


Performance Evaluation of Wireless Communication in a Modular Heliostat Field

by

Andreas Denan Liebenberg



*Thesis presented in partial fulfilment of the requirements for the
degree of Master of Engineering (Mechatronic) in the Faculty of
Engineering at Stellenbosch University*

Supervisor: Dr. W.J. Smit

March 2021

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.

Date: 2020/11/20

Copyright © 2021 Stellenbosch University
All rights reserved.

Abstract

Performance Evaluation of Wireless Communication in a Modular Heliostat Field

A.D. Liebenberg

*Department of Mechanical and Mechatronic Engineering,
University of Stellenbosch,
Private Bag X1, Matieland 7602, South Africa.*

Thesis: MEng (Mechatronic)

March 2021

Concentrating Solar Thermal (CST) is differentiated from other renewable energy technologies by the ability to store large amounts of thermal energy cost-effectively. CST, and especially Central Receiver Systems (CRS), still require further development to drive down the Levelised Cost of Electricity (LCOE) to remain economically competitive in the renewable energy market.

Heliostats are used in CRS to reflect sunlight onto a central receiver. Data transmission is required between a central server and Local Control Units (LCUs) to enable high precision tracking, over-the-air calibration updates, and swift emergency reaction. Wireless communication technologies have shown both capital cost and construction time saving potential for CST plants, but practical performance knowledge of operation in heliostat fields is limited.

The purpose of this study is to evaluate the feasibility of wireless communication within a CST field by means of practical tests aimed to determine the influence of interference, dense networks, and transmission range on the network performance. The project forms part of the H2020 PREMA project, which aims to advance novel energy systems in the drying and pre-heating of furnace manganese ore.

A communication network consisting of 50 LCUs, capable of controlling 300 heliostats, was designed and implemented to simulate data transmission in a CST network. Performance experiments were conducted at the Helio100 site to determine the effect of various performance limiting factors. Experimental results showed a high performing, robust wireless communication system with a worst-case polling rate of 0.79 s, which satisfies a major requirement of the communication network. The thesis concludes that wireless communication is feasible in a

ABSTRACT

iii

CST field, which could lead to the reduction of capital costs and construction time, reducing the LCOE of CST plants.

Uittreksel

Prestasie Evaluering van Draadlose Kommunikasie in 'n Modulêre Heliostaatveld

(“Performance Evaluation of Wireless Communication in a Modular Heliostat Field”)

A.D. Liebenberg

Departement Meganiese en Megatroniese Ingenieurswese,

Universiteit van Stellenbosch,

Privaatsak X1, Matieland 7602, Suid Afrika.

Tesis: MEng (Megatronies)

March 2021

Gekonsentreerde Sonkrag (GSK) word onderskei van ander hernubare energie tegnologieë deur die vermoë om groot hoeveelhede termiese energie koste effektief te stoor. Gekonsentreerde sonkrag, en veral sentrale ontvangersisteme, benodig nog verdere ontwikkeling om die gelykte koste van elektrisiteit te verlaag om ekonomies kompetend te bly in die hernubare energie mark.

Heliostate word in sentrale ontvangersisteme gebruik om sonlig op 'n sentrale ontvanger te weerkaats. Data-oordrag is nodig tussen 'n sentrale bediener en plaaslike beheer eenhede om hoë akkuraatheid volgingsbeweging, oor die lug kalibrasie-opdaterings en vinnige noodreaksie moontlik te maak. Draadlose kommunikasietegnologie het die potensiaal vir kapitaal kostebesparings en konstruksie tydbesparings vir GSK-aanlegte getoon, maar praktiese kennis oor die netwerk gedrag in heliostaat velde is beperk.

Die doel van hierdie studie is om die uitvoerbaarheid van draadlose kommunikasie binne 'n sentrale ontvangersisteem te evalueer aan die hand van praktiese toetse om die invloed van interferensie, digte netwerke en die transmissie afstand, op die netwerk werking te bepaal. Die projek vorm deel van die H2020 PREMA-projek, wat streef om nuwe energie stelsels te bevorder vir die droog en voorverhitting prosesse van mangaanerts.

'n Kommunikasienetwerk bestaande uit 50 plaaslike beheer eenhede, wat 300 heliostate kan beheer, is ontwerp en geïmplementeer om data-oordrag in 'n gekonsentreerde sonkrag netwerk te simuleer. 'n Eksperimentele ondersoek is op die Helio100-terrein gedoen om die effek van verskillende netwerk beperkende faktore te bepaal.

Eksperimentele resultate toon 'n hoë prestasie, robuuste draadlose kommunikasiesistelsel wat, in die slegste geval, die hele netwerk binne 0.79 s kan bedien en so 'n hoofvereiste van die kommunikasie netwerk bevredig. Die tesis het tot die gevolgtrekking gekom dat draadlose kommunikasie binne 'n GSK heliostaat veld uitvoerbaar is en dat dit moontlik tot 'n verlaging in kapitaal koste en konstruksie tyd in sentrale ontvangersisteme kan lei en die gelykte koste van elektrisiteit van gekonsentreerde sonkrag aanlegte kan verminder.

Acknowledgements

The author would like to express his sincere gratitude to the following people and organisations:

A special thanks to PREMA who made this project a reality. I thank you for the continuous financial support.

A special thanks to CSIR for their continuous support throughout the span of the project.

The solar thermal energy research group for their constant support with facilities, hardware and advice.

My parents, Denny and Andreas, for their unconditional love, support and prayer.

Dr Willie Smit for his knowledge and guidance throughout the project.

Contents

Declaration	i
Abstract	ii
Uittreksel	iv
Acknowledgements	vi
Contents	vii
List of Figures	ix
List of Tables	xiii
Nomenclature	xv
1 Introduction	1
1.1 Solar energy	1
1.2 Problem statement	7
1.3 Objectives	7
1.4 Methodology	8
1.5 Scope	8
1.6 Document outline	9
2 Literature Study	10
2.1 Wireless communication technologies	10
2.2 Wireless communication development in CSP plants	21
2.3 Hindrances of wireless communication in a CST field	23
2.4 ZeroMQ	26
2.5 Measuring network performance	28
3 System Requirements	31
3.1 Field sizing	32
3.2 Control interval	33
3.3 Message structure	35

CONTENTS**viii**

4	Hardware Setup	37
4.1	Proposed topology and communication standard	37
4.2	Hardware overview	40
5	Software Setup	43
5.1	ZeroMQ message packet structure	43
5.2	Network architecture and processes	45
5.3	Determining the time offset	50
6	Experimental Setup	55
6.1	Test environments	55
6.2	Test parameters	59
7	Results	61
7.1	Initial test results	61
7.2	Field tests	65
7.3	Discussion	75
8	Conclusion	79
8.1	Summary of work	79
8.2	Conclusions	80
8.3	Summary of contributions	81
8.4	Recommendations for further research	81
	Appendices	82
A	Statistical Analysis of Resulting Field Test Data	83
	Bibliography	85

List of Figures

1.1	DNI) world map [1]. It can be seen that DNI levels in the Northern Cape exceed 3000 kWh/(m ² a).	2
1.2	Solar radiation from the sun is focussed on a solar receiver by solar concentrators. The harvested thermal energy can be used in a cycle or stored as thermal energy. The entire system is referred to as the collector. Adapted from [2].	3
1.3	Classification of solar power plants (adapted from [3]).	3
1.4	Line focus solar collectors: Parabolic trough (left) and linear Fresnel (right) [4].	4
1.5	Point focus solar collectors: Parabolic dish (left) and central receiver system (CRS) (right) [4].	4
1.6	Thermal energy storage tank for molten salt CSP applications [5]. . . .	5
1.7	An aerial view of South Africa's first CSP CRS, Khi Solar One, located in the Northern Cape [6].	5
1.8	Aerial view of the construction of the Helio100 facility located outside the town of Stellenbosch. The receiver on top of the tower, the HelioPods and the control container can be seen.	6
1.9	HelioPod structure as seen from behind (left) and from the front (right). A HelioPod consists of six heliostats mounted on a triangular frame that is not fixed to the ground.	7
2.1	NB-IoT modes for operation [7]. NB-IoT can operate inside the LTE narrow bands that not used for transmission, inside the guard bands, or deployed within unused GSM spectrum.	12
2.2	A typical BLE mesh network adapted from [8]. Various node types are presented in their typical arrangement.	15
2.3	CSMA/CA in a dense network scenario, adapted from [8]. It can be seen that even though WLAN 3 operates in the widest channel, interference causes a deteriorating network performance.	19
2.4	Node placement of the HELIOMESH project by Kubish et al. [9]. Nodes are placed to tests dense networks and long, multi-hop, transmissions.	22
2.5	Friss free space signal attenuation model. The Friss free space equation describes a lower received power with an increase in transmission distance [10].	24

LIST OF FIGURES

x

2.6	Ground effect of reflected electromagnetic waves. The receiver antenna receives undesired copies of the same message that can be delayed in time and shifted in phase and frequency [11].	25
2.7	Message overhead for a 30 byte message sent over Wi-Fi with the ZeroMQ messaging protocol. It can be seen that 64 bytes of overhead data is required to transmit a 30 byte data payload [12].	25
2.8	Simple ZeroMQ REQ/REP messaging structure [13].	27
2.9	ZeroMQ Transfer delay tests from Zhou et al. [14].	28
3.1	Proposed CST plant for PREMA H2020 project [15].	31
3.2	Proposed field layout as seen from above with tower position at (0,0). It can be seen that the maximum data transmission range is approximately 80 m.	32
3.3	Solar subtending angle (adapted from [16]).	33
3.4	Area of intersection of a solar image from an offset heliostat and its intended receiver target (adapted from [3]).	33
3.5	Heliostat ratio of interception versus normal vector angular tracking error	34
4.1	Common network topologies (adapted from [17]).	38
4.2	The AP was positioned on top of the control container, behind the heliostat field. The receiver can be seen located on the tower which is visible above the horizon.	41
5.1	Simple ZeroMQ request-reply envelope (adapted from [13]). The empty delimiter frame and the data frame can be seen in the message envelope.	44
5.2	ZeroMQ extended request reply pattern (adapted from [13]). The router-dealer combination is located between a request-reply pattern. This allows messages to be routed to specified destinations.	44
5.3	ZeroMQ message envelope sent from a REQ socket to a ROUTER socket (adapted from [13]). The ROUTER socket adds the connection identity frame of the sender on top of the empty delimiter frame.	44
5.4	Network architecture overview. The network was designed using ZeroMQ and TCP communication. A total of three socket pairs are present between the server and the broker process (both hosted on the same machine). Two socket pairs are present between the broker and LCU processes and a publish socket is used for emergency situations.	46
5.5	Server process logic flow diagram.	48
5.6	Broker process logic flow diagram.	49
5.7	LCU process logic flow diagram.	50
5.8	Analysis of taking a message timestamp versus duration on two different devices.	51
5.9	Offset process messaging sequence	53
5.10	Latency measurements for five LCUs connected to the server. The latency measurements drift due to the clocks of the various processors drifting.	53

5.11	Latency measurement with NTP protocol enabled. No drift in latency is present as the clock speeds of the LCUs are updated to match the server.	54
6.1	Placement of the LCUs for the baseline test. LCUs were placed in an noise free open field with LoS to the AP at a range of 10-35 m.	56
6.2	HelioStat LoS test LCU placement. LCUs fitted on the frame of the heliostat with velcro ensuring LoS to the AP located on top of the control container.	58
6.3	HelioStat non-LoS test LCU placement. LCUs fitted below the frame of the heliostat with velcro blocking LoS to the AP located on top of the control container.	58
6.4	Dense network LCU deployment in heliostat field. Two heliostats were placed on each heliostat facet to total 12 LCUs per HelioPod.	59
7.1	High uplink latencies found during initial testing with ten LCUs showing a recurring latency around 610 ms.	62
7.2	Moving average plot of the last 100 messages showing the relationship between payload size, latency and throughput. Five LCUs were used in the test.	63
7.3	Uplink latency of initial tests, including five LCUs, with a varying polling strategy. The polling strategies included a long, short, and optimized design. The optimized solution outperforms both the initial long and short polling rates.	65
7.4	Baseline test UL and DL latencies. Latencies plotted per message for the duration of the test.	66
7.5	Baseline UL and DL latency. Latencies plotted on a normalised histogram with latencies grouped in 1 ms bins.	67
7.6	Baseline UL and DL throughput. Throughput plotted on a normalised histogram with throughput grouped in bins of 10 kbps.	68
7.7	Baseline RTT latency. RTT latencies plotted on a normalised histogram with latencies grouped in 1 ms bins.	68
7.8	RTT latency for a network capacity of 10, 30, and 50 (baseline) respectively. RTT latencies plotted on a normalised histogram with latencies grouped in 1 ms bins.	70
7.9	DL latency for the 60 m single LCU range test, 110 m single LCU range test and the 85 m-110 m full capacity (50 LCU) range test. DL latencies are plotted on a normalised histogram with latencies grouped in 1 ms bins.	71
7.10	RTT latency of the single LCU 60 m and 110 m range test and the 85 m-110 m full capacity (50 LCUs) range test. RTT latencies are plotted on a normalised histogram with latencies grouped in 1 ms bins.	72
7.11	RTT latency of the heliostat LoS test and the heliostat non-LoS test. RTT latencies are plotted on a normalised histogram frequency plot with latencies grouped in 1 ms bins.	73

*LIST OF FIGURES***xii**

7.12	RTT latency of the heliostat low density test and the heliostat high density test. RTT latencies are plotted on a normalised histogram with latencies grouped in 1 ms bins.	75
7.13	Gaussian distribution of the RTT latency for the heliostat non-LoS test following sampling from the original dataset as described by the central limit theorem.	76
7.14	Box and whisker plots of normalised data to visualise the influence of different testing environments on the RTT latency results.	77

List of Tables

3.1	A typical polling and ACK process message structure	35
4.1	Wireless communication technology comparison. Each wireless standard is specified a topology as it would be implemented in a CST plant. The main theoretical performance metrics are compared.	39
4.2	Overview of the main hardware specifications of the three primary components of the network.	42
6.1	Test environment overview. This table provides the capacity, range, placement and reasoning for each of the 10 field tests.	56
6.2	Test parameters and AP settings for field tests.	60
7.1	Baseline test latency result.	66
7.2	Baseline test throughput result.	66
7.3	Latency test results comparing the 10 LCU network capacity, 30 LCU network capacity, and baseline test.	69
7.4	Throughput test results comparing the 10 LCU network capacity, 30 LCU network capacity, and baseline test.	69
7.5	Latency test results comparing the single LCU at 60 m range, LCU at 110 m range, full capacity range test and baseline test.	70
7.6	Throughput test results comparing the single LCU at 60 m range, LCU at 110 m range, full capacity range test and baseline test.	70
7.7	Latency test results comparing the LoS heliostat test, non-LoS heliostat test and baseline test.	72
7.8	Throughput test results comparing the LoS heliostat test, non-LoS heliostat test and baseline test.	73
7.9	Latency test results comparing the dense LoS heliostat test, the single LCU LoS heliostat test, the LoS heliostat test, the 60 m single LCU range test, and the baseline test.	74
7.10	Throughput test results comparing the dense LoS heliostat test, the single LCU LoS heliostat test, the LoS heliostat test, the 60 m single LCU range test, and the baseline test.	74

*LIST OF TABLES***xiv**

A.1	Results of t-tests for each test combination. A negative t-value indicates that the mean of data set 2 is larger compared to data set 1. To reject the null hypothesis the absolute value of the resulting t-value should exceed 2.576.	84
-----	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	----

Nomenclature

Variables

$A_{\text{intercept}}$	Area of interception	[m ²]
d	Distance	[m]
E_{tracking}	Heliostat's normal vector angular tracking error	[rad]
G_R	Gain of receiving antenna	[]
G_T	Gain of transmitting antenna	[]
L_n	Latency of the n th message	[s]
L_{slant}	Slant range	[m]
N	Number of messages	[]
n	Path loss exponent	[]
P_R	Power of receiving antenna	[W]
P_{size}	Data payload size	[bytes]
P_T	Power of transmitting antenna	[W]
R	Ideal heliostat reflected image radius	[m]
R_i	Ratio of interception	[]
R_n	Receive timestamp of the n th message	[s]
S_n	Sending timestamp of the n th message	[s]
T_R	Timestamp when message was received	[s]
T_S	Timestamp when message was sent	[s]
T_1	Time before first message was sent	[s]
T_2	Time after last message was received	[s]
$T_{\text{SinglePacket}}$	Average single packet transfer time	[s]
TP_{avg}	Average throughput over a selected period	[kbps]
TP_{ota}	Over-the-air throughput	[kbps]
t_c	Control interval	[s]
α	Two tail t test significance level	[]
δ_t	Time offset	[s]
θ_{delay}	Tracking delay angle	[rad]
θ	Standard deviation	[]

λ	Wavelength of electromagnetic wave [m]
μ	mean []
σ slope	mean slope error [rad]
Ω_s	Sun's subtending angle [rad]
ω_{earth}	Angular velocity of the earth [rad/s]

Abbreviations

3GPP	Third Generation Partnership Project
ACK	Acknowledgment
AODV	Ad-hoc On-demand Distance Vector
AP	Access Point
API	Application Interface
BPSK	Binary Phase Shift Keying
BSS	Basic Service Set
CCA	Clear Channel Assessment
CRC	Cyclic Redundancy Check
CRS	Central Receiver System
CSMA/CA	Carrier Sense Multiple Access with Collision Avoidance
CSP	Concentrating Solar Power
CSS	Chirp Spread Spectrum
CST	Concentrating Solar Thermal
CTS	Clear To Send
DL	Downlink
DLR	German Aerospace Center
DNI	Direct Normal Irradiation
EMI	Electromagnetic Interference
FCU	Field Control Unit
FHSS	Frequency Hopping Spread Spectrum
GATT	Generic Attribute
GI	Guard Interval
GFSK	Gaussian Frequency Shift Keying
GPRS	General Packet Radio Service
GSM	Global System for Mobile Communications
IEEE	Institute of Electrical and Electronic Engineers
IoT	Internet of Things
IRP	Integrated Resource Plan
ISM	Industrial, Scientific and Medical

LCOE	Levelised Cost of Energy
LCU	Local Control Unit
LoRa	Long Range
LoRaWAN	LoRa Wide Area Network
LPWAN	Low Power Wide Area Network
LTE	Long Term Evolution
M2M	Machine to Machine
MAC	Medium Access Control
MIMO	Multiple-Input Multiple-Output
MTU	Maximum Transmission Unit
MU-MIMO	Multi User Multiple-Input Multiple-Output
NB-IoT	Narrow Band Internet of Things
NREL	National Renewable Energy Laboratory
OFDM	Orthogonal Frequency Division Multiplexing
OFDMA	Orthogonal Frequency Division Multiple Access
OQPSK	Offset Quadrature Phase Shift Keying
PAN	Personal Area Network
PER	Package Error Ratio
PHY	Physical
PV	Photovoltaic
QAM	Quadrature Amplitude Modulation
QoS	Quality of Service
RAW	Restricted Access Window
REI4P	Renewable Energy Independent Power Producers Procurement Program
RF	Radio frequency
RFI	Radio Frequency Interference
RPi	Raspberry Pi
RPS	RAW Parameter Set
RSS	Received Signal Strength
RTS	Request to Send
RTT	Round Trip Time
RU	Resource Unit
SC-FDMA	Single Carrier Frequency Division Multiple Access
SCP	Secure Copy
SD	Standard Deviation
SNR	Signal to Noise Ratio

*NOMENCLATURE***xviii**

STERG	Solar Thermal Energy Research Group
TCP	Transmission Control Protocol
TDMA	Time Division Multiple Access
TIM	Traffic Identification Map
TTL	Time To Live
TWT	Target Wake Time
UL	Uplink
WLAN	Wireless Local Area Network
WMN	Wireless Mesh Network
WSN	Wireless Sensor Network
ZMQ	ZeroMQ Message Transfer Protocol

Chapter 1

Introduction

Thomas Edison famously stated, "I'd put my money on the sun and solar energy. What a source of power! I hope we don't have to wait until oil and coal run out before we tackle that" [18]. Climate change has driven a worldwide movement to accelerate the implementation of renewable and sustainable energy generation. This is due to the limited amount of fossil fuels being exhausted at a rapid rate and the destructive environmental impact thereof. Renewable energy is harvested either directly from the sun (in forms such as thermal, photo-chemical, and photo-electric) or indirectly from the sun (in forms such as wind, hydropower, and photosynthetic energy stored in biomass). It can also be harnessed from other natural mechanisms such as tidal and geothermal energy [19]. In addition to providing environmental, social, and economic benefits, renewable energy is also continuously replenished by nature. This ensures the planet's future energy production as it is predicted that global oil, coal, and gas reserves will be depleted by the end of the century [20].

In 2018 South Africa emitted 8.2 tons of CO₂ per capita, making it the 14th highest CO₂ emitting country in 2018 [21]. The implementation of wind power and Photovoltaic (PV) panels help reduce the dependency on fossil fuels. However, the intermittent nature of wind and PV combined with the inability to cost-effectively store electrical energy at a significant scale, render these technologies inadequate for baseload or peaking power applications.

This chapter starts by giving a background for solar energy and shows where Concentrating Solar Thermal (CST) and Heliopod technologies fit into this category. The project's problem statement, objectives, methodology, and scope are discussed before concluding the chapter with a brief overview of the rest of the document.

1.1 Solar energy

1.1.1 The sun as a source of energy

The sun is critical for life on earth as it is the primary source of energy. The sun provides the earth with more than 99.9 percent of its energy [22]. Figure 1.1 shows

the Direct Normal Irradiation (DNI) world map. DNI is the flux density of solar energy received per unit area of a flat surface held perpendicular to the sun's rays [23]. The map shows that South Africa, particularly the Northern Cape, is a prime spot for solar energy harvesting due to DNI values exceeding 3000 kWh/(m²a)).




Figure 1.1: DNI world map [1]. It can be seen that DNI levels in the Northern Cape exceed 3000 kWh/(m²a).

1.1.2 Harvesting energy from the sun

Two common methods to harvest the energy from the sun are PV panels and concentrating solar systems. PV panels use semiconductors to convert energy from light photons directly into electricity. CST technology focuses direct sunlight to create thermal energy. This thermal energy can be used in a power cycle to produce electricity, known as Concentrating Solar Power (CSP). A CST system consists of the following parts as seen in Figure 1.2 [24]:

- **Concentrator:** An arrangement of optical components that focus the incident solar energy on to a smaller receiver area.
- **Receiver:** Component that absorbs the concentrated solar energy from the concentrator and converts it into thermal energy.
- **Collector:** Refers to the entire system which includes the receiver and the concentrator.

Figure 1.3 shows the classification of solar energy into solar thermal and PV systems. A distinction is made between concentrating and non-concentrating systems. Within the scope of CST systems, the leading technologies are listed as linear Fresnel, parabolic trough, Central Receiver Systems (CRS), and a parabolic dish.

Solar collectors are used when high temperatures are required. A solar collector's primary purpose is to focus solar energy falling on a large surface area to a




Figure 1.2: Solar radiation from the sun is focussed on a solar receiver by solar concentrators. The harvested thermal energy can be used in a cycle or stored as thermal energy. The entire system is referred to as the collector. Adapted from [2].

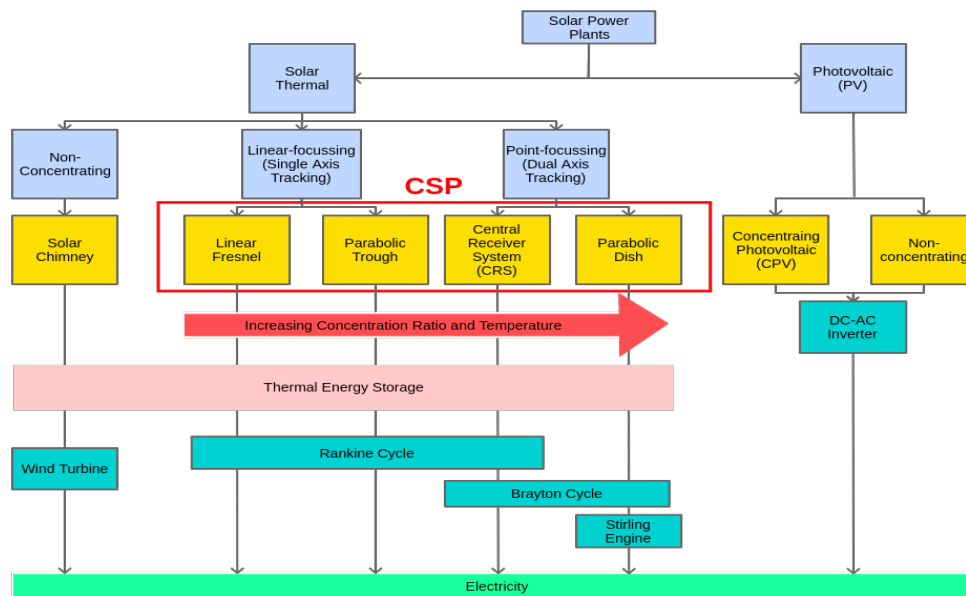


Figure 1.3: Classification of solar power plants (adapted from [3]).

smaller area before it is converted to heat. This limits heat loss due to convection and radiation. There are two common types of solar concentrators: line focus collectors, shown in Figure 1.4, and point focus collectors, shown in Figure 1.5. Line focus collectors such as parabolic trough and Fresnel collectors concentrate solar radiation onto a receiver tube running the length of the concentrator. Line focus collectors require single-axis tracking of the sun and can achieve temperatures around 400°C without significant heat loss.

Point focus receivers, such as a parabolic dish and CRS, focus all incident solar energy to a point. This allows point focus collectors to achieve higher temperatures, but they require a more complex dual-axis sun tracking system. A parabolic dish concentrates all incident sunlight onto an insulated cavity containing a heat transfer device. A CRS uses multiple dual-axis tracking mirrors, called heliostats, to reflect sunlight on to a central receiver located on top of a tower in the heliostat field. The receiver acts as a heat exchanger by absorbing the reflected solar energy and using

it to heat a heat transfer fluid, typically molten salt [25].




Figure 1.4: Line focus solar collectors: Parabolic trough (left) and linear Fresnel (right) [4].




Figure 1.5: Point focus solar collectors: Parabolic dish (left) and central receiver system (CRS) (right) [4].

1.1.3 Concentrating Solar Power

In 2018 around 5.5 GW of CSP was installed worldwide, which is 4.3 times more than in 2010. Although CSP installation has increased dramatically in the last decade, it is still far less than solar PV worldwide. To put this into perspective, the global solar PV installation passed 5.5 GW by 2005. Currently, PV also offers more cost-effective electricity generation than CSP. The weighted-average Levelised Cost of Energy (LCOE) for a CSP plant commissioned in 2018 was 26 % lower than in 2017, showing significant progress to deliver cost-competitive electricity. This was, however, still 117 % more expensive than PV [26].

CSP's main advantage is the relative ease of storing thermal energy cost-effectively in a suitable medium such as molten salt in large storage tanks for extended periods (seen in Figure 1.6). This allows a CSP plant to operate as a dispatchable power station and provide electricity well after sunset [27]. The capacity factor, which is a measure of how much energy is produced by a plant compared to its maximum possible output, of CSP CRSs commissioned in 2018 are as high as 63 % [26].

This advantage becomes increasingly more valuable as baseload coal facilities are phased out.



Figure 1.6: Thermal energy storage tank for molten salt CSP applications [5].

CSP has been proposed, among others, by the Renewable Energy Independent Power Producers Procurement Program (REI4P) as a source of renewable energy to assist the transition to a low carbon economy [28]. Due to South Africa's high DNI values, seven CSP plants with a combined capacity of 600 MW have been commissioned due to the REI4P. Included in the 600 MW capacity is Khi Solar One, a 50 MW plant located outside Upington in the Northern Cape. Khi Solar One, shown in Figure 1.7, was South Africa's first CSP CRS when it was commissioned in 2016. However, the 2018 South African draft Integrated Resource Plan (IRP) excluded any CSP projects from 2020 onwards [29].



Figure 1.7: An aerial view of South Africa's first CSP CRS, Khi Solar One, located in the Northern Cape [6].

1.1.4 HelioPod

According to the national energy balance of 2015, provided by the Department of Energy, 36 % of the total final consumption of energy in South Africa is consumed by the industry sector [30]. Furthermore, 62 % of the industry sector's energy consumption was due to process heat [31]. Process heat in the industry is used for operations such as desalination, evaporation, drying, melting, and pre-heating. Due to CSP's omission from the 2018 draft IRP, the market now exists to use CST technology to provide process heat and reduce carbon emissions.

The Solar Thermal Energy Research Group (STERG) launched a CST project called Helio100 (seen in Figure 1.8), which was completed in 2015, to introduce their HelioPod technology. A HelioPod, as shown in Figure 1.9, is a common mounting structure that supports six heliostats, each with an optical aperture of 1.83 m x 1.22 m.



Figure 1.8: Aerial view of the construction of the Helio100 facility located outside the town of Stellenbosch. The receiver on top of the tower, the HelioPods and the control container can be seen.

Each heliostat is equipped with a battery/PV combination to supply power for actuators and processing. The HelioPod is not fixed to the ground, thus expensive groundwork costs are eliminated, and it lowers the risk of an energy service company as it can use the CST plant to sell heat to a process rather than the process plant owning and operating its own CST plant [32]. The Helio100 project lays the research foundation to develop the technology to produce a fully autonomous CST plant to advance novel energy systems in the drying and pre-heating of furnace manganese ore as part of the H2020 PREMA project.



Figure 1.9: HelioPod structure as seen from behind (left) and from the front (right). A HelioPod consists of six heliostats mounted on a triangular frame that is not fixed to the ground.

1.2 Problem statement

The HelioPod is a novel technology that revolutionized heliostat fields. It aims to reduce capital costs as well as enable the rapid deployment of a CST plant. HelioPods can also be added or removed from an existing heliostat field with ease, giving it a modular nature. To fully utilize all the advantages that the HelioPod technology offers, a wireless communication system is required. Wireless communication will decrease costs by eliminating trenching for cabling and allow unconstrained placement of HelioPods within the heliostat field. New HelioPods can also seamlessly join the wireless network if needed. The wireless communication system should, however, have sufficient capacity, range, throughput, and latencies to enable a high level of control over each heliostat.

A heliostat field poses multiple challenges for wireless communication, such as interference caused by metal structures, dense network operation, high bandwidth, and increased range requirements. Other factors, such as small data packets, low power requirements, and reliability, should also be considered. As a dense, small, and modular heliostat field is still a novel idea, limited research is available to assess the feasibility of wireless communication in such a CST field. The lack of practical in-field data can result in uncertainties during the planning phase of a CST plant, which can result in costly alterations later in the project.

1.3 Objectives

This study's primary objective is to evaluate the feasibility of wireless communication within the context of a CST field by means of practical tests aimed to determine the influence of interference, dense networks, and range on the network performance. The secondary objectives are:

1. To investigate different wireless communication technologies, including Internet of Things (IoT) solutions, to propose an optimal technology for a next-generation plant.

2. To compare different network topologies and identify the most suitable for future projects.
3. To set up a wireless communication network consisting of off the shelf bought items to simulate a CST communication network.
4. To use in field tests to gather experimental data that can be used to design a next-generation plant with lower amounts of uncertainties.
5. To make recommendations for future work that can be carried out under this topic of study.

1.4 Methodology

Stemming from the objectives, the outline of the overall research methodology is summarised as:

1. Investigate relevant literature to understand previous work done in the field.
2. Examine the communication network of the current field.
3. Design the network architecture for a network test.
4. Plan and perform in field tests to analyse the network performance in different environments.
5. Evaluate the test results to determine if wireless communication is feasible in a heliostat field.

1.5 Scope

This project represents the first study in heliostat communication networks in STERG. The project focuses on network design and performance measurement by using practical in field tests. In this thesis, fifty Raspberry Pi Zero WH microcontrollers are used to simulate nodes in a CST network to determine the influence of uncertainties from the surrounding environment. Included in the scope is the programming of all network nodes using ZeroMQ message libraries. The data capturing and data processing procedures are also included along with the analysis of different modern communication technologies such as IoT solutions and different Wi-Fi standards. It excludes all work associated with the heliostat's control and aiming and is only concerned with determining if wireless communication is feasible within a CST field.

1.6 Document outline

- Chapter 2 contains a literature study of existing wireless communication technologies and standards, wireless communication deployment in CST plants, and the possible hindrances to wireless communication performance in CST plants. This is based on a previous study conducted by the author [33]. ZeroMQ middleware is introduced, and the chapter concludes with an overview of network performance measurement.
- Chapter 3 derives the system requirements for the next generation heliostat field. The size requirements of the field is given, and the required control interval is determined. The chapter concludes with a description of a typical message structure.
- Chapter 4 describes the chosen wireless communication standard and provides an overview of all hardware used in the project.
- Chapter 5 develops the software setup of the project. The chapter describes the ZeroMQ packet structure as well as giving an overview of the network architecture and processes. The chapter concludes with an explanation of the time offset procedure.
- Chapter 6 gives an overview of the experimental setup for the project. The field test scenarios are described, and the test parameters are derived.
- Chapter 7 presents and discusses the initial results and findings as well as the field test results. Statistical analysis is also presented to validate the results.
- Chapter 8 concludes the thesis and provides recommendations for future work.

Chapter 2

Literature Study

This chapter reviews work from literature relevant to communication in heliostat fields. A description of various wireless communication standards is given, including IoT (Internet of Things) technologies such as LoRa, NB-IoT, ZigBee, BLE, and various Wi-Fi standards. Previous wireless communication implementation in solar heliostat fields is researched, and possible performance impediments for wireless data transmission in CST plants are discussed. The chapter concludes with a brief overview of ZeroMQ and methods of wireless communication performance measurements.

2.1 Wireless communication technologies

There are currently a variety of wireless communication standards and technologies, including Low Powered Wide Area Network (LPWAN) technologies, IoT, and IEEE 802.11 (Wi-Fi) variants. In this section, the following technologies are examined: LoRa and LoRaWAN, NB-IoT, ZigBee, BLE, and Wi-Fi and its variants. This section is based on previous research done by the author [33].

2.1.1 LoRa and LoRaWAN

Long Range (LoRa), a technology developed by Semtech, is a proprietary spread spectrum modulation scheme based on Chirp Spread Spectrum (CSS) modulation. The CSS modulation spreads a narrow-band signal over a wider channel bandwidth resulting in low noise levels and high interference resilience. The efficient modulation scheme allows transmission distances of up to 15 km while maintaining a low power usage, making the technology increasingly popular for LPWANs [34].

To adjust the trade-off between range and data rate, LoRa implements the use of six spreading factors (SF7 to SF12). Lower spreading factors result in higher data rates at the expense of transmission range and vice versa. Depending on the chosen spreading factor and bandwidth, LoRa can have a data rate of up to 50 kbps, with a maximum payload length of 243 bytes per message [7].

LoRa Wide Area Network (LoRaWAN) includes the Medium Access Control (MAC)-layer and network layer protocol built on top of the LoRa modulation scheme. By enabling all base stations in the transmission range of a LoRa end device to receive a message, LoRaWAN reduces the number of unsuccessful messages. This also eases handover, especially if a node is mobile. LoRaWAN classifies end devices into three different classes depending on the requirements of the device [35].

- *Class A devices:* Class A devices are the most energy-efficient end device. This device only leaves a low powered state to transmit data at a scheduled time or due to an event-driven process. Communication from the server to the end device can only happen in two short receive windows following the transmission from the end device.
- *Class B devices:* Class B devices open additional receive windows at scheduled times. They do this by receiving a time-synchronized beacon from the base station. The network server can then use these receive windows to transmit data to the end device at scheduled times.
- *Class C devices:* Class C devices are the least power-efficient end devices. They have continuously open receive windows and can thus receive data at any time, except when transmitting.

LoRa uses the unlicensed Industrial, Scientific and Medical (ISM) spectrum for communication, i.e., 433 MHz in Asia, 868 MHz in Europe, and 915 MHz in North America [7]. LoRa incorporates the 915 MHz ISM band in South Africa as per the Independent Communications Authority of South Africa's resource allocation [36]. These sub-1 GHz frequency bands are free of licensing fees and are also commonly less congested than the popular 2.4 GHz and 5 GHz frequency bands used by multiple technologies. The ISM bands, however, limit devices to a 1% time-on-air duty cycle and to a maximum transmission power of 14 dBm to accommodate many devices [37]. When accounting for network constraints, such as the 1% duty cycle and a 13-byte preamble, simulations from [37] have shown an effective throughput data rate over a 24-hour period of as low as 0.544 bps. Thus, it is clear that the maximum duty cycle is a crucial constraint when operating in the unlicensed ISM band.

2.1.2 NB-IoT

Narrow Band Internet of Things (NB-IoT), also known as LTE Cat NB1, is a communication technology designed by the Third Generation Partnership Project (3GPP). NB-IoT is designed to meet the requirements for low-cost massive machine-type communication with a large number of devices in a cell. NB-IoT stresses the importance of low module cost, low power consumption and a sizeable coverage improvement over GPRS (General Packet Radio Service) [38].

NB-IoT coexists with GSM (Global System for Mobile communications) and LTE (Long Term Evolution) in licensed frequency bands, i.e., 700 MHz, 800 MHz, and 900 MHz. NB-IoT is supported in a bandwidth of 180 kHz, allowing operation

in a single resource block in LTE and GSM transmission [7]. Network operators have three different NB-IoT deployment options, as seen in Figure 2.1:

- *In-band operation:* The 180 kHz NB-IoT spectrum is placed inside LTE narrow bands that are not used for LTE transmission.
- *Guard-band operation:* The NB-IoT spectrum is placed in the unused resource blocks in an LTE carrier's guard band that is designed to limit interference.
- *Stand-alone operation:* NB-IoT is deployed using any available GSM spectrum. The 200 kHz GSM channel bandwidth allows for a 10 kHz guard buffer on both sides of the neighboring GSM channels.




Figure 2.1: NB-IoT modes for operation [7]. NB-IoT can operate inside the LTE narrow bands that not used for transmission, inside the guard bands, or deployed within unused GSM spectrum.

NB-IoT uses Single Carrier-Frequency Division Multiple Access (SC-FDMA) modulation for uplink communication and Orthogonal Frequency-Division Multiplexing (OFDM) modulation for downlink communication, which allows a single base station to serve 100-200k devices [39]. As NB-IoT uses licensed frequency bands, no duty cycle limitations apply, which improves the Quality of Service (QoS). There are also no restrictions on output power, with devices typically transmitting at 23 dBm [37]. The maximum downlink and uplink data rates are limited to 200 kbps and 180 kbps respectively, and the maximum payload size per message is 1600 bytes [7; 39]. Furthermore, an NB-IoT end device can achieve a battery lifetime of 10 years when transmitting 200 bytes per day on average, making it an ideal technology for a LPWAN.

The major disadvantage is that NB-IoT is not deployable where existing cellular networks are absent. This creates a challenge as most CST plants in South Africa are situated in the Northern Cape, a sparsely populated region with limited cellular network coverage. Without existing infrastructure, the network and deployment cost can be as high as \$15 000/base station and spectrum costs over \$500 000/MHz [40].

2.1.3 ZigBee

ZigBee is a communication technology based on the IEEE 802.15.4 standard. It is one of the most popular technologies available for a Wireless Sensor Network (WSN) due to its low cost and low power usage [41]. ZigBee incorporates the IEEE 802.15.4 standard as its physical (PHY) and MAC layers and improves the interoperability from the PHY layer to the network, application, and security services [42]. ZigBee uses the 868 MHz (Europe), 915 MHz (North America), and 2.4 GHz (worldwide) ISM frequency spectrum. ZigBee uses Binary Phase Shift Keying (BPSK) modulation in the 868 MHz and 915 MHz frequency ranges and Offset Quadrature Phase Shift Keying (OQPSK) modulation in the 2.4 GHz spectrum. ZigBee can achieve an over-the-air data rate of 20 kbps at 868 MHz (1 channel), 40 kbps at 915 kbps (10 channels), and 250 kbps at 2.4 GHz (16 channels) [41].

The nodes within a ZigBee network are categorized into three categories based on the required tasks they perform [43]:

1. *Coordinator*: A full functioning device that is responsible for setting up the network with a given Personal Area Network (PAN) identifier to which other ZigBee devices can join. The coordinator is also responsible for channel selection, issuing network IDs, and storing network information such as security keys. There is only one coordinator in a ZigBee network.
2. *Router*: A full functioning device that can act as a coordinator if the original coordinator node fails. The router is responsible for the relay of packages within the network. A router can accept connections from nodes in an existing ZigBee network and relay them to extend the network.
3. *End device*: A reduced function device that is primarily used to collect information from sensors. An end device has the least processing capabilities and memory of all ZigBee devices and is thus the most power-efficient. An end device can only communicate to the device that controls it, i.e., a router or coordinator, and cannot relay a message to another node.

ZigBee accommodates both tree and mesh routing. Tree routing is a memory-less routing procedure as routing paths can be inferred directly from the network addresses. Tree routing is computationally inexpensive and fast but can have long, unoptimized routing paths. Another drawback is that if one node in a tree routing scheme fails, no message can be delivered to any nodes downstream of the failed node. ZigBee also incorporates mesh routing, which uses an Ad-hoc On-demand Distance Vector (AODV) algorithm to find the shortest routing path. Mesh routing results in more routing overheads but allows for a robust network design. If a node in the network fails, the AODV algorithm will reroute the messages around the failed node, creating a self-healing network [42].

As ZigBee has a limited transmission range of 10-20 m, multiple hops are required to create a network covering larger areas [35]. Up to 65 000 nodes can exist

in a single ZigBee network, allowing a scalable implementation of the technology [35]. A Cyclic Redundancy Check (CRC) is also implemented to limit packet errors in large networks [44]. However, the true data throughput will be significantly less than the over-the-air maximum data rate due to package overheads and processing delays introduced due to a multiple hop network topology.

2.1.4 Bluetooth low energy

Bluetooth Low Energy (BLE) was introduced when Bluetooth version 4 was announced in 2010 by the Bluetooth special interest group [45]. BLE was designed to meet the requirements of low duty cycle, low power, and low-cost devices. BLE uses the Frequency Hopping Spread Spectrum (FHSS) technology to limit interference and fading during communication [45]. BLE operates in the 2.4 GHz ISM frequency band and uses 40 channels with center frequencies 2402 MHz - 2480 MHz, each channel separated by 2 MHz. The 40 channels are broken into 37 data channels and 3 advertising channels [45]. BLE uses Gaussian Frequency Shift Keying (GFSK) modulation, which allows an over the air data rate of 1 Mbps and throughput of around 0.27 Mbps [45; 35]. Typically, a BLE device will have a transmit power of between -20 dbm and 10 dbm [45], which leads to a power consumption of 1 mW - 10 mW while transmitting [35].

BLE has a maximum range of 50 m; thus, multi hopping is needed to relay a message over larger distances [45]. When Bluetooth version 4.0 was introduced in 2010, the basic topology of the network was described as a piconet. A piconet resembles a star network topology consisting of a single master and a set of slaves. Devices were not allowed to operate in different piconets at the same time [8]. Using a Time Division Multiple Access (TDMA) like approach, the master transmits to each of its slaves in the piconet. The slave can reply depending on the original message's context; thus, all access in the piconet is controlled by the master [45]. The introduction of Bluetooth version 4.1, introduced in 2013, allowed devices to run multiple link-layer instances simultaneously, allowing a single node to bridge two piconets. The device can have a slave/master or a slave/slave combination within two respective piconets. This allowed the formation of Bluetooth scatternets, which are piconets bridged by common nodes [46].

Bluetooth version 4.2 added internet connectivity and enhanced security measures coupled with higher throughput. No changes were made with regard to the network roles of different nodes. Lastly, Bluetooth version 5.0 and 5.1 added, along with increased range and throughput, added advertising message capacity [46]. This allowed devices to still use the three primary advertising channels as specified in version 4.0, but also added the use of the remaining 37 channels as secondary advertising channels [46]. Although all modern versions of Bluetooth allowed the formation of scatternets, there was no official standard that existed for mesh network routing [8].

The BLE mesh profile was released in 2017, which enables the operation of connectionless Bluetooth multi-hop networks [8]. Each device in a BLE mesh can

communicate with every other device providing that the topology allows for it. Each device associated with a particular BLE mesh network is called a node. All nodes within the mesh network are able to send and receive messages, but nodes can be configured to allow for additional features. The BLE Mesh Profile defines four different features [8]:

- *Relay*: Nodes with relay capabilities enabled are able to forward received messages, thus allowing for a multi-hop network.
- *Proxy*: Nodes with proxy enabled are able to communicate with BLE devices that do not support BLE Mesh Profile by providing a Genetic ATtribute (GATT) profile interface.
- *Low power*: Nodes with low power mode activated enter a low duty cycle state, conserving energy.
- *Friend*: Nodes with friend mode activated work hand in hand with low power nodes. A friend node receives and stores messages destined to a low power node. The messages are then forwarded to the low power node when it becomes active.

During basic mesh operation, the BLE nodes advertise and scan the network and, thus, are not connected in a conventional Bluetooth sense. Bluetooth mesh networks do not take advantage of the FHSS and TDMA MAC scheme used in a conventional piconet topology [8]. A typical BLE mesh network topology is depicted in Figure 2.2.

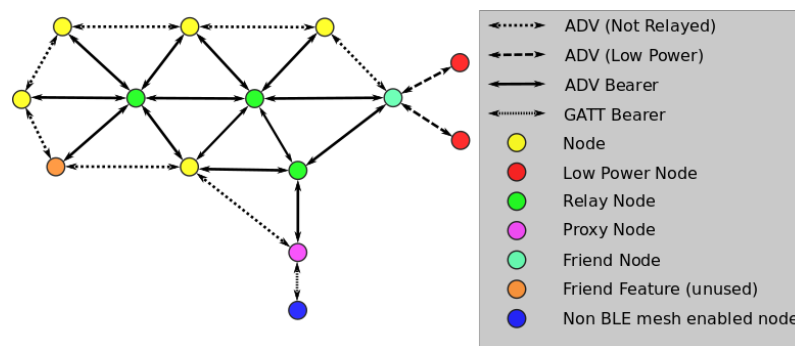


Figure 2.2: A typical BLE mesh network adapted from [8]. Various node types are presented in their typical arrangement.

A flooding technique is used to send messages in a BLE mesh network. The message is sent to all nodes within the sending node's radio range, and nodes with relay features forward the received message to all nodes in their vicinity. To limit overhead, BLE mesh network implements a heartbeat strategy, where each node periodically indicates its presence. The heartbeat can be used by the receiver to indicate the number of hops to the heartbeat sender. Furthermore, each message

contains a Time-To-Live (TTL) counter, which defines the maximum number of hops a message can pass. By using the combination of the heartbeat and the TTL counter, the network overhead can be significantly reduced [8].

2.1.5 Wi-Fi

In 1997 the first 802.11 standard was ratified by the Institute of Electrical and Electronics Engineers (IEEE) to create the world's first Wireless Local Area Network (WLAN). Since then, many IEEE 802.11 standard evolutions have increased transmission rates, coverage area, and QoS. Today the IEEE 802.11 standards, commonly known as Wi-Fi, have become synonymous with WLANs, with more than 11 billion Wi-Fi devices already expected to be in use by the end of 2017 [47].

2.1.5.1 IEEE 802.11n

In 2009 the IEEE 802.11n standard, also referred to as Wi-Fi 4, was launched, which used the 2.4 GHz and 5 GHz ISM frequency bands for communication. The 802.11n standard introduced new PHY and MAC features that resulted in a significant increase in throughput. The 802.11n standard adopted OFDM, which divided each 20 MHz channel into 52 subcarriers. This reduced interference and improves efficiency to allow a peak data rate per stream of 65 Mbps when using an 800 ns Guard Interval (GI) and 72.2 Mbps when using a 400 ns GI [48]. 802.11n was also the first Wi-Fi standard to include channel bonding. This allowed two adjacent 20 MHz channels to be bonded to produce a 40 MHz channel, which boasted 108 subcarriers that resulted in even higher data rates compared to legacy 802.11 standards at the time. Wi-Fi 4 uses 64-Quadrature Amplitude Modulation (64-QAM) scheme, which contributes to a theoretical maximum data rate of 600 Mbps when operating with 40 MHz channels and 400 ns GI.

Multiple-Input Multiple-Output (MIMO) technology was introduced with the 802.11n standard, which uses spatial multiplexing to transmit parallel data streams in the same frequency channel. For MIMO to work, the transmitter and receiver must each have multiple radio/antenna chains. Each spatial stream is transmitted from a different radio/antenna chain on the transmitter and received by the identical radio/antenna chains on the receiver. By using the known phase offsets of its antennas, the receiver can reconstruct the original data streams. MIMO increased data rates but is constrained by the number of spatial streams supported by the transmit/receive pair, i.e., if an Access Point (AP) supports sending Radio Frequencies (RF) on one stream and receiving on another (1x1), it will be able to do so with a client that supports 1x1 connectivity. If the AP supports 3x3 connectivity and the client 1x1, the performance will be limited by the client device [49].

2.1.5.2 IEEE 802.11ac

The IEEE 802.11ac Wave2 standard, commonly known as Wi-Fi 5, was released in 2013, offering theoretical data rates over 1 Gbps. The 802.11ac standard operates

solely in the 5 GHz frequency range and expands channel bonding, as seen with Wi-Fi 4, to include 20,40,80 and 160 MHz channels. 802.11ac introduced Downlink (DL) Multi-User MIMO (MU-MIMO) with eight spatial streams. MU-MIMO allows simultaneous downlink transmission of separate data streams to multiple clients in the same channel [48].

Wi-Fi 5 also incorporates OFDM but introduces a higher-order modulation scheme called 256-QAM. 256-QAM is 33 % more efficient than the 64 QAM offered in 802.11n but can only be used in environments with a high Signal to Noise Ratio (SNR). The 256-QAM increases the complexity by another order of magnitude and can now represent 8 bits with each constellation point. However, it is necessary that the transmitter and receiver are designed to ensure that the transmit and receive error vector magnitude can support the higher constellation. 802.11ac will thus only utilize 256-QAM in low interference scenarios and short ranges (typically 10-20 m). Over longer distances, the radio will switch to narrower channels and a lower modulation scheme [48; 50].

To improve the SNR 802.11ac incorporates a feature called beamforming. Most antennas, especially before the introduction of 802.11ac, were omnidirectional antennas, meaning that the transmitted signal from the antenna is spread out in all directions. This keeps the wireless medium busy in all directions for each transmission. Beamforming includes the use of multiple directional antennas on an AP to send numerous signals to clients. This can be used to determine the optimal transmission path to create an intense directional RF beam, limiting interference in other directions. Beamforming allows higher throughput, extended range, and more spatial reuse. Although the 802.11ac standard promised data rates up to 1.3 Gbps, tests by [51] showed a throughput of 207.5 Mbps with one spatial stream, 514.9 Mbps with two spatial streams, and 572.3 Mbps with three spatial streams. The 207.5 Mbps result represents a practical scenario as most client devices only support one spatial stream and will thus limit the throughput.

2.1.5.3 IEEE 802.11ah

The 802.11ah standard, also known as Wi-Fi HaLow, was released in 2016 and promised to bridge the gap between LPWANs and WLANs. 802.11ah offers significantly higher throughput compared to LPWAN devices with lower power consumption than traditional WLAN devices. The 802.11ah standard thus targeted the specific needs of higher throughput IoT devices as well as Machine-to-Machine (M2M) communication. Wi-Fi HaLow operates in the unlicensed sub 1 GHz frequency bands (863-868 MHz in Europe, 755-787 MHz in China, and 902-928 MHz in North America) and offers channel bandwidths of 1-16MHz, with 1 and 2 MHz support being mandatory, allowing transmission ranges up to 1 km while maintaining data rates of 150 kbps [52].

Hierarchical association identification technology allows up to 8192 devices to connect to a single AP. To help reduce interference and collisions, the 802.11ah standard introduces MAC layer functions such as Restricted Access Window (RAW) and Traffic Identification Map (TIM). The RAW mechanism is used to restrict the

number of devices that can access the channel simultaneously by grouping devices together and only allowing access to a single group per RAW interval. Each RAW interval is preceded by a beacon containing a RAW Parameters Set (RPS) information element that identifies all devices in the group and the RAW interval start time [53]. RAW intervals can also be divided into slots that can be allocated to designated devices on the network. Devices with the TIM mechanism activated can receive the transmitted beacon containing the RPS. If the device is listed in the group of the RPS, the device will be allowed both uplink and downlink access to the medium within its designated RAW slot. If the device is not listed in the RPS group, it will enter a low power mode and wake for the next beacon to be transmitted.

The 802.11ah standard also allows for non-TIM enabled devices which act as power-saving devices. These devices do not wake periodically to receive the beacon transmission from the AP but transmit one power saving poll or trigger a frame transmission to the associated AP every listen interval [52]. These devices typically also incorporate a feature called Target Wake Time (TWT) that allows devices to remain in a low-powered state and only wake to transmit a small amount of data allowing extended battery lifetime in IoT devices.

Although the 802.11ah standard was introduced in 2016, the availability of APs and clients remain limited. This could be due to the interoperability challenges between the 900 MHz spectrum and the existing Wi-Fi infrastructure, the absence of a global 900 MHz standard, or the high costs associated with the development of first-generation 802.11ah supporting platforms.

2.1.5.4 IEEE 802.11ax

The IEEE 802.11n and 802.11ac standards have been widely implemented in Wi-Fi applications worldwide. Partly due to the success of these standards, two common problems have arisen in the use of Wi-Fi networks:

1. The development of IoT and smart devices, coupled with the upsurge of users, has caused an increase in the number of nodes on each network, which leads to a *dense network* scenario. A dense network causes higher interference as there are more devices in the RF range contesting for the medium. This leads to the deterioration of the performance of the network.
2. Users require *higher throughput* from their network for applications such as high definition streaming, Voice over Internet Protocol (VoIP), and cloud storage.

Wi-Fi WLANs use a MAC protocol called Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA). When a node has a packet to send, it listens to determine if a channel is detected as free, i.e., the channel's detected energy level is lower than the Clear Channel Assessment (CCA) threshold. The node then selects a random initial value to start a backoff procedure. The backoff counter, or NAV counter, is reduced if the channel is free and paused when a transmission, either

from within the WLAN or a neighboring WLAN, is detected on the channel. The node starts transmitting when the backoff counter reaches zero [50]. In a dense network, there are many devices on the WLAN that are contesting to transmit, resulting in long waiting periods per device as the backoff timer is often paused while another node in the network transmits. In scenarios where multiple APs are placed in close proximity, such as in a stadium or office, interference from the neighboring WLANs further reduce the number of concurrent transmissions and decrease network performance. This is a common problem with the previous 802.11 standards.

Figure 2.3a shows three neighboring WLANs with all APs within each other's carrier sense range. Figure 2.3b shows their respective operating channels, and it can be seen that the channels of WLAN 3 overlap with both WLAN 1 and WLAN 2. A transmission by WLAN 1 will thus cause a pause in the backoff procedure within WLAN 3, with the same being true between WLAN 2 and WLAN 3. In Figure 2.3c, it can be seen that even though WLAN 3 has the widest channel, it achieves the lowest throughput. This is due to the interference from its neighboring APs. The throughput of WLAN 1 and WLAN 2 is significantly higher as they operate independently of one another and is only influenced by WLAN 3.

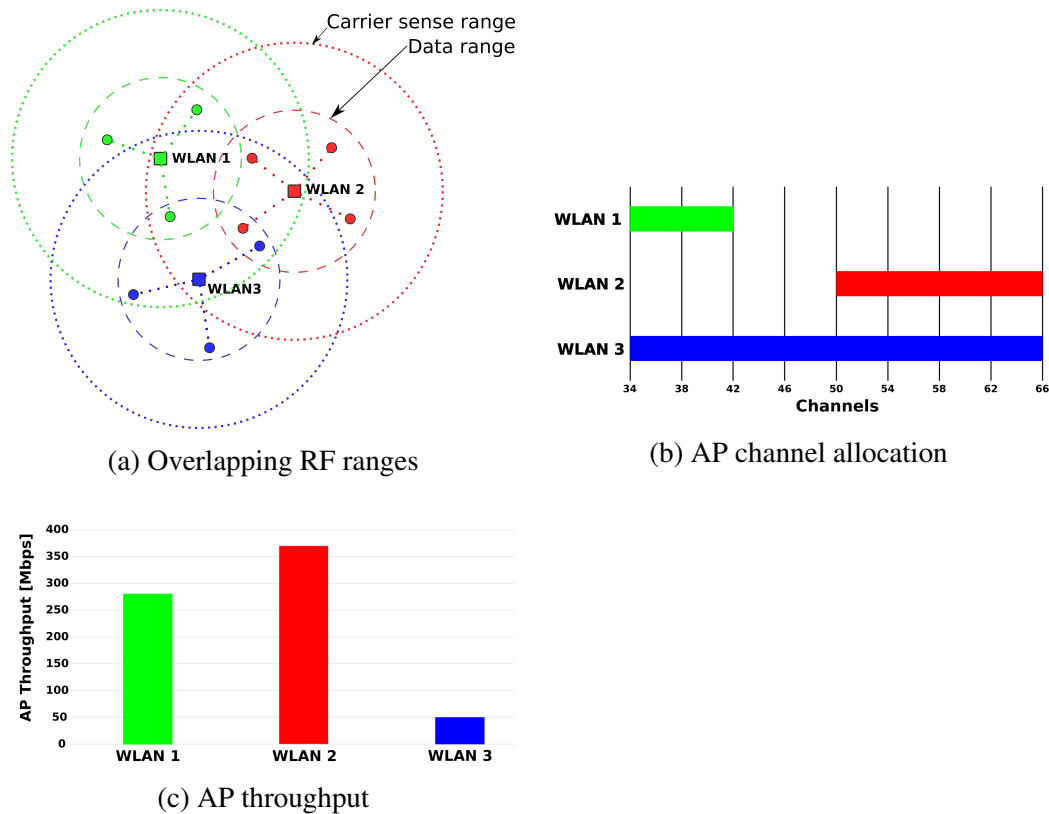


Figure 2.3: CSMA/CA in a dense network scenario, adapted from [8]. It can be seen that even though WLAN 3 operates in the widest channel, interference causes a deteriorating network performance.

The IEEE 802.11ax standard, also referred to as Wi-Fi 6, is due to be finalized in Q4 2020 [54], and is designed to meet the unique requirements of high-density networks and deliver a fourfold increase in average throughput per user [55]. In dense network deployments, the CCA threshold can be increased to require a higher energy level on the channel to qualify it as busy. This will reduce the pausing of the backoff timer but render the sent message susceptible to interference when two clients from neighboring stations transmit simultaneously [50]. The transmission power can also be lowered to encourage more spatial reuse and minimize the overlapping of the APs' carrier sense range. The 802.11ax standard dynamically adapts the CCA and transmission power to reduce the interference on the network and optimize performance. A Basic Service Set (BSS) color mechanism is also introduced in 802.11ax, which enables a node to determine whether the received frame originated from within its own BSS or from a neighboring BSS. Each node in the network is required to have two NAV counters, one for intra-BSS messages and one for inter-BSS messages, which further optimizes transmissions in dense deployments [56].

802.11ax introduces Orthogonal Frequency Division Multiple Access (OFDMA), which acts as a multi-user version of OFDM used in legacy 802.11 standards. In OFDM, each 20 MHz channel is subdivided into 64 subcarriers, each with a length of 312.5 kHz. This equates to 52 data, four pilot and eight unused subcarriers. The 52 data subcarriers can only be used to transmit data to a single client using OFDM. Each client thus uses the entire 20 MHz channel to transmit or receive data.

OFDMA reduces the subcarrier size from 312.5 kHz to 78.125 kHz by increasing the symbol time from 3.2 μ s to 12.8 μ s. This increases the number of subcarriers to 256 [56]. These subcarriers are grouped into smaller sub-channels referred to as Resource Units (RUs), with each RU having 26, 52, 106, or 242 data subcarriers. Each RU can be used to transmit data to a different user allowing up to 9 parallel simultaneous transmissions of small frames to multiple nodes. This reduces the contention time on the network and increases network performance.

802.11ax operates in both the 2.4 GHz and the 5 GHz ISM band and allows channel bandwidths of 20, 40, 80, and 160 MHz. MU-MIMO is expanded from 802.11ac to enable both uplink MU-MIMO (UL MU-MIMO) and DL MU-MIMO with eight spatial streams, allowing more devices to receive simultaneous data. OFDMA can also be used to further split each MU-MIMO stream allowing even more simultaneous transmission [55]. This is especially advantageous for high-density IoT development to enable more frequent small frame transmissions. Furthermore, 802.11ax also includes the TWT feature that was introduced by 802.11ah to extend the battery life of IoT devices. The inclusion of beamforming allows directional transmissions, which increases range and spatial reuse resulting in advancements in performance in dense deployments [50].

2.2 Wireless communication development in CSP plants

Most heliostats in a commercial CST plant are controlled and powered by dedicated wired links and bus systems [9]. Atabaev et al. [57] describes such a wired modern automated heliostat communication network based on the RS485 interface. The RS485 interface propagates signals via 5FTP or UTP cables and is capable of transfer distances of 1200 m with a data rate of 2 Mbps. Atabaev et al. utilizes a bus topology with a central computer master node and 62 heliostat slave nodes. The RS485 network is used to transfer essential parameters from each slave node to the central server upon request. By using the RS485, Atabaev et al. explains that an expansion of the field will require only two cables, a data and a power supply cable.

Having a wired network in a CST plant has the advantages of high data rates and reduced interference, but it is at a high expense. Some studies have estimated field-wiring costs of around \$400 per 20 m² heliostat [58]. The main contributing factors to these high costs are the expense of the cables themselves, trenching, protection from rodents, and the installation of lightning arrestors [9]. Wireless communication combined with self-powered heliostats offers a significant reduction in these costs. An investigation launched by the National Renewable Energy Laboratory (NREL) claimed cost savings potential of over 20 % when installing wireless communication in heliostat fields [59]. Additionally, the plant manufacturer also gains the flexibility to place and install heliostats in the field without the need for prepared dedicated wiring [60]. However, wireless communication does present challenges of error-prone channels due to noise, interferences, and shadowing effects [9].

The first autonomous and self-powered wireless control heliostats were developed by Gacia et al. [61] as part of the PCHA project. The field included 92 fully autonomous heliostats capable of tracking the sun via solar position calculations accurate to 0.5 arcmin. Wireless communication was enabled by a 430 MHz radio modem, which allowed data rates of 9.6 kbps. As each heliostat was equipped with solar tracking algorithms, the wireless communication medium was only used to poll the heliostat to confirm communication and change the heliostat's state between tracking, stow, and emergency states. By adding the intelligence of solar position calculations capable of independently tracking the sun to each node, it removed the load off the wireless medium and allowed the central server to poll up to 20 heliostats per second. The disadvantage was that outfitting each node with a higher level of intelligence is costly. Further development in wireless communication for a heliostat field was needed to fully utilize the available cost-saving potential.

Perhaps the most valuable contribution to the development of wireless self-powered heliostats was made by Kubish et al. [9] as part of the HELIOMESH project. During this project, 93 nodes were placed within the heliostat field of the German Aerospace Center (DLR) in Jülich, Germany. All nodes contained an IEEE 802.15.4-2006 compliant transceiver used within a Wireless Mesh Network (WMN) for communication. The majority of the nodes were placed to ensure a distribution

over the entire heliostat field, with the most distant node five intermediate transmissions away from the gateway, testing the transmission distances and latencies due to multiple hop communications. A group of nodes in the middle of the heliostat field was also placed in close proximity to test the interferences due to a high-density network. The exact placements of the nodes can be seen in Figure 2.4.




Figure 2.4: Node placement of the HELIOMESH project by Kubish et al. [9]. Nodes are placed to tests dense networks and long, multi-hop, transmissions.

The IEEE 802.15.4-2006 standard utilizes the 2.4 GHz spectrum with 16 non-overlapping channels to limit interference and enabling concurrent communication when spatially separated gateways are used. The field tests reported average round trip latencies of 200 ms with a packet error rate of between 1-20 %. Little information is shared to determine the performance of the individual nodes. However, it is assumed that the higher packet error rate can result from the CSMA scheme that was applied and the challenges with this approach in dense networks, as stated in Section 2.1.5.4. In an emergency situation, a broadcast message can be transmitted in 50 ms and covers 98 % of all nodes. The limit in reliability is due to the absence of acknowledgments sent by the nodes, but this can be mitigated by transmitting multiple broadcast messages to the entire field. This research was expanded on by Unterschütz [60] that investigated the use of the IEEE 802.11 standard as a backbone for the WMN in a full-scale deployment.

A variety of research delves into the use of WSN within a heliostat field. Freeman et al. [62] propose a network comprising of Memsic MizaZ ZigBee compliant nodes operating in the 2.4 MHz frequency band. The network structure requires a single heliostat head node responsible for the solar position calculations using a high precision NREL algorithm. These angles are then transmitted via the ZigBee network, including multiple transmission hops if the distance exceeds the range of a single node, to heliostat control nodes capable of controlling up to 7 individual heliostats each by using a multiplexer. The computationally expensive NREL algorithm is confined to only the heliostat head node. The resulting solar azimuth angle, zenith angle, and time of day are sent over the network as a string. This allows

small data packets of 20-30 bytes. Freeman et al. [62] completed field tests with a heliostat head node and a one heliostat control node operating a single-axis tracking reflector. This study thus does not investigate the influence of dense network operation, interference due to heliostat steel structures, latencies introduced due to multiple hop transmissions, and capacity constraints of the network. Further WSN research has been conducted by Younis et al. [63] and Mabusela et al. [64]. However, these studies also lack the practical data needed to conclude whether a WSN will be feasible in a full-scale heliostat field.

2.3 Hindrances of wireless communication in a CST field

Wireless communication faces many challenges in a CST field that can hinder the performance of the network, such as interference from other devices in a dense network, obstacles, and signals that reflect, scatter, attenuate and diffract [65]. This section will describe the factors that can reduce the performance of wireless communication in a CST field.

The first factor to take into account is *range*. Many CSP plants have heliostat fields that stretch over 1000 m, which will result in long-range communication requirements. Wireless communication is achieved by propagating electromagnetic waves from a transmitter to a receiver through a communication medium such as air. As the electromagnetic wave moves through the air, it is attenuated until the point where the signal is too weak for a receiver to interpret the signal accurately. When the Received Signal Strength (RSS) is too weak, a receiver will likely encounter bit errors, which will cause the CRC to fail and the receiver to drop the message. The range of a wireless signal in free space can be estimated using the Friss free-space equation,

$$P_R = P_T G_T G_R \left(\frac{\lambda}{4\pi}\right)^2 \left(\frac{1}{d}\right)^n, \quad (2.1)$$

where P_R is the power available from the receiving antenna, P_T is the power supplied to the transmitting antenna. G_T and G_R are the gains of the transmitting and receiving antennas, respectively. λ is the wavelength of the electromagnetic wave in the transmission medium, d is the distance between the transmitter and receiver, and n is a path loss exponent that is determined experimentally ($n = 2$ in free space) [65]. Figure 2.5 shows a decline in the receiving antenna power with an increase in transmission range for a typical Wi-Fi application operating in the 2.4 GHz ISM band. In real-world applications, the transmission range will be less than the free space model due to other obstacles and interferences.

The second factor to consider is *Radio Frequency Interference (RFI)*. RFI is commonly found in dense networks with multiple nodes and APs on the network, all operating in the same frequency. As the ISM bands are free, it is popular with multiple wireless technologies. For example, Wi-Fi, Bluetooth, Zigbee, Walkie-




Figure 2.5: Friss free space signal attenuation model. The Friss free space equation describes a lower received power with an increase in transmission distance [10].

talkies, and even microwave ovens operate in the popular 2.4 GHz frequency band. When two devices transmit simultaneously, collisions can occur, leading to dropped packets and lower network performance. Modern standards include collision avoidance technology such as CSMA/CA discussed in Section 2.1.5.4. Frequency hopping algorithms can also be implemented to transmit in the channel with the lowest noise. Furthermore, transmitters can adjust their transmit power to limit their carrier sense range to encourage spatial reuse [66].

The third factor is *Electromagnetic Interference (EMI)*. EMI is unconstrained noise that is propagated over a wide range of frequencies [67]. Devices such as a DC motor have slip rings and brushes that "arc and spark," which causes high levels of EMI that can influence the performance of the wireless network. Most heliostats are moved by means of an actuator using a DC motor, which can cause EMI. Other devices, such as variable speed drives, use DC pulses to induce motor action that can also interfere with an electromagnetic wave. The best practice is to place devices that cause EMI in a Faraday cage that limits the effects.

The fourth factor to consider is *shadowing effects*. Shadowing effects cause the power of the received electromagnetic signal to deviate from the average value due to obstacles in the path of the propagated wave. In a CST field, these obstacles take the form of metal heliostat structures. As RF waves are electromagnetic, they are captured by a conductive material such as metal. For comparison, a wireless signal propagating through a wooden door will have a 3 dBm loss versus a 6 dBm loss when the same wireless signal is propagated through a metal door [68]. The metal heliostat structures will thus shadow the receiver from the transmitter, which could degrade the performance of the wireless network.

The fifth factor that can hamper performance is *multipath effects*. Electromagnetic waves can be reflected, diffracted, or even scattered from surfaces between the transmitter and receiver. The receiver antenna may then receive undesired copies of the same transmission. These signals can be delayed in time as well as be shifted in phase and frequency [65]. Figure 2.6 depicts a typical scenario of ground effect, where the ground surface reflects the electromagnetic wave, and the receiving

antenna receives both the reflected and direct signals. The shift in phase and reduction in magnitude can cause misinterpretation by the receiver, which will lead to bit errors that will cause messages to drop.




Figure 2.6: Ground effect of reflected electromagnetic waves. The receiver antenna receives undesired copies of the same message that can be delayed in time and shifted in phase and frequency [11].

The sixth factor to consider is *message overhead*. Wireless messaging standards send frames over the air. Within these frames are various overheads such as the destination and source MAC addresses, type of message, CRC, and various headers. These overheads are used to route the message to the correct receiver, check that the data received is correct, define which version is used, define the message's size, etc. A typical heliostat controller message will request the tracking parameters such as the solar position, desired motor counts, desired aimpoint, etc. and will be a relatively small message not exceeding a few tens of bytes. Figure 2.7 shows a typical frame of a short control message sent over Wi-Fi with the ZeroMQ messaging protocol that will be discussed later in the chapter. It can be seen that 66 % of the frame consists of overhead data. If the payload size was increased to 1000 bytes, only 5.5 % of the frame would consist of overhead data. The short control messages sent to the heliostat controllers thus decrease the throughput potential of the wireless network. Combining multiple messages into a single transmission could enhance the performance of the network.




Figure 2.7: Message overhead for a 30 byte message sent over Wi-Fi with the ZeroMQ messaging protocol. It can be seen that 64 bytes of overhead data is required to transmit a 30 byte data payload [12].

The seventh factor to consider is *the hidden node problem*. When an AP is located in the middle of a wireless network, all nodes should have sufficient communication range to exchange messages with the AP. The hidden node problem arises when two nodes, located on opposite sides of the AP, are not within each other's transmission range. The nodes are thus unaware of the communication originating from one another which causes concurrent transmission to the AP and corrupts the messages received by the AP. Mechanisms such as Request To Send and Clear To Send (RTS/CTS) are implemented on the AP to reduce the risk of collisions, but this increases the overhead traffic on the network hindering network performance.

2.4 ZeroMQ

ZeroMQ is a high performance asynchronous open-source networking library that was originally developed by iMatrix under the LGPLv3 license [69]. An active community currently drives the evolution and maintenance of the ZeroMQ library. The library supports a socket style Application Interface (API) and supports different message-passing communication patterns such as in-process, inter-process, Transmission Control Protocol (TCP) inter-host, and multicast messaging. On the wire, ZeroMQ uses the ZeroMQ Message Transfer Protocol (ZMTP) [70]. ZeroMQ offers multiple socket options that allow different network configurations and characteristics. Some of the most common possible socket combinations as described by Sliwinski and Lauener [69] are:

- REQ/REP - The request-reply socket combination is the most common and simplest. For each request sent from the client, the server must send a reply. The client cannot send a second request until a reply to the first request is received. This messaging combination does thus not support asynchronous messaging.
- DEALER/DEALER - The dealer-dealer socket combination acts as an asynchronous REQ/REP pattern. The dealer node on both the server and the client-side can send multiple request messages without receiving a reply. However, this socket combination is mostly used for a one to one network, as there is no routing involved in the messaging strategy.
- ROUTER/DEALER - The router-dealer socket combination allows asynchronous messaging in a one to many network configuration. The first frame sent from a router socket contains the client's identification that the message is to be routed to.
- ROUTER/ROUTER - The router-router socket combination allows asynchronous messaging in a many to many network configuration. As both sockets add an identification frame to the message, any message can be routed to a specific client on the network.

- ROUTER/REQ - The router socket can be paired with a request (or reply) socket. This allows the router, typically a network server, to asynchronously receive and send multiple requests and replies to clients on the network. The client can only send a new request once the router has replied with the previous request, allowing the server to exhibit a high level of control over each node on the network.

A typical hello world message exchange between a server and client using the REQ/REP socket combination is shown in Figure 2.8. Firstly, a socket context is created on both nodes, which is used to create the request and reply sockets on the client and server, respectively. The server binds its REP socket to the TCP port while the client connects its REQ socket to the same port. The client can then send a request message via the REQ socket that can be replied to by the server's REP socket.





Figure 2.8: Simple ZeroMQ REQ/REP messaging structure [13].

ZeroMQ acts as middleware between the operating system and the network. This allows the user to relatively easily design and configure the network using the ZeroMQ API that is provided on many software platforms, including C/C++, Java, and Python. ZeroMQ decodes the user instructions to form complex transfer functions used to transmit data via different message passing patterns such as TCP.


Zhou et al. [14] investigated the performance of popular middleware platforms by comparing omniORB, Binder, and ZeroMQ. During each test, 1 million messages were sent over the network between a server and client. In each test, the time before the first message was sent (T_1) and the time after the last message was received (T_2) were noted. The average single packet transfer time was determined with equation 2.2,

$$T_{\text{Single Packet}} = \frac{T_2 - T_1}{1e6}. \quad (2.2)$$

For the first tests, the packet size was kept at 1024 bytes, and the number of nodes was increased from 2 to 10. For the second test, the number of nodes was kept at two while the packet size was increased from 128 bytes to 8192 bytes. The results from the tests can be seen in Figure 2.9.



(a) Transfer delay with a 1024 bytes packet size and increasing number of nodes



(b) Transfer delay with 2 nodes and increasing packet size

Figure 2.9: ZeroMQ Transfer delay tests from Zhou et al. [14].

It is clear that ZeroMQ outperforms the other middleware platforms, and this is partly due to the zero message copy structure within ZeroMQ. ZeroMQ thus provides a high-performing messaging library that simplifies network design and implementation and can therefore be regarded as an optimal solution for communication in a CST heliostat field.

The ZeroMQ library was used to design the original Helio100 project at Stellenbosch University. It is thus a logical choice to utilize ZeroMQ in this project as it allows work developed during the course of the thesis to be seamlessly incorporated into the next generation heliostat field.

2.5 Measuring network performance

Network performance is mainly measured for two different reasons. The first is to provide figures for marketing purposes, such as to define the number of messages a network is capable of delivering in a finite period. Secondly, network performance measurement is used for diagnostic purposes, such as to identify latencies in the system. More specifically, network tests are aimed to identify patterns in the network performance, such as periodic latency spikes or packet loss within a specific scenario. This section will explain the different network performance metrics and how to measure them.

It is essential to understand that the system to be measured is discrete, i.e., messages are sent and received at a specific point in time. Thus, we can not assume that all the network performance graphs are continuous, and therefore latency and throughput cannot be measured for a specific point in time. It is, however, possible to measure latencies and throughput for a specific message(s).

An important network performance metric is latency. Latency is the interval between sending and receiving a message. The latency for the n^{th} message (L_n) can be calculated as

$$L_n = R_n - S_n, \quad (2.3)$$

where R_n and S_n are the timestamps that the message was received and sent, respectively. Latencies can be measured in a single direction, i.e., uplink or downlink latencies, or as a Round Trip Time (RTT) latency, which measures the time duration of when a message was sent until a reply is received. The presence of high latencies indicates that the network is performing slowly as packets take a long time to reach their intended destination. It is important to realize that timestamps should be taken using the same machine as clock times can differ on two different machines. The alternative is to either synchronize the clocks on different machines or measure time durations rather than timestamps.

Bandwidth is a performance measure to describe the maximum amount of data that can be transferred throughout the network [71]. The bandwidth of a network is limited by the communication standard, network devices, and the standard of the connection. Throughput is defined by the actual amount of messages (or data) sent over a period of time. In short, bandwidth represents the theoretical limit for the amount of data to be transferred, while throughput represents the actual data transferred over a specified period.

Throughput can be measured as either an average over a specified time (TP_{avg}) or per message sent over-the-air (TP_{ota}). TP_{avg} considers the total amount of data (excluding overheads used for addressing, CRC, acknowledgments, etc.) that was sent over the entire period of a test. This will include time spent in queues, polling for messages, and during the encoding and decoding of messages. TP_{avg} can be measured in the uplink and downlink as well as the combined total. This produces the most realistic measurement of throughput. The average throughput can be calculated as

$$TP_{avg} = \frac{N \times P_{size}}{R_n - S_1}, \quad (2.4)$$

where N represents the number of messages, and P_{size} is the size of the data payload. TP_{ota} considers only the throughput associated with a single message. The time a message is sent by the sender (T_S) and the time the message is received by the receiver (T_R), is used to calculate the over the air throughput. Over-the-air throughput can be calculated in the uplink and downlink direction. This method does not consider polling durations, encoding, or decoding time of the message and is thus expected to be higher than the average throughput. The over-the-air throughput can be calculated as

$$TP_{ota} = \frac{P_{size}}{T_R - T_S}, \quad (2.5)$$

Furthermore, it is vital to measure the reliability of the network. This is done by measuring the Packet Error Ratio (PER), which is defined as the ratio, in percent, of the number of packets that were not received successfully. This can be done by capturing network traffic on both the sender and receiver side and determining which

packets were sent unsuccessfully. It is also possible to incorporate a timeout on the network before a packet is dropped and retransmitted. Retransmissions ensure reliability in a network but have a negative effect on network throughput.

Chapter 3

System Requirements

Stellenbosch University launched the Helio100 project in 2015, introducing their novel HelioPod technology, as described in Section 1.1.4. The use of the HelioPod technology decreases capital costs by eliminating the need for expensive foundations for the heliostats. If a fully operational wireless communication system is installed, further cost savings will be incurred by removing the need for trenching for cabling.

The Helio100 project's success resulted in the inclusion of the HelioPod technology in the PREMA H2020 project, which aims to advance novel energy systems in the drying and preheating of furnace manganese ore. The PREMA H2020 project aims to incorporate the HelioPods to supply solar thermal energy to the DLR's 2.5 MW_{th} CentRec[®] particle receiver. The particles are to be heated to 900 °C and transported to a 12-hour thermal energy storage system. Air, heated by the particles to 800 °C, will be used to preheat manganese ore to 600 °C. The incorporation of the solar energy preheating system aims to reduce CO₂ emissions by up to 20 % and incur a cost-saving of at least 15 %. The high level proposed plant design can be seen in Figure 3.1.




Figure 3.1: Proposed CST plant for PREMA H2020 project [15].

3.1 Field sizing

The heliostat field was designed to meet the $2.5 \text{ MW}_{\text{th}}$ design point of the particle receiver. Hockaday et al. [15] published research on the sizing of the heliostat field for locations with varying levels of average annual DNI levels. For this research, four locations were used, two in South Africa, one in Spain, and one in China. It was found that the location with the lowest annual DNI, China, required 270 HelioPods (1620 heliostats), 23 towers, and 16 hours of thermal energy storage to supply the required thermal energy for the drying and preheating process. Each HelioPod will be fitted with a single Local Control Unit (LCU) responsible for all communication and motor control for each facet on the HelioPod. The heliostat field will thus consist of 270 LCUs. The proposed field layout, as seen from above, is depicted in Figure 3.2. From the proposed layout, the following network design requirements can be identified:

- With the tower positioned at (0,0), it is clear that the communication range of each LCU should exceed 80 m Line of Sight (LoS).
- The network capacity limit should exceed 270 nodes.
- The specified wireless communication technology should be able to operate in a dense network.

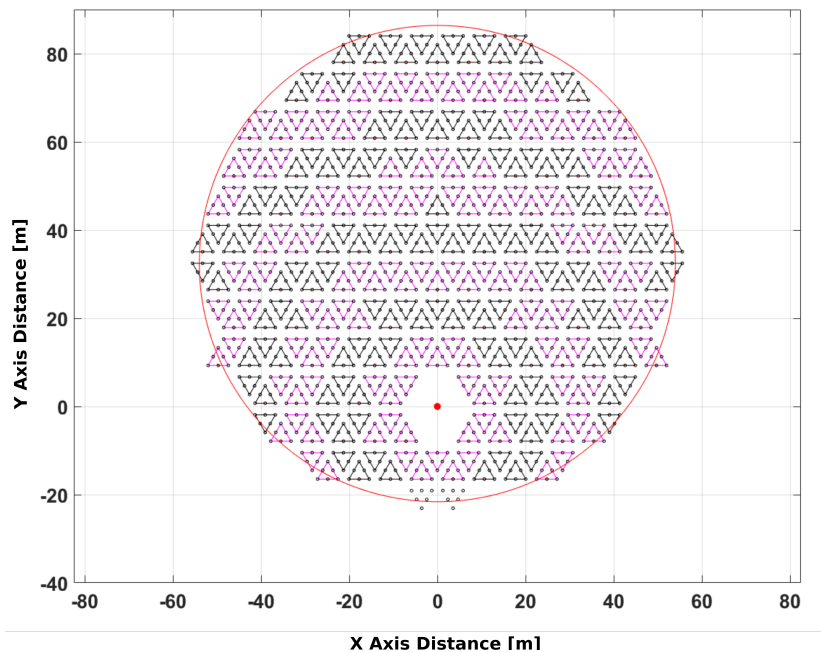


Figure 3.2: Proposed field layout as seen from above with tower position at (0,0). It can be seen that the maximum data transmission range is approximately 80 m.

3.2 Control interval

To an observer on earth, the sun appears as a round disk with near-constant radiation distribution over its surface [3]. A heliostat's ability to reflect the incident sunlight is limited by the subtending angle (Ω_s) of the sun, as can be seen in Figure 3.3.




Figure 3.3: Solar subtending angle (adapted from [16]).

The sun's subtending angle (≈ 9.3 mrad) as well as the distance from the heliostat to receiver or slant range (L_{slant}) limits the smallest image that an ideal heliostat can reflect on to a receiver. As the slant range increases, the reflected image area increases. Therefore, it is common to use the furthest heliostat's smallest theoretical image to specify the receiver aperture size [3]. Furthermore, the required tracking accuracy of a heliostat can be related to the fraction of the reflected image that falls onto the intended receiver target, as depicted in Figure 3.4.




Figure 3.4: Area of intersection of a solar image from an offset heliostat and its intended receiver target (adapted from [3]).

The area of interception can be expressed as a function of the ideal heliostat reflected image radius (R) and the offset distance (d) between the intended receiver target and the actual reflected target as a result of slope errors and tracking errors [3; 72]:

$$A_{\text{intercept}} = 2R^2 \cos^{-1} \left(\frac{d}{2R} \right) - \frac{d}{2} \sqrt{4R^2 - d^2}, \quad (3.1)$$

where,

$$R = 2L_{\text{slant}} \left(\frac{\Omega_s}{2} + \sigma_{\text{slope}} \right) \quad \text{and} \quad d = 2L_{\text{slant}} E_{\text{tracking}}, \quad (3.2)$$

and where Ω_s is the subtending angle of the sun as expressed in Figure 3.3, σ_{slope} is the mean slope error of the heliostat and E_{tracking} is the heliostat's normal vector angular tracking error. Equation 3.1 is rewritten to determine the ratio of interception (R_i) normalised for the slant range [3]:

$$R_i = \frac{4k^2 \cos^{-1} \left(\frac{E_{\text{tracking}}}{2k} \right) - E_{\text{tracking}} \sqrt{4k^2 - E_{\text{tracking}}^2}}{2\pi k^2} \quad \text{with} \quad k = \frac{\Omega_s}{2} + \sigma_{\text{slope}}. \quad (3.3)$$

Equation 3.3 is plotted in Figure 3.5, with $\Omega_s = 9.3$ mrad and $\sigma_{\text{slope}} = 1$ mrad which is equal to the slope error specified for the HelioPod facets [32].




Figure 3.5: Heliostat ratio of interception versus normal vector angular tracking error

The HelioPod facets are specified to have a 0.5 mrad tracking error [32] which relates to a 94 % intercept ratio by using Equation 3.3 with $\Omega_s = 9.3$ mrad and $\sigma_{\text{slope}} = 1$ mrad. A 0.5 mrad tracking accuracy is low compared to commercial heliostats, but the small size of the facets on the HelioPod will result in more accurate manufacturing.

A heliostat's aiming position will be adjusted periodically by sending a control message to the heliostat. The duration between each heliostat position adjustment is defined as the control interval (t_c). As the heliostat tracking point will lag the

position of the sun it creates a tracking delay angle (θ_{delay}) of $\omega_{\text{earth}} \times t_c$, with ω_{earth} representing the angular velocity of the earth ¹ [3]. By taking into account the specified minimum desired tracking error (0.5 mrad) and the angular velocity of the earth, ω_{earth} , the maximum control interval is calculated as 6.8 seconds ².

A control interval of 6.8 seconds will require zero tracking losses resulting from actuator inaccuracy, foundation structures, solar position calculations, etc. The presence of tracking losses outside of the control interval scope will result in a tracking inaccuracy and thus not meet the 0.5 mrad requirement. To account for these tracking errors, the chosen tracking error should thus be significantly lower than the upper boundary of 6.8 s. Furthermore, the heliostat tracking algorithm can be altered so that the heliostat can move to a position that leads the exact position required to reflect the incoming solar flux on the intended target. In other words, at any time when a heliostat should update its aiming point, it can update the position to the current timestamp + $\frac{t_c}{2}$ to ensure a lower average offset distance over the control interval.

3.3 Message structure

In the current Helio100 project, a control interval of 1.5 s is used for a field of 17 HelioPods. For the next-generation plant, a target control interval of 1 s is proposed to ensure a quick response to any emergency situation. During each control interval, the Field Control unit (FCU) will poll each LCU with the updated motor counts, heliostat status, and current time. Upon receiving this control message, each LCU will be required to send an Acknowledgement (ACK) message back to the FCU containing the updated battery level, heliostat state, and error flags. The structure of a typical polling and ACK message can be seen in Table 3.1.

Table 3.1: A typical polling and ACK process message structure

Message Type	Data Payload	
	Data	Size [bytes]
polling message	Message ID	6
	Date and Time	14
	Heliostat status	6×2
	Motor counts	6×12
ACK message	Error flags	6×1
	Heliostat status	6×2
	Battery level	6×2

¹ $\omega_{\text{earth}} = \frac{2\pi}{24 \times 3600} = 7.272\text{e-}5 \text{ rad s}^{-1}$

² $\frac{0.5 \text{ [mrad]}}{0.072722 \text{ [mrad/s]}} \approx 6.88 \text{ s}$

The polling message includes a timestamp from the FCU to act as a heartbeat timer for the LCU. If the LCU does not receive an update message for a specified amount of time, a heartbeat timer would be triggered, which will broadcast a loss of communication, after which the heliostat will move to a defocus state. The polling message also includes two bytes to change the state of each of the six heliostats on the HelioPod. The different states can include track, freeze, setup, calibrate, stow, etc. The ACK message from the LCU will confirm the new status of each of the six heliostats and confirm the battery level (0-99). The ACK message contains a single byte per heliostat for error flags such as loss of communication to a single heliostat from the LCU, parameter errors, low battery levels, etc. Over-the-air, each number or character is sent as a single byte char; thus, a six-digit number (motor count) would result in a 6-byte message portion.

The round trip polling and ACK message exchange requires 134 bytes to be sent over the network between the FCU and LCU. If a 1-second control interval is required, a 134-byte message is sent to each of the 270 LCUs in the network per second. This will result in a total network throughput requirement of 289.4 kbps³.

³ $134 \text{ bytes} \times \frac{8 \times 10^{-3} \text{ [kb]}}{1 \text{ [byte]}} \times \frac{270 \text{ LCUs}}{1 \text{ [s]}} = 289.44 \text{ kbps}$

Chapter 4

Hardware Setup

This chapter will give an overview of common network topologies and the performance influence associated with the use of each. A comparison of each wireless technology described in Section 2.1 follows with a specific evaluation of each technology with regards to its use in a CST heliostat field. The wireless communication standard to be used in the network tests is selected based on the evaluation. Finally, the hardware components used within the network are specified and described.

4.1 Proposed topology and communication standard

4.1.1 Topology

The chosen network topology within a designed network significantly impacts network performance measures such as latency, throughput, and battery lifespan. Some common network topologies that will be discussed include star, cluster-tree, and mesh topologies. These topologies can be seen in Figure 4.1. Each topology enables unique routing paths and also determines whether broadcast or unicast can be used. The specified topology can also influence the packet and overhead size as routing information can be stored as additional overhead [73]. A network topology can be evaluated according to scalability, latency, redundancy, energy consumption, reaction time and ease of deployment.

4.1.1.1 Star topology

A star topology is an example of a point-to-multipoint connection. A single central gateway node that acts as an FCU to the network within a star topology allows direct connection to corresponding LCUs in the network. Processed or stored information on each LCU can be transmitted directly to the FCU, where the end-user can access it. Each LCU in the system is only responsible for communication between the FCU and itself and can thus periodically enter a low power state when not transmitting and receiving. This enables star topologies to have low overall power consumption [73]. The FCU is the core limiting factor within a star network. The scalability of the network depends on the range of the FCU and the capacity constraint of the FCU




Figure 4.1: Common network topologies (adapted from [17]).

gateway. The FCU is also a single point of failure on the network, but redundancy in the network can be built in by incorporating a backup FCU to trigger if the primary fails. With the exception of an FCU failure, a failure of a node in the network will not influence the performance of any other node, and thus troubleshooting and fault finding in a star topology is relatively simple.

4.1.1.2 Cluster-tree topology

A cluster-tree topology is also referred to as a hierarchical topology. The deployment of the nodes resembles a logical tree. Each node in the network, with the exception of the FCU node, has a parent node and potentially a child node. Communication can only occur between a parent and child node pair, introducing hierarchy into the network. Cluster-tree topologies primarily utilize unicast messages instead of broadcast messages to prevent flooding of the network [73]. Unicast messaging allows nodes that are not involved in the current communication branch to enter a low power state to conserve battery life. The primary advantage of a cluster tree topology is the network's ability to be scaled to cover larger areas. However, this will introduce multiple hops in the communication path, which increases latency as the message is relayed between the parent-child pairs. A cluster-tree topology's primary disadvantage is that when a parent node fails, all communication to the downstream nodes also fails, thus introducing multiple failure points. This higher complexity and multiple points of failure could lead to time-consuming and costly network maintenance.

4.1.1.3 Mesh topology

A WMN typically incorporates a central network coordinator (or FCU) responsible for initiating the network with a network ID and various security protocols depending on the wireless standard used. The rest of the wireless nodes in the network are either classified as router or end nodes. A router node is able to relay a message

within the field, whereas an end node's communication is restricted to a parent-child relationship. The main difference between a cluster-tree and a WMN is that all router nodes can communicate with all other nodes within its communication range. Algorithms such as the AODV used by ZigBee ensure efficient routing of messages throughout the network. The self-healing feature of a WSN ensures a high network redundancy as messages can be routed around nodes that failed. A WMN is also highly scalable as nodes can be added to the network without additional gateways. The major disadvantage of a WMN is that multiple transmission hops introduce significant latencies. Another disadvantage of a WMN is that all nodes in the field do not drain their batteries at the same rate. Nodes closer to the FCU will relay more packets and thus drain their batteries faster. WMNs are commonly implemented where network redundancy is essential and is complicated to set up and troubleshoot.

4.1.2 Wireless communication standard comparison

The different wireless communication standards, as discussed in Section 2.1, are compared in Table 4.1. Each technology is specified a network topology that would render the standard suitable for use in a CST plant. The chosen topology is influenced by the range of an individual node as well as the capacity constrain for the central gateway of the network, e.g., a ZigBee device has a range of 20 m. Thus the use of a mesh network is necessary to ensure a sufficient range for the network.

Table 4.1: Wireless communication technology comparison. Each wireless standard is specified a topology as it would be implemented in a CST plant. The main theoretical performance metrics are compared.

Parameter	LoRa	NB-IoT	ZigBee	BLE	Wi-Fi 4	Wi-Fi 5	Wi-Fi 6	Wi-Fi HaLow
Topology	Star	Star	Mesh	Mesh	Tree	Tree	Tree	Star
Range of individual module	15 km	10 km	20 m	50 m	250 m	120 m	250 m	1 km
Theoretical throughput	50 kbps	200 kbps	250 kbps	1 Mbps	600 Mbps	1.3 Gbps	3.5 Gbps	347 Mbps
Energy consumption	Very low	Very low	Low	Low	High	High	Moderate	Low
Single transmission latency	Classification dependant	<10 s	<200 ms	<200 ms	<10 ms	<10 ms	<10 ms	<100 ms
Complexity of implementation		High	Moderate	Moderate	Low	Low	High	High
Cost	Low	High	Low	Low	Low	Low	High	High

When comparing the technologies, both LoRa and NB-IoT do not meet the 289.4 kbps network throughput requirement. NB-IoT also has the added disadvantage that the use of a licensed spectrum increases the cost and the complexity of installation. Both ZigBee and BLE have a relatively low theoretical range and would require a WMN, which, as discussed previously, increases the latency on the network.

Wi-Fi HaLow, on the surface, seems like the ideal technology to be incorporated into a CST heliostat field. The sub 1 GHz operating frequency allows long-range communications. The technology also boasts a high throughput rate coupled with battery-saving features such as TWT. The major drawback of Wi-Fi HaLow is the

limited availability of devices to be incorporated on LCUs. The lack of implementation of the technology on a large scale renders the available devices extremely expensive and thus inadequate for use in a heliostat field that aims to reduce the cost of CST technology. Wi-Fi 6 suffers from the same limitation. The 802.11ax standard, at the time of writing, has not yet been ratified by the IEEE, with ratification dates set to Q4 in 2020. Unlike Wi-Fi HaLow, chipmakers have already started developing routers and limited end devices based on the draft specifications of the 802.11ax standard as set out by the Wi-Fi Alliance. Wi-Fi 6 also promises to provide high-performance wireless communication to future CST plants. It boasts technology such as beamforming to increase range, TWT to extend battery lifespan, and OFDMA and BSS coloring to reduce collisions in a dense network. A previous study by the authors [33] concluded that Wi-Fi 6 is the optimal technology to be used in future modular heliostat plants such as the H2020 PREMA project, but the lack of end devices to incorporate in an LCU eliminated the use of this standard for the project in question.

The two remaining wireless communication standards include the Wi-Fi-4 and Wi-Fi 5 standards. When comparing these two technologies, Wi-Fi 4 incorporates the 2.4 GHz frequency band instead of the 5 GHz band, which increases the theoretical transmission range. The maximum theoretical throughput of 600 Mbps also satisfies the minimum plant throughput requirement. The technology is also popular, which ensures a wide range of off-the-shelf devices compatible with the standard. Based on the aforementioned reasons, the 802.11n standard or Wi-Fi 4 was chosen as the preferred wireless standard. As a Wi-Fi 4 AP has a maximum capacity of 255 devices on the network (limited by IP addresses on the BSS), a cluster-tree topology will be required in a plant of 270 nodes. The number of transmission hops can be limited to 2 to avoid high latencies.

By utilizing the Wi-Fi 4 standard and a cluster-tree topology, several clusters can operate in non-overlapping frequency bands, which will allow parallel computation and messaging. Thus, the control interval specification can be focussed to ensure a 1 s control interval within each cluster and not the entire field. In other words, if each cluster operates independently in a non-overlapping frequency band and ensures a sub 1 s control interval, the field-wide control interval will satisfy the 1 s requirement.

4.2 Hardware overview

The objective of this study is to replicate a heliostat field communication network as accurately as possible. A Lenovo ThinkPad i5 laptop was used as an FCU, which will house most of the high-performance testing procedures. The Lenovo laptop was connected via a wired ethernet link to a TP-Link EAP 110 802.11n outdoor AP. The Lenovo laptop was placed inside the current Helio100 site's control container to eliminate the possibility of weather damage. The AP was placed on top of the control container at the height of approximately 6 m, allowing optimal field of view

to the heliostat field, as shown in Figure 4.2. The elevated position also ensures a better quality LoS path to each LCU.



Figure 4.2: The AP was positioned on top of the control container, behind the heliostat field. The receiver can be seen located on the tower which is visible above the horizon.

A total of 50 Raspberry Pi (RPi) Zero WH microcontrollers were used to replicate the LCUs in the field. Each RPi Zero WH microcontroller was powered by an 8000 mAh power bank supplying 2 A via USB. As each LCU will be used to control all six heliostats mounted on the HelioPod, the 50 RPi microcontrollers replicated communication to 300 heliostats. This represents a realistic replication of a single cluster within the cluster-tree topology as specified for the 802.11n standard in Table 4.1.

The RPi Zero WH microcontroller is an off-the-shelf bought item with onboard 802.11n wireless connectivity. Furthermore, the RPi Zero WH is a small, cheap, and low power option compared to other RPi products. A BCM 2835 (32-bit ARM-based processor) SoC also allows the RPi Zero WH to compute solar angles with more than sufficient precision accurately. A comprehensive overview of the main hardware components and specifications are shown in Table 4.2.

Table 4.2: Overview of the main hardware specifications of the three primary components of the network.

Component	Specification description	
TP-Link EAP 110-Outdoor Wireless N Access Point	Wireless standards	IEEE 802.11n/g/b
	Transmit power	FCC: <22 dBm
	Antenna type	2*3 dBi External Omni waterproof
Lenovo ThinkPad laptop	Processor	Intel® Core™ i5-8250U CPU @ 1.6 GHz x 8
	RAM	8 GiB
	OS type	Ubuntu 18.04.5 LTS 64-bit
Raspberry Pi Zero WH	Processor	32-bit BCM 2835 SoC @ 1 GHz
	RAM	512 MB
	Wireless Communication	BCM43143 SOC supporting 802.11n, Bluetooth 4.1 and BLE

Chapter 5

Software Setup

The network architecture was designed using the ZeroMQ library, an open-source high performance asynchronous messaging library that utilizes a socket style application interface and supports message patterns such as in-process, inter-process, TCP, and multicast messaging as described in Section 2.4. The high-performance characteristics and interoperability with different programming languages rendered ZeroMQ the ideal middleware for the network's backbone. For this project, Python was selected as the preferred programming language because of the ease of use of the pyzmq libraries with user-friendly creation and closing of TCP ports. The RPi zero WH microcontrollers are also loaded with the Raspbian operating system, which has a built-in Python interpreter.

In this chapter, the ZeroMQ message structure is described, giving background to the routing of messages through the network. The network architecture is also illustrated, and further elaboration on the different processes in the network is given. The chapter concludes with an explanation of offset time determination and some practical challenges experienced with time drift during the initial phase of testing.

5.1 ZeroMQ message packet structure

This section will give an overview of the ZeroMQ message packet structures relevant to this thesis. More detailed information on ZeroMQ message packet structures can be found in the ZeroMQ handbook [13].

ZeroMQ uses multipart messages, or message envelopes, to relay data from one user to the next. Each data part within a message envelope is a separate `zmq_msg_item` and is separated by a frame boundary. When sending a message envelope from node A to node B, node A will queue and "build" the message envelope with all the frames entirely before sending the message over-the-air. Node B will then receive all parts of the message envelope or none at all. Thus the frames are encapsulated in the envelope and are not sent individually over the network.

Message envelopes are used even in the simplest request-reply communication between a REQ/REP socket pair. When the REQ socket initiates communication by sending "Hello" to a REP socket, the simplest form of a message envelope is

created. The envelope consists of an empty delimiter frame and a frame containing the data, as shown in Figure 5.1. Upon receiving the envelope, the REP socket saves all frames up to the empty delimiter frame (in this case, only the empty delimiter frame is saved). The remaining frames are sent to the application. In this case, this will only include the "Hello" frame. When replying, the REP socket adds the frames containing the reply data to the previously saved frames (the empty delimiter frame) and sends the data back. In this case, the message structure will be identical to Figure 5.1, but with a possible change to the data frame.




Figure 5.1: Simple ZeroMQ request-reply envelope (adapted from [13]). The empty delimiter frame and the data frame can be seen in the message envelope.

Message envelopes are crucial when more complex socket pairs are used in a network. Figure 5.2 shows an example of a ROUTER/DEALER process located "in-between" the REQ and REP sockets.



Figure 5.2: ZeroMQ extended request reply pattern (adapted from [13]). The router-dealer combination is located between a request-reply pattern. This allows messages to be routed to specified destinations.

As stated in Section 2.4, the ROUTER socket allows routing within the network, and the DEALER socket acts as an asynchronous REQ or REP socket. A ROUTER socket tracks every connection it has by attaching a connection identity frame on top of the empty delimiter frame for every message envelope that it receives. The connection identity frame is simply a binary string unique to a specific node on the network. A message pattern received by the ROUTER from the REQ socket (with identity "XYZ") is shown in Figure 5.3.




Figure 5.3: ZeroMQ message envelope sent from a REQ socket to a ROUTER socket (adapted from [13]). The ROUTER socket adds the connection identity frame of the sender on top of the empty delimiter frame.

The ROUTER/DEALER acts as a proxy loop in the network used to read on one socket and write on another. The three frames received by the router are thus relayed to the DEALER. The DEALER, in turn, sends all three frames to the REP socket. The REP socket saves all the frames up to the empty delimiter (the connection identity and the delimiter frame) and sends the remaining data frames to the application. The reply message is structured in the same manner from the REP socket to the ROUTER socket via the DEALER. The ROUTER socket strips off the connection identity frame before sending the reply to the REQ socket, which only receives an empty delimiter and data frame.

This network configuration allows multiple REQ nodes to connect to a single server via the network proxy. The saved connection identity allows the ROUTER socket to route the message back to the original sender. The network configuration in Figure 5.2 can be altered by replacing the DEALER socket with another ROUTER socket. This will allow multiple REQ processes to connect to multiple REP processes as both ROUTER sockets have the ability to route messages to the appropriate process. In the context of a heliostat field communication system, this allows the network architect to connect multiple LCUs to different processes running on the FCU (which can include one or more servers connected to the network) via the ROUTER/ROUTER proxy. The processing of the field on the FCU side can thus easily be shared among different processing machines.

5.2 Network architecture and processes

The network was designed with network performance testing in mind rather than in field daily operation. The network consists of three main processes, the server, the broker, and the LCU. Both the server and the broker processes run on the FCU, with each LCU in the field running its own process. An overview of the network architecture can be seen in Figure 5.4. It can be seen that there are three socket connections between the server and the broker:

- *Port 5555*: A DEALER/ROUTER socket pair used to send polling messages to each LCU to update the aiming positions of the heliostats.
- *Port 5557*: A REP/REQ socket pair used to associate new clients on the network before a test.
- *Port 5559*: A DEALER/ROUTER socket pair used to determine the time offset between the server and an LCU.

There are a further two socket connections between the broker and an LCU:

- *Port 5556*: A ROUTER/REQ socket pair used to relay the polling message updates from the broker to the respective LCU.
- *Port 5560*: A ROUTER/REQ socket pair used to determine the time offset between the server and the LCU.

Furthermore, a PUB/SUB socket on port 5558 is used to relay an emergency message from the server to all other network components in order to stop the test at any time.

All ports referenced refer to TCP ports. TCP is the chosen messaging standard of the ZeroMQ middleware as it ensures reliable message delivery and prevents any data loss. The TCP standard uses various other protocols such as IPv4 and Ethernet to relay messages on the network, as shown in Figure 2.7. The TCP standard incorporates a CRC to ensure there is no loss or corruption in the received data, which, if detected will trigger a data retransmission. TCP also uses ACK messages to ensure data was received by the intended receiver on the network, ensuring message delivery reliability.

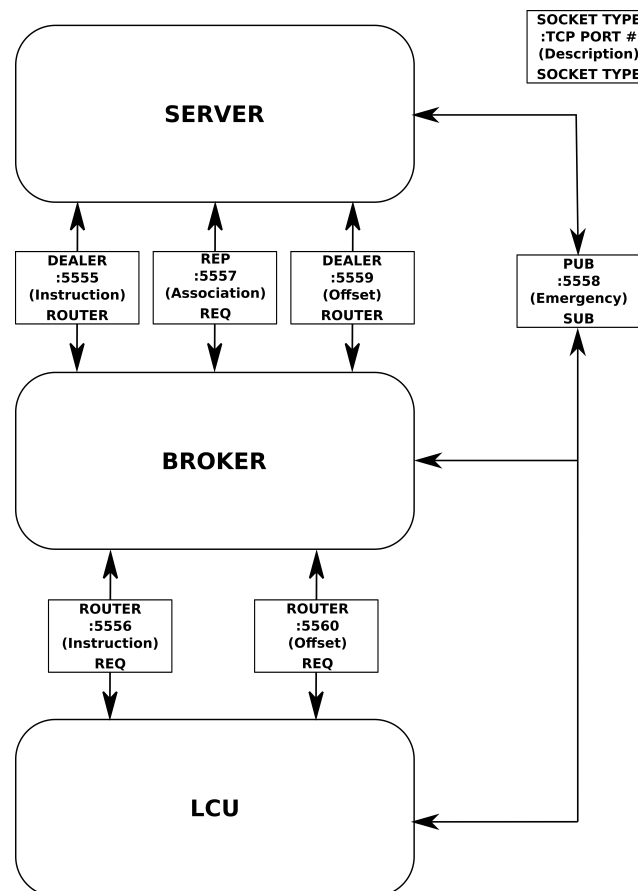


Figure 5.4: Network architecture overview. The network was designed using ZeroMQ and TCP communication. A total of three socket pairs are present between the server and the broker process (both hosted on the same machine). Two socket pairs are present between the broker and LCU processes and a publish socket is used for emergency situations.

5.2.1 Server process

The server process is responsible for the majority of the processing within the test. The flow of the server logic can be seen in Figure 5.5. A network test is initialized, and the specific LCUs to be included in the test are specified. The server polls the incoming messages from the broker for request to join messages from the LCUs. If all specified LCUs have not joined before a timeout procedure, the user is prompted with an option to proceed without the missing LCUs or to wait longer for their connection. If all are connected, or the test is started manually, a procedure designed to measure the offset is triggered, which will be described in Section 5.3.

Once the main loop is started, the server polls the broker (backend) for any incoming messages. Incoming messages have a higher priority than outgoing messages to avoid queuing of messages, which will skew the latency measurements. If a message is available from the broker, the server will decode and unpack the message and log the time that the message was received and the corresponding message ID and save it in memory. The LCU from which the message was received is added to an "available to receive instruction" table to ensure the next instruction is sent as soon as possible. If no message is available from the broker within the 2 ms polling interval, an instruction message is built, encoded, and sent to the LCU that received updated information the longest ago. Again the time and message ID is logged. The corresponding LCU is removed from the "available to receive instruction" to ensure a new instruction is not sent without receiving a reply.

After the test is completed (600 s), a shutdown procedure is initiated on all LCUs to send all their logged information to the server and exit the test. The server saves all message logs (from the server itself and every individual LCU that participated in the test) and metadata before exiting the procedure itself.

5.2.2 Broker process

The broker process acts as a proxy in the network between the LCU and server processes. It allows a single connection between the LCU and broker process compared to a dedicated connection between each LCU and server process, which scales poorly. For example, when subscribers are connected to a publisher directly, a new subscriber can be added by manually adding the connection to the publisher on the new subscriber. Suppose a new publisher was to be added to the system. In that case, the direct connections on each of the subscribers have to be altered manually, which can be a timely and complicated exercise. By incorporating a broker process, each subscriber and publisher only have to establish a proxy connection. When further processes are added, no new manual connections are required as all messages move through the broker process and thus through the already established connection. Processes can thus establish a connection via dynamic discovery using ROUTER sockets and specified identities. The broker process can be seen in Figure 5.6.

The broker process starts by connecting all the sockets to the server and LCUs and registering a polling process to poll for incoming messages from the LCU side.

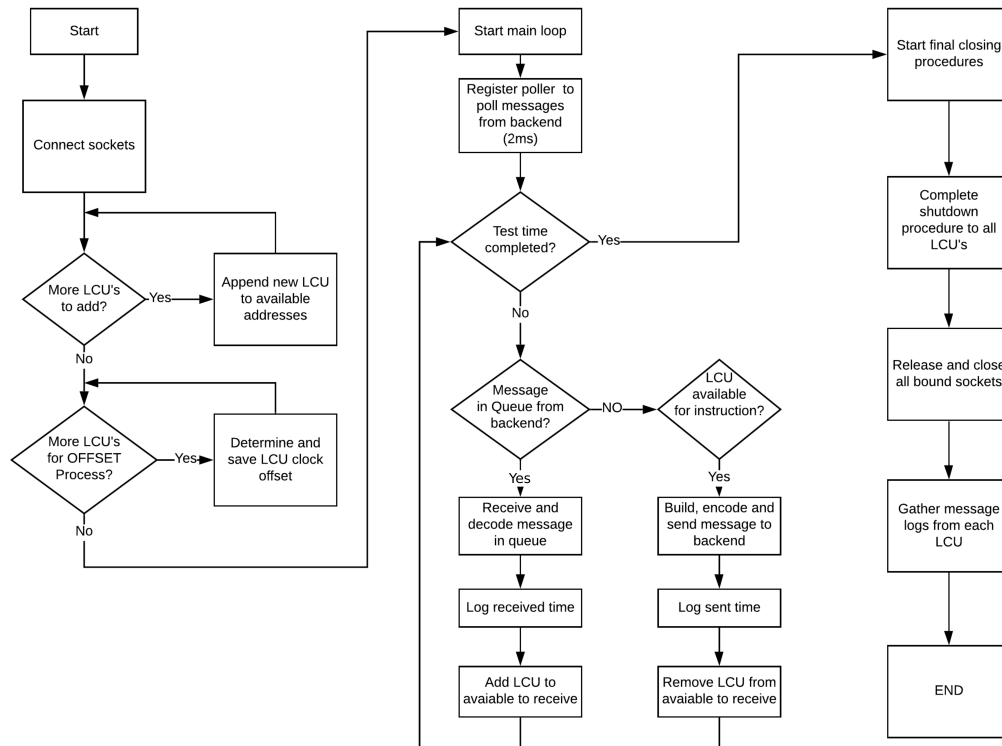


Figure 5.5: Server process logic flow diagram.

The broker will poll the LCUs for requests to be associated with the test and forward these messages to the server until the server has reached capacity and enters the offset procedure. During the server's offset procedure, the broker continuously forwards offset data between the server and the associated LCU. Upon completing the offset procedure, the broker enters the main loop and registers both the frontend (server process) and backend (LCU process) pollers. The broker monitors both up-link and downlink sockets to detect incoming messages. If a message is detected on the backend, it is forwarded to the server for processing. If a message is detected on the frontend, it is forwarded to the specified LCU for processing. This main loop executes until the server triggers the shutdown procedure after a specified time duration. The broker process's purpose is mainly to relay messages from the backend to the frontend and vice versa. However, more complex filters and message priority algorithms can be implemented on the broker when designing a long term fully operational heliostat network.

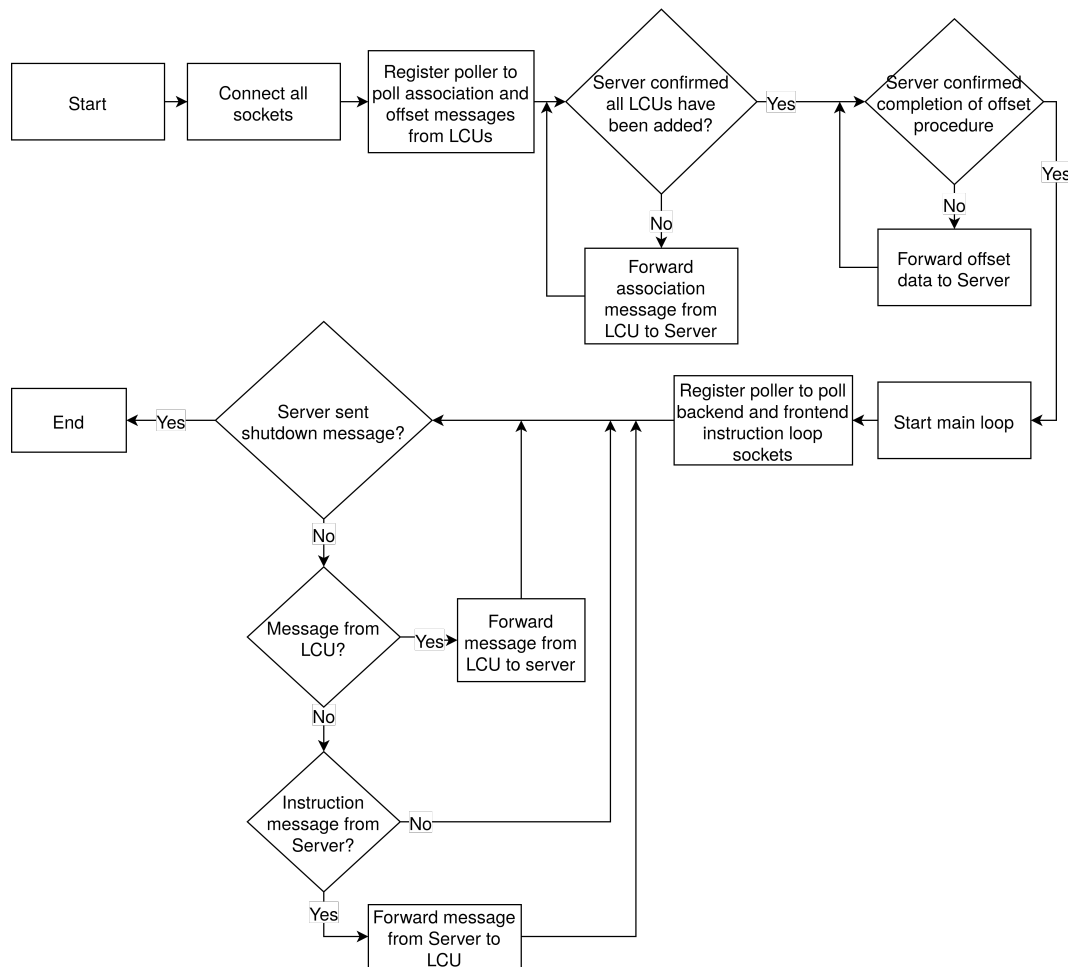


Figure 5.6: Broker process logic flow diagram.

5.2.3 LCU process

Each LCU in the network has an LCU process running on it. The LCU process logic flow can be seen in Figure 5.7. The LCU process is designed to be as computationally inexpensive as possible to limit skewed network test results originating from slow processing speeds.

The LCU process starts by connecting all the required sockets to the broker process. An association message is sent to the broker, which is either accepted or declined by the server. If the association fails, the node will exit the process and wait for future tests to start. If the association to the test is accepted, the LCU will wait for an offset message from the server via the broker, to which it will periodically reply until the offset procedure is completed. Within the main loop, the LCU waits for a message from the broker, which it decodes. The time at which the message was received is logged in memory. A reply message is then built and encoded. As the message is sent to the broker, the time is logged in memory. After the message has been sent, the logged sent and received times are written to a text file on the LCU. This process is done while waiting for the next instruction to minimize the

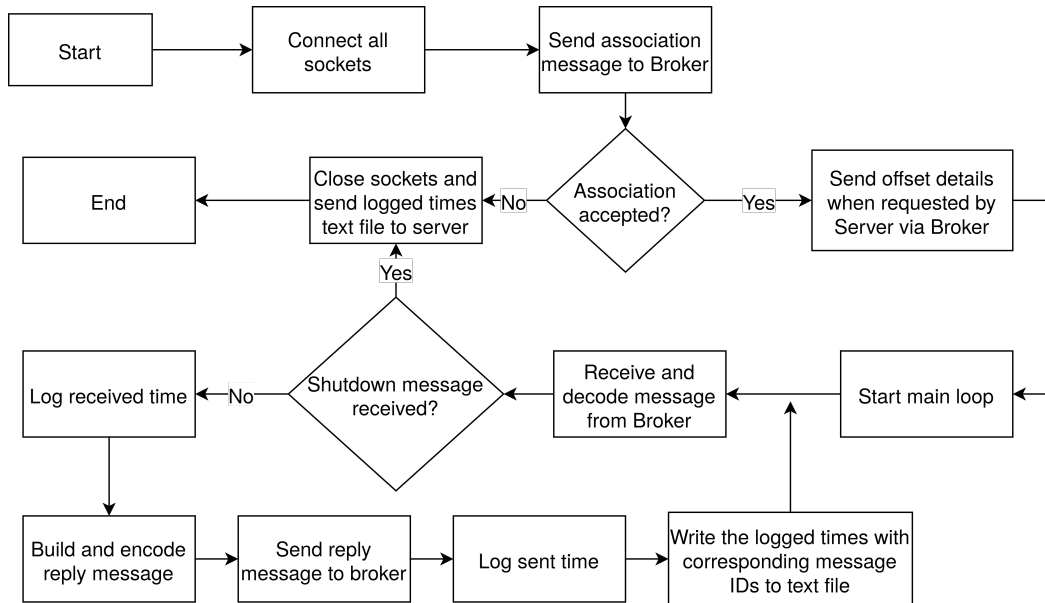


Figure 5.7: LCU process logic flow diagram.

computational load. The RPi Zero WH only has access to 512 MB of RAM, and it is thus required to move the logged times out of the memory to avoid a system crash. Once the server issues the shutdown instruction, the text file is sent to the server via Secure Copy (SCP) protocols before exiting closing all sockets and exiting the process.

5.3 Determining the time offset

For the purpose of performance tests, it is essential to log accurate message sent and received times to compute uplink and downlink network performance statistics using Equations 2.3 - 2.5. When measuring RTT latencies, all time measurements are taken on the same computing machine, thus leading to high precision timestamps. RTT latency thus represents the most accurate network performance parameter. However, when it is essential to analyse the uplink or downlink performance trends in a network, time stamps of when a message was sent and received between two nodes on the network are required. This can be achieved by two methods:

1. Both nodes are connected to a data acquisition device that measures a message's timestamp using signal triggers from each node. All timestamps are taken on the same machine and can thus be used to determine uplink and downlink performance.
2. Timestamps are to be logged by the two nodes themselves to determine uplink and downlink performance. A discrepancy can, however, exist between the time stamps taken on either device.

As connecting each of the fifty LCUs to a data acquisition device is not feasible, it was required to take time stamps on the RPi Zero WH themselves. The problem is that timestamps taken on two different devices are not comparable to one another and can be thought of as having different units. Thus no performance statistic calculations could be performed with these timestamps. This is because each device can either lag or lead the actual time, which will produce inaccuracy in the presented data. However, it is possible to measure time durations, rather than timestamps, on each device for performance statistic calculations. Figure 5.8 shows a simplified illustration of a message sent from device 2 to device 1. At the start of the process, device 2 leads device 1 by 10 seconds as there is inaccuracy in the timekeeping procedure. If timestamps are used to determine the uplink (device 2 to device 1) network performance, the procedure will follow:

1. T_1 will be logged as 12:00:11.
2. T_2 will be logged as 12:00:06.
3. The uplink latency will be calculated as -5 seconds ($T_2 - T_1$).

This procedure is inaccurate and will even lead to impossible readings i.e. negative latency measurements as a message can not be received before it is sent. It is thus essential to determine latencies via time duration measurements if two different devices are used. The procedure will then follow:

1. T_1 will be logged as $\Delta = 1$ s.
2. T_2 will be logged as $\Delta = 6$ s.
3. The uplink latency will be calculated as 5 seconds ($T_2 - T_1$), which is correct.




Figure 5.8: Analysis of taking a message timestamp versus duration on two different devices.

5.3.1 Offset procedure

As can be seen from Figure 5.8 it is not essential that every device's clock is synchronized when starting a test, as duration measurements counter the clock deviations. However, it is essential to synchronize the start of offset measurement, i.e., the beginning of the test in Figure 5.8. If both durations are not measured from the same reference in time, it is impossible to ensure accurate performance measurement.

Within the scope of the network tests, it is impossible to accurately start each LCU and server process at the exact same time. An offset procedure is thus incorporated into the testing process to synchronize the starting reference point. Once all the LCU devices are associated on the network, the server and LCUs start a high precision timer (`time.perf_counter` in Python 3) to measure time durations. The server cycles through all connected LCUs to determine their offset individually. An offset cycle can be seen in Figure 5.9. The procedure for a single offset procedure follows:

- The server reads a duration timestamp, T_1 , and immediately sends it to the LCU.
- The LCU receives the message and logs the duration timestamp that the message is received, T_2 .
- The LCU logs the duration timestamp T_3 and replies to the server with the duration time stamps T_2 and T_3 .
- The server received the message and logs the duration time stamp T_4 .

It is now clear that the server has four timestamps to analyze. Two timestamps, T_1 and T_4 , are taken on the server itself, and two timestamps, T_2 and T_3 , are taken on the LCU, which does not have the same starting reference point as the server. The two different sets of time duration stamps should then be treated as if they have different units. By assuming that the uplink latency is equal to the downlink latency ($T_{12} = T_{34}$), the offset (Δt) can then be determined :

$$T_2 = T_1 + T_{12} + \Delta t \quad (5.1)$$

$$T_3 = T_4 - T_{34} + \Delta t \quad (5.2)$$

$$\Delta t = 0.5 [(T_2 - T_1) + (T_3 - T_4)] \quad (5.3)$$

The offset sequence in Figure 5.9 is run 100 times for each LCU, with a positive offset value indicating that the LCU time is leading the server time. The average offset for all results within one standard deviation is saved on the server. This saved offset can then be used to relate time measurements on the LCU to those taken on the server, allowing uplink and downlink network performance calculation.




Figure 5.9: Offset process messaging sequence

5.3.2 Time drift

During the initial testing phase, it was observed that the system clock speeds of the Raspberry Pi and the server drifted quite significantly. Figure 5.10 depicts the latency measurement for a server connected to five LCUs. It can be seen that the offset procedure successfully synchronizes the clocks at the start of the test, but due to the drift in the clocks, the latency measurement also drifts in time. This is due to the accuracy of the system clock. The most extreme drift recorded during the initial testing phase amounted to 50 ms drift over a 10 min period. This is equivalent to an accuracy of 83 ppm¹ which will result in a 7.2 s drift over a 24-hour period.




Figure 5.10: Latency measurements for five LCUs connected to the server. The latency measurements drift due to the clocks of the various processors drifting.

¹ $(0.05 \text{ s} / (10 \times 60)) \times 1\text{e}6 = 83.333 \text{ ppm}$

This drift in time renders it impossible to record accurate timestamps for network performance calculations. It was required to synchronize the clock speeds to one another to allow for more accurate time measurements. The Network Time Protocol (NTP) was used to allow clock synchronization. NTP uses the same basic offset calculations as expressed in Equations 5.1 - 5.3, except with the LCU polling the server for time. Time offsets are taken periodically depending on the stability of the system clock (every 64 seconds before accurate synchronization and up to every 1024 seconds if accurate synchronization is achieved). The NTP process adjusts the clock speed of the LCU after every offset polling using one of two methods:

1. *Stepping*: If the LCU system clock is offset from the server by more than 128 ms, the NTP process steps the LCU clock. During this process, the time is adjusted on the system clock by resetting the system time (on the LCU) to the server time.
2. *Slewing*: If the LCU system clock is below the step threshold, the NTP process will slew the LCU clock. During this process, the clock speed of the LCU is adjusted (to a maximum of 500 ppm) to meet that of the server. These time adjustments are so minute that it does not interfere with applications running on the system.

Using this process, the NTP can adjust the time on the LCU to within a millisecond on very stable local area networks. This limits the clock drift in a network and allows more accurate time measurements. The ThinkPad laptop used in the network test (also running the server and broker processes) is set up as an offline stratum 4 server to which the LCUs can connect. The laptop is not synchronized to the internet to avoid time updates to the laptop, which will require all LCU devices to restart the synchronization process. It was observed that to achieve a synchronization accuracy of under 2 ms, the LCUs should be allowed to synchronize for at least 90 minutes. Figure 5.11 shows a latency measurement with five LCUs that are synchronized via NTP.




Figure 5.11: Latency measurement with NTP protocol enabled. No drift in latency is present as the clock speeds of the LCUs are updated to match the server.

Chapter 6

Experimental Setup

Wireless communication has shown the potential to reduce capital costs and construction time in CST plants. However, the implementation of wireless communication standards in CST plants is still in the development phase, and limited research is available describing the expected practical performance in heliostat fields. Section 4.1.2 specified Wi-Fi 4 as the most suitable current technology available to implement in the next generation H2020 PREMA CST plant. This study forms part of research in application which aims to analyze the performance and feasibility of Wi-Fi 4, and thus wireless communication, in CST plants. The main purpose of the field tests is to measure performance metrics to provide future plant designers with a better understanding of practical performance expectations.

The Helio100 CST site, outside the city of Stellenbosch in the Western Cape, was the chosen site to perform the network performance tests. Prior to testing, the surrounding environment was analyzed to ensure that no external signals were present in the 2.4 Ghz frequency to avoid RFI from transmitting devices outside the scope of this experiment. Field tests were conducted by placing LCUs in different configurations to determine wireless communication's suitability for a heliostat field. Each test condition was chosen to analyze the influence of various impediments on the network's performance, as described in Section 2.3. This chapter describes each of the test environments with the expected performance limiting factors in each scenario.

6.1 Test environments

Eight main test scenarios were considered to determine the performance of wireless communication in a heliostat field. A summary of all eight tests can be seen in Table 6.1. The main factors that are changed between tests include:

- The number of LCUs that are part of the test.
- The transmission range between the LCUs and AP.
- Whether LoS is present between the LCUs and the AP.

- Whether the LCU is placed within the heliostat field or not.

Table 6.1: Test environment overview. This table provides the capacity, range, placement and reasoning for each of the 10 field tests.

Test name	# LCUs	Range [m]	Placement	Reason
Baseline	50	10-35	Open field with LoS	Interference and noise free environment to compare all other tests to
Capacity	10 and 30	10-35	Open field with LoS	Analyse the influence of network capacity and RFI
Range Single LCU	1	60/110	Open field with LoS	Analyse the influence of signal attenuation over large distances
Range Full Capacity	50	85-110	Open field with LoS	Analyse the influence of signal attenuation and RFI in a low RSS environment
Heliostat LoS	50	45-80	Heliostat field with LoS	Analyse the influence of signal reflection, EMI and RFI
Heliostat non-LoS	50	45-80	Heliostat field with no LoS	Analyse the influence of shadowing effects, EMI and RFI
Heliostat High Density	50	80	Heliostat field with LoS	Analyse the influence of dense networks in a heliostat field
Heliostat Low Density	1	1	Heliostat field with LoS	Analyse the influence of space networks in a heliostat field

6.1.1 Baseline test

A reliable control was needed to compare different tests; therefore, a baseline test was conducted. Interference in the baseline test was limited as much as possible to create a reliable reference. Fifty LCUs were placed between 10 m and 35 m from the AP, which is well within the range of the 802.11n standard. The AP was located at a height of 6 m on the control container to ensure LoS to each LCU. The LCUs were arranged approximately 3 m apart from one another. The LCUs were placed in an open field; thus, there was no interference or shadowing caused by nearby steel structures. Figure 6.1a shows an LCU powered by an 8000 mAh solar power bank in an open field, and Figure 6.1b depicts the general placement of the LCUs in an open field during the baseline tests.



(a) LCU in field placement example.



(b) LCU Baseline test placement.

Figure 6.1: Placement of the LCUs for the baseline test. LCUs were placed in an noise free open field with LoS to the AP at a range of 10-35 m.

6.1.2 Capacity test

It was required to adjust the network's capacity (number of LCUs) and analyze the effect on network performance. To accomplish this, tests with a capacity of 10 LCUs and 30 LCUs were conducted and compared to the 50 LCU baseline test. The LCUs were placed in the same field and at the same distance as the baseline test. Thus the only differentiating factor between the tests was the capacity of the network. Once again, all LCUs had direct LoS to the AP with no shadowing effects or influence from steel structures. It was expected that an increase in the number of LCUs participating in the network would increase network traffic. With more nodes contending for the medium, an increase in RFI within the same BSS was expected, which will increase latencies and lead to a higher amount of dropped messages. The performance of the network was thus expected to degrade with an increase in capacity.

6.1.3 Range tests

First, the range tests were conducted with only a single LCU participating in the network tests. The range measurements were taken with a Leica Disto D510 outdoor laser distance meter, which has a range of 200 m and a typical accuracy of approximately 1 mm. The single LCU was placed, with LoS to the AP, at a range of 60 m and again at 110 m. LCU was placed in the open field, similar to the baseline test, with no shadowing effects or influence from nearby steel structures. The Friis free space signal attenuation model, as shown in Figure 2.5, describes a loss of received power as the transmission distance is increased. The transmission range is limited by the power of the transmitter and the sensitivity of the receiver. A shorter range than that specified by the 802.11n standard was expected as the RPi Zero WH antenna is etched on to the board itself, producing a lower transmitter output than an external antenna could.

A full capacity range test was conducted by placing all 50 LCUs in the same open field at a range between 85 m and 110 m. It was expected that the signal attenuation would result in a lower LCU RSS, which will lead to lower network performance. During the full capacity range tests, the increase of network capacity was expected to introduce further network contention and RFI, which will compound the network performance degradation.

6.1.4 Heliostat LoS tests

All 50 LCUs were placed on the frames of the heliostats throughout the existing 17 HelioPod field, as seen in Figure 6.2. The placement of the LCUs allowed LoS to the AP at ranges from 45 m to 80 m. The arrangement of the LCUs is denser than will be expected in a fully operational field. In the next generation plant, a single LCU will be placed per HelioPod, but the limited HelioPods (17) available at the Helio100 site, resulted in three LCUs placed on a single HelioPod. The heliostat field will introduce multipath effects as the received signals are reflected off nearby

heliostats. Furthermore, the linear actuators driven by brushed DC motors will result in EMI. These factors, coupled with the RFI from the LCUs near one another, could significantly hinder performance compared to the baseline test.



(a) LCU LoS placement on heliostat frame. (b) LCU LoS placement relative to AP located on top of the control tower.

Figure 6.2: Heliostat LoS test LCU placement. LCUs fitted on the frame of the heliostat with velcro ensuring LoS to the AP located on top of the control container.

6.1.5 Heliostat non-LoS tests

The LCUs were placed behind the heliostats (relative to the AP), as seen in Figure 6.3, to analyze the shadowing effect of the heliostats themselves. There was thus no LoS between the AP and the LCUs. The nodes were placed at the same range and density as the heliostat LoS test. Thus, the only differentiating factor between the heliostat LoS and heliostat non-LoS tests was the heliostat facets' shadowing effects. The shadowing effect was expected to result in a significant loss in RSS and lead to an increase in dropped messages and retransmissions. The single antenna of the RPi Zero WH was also expected to be hampered by the multipath effects in the heliostat field, especially due to the low RSS caused by shadowing effects.



(a) Heliostat non-LoS LCU placement top view. (b) Heliostat non-LoS LCU placement bottom view.

Figure 6.3: Heliostat non-LoS test LCU placement. LCUs fitted below the frame of the heliostat with velcro blocking LoS to the AP located on top of the control container.

6.1.6 Heliostat density tests

The effect of network density in a heliostat field was analysed. Two different cases were considered:

1. The dense placement of all 50 LCUs in the heliostat field.
2. A single LCU placed in the heliostat field.

For the dense placement, all LCUs were placed on the furthest four HelioPods. Approximately two LCUs were placed per heliostat facet resulting in 12 LCUs placed per HelioPod, as shown in Figure 6.4. The LCUs had an average range to the AP of approximately 80 m and were placed on top of the facets to ensure LoS with the AP.

A single LCU was also placed in the heliostat field with LoS to the AP at a range of approximately 80 m. This test, along with the sparser deployment of the LCUs in the heliostat LoS test and the baseline test, will be compared to the results of the dense placement of the LCUs. It was expected that the dense placement would result in a higher number of dropped messages and lower network throughput due to a high level of RFI, multipath effects, and higher network capacity.



Figure 6.4: Dense network LCU deployment in heliostat field. Two heliostats were placed on each heliostat facet to total 12 LCUs per HelioPod.

6.2 Test parameters

The test parameters can be seen in Table 6.2. It can be noted that the downlink payload size exceeds the 104 byte size as specified by Table 3.1. In real-world operation, processes such as calibration of the heliostat and manual inputs would require additional information to be sent to the heliostat. This could include an exact aimpoint, an update to the heliostat's error parameters, and could even include

over-the-air firmware updates. Provision is made for such processes by allocating an additional 88 bytes in the downlink message. This equates to a 192 byte downlink message and a 30 byte uplink message.

Furthermore, it can be seen that a message is dropped and retransmitted on the ZeroMQ level if it is not successfully transmitted in 50 ms. A maximum retransmission limit of three is also specified to ensure the system does not utilize all the resources trying to deliver a message to a faulty LCU. If the retransmission limit of three is reached, the LCU is overlooked for the current polling cycle and added to the back of the queue.

Table 6.2: Test parameters and AP settings for field tests.

Test parameter	Value
Test time	600 s
# Offset loops per LCU	100
Uplink payload size	30 Bytes
Downlink payload size	192 Bytes
Message drop timeout	50 ms
Retransmission limit	3
Access point settings	
Frequency range	2451 MHz - 2473 MHz
Wi-Fi channel	11
Tx power	20 dBm
Channel width	20 MHz
RTS threshold	1

Within the AP settings, it can be seen that a 20 MHz channel width is chosen. In the next generation plant, the cluster-tree topology, specified in Table 4.1, would require multiple BSSs to operate within the carrier sense range. Thus utilizing 20 MHz channels would allow 3 APs to operate in non-overlapping channels in the 2.4 GHz frequency. Furthermore, the Request To Send (RTS) threshold is set to 1, which will force an RTS/CTS process. This process adds more overhead to the transmission but reduces the number of retransmission on the network.

Chapter 7

Results

Wireless communication tests, using the Wi-Fi 4 standard, were conducted at the Helio100 site, outside Stellenbosch. The field tests were conducted to measure various performance metrics in different experimental scenarios. The measured performance metrics included over-the-air Uplink (UL) and Downlink (DL) latency and throughput, RTT latency, average test throughput, and package error rates as described in Section 2.5. These performance metrics are essential to ensure that the next-generation H2020 PREMA CST can operate with sufficient accuracy and level of control using wireless communication networks. Future plant designers can use the measured performance metrics to understand the practical expectations of wireless performance in heliostat fields.

This chapter starts by describing the initial test results, which led to an investigation into high periodic latencies, the relationship between payload size and throughput, and optimal polling rates. The results of the field tests are then presented with a summary of each experimental scenario outcome. The chapter concludes with a discussion of the statistical significance of the results and observations in the results.

7.1 Initial test results

Before network field tests were conducted, it was essential to understand the basic operation of the ZeroMQ network. Initial tests were conducted with the TP-Link AP and ten LCUs. All the hardware components were placed in the same room, and the network was set to operate in a noise-free channel. The simplified setup allowed test parameters to be altered to investigate the relationship between payload size, latency, and throughput. Furthermore, the initial tests resulted in a practical understanding of the polling rate's influence on the throughput level, and to determine the time it takes to sufficiently synchronize the clock speeds of the server and LCUs with the NTP process.

7.1.1 High latencies during initial testing

Initially, a server process polled each of the ten LCUs synchronously, i.e., the server would send an instruction to an LCU and wait for the reply before moving on to the next LCU. Figure 7.1 depicts the uplink latencies recorded during the test. It can be seen that multiple uplink messages had latencies between 600 ms and 700 ms. The latency pattern seems to have an almost equal 610 ms latency, which suggests a background process, either on the LCU or server, is causing the delay.




Figure 7.1: High uplink latencies found during initial testing with ten LCUs showing a recurring latency around 610 ms.


On the server side, the process was granted a higher priority by changing the nice value of the process to -20 (Ubuntu 18.04 LTS OS). The process was also moved to an isolated core to reduce the risk of background processes being allocated processor time in front of the server process. The LCU only has a single CPU core, and thus OS processes such as garbage collection would be allocated CPU time, which can cause larger latencies in the network.

The LCU process was also allocated a higher priority to minimize background processes interrupting the messaging process. The RPi Pi Zero WH also has built-in power-saving features that suspend Wi-Fi connectivity periodically, which increases the latency due to the time needed to reconnect to the network. This feature was disabled on startup, and combined with the higher priorities of the processes and isolated cores, the 610 ms latency delay was eliminated.


7.1.2 Relationship between payload size, latency and throughput

The payload of a message is the data portion of the message transmitted from the sender to the receiver and excludes overhead data such as CRC, routing addresses, and headers. As the payload size is increased, a larger amount of data needs to be transferred over-the-air. The increase in the amount of data to be transferred will result in longer latencies per message. As latency and throughput are inversely related, a larger latency tends to lead to lower throughput. However, as the message overhead is constant, an increasing payload size results in a larger portion of the message sent over-the-air to be regarded as throughput, and thus a larger payload will result in higher throughput.

It was expected that the payload increase would have a more significant effect on the throughput than the latency increase. A network test was designed to increase the payload size periodically and measure the resulting latencies and throughput to investigate this hypothesis. The payload size was increased at a rate of 10 bytes per second for both the UL and DL messages. The network test was designed using five LCUs. Figure 7.2 shows the latency and throughput moving average plot of the last 100 messages for the DL communication.



(a) Downlink latency moving average plot of the last 100 messages.



(b) Downlink throughput moving average plot of the last 100 messages.

Figure 7.2: Moving average plot of the last 100 messages showing the relationship between payload size, latency and throughput. Five LCUs were used in the test.

It can be seen that the increase in payload size results in higher latencies with a near 300 % increase in average latencies over the test. It can also be seen that the payload size increase resulted in higher throughput despite the increase in latencies, thus confirming the hypothesis. It can be noted that at around 150 s into the test, a significant increase in latency resulted in reduced throughput. When examining this phenomenon, it was found that at the 150 s mark, the packet size is approximately 1500 bytes as it was increased at a rate of 10 bytes per second.

The ZeroMQ protocol in this test uses the TCP transport layer, IPV4 internet layer, and an ethernet network access layer, as described in Figure 2.7. The ethernet

protocol specifies a Maximum Transmission Unit (MTU) of 1518 bytes (including the 14-byte Ethernet II header and a 4-byte frame check sequence). The MTU represents the largest data packet that a network-connected device will accept, i.e., the maximum amount of data available to the IP, TCP, and ZeroMQ protocols to be sent in a single frame. If a packet exceeds the MTU limit, it is fragmented into two or more packets. Each fragmented packet's IP header includes information that the original packet is fragmented and requires reassembly, how many fragmented packets there are, and in what order they are sent.

Thus, this spike in latency is due to the fragmentation of the original packet, which requires multiple packets to be sent over-the-air. A small spike in latency can be seen at 300 s, which correlates to the fragmentation from two to three packets. The effect is less prominent at 450 s, which correlates to a 4.5 MB payload size; however, this message size would be uncommon in a heliostat field. Thus, it is apparent that the payload size must be kept under the MTU limit to ensure low latency, high throughput network communication.

7.1.3 Investigation into optimal polling rates

The ZeroMQ polling function is used to monitor incoming messages from multiple sockets. The poll function is used within the server and broker process. The poller on the broker process reads messages from the backend (LCUs) and frontend (Server) and relays the message to the appropriate destination. As the broker does not generate new commands, the poller is designed to continuously listen for messages and execute the relay when a message is detected. When a message is detected, the process will thus exit the polling process, relay the message and return to the polling process.

The server polls the broker for messages on multiple sockets. If no messages are available to read within a set time, called the polling rate, a new instruction is generated and sent to the next LCU in the list. The poll timer is reset when the new instruction is sent over the air. Thus, the polling rate of the server influences the number of messages processed by the LCUs in parallel. Three different strategies for polling on the server process were investigated, and the uplink latency results of the three strategies can be seen in Figure 7.3. The polling tests were conducted with five LCUs. Initially, the poller was set up to monitor the sockets for the entire duration of the polling rate, creating a queue of all incoming messages to be processed after the poll rate expired. This approach was tested with both a long (10 ms) and short (1 ms) poll rate.

Later, an optimized approach similar to the strategy on the broker process was implemented by setting the poll rate to 2 ms. If a message is detected within the 2 ms time frame, the polling process will terminate, and the message is processed as it is detected. Thus, the polling process will exit before the 2 ms limit if a message is detected. A new instruction is sent if no message was detected during the 2 ms poll rate.

It can be seen that the initial approaches can lead to a large uplink latency as

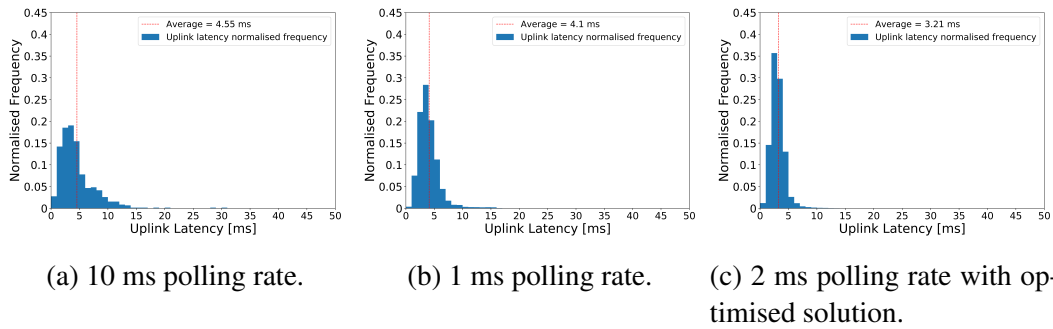


Figure 7.3: Uplink latency of initial tests, including five LCUs, with a varying polling strategy. The polling strategies included a long, short, and optimized design. The optimized solution outperforms both the initial long and short polling rates.

messages are kept in queues until the polling time limit is exceeded. For the 10 ms polling interval, messages are kept in the queue for the most prolonged period, thus resulting in the longest average uplink latency (4.55 ms) with uplink latency peaks exceeding 15 ms. The shorter 1 ms polling rate resulted in lower latencies compared to the 10 ms poll, with fewer peaks. The optimized solution proved to result in the lowest average uplink latencies as the polling procedure exited instead of waiting for the timer to exceed the limit. This allowed messages to be processed as they were detected by the poller and avoid queuing.

The 10 ms poll rate resulted in approximately 43 000 round trip messages delivered in 600 s, whereas both the other strategies resulted in message numbers exceeding 150 000. This is because if the poll rate is 10 ms, a new instruction will only be sent to an LCU every 10 ms if the poller does not detect any incoming messages. If all messages thus had a sub-10 ms RTT, this will only result in one completed instruction per 10 ms or an upper limit of 60 000 messages in the 600 s test duration. An optimized 2 ms polling rate was thus used for further testing.

7.2 Field tests

After the initial tests were completed successfully, the LCUs were placed in the field, as described in Chapter 6. For each test environment, a minimum of six tests were conducted to produce 60 min worth of data per test environment. The data was analyzed to determine the test performance parameters such as RTTs, UL and DL latencies, over the air throughput (TP_{ota}), average test throughput (TP_{avg}), and message drop percentage as described in Section 2.5. This section will provide more information on the performance results of each different field test.

7.2.1 Baseline test

As described in Section 6.1.1, the baseline test was designed to be a reference point to which all other test results can be compared. The LCUs were placed in an open

field at a range of 10 m - 35 m with direct LoS to the AP. The latency results of the baseline latency tests can be seen in Table 7.1, which presents the UL and DL latency mean and standard deviation (SD), as well as the RTT latency mean and SD. The throughput results of the baseline latency tests can be seen in Table 7.2, which presents the UL and DL over-the-air throughput mean and SD, as well as the average test UL and DL throughput.

Table 7.1: Baseline test latency result.

Testname	# Nodes	# Messages	% Dropped	UL latency		DL latency		RTT latency	
				Mean	SD	Mean	SD	Mean	SD
Baseline	50	48470.4	0.82	2.28 ms	2.11 ms	7.61 ms	2.89 ms	11.45 ms	3.86 ms

Table 7.2: Baseline test throughput result.

Testname	# Nodes	# Messages	% Dropped	UL TP_{ota}		DL TP_{ota}		TP_{avg}	
				Mean	SD	Mean	SD	UL	DL
Baseline	50	48470.4	0.82	174.6 kbps	65.5 kbps	229.7 kbps	63.4 kbps	25.7 kbps	130.4 kbps

Figure 7.4 shows the UL and DL latency results for one of the baseline tests. Latency is measured per message and plotted midway between the sent and received logged times on the x-axis, as seen in Figure 7.1. This representation of the data allows the user to analyze patterns in the data but renders a noisy presentation of the data if no apparent pattern is evident.




Figure 7.4: Baseline test UL and DL latencies. Latencies plotted per message for the duration of the test.

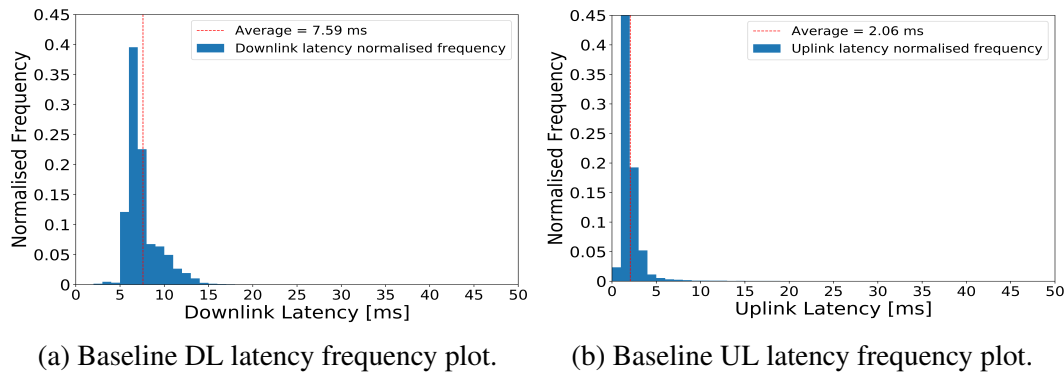



Figure 7.5: Baseline UL and DL latency. Latencies plotted on a normalised histogram with latencies grouped in 1 ms bins.

Figure 7.5 shows a histogram plot of the same data. Each x-axis bin spans a width of 1 ms, and the normalised frequency is plotted on the y-axis. This representation allows the user to analyze the data's spread and can easily be used to compare tests in various scenarios.

When analyzing Figure 7.5, it is evident that the DL latencies far exceed the UL latencies, with the majority of the DL latencies between 5 ms and 10 ms compared and the UL latencies between 2 ms and 5 ms. This can be due to multiple reasons, such as:

- The DL payload size, excluding the message ID, is 6.4 times the size of the UL payload size, and it is evident from Figure 7.2 that latency increases with an increase in payload size.
- The DL messages have to be routed to the correct LCU, which creates two problems. Firstly, routing overhead is added to the packet to ensure delivery to the correct LCU, which increases the message size. Secondly, the routing algorithm needs to match an LCU's ID to a routing table destination, which could take valuable processing time, adding to the latency of the message.
- As stated in Section 5.3.1, an assumption is made during the offset procedure that the uplink and downlink latencies are equal. This assumption is made as there is no practical way to measure the difference in these latencies, without uncertainties, except by connecting both the server and LCU to a data acquisition device that would not be practical in this project's scope.

The UL and DL over-the-air throughput results can be seen in Figure 7.6. Thus, these graphs represent the data transfer speed from the point in time where a message was sent until it is received. It can be noted that the DL throughput is significantly larger than the UL throughput despite the higher DL latencies compared to the UL latencies. Referring to Table 3.1, the DL payload size, excluding the message ID, is 192 bytes compared to the 30 byte UL payload. As seen in Figure 7.2, an increase in payload size leads to an increase in throughput, which explains the higher DL throughput found in the test.



(a) Baseline DL throughput frequency plot. (b) Baseline UL throughput frequency plot.

Figure 7.6: Baseline UL and DL throughput. Throughput plotted on a normalised histogram with throughput grouped in bins of 10 kbps.

The RTT latency for the baseline test can be seen in Figure 7.7. As the RTT latency is measured on the same machine, it is the most accurate measurement to compare tests in different environments. The RTT latency includes the message encoding and decoding times, over-the-air transfer times (UL and DL), and the message processing time on the LCU. The average RTT latency for the baseline tests was measured at 11.45 ms, with the majority of the messages having an RTT latency of 9 ms to 15 ms. There is a presence of messages with higher than 15 ms RTT latencies, although this does not equate to more than five percent of the messages. An 11.45 ms RTT latency will result in 87 messages per second provided no processing time between messages and no dropped messages ¹.




Figure 7.7: Baseline RTT latency. RTT latencies plotted on a normalised histogram with latencies grouped in 1 ms bins.

7.2.2 Capacity test

As described in Section 6.1.2, the capacity tests were designed to measure the network's performance with a variable amount of LCUs participating in the test. The LCUs were placed in an open field at a range of 10 m to 35 m with direct LoS to the AP. The capacity is thus the only differentiating factor when compared to the

¹ $\frac{1}{11.45 \text{ [ms]}} \approx 87 \text{ messages/s}$

baseline test. The tests were run for 10 LCUs and 30 LCUs participating in the tests, and the latency results compared to the baseline test can be seen in Table 7.3 and the throughput results can be seen in Table 7.4.

Table 7.3: Latency test results comparing the 10 LCU network capacity, 30 LCU network capacity, and baseline test.


Testname	# Nodes	# Messages	% Dropped	UL latency		DL latency		RTT latency	
				Mean	SD	Mean	SD	Mean	SD
Capacity 10	10	69752.9	0.015	2.31 ms	1.20 ms	7.24 ms	1.31 ms	10.09 ms	1.72 ms
Capacity 30	30	56176.0	0.4	2.52 ms	2.35 ms	7.11 ms	1.82 ms	10.51 ms	2.67 ms
Baseline	50	48470.4	0.82	2.28 ms	2.11 ms	7.61 ms	2.89 ms	11.45 ms	3.86 ms

Table 7.4: Throughput test results comparing the 10 LCU network capacity, 30 LCU network capacity, and baseline test.

Testname	# Nodes	# Messages	% Dropped	UL TP_{ota}		DL TP_{ota}		TP_{avg}	
				Mean	SD	Mean	SD	UL	DL
Capacity 10	10	69752.9	0.015	165.8 kbps	32.1 kbps	257.6 kbps	35.9 kbps	27.9 kbps	178.6 kbps
Capacity 30	30	56176.0	0.4	173.6 kbps	47.2 kbps	233.3 kbps	41.93 kbps	26.5 kbps	143.8 kbps
Baseline	50	48470.4	0.82	174.6 kbps	65.5 kbps	229.7 kbps	63.4 kbps	25.7 kbps	130.4 kbps

The latency and over-the-air throughput did not vary significantly between the three tests. However, it is interesting to note that both capacity tests have a higher UL latency and a lower DL latency compared to the baseline. Even though the variation is slight, it can be attributed to the uncertainties during the offset procedure. The major difference between the tests is the packet error ratio. An increase in capacity leads to a higher percentage of dropped messages. This coincides with the hypothesis expressed in Section 6.1.2, where an increase in RFI due to more of LCUs participating in the network test would result in more network contention and lead to dropped messages. A lower dropped message statistic will result in fewer retransmissions, and thus a higher message delivery rate. This influences the number of messages delivered successfully throughout the 600 s test, which affects the average test throughput.

The RTT latencies for the capacity tests and baseline are shown in Figure 7.8. The plots do have similar shapes and are comparable to the baseline RTT results. It can be seen that the capacity of a network and the RTT latencies on the network are directly related. Compared to the 10 LCU capacity test's RTT latency, the 30 LCU capacity test had an RTT latency increase of 4.16 %, and the baseline test around 13.48 %. It can thus be inferred that the capacity of the network has a critical influence on network performance.



(a) Capacity 10 RTT latency. (b) Capacity 30 RTT latency. (c) Baseline RTT latency.

Figure 7.8: RTT latency for a network capacity of 10, 30, and 50 (baseline) respectively. RTT latencies plotted on a normalised histogram with latencies grouped in 1 ms bins.

7.2.3 Range tests

As noted in Section 6.1.3, range tests were conducted with both a single LCU and at full capacity (50 LCUs). The nodes were again placed in an open field with direct LoS to the AP. During the single node tests, the LCUs were placed at a range of 60 m and 110 m.

The full capacity range tests were conducted with LCUs ranged between 80 m and 110 m. The latency results of the range tests compared to the baseline tests can be seen in Table 7.5 and the throughput results can be seen in Table 7.6.

Table 7.5: Latency test results comparing the single LCU at 60 m range, LCU at 110 m range, full capacity range test and baseline test.

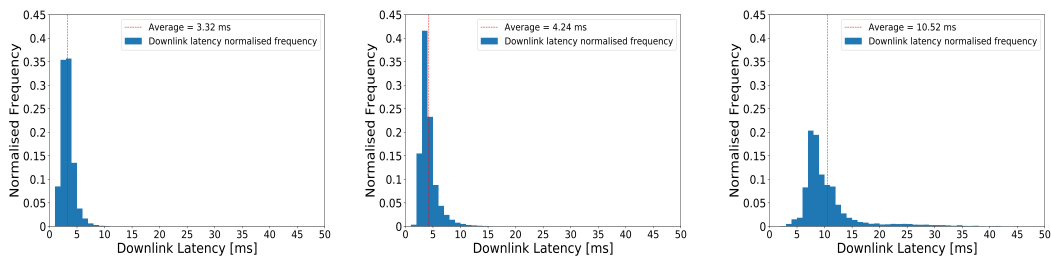
Testname	# Nodes	# Messages	% Dropped	UL latency		DL latency		RTT latency	
				Mean	SD	Mean	SD	Mean	SD
Single 60 m	1	75834.1	0.13	2.45 ms	1.46 ms	3.32 ms	1.32 ms	7.60 ms	1.96 ms
Single 110 m	1	34863.4	1.54	3.68 ms	3.3 ms	4.24 ms	1.71 ms	9.47 ms	3.85 ms
Full Cap range	50	11349.6	9.29	4.09 ms	16.0 ms	11.06 ms	26.5 ms	15.71 ms	8.05 ms
Baseline	50	48470.4	0.82	2.28 ms	2.11 ms	7.61 ms	2.89 ms	11.45 ms	3.86 ms

Table 7.6: Throughput test results comparing the single LCU at 60 m range, LCU at 110 m range, full capacity range test and baseline test.

Testname	# Nodes	# Messages	% Dropped	UL TP_{ota}		DL TP_{ota}		TP_{avg}	
				Mean	SD	Mean	SD	UL	DL
Single 60 m	1	75834.1	0.13	176.2 kbps	43.11 kbps	543.6 kbps	180.2 kbps	40.3 kbps	204.1 kbps
Single 110 m	1	34863.4	1.54	157.7 kbps	39.6 kbps	423.0 kbps	121.8 kbps	18.4 kbps	93.8 kbps
Full Cap range	50	11349.6	9.29	147.6 kbps	84.9 kbps	176.4 kbps	72.8 kbps	5.93 kbps	30.4 kbps
Baseline	50	48470.4	0.82	174.6 kbps	65.5 kbps	229.7 kbps	63.4 kbps	25.7 kbps	130.4 kbps

Firstly it can be seen in Table 7.5 and Figure 7.9 that the DL latencies for the single node tests are significantly lower than higher capacity tests. This lead to a notable increase in the DL throughput. This phenomenon can be due to:

- The routing algorithm, as less processing for routing is needed if only a single node is connected to the socket.
- Fewer LCUs are active on the network, which lowers the medium contention and, thus, the number of collisions on the network.
- The AP hardware specifications, as a higher network capacity can result in the saturation of the AP's RAM or processing power.



(a) 60 m range DL latency. (b) 110 m range DL latency. (c) Full capacity range test DL latency.

Figure 7.9: DL latency for the 60 m single LCU range test, 110 m single LCU range test and the 85 m-110 m full capacity (50 LCU) range test. DL latencies are plotted on a normalised histogram with latencies grouped in 1 ms bins.

Secondly, when comparing the two single node tests, the increase in range led to a higher percentage of dropped messages. This is due to signal attenuation resulting in lower RSS leading to a failed CRC. The retransmission of the dropped messages led to a significant drop in message throughput as the 60 m test yielded 117 % more successfully delivered messages during the span of the test compared to the 110 m test.

Lastly it can be seen that the full capacity range test resulted in remarkably high latencies and a decrease in throughput compared to the baseline test. It is also evident from Table 7.5 that the full capacity range test data has large UL and DL latency standard deviations. This indicates that LCUs do not have a stable connection to the AP, and multiple high latency messages are sent. The increase in medium contention and thus RFI coupled with the signal attenuation, played a substantial role in the high message drop rate, which resulted in a lower message throughput. The lack of an external high power antenna on the RPi Pi Zero WH could have contributed to the difficulty of transmitting over large distances through a noisy medium.

When analyzing the RTT latencies for the different range tests, seen in Figure 7.10, it is apparent that range and RTT latency have a direct relationship. Furthermore, it can be noted that the introduction of RFI due to an increase in LCUs on the network results in deteriorating network performance as a significant amount of RTT latencies greater than 15 ms are apparent in the full capacity range test. It can thus be concluded that the network performance would deteriorate beyond effectiveness at transmission ranges exceeding 110 m.

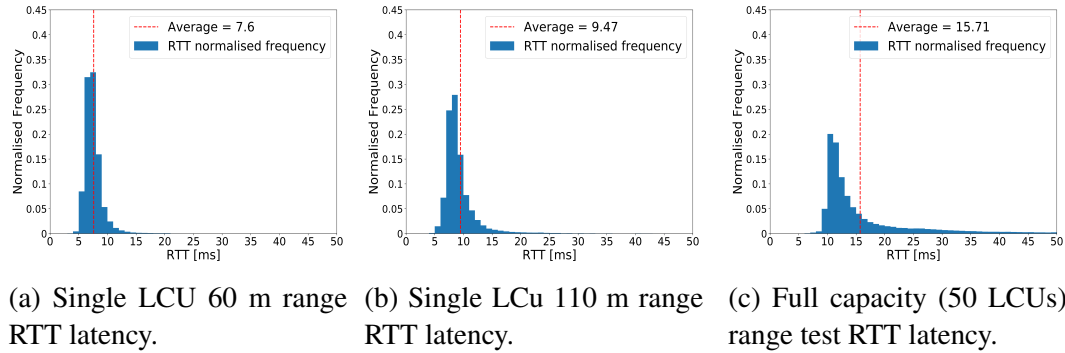


Figure 7.10: RTT latency of the single LCU 60 m and 110 m range test and the 85 m-110 m full capacity (50 LCUs) range test. RTT latencies are plotted on a normalised histogram with latencies grouped in 1 ms bins.

7.2.4 Heliostat LoS/non-LoS tests

Two different heliostat tests were conducted: one where all LCUs have LoS to the AP as described by Section 6.1.4 and another where the LCUs' LoS is obstructed by the mirror facet as described by Section 6.1.5. The LCUs were placed on the heliostat at the Helio100 plant at ranges varying from 45 m to 80 m from the AP. The latency results of the heliostat tests compared to the baseline can be seen in Table 7.7, and the throughput results can be seen in Table 7.8.

Table 7.7: Latency test results comparing the LoS heliostat test, non-LoS heliostat test and baseline test.

Testname	# Nodes	# Messages	% Dropped	UL latency		DL latency		RTT latency	
				Mean	SD	Mean	SD	Mean	SD
Heliostat LoS	50	47530.6	0.81	2.07 ms	1.7 ms	7.91 ms	3.1 ms	11.29 ms	3.55 ms
Heliostat non-LoS	50	39086.2	1.65	2.62 ms	2.65 ms	8.65 ms	4.18 ms	12.82 ms	5.1 ms
Baseline	50	48470.4	0.82	2.28 ms	2.11 ms	7.61 ms	2.89 ms	11.45 ms	3.86 ms

Table 7.8: Throughput test results comparing the LoS heliostat test, non-LoS heliostat test and baseline test.

Testname	# Nodes	# Messages	% Dropped	UL TP_{ota}		DL TP_{ota}		TP_{avg}	
				Mean	SD	Mean	SD	UL	DL
Heliostat LoS	50	47530.6	0.81	163.5 kbps	81.9 kbps	220.1 kbps	42.7 kbps	23.5 kbps	119.5 kbps
Heliostat non-LoS	50	39086.2	1.65	165.3 kbps	68.9 kbps	207.2 kbps	51.3 kbps	20.7 kbps	105.1 kbps
Baseline	50	48470.4	0.82	174.6 kbps	65.5 kbps	229.7 kbps	63.4 kbps	25.7 kbps	130.4 kbps

When comparing the heliostat LoS test to the baseline test, it is evident that both tests yield similar results. Compared to the baseline test, the LoS heliostat test yielded a 1.9 % reduction in the number of round trip messages, 6.4 % lower UL over-the-air throughput, and a 4.2 % reduction in DL over-the-air throughput. Both the baseline and LoS heliostat test had a drop rate of approximately 0.8 %. It can thus be inferred that the multipath effects due to reflection from neighboring metal structures and EMI induced by the brushed motors of the actuators have a marginal effect on the network performance if the LCU has LoS to the AP.

The RTT latency results can be seen in Figure 7.11. The LoS heliostat tests yield an average RTT latency of 11.29 ms, which is lower than that of the baseline. Care must thus be taken to ensure that, as much as possible, an LCU has direct LoS to the AP. While shadowing effects in the non-LoS test may hamper the performance of the network, the results showed an average of 12.82 ms RTT latency equating to 78 round trip messages per second². Thus, it can be concluded that shadowing effects significantly hamper network performance but are not fatal to the functioning of the network.

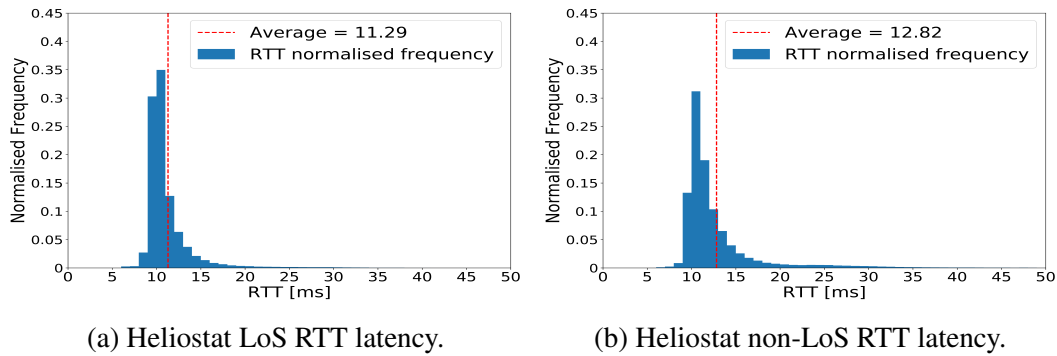


Figure 7.11: RTT latency of the heliostat LoS test and the heliostat non-LoS test. RTT latencies are plotted on a normalised histogram frequency plot with latencies grouped in 1 ms bins.

7.2.5 Heliostat Density test

To determine the influence of dense networks on the network performance in the heliostat field, the heliostat density test was conducted as described in Section 6.1.6.

$$2 \frac{1 \text{ [message]}}{12.82 \text{ [ms]}} \approx 78 \text{ messages/s}$$

The performance of 50 LCUs in an exceptionally dense placement in the heliostat field was compared to the performance of a single LCU placed at the same range in the heliostat field. All LCUs were placed with LoS to the AP at a range of approximately 80 m. The result of the two tests, as well as the more sparsely placed heliostat LoS test result, 60 m single LCU range test, and baseline test results, can be seen in Table 7.9 and Table 7.10.

Table 7.9: Latency test results comparing the dense LoS heliostat test, the single LCU LoS heliostat test, the LoS heliostat test, the 60 m single LCU range test, and the baseline test.

Testname	# Nodes	# Messages	% Dropped	UL latency		DL latency		RTT latency	
				Mean	SD	Mean	SD	Mean	SD
Dense LoS	50	38453.5	1.54	2.93 ms	2.75 ms	8.34 ms	3.38 ms	12.69 ms	4.32 ms
Single LCU heliostat LoS	1	80574.7	0.11	2.49 ms	1.13 ms	2.88 ms	1.24 ms	7.15 ms	1.5 ms
Heliostat LoS	50	47530.6	0.81	2.07 ms	1.7 ms	7.91 ms	3.1 ms	11.29 ms	3.55 ms
Single 60 m	1	75834.1	0.13	2.45 ms	1.46 ms	3.32 ms	1.32 ms	7.60 ms	1.96 ms
Baseline	50	48470.4	0.82	2.28 ms	2.11 ms	7.61 ms	2.89 ms	11.45 ms	3.86 ms

Table 7.10: Throughput test results comparing the dense LoS heliostat test, the single LCU LoS heliostat test, the LoS heliostat test, the 60 m single LCU range test, and the baseline test.

Testname	# Nodes	# Messages	% Dropped	UL TP_{ata}		DL TP_{ata}		TP_{avg}	
				Mean	SD	Mean	SD	UL	DL
Dense LoS	50	38453.5	1.54	144.0 kbps	64.7 kbps	213.2 kbps	81.22 kbps	20.4 kbps	103.4 kbps
Single LCU heliostat LoS	1	80574.7	0.11	187.1 kbps	50.7 kbps	653.8 kbps	193.3 kbps	42.8 kbps	216.9 kbps
Heliostat LoS	50	47530.6	0.81	163.5 kbps	81.9 kbps	220.1 kbps	42.7 kbps	23.5 kbps	119.5 kbps
Heliostat non-LoS	50	39086.2	1.65	165.3 kbps	68.9 kbps	207.2 kbps	51.3 kbps	20.7 kbps	105.1 kbps
Baseline	50	48470.4	0.82	174.6 kbps	65.5 kbps	229.7 kbps	63.4 kbps	25.7 kbps	130.4 kbps

It can once again be noted that the single LCU deployment results in a significant reduction in DL latency and an increase in DL throughput. Interestingly the single LCU heliostat LoS deployment outperforms the 60 m single LCU range test with approximately 6.2 % higher test average UL and DL throughput. This could be due to ground effects, described in Figure 2.6, that resulted in the slight increase of dropped packets and retransmissions. When comparing the dense heliostat test and the more sparse heliostat LoS test, it is evident that the dense deployment resulted in a higher number of dropped messages as 1.54 % of the messages were dropped compared to the 0.11 % drop message rate in the heliostat LoS test. The message drop rate resulted in reduced message throughput and, thus, lower average throughput. The over-the-air latency is higher in the dense deployment as medium contention can result in backoff procedures to avoid collisions, which will increase latency.

The RTT latency for the single LCU placed in the heliostat field, as well as the dense deployment tests, can be seen in Figure 7.12. As expected, it can be seen that the single LCU test outperforms the dense test. It can also be noted that the RTT latency for the single node heliostat deployment was 5.9 % less than the single LCU 60 m range test, which contributes to the higher successful message count seen in Table 7.9.

The dense test resulted in an average RTT latency 12.4 % larger than that measured in the sparser LoS Heliostat test due to increased RFI. The average RTT latency of 12.69 ms equates to approximately 79 messages per second³. As this deployment in the test is an over exaggeration of the density expected in the next generation plant, the results are promising as all 50 LCUs can be polled in 0.64 s⁴.

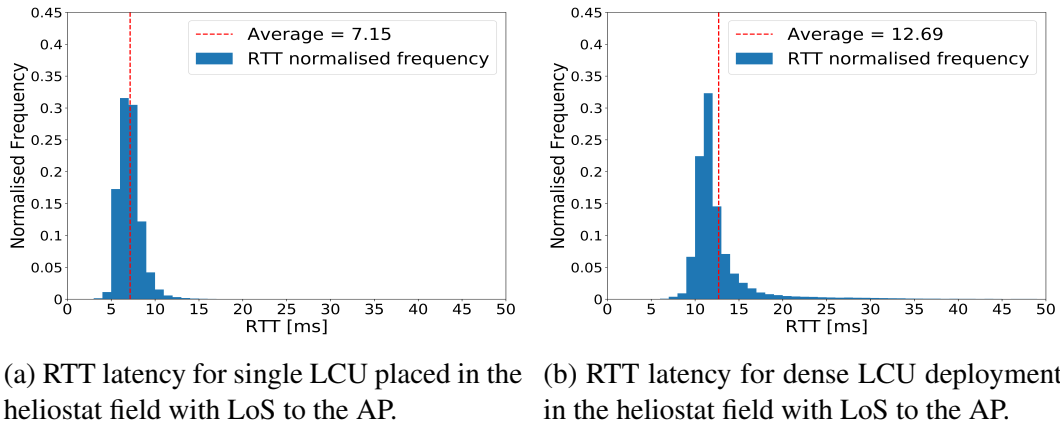


Figure 7.12: RTT latency of the heliostat low density test and the heliostat high density test. RTT latencies are plotted on a normalised histogram with latencies grouped in 1 ms bins.

7.3 Discussion

7.3.1 Statistical significance of resulting data

It was necessary to determine if the data from each test was indeed statistically significant and that the environment of the test was indeed the driving factor for the variation in results and not chance. An independent t -test can be performed to determine the statistical significance of the data. The t -test assumes a Gaussian-like distribution of the data. As can be seen from Figure 7.12, the distribution of the RTT latency tends to have a lagging tail resembling a gamma-like distribution, which is not suitable for the t -test.

$$^3 \frac{1 \text{ [message]}}{12.69 \text{ [ms]}} \approx 78.8 \text{ messages/s}$$

$$^4 50 \text{ [messages]} \times 12.69 \text{ [ms]} = 0.6345 \text{ [s]}$$

According to the central limit theorem, if sufficient samples are taken from a population with mean μ and standard deviation σ , the distribution of the means of the samples will resemble a Gaussian-like distribution [74]. The original data was randomly shuffled, and samples of 100 RTT latency data points were taken without replacement. The mean of each sample was used to determine the new data points, which, according to the central limit theorem, will resemble a Gaussian-like distribution. Figure 7.13 depicts the result obtained from sampling the heliostat non-LoS test, and it can be seen that the data resembles a Gaussian-like distribution.




Figure 7.13: Gaussian distribution of the RTT latency for the heliostat non-LoS test following sampling from the original dataset as described by the central limit theorem.

The normally distributed data allowed t -tests to be performed for each test environment combination. The null hypothesis for the independent t -test is that the population means resulting from the data of two different test environments are equal. Thus, it is required to reject the null hypothesis and accept the alternative hypothesis, which is that the population means are not equal, to prove that the test environments are the drivers for the change seen across different RTT latency tests. Each test's normally distributed data had over 1600 entries, and thus a critical t -value was chosen with infinite degrees of freedom and a significance level, $\alpha = 0.01$. This equated to a critical t -value of 2.576, read from the two tails t -distribution table. The calculated t -value results of all t -tests can be seen in Appendix A. The calculated t -values ranged from 13.2 to 838.8, which is significantly higher than the critical t -value and produced p -values lower than 0.001 %. Thus the null hypothesis can be rejected for each case. This proved that the test data is statistically significant, indicating that the change in the testing environment is the core driving factor for variation between RTT test results and was not up to chance.

A box and whisker plot of the normally distributed data can be seen in Figure 7.14. The plot visualizes the effect of the changing test environment on the RTT latency performance. It can be noted that the results of each test have a normal distribution; thus, the central limit theory holds. The results also show little to no overlap, providing further proof that the data is statistically significant.




Figure 7.14: Box and whisker plots of normalised data to visualise the influence of different testing environments on the RTT latency results.

7.3.2 Observation in results

Following the field tests, a few significant results were noted. It was found that the capacity of the network was linked to the percentage of dropped messages, the maximum feasible transmission range of the network is around 110 m, shadowing effects significantly hamper network performance, and extreme network density can result in up to 12.4 % increase in network latencies.

It was evident that an increase in the number of LCUs participating in a test resulted in a higher number of messages exceeding the 50 ms timeout limit and subsequently dropped. The dropped messages lead to retransmissions, which deteriorates network performance. Messages were dropped if the intended recipient never receives the message, and thus the sender never receives an ACK message, which triggers a retransmission. A message can also be retransmitted on the TCP level if bits were lost or scrambled over the air leading to a CRC failure. The presence of more LCUs on the network introduces RFI and busy channels. When an LCU intends to send a message, a CCA procedure is triggered to determine if a channel is free to transmit; if not, a random backoff counter is triggered. If this process continually repeats itself, the degradation of the network performance can be expected, as can be seen by comparing the capacity test results to the baseline tests in Figure 7.14.

It was noted that the worst-case RTT result, the full capacity range test, produced an average RTT latency of 15.71 ms which equated to 63 messages per second ⁵ or

$$^5 \frac{1 \text{ [message]}}{15.71 \text{ [ms]}} \approx 63 \text{ messages/s}$$

a control interval of 0.79 s for 50 Heliopods⁶. As stated in Section 3.1, the largest next-generation field will require 270 Heliopods, 23 receiver towers, and 16 hours of thermal storage. As the Wi-Fi 4 standard was specified to utilize a cluster-tree topology, as stated in Table 4.1, this will require six clusters of the size used in the tests to service the 270 LCUs in the next generation plant. As there are only three non-overlapping frequency bands in the 2.4 GHz spectrum, this will create RFI between the different clusters or BSSs. To reduce the RFI influence, the following can be assessed:

- The AP's transmission power can be adjusted to lower the carrier sense range of the network, allowing other networks to coexist closer to the transmitter.
- The arrangement of the APs can be adjusted to spatially separate networks operating on the same channel as far as possible.
- The use of directional antennas can be investigated.
- Next-generation Wi-Fi 6 APs and LCUs can be implemented, which uses advanced features such as BSS coloring to identify overlapping BSSs and differentiate between transmissions originating from within each BSS, which reduces RFI.

Furthermore, the single node range tests showed a 24.6 % increase in the RTT latency when the LCU was moved to a 110 m range from a 60 m range. The result of the full capacity range test indicated that the addition of high network capacity coupled with long-range transmission resulted in poor network performance. It can be assumed that the network performance will deteriorate significantly if the range of the network should exceed 110 m. As noted by Figure 3.2, the required range for the next generation field will be approximately 80 m, and thus the range results satisfied the requirement.

From the tests conducted within the heliostat fields, it can be observed that two additional factors significantly affect the performance of the wireless network, namely LoS and network density. Compared to the heliostat LoS deployment, the heliostat non-LoS tests resulted in an average RTT latency increase of 13.6 %, and the dense heliostat deployment resulted in an RTT latency increase of 12.4 %. It is thus essential to mount the LCUs with LoS to the AP to ensure optimal network performance. Lastly, the increase in latency for dense networks can be further mitigated by migrating to the next generation Wi-Fi 6 standard, which introduces features such as OFDMA to increase wireless performance significantly in dense network deployments such as the next generation PREMA project.

⁶50 [messages] \times 15.71 [ms] = 0.7855 [s]

Chapter 8

Conclusion

This thesis focussed on evaluating the performance of wireless communication in a Concentrating Solar Thermal (CST) Central Receiver System (CRS). This chapter summarises the project workflow, findings, and contributions. Conclusions are made based on the results of field tests, and recommendations are made for future work.

8.1 Summary of work

Research showed that the implementation of wireless communication in the heliostat field had cost and time saving potential, although research delving into the implementation and practical in-field performance is limited. Due to wireless communication's savings potential in a CST plant, a high-speed, low latency, and robust wireless communication network is proposed for the next-generation CST plant as part of the H2020 PREMA project.

An investigation into various wireless technologies, including Internet of Things (IoT) standards such as LoRa, NB-IoT, BLE, ZigBee, and various Wi-Fi standards, was conducted. Furthermore, an analysis of different network topologies required for the implementation of the technologies was completed. The possible hindrances to network performance in a heliostat field were identified and discussed.

A network architecture was designed using the 802.11n wireless standard, commonly known as Wi-Fi 4, and ZeroMQ, which is an open-source high performance asynchronous messaging library that utilizes a socket-based application interface. The architecture includes a server process responsible for major network tasks, a broker process used to route messages to the desired destinations, and a Local Control Unit (LCU) process to ensure a high level of control over each heliostat in the field. The network architecture was implemented on off-the-shelf bought items to simulate a fully operational CST communication network used to measure network performance in various test scenarios. The software and hardware platforms were fine-tuned with network time protocols, time offset procedures, and higher priority settings during the initial phase of testing to ensure reliable time sampling, which was used to determine the network performance.

Field tests were conducted at the Helio100 facilities at Mariendal, outside the town of Stellenbosch. The network performance was tested with LCUs placed in an interference-free, open field as well as within the heliostat field. Various tests were conducted to evaluate the impact of performance limiting factors such as transmission range, network capacity, network density, and shadowing. The network communication system successfully demonstrated the performance possibilities of wireless communication within a heliostat field and provided an understanding of the practical hindrances and limitations of wireless communication that will aid future plant designers.

8.2 Conclusions

All the objectives of this thesis have been successfully met. The implemented wireless communication network benefits from the simplicity of the ZeroMQ platform and off-the-shelf bought items while demonstrating that it can meet the stringent communication demands of high precision tracking heliostats in a CST plant. It has been demonstrated that low latency polling is achievable in a network of 50 LCUs, which simulates a cluster of the next generation heliostat field, proving the potential of wireless communication in CST fields.

The design requirement to poll the entire heliostat field's LCUs under one second was met as it was shown that the worst-case test scenario resulted in a polling rate of 0.79 s. Furthermore, a better understanding of the influence of different factors limiting wireless communication performance, such as network capacity, transmission range, network density, and shadowing effects, will pave the way for future CST plant design incorporating wireless communication.

During field testing, no results discredited the scaling potential of wireless communication for the next generation plant as the range limit of the network was found to be around 110 m, exceeding the network requirement of 80 m. The nature of the cluster-tree topology that the full-scale plant will implement requires several clusters operating concurrently. Care should be taken to operate neighboring Basic Service Sets (BSSs) in non-overlapping frequency bands and spatially separate the BSSs that operate in the same frequency band. The LCUs should also be placed with direct Line of Sight to the Access Point as shadowing effects in the heliostat field were found to degrade the network performance by 12.4 %.

This study's findings proved that if care is taken with the placement of LCUs to ensure LoS and radio frequency interference is limited by operating in non-overlapping frequency bands, wireless communication is feasible in CST plants and could drive down costs for future projects. Furthermore, the next generation Wi-Fi standard, 802.11ax, provides exciting features such as BSS coloring, target wake time, and orthogonal frequency division multiple access, which will further improve performance measurements presented in this thesis.

Although the tests were limited to a single site, the results can be generalized as the site is representative of an operational heliostat field, with limited external radio interference, dense LCU placement, surrounding steel structures, and large

capacity networks. Thus, the research question has been answered satisfactorily as it was proven that, yes, a wireless communication system is feasible to use in the next-generation CST plant.

8.3 Summary of contributions

This was the first project relating to network communication in heliostat fields in STERG. All network components, methods, hardware, and software were developed originally for this project. The resulting network and network performance tests satisfied the project's objectives but will also serve as a reference for practical performance of wireless communication in heliostat fields for future CST plant designers. The ease of future adaptation and implementation was a key consideration during the selection of hardware components and software platforms. The use of off the shelf bought items such as the Raspberry Pi Zero WH microprocessors illustrates this point. Further contributions included:

- A peer reviewed research paper was presented at the SAUPEC international conference in Cape Town in January 2020, with publication in the conference proceedings [33].
- A peer reviewed research paper was presented at the SolarPACES international conference that took place online in October 2020. The publication of the conference proceedings are expected in January 2021.

8.4 Recommendations for further research

Although the research question has been answered this section highlights areas, which according to the author's opinion, could benefit from future work to offer significant improvements in the network's performance:

- A network consisting of multiple clusters can be investigated to determine the operation in non-overlapping frequency bands. Spatially separated BSSs that operating on the same frequency band can also be investigated to limit inter-cluster RFI to ensure a high performing network.
- The next generation Wi-Fi 6 standard can be tested to determine if the new features discussed in Section 8.2 are indeed as promising as stated. Furthermore a test to utilise both the 2.4 GHz and 5 GHz spectrum offered by the 802.11ax standard can be explored to limit overlapping BSSs in the network.
- The security implications of the network design have not been considered, due to being out of scope. The final product will benefit from an in depth investigation into security measures as this is an important aspect of an industrial system implementing wireless communication.

Appendices

Appendix A

Statistical Analysis of Resulting Field Test Data

To determine if the results obtained in the field tests were driven by the change in test environment and conditions it was necessary to ensure that the obtained data was statistically significant. To determine the statistical significance, an independent t -test was conducted between each of the field test pair combinations to determine if any resulting data could be due to chance. The null hypothesis of the t -test states that the two populations have equal means. The alternative hypothesis states that the means of the populations are not equal, and thus the data is statistically significant.

The t -test assumes normally distributed data or a Gaussian-like distribution. The RTT latency test results, shown in Section 7.2, can generally have a thick tail and will thus reduce the power of the t -test. Thus, it was necessary to display the data in a normal distribution, for which the central limit theorem was used. The central limit theorem states that for a population with a specific mean and standard deviation, sufficient samples can be taken, and the means of these samples will resemble a Gaussian-like distribution [74].

Following the central limit theorem, the dataset was randomly shuffled, and samples of 100 RTT latency measurements were taken. The samples were not replaced into the dataset before taking the next sample, as the dataset does not resemble the population but rather a larger sample from the population. Each sample's mean was calculated to resemble a new data point in the normalized data, i.e., each data point in the normalized data resembles the average of 100 randomly sampled data points from the original data.

The normally distributed data can be used to perform t -tests. As each normally distributed dataset exceeds 1600 entries, an infinite degree of freedom is assumed. An infinite degree of freedom and a two tail significance level (α) of 0.01 results in a critical t -value of 2.576. If the absolute value of the result of the t -test exceeds the critical t -value, the null hypothesis should be rejected, and the alternative hypothesis assumed. Table A.1 shows the results of the corresponding t -tests. The test in the first column was taken as data set 1 and the test in the first row as data set 2. Thus a positive t value indicates that the mean of data set 2 is smaller than that of data set

1, and a negative t -value represents the opposite.

Table A.1: Results of t -tests for each test combination. A negative t -value indicates that the mean of data set 2 is larger compared to data set 1. To reject the null hypothesis the absolute value of the resulting t -value should exceed 2.576.

	Base	Cap ₁₀	Cap ₃₀	R ₆₀	R ₁₁₀	R _{FC}	H _{LoS}	H _{non-LoS}	H _{Dense}	H _{Sparse}
Base	X	163.3	106.4	606.0	235.5	-307.9	20.9	150.2	-153.6	723.5
Cap ₁₀	-163.3	X	-64.3	513.4	71.9	-343.8	-155.7	-252.6	-281.3	740.0
Cap ₃₀	-106.4	64.3	X	507.2	111.9	-303.0	-93.9	-203.5	-222.9	677.8
R ₆₀	-606.0	-513.4	-507.2	X	-277.8	-649.4	-612.9	-636.6	-721.1	124.5
R ₁₁₀	-235.5	-71.9	-111.9	277.8	X	-384.6	-219.7	-314.2	-344.9	374.9
R _{FC}	308.9	343.8	303.1	649.4	384.6	X	313.8	168.0	198.4	729.1
H _{LoS}	-20.9	155.7	93.9	612.9	219.7	-313.8	X	-164.1	-170.8	744.3
H _{non-LoS}	150.3	252.6	203.5	636.6	314.2	-168.0	164.1	X	13.2	735.6
H _{Dense}	153.6	281.3	222.9	721.1	344.9	-198.4	170.8	-13.2	X	838.8
H _{Sparse}	-723.5	-740.0	-677.8	-124.5	-374.9	-729.1	-744.3	-735.6	-838.8	X

It is clear to see that the t -values range from 13.2 to 838.8, which is significantly higher than the critical t -value of 2.576. The lowest t -value of 13.2 results in a p -value less than 0.0001. For this reason, the null hypothesis should be rejected, and the alternative hypothesis accepted. This proves that the datasets are statistically significant and that the driving forces for the shift in results are the test environments themselves.

Bibliography

- [1] The World Bank: Solar resource data: Solargis. 2017.
Available at: <https://globalsolaratlas.info/downloads/world>
- [2] Cumpston, J. and Mitsos, A.: Concentrating solar thermal overview. *Chemical Engineering Progress*, 2017. ISSN 03607275.
- [3] Malan, K.J.: A heliostat field control system. 2014.
- [4] Romero, M. and Steinfeld, A.: Concentrating solar thermal power and thermochemical fuels. *Energy & Environmental Science*, vol. 5, no. 11, pp. 9234–9245, 2012. ISSN 1754-5692.
- [5] Energy, T.: What is concentrated solar power? 2020.
Available at: <https://torresolenergy.com/en/technologies/what-is-concentrated-solar-power/>
- [6] Mosaic Slider 2. 2017.
Available at: <https://mosaic-h2020.eu/2017/02/12/slider-2/>
- [7] Mekki, K., Bajic, E., Chaxel, F. and Meyer, F.: A comparative study of LPWAN technologies for large-scale IoT deployment. *ICT Express*, 2019. ISSN 24059595.
- [8] Todtenberg, N. and Kraemer, R.: A survey on Bluetooth multi-hop networks. *Ad Hoc Networks*, 2019. ISSN 15708705.
- [9] Kubisch, S., Randt, M., Buck, R., Pfahl, A. and Unterschütz, S.: Wireless heliostat and control system for large self-powered heliostat fields. *SolarPaces Conference*, 2011.
- [10] Vergès, F.: Free Space Path Loss Diagrams. 2018.
Available at: <https://www.semfonetworks.com/blog/free-space-path-loss-diagrams>
- [11] Janek, J.F. and Evans, J.J.: Predicting Ground Effects of Omnidirectional Antennas in Wireless Sensor Networks. *Wireless Sensor Network*, 2010. ISSN 1945-3078.
- [12] Legare, C.: Which IoT protocol should I use for my system? 2017.
Available at: <https://www.embeddedcomputing.com/technology/software-and-os/which-iot-protocol-should-i-use-for-my-system>
- [13] Hintjens, P.: ØMQ-The Guide. *Online: http://zguide.zeromq.org/page: all, ..., 2011.*

- [14] Zhao, Y., Zhou, L., Tang, Q., Ma, D., Zhao, H., Wang, S. and Wei, J.: Investigation of High-Efficient Transfer Mechanisms for SCA 4.1. 2017.
- [15] Hockaday, L., McKechnie, T., von Puttkamer, M.N. and Lubkoll, M.: The Impact of Solar Resource Characteristics on Solar Thermal Pre-heating of Manganese Ores. In: *Minerals, Metals and Materials Series*. 2020. ISBN 9783030368296. ISSN 23671696.
- [16] Duffie, J.A., Beckman, W.A. and McGowan, J.: Solar Engineering of Thermal Processes . *American Journal of Physics*, 1985. ISSN 0002-9505.
- [17] Bondo, A.: Industrial or Commercial Wireless Mesh Technologies. 2019.
Available at: <https://www.datarespons.com/industrial-or-commercial-wireless-mesh-technologies/>
- [18] Newton, J.D.: *Uncommon friends*. Harcourt Brace Jovanovich, 1989.
- [19] Ellabban, O., Abu-Rub, H. and Blaabjerg, F.: Renewable energy resources: Current status, future prospects and their enabling technology. *Renewable and Sustainable Energy Reviews*, vol. 39, pp. 748–764, 2014. ISSN 1364-0321.
Available at: <http://www.sciencedirect.com/science/article/pii/S1364032114005656>
- [20] Ecotricity The end of fossil fuels.
Available at: <https://www.ecotricity.co.uk/our-green-energy/energy-independence/the-end-of-fossil-fuels>
- [21] Crippa, M., Oreggioni, G., D, G., Muntean, M., Schaaf, E., Lo Vullo, E., Solazzo, E., Monforti-Ferrario, F., Olivier, J. and Vignati, E.: *Fossil CO₂ and GHG emissions of all world countries*. 2019. ISBN 9789276111009.
- [22] Da Rosa, A.V.: *Fundamentals of renewable energy processes*. Elsevier Academic Press, Amsterdam, 2005. ISBN 0120885107.
- [23] Lovegrove, K. and Pye, J.: 2 - Fundamental principles of concentrating solar power (CSP) systems. In: Lovegrove, K. and Stein, W. (eds.), *Concentrating Solar Power Technology*, Woodhead Publishing Series in Energy, pp. 16–67. Woodhead Publishing, 2012. ISBN 978-1-84569-769-3.
Available at: <http://www.sciencedirect.com/science/article/pii/B9781845697693500029>
- [24] Duffie, J.A. and Beckman, W.A.: *Solar Engineering of Thermal Processes*. Wiley, 2013. ISBN 9781118415412.
Available at: <https://books.google.co.za/books?id=5uDdUfMgXYQC>
- [25] B. Stine, W. and Geyer, M.: *Power From The Sun by William B. Stine and Michael Geyer*. 2001.
- [26] International Renewable Energy Agency (IRENA): *Renewable Power Generation Costs in 2018*. 2019. ISBN 978-92-9260-040-2. arXiv:1011.1669v3.
- [27] Power, C.S.: Technology Roadmap Concentrating Solar Power. *Current*, 2010. ISSN 1557170X.

- [28] South Africa : Energy releases State of Renewable Energy in SA report. *Mena Report*, 2015. ISSN 2219-0112.
- [29] Department of Energy: Integrated Resource Plan. *Department of Energy*, 2018.
Available at: <http://www.energy.gov.za/IRP/irp-update-draft-report2018/IRP-Update-2018-Draft-for-Comments.pdf>
- [30] DoE: Commodity Flows and Energy Balances. 2015.
Available at: http://www.energy.gov.za/files/media/Energy{_}Balances.html
- [31] Assumptions and Methodologies in the South African TIMES (SATIM) Energy Model. 2014.
Available at: http://www.erc.uct.ac.za/sites/default/files/image{_}tool/images/119/Researchdocs/Satim/SATIMMethodology-v2.1.pdf
- [32] Lubkoll, M., Hockaday, S.A.C., Harms, T.M., Backström, T.W.V., Amsbeck, L. and Buck, R.: Integrating solar process heat into manganese ore pre-heating. In: *5th Southern African Solar Energy Conference, SASEC2018*. 2018. ISBN 0)218084242.
- [33] Liebenberg, A. and Smit, W.J.: Wireless Communication For A Modular Heliostat Field. pp. 1–6. 2020.
- [34] Alliance, L.: A technical overview of LoRa and LoRaWAN. *White Paper, November*, 2015.
- [35] Elkhodr, M., Shahrestani, S. and Cheung, H.: Emerging Wireless Technologies in the Internet of Things : A Comparative Study. *International Journal of Wireless & Mobile Networks*, vol. 8, 2016.
- [36] ICASA: National Radio Frequency Plan 2018. *Government Gazette*, 2018.
- [37] Durand, T., Visagie, L. and Booysen, M.J.T.: Evaluation of next-generation low-power communication technology to replace GSM in IoT-applications. *IET Communications*, p. 1, 2019.
- [38] Liberg, O., Sundberg, M., Wang, Y.-P.E., Bergman, J. and Sachs, J.: The Cellular Internet of Things. In: *Cellular Internet of Things*. 2018.
- [39] Vodafone: White paper: Narrowband-IoT: pushing the boundaries of IoT. Tech. Rep., 2017.
Available at: https://www.vodacombusiness.co.za/cs/groups/public/documents/document/vodafone{_}nb\~{I}ot{_}white{_}paper{_}fi.pdf
- [40] Sinha, R.S., Wei, Y. and Hwang, S.H.: A survey on LPWA technology: LoRa and NB-IoT. 2017.
- [41] Benitez, V.H., Armas-Flores, R.V. and Pacheco-Ramirez, J.H.: Experimental Study for the Development of a Wireless Communication System in a Solar Central Tower Facility. *International Journal of Electrical, Computer, Energetic, Electronic and Communication Engineering*, 2016. ISSN eISSN:1307-6892.

- [42] Tsai, C.-H., Pan, M.-S., Lu, Y.-C. and Tseng, Y.-C.: Self-Learning Routing for ZigBee Wireless Mesh Networks. 2009.
- [43] Farahani, S.: *ZigBee Wireless Networks and Transceivers*. 2008. ISBN 9780750683937.
- [44] Georgakakis, E., Nikolidakis, S.A., Vergados, D.D. and Douligeris, C.: An analysis of bluetooth, Zigbee and bluetooth low energy and their use in WBANs. In: *Lecture Notes of the Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering*. 2011. ISBN 9783642208669. ISSN 18678211.
- [45] Lin, J.-R., Talty, T. and Tonguz, O.K.: On the Potential of Bluetooth Low Energy Technology in Vehicular Applications. *IEEE Communications Magazine*, vol. 53, pp. 267–275, 2015.
- [46] Darroudi, S.M. and Gomez, C.: Bluetooth Low Energy Mesh Networks: A Survey. In: *Sensors*. 2017.
- [47] Wi-Fi Alliance® publishes 7 for '17 Wi-Fi® predictions. 2017.
Available at: <https://www.wi-fi.org/news-events/newsroom/wi-fi-alliance-publishes-7-for-17-wi-fi-predictions>
- [48] Siddiqui, F., Zeadally, S. and Salah, K.: Gigabit Wireless Networking with IEEE 802.11ac: Technical Overview and Challenges. *Journal of Networks*, 2015. ISSN 1796-2056.
- [49] What is MU-MIMO? How does it work and why do I need it? 2020.
Available at: <https://www.digitalairwireless.com/articles/blog/what-mu-mimo-how-does-it-work-and-why-do-i-need-it>
- [50] Bellalta, B.: IEEE 802.11ax: High-efficiency WLANS. *IEEE Wireless Communications*, 2016. ISSN 15361284.
- [51] Aruba Networks: Key Findings and Solutions. Tech. Rep., Miercom, 2013.
Available at: https://www.arubanetworks.com/pdf/technology/TR{_}Miercom{_}020ct2013.pdf
- [52] Šljivo, A., Kerkhove, D., Tian, L., Famaey, J., Munteanu, A., Moerman, I., Hoebeke, J. and De Poorter, E.: Performance evaluation of IEEE 802.11ah networks with high-throughput bidirectional traffic. *Sensors (Switzerland)*, 2018. ISSN 14248220.
- [53] Tian, L., Deronne, S., Latré, S. and Famaey, J.: Implementation and validation of an IEEE 802.11ah Module for ns-3. In: *ACM International Conference Proceeding Series*. 2016. ISBN 9781450342162.
- [54] OFFICIAL IEEE 802.11 WORKING GROUP PROJECT TIMELINES. 2020.
Available at: http://www.ieee802.org/11/Reports/802.11{_}Timelines.htm
- [55] Weinberg, N.: What is 802.11ax Wi-Fi, and what will it mean for 802.11ac.(Brief article). *Network World*, 2018. ISSN 0887-7661.

- [56] Qu, Q., Li, B., Yang, M., Yan, Z., Yang, A., Deng, D.J. and Chen, K.C.: Survey and Performance Evaluation of the Upcoming Next Generation WLANs Standard - IEEE 802.11ax. *Mobile Networks and Applications*, vol. 24, no. 5, pp. 1461–1474, oct 2019. ISSN 15728153. 1806.05908.
- [57] Atabaev, I.G., Akhatov, Z.S., Mukhamediev, E.D. and Zievaddinov, Z.: Modernization of an automated controlling system for heliostat field of big solar furnace. *Applied Solar Energy (English translation of Geliotekhnika)*, 2016. ISSN 0003701X.
- [58] Kolb, G.J., Jones, S.a., Donnelly, M.W., Gorman, D., Thomas, R., Davenport, R. and Lumia, R.: Heliostat Cost Reduction Study. *Test*, p. 115, 2007.
- [59] Kutscher, C., Gray, A., Ihas, B., Netter, J., Wendelin, T. and Zhu, G.: Low-cost heliostat for modular systems. *Concentrating Solar Power program review 2013*, pp. 123–124, 2013.
Available at: <https://www.nrel.gov/docs/fy13osti/58484.pdf>
- [60] Unterschütz, S. and für Telematik, I.: *Methodologies and Protocols for Wireless Communication in Large-Scale, Dense Mesh Networks*. Cuvillier, 2014. ISBN 9783954047079.
Available at: <https://books.google.co.za/books?id=SD5ooAEACAAJ>
- [61] García, G., Egea, A. and Romero, M.: Performance evaluation of the first solar tower operating with autonomous heliostats: PCHA project. 2004.
- [62] Freeman, J.D., Varghese, J.T. and Pullarkatt, D.: WSN based tracking for a concentrating solar thermal energy system. In: *2013 1st IEEE Conference on Technologies for Sustainability (SusTech)*, pp. 203–207. aug 2013.
- [63] Younis, M., Al-Shehhi, H., Hamar, N.A. and Meribout, M.: A wireless sensor network-based heliostat system using real-time image processing techniques. In: *2011 IEEE GCC Conference and Exhibition (GCC)*, pp. 465–468. feb 2011.
- [64] Mabusela, K., Kruger, C.P., Silva, B.J. and Hancke, G.P.: Design of a wireless heliostat system. In: *IEEE AFRICON Conference*, vol. 2015-Novem. Institute of Electrical and Electronics Engineers Inc., nov 2015. ISBN 9781479974986. ISSN 21530033.
- [65] de Paz, J.F., Tapia, D.I., Alonso, R.S., Pinzón, C.I., Bajo, J. and Corchado, J.M.: Mitigation of the ground reflection effect in real-time locating systems based on wireless sensor networks by using artificial neural networks. *Knowledge and Information Systems*, 2013. ISSN 02191377.
- [66] Kunavut, K.: The Effect of Transmission Range in Multi-hop Wireless Networks. In: *Proceedings of the 20th Asia-Pacific Conference on Communication, APCC 2014*. 2014.
- [67] Daniel E. Capano: Control Engineering | Understanding RFI and EMI's effects. 2015.
Available at: <https://www.controleng.com/articles/understanding-rfi-and-emis-effects/>

- [68] White paper: The basics of signal attenuation. 2016.
Available at: https://www.dataloggerinc.com/wp-content/uploads/2016/11/16{_}Basics{_}of{_}signal{_}attenuation.pdf
- [69] Lauener, J. and Sliwinski, W.: HOW TO DESIGN & IMPLEMENT A MODERN COMMUNICATION MIDDLEWARE BASED ON ZeroMQ. 2017.
- [70] Hintjens, P.: 23/ZMTP. 2020.
Available at: <https://rfc.zeromq.org/spec/23/>
- [71] Keary, T.: Latency vs Throughput – Understanding the Difference. 2020.
Available at: <https://www.comparitech.com/net-admin/latency-vs-throughput/>
- [72] Weisstein, E.W.: "Circle-Line Intersection." From MathWorld–A Wolfram Web Resource.
Available at: <http://mathworld.wolfram.com/Circle-LineIntersection.html>
- [73] Mamun, Q.: A qualitative comparison of different logical topologies for Wireless Sensor Networks. *Sensors (Switzerland)*, 2012. ISSN 14248220.
- [74] Sullivan, L.M.: *Essentials of biostatistics in public health*. 2nd edn. Jones & Bartlett Learning, 2012.