D. Schneider, J. Lin, S. Polonsky, V. Semenov, and C. Hamilton, "Broadband interfacing of superconducting digital systems to room temperature electronics," *IEEE Transactions on Applied Superconductivity*, vol. 5, no. 2, pp. 3152–3155, 1995.
- [69] H. Van Heerden, "The design and testing of a superconducting programmable gate array," Masters thesis, Stellenbosch University, 2005.
- [70] C. J. Fourie and H. van Heerden, "An RSFQ superconductive programmable gate array," *IEEE Transactions on Applied Superconductivity*, vol. 17, no. 2, pp. 538–541, 2007.
- [71] N. K. Katam, O. A. Mukhanov, and M. Pedram, "Superconducting magnetic field programmable gate array," *IEEE Transactions on Applied Superconductivity*, vol. 28, no. 2, 2018, Art. no. 1300212.
- [72] T. Hosoya, Y. Yamanashi, and N. Yoshikawa, "Compact superconducting lookup table composed of two-dimensional memory cell array reconfigured by external dc control currents," *IEEE Transactions on Applied Superconductivity*, vol. 31, no. 3, 2021, Art. no. 1300406.
- [73] D. Takahashi, N. Takeuchi, Y. Yamanashi, and N. Yoshikawa, "Design and demonstration of a superconducting field-programmable gate array using adiabatic quantum-flux-parametron logic and memory," *IEEE Transactions on Applied Superconductivity*, vol. 32, no. 7, 2022, Art. no. 1301207.
- [74] C. C. Maree and C. J. Fourie, "Development of an all-SFQ superconducting field-programmable gate array," *IEEE Transactions on Applied Superconductivity*, vol. 29, no. 4, Dec. 2019, Art. no. 1300212.
- [75] V. Betz and J. Rose, "VPR: A new packing, placement and routing tool for FPGA research," in *Proceedings of the 7th International Workshop on Field-Programmable Logic and Applications*, ser. FPL '97, Berlin, Heidelberg: Springer-Verlag, 1997, pp. 213–222.
- [76] J. Luu, J. Goeders, M. Wainberg, A. Somerville, T. Yu, K. Nasartschuk, M. Nasr, S. Wang, T. Liu, N. Ahmed, K. B. Kent, J. Anderson, J. Rose, and V. Betz, "VTR 7.0: Next generation architecture and CAD system for FPGAs," vol. 7, no. 2, Jul. 2014, Art. no. 6.
- [77] Z. Deng, N. Yoshikawa, S. Whiteley, and T. Van Duzer, "Data-driven self-timed RSFQ digital integrated circuit and system," *IEEE Transactions on Applied Superconductivity*, vol. 7, no. 2, pp. 3634–3637, 1997.
- [78] S. J. E. Wilton, "Architectures and algorithms for field-programmable gate arrays with embedded memory," Ph.D. dissertation, University of Toronto, 1997.
- [79] P. Jamieson, K. B. Kent, F. Gharibian, and L. Shannon, "Odin II - An open-source Verilog HDL synthesis tool for CAD research," in *2010 18th IEEE Annual International Symposium on Field-Programmable Custom Computing Machines*, 2010, pp. 149–156.
- [80] A. Mishchenko. (2012). "ABC: A system for sequential synthesis and verification," [Online]. Available: <https://people.eecs.berkeley.edu/~alanmi/abc/>. (accessed: 25.10.2022).

- [81] J. Luu, I. Kuon, P. Jamieson, T. Campbell, A. Ye, W. M. Fang, and J. Rose, "VPR 5.0: FPGA cad and architecture exploration tools with single-driver routing, heterogeneity and process scaling," in *Proceedings of the ACM/SIGDA International Symposium on Field Programmable Gate Arrays*, ser. FPGA '09, Monterey, California, USA: Association for Computing Machinery, 2009, pp. 133–142.
- [82] S. Yang, *Logic Synthesis and Optimization Benchmarks User Guide: Version 3.0*. Microelectronics Center of North Carolina (MCNC), 1991.
- [83] M. Dorojevets, C. L. Ayala, and A. K. Kasperek, "Data-flow microarchitecture for wide datapath RSFQ processors: Design study," *IEEE Transactions on Applied Superconductivity*, vol. 21, no. 3, pp. 787–791, Jun. 2011.
- [84] T. Filippov, M. Dorojevets, A. Sahu, A. Kirichenko, C. Ayala, and O. Mukhanov, "8-bit asynchronous wave-pipelined RSFQ Arithmetic-Logic Unit," *IEEE Transactions on Applied Superconductivity*, vol. 21, no. 3, pp. 847–851, Jun. 2011.
- [85] M. H. Volkmann, I. V. Vernik, and O. A. Mukhanov, "Wave-pipelined eSFQ circuits," *IEEE Transactions on Applied Superconductivity*, vol. 25, no. 3, Jun. 2015, Art. no. 1301005.
- [86] H. Gerber, C. Fourie, and W. Perold, "RSFQ-asynchronous timing (RSFQ-AT): A new design methodology for implementation in CAD automation," *IEEE Transactions on Applied Superconductivity*, vol. 15, no. 2, pp. 272–275, 2005.
- [87] M. Maezawa, I. Kurosawa, M. Aoyagi, H. Nakagawa, Y. Kameda, and T. Nanya, "Rapid single-flux-quantum dual-rail logic for asynchronous circuits," *IEEE Transactions on Applied Superconductivity*, vol. 7, no. 2, pp. 2705–2708, 1997.
- [88] Z. Deng, N. Yoshikawa, J. Tierno, S. Whiteley, and T. van Duzer, "Asynchronous circuits and systems in superconducting RSFQ digital technology," in *Proceedings Fourth International Symposium on Advanced Research in Asynchronous Circuits and Systems*, 1998, pp. 274–285.
- [89] L. C. Müller, H. R. Gerber, and C. J. Fourie, "Review and comparison of RSFQ asynchronous methodologies," *Journal of Physics: Conference Series*, vol. 97, no. 1, Feb. 2008, Art. no. 012109.
- [90] P. Patra, S. Polonsky, and D. Fussell, "Delay insensitive logic for RSFQ superconductor technology," in *Proceedings Third International Symposium on Advanced Research in Asynchronous Circuits and Systems*, 1997, pp. 42–53.
- [91] H. Gerber, C. Fourie, and W. Perold, "Specification of a technology portable logic cell library for RSFQ: An automated approach," *IEEE Transactions on Applied Superconductivity*, vol. 15, no. 2, pp. 368–371, 2005.
- [92] C. J. Fourie, K. Jackman, M. M. Botha, S. Razmkhah, P. Febvre, C. L. Ayala, Q. Xu, N. Yoshikawa, E. Patrick, M. Law, Y. Wang, M. Annavaram, P. Beerel, S. Gupta, S. Nazarian, and M. Pedram, "Coldflux superconducting EDA and TCAD tools project: Overview and progress," *IEEE Transactions on Applied Superconductivity*, vol. 29, no. 5, Aug. 2019, Art. no. 1300407.
- [93] S. N. Shahsavani, T.-R. Lin, A. Shafaei, C. J. Fourie, and M. Pedram, "An integrated row-based cell placement and interconnect synthesis tool for large SFQ logic circuits," *IEEE Transactions on Applied Superconductivity*, vol. 27, no. 4, Jun. 2017, Art. no. 1302008.

- [94] S. Polonsky, V. Semenov, and D. Schneider, “Transmission of single-flux-quantum pulses along superconducting microstrip lines,” *IEEE Transactions on Applied Superconductivity*, vol. 3, no. 1, pp. 2598–2600, 1993.
- [95] M. Tanaka, T. Kondo, N. Nakajima, T. Kawamoto, Y. Yamanashi, Y. Kamiya, A. Akimoto, A. Fujimaki, H. Hayakawa, N. Yoshikawa, H. Terai, Y. Hashimoto, and S. Yorozu, “Demonstration of a single-flux-quantum microprocessor using passive transmission lines,” *IEEE Transactions on Applied Superconductivity*, vol. 15, no. 2, pp. 400–404, 2005.
- [96] H. Akaike, M. Tanaka, K. Takagi, I. Kataeva, R. Kasagi, A. Fujimaki, K. Takagi, M. Igarashi, H. Park, Y. Yamanashi, N. Yoshikawa, K. Fujiwara, S. Nagasawa, M. Hidaka, and N. Takagi, “Design of single flux quantum cells for a 10-Nb-layer process,” *Physica C: Superconductivity*, vol. 469, no. 15, pp. 1670–1673, 2009.
- [97] C. J. Fourie, S. Miyanishi, and N. Yoshikawa, “Grounding methods to reduce stray coupling in multi-layer layouts,” in *2015 15th International Superconductive Electronics Conference (ISEC)*, Jul. 2015, pp. 1–3.
- [98] C. J. Fourie, C. L. Ayala, L. Schindler, T. Tanaka, and N. Yoshikawa, “Design and characterization of track routing architecture for RSFQ and AQFP circuits in a multilayer process,” *IEEE Transactions on Applied Superconductivity*, vol. 30, no. 6, Sep. 2020, Art. no. 1301109.
- [99] L. Schindler, J. A. Delport, and C. J. Fourie, “The ColdFlux RSFQ cell library for MIT-LL SFQ5ee fabrication process,” *IEEE Transactions on Applied Superconductivity*, vol. 32, no. 2, Mar. 2022, Art. no. 1300207.
- [100] H. R. Gerber, C. J. Fourie, W. J. Perold, and L. C. Muller, “Design of an asynchronous microprocessor using RSFQ-AT,” *IEEE Transactions on Applied Superconductivity*, vol. 17, no. 2, pp. 490–493, 2007.
- [101] N. Yoshikawa and Y. Kato, “Reduction of power consumption of rsfq circuits by inductance-load biasing,” *Superconductor Science and Technology*, vol. 12, no. 11, pp. 918–920, Nov. 1999.
- [102] 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, no. 2, pp. 150–153, 2007.
- [103] T. Ortlepp, O. Wetzstein, S. Engert, J. Kunert, and H. Toepfer, “Reduced power consumption in superconducting electronics,” *IEEE Transactions on Applied Superconductivity*, vol. 21, no. 3, pp. 770–775, 2011.
- [104] 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, May 2012, Art. no. 053102.
- [105] A. V. Ustinov and V. K. Kaplunenko, “Rapid single-flux quantum logic using π -shifters,” *Journal of Applied Physics*, vol. 94, no. 8, pp. 5405–5407, 2003.
- [106] Y. Yamanashi, S. Nakaishi, A. Sugiyama, N. Takeuchi, and N. Yoshikawa, “Design methodology of single-flux-quantum flip-flops composed of both 0- and π -shifted Josephson junctions,” *Superconductor Science and Technology*, vol. 31, no. 10, Aug. 2018, Art. no. 105003.

- [107] F. Li, Y. Takeshita, D. Hasegawa, M. Tanaka, T. Yamashita, and A. Fujimaki, “Low-power high-speed half-flux-quantum circuits driven by low bias voltages,” *Superconductor Science and Technology*, vol. 34, no. 2, p. 025 013, Jan. 2021, Art. no. 025013.
- [108] C. J. Fourie, O. Wetzstein, J. Kunert, H. Toepfer, and H.-G. Meyer, “Experimentally verified inductance extraction and parameter study for superconductive integrated circuit wires crossing ground plane holes,” *Superconductor Science and Technology*, vol. 26, no. 1, 2013, Art. no. 015016.
- [109] C. Shawawreh, D. Amparo, J. Ren, M. Miller, M. Y. Kamkar, A. Sahu, A. Inamdar, A. F. Kirichenko, O. A. Mukhanov, and I. V. Vernik, “Effects of adaptive DC biasing on operational margins in ERSFQ circuits,” *IEEE Transactions on Applied Superconductivity*, vol. 27, no. 4, pp. 1–6, 2017, Art. no. 1301606.
- [110] A. F. Kirichenko, M. Y. Kamkar, J. Walter, and I. V. Vernik, “ERSFQ 8-bit Parallel Binary Shifter for Energy-Efficient Superconducting CPU,” *IEEE Transactions on Applied Superconductivity*, vol. 29, no. 5, pp. 1–4, 2019, Art. no. 1302704.
- [111] M. H. Volkmann, A. Sahu, C. J. Fourie, and O. A. Mukhanov, “Implementation of energy efficient single flux quantum digital circuits with sub-aJ operation,” *Superconductor Science and Technology*, vol. 26, no. 1, 2013, Art. no. 015002.
- [112] M. H. Volkmann, “A superconducting software-defined radio frontend with application to the Square-Kilometre Array,” Ph.D. dissertation, Stellenbosch University, 2013.
- [113] S. Kaplan and O. Mukhanov, “Operation of a superconductive demultiplexer using rapid single flux quantum (RSFQ) technology,” *IEEE Transactions on Applied Superconductivity*, vol. 5, no. 2, pp. 2853–2856, 1995.
- [114] M. H. Volkmann, A. Sahu, C. J. Fourie, and O. A. Mukhanov, “Experimental investigation of energy-efficient digital circuits based on eSFQ logic,” *IEEE Transactions on Applied Superconductivity*, vol. 23, no. 3, Jun. 2013, Art. no. 1301505.
- [115] I. V. Vernik, S. B. Kaplan, M. H. Volkmann, A. V. Dotsenko, C. J. Fourie, and O. A. Mukhanov, “Design and test of asynchronous eSFQ circuits,” *Superconductor Science and Technology*, vol. 27, no. 4, 2014, Art. no. 044030.
- [116] R. Landauer, “Irreversibility and heat generation in the computing process,” *IBM Journal of Research and Development*, vol. 5, no. 3, pp. 183–191, 1961.
- [117] R. W. Keyes and R. Landauer, “Minimal energy dissipation in logic,” *IBM Journal of Research and Development*, vol. 14, no. 2, pp. 152–157, 1970.
- [118] K. Likharev, “Dynamics of some single flux quantum devices: I. Parametric quantron,” *IEEE Transactions on Magnetics*, vol. 13, no. 1, pp. 242–244, 1977.
- [119] N. Shimizu, Y. Harada, N. Miyamoto, and E. Goto, “A new A/D converter with quantum flux parametron,” *IEEE Transactions on Magnetics*, vol. 25, no. 2, pp. 865–868, 1989.
- [120] M. Hosoya, W. Hioe, J. Casas, R. Kamikawai, Y. Harada, Y. Wada, H. Nakane, R. Suda, and E. Goto, “Quantum flux parametron: A single quantum flux device for Josephson supercomputer,” *IEEE Transactions on Applied Superconductivity*, vol. 1, no. 2, pp. 77–89, 1991.

- [121] N. Takeuchi, Y. Yamanashi, and N. Yoshikawa, “Energy efficiency of adiabatic superconductor logic,” *Superconductor Science and Technology*, vol. 28, no. 1, Nov. 2014, Art. no. 015003.
- [122] N. Takeuchi, S. Nagasawa, F. China, T. Ando, M. Hidaka, Y. Yamanashi, and N. Yoshikawa, “Adiabatic quantum-flux-parametron cell library designed using a 10 ka cm⁻² niobium fabrication process,” *Superconductor Science and Technology*, vol. 30, no. 3, Jan. 2017, Art. no. 035002.
- [123] N. Takeuchi, H. Suzuki, C. J. Fourie, and N. Yoshikawa, “Impedance design of excitation lines in Adiabatic Quantum-Flux-Parametron logic using InductEx,” *IEEE Transactions on Applied Superconductivity*, vol. 31, no. 5, Aug. 2021, Art. no. 1300605.
- [124] C. J. Fourie, N. Takeuchi, K. Jackman, and N. Yoshikawa, “Evaluation of flux trapping moat position on AQFP cell performance,” *Journal of Physics: Conference Series*, vol. 1975, no. 1, Jul. 2021, Art. no. 012027.
- [125] S. K. Tolpygo, E. B. Golden, T. J. Weir, and V. Bolkhovsky, “Mutual and self-inductance in planarized multilayered superconductor integrated circuits: Microstrips, striplines, bends, meanders, ground plane perforations,” *IEEE Transactions on Applied Superconductivity*, vol. 32, no. 5, 2022, Art. no. 1400331.
- [126] F. London and H. London, “The electromagnetic equations of the supraconductor,” *Proceedings of the Royal Society A*, vol. 149, pp. 71–88, 1935.
- [127] Y. I. Ismail, E. G. Friedman, and J. L. Neves, “Figures of merit to characterize the importance of on-chip inductance,” *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 7, no. 4, pp. 442–449, 1999.
- [128] G. Serval, L. Kenmei, F. Huret, E. Paleczny, P. Kennis, and D. Deschacht, “Inductance effect for interconnection timing analysis in submicronic circuits,” in *ICECS’99. Proceedings of ICECS ’99. 6th IEEE International Conference on Electronics, Circuits and Systems (Cat. No.99EX357)*, vol. 3, 1999, 1407–1410 vol.3.
- [129] L.-F. Chang, K.-J. Chang, and R. Mathews, “When should on-chip inductance modeling become necessary for VLSI timing analysis?” In *Proceedings of the IEEE 2000 International Interconnect Technology Conference (Cat. No.00EX407)*, 2000, pp. 170–172.
- [130] S. Roy and A. Dounavis, “Closed-form delay and crosstalk models for RLC on-chip interconnects using a matrix rational approximation,” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 28, no. 10, pp. 1481–1492, 2009.
- [131] T. Chen, “On the impact of on-chip inductance on signal nets under the influence of power grid noise,” *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 13, no. 3, pp. 339–348, 2005.
- [132] Y. Ismail, “On-chip inductance cons and pros,” *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 10, no. 6, pp. 685–694, 2002.
- [133] J. L. López, P. Zumel, S. O’Driscoll, Z. Pavlovic, R. Murphy, C. O’Mathuna, and C. Fernandez, “Comprehensive design procedure for racetrack microinductors,” *IEEE Journal of Emerging and Selected Topics in Power Electronics*, vol. 9, no. 6, pp. 6912–6923, 2021.

- [134] D. V. Harburg, A. J. Hanson, J. Qiu, B. A. Reese, J. D. Ranson, D. M. Otten, C. G. Levey, and C. R. Sullivan, "Microfabricated racetrack inductors with thin-film magnetic cores for on-chip power conversion," *IEEE Journal of Emerging and Selected Topics in Power Electronics*, vol. 6, no. 3, pp. 1280–1294, 2018.
- [135] F. Tounsi, D. Flandre, L. Rufer, and L. A. Francis, "Performances evaluation of on-chip large-size-tapped transformer for MEMS applications," *IEEE Transactions on Instrumentation and Measurement*, vol. 69, no. 9, pp. 7051–7060, 2020.
- [136] A. Mezhiba and E. Friedman, "Inductive properties of high-performance power distribution grids," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 10, no. 6, pp. 762–776, 2002.
- [137] N. Srivastava, X. Qi, and K. Banerjee, "Impact of on-chip inductance on power distribution network design for nanometer scale integrated circuits," in *Sixth international symposium on quality electronic design (isqed'05)*, 2005, pp. 346–351.
- [138] W.-Y. Ding, X.-C. Wei, Y.-F. Shu, D. Yi, and T.-M. Xiang, "Inductance extraction of grid power distribution network," *IEEE Transactions on Components, Packaging and Manufacturing Technology*, vol. 8, no. 6, pp. 1066–1072, 2018.
- [139] C. Qu, Z. Zhu, Y. En, L. Wang, and X. Liu, "Area-efficient extended 3-D inductor based on TSV technology for RF applications," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 29, no. 2, pp. 287–296, 2021.
- [140] M. W. Johnson, Q. P. Herr, D. J. Durand, and L. A. Abelson, "Differential SFQ transmission using either inductive or capacitive coupling," *IEEE Transactions on Applied Superconductivity*, vol. 13, no. 2, pp. 507–510, Jun. 2003.
- [141] J. H. Kang and S. B. Kaplan, "Current recycling and SFQ signal transfer in large scale RSFQ circuits," *IEEE Transactions on Applied Superconductivity*, vol. 13, no. 2, pp. 547–550, Jun. 2003.
- [142] M. Igarashi, Y. Yamanashi, N. Yoshikawa, K. Fujiwara, and Y. Hashimoto, "SFQ pulse transfer circuits using inductive coupling for current recycling," *IEEE Transactions on Applied Superconductivity*, vol. 19, no. 3, pp. 649–652, Jun. 2009.
- [143] S. B. Kaplan, "Serial biasing of 16 modular circuits at 50 Gb/s," *IEEE Transactions on Applied Superconductivity*, vol. 22, no. 4, Aug. 2012, Art. no. 1300103.
- [144] K. Ehara, A. Takahashi, Y. Yamanashi, and N. Yoshikawa, "Development of pulse transfer circuits for serially biased SFQ circuits using the Nb 9-layer 1- μ m process," *IEEE Transactions on Applied Superconductivity*, vol. 23, no. 3, 2013, Art. no. 1300504.
- [145] S. S. Wong, P. Yue, R. Chang, S.-Y. Kim, B. Kleveland, and F. O'Mahony, "On-chip interconnect inductance - friend or foe," in *Fourth International Symposium on Quality Electronic Design, 2003. Proceedings.*, Mar. 2003, pp. 389–394.
- [146] M. Horowitz and R. W. Dutton, "Resistance extraction from mask layout data," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 2, no. 3, pp. 145–150, Jul. 1983.
- [147] M. Harbour and J. Drake, "Calculation of multiterminal resistances in integrated circuits," *IEEE Transactions on Circuits and Systems*, vol. 33, no. 4, pp. 462–465, Apr. 1986.

- [148] T. Mitsuhashi and K. Yoshida, "A resistance calculation algorithm and its application to circuit extraction," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 6, no. 3, pp. 337–345, May 1987.
- [149] Z. Wang and Q. Wu, "A two-dimensional resistance simulator using the boundary element method," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 11, no. 4, pp. 497–504, Apr. 1992.
- [150] D. Sitaram, Yu Zheng, and K. L. Shepard, "Full-chip, three-dimensional shapes-based RLC extraction," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 23, no. 5, pp. 711–727, May 2004.
- [151] K. Nabors and J. White, "Fastcap: A multipole accelerated 3-d capacitance extraction program," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 10, no. 11, pp. 1447–1459, Nov. 1991.
- [152] S. Yan, V. Sarin, and W. Shi, "Fast 3-d capacitance extraction by inexact factorization and reduction," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 25, no. 10, pp. 2282–2286, Oct. 2006.
- [153] F. Grover, *Inductance Calculations*, ser. Dover Books on Electrical Engineering. Dover Publications, 2013.
- [154] N. H. Meyers, "Inductance in thin-film superconducting structures," *Proceedings of the IRE*, vol. 49, no. 11, pp. 1640–1649, Nov. 1961.
- [155] Y. Massoud and Y. Ismail, "Grasping the impact of on-chip inductance," *IEEE Circuits and Devices Magazine*, vol. 17, no. 4, pp. 14–21, Jul. 2001.
- [156] M. W. Beattie and L. T. Pileggi, "IC analyses including extracted inductance models," in *Proceedings 1999 Design Automation Conference (Cat. No. 99CH36361)*, Jun. 1999, pp. 915–920.
- [157] G. V. Kopsay, B. Krauter, D. Widiger, A. Deutsch, B. J. Rubin, and H. H. Smith, "A comprehensive 2-D inductance modeling approach for VLSI interconnects: Frequency-dependent extraction and compact circuit model synthesis," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 10, no. 6, pp. 695–711, Dec. 2002.
- [158] S. Yu, D. M. Petranovic, S. Krishnan, Kwyro Lee, and C. Y. Yang, "Loop-based inductance extraction and modeling for multiconductor on-chip interconnects," *IEEE Transactions on Electron Devices*, vol. 53, no. 1, pp. 135–145, Jan. 2006.
- [159] M. W. Beattie and L. T. Pileggi, "On-chip induction modeling: Basics and advanced methods," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 10, no. 6, pp. 712–729, Dec. 2002.
- [160] K. L. Shepard and Zhong Tian, "Return-limited inductances: A practical approach to on-chip inductance extraction," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 19, no. 4, pp. 425–436, Apr. 2000.
- [161] So Young Kim, Y. Massoud, and S. S. Wong, "On the accuracy of return path assumption for loop inductance extraction for 0.1 μm technology and beyond," in *Fourth International Symposium on Quality Electronic Design, 2003. Proceedings.*, Mar. 2003, pp. 401–404.

- [162] A. E. Ruehli, "Inductance calculations in a complex integrated circuit environment," *IBM Journal of Research and Development*, vol. 16, no. 5, pp. 470–481, Sep. 1972.
- [163] M. Mondal and Y. Massoud, "Accurate loop self inductance bound for efficient inductance screening," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 14, no. 12, pp. 1393–1397, Dec. 2006.
- [164] 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, no. 9, pp. 1750–1758, Sep. 1994.
- [165] M. Beattie, B. Krauter, L. Alatan, and L. Pileggi, "Equipotential shells for efficient inductance extraction," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 20, no. 1, pp. 70–79, Jan. 2001.
- [166] A. Devgan, Hao Ji, and W. Dai, "How to efficiently capture on-chip inductance effects: Introducing a new circuit element k," in *IEEE/ACM International Conference on Computer Aided Design. ICCAD - 2000. IEEE/ACM Digest of Technical Papers (Cat. No.00CH37140)*, Nov. 2000, pp. 150–155.
- [167] T.-H. Chen, C. Luk, and C. C. .-. Chen, "INDUCTWISE: Inductance-wise interconnect simulator and extractor," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 22, no. 7, pp. 884–894, Jul. 2003.
- [168] H. Hu and S. Sapatnekar, "Efficient inductance extraction using circuit-aware techniques," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 10, no. 6, pp. 746–761, 2002.
- [169] Xiaoning Qi, Gaofeng Wang, Zhiping Yu, R. W. Dutton, Tak Young, and N. Chang, "On-chip inductance modeling and RLC extraction of VLSI interconnects for circuit simulation," in *Proceedings of the IEEE 2000 Custom Integrated Circuits Conference (Cat. No.00CH37044)*, May 2000, pp. 487–490.
- [170] S.-W. Tu, W.-Z. Shen, Y.-W. Chang, and T.-C. Chen, "Inductance modeling for on-chip interconnects," in *2002 IEEE International Symposium on Circuits and Systems (ISCAS)*, vol. 3, 2002, p. III.
- [171] Q. Herr, M. Wire, and A. Smith, "Ballistic SFQ signal propagation on-chip and chip-to-chip," *IEEE Transactions on Applied Superconductivity*, vol. 13, no. 2, pp. 463–466, 2003.
- [172] Y. Kameda, S. Yorozu, and Y. Hashimoto, "A new design methodology for single-flux-quantum (SFQ) logic circuits using passive-transmission-line (PTL) wiring," *IEEE Transactions on Applied Superconductivity*, vol. 17, no. 2, pp. 508–511, 2007.
- [173] W. H. Chang, "The inductance of a superconducting strip transmission line," *Journal of Applied Physics*, vol. 50, no. 12, pp. 8129–8134, 1979.
- [174] S. Nagasawa and M. Hidaka, "Run-to-run yield evaluation of improved Nb 9-layer advanced process using single flux quantum shift register chip with 68,990 Josephson junctions," *Journal of Physics: Conference Series*, vol. 871, Jul. 2017, Art. no. 012065.
- [175] A. E. Ruehli, "Equivalent circuit models for three-dimensional multiconductor systems," *IEEE Transactions on Microwave Theory and Techniques*, vol. 22, no. 3, pp. 216–221, Mar. 1974.

- [176] C. J. Fourie, O. Wetzstein, T. Ortlepp, and J. Kunert, “Three-dimensional multi-terminal superconductive integrated circuit inductance extraction,” *Superconductor Science and Technology*, vol. 24, no. 12, 2011, Art. no. 125015.
- [177] H. Wei, W. Yu, and Z. Wang, “A fast algorithm for 3-D inductance extraction based on investigation of open-circuit current,” in *2005 International Conference On Simulation of Semiconductor Processes and Devices*, 2005, pp. 203–206.
- [178] S. Zeng, W. Yu, X. Hong, and Z. Wang, “An efficient 3D reluctance extractor for on-chip interconnects,” in *2006 8th International Conference on Solid-State and Integrated Circuit Technology Proceedings*, 2006, pp. 357–359.
- [179] P. Xiao, E. Charbon, A. Sangiovanni-Vincentelli, T. Van Duzer, and S. Whiteley, “INDEX: An inductance extractor for superconducting circuits,” *IEEE Transactions on Applied Superconductivity*, vol. 3, no. 1, pp. 2629–2632, Mar. 1993.
- [180] O. Brandel, O. Wetzstein, T. May, H. Toepfer, T. Ortlepp, and H.-G. Meyer, “RSFQ electronics for controlling superconducting polarity switches,” *Superconductor Science and Technology*, vol. 25, no. 12, Oct. 2012, Art. no. 125012.
- [181] M. Maezawa, M. Suzuki, H. Sasaki, and A. Shoji, “Analog-to-digital converter based on RSFQ technology for radio astronomy applications,” *Superconductor Science and Technology*, vol. 14, no. 12, pp. 1106–1110, Nov. 2001.
- [182] S. Polonsky, J. C. Lin, and A. Rylyakov, “RSFQ arithmetic blocks for DSP applications,” *IEEE Transactions on Applied Superconductivity*, vol. 5, no. 2, pp. 2823–2826, 1995.
- [183] P. Bunyk, M. Leung, J. Spargo, and M. Dorojevets, “Flux-1 RSFQ microprocessor: Physical design and test results,” *IEEE Transactions on Applied Superconductivity*, vol. 13, no. 2, pp. 433–436, 2003.
- [184] A. Inamdar, A. Sahu, and V. Semenov, “Decimation filter with improved DC biasing and data transfer,” *IEEE Transactions on Applied Superconductivity*, vol. 17, no. 2, pp. 482–485, 2007.
- [185] M. M. Khapaev, A. Y. Kidiyarova-Shevchenko, P. Magnelind, and M. Y. Kupriyanov, “3D-MLSI: Software package for inductance calculation in multilayer superconducting integrated circuits,” *IEEE Transactions on Applied Superconductivity*, vol. 11, no. 1, pp. 1090–1093, 2001.
- [186] M. M. Khapaev, M. Y. Kupriyanov, E. Goldobin, and M. Siegel, “Current distribution simulation for superconducting multi-layered structures,” *Superconductor Science and Technology*, vol. 16, no. 1, pp. 24–27, Nov. 2003.
- [187] M. M. Khapaev and M. Y. Kupriyanov, “Sheet current model for inductances extraction and josephson junctions devices simulation,” *Journal of Physics: Conference Series*, vol. 248, Nov. 2010, Art. no. 012041.
- [188] M. M. Khapaev and M. Y. Kupriyanov, “Inductance extraction of superconductor structures with internal current sources,” *Superconductor Science and Technology*, vol. 28, no. 5, Apr. 2015, Art. no. 055013.
- [189] E. J. Romans, S. Rozhko, L. Young, A. Blois, L. Hao, D. Cox, and J. C. Gallop, “Noise performance of niobium nano-SQUIDs in applied magnetic fields,” *IEEE Transactions on Applied Superconductivity*, vol. 21, no. 3, pp. 404–407, 2011.

- [190] L. Longobardi, D. Stornaiuolo, G. Papari, and F. Tafuri, “Feasibility of a high temperature superconductor rf-SQUID based on biepitaxial Josephson junction technology,” *IEEE Transactions on Applied Superconductivity*, vol. 21, no. 3, pp. 151–155, 2011.
- [191] J. Nagel, K. B. Konovalenko, M. Kemmler, M. Turad, R. Werner, E. Kleisz, S. Menzel, R. Klingeler, B. Büchner, R. Kleiner, and D. Koelle, “Resistively shunted YBa₂Cu₃O₇ grain boundary junctions and low-noise SQUIDs patterned by a focused ion beam down to 80 nm linewidth,” *Superconductor Science and Technology*, vol. 24, no. 1, Dec. 2010, Art. no. 015015.
- [192] S. Whiteley. (Feb. 2020). “WRspice,” [Online]. Available: <http://www.wrcad.com/wrspice.html>.
- [193] Y. Mizugaki, R. Kashiwa, M. Moriya, K. Usami, and T. Kobayashi, “Grounding positions of superconducting layer for effective magnetic isolation in Josephson integrated circuits,” *Journal of Applied Physics*, vol. 101, no. 11, 2007, Art. no. 114509.
- [194] B. Guan, M. Wengler, P. Rott, and M. Feldman, “Inductance estimation for complicated superconducting thin film structures with a finite segment method,” *IEEE Transactions on Applied Superconductivity*, vol. 7, no. 2, pp. 2776–2779, 1997.
- [195] C. J. Fourie and W. J. Perold, “Simulated inductance variations in RSFQ circuit structures,” *IEEE Transactions on Applied Superconductivity*, vol. 15, no. 2, pp. 300–303, 2005.
- [196] C. K. Teh, M. Kitagawa, and Y. Okabe, “Inductance calculation of 3D superconducting structures with ground plane,” *Superconductor Science and Technology*, vol. 12, no. 11, pp. 921–924, Nov. 1999.
- [197] C. J. Fourie and W. J. Perold, “On using finite segment methods and images to establish the effect of gate structures on inter-junction inductances in RSFQ circuits,” *IEEE Transactions on Applied Superconductivity*, vol. 13, no. 2, pp. 539–542, 2003.
- [198] C. J. Fourie and W. J. Perold, “Reflection plane placement in numerical inductance calculations using the method of images for thin-film superconducting structures,” *Transactions of the South African Institute of Electrical Engineers*, vol. 94, no. 2, pp. 18–24, 2003.
- [199] N. Deo, *Graph Theory with applications to engineering and computer science*. Prentice-Hall, Inc., 1974.
- [200] O. Mielke, T. Ortlepp, J. Kunert, H.-G. Meyer, and H. Toepfer, “Controlled initialization of superconducting pi-phaseshifters and possible applications,” *Superconductor Science and Technology*, vol. 23, no. 5, Apr. 2010, Art. no. 055003.
- [201] C. J. Fourie, “Calibration of inductance calculations to measurement data for superconductive integrated circuit processes,” *IEEE Transactions on Applied Superconductivity*, vol. 23, no. 3, Jun. 2013, Art. no. 1301305.
- [202] C. J. Fourie, “Full-gate verification of superconducting integrated circuit layouts with InductEx,” *IEEE Transactions on Applied Superconductivity*, vol. 25, no. 1, Feb. 2015, Art. no. 1300209.

- [203] Y. Saad and M. H. Schultz, “GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems,” *SIAM Journal on Scientific and Statistical Computing*, vol. 7, no. 3, pp. 856–869, 1986.
- [204] L. Greengard and V. Rokhlin, “A fast algorithm for particle simulations,” *Journal of Computational Physics*, vol. 73, no. 2, pp. 325–348, 1987.
- [205] M. Benzi, “Preconditioning techniques for large linear systems: A survey,” *Journal of Computational Physics*, vol. 182, no. 2, pp. 418–477, 2002.
- [206] I. Haverkamp, O. Wetzstein, J. Kunert, T. Ortlepp, R. Stolz, H.-G. Meyer, and H. Toepfer, “Optimization of a digital SQUID magnetometer in terms of noise and distortion,” *Superconductor Science and Technology*, vol. 25, no. 6, p. 065 012, Apr. 2012.
- [207] N. Takeuchi, T. Ortlepp, Y. Yamanashi, and N. Yoshikawa, “Novel latch for adiabatic quantum-flux-parametron logic,” *Journal of Applied Physics*, vol. 115, no. 10, 2014, Art. no. 103910.
- [208] M. Kamon, “Efficient techniques for inductance extraction of complex 3-D geometries,” Master of Science, Massachusetts Institute of Technology, 1994.
- [209] K. Jackman and C. J. Fourie, “Fast multicore FastHenry and a tetrahedral modeling method for inductance extraction of complex 3D geometries,” in *2015 15th International Superconductive Electronics Conference (ISEC)*, Jul. 2015, pp. 1–3.
- [210] T. A. Davis, *Direct methods for sparse linear systems*. Society for Industrial and Applied Mathematics, 2006.
- [211] G. Strang, *Introduction to linear algebra*. Wesley-Cambridge Press, 1993.
- [212] J. Demmel, J. R. Gilbert, and X. S. Li, “Superlu users” guide,” *Lawrence Berkeley National Laboratory*, 1997.
- [213] J. W. Demmel, J. R. Gilbert, and X. S. Li, “An asynchronous parallel supernodal algorithm for sparse gaussian elimination,” *SIAM Journal on Matrix Analysis and Applications*, vol. 20, no. 4, pp. 915–952, 1999.
- [214] K. Jackman and C. J. Fourie, “Tetrahedral modeling method for inductance extraction of complex 3-D superconducting structures,” *IEEE Transactions on Applied Superconductivity*, vol. 26, no. 3, Apr. 2016, Art. no. 0602305.
- [215] K. Jackman, “Fast multi-core CEM solvers and flux trapping analysis for superconducting structures,” Ph.D. dissertation, Stellenbosch University, 2017.
- [216] Y. Schols and G. A. Vandenbosch, “Separation of horizontal and vertical dependencies in a surface/volume integral equation approach to model quasi 3-D structures in multilayered media,” *IEEE transactions on antennas and propagation*, vol. 55, no. 4, pp. 1086–1094, 2007.
- [217] J. Markkanen, P. Ylä-Oijala, and S. Järvenpää, “Volume integral equation methods in computational electromagnetics,” in *Electromagnetics in Advanced Applications (ICEAA), 2013 International Conference on*, 2013, pp. 880–883.
- [218] J. Markkanen, P. Ylä-Oijala, and A. Sihvola, “Discretization of volume integral equation formulations for extremely anisotropic materials,” *IEEE Transactions on Antennas and Propagation*, vol. 60, no. 11, pp. 5195–5202, 2012.

- [219] L.-M. Zhang and X.-Q. Sheng, "Solving volume electric current integral equation with full-and half-SWG functions," *IEEE Antennas and Wireless Propagation Letters*, vol. 14, pp. 682–685, 2015.
- [220] M. Li and W. C. Chew, "Applying divergence-free condition in solving the volume integral equation," *Progress In Electromagnetics Research*, vol. 57, pp. 311–333, 2006.
- [221] D. Schaubert, D. Wilton, and A. Glisson, "A tetrahedral modeling method for electromagnetic scattering by arbitrarily shaped inhomogeneous dielectric bodies," *IEEE Transactions on Antennas and Propagation*, vol. 32, no. 1, pp. 77–85, 1984.
- [222] R. F. Harrington and J. L. Harrington, *Field computation by moment methods*. Oxford University Press, 1996.
- [223] C. Geuzaine and J.-F. Remacle, "Gmsh: A 3-D finite element mesh generator with built-in pre-and post-processing facilities," *International journal for numerical methods in engineering*, vol. 79, no. 11, pp. 1309–1331, 2009.
- [224] H.-Y. R. Chao, *A multilevel fast multipole algorithm for analyzing radiation and scattering from wire antennas in a complex environment*. University of Illinois at Urbana-Champaign, 2002.
- [225] H.-G. Wang and Z. Peng, "Combing multilevel green's function interpolation method with volume loop bases for inductance extraction problems," *Progress In Electromagnetics Research*, vol. 80, pp. 225–239, 2008.
- [226] M. Khapaev, "Inductance extraction of multilayer finite-thickness superconductor circuits," *IEEE Transactions on Microwave Theory and Techniques*, vol. 49, no. 1, pp. 217–220, 2001. DOI: 10.1109/22.900014.
- [227] M. Khapaev and M. Y. Kupriyanov, *Sparse Approximation of FEM Matrix for Sheet Current Integro-Differential Equation*, pp. 510–522.
- [228] W. H. Chang, "Measurement and calculation of Josephson junction device inductances," *Journal of Applied Physics*, vol. 52, no. 3, pp. 1417–1426, 1981.
- [229] S. Rao, D. Wilton, and A. Glisson, "Electromagnetic scattering by surfaces of arbitrary shape," *IEEE Transactions on Antennas and Propagation*, vol. 30, no. 3, pp. 409–418, 1982.
- [230] M. M. Khapaev, "Extraction of inductances of plane thin film superconducting circuits," *Superconductor Science and Technology*, vol. 10, no. 6, pp. 389–394, Jun. 1997.
- [231] R. S. Bakolo, "Magnetic modelling, analysis and on-chip shielding of SFQ circuits," Ph.D. dissertation, Stellenbosch University, 2018.
- [232] R. S. Bakolo, R. Van Staden, F. P., and C. J. Fourie, "Modelling magnetic fields and shielding efficiency in superconductive integrated circuits," *Journal of Superconductivity and Novel Magnetism*, vol. 30, no. 6, pp. 1649–16531, 2017.
- [233] C. J. Fourie and K. Jackman, "High-fidelity circuit simulation of AQFP circuits through compact models extracted from layout," *Journal of Physics: Conference Series*, vol. 2323, no. 1, Aug. 2022, Art. No. 012034.
- [234] N. Deo, *Graph theory with applications to engineering and computer science*. Prentice Hall, 1974.

- [235] A. I. Braginski and J. Clarke, *The SQUID handbook Vol. I Fundamentals and technology of SQUIDs and SQUID systems*. Wiley-VCH Verlag GmbH and Co., 2004.
- [236] W. H. Henkels, “Accurate measurement of small inductances or penetration depths in superconductors,” *Applied Physics Letters*, vol. 32, no. 12, pp. 829–831, 1978.
- [237] S. K. Tolpygo, E. B. Golden, T. J. Weir, and V. Bolkhovsky, “Inductance of superconductor integrated circuit features with sizes down to 120 nm,” *Superconductor Science and Technology*, vol. 34, no. 8, Jun. 2021, Art. no. 085005.
- [238] S. K. Tolpygo, V. Bolkhovsky, D. E. Oates, R. Rastogi, S. Zarr, A. L. Day, T. J. Weir, A. Wynn, and L. M. Johnson, “Superconductor electronics fabrication process with MoN_x kinetic inductors and self-shunted Josephson junctions,” *IEEE Transactions on Applied Superconductivity*, vol. 28, no. 4, 2018, Art. no. 1100212.
- [239] C. J. Fourie, X. Peng, R. Numaguchi, and N. Yoshikawa, “Inductance and coupling of stacked vias in a multilayer superconductive IC process,” *IEEE Transactions on Applied Superconductivity*, vol. 25, no. 3, Jun. 2015, Art. no. 1101104.
- [240] D. T. Yohannes, A. Inamdar, and S. K. Tolpygo, “Multi-Jc (Josephson critical current density) process for superconductor integrated circuits,” *IEEE Transactions on Applied Superconductivity*, vol. 19, no. 3, pp. 149–153, Jun. 2009.
- [241] C. Fourie, X. Peng, A. Takahashi, and N. Yoshikawa, “Modeling and calibration of ADP process for inductance calculation with InductEx,” in *2013 IEEE 14th International Superconductive Electronics Conference (ISEC)*, 2013, pp. 1–3.
- [242] C. J. Fourie, A. Takahashi, and N. Yoshikawa, “Fast and accurate inductance and coupling calculation for a multi-layer Nb process,” *Superconductor Science and Technology*, vol. 28, no. 3, Feb. 2015, Art. no. 035013.
- [243] C. J. Fourie, C. Shawawreh, I. V. Vernik, and T. V. Filippov, “High-accuracy InductEx calibration sets for MIT-LL SFQ4ee and SFQ5ee processes,” *IEEE Transactions on Applied Superconductivity*, vol. 27, no. 2, Mar. 2017, Art. no. 1300805.
- [244] C. J. Fourie and M. H. Volkmann, “Status of superconductor electronic circuit design software,” *IEEE Transactions on Applied Superconductivity*, vol. 23, no. 3, Jun. 2013, Art. no. 1300205.
- [245] K. Gaj, Q. Herr, V. Adler, A. Krasniewski, E. Friedman, and M. Feldman, “Tools for the computer-aided design of multigigahertz superconducting digital circuits,” *IEEE Transactions on Applied Superconductivity*, vol. 9, no. 1, pp. 18–38, 1999.
- [246] M. N. Muchuka, J. A. Delpont, and C. J. Fourie, “Superconducting digital circuit design with an open source and freeware tool chain,” *IEEE Transactions on Applied Superconductivity*, vol. 26, no. 8, Dec. 2016, Art. no. 1302008.
- [247] C. J. Fourie, “Digital superconducting electronics design tools - status and roadmap,” *IEEE Transactions on Applied Superconductivity*, vol. 28, no. 5, Aug. 2018, Art. no. 1300412.
- [248] C. Fourie, “Single flux quantum circuit technology and CAD overview,” in *2018 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, Nov. 2018, pp. 1–6.
- [249] C. J. Fourie, “Electronic design automation tools for superconducting circuits,” *Journal of Physics: Conference Series*, vol. 1590, no. 1, Jun. 2020, Art. no. 012040.

- [250] V. K. Semenov and V. P. Zavaleev, “Automation of numerical analysis of the superconducting circuits,” presented at ASC 84: Applied Superconductivity Conference, San Diego, USA, 1984.
- [251] V. K. Semenov, A. A. Odintsov, and A. B. Zorin, “Automation of numerical analysis of circuits with Josephson tunnel junctions,” in *SQUID '85, superconducting quantum interference devices and their applications*, 1985, pp. 71–75.
- [252] N. R. Werthamer, “Nonlinear self-coupling of Josephson radiation in superconducting tunnel junctions,” *Physical Review*, vol. 147, pp. 255–263, 1 Jul. 1966.
- [253] P. Shevchenko. (Feb. 2020). “PSCAN2 - Superconductor Circuit Simulator,” [Online]. Available: <http://pscan2sim.org/>.
- [254] W. C. Stewart, “Current-voltage characteristics of Josephson junctions,” *Applied Physics Letters*, vol. 12, no. 8, pp. 277–280, 1968.
- [255] D. E. McCumber, “Effect of ac impedance on dc voltage-current characteristics of superconductor weak-link junctions,” *Journal of Applied Physics*, vol. 39, no. 7, pp. 3113–3118, 1968.
- [256] S. Whiteley, “Josephson junctions in SPICE3,” *IEEE Transactions on Magnetics*, vol. 27, no. 2, pp. 2902–2905, 1991.
- [257] E. S. Fang, “A Josephson integrated circuit simulator (JSIM) for superconductive electronics application,” in *Extended Abstracts of 1989 International Superconductivity Electronics Conf. (The Japan Society of Applied Physics, Tokyo, 1989)*, 1989.
- [258] J. Satchell, “Stochastic simulation of SFQ logic,” *IEEE Transactions on Applied Superconductivity*, vol. 7, no. 2, pp. 3315–3318, 1997.
- [259] C.-W. Ho, A. Ruehli, and P. Brennan, “The modified nodal approach to network analysis,” *IEEE Transactions on Circuits and Systems*, vol. 22, no. 6, pp. 504–509, 1975.
- [260] J. A. Delport, “Simulation and verification software for superconducting electronic circuits,” Ph.D. dissertation, Stellenbosch University, 2019.
- [261] J. A. Delport, K. Jackman, P. le Roux, and C. J. Fourie, “JoSIM—superconductor SPICE simulator,” *IEEE Transactions on Applied Superconductivity*, vol. 29, no. 5, Aug. 2019, Art. no. 1300905.
- [262] C. F. Curtiss and J. O. Hirschfelder, “Integration of stiff equations,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 38, no. 3, pp. 235–243, 1952.
- [263] T. A. Davis and E. Palamadai Natarajan, “Algorithm 907: KLU, a direct sparse solver for circuit simulation problems,” *ACM Transactions on Mathematical Software*, vol. 37, no. 3, 2010.
- [264] T. Weingartner, N. Pokhrel, M. Sulangi, L. Bjorndal, E. Patrick, and M. E. Law, “Modeling process and device behavior of Josephson junctions in superconductor electronics with TCAD,” *IEEE Transactions on Electron Devices*, vol. 68, no. 11, pp. 5448–5454, 2021.
- [265] H. F. Herbst, P. le Roux, K. Jackman, and C. J. Fourie, “Improved transmission line parameter calculation through TCAD process modeling for superconductor integrated circuit interconnects,” in *2019 IEEE International Superconductive Electronics Conference (ISEC)*, Aug. 2019, pp. 1–5.

- [266] H. F. Herbst, P. le Roux, K. Jackman, and C. J. Fourie, “Improved transmission line parameter calculation through TCAD process modeling for superconductor integrated circuit interconnects,” *IEEE Transactions on Applied Superconductivity*, vol. 30, no. 7, Oct. 2020, Art. no. 1100504.
- [267] W. Perold and C. Fourie, “Modeling superconducting components based on the fabrication process and layout dimensions,” *IEEE Transactions on Applied Superconductivity*, vol. 11, no. 1, pp. 345–348, 2001.
- [268] M. Jeffery, W. Perold, Z. Wang, and T. van Duzer, “Monte carlo optimization of superconducting complementary output switching logic circuits,” *IEEE Transactions on Applied Superconductivity*, vol. 8, no. 3, pp. 104–119, 1998.
- [269] M. Johnson, Q. Herr, and J. Spargo, “Monte-carlo yield analysis,” *IEEE Transactions on Applied Superconductivity*, vol. 9, no. 2, pp. 3322–3325, 1999.
- [270] C. Fourie, W. Perold, and H. Gerber, “Complete Monte Carlo model description of lumped-element RSFQ logic circuits,” *IEEE Transactions on Applied Superconductivity*, vol. 15, no. 2, pp. 384–387, 2005.
- [271] C. Hamilton and K. Gilbert, “Margins and yield in single flux quantum logic,” *IEEE Transactions on Applied Superconductivity*, vol. 1, no. 4, pp. 157–163, 1991.
- [272] Q. Herr and M. Feldman, “Multiparameter optimization of RSFQ circuits using the method of inscribed hyperspheres,” *IEEE Transactions on Applied Superconductivity*, vol. 5, no. 2, pp. 3337–3340, 1995.
- [273] S. Polonsky, P. Shevchenko, A. Kirichenko, D. Zinoviev, and A. Rylyakov, “PSCAN’96: new software for simulation and optimization of complex RSFQ circuits,” *IEEE Transactions on Applied Superconductivity*, vol. 7, no. 2, pp. 2685–2689, 1997.
- [274] C. J. Fourie and W. J. Perold, “Comparison of genetic algorithms to other optimization techniques for raising circuit yield in superconducting digital circuits,” *IEEE Transactions on Applied Superconductivity*, vol. 13, no. 2, pp. 511–514, 2003.
- [275] C. J. Fourie and W. J. Perold, “Yield optimization of high frequency superconducting digital circuits with genetic algorithms,” *Transactions of the South African Institute of Electrical Engineers*, vol. 94, no. 2, pp. 11–17, 2003.
- [276] J. R. Koza, *Genetic programming: on the programming of computers by means of natural selection*. MIT Press, 1993.
- [277] T. Back, U. Hammel, and H.-P. Schwefel, “Evolutionary computation: Comments on the history and current state,” *IEEE Transactions on Evolutionary Computation*, vol. 1, no. 1, pp. 3–17, 1997.
- [278] N. Mori, A. Akahori, T. Sato, N. Takeuchi, A. Fujimaki, and H. Hayakawa, “A new optimization procedure for single flux quantum circuits,” *Physica C: Superconductivity*, vol. 357-360, pp. 1557–1560, 2001.
- [279] 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, 1997.
- [280] 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, 2011.

- [281] P. le Roux and C. J. Fourie, “Distance-to-failure-maximization optimization algorithm for SFQ logic cells,” *IEEE Transactions on Applied Superconductivity*, vol. 30, no. 7, Oct. 2020, Art. no. 1301405.
- [282] L. C. Müller and C. J. Fourie, “Automated state machine and timing characteristic extraction for RSFQ circuits,” *IEEE Transactions on Applied Superconductivity*, vol. 24, no. 1, Feb. 2014, Art. no. 1300110.
- [283] A. Inamdar, J. Ren, and D. Amparo, “Improved model-to-hardware correlation for superconductor integrated circuits,” *IEEE Transactions on Applied Superconductivity*, vol. 25, no. 3, 2015, Art. no. 1300308.
- [284] P. le Roux, K. Jackman, J. A. Delpont, and C. J. Fourie, “Modeling of superconducting passive transmission lines,” *IEEE Transactions on Applied Superconductivity*, vol. 29, no. 5, Aug. 2019, Art. no. 1101605.
- [285] P. le Roux, C. Fourie, S. Razmkhah, and P. Febvre, “Accurate small signal simulation of superconductor interconnects in SPICE,” *IEEE Transactions on Applied Superconductivity*, vol. 31, no. 5, Aug. 2021, Art. no. 1303006.
- [286] L. Schindler, P. le Roux, and C. J. Fourie, “Impedance matching of passive transmission line receivers to improve reflections between RSFQ logic cells,” *IEEE Transactions on Applied Superconductivity*, vol. 30, no. 2, Mar. 2020, Art. no. 1300607.
- [287] M. Pedram, “Superconductive Single Flux Quantum logic devices and circuits: Status, challenges, and opportunities,” in *2020 IEEE International Electron Devices Meeting (IEDM)*, 2020, pp. 25.7.1–25.7.4.
- [288] J. F. De Villiers, “Automated synthesis, placement and routing of large-scale RSFQ integrated circuits,” Masters thesis, Stellenbosch University, 2021.
- [289] E. Verburg, J. F. De Villiers, and C. J. Fourie, “Viper - a tool chain for synthesis placement and routing of combinational rsfq circuits,” *IEEE Transactions on Applied Superconductivity*, Submitted for publication.
- [290] J. A. Delpont and C. J. Fourie, “A static timing analysis tool for RSFQ and ERSFQ superconducting digital circuit applications,” *IEEE Transactions on Applied Superconductivity*, vol. 28, no. 5, Aug. 2018, Art. no. 1300705.
- [291] B. Zhang and M. Pedram, “qSTA: a static timing analysis tool for superconducting single-flux-quantum circuits,” *IEEE Transactions on Applied Superconductivity*, vol. 30, no. 5, 2020, Art. no. 1700309.
- [292] V. Adler, C.-H. Cheah, K. Gaj, D. Brock, and E. Friedman, “A cadence-based design environment for single flux quantum circuits,” *IEEE Transactions on Applied Superconductivity*, vol. 7, no. 2, pp. 3294–3297, 1997.
- [293] R. M. C. Roberts and C. J. Fourie, “Layout-versus-schematic verification for superconductive integrated circuits,” *IEEE Transactions on Applied Superconductivity*, vol. 25, no. 3, Jun. 2015, Art. no. 1200105.
- [294] H. Terai, Y. Kameda, S. Yorozu, A. Fujimaki, and Z. Wang, “The effects of DC bias current in large-scale SFQ circuits,” *IEEE Transactions on Applied Superconductivity*, vol. 13, no. 2, pp. 502–506, 2003.

- [295] R. van Staden, K. Jackman, C. J. Fourie, and P. Febvre, "Influence of the superconducting ground plane on the performance of RSFQ cells," *IEEE Transactions on Applied Superconductivity*, vol. 27, no. 4, Jun. 2017, Art. no. 1300704.
- [296] H. Akaike, A. Fujimaki, T. Satoh, K. Hinode, S. Nagasawa, Y. Kitagawa, and M. Hidaka, "Effects of a DC-power layer under a ground plane in SFQ circuits," *IEEE Transactions on Applied Superconductivity*, vol. 17, no. 2, pp. 466–469, 2007.
- [297] H. Akaike, K. Shigehara, A. Fujimaki, T. Satoh, K. Hinode, S. Nagasawa, and M. Hidaka, "The effects of a DC power layer in a 10-Nb-layer device for SFQ LSIs," *IEEE Transactions on Applied Superconductivity*, vol. 19, no. 3, pp. 594–597, 2009.
- [298] K. Fujiwara, S. Nagasawa, Y. Hashimoto, M. Hidaka, N. Yoshikawa, M. Tanaka, H. Akaike, A. Fujimaki, K. Takagi, and N. Takagi, "Research on effective moat configuration for nb multi-layer device structure," *IEEE Transactions on Applied Superconductivity*, vol. 19, no. 3, pp. 603–606, Jun. 2009.
- [299] C. J. Fourie, N. Takeuchi, and N. Yoshikawa, "Inductance and current distribution extraction in Nb multilayer circuits with superconductive and resistive components," *IEICE Transactions on Electronics*, vol. E99-C, no. 6, pp. 683–691, Jun. 2016.
- [300] K. Jackman and C. J. Fourie, "Layout strategies for connecting multiple superconducting ground planes with ground pillars," in *2019 IEEE International Superconductive Electronics Conference (ISEC)*, Aug. 2019, pp. 1–4.
- [301] K. Jackman and C. J. Fourie, "Multipole accelerated magnetic field calculations for superconducting circuits," *Superconductor Science and Technology*, vol. 32, no. 1, 2019, Art. no. 015011.
- [302] C. Fourie, S. M. Anton, and J. Clarke, "Magnetic field calculations in the vicinity of superconductive circuit structures," in *2013 IEEE 14th International Superconductive Electronics Conference (ISEC)*, 2013, pp. 1–3.
- [303] R. Collot, P. Febvre, J.-L. Issler, T. Robert, C. Fourie, J. Kunert, R. Stolz, and H.-G. Meyer, "Influence of external magnetic fields on the inductive properties of rapid single-flux-quantum digital circuits," in *2013 IEEE 14th International Superconductive Electronics Conference (ISEC)*, 2013, pp. 1–3.
- [304] R. S. Bakolo, J. A. Delport, P. Febvre, and C. J. Fourie, "Analysis of a shielding approach for magnetic field tolerant SFQ circuits," *IEEE Transactions on Applied Superconductivity*, vol. 27, no. 4, Jun. 2017, Art. no. 1301305.
- [305] C. J. Fourie and K. Jackman, "Software tools for flux trapping and magnetic field analysis in superconducting circuits," *IEEE Transactions on Applied Superconductivity*, vol. 29, no. 5, Aug. 2019, Art. no. 1301004.
- [306] R. Collot, P. Febvre, J. Kunert, and H.-G. Meyer, "Operation of low- T_c circuits in a magnetic environment," *IEEE Transactions on Applied Superconductivity*, vol. 23, no. 3, 2013, Art. no. 1700404.
- [307] T. V. Filippov, A. Sahu, S. Sarwana, D. Gupta, and V. K. Semenov, "Serially biased components for digital-RF receiver," *IEEE Transactions on Applied Superconductivity*, vol. 19, no. 3, pp. 580–584, Jun. 2009.
- [308] V. K. Semenov and Y. Polyakov, "Current recycling: New results," *IEEE Transactions on Applied Superconductivity*, vol. 29, no. 5, Aug. 2019, Art. no. 1302304.

- [309] L. P. Lee, J. Longo, V. Vinetskiy, and R. Cantor, “Low-noise $\text{YBa}_2\text{Cu}_3\text{O}_{7-\delta}$ direct-current superconducting quantum interference device magnetometer with direct signal injection,” *Applied Physics Letters*, vol. 66, no. 12, pp. 1539–1541, 1995.
- [310] D. Koelle, A. H. Miklich, F. Ludwig, E. Dantsker, D. T. Nemeth, and J. Clarke, “Dc squid magnetometers from single layers of $\text{YBa}_2\text{Cu}_3\text{O}_{7-\delta}$,” *Applied Physics Letters*, vol. 63, no. 16, pp. 2271–2273, 1993.
- [311] M. Ketchen, “Integrated thin-film dc squid sensors,” *IEEE Transactions on Magnetism*, vol. 23, no. 2, pp. 1650–1657, 1987.
- [312] J. Jaycox and M. Ketchen, “Planar coupling scheme for ultra low noise DC SQUIDS,” *IEEE Transactions on Magnetism*, vol. 17, no. 1, pp. 400–403, 1981.
- [313] S. M. Anton, I. A. B. Sognaes, 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, no. 7, 2013, Art. no. 075022.
- [314] M. A. Washington and T. A. Fulton, “Observation of flux trapping threshold in narrow superconducting thin films,” *Applied Physics Letters*, vol. 40, no. 9, pp. 848–850, 1982.
- [315] S. Bermon and T. Gheewala, “Moat-guarded Josephson SQUIDS,” *IEEE Transactions on Magnetism*, vol. 19, no. 3, pp. 1160–1164, May 1983.
- [316] Y. Polyakov, S. Narayana, and V. K. Semenov, “Flux trapping in superconducting circuits,” *IEEE Transactions on Applied Superconductivity*, vol. 17, no. 2, pp. 520–525, Jun. 2007.
- [317] 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. 2, pp. 1769–1771, Sep. 1995.
- [318] R. Robertazzi, I. Siddiqi, and O. Mukhanov, “Flux trapping experiments in single flux quantum shift registers,” *IEEE Transactions on Applied Superconductivity*, vol. 7, no. 2, pp. 3164–3167, Jun. 1997.
- [319] S. Narayana, Y. A. Polyakov, and V. K. Semenov, “Evaluation of flux trapping in superconducting circuits,” *IEEE Transactions on Applied Superconductivity*, vol. 19, no. 3, pp. 640–643, Jun. 2009.
- [320] Y. Yamanashi, H. Imai, and N. Yoshikawa, “Influence of magnetic flux trapped in moats on superconducting integrated circuit operation,” *IEEE Transactions on Applied Superconductivity*, vol. 28, no. 7, Oct. 2018, Art. no. 1301105.
- [321] K. Jackman and C. J. Fourie, “Flux trapping analysis in superconducting circuits,” *IEEE Transactions on Applied Superconductivity*, vol. 27, no. 4, Jun. 2017, Art. no. 1300105.
- [322] K. Jackman and C. J. Fourie, “Flux trapping experiments to verify simulation models,” *Superconductor Science and Technology*, vol. 33, no. 10, 2020, Art. no. 105001.
- [323] C. J. Fourie and K. Jackman, “Experimental verification of moat design and flux trapping analysis,” *IEEE Transactions on Applied Superconductivity*, vol. 31, no. 5, Aug. 2021, Art. no. 1300507.