Time and Frequency Transfer in a Coherent Multistatic Radar using a White Rabbit Network

Author: Simon A. C. Lewis
Supervisor: Emeritus Prof. Mike Inggs

A thesis submitted to the Department of Electrical Engineering,
University of Cape Town, in fulfilment of the requirements
for the degree of Doctor of Philosophy in Engineering.

Cape Town, March 2021

Declaration

I declare that this dissertation is my own, unaided work. It is being submitted for the degree of PhD in Electrical Engineering in the University of Cape Town. It has not been submitted before for any degree or examination in any other university.

Signature of Author . . . Signed by candidate . . .

Cape Town
September 2, 2021

Abstract

Networks of coherent multistatic radars require accurate and stable time and frequency transfer (TFT) for range and Doppler estimation. TFT techniques based on global navigation satellite systems (GNSS), have been favoured for several reasons, such as enabling node mobility through wireless operation, geospatial referencing, and atomic clock level time and frequency stability. However, such systems are liable to GNSS-denial, where the GNSS carrier is temporarily or permanently removed. A denial-resilient system should consider alternative TFT techniques, such as the White Rabbit (WR) project. WR is an Ethernet based protocol, that is able to synchronise thousands of nodes on a fibre-optic based network with sub-nanosecond accuracy and picoseconds of jitter. This thesis evaluates WR as the TFT network for a coherent multistatic pulse-Doppler radar – NeXtRAD. To test the hypothesis that WR is suitable for TFT in a coherent multistatic radar, the time and frequency performance of a WR network was evaluated under laboratory conditions, comparing the results against a network of multi-channel GPS-disciplined oscillators (GPSDO). A WR-disciplined oscillator (WRDO) is introduced, which has the short-term stability of an ovenised crystal (OCXO), and long-term stability of the WR network. The radar references were measured using a dual mixer time difference technique (DMTD), which allows the phase to be measured with femtosecond level resolution. All references achieved the stringent time and frequency requirements for short-term coherent bistatic operation, however the GPSDOs and WRDOs had the best short-term frequency stability. The GPSDOs had the highest amount of long-term phase drift, with a peak-peak time error of 9.6 ns, whilst the WRDOs were typically stable to within 0.4 ns, but encountered transient phase excursions to 1.5 ns.

The TFT networks were then used on the NeXtRAD radar, where a lighthouse, Roman Rock, was used as a static target to evaluate the time and frequency performance of the references on a real system. The results conform well to the laboratory measurements, and therefore, WR can be used for TFT in coherent radar.

*“All I know, Time is a valuable thing. Watch it fly by as the pendulum swings.
Watch it count down to the end of the day. The clock ticks life away.”*

In the end, Linkin Park

Acknowledgements

I don't think I would have ever made it to writing this page, if it wasn't for my colleagues, friends, and family. Firstly, I'd like to thank UCT for giving me the best view of Cape Town to drink my 'early' morning coffee, I really cannot imagine a more picturesque setting to finish my studies. I'd also like to thank all of my University friends who joined me for said coffees, and who offered deep, but often ridiculous, conversations.

To the UCT personnel who helped me along my degree, starting with the stalwart of radar engineering, that is my supervisor, Emeritus Professor, Michael Inggs. Prof. Inggs has always been a helping hand, offering advice, support, and a relaxed sense of supervision. His years and years of experience, not only in theoretical work, but also practical engineering, really helped me come to grips with my project, and reignite my passion for hands on engineering. To Stephan Sandenbergh, who had a profound role in my project, and his work shaped much of my own. Without Stephan's guidance and incredible experience in electronics and time metrology, I would have been at a complete loss. I'd also like to thank the RRSG, UCT admin staff, and NeXtRAD teams, old and new. Of course, there are many, many people here, so I'll leave the individual mentions out, for fear of overlooking anyone.

Lastly, I'd like to thank my close friends and family, who pushed me to *just finish*, and who always believed that I would get there, even when I did not. This project would not have been possible without their support. I'd also like to single out my love, Ash, who has always been by my side, and gave me the confidence to complete this mammoth undertaking.

Contents

Declaration i

Abstract ii

Acknowledgements v

Contents vi

List of Figures x

List of Tables xiii

List of Symbols xiv

Publications xv

1 Introduction 1

- 1.1 TFT in Multistatic Radar . . . 7
  - 1.1.1 White Rabbit . . . 13
- 1.2 Research Motivation . . . 14
- 1.3 Hypothesis . . . 15
- 1.4 Executive Summary . . . 17

2 Background 26

- 2.1 Multistatic Time and Frequency Requirements . . . 27
  - 2.1.1 Multistatic Model . . . 28
  - 2.1.2 Time Accuracy . . . 30

2.1.3 Frequency Accuracy 32
2.1.4 Frequency Stability 32
2.2 Time and Frequency Transfer 36
2.2.1 TFT Methods 37
2.2.2 TFT Techniques 39
2.3 The White Rabbit Project 46
2.3.1 White Rabbit Overview 46
2.3.2 Hardware 49
2.3.3 Applications 51
2.3.4 Limits of White Rabbit Performance 53
2.4 Conclusion 55

3 Design of a White Rabbit Disciplined Oscillator Network 57
3.1 NeXtRAD Time and Frequency System 58
3.1.1 UCT GPSDOs 59
3.1.2 Frequency Distribution Unit (FDU) 63
3.1.3 Timing and Control Unit (TCU) 64
3.2 NeXtRAD Requirements 65
3.2.1 Time 65
3.2.2 Frequency 65
3.3 White Rabbit and GPSDO Setup 67
3.3.1 White Rabbit Network 67
3.3.2 Updates and Calibration 68
3.4 A White Rabbit Disciplined Oscillator 69
3.4.1 White Rabbit Fine Delay Design 70
3.4.2 White Rabbit Disciplined Oscillator 72
3.4.3 GPSDOs 73
3.5 Conclusion 74

4 Time-Domain Measurements of Coherent Frequency References 76
4.1 DMTD Test Bench 77
4.1.1 DMTD Technique 77
4.1.2 Test Setup 79

4.1.3 Sampling time and recording length 79
4.1.4 Cabling 81
4.2 Stability Analysis 82
4.2.1 Data Recording 82
4.2.2 Preprocessing 82
4.2.3 Data Interpretation 83
4.3 DMTD Calibration 83
4.3.1 DMTD Noise Floor 84
4.4 Results 86
4.4.1 Phase and Fractional Frequency 86
4.4.2 Time and Frequency Stability 91
4.5 Conclusion 96

5 Coherent Multistatic Radar TFT using WRDOs and Multichannel GPSDOs 98
5.1 NeXtRAD Setup 99
5.1.1 Radar Geometry 99
5.1.2 Target 99
5.1.3 Radar Parameters 101
5.1.4 GPSDO and WRDO setup 102
5.2 Measurement Testbed 103
5.2.1 DMTD Measurement and Radar Data 104
5.2.2 Performance Metrics and Processing 104
5.2.3 Radar Data Processing 105
5.3 Results 107
5.3.1 1 kHz PRF Bistatic Data 107
5.3.2 DMTD Phase and 100 Hz Bistatic PRF Data 111
5.3.3 Time Stability and Maximum Error 112
5.3.4 Phase Correction using DMTD Phase 116
5.4 Conclusions 118

6 Conclusions 121
6.1 Summary 121

6.2 Conclusion and Recommendations 125

A NeXtRAD 126

A.0.1 Active and Passive Nodes 127
A.0.2 Network 131
A.0.3 Operational Overview 132

B Reference Phase Noise Measurements 134

C Time and Frequency Stability 139

C.1 Phase Noise 139
C.2 Time-Domain Frequency Stability 140

Bibliography 143

List of Figures

1.1 Taxonomy of EM Sensors 2

2.1 Multistatic Model 29

2.2 White Rabbit Network Diagram 47

2.3 WR syntonisation and synchronisation 48

2.4 WR synchronisation flow 49

3.1 NeXtRAD time and frequency system 58

3.2 GPSDO block diagram 61

3.3 GPSDO MDEVs 62

3.4 GPSDO alarm and triggering system (images from [1]). 63

3.5 WR network setup 68

3.6 WR synchronised network 69

4.1 DMTD diagram 77

4.2 Reference DMTD measurement diagram 80

4.3 Image of the DMTD testbed 81

4.4 Noise Floor Phase 85

4.5 Noise Floor MDEV 85

4.6 WR fine delay phase and frequency 87

4.7 10s WRDO phase and frequency 87

4.8 100s WRDO phase and frequency 88

4.9 WRDO transients 89

4.10 GPSDO phase and frequency 90

4.11 Phase histograms 91

4.12 Fractional frequency histograms 92

4.13 ODEV of all references 94

4.14 MDEV of all references 94

4.15 TDEV of all references 95

4.16 MTIE of all references 96

5.1 NeXtRAD quasi-monostatic setup 100

5.2 Image of antennas and geometry setup 101

5.3 Image of False Bay and Roman Rocks 102

5.4 Control centre 103

5.5 Fast-time magnitude of Roman Rock 107

5.6 1kHz phase series 108

5.7 GPSDO frequency offset 109

5.8 PSD of phase for radar data 110

5.9 Coherent radar loss vs number of pulses 110

5.10 100Hz radar and DMTD phase 113

5.11 100Hz radar and DMTD frequency error 114
5.12 100Hz radar and DMTD MDEV 115
5.13 TDEV and MTIE of DMTD 116
5.14 DMTD corrected radar phase 117
5.15 DMTD corrected radar MDEV 117

A.1 NeXtRAD multistatic network 128

B.1 FDU phase noise 135
B.2 UCT GPSDO phase noise 136
B.3 Free-running WR fine delay phase noise 137
B.4 M400 phase noise 138

List of Tables

3.1 Phase noise requirement for REX 100 MHz reference and measured FDU output. Wenzel specifications given in brackets, and measurement noise floor given in asterisks. 66

3.2 WR network Specifications 67

4.1 Peak-peak and standard deviation of phase and frequency errors. . 90

4.2 Fractional frequency results over 8 hours. 91

4.3 Time domain stability results 93

5.1 Results of DMTD and 100 Hz Bistatic Recordings 106

5.2 Requirements vs. Results 118

A.1 NetRAD and NeXtRAD Specifications 126

C.1 Noise Power Laws 141

List of Symbols

B — Transmitted RF bandwidth

B_n — Noise bandwidth of receiver's pre-detection filter

c — Speed of light

f_c — Radar centre transmit frequency

f_{prf} — Pulse repetition frequency

F_R — Pattern propagation factor for Receiver-Target

F_T — Pattern propagation factor for Transmitter-Target

G_R — Receiver gain

G_T — Transmitter gain

K — Boltzmann's constant

λ — Wavelength

L_R — Receiver losses

L_T — Transmitter losses

R_R — Receiver-target range

R_T — Transmitter-target range

P_T — Transmitter power

σ_b — Bistatic radar cross-section of target

S/N — Signal-to-noise ratio

τ — Averaging interval

T_{int} — Coherent integration time

T_s — Receiver's system temperature

T_U — Update rate of clock synchronisation

Publications

White Rabbit in coherent multistatic radar

[2] S. Lewis and M. Inggs, Synchronisation of coherent netted radar using white rabbit compared with one-way multi-channel GPSDOs, in *IEEE Transactions on Aerospace and Electronic Systems*, pp. 110, 2020.

[3] S. Lewis, J. S. Sandenbergh, and M. R. Inggs, Evaluating an off-the-shelf white rabbit system to synchronise network radar via optic fibre, in *IEEE Radar Conference, RadarConf 2017*, pp. 16571662, 2017.

[4] S. Lewis, M. Inggs, and J. Sandenbergh, Investigation of white rabbit for synchronization and timing of netted radar, in *2015 IEEE Radar Conference*, pp. 214217, 2015.

NeXtRAD

[5] P. Beasley, M. Ritchie, H. Griffiths, W. Miceli, M. Inggs, S. Lewis, and B. Kahn, Multistatic radar measurements of UAVs at X-band and L-band, in *2020 IEEE Radar Conference (RadarConf20)*, pp. 16, 11 2020.

[6] R. Palama, F. Fioranelli, M. Ritchie, M. Inggs, and S. Lewis, Measurements of multistatic X- and L-band radar signatures of UAVs, in *2019 International Radar Conference (RADAR)*, pp. 15, 09 2019.

[7] M. Inggs, S. Lewis, R. Palama, M. Ritchie, and H. Griffiths, Report on the 2018 trials of the multistatic NeXtRAD dual band polarimetric radar, in *2019 IEEE Radar Conference (RadarConf)*, pp. 16, 04 2019.

Acronyms

5G fifth generation.

ADC analog-to-digital converter.
ADEV Allan deviation.
AOM acousto-optical modulator.
ARM Advanced RISC Machines.
AVAR Allan variance.

BC boundary clock.
BiSAR bistatic SAR.
BMCA best master clock algorithm.

CDR clock and data recovery.
CERN European Organisation for Nuclear Research.
CMOS complementary metaloxidesemiconductor.
CNC command and control.
CNGS CERN neutrinos to Grand Sasso.
COTS commercial off-the-shelf.
CPI coherent processing interval.
CW continuous-wave.

DAC digital-to-analog converter.
DDC digital downconversion.
DDMTD digital DMTD.
DDS direct digital synthesis.
DMTD dual mixer time difference.

DRO dielectric resonator oscillator.
DUC digital upconverter.
DUT device under test.

EM electromagnetic.
EMI electromagnetic interference.
ERP effective radiated power.
ESMA European Securities and Markets Authority.
EVN European VLBI Network.
EW electronic warfare.

FD fine-delay.
FDU frequency distribution unit.
FIFO first-in, first-out.
FM frequency modulated.
FMC fast mezzanine carrier.
FPGA field programmable gate array.

GDOP geometric dilution of precision.
GLONASS Global Navigation Satellite System.
GM grandmaster.
GMT general machine timing.
GNSS global navigation satellite service.
GNSSDO GNSS disciplined oscillator.
GPIO general purpose input and output.
GPS global positioning system.
GPSDO GPS disciplined oscillator.
GSI GSI Helmholtz Center for Heavy Ion Research.
GUI graphical user interface.

HDD hard disk drive.
HPA high powered amplifier.
HPBW half-power beamwidth.
HPSEC high precision second.
HVAR Hadamard variance.

IEEE Institute of Electrical and Electronics Engineers.
IF intermediate frequency.
IOO illuminators of opportunity.
IP internet protocol.
IQ in-phase and quadrature.

L1 layer 1.
LAN local area network.
LEO low earth orbit.
LFM linear frequency modulation.
LHC Large Hadron Collider.
LNA low-noise amplifier.
LO local oscillator.
LoS line-of-sight.
LVDS low-voltage differential signalling.

MAD mean absolute deviation.
MDEV modified Allan deviation.
MEO medium earth orbit.
MIKES Finnish centre for metrology.
MIMO multiple-input multiple-output.
MISO multiple-input single-output.
MMCM mixed-mode clock manager.
MO master oscillator.
MSRS multisite radar system.
MTIE maximum time interval error.
MVAR modified Allan variance.

NMEA National Marine Electronics Association.
NTP network time protocol.

OC ordinary clock.
OCXO ovenised crystal oscillator.
ODEV overlapping Allan deviation.
OHL open-hardware licence.

OHWR open-hardware repository.
OSI Open Systems Interconnection.
OTH over-the-horizon.
OVAR overlapping Allan variance.

P2P point-to-point.
PCB printed circuit board.
PCIe peripheral component interconnect express.
PCL passive coherent location.
PHY physical.
PLL phase-locked loop.
PLO phase-locked oscillator.
PM phase modulated.
PNT position, navigation, and timing.
PoE power-over-Ethernet.
PPP precise point positioning.
PPS pulse-per-second.
PRF pulse repetition frequency.
PTP precision time protocol (IEEE1588).

RCS radar cross section.
REX receiver and exciter subsystem.
RF radio frequency.
RFI radio frequency interference.
RFoF RF-over-fibre.
RHINO Reconfigurable Hardware Interface for computatioN and radiO.
RMS root mean square.
RTC real-time clock.

SA selective availability.
SCV sub-clutter visibility.
SDF software defined filter.
SDR software defined radio.
SFP small form-factor pluggable.

SIMO single-input, multiple-output.
SISO single-input single-output.
SKA Square Kilometre Array.
SNR signal-to-noise ratio.
SPEC simple PCIe carrier card.
SSD solid state drive.
STALO stable local oscillator.
SyncE synchronous ethernet.

TCU timing and control unit.
TCXO temperature compensated crystal oscillator.
TDC time-to-digital converter.
TDEV time deviation.
TF time and frequency.
TFT time and frequency transfer.
TIA international atomic time.
TIC time interval counter.
TOA time of arrival.
TTL transistor-transistor logic.
TVAR time variance.
TWOTFT two-way optical time and frequency transfer.
TWSTFT two-way satellite TFT.

UCL University College London.
UCT University of Cape Town.
UTC universal coordinated time.
UTP unshielded twisted pair.

VCO voltage controlled oscillator.
VHF very-high frequency.
VLA Very Large Array.

w.r.t with respect to.
WR White Rabbit.
WR-LEN WR light embedded node.

WRDO	White Rabbit disciplined oscillator.
WRN	White Rabbit network.
WRPC	WR PTP core.
WRS	White Rabbit switch.
ZCD	zero-crossing detector.

Chapter 1

Introduction

Since the invention of the *telemobiloscope* by Christian Hülsmeyer in 1904[1], the world of remote sensing has been revolutionised by the explosive growth of microwave sensor technology. Microwave sensing has allowed us to accurately map the topography of our planet [8], track space debris smaller than a baseball [9], monitor coastlines for illegal fishing or hostile agents [10], and image the distribution of neutral hydrogen in the Milky Way galaxy [11].

Despite the wide variety of applications and hardware designs, all microwave sensors have common features, and can therefore be grouped into categories. In Inggs et al's taxonomy of microwave sensors [12] (figure 1.1), a sensor is primarily defined by three features, namely,

1. *Transmitter Dedication*: Active sensors make use of dedicated transmitters, which are primarily committed to their sensing application, a feature common to conventional radar. Where a transmitter or source of radiation is already available, one can use passive, or receive only sensors. Passive sensing typically reduces the cost of the sensor significantly, and often provides additional benefits like covert operation, weight reduction, and reduced power consumption.

---

[1] Hülsmeyer was the first inventor to patent a device that used electromagnetic radiation to detect ships hidden in heavy fog.

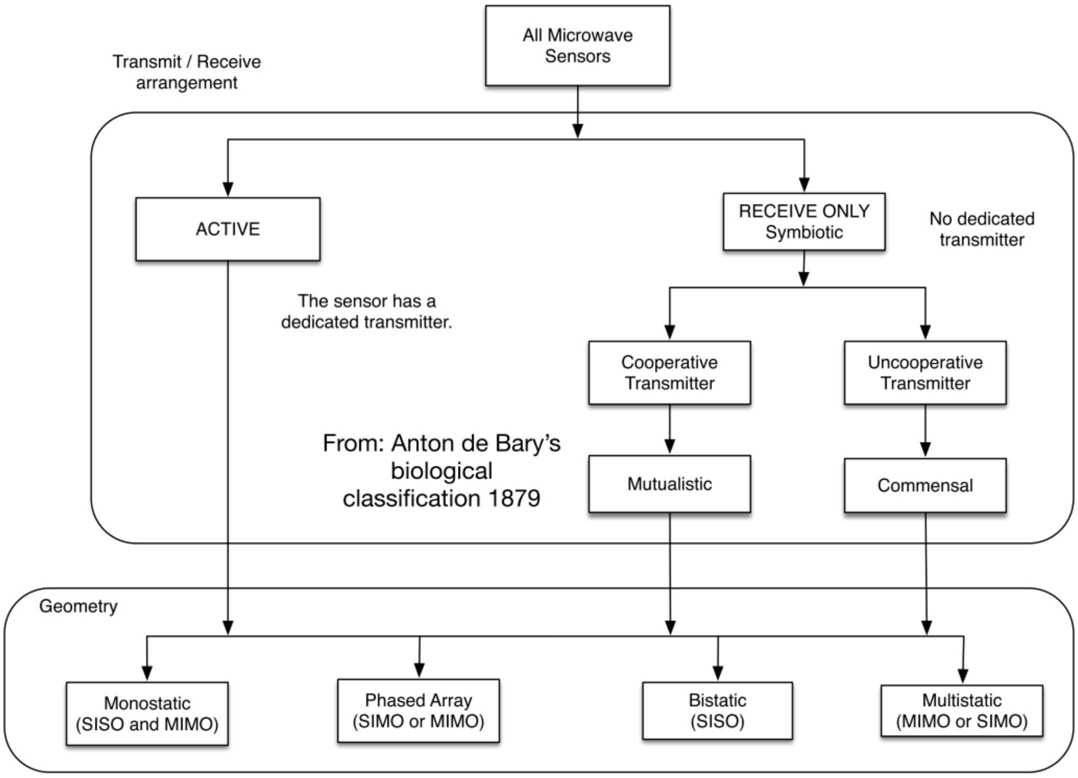

Figure 1.1: Taxonomy of EM sensors from [12] (used with permission)

2. *Cooperation*: Sensors can be *cooperative*, where each node in the network is aware of every other node, or *independent*, where nodes act autonomously, and do not need to directly communicate with other nodes. Cooperative transmitters are required if one needs to control the transmission properties, such as the effective radiated power (ERP), carrier frequency, polarisation, modulation, and beam pattern.

   In passive only sensing, Inggs et al further describe the level of cooperation within the sensor by the type of symbiotic relationship that exists between a transmitter and receiver. Using Anton de Bary's biological classification nomenclature, passive sensors can be *mutualistic*, where the transmitter and receiver share common goals, or *commensal*, where only the receiver gains a benefit from the relationship.

   Mutualistic transmitters either reserve a portion of the transmission time to the sensor's goals, or design their transmissions to simultaneously serve the sensor and primary application. In the first case, during a short segment or allotted time, the transmitter radiates a specific waveform designed solely for sensing, whilst during normal operation the transmitter serves its primary purpose. The second scenario has gained more attention in recent years, specifically in the case of joint radar and communications [13]. The rise in popularity is due to increasing electromagnetic (EM) spectrum congestion, and the increase in research into autonomous vehicles, which require radar sensors and vehicle-to-vehicle communications.

Commensal sensing does not require that the source of radiation is aware of the receiver, and its transmissions are solely used for its own purposes. Commensal sensors therefore make use of illuminators of opportunity (IOO), where the received radiation is either from man-made sources, like a communications tower, or from environmental sources, like solar radiation. Commensal sensing has become popular because one does not need to construct or operate the transmitter, and the sensor can better utilise the already crowded EM spectrum. However, the sensor has no control over the transmission properties, which may constrain the performance of such systems.

3. *Geometry*: The spatial distribution of the sensor's hardware defines its geometry. Monostatic sensors are composed of a single co-located transceiver, and generally use a single antenna with a transmit-receive switch to isolate between the much higher transmission power and the receiver's sensitive front-end. Bistatic sensors have the transmitter and receiver physically separated along a baseline, allowing the phenomena or target to be viewed from an aspect angle that is off-axis to the transmitter. The angle ranges from zero, or quasi-monostatic, to 180°, or forward-scatter. Phased arrays function as monostatic or bistatic systems, but distribute the antenna's aperture into multiple transmit and/or receive elements, such that the antenna's directivity can be electronically steered.

   Multistatic sensors extend the concept of bistatic sensing, where the sensor is composed of multiple cooperative bistatic pairs. This allows the sensors

> to simultaneously observe the phenomenon or target from multiple perspectives, thereby enhancing the amount of available spatial information. Unfortunately, there is a fair amount of ambiguity in the naming of multistatic systems. For instance, the term *multistatic sensor* is specifically used to describe a system of multiple cooperative bistatic pairs, which requires some form of communications network to fuse the data. However, a *netted sensor* describes a network of cooperative or independent monostatic sensors, able to simultaneously transmit and receive radiation from any other transmitter, acting somewhat like a bistatic radar. In radar, Chernyak coined the term multisite radar system (MSRS), which defines the most general form of distributed radar. A MSRS can be composed of an arbitrary arrangement of cooperative or non-cooperative bistatic and monostatic radars, thereby including multistatic and netted radars.
>
> Inggs et al also include the input-output relationship, which is often where some confusion arises. Monostatic and bistatic systems are single-input single-output (SISO), whilst multistatic systems are multiple-input multiple-output (MIMO), single-input, multiple-output (SIMO), or multiple-input single-output (MISO), depending on the number of transmitting and receiving elements. However, one can also have monostatic or bistatic MIMO, where the antenna is composed of a specialised phased array, that uses multiple orthogonal transmit and receive channels to improve angular resolution.

Whilst this taxonomy essentially covers all microwave sensors, it can be extended to include the level of *coherence* that exists between nodes. Coherence in microwave sensing is related to the degree of phase stability that exists between nodes during some finite amount of time, known as the coherent processing interval (CPI). In a more concrete, mathematical description, coherence is a normalised measure of the cross-correlation between the transmitted and received waveforms.

The sensor's spatial coherence is governed by the phase stability due to the propagation of the microwave signal through physical hardware and the transmission medium, the antenna phase centres, as well as the dynamics of the transmitter, receiver, and phenomenon or target. If the spatially induced phase is constant, or at least nominally stable, the phase presented to the receiver's sampler is a function of the underlying frequency stability of the frequency references, like the local oscillator (LO), and the aperture jitter of the digitiser's analog-to-digital converter (ADC).

Considering that most modern transmitters and receivers use heterodyning techniques to convert between a baseband signal and the microwave carrier, the phase stability in this case is dominated by the frequency stability and accuracy of the sensor's stable local oscillator (STALO)[2], which in turn is typically locked to a lower frequency reference, known as the sensor's master oscillator (MO). The MO provides a frequency reference for the local clock, which gives the sensor a local notion of time, and provides a timebase for digital hardware and time-triggered control operations.

Monostatic systems employ a single STALO, which is used for both the up- and downconversion. Using the same STALO ensures that there is always a reference value of phase for the receiver, and thus absolute measurements of phase are inherent to the sensor. Additionally, low frequency phase instability, or *close-in* phase noise, due to the oscillator on transmit, is highly correlated with the phase noise on receive. This strong correlation allows a significant portion of the low-frequency phase instability to be partially cancelled, which assists in sensing phenomena with low Doppler shifts, that would otherwise be swamped by the sensor's phase noise.

Bi- and multistatic systems on the other-hand, are physically separated, and therefore generally do not share a common STALO. Because different STALOs are now required, and no two clocks are ever precisely the same, relative time and frequency (TF) errors between the transmitter and receiver cause proportional phase errors in the bistatic receiver. A drift in phase over time, due to a small offset in reference frequency, leads to a loss of coherence, and subse-

---

[2] The phase uncertainty due to ADC aperture jitter is typically negligible in comparison to the STALO phase noise.

quently sets a maximum time for the CPI. The spatial coherence in a distributed sensor is therefore dependent on the relative phase stability between the two nodes. Having a phase reference for absolute phase measurements is now also more challenging, as the phase between the STALOs is mutually independent. Furthermore, the uncorrelated nature of distributed STALOs, means that there is no cancellation of phase noise in the receiver, impacting the spectral performance of these sensors. Systematic frequency offsets cause an apparent shift in the spectrum of the baseband signal, whilst time errors between the local clocks impacts the temporal ordering of events, which causes errors in the triggering, control, and communications between sensor nodes. Distributed nodes therefore need to be synchronised in time and phase, and syntonised in frequency.

Syntonisation, which is achieved through frequency transfer, is the process of setting the reference frequency of a sensor node to common network frequency, where differences between the two node's absolute phase and time are ignored. Synchronisation involves the alignment of time, and/or phase to a master source by using time transfer techniques. Some confusion often arises between time and phase, because in essence, *time* is a continuous property. However, in its common usage, time is a measured property that takes on a discrete value, with a resolution equal to one complete period of the clock's reference frequency. The phase, and in particular the *instantaneous phase*, is a radial value between 0 and 2π that exists between each clock increment.

The degree to which nodes can be synchronised and syntonised, ultimately impacts the type of coherent processing that can be achieved. In terms of MSRSs, Chernyak defines three types of spatial coherence, ordered from the most challenging fully coherent network, to the simplest noncoherent network.

1. *Fully Coherent*: Each node has an absolute notion of phase, and the phase is stable throughout the CPI. This requires accurate carrier phase synchronisation and syntonisation, as well as an explicit knowledge of the position and dynamics of each node. Time synchronisation is also required for cooperative reception and control.

   By measuring the absolute phase and time of arrival (TOA), a fully spatially coherent network can operate in such a way that EM field distribution can be measured. This is a requirement in distributed beamforming, interferometers, sparse arrays, and fully coherent netted sensors, where data is fused at the raw, or radio signal level.

2. *Short-term Coherent*: Each node is only phase stable during the CPI, giving magnitude and relative phase information, and phases at the beginning of a CPI are random and mutually independent. This requires only that there is accurate carrier syntonisation, and time synchronisation for cooperative sensing. The relaxation in phase requirements, means that the field distribution can no longer be measured on timescales longer than a few seconds, but relative changes between stations, such as differential Doppler information and TOA, are still valid.

3. *Noncoherent*: There is no carrier phase synchronisation or syntonisation, and consequently only magnitude information is available for multistatic data fusion. However in systems of independent active sensors, the nodes will still have high temporal coherence, so independent sensors are able to measure a change in their radial frame, such as local Doppler information. A network of noncoherent, but cooperative sensors can therefore fuse track and plot data, and a majority voting scheme can be used to discern Doppler and TOA. However, the cooperative nature still implies that time synchronisation is needed.

1.1 Multistatic Radar Time and Frequency Transfer

The very first radars operated in continuous-wave (CW) bistatic geometries, where a receiver measured the beat-frequency produced by the interference of a Doppler shifted echo returned from a target in motion, and the direct path signal from a transmitting station situated nearby [14, 15]. These radars provided no sense of target range, but were instrumental in detecting approaching hostile aircraft during World War 2. With the inventions of the transmit-receive switch and coherent microwave transmitters, monostatic systems gained popularity, and have historically been favoured due to their lower cost and design simplicity. However, monostatic radars can only estimate target range and Doppler in the radial direction to the sensor, limiting their ability to determine precise measurements of target position and range. This limitation also makes them vulnerable to targeted jamming, and weak to stealth techniques that reflect the radar signal in non-radial directions.

The geometric diversity of multistatic radars offer an effective means of countering hostile stealth targets in homeland security applications [16], mitigating sea clutter to detect low radar cross section (RCS) targets in high sea states [17]; and the ability to calculate target heading using the additional spatial vectors, assisting in tracking applications. Additionally, with the use of covert passive receive only nodes, the radar is resilient to targeted hostile jamming. This is because passive nodes do not radiate, and therefore hostile signal interception only affects active transmitters. The concept of graceful failure is also inherent in the network, as failing or compromised receivers reduce the overall coverage of the radar, but don't stop the radar from operating completely. However, the design of coherent multistatic radars is more complex than monostatic or incoherent bistatic radar, with previous attempts being limited by contemporary hardware and data processing.

There is once again a growing interest in multistatic radar, propelled by improvements in reconfigurable hardware and software defined radio (SDR), parallel computing, advances in digital signal processing, high-bandwidth data links, and low-cost, accurate and stable time and frequency transfer (TFT) systems. These improvements have also assisted the growth of other remote sensing fields, notably: cognitive radar [18], interferometric array astronomy [19], shared spectrum radar and communications [20, 21], wireless beamforming using distributed phased arrays [22], and commensal sensing [23].

However, even with the state-of-the-art, designing a high resolution coherent multistatic radar is still not a trivial endeavour. Fully coherent or short-term coherent pulse-Doppler multistatic radars are especially difficult, because they have strict requirements on the level of time and phase synchronisation, and frequency syntonisation. To motivate the need for TFT in multistatic radar, consider the case of a single cooperative radio frequency (RF) transmitter and several identical passive receivers placed equidistant from the transmitter. In this perfect example case, let us assume that the reference is a pure tone, with no phase noise, and that there is no noise induced in the frequency transfer.

In this idealised scenario, each node of the radar shares a single reference frequency from which all other ensemble frequency sources and timing and control signals are derived. Assuming a lossless transfer medium, the LOs and transmitter RF could be synthesised locally at the reference node, and then distributed to each remote node, or alternatively, the MO's frequency could be distributed directly, and the respective LOs synthesised remotely. Distributing the same frequency reference ensures that the measured phase of each echo has a consistent offset with the transmitter's phase, giving the network full spatial coherence.

Rephasing the timebase of the nodes, and thus phase synchronising them, is simple because the geometric displacement is known, and thus one can calculate the one-way phase difference of the frequency reference. Whilst this does not give an absolute notion of time, one could use a simple one-way time transfer, such as an epoch pulse sent from the transmitter, that takes into account the one-way delay. Each node therefore shares the network time-base, allowing precision control over time-sensitive operations such as developing a stable pulse repetition frequency (PRF), controlling transmit/receive switches and effectively triggering the start of a measurement.

However, in a real world scenario, a number of problems with the idealisation become apparent. For instance, the baseline between the transmitter and receiver nodes may be several kilometres, meaning that additive noise, interference, dispersion, and signal attenuation is incurred by transferring the reference from the master to the slave site. Phase noise of the reference is inescapable, and only low-frequency components and oscillator drift are typically correlated within the transfer time. The phase stability of the transfer medium is also not timeinvariant, which is mostly due to path length changes from thermal expansion in the case of cables, or atmospheric turbulence in microwave or optical signals. The simplistic one-way time transfer signals will also incur additive jitter due to the noisy transmission channel, finite channel bandwidth, and active components, adding uncertainty to the epoch trigger time.

If transferring the reference to remote sites is therefore difficult, one could attempt to use individual stable MOs at each node. Whilst this may alleviate problems associated with transferring the TF reference directly, there will be a new set of challenges to overcome. For starters, the STALOs are now matched, but have uncorrelated phase noises. This is essentially the case of an infinite delay in the transmit-receive correlation, meaning that close-in phase noise components are no longer cancelled in downconversion. This increase in low-frequency phase instability leads to broadening of the clutter spectrum, thus reducing the sub-clutter visibility (SCV). Frequency offsets between the transmitter and a receiver will impose a frequency offset on the demodulated signal, leading to an apparent Doppler shift. Coherent integration is now impaired by relative oscillator phase drift, which limits the maximum CPI. If each node now has a local representation of time, as opposed to a network time, there will be ambiguity in the start time of an experiment, even if the times were initially synchronised. Their notion of time will diverge according to their time stability. Time errors between the receiver and transmitter are manifested as proportional range errors, and a mere 10 ns difference translates to a bistatic target range inaccuracy of 1.5 m (in the case of a quasi-monostatic bistatic radar). Mitigating this range drift and inaccuracy will require accurate time synchronisation between the receiver and reference clocks, with clock frequency syntonisation to reduce drift during the clock holdover. Cooperative reception also requires knowledge of each node's position and staring vector.

The bistatic pairs must therefore have some means of synchronisation in time, and syntonisation of the frequency references. Full coherence further requires absolute phase synchronisation. Therefore, accurate and stable TFT is a critical feature of any coherent multistatic radar. Achieving short-term and full coherence is possible with one of the numerous TFT methods and techniques. Common methods used in multistatic sensing, include simple one-way coaxial cables dedicated to time or frequency transfer [24], networks of multi-channel GNSS disciplined oscillator (GNSSDO)s [1], two-way TFT using microwave links, or one-way direct-path RF synchronisation. However, the most commonly used methods are those based on global navigation satellite service (GNSS) TFT, and have been favoured for a number of reasons.

Firstly, they simplify the spatial requirements for multistatic systems, inherently providing a GNSS location with respect to a standard geoid model (typically WGS84). Nodes can therefore be placed arbitrarily (under normal GNSS required conditions), forming baselines on the order of several kilometres without the need for fixed infrastructure. Highly mobile nodes can therefore be designed, allowing arbitrary multistatic geometries to be formed. Secondly, GNSS satellites are equipped with high-performance atomic clocks that are constantly steered to the universal atomic timebase, international atomic time (TIA)/universal coordinated time (UTC). The need for accurate and stable time underpins the operation of any GNSS system, thus the terrestrial receivers inherently gain a high degree of time accuracy (below 100ns) and long-term stability relative to UTC. On typical baselines seen in multistatic systems, the GNSS receivers share common sky coverage and can thus be operated collaboratively in common-view, which reduces the relative time and phase errors of the GNSS signal due to atmospheric phase instability.

Locking an averaged version of the pulse-per-second (PPS) output from a standard GNSS receiver to the output of a high quality oscillator, typically an ovenised crystal oscillator (OCXO), provides a cost-effective frequency reference with the short-term frequency stability of the OCXO and the long-term stability of the GNSS constellation - known as a GNSSDO. Finally, GNSS receivers are increasingly getting cheaper, and better in both positional and temporal accuracy due to additional satellites being added to the various constellations, and receivers simultaneously using multiple GNSS providers.

For similar reasons, GNSS has revolutionised almost all scientific, commercial and social sectors for any application that requires position, navigation, and timing (PNT). The services are offered for free, and are always available, being treated more as utilities, akin to water and electricity. However, reliance on the availability of GNSS services comes with a conceptual danger. Their ubiquity means that if the respective GNSS service is interrupted, even for short periods, there may be severe and wide-reaching consequences such as property damage, failures in the stock market or loss of life [25]. In the United States, denial of the global positioning system (GPS) carrier has been labelled as a "single point of failure for much of the US economy and critical infrastructure" [26].

This impairment is now collectively referred to as *GNSS-Denial*, where the respective GNSS signal is temporarily, or permanently removed. In the case of multistatic radar, GNSS-denial removes the inter-node synchronisation, which affects the relative TF accuracy and stability of the clocks and STALOs. This proportionally affects the range and Doppler characteristics of the radar, leading to poor estimates of target position and velocity. Notably, this issue is a threat in the defence community, where intentional GNSS jamming as a form of electronic warfare (EW) is being pursued using low-power, commercially available radio jammers and spoofers [27]. Over the last few years there has been an increase in the use of these devices, with several cases of EW exploiting this vulnerability in both military and civilian contexts [28, 29].

GNSS governing bodies also control access to their respective GNSS service, and have the means to deny or deteriorate the service's signal quality if there is significant threat. In the case of GPS, an added noise signal, known as selective availability (SA) [30], was intentionally included in the transmission to reduce the positional accuracy of the civilian carrier for security concerns. As of the year 2000, SA has been deactivated (with no intention of reactivation), but could plausibly be included in times of conflict.

Probably the most underrated concerns, yet most catastrophic, are the effects of space-weather and space-junk on the performance of GNSS constellations. Solar flares have the potential to significantly impair GNSS coverage during peak solar flux. This was proven during a study by Sreeja et al [31], where a reduction of several decibels in the carrier to noise ratio of all GPS signals was measured for a moderate solar flare. In terms of space-junk, the Kessler Syndrome [32] - a hypothetical cascade of space-junk - may permanently remove access to GNSS. Whilst GNSS satellites are typically in medium earth orbit (MEO), direct satellite damage is still plausible. Furthermore, even if the satellites are unharmed, low earth orbit (LEO) space debris will hamper the free-space propagation of the RF carrier.

Denial-resilience is therefore a necessity in applications where significant harm or loss is possible. As such, a number of *Resilient PNT* architectures have emerged, from rebuilding a modern LORAN system called eLoran [26] to using private sector LEO satellites [33]. If the time and frequency system is the only concern, there are several other options available, such as using chip-scale atomic clocks for improved holdover [34] or the use of dedicated TFT hardware ranging from microwave links or fixed cable installations for medium baselines, to two-way satellite TFT (TWSTFT) or dedicated atomic clocks for very large baselines. However, many of these options are not feasible due to performance or cost.

1.1.1 White Rabbit

One option for coherent pulse-Doppler multistatic radar, which has seemingly been overlooked, is the use of the White Rabbit (WR) protocol [35]. The protocol, originally developed by the European Organisation for Nuclear Research (CERN) [36] and a number of research and industrial partners (notably GSI Helmholtz Center for Heavy Ion Research (GSI) [37]), aims to create a standardised, deterministic Ethernet network with sub-nanosecond time synchronisation and picoseconds of jitter. WR is able to synchronise thousands of nodes on a single network using a hierarchical clock structure, with inter-node single-mode fibre lengths up to 10 km. The hardware and software is licensed under an Open Source licence, allowing designers and developers to cater a WR approach to their specific application, avoiding any vendor lock-in.

Whilst initially developed to replace CERN's general machine timing (GMT) network for the accelerator complex (which includes the Large Hadron Collider

(LHC)), WR has garnered the attention of users requiring accurate time dissemination and deterministic data transfer. A list of disclosed projects is recorded on the WR project site [38], such as a modified WR link implemented by Finnish centre for metrology (MIKES) spanning 950 km in Finland for UTC(k) comparisons [39, 40]; the timing system for the Cherenkov Gamma-Ray telescope array [41]; and the timing and control system for the CERN neutrinos to Grand Sasso (CNGS) [42] experiment.

WR is able to provide sub-nanosecond time synchronisation by synchronising the nodes with a modified version of the precision time protocol (IEEE1588) (PTP), that additionally syntonises the nodes to a common network frequency using layer 1 (L1) frequency transfer. The enhancement in time synchronisation accuracy is due to a novel method for phase measurements, that additionally synchronises the phase of the nodes. These features are highly attractive to realising a fully coherent multistatic pulse-Doppler radar.

1.2 Research Motivation

The NeXtRAD radar is a joint project between University of Cape Town (UCT) and University College London (UCL), which aims to create a fully coherent multistatic networked radar. The radar successfully incorporates multi-channel GPS disciplined oscillator (GPSDO)s for synchronisation and syntonisation, which are capable of synchronising TF to acceptable levels for multistatic radar [1], with relative peak-to-peak time accuracy on the order of 10 ns [43]. Their long-term frequency stability is comparable to caesium atomic references, but are a fraction of the cost. However, their relative phase drift limits the CPI to the order of a few seconds, meaning that NeXtRAD can only be operated as a short-term coherent system.

These systems are also prone to GNSS related issues, specifically jamming of the GNSS carrier, atmospheric instability or loss of GNSS signal coverage. Mitigating the effects of GNSS-denial is critical to system robustness, especially in cases where significant damage to social and economic sectors is plausible. Whilst fixed infrastructure is not ideal for NeXtRAD specifically, little research has focused on replacing vulnerable multistatic radars with alternative TFT networks. In the case of stationary multistatic radar, such as a coastal defence network or dispersed array, one can make use of fixed infrastructure.

In this sense, fibre optic based TFT, as is the case with WR, is considered as an important step to improving both the robustness and performance of NeXtRAD. Fibre based systems are robust against electromagnetic interference (EMI), can carry communication data in a secure channel, and typically provide a premium on TFT performance for long-range TF dissemination.

This thesis will explore the use of WR in the remote microwave sensing field, with an emphasis on TFT on fully coherent multistatic pulse-Doppler radar.

1.3 Hypothesis

A fibre-optic based WR network installation presents a unique opportunity to avoid radio frequency interference (RFI) related problems, whilst improving the time synchronisation accuracy of a multistatic radar to the sub-nanosecond regime, and opening the door to fully coherent multistatic radar with improved phase synchronisation performance. Therefore the primary hypothesis of this thesis is:

The White Rabbit protocol can be used in a fully coherent pulse-Doppler multistatic radar, enabling sub-nanosecond time synchronisation, and improved phase stability for coherent processing, when compared to a network of multi-channel GPSDOs.

To test this hypothesis, the following research questions are drawn:

1. *Are commercial White Rabbit systems capable of fulfilling the requirements for time and frequency accuracy and stability in multistatic radar?*

   The TFT performance of WR has been extensively evaluated over the last few years. However, the protocol's focus is primarily concerned with accurate time dissemination. Frequency accuracy and stability are expected to trace the Grandmaster clock in the network, however it is hypothesised that the phase noise of commercially available WR equipment may be higher than precision OCXOs, due to their lower quality temperature compensated crystal oscillator (TCXO)s.

2. *How does the TF accuracy and stability of a WR network compare to a network of multi-channel GPSDOs in an operational multistatic radar?*

   WR is evaluated as an alternative to GPSDOs, which are typically a popular choice for distributed sensors due to their mobility and spatial referencing. The performance of a network of multi-channel GPSDOs was evaluated by Sandenbergh using the Netrad radar [1], and found to be acceptable for short-term coherent multistatic radar synchronisation. To validate this research question, the WR based system's TF performance must meet or exceed that of the GPSDOs.

1.4 Executive Summary

This thesis concerns TFT for networks of coherent sensors, with a specific focus on time synchronisation and frequency syntonisation of a coherent multistatic radar using the White Rabbit protocol.

Multistatic radars provide several benefits over traditional monostatic systems, from increased spatial coverage to low-cost and covert passive receivers. Historically however, monostatic radar has been favoured due to their simplicity. For a radar to make use of coherent processing, a stable source of frequency is required for phase referencing. In the monostatic case, all frequencies (from the digital clocks to the high frequency STALOs) are derived from a single stable MO. This fact means that there is a system-wide notion of time and frequency, traceable to a single source. As a benefit of this scheme, there is a direct correlation to sources of frequency instability, specifically for long-term effects like MO drift and wander. More importantly, because the same STALO is used for up- and downconversion, there is an innate sense of self-coherence between the transmitted and received waveforms, which assists in reducing STALO phase noise.

In addition to requiring more equipment and resources, physically displacing the transmitter and receiver for bi- or multistatic radar, introduces a number of complications. In particular, the time and frequency system is more complex because baselines of several kilometres are typical, meaning that sharing a common frequency reference is non-trivial. Individual MOs are therefore typically used as sources of time and frequency for each node. However, no two clocks run in perfect synchronicity due to their independent noise processes. Nodes therefore require time synchronisation and frequency syntonisation. The network timebase is used to order and control events, creating a common epoch for bistatic range estimation; whilst frequency syntonisation is required for the radar to operate coherently. Relative errors in time lead to proportional bistatic range errors, whilst frequency offsets lead to errors in Doppler estimation. Relative frequency instability becomes a major limiting factor in the radar performance, as the long-term stability determines the coherence time between the node references, limiting the maximum CPI; whilst downconversion is no longer performed using the same STALO, causing the uncorrelated phase noise in the transmitted waveform to add in the receiver.

Accurate and stable TFT is therefore pivotal in distributed coherent sensing. Whilst several TFT options are available, distributed sensors have benefited from the technological advances in GNSS technology. In particular, GNSSDOs have become ubiquitous in applications that require accurate and stable time and frequency references synchronised to UTC, from lab primary references to telecommunication base-stations. This factor, along with their low-cost, and inherent geospatial referencing, suit the requirements for multistatic radar, specifically when operated on baselines where all-in-view TFT can be exploited.

Despite these advantages, GNSS is susceptible to denial due to jamming and signal interference. This is particularly a concern in the defence community, where low cost jammers are simple to construct, and are even commercially available. A loss in synchronisation between the radar nodes leads to proportional range and frequency drift. Mitigation of damage due GNSS-Denial is therefore a necessity, and thus resilience to it is a hot-topic, with solutions ranging from using atomic clocks for holdover to rebuilding a modern LORAN system. Alternatively, one could use a fibre-optic time and frequency network, assuming fixed infrastructure is available. However, such systems are often costly and use proprietary hardware, making it difficult to build the TFT system around the radars requirements.

This thesis explores the use of WR technology, which offers a promising opportunity for low-cost, TFT for networked radars. WR is a fully open-source and open-hardware network protocol designed by CERN and its partners, which provides time synchronisation at the sub-nanosecond level. Whilst initially designed to meet the requirements of CERNs LHC timing network, it has since been used for hundreds of projects, ranging from international UTC comparisons to interferometric array synchronisation. It is therefore hypothesised that WR can be used in the multistatic radar domain, albeit with potential alterations to the hardware to improve the high frequency instability of standard WR equipment.

Chapter 2 provides a literature review on the present state of TFT systems in the context of multistatic radar. This chapter serves three primary goals, namely elucidating the time and frequency requirements for multistatic radar; discussing the past and current trends for TFT, with example bi- and multistatic systems; and a discussion of WR, focusing on the current state of the technology, its primary use-cases and limits to its performance.

TFT has a long and rich history in the fields of metrology, distributed arrays and communications. However, likely due to focus on monostatic radar systems, there has been little published directly on the topic for bi- and multistatic radar. Nevertheless, a few rules of thumb have emerged as guides for multistatic radar designers. Willis [44] laid the foundation for the commonly used first order approximations of bistatic time, frequency and phase synchronisation, as well as the required phase stability for coherent processing. A further summary of synchronisation for bistatic radars is given by Weiss [45], additionally discussing platform motion compensation. The authors state that range estimation requires time accuracy on the order of a tenth of the compressed pulse width, whilst phase stability is required up to a quadrature phase error over the CPI. Sandenbergh further developed synchronisation requirements, modelling bistatic pulse-Doppler radars with imperfect synchronisation. Using his model he was able to derive expressions for errors in phase, Doppler and range as functions of the MO synchronisation, and predict the effects of phase drift on coherent and non-coherent pulse integration. The author showed both analytically and with real radar data that given adequate line-of-sight (LoS) and synchronised receivers, the direct path signal can be used to reduce oscillator induced phase errors in the bistatic node.

The review then discusses methods for TFT for bi- and multistatic systems. The discussion is divided into wired and wireless approaches. For wired systems, two transfer mediums are considered - low-loss coaxial cable and fibre-optics. Fibre systems are particularly appealing due to their ultra high bandwidth, low attenuation and EMI resilience. Wireless systems are then discussed, which offer greater flexibility in terms of node placement and manoeuvrability, but come at the cost of increased susceptibility to EMI related issues. Wireless systems can further be separated into several overarching technologies, specifically: atomic clocks, microwave links, GNSS systems and direct and Scattered RF TFT. Whilst every method has its merit, current literature has pursued GNSS based TFT due to its low-cost and simplicity, especially when nodes are placed on large baselines with minimal line-of-sight.

The most comprehensive study on the use of GPSDOs acting under all-in-view time transfer for coherent multistatic radar is given by Sandenbergh. Sandenbergh synchronised the three node multistatic radar, NetRAD, using custom designed GPSDOs. His system aided in the design of a novel adaptive quick-locking phase-locked loop (PLL) filter, and a sophisticated triggering mechanism synchronised to the UTC second transition. The peak to peak error in triggering two nodes was on the order of 6 ns for a pair of co-located GPSDOs.

Lastly, the review focuses on the WR protocol and hardware, looking at current implementations of the technology in industry, and its limitations in the radar context. In essence, a standard WR network uses the WR protocol to synchronise a hierarchical Gigabit Ethernet network to within a nanosecond of the Grandmaster clock, which is usually a GNSS or atomic clock. The network can consist of thousands of nodes (switches and end-nodes) with links between nodes typically spanning up to 10 km using single-mode fibre and bidirectional small form-factor pluggable (SFP) transceivers, although long-range links are also possible [46]. The protocol was developed using existing networking standards, specifically PTP/IEEE1588 for coarse time transfer, and L1 syntonisation similar to synchronous ethernet (SyncE). The sub-nanosecond accuracy is achieved due to the inclusion of a novel digital variant of the digital DMTD (DDMTD) circuit, which makes sub-period phase measurements possible, improving the time resolution to the picosecond regime.

The technology has been adopted, or is being investigated, by hundreds of organisations [47] requiring synchronisation in the nanosecond regime. Whilst there are alternative systems offering similar, if not better TFT performance, mass adoption is likely due to the open source/hardware nature of the project, scalable and flexible design, commercial availability, and backwards compatibility with existing PTP hardware. Lipinski et al [48] catalogued current and future applications based on a survey of the WR users section of CERNs open hardware repository. Their report found six major categories, namely: TFT, time triggered control, precise timestamping, trigger distribution, fixed latency data transfer and RF-over-WR. In relation to remote sensing, all of these applications show great potential for coherent multistatic systems.

The performance of WR synchronisation is ultimately limited by the calibration of hardware delays, path length instabilities, additive phase noise in the DDMTD field programmable gate array (FPGA) circuitry and the performance of the Gigabit transceivers. Additionally, the oscillators used in the standard WR switch are of low quality, thus a low-jitter daughterboard is available commercially for enhanced jitter performance, improving the Allan deviation (ADEV) from 1×10^{-11} to 1×10^{-12} ($\tau = 1s$) with cumulative root mean square (RMS) jitter improving from 11 ps to 1.1 ps over a 100 kHz bandwidth. WR has also recently been shown to be achievable over sub-millimetre point-to-point (P2P) links [49], opening up the possibilty for WR TFT without the need for fixed infrastructure.

Chapter 3 details the design of a WR network for use in an operational multistatic radar, paying attention to the requirements posed by a real system, NeXtRAD. NeXtRAD is a fully polarimetric X- and L-Band radar designed by UCT and UCL. The radar is composed of three nodes, an active monostatic node, and two passive receivers, that are connected in an Ethernet network using WiFi links. Each node has a TF system, that is composed of a timing and control unit (TCU), frequency distribution unit (FDU), and GPSDO. The GPSDOs, which are responsible for synchronisation and syntonisation in NeXtRAD, output a low phase noise 10 MHz frequency reference from an Oscilloquartz 8788 SC-cut single oven OCXO, and are controlled by a Matlab graphical user interface (GUI) that allows a software defined PLL to be implemented. An output trigger is programmed to fire at selected UTC time, which is distributed from a command and control (CNC) node.

The section then focuses on NeXtRAD's TF requirements. For time synchronisation, the TFT must synchronise the nodes to an accuracy below 2.2 ns, with a holdover stability of 1.1×10^{-9} for a 1 s holdover. The relative frequency error between the transmitter and receiver for coherent processing is required to be below 1.94×10^{-10}

The chapter then details the design of a WR based TFT network, paying attention to NeXtRADs specific TFT requirements. Two WR options were considered: a fully commercial off-the-shelf (COTS) WR system using Fine Delay (FD) cards; and a hybrid WR and GPSDO design – a White Rabbit disciplined oscillator (WRDO). Based on preliminary research, the lower-quality TCXOs within the FD nodes were hypothesised to have unsatisfactory short-term stability compared to the OCXOs available in the GPSDO system. On the other hand, the relative long-term stability is comparable to a Caesium clock, and thus cleaning up the short-term noise with a high quality oscillator was hypothesised to improve the close in noise performance. Thus the second option, the WRDO, was designed such that an alteration to the previous UCT GPSDOs would allow them to use WR as the disciplining mechanism. Essentially, two GPSDOs are provided an external PPS signal from individual WR nodes, as opposed to their respective GPS receiver's PPS. The rationale behind this is three-fold: firstly, the WR PPS has picoseconds of jitter; secondly, unlike a GNSS PPS, there is no saw-tooth error to correct for in the phase detection, leading to lower noise from quantisation; and lastly having WR as an external reference, implies that with a few alterations to the control circuitry/software, it should be possible to use the GPS receiver as a redundant timing source if failure to the WR system is detected, as well as integrating the GPS spatial referencing into the design.

The WRDO therefore combines the short-term stability from the OCXO of the GPSDO, but improves the frequency stability with respect to the GPSDOs at longer averaging times. Using a time interval counter (TIC) method, the GPSDOs had been previously evaluated by Sandenbergh to have an optimum damped PLL loop time constant of 1000 s. The optimum WRDO PLL time constant was an unknown, but hypothesised to be on the order of a 100 s. Two loop constants were investigated for the WRDO configuration - 10s and 100s. Four TFT options were therefore evaluated: UCT GPSDOs in an all-in-view configuration, WR fast mezzanine carrier (FMC) and WRDOs with PLL loop-constants of 10 s and 100 s. However, NeXtRAD uses 100 MHz MOs for both the transceiver and passive receivers, thus the 10 MHz outputs were further used to steer ultra-low noise Wenzel PLOs with a PLL bandwidth of 5 Hz.

Chapter 4 discusses the measurements of the TF references in a laboratory setting. Whilst trigger accuracy ultimately determines the accuracy of the start pulse between each node, the systematic time offsets can be nulled to a resolution of 1 ns using settable delays in NeXtRAD's TCU and GPS reveiver. The main predictor of radar performance is then the relative phase and frequency stability between the transmitter and receiver MOs. A dual mixer time difference (DMTD) technique was used to measure the phase difference between two units with femtosecond level resolution, where the reference oscillator was set to 1 Hz above the 100 MHz MOs. The quantitative metrics chosen for this evaluation were the phase, fractional frequency, and TF stability metrics, notably the overlapping Allan deviation (ODEV), modified Allan deviation (MDEV), time deviation (TDEV) and maximum time interval error (MTIE). The data were processed using Python and the AllanTools library [50].

The test bed was first calibrated to characterise the noise floor of the DMTD system. The measurement noise is more than an order of magnitude lower than the frequency stability of any of the references. Time-domain frequency stability metrics were then computed for each experiment to determine the TF stability for each reference. The WRDOs show excellent short-term frequency stability, comparable to the GPSDOs - 1.9×10^{-11} vs 2.2×10^{-11} at 1 s averaging times. There is also an improvement in frequency stability for long-term averages - 1.3×10^{-14} vs 2.3×10^{-12} at 1000 s. The drift of the WR based TF systems are typically below 1 ns, in comparison to the GPSDOs, which have a peak-peak drift of nearly 10 ns during the experiment. A comparison to the NeXtRAD time and frequency requirements presented in chapter 3, shows that a WRDO system comfortably fulfils the TFT requirements for short-term coherent multistatic radar.

Given that the results presented in chapter 4 achieved the multistatic requirements, chapter 5 discusses the use of a WRDO network for TFT in a real multistatic radar. The radar geometry was set in a quasi-monostatic bistatic configuration on a test facility in Simon's Town, South Africa. This geometry was chosen to avoid unspooling a 5 km reel of fibre to form a larger baseline, and to accommodate the DMTD measurement equipment. However, owing to the fibre length, the timing baseline between the nodes is 5 km, despite the nodes being co-located. Only two nodes were used due to available resources, but the principle easily extends to the multistatic case. The target of interest was an offshore lighthouse, called Roman Rock, roughly 1.83 km away from the radar.

Two radar PRFs were chosen to inspect the short-term and long-term phase performance of the radar under GPSDO or WRDO. A medium 1 kHz PRF was used to verify the phase noise and coherent integration loss of each TF system. The bistatic receiver phase noise was higher than the monostatic node due to the uncorrelated STALOs and PRF aliasing. Spurious powerline modulation is significantly higher in the bistatic receiver. The initial frequency error and frequency drift between the transmitter and receiver was higher than predicted, leading to a coherent integration loss of 3 dB within 2 s. However, this does not reflect the typical GPSDO performance when fully synchronised and syntonised. The error was likely due to setup time, or transient GPS RFI. A low 100 Hz PRF was also selected to analyse the long-term phase phase and frequency stability. Simultaneous measurements of the MO phase errors and Roman Rock phase were recorded, to show the influence of the MO frequency stability on the bistatic phase. There was surprigingly more frequency noise in the WRDOs, compared to the GPSDOs, but this is likely due interference in the single-ended to differential converter. No transients were found in WRDO, leading to a peak-peak time error of 0.64 ns, compared to the 7.64 ns of the GPSDOs.

The thesis concludes with 6, where the major results are discussed and recommendations for future applications are given. This thesis did not delve into using WR for data transfer, and an expansion to the current NeXtRAD system could benefit from the Gigabit network, albeit purely for CNC operations. Development of a complete WR radar is the next logical step, where the entire RF system could be implemented using RF-over-WR, potentially opening up the door for fully coherent multistatic MIMO radar.

Chapter 2

Background

This chapter reviews and summarises the literature required for the design and implementation of a WR based TFT network for multistatic radar. The chapter has three important goals in the context of this thesis, namely: 1) determining the time and frequency requirements for coherent multistatic radars, paying attention to phase and time synchronisation, and frequency syntonisation; 2) providing an overview of current methods of TFT, and reviewing bi- and multistatic radars that explicitly discuss their TFT system; and, 3) reviewing the WR protocol as used in other fields, the limitations and future of the protocol, and why WR may be an excellent candidate for multistatic radar.

The scope of this review only considers the TF properties of bi- and multistatic radars, and consequently to the spatial coherence of the network. Radars with platform motion introduce additional complications, which will not be considered in this thesis. The performance of the information fusion depends on the coherence of the network, however these algorithms will not be discussed either, as they are outside the scope of this work. For a detailed treatment of coherent multistatic data fusion, multistatic ambiguity functions, and optimal detection algorithms, see references [51–54]. Additionally, in an operational system there will be requirements on the data network and multistatic processor, which will also not be covered.

The review will first discuss TF requirements for a fully coherent multistatic radar, treating the multistatic radar as a network of monostatic and bistatic radars. Coherent multistatic radars require time synchronisation for bistatic range estimation, frequency accuracy, and phase stability during coherent intervals for coherent integration gain and Doppler processing. Because the concept of TF stability is important to this section, the reader is encouraged to see Appendix C for an understanding of phase noise and the Allan variance.

Following these requirements, an overview of TFT methods and techniques is given, highlighting their potential strengths and weaknesses in the multistatic radar context. Where applicable, operational systems employing each method are discussed. Both wired and wireless systems are reviewed. Lastly, the WR protocol is reviewed, and in particular, the review focuses on the limitations of WR, and how WR offers the necessary characteristics for TFT in a fully coherent multistatic radar.

2.1 Multistatic Time and Frequency Requirements

The theory and design of multistatic and netted radars, or MSRS (using Chernyak's definition [51]), has been an active topic for several decades, and has therefore generated a number of seminal books and articles. In particular, Willis provides the quintessential text on bistatic radar theory [44], whilst Chernyak extends the theory into MSRS's [51]. Derham [52], and later Doughty [55], design and develop a coherent multistatic radar called NetRAD, which is the precursor to the radar used in this study.

Multistatic radars typically differ in scope and architecture, but as discussed previously in chapter 1, have several common design considerations, namely: if there are one or more dedicated transmitters for the network; the number of nodes and their geometric arrangement; the level of cooperation between nodes; the type of spatial coherence and information fusion; and the TF requirements of the multistatic sensor.

Willis [14], and later Weiss [45], provide an overview of synchronisation for bistatic radars, and give first-order estimates of the suggested TF accuracy and stability between the transmitter and receivers. Sandenbergh extends this theory by modelling a bistatic radar, and provides a practical application of this model using a network of multi-channel GPSDOs, which were used to synchronise the NetRAD radar [1, 56, 57]. Other authors, notably Herbert [58], Wang [59], and Krieger [60] have modelled and simulated the effects of synchronisation systems. Wang in particular, provides a thorough treatment of the phase errors due to an active synchronisation system, discussing numerous sources of phase error, and a solid stochastic model to predict the impact of the synchronisation system on the performance of a bistatic SAR (BiSAR) system.

2.1.1 Multistatic Model

In figure 2.1 the multistatic model used in this thesis is depicted. The model is comprised of a monostatic and bistatic radar acting cooperatively, connected via a data communications network and some form of TFT network. The model can easily be scaled to any number of active or passive receivers.

In this model, each node has it's own TF reference, which has a MO that produces a constant and stable tone at some lower frequency, typically 5-100 MHz. The MO is used as the clock timebase, which serves as a reference for the digital systems, and iterates a real-time clock (RTC) for absolute time. The RTC of the master reference, which in most cases is a transmitter, represents the network time. Whilst not always required, the network time can also gain traceability to a civil standard or UTC by synchronising to a local timeserver. Each other node periodically synchronises to the network time using a time transfer technique, giving each node in the network a notion of absolute time. Because the clock is typically derived from the MO, the phase error between the nodes is directly related to the time error. Time error, $x(t)$, is therefore synonymous with phase error, and the time error of frequency source is related by the phase change

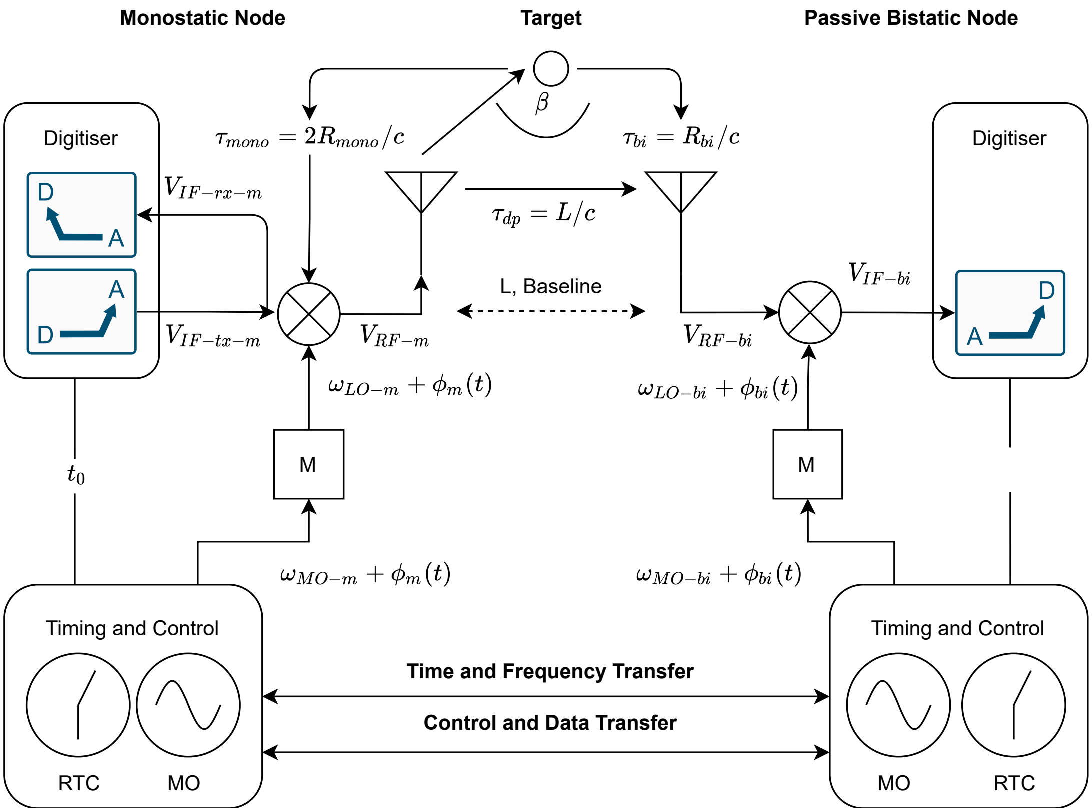

Figure 2.1: A model of a multistatic radar.

between the MOs, by scaling the phase error by the nominal frequency.

$$x(t) = \frac{\phi(t)}{\omega_0} = \frac{\phi(t)}{2\pi f_0} \qquad [\text{s}] \qquad (2.1)$$

The MOs are also used to synthesise the higher frequency STALOs. A simple model of the STALO frequency, ω_{STALO}, is just a multiplication, M, of the MO frequency. However, the phase noise of the MO increases as the square of the multiplication. As general rule, the phase noise spectral density due to multiplication is,

$$S_{\text{STALO}}(f) = S_{\text{MO}}(f) + 20\log_{10}(M) \qquad [\text{dBc/Hz}] \qquad (2.2)$$

That is, a multiplication of 10 in frequency, causes the output to have roughly 20 dB higher phase noise than the reference.

The time/phase error is then used to calculate the fractional frequency error, $y(t)$. The fractional frequency is a unit-less measure of frequency error normalised by the nominal carrier frequency. Strictly speaking, we measure the average fractional frequency, which is the rate of change in the time error during an averaging interval, τ. The averaging interval is based on positive integer multiples of the sampling time, τ_0.

$$y(t, \tau = m\tau_0) = \frac{x(t+\tau) - x(t)}{\tau}, m \in \mathbb{N} \qquad [\,] \qquad (2.3)$$

This has practical importance, because one can then calculate measures of time-domain TF stability for the MOs and STALOs using the phase or frequency, as explained in Appendix C.

2.1.2 Time Accuracy

A monostatic radar, can be modelled as a perfectly synchronous device, and subsequently only has a range induced error based on the frequency error and sampled instant, which is related to the signal-to-noise ratio (SNR) of the received pulse.

In bistatic radar, the range error additionally includes the synchronisation error between the transmitter and receiver. The error is directly proportional to the synchronisation error, $\Delta\tau$,

$$\Delta R = c\Delta\tau \qquad [\mathrm{m}] \qquad (2.4)$$

From Willis [44], a typical first-order requirement for bistatic radar is that the synchronisation error is a 10^{th} of the compressed pulse-width, τ_{PC}.

$$\Delta\tau = \frac{\tau_{PC}}{10} = \frac{1}{10B} \qquad [\mathrm{s}] \quad (2.5)$$

Where the compressed pulse-width is a function of the transmitted pulse bandwidth [61], B. The clock's frequency stability during holdover, $\sigma_y(\tau_h)$, needed to maintain sufficient range accuracy is then a function of the receiver clock's update interval, or oscillator holdover period, τ_h, with respect to (w.r.t) the transmitting node. If direct synchronisation methods are used, the required stability is,

$$\sigma_y(\tau_h) = \frac{\Delta\tau}{\tau_h} \qquad [\,] \quad (2.6)$$

A factor of 1/2 is included for the more stringent case of indirectly synchronised nodes, such as GPSDOs, where the transmitter and receiver each require the stipulated frequency stability during the holdover period.

$$\sigma_y(\tau_h) = \frac{\Delta\tau}{2\tau_h} \qquad [\,] \quad (2.7)$$

Drift Prediction

Equation 2.7 assumes that there is no initial time offset (perfect synchronisation), and that the time drift is purely due to a constant frequency offset. Sandenbergh therefore includes the time drift based on the initial time offset, x_0, and the predicted time error accumulated during the holdover period, τ_h,

$$\Delta\tau \leq x(t_0, \tau_\mathrm{h}) = x_0(t_0) + x_\mathrm{drift}(\tau_h) \qquad [\mathrm{s}] \quad (2.8)$$

Where the predicted time drift $x_\mathrm{drift}(\tau_h)$ includes the initial frequency offset, y_0, temperature stability, $\Delta T K_T$, ageing, D, and oscillator time stability $\sigma_x(\tau_h)$.

$$x_{\text{drift}}(\tau_h) = y_0\tau_h + \Delta T K_T \tau_{\text{h}} + \frac{1}{2}D\tau_h^2 + \sigma_x(\tau_h) \qquad \text{[s]} \qquad (2.9)$$

This equation also has practical importance in predicting the time and phase drift due denial of the synchronisation system, given the ageing rate of the holdover oscillator.

2.1.3 Frequency Accuracy

An absolute frequency error with respect to UTC creates a proportional Doppler error. Relative errors in the frequency between the transmitter and receiver create an apparent Doppler shift in the bistatic receiver. The requirement on the relative bistatic frequency accuracy in therefore depends on the required velocity accuracy, V_{acc}. From [1], the maximum allowable frequency error, y_ϵ, required is,

$$y_\epsilon \leq |V_{\text{acc}}/c| \qquad [\,] \qquad (2.10)$$

2.1.4 Frequency Stability

Short-term Stability (Phase Noise)

Detection of moving targets in stationary clutter is referred to as the SCV, and is directly dependent on the receiver phase noise. In the baseband frequency-domain, phase noise spreads the clutter power into non-zero Doppler bins, and therefore places a bound on the detectability of slow moving targets in strong clutter environments. In scenarios where the Doppler shift is on the order of a few Hertz, the target's power spectrum can easily be swamped by the high levels of close-in phase noise. This is particularly challenging in high clutter environments with slow moving targets, such as small boats in high sea-states.

In monostatic radars, the use of a single STALO significantly assists in reducing the phase noise seen at the receiver's intermediate frequency (IF) due to its self-coherence. The receiver phase noise is based on the return-trip time and the phase noise of the STALO, where low frequency phase noise components are cancelled within the signal delay, due to a first difference in the phase. Auterman [62] estimates the monostatic phase noise at the receiver IF as,

$$S_{\phi-\mathrm{IF-Mono}}(f) = 2M^2(2\sin\pi f\tau)^2\mathcal{L}_{\mathrm{MO}}(f) \quad [\mathrm{dBc/Hz}] \quad (2.11)$$

where $S_{\phi-IF}$ is the phase noise at the receiver IF, $S_{\phi-STALO}$ is the STALO phase noise, f is the phase noise frequency offset, and τ is the return-trip delay. The estimate can also be computed given the one-sided MO phase noise as often specified in oscillator data-sheets, $L_{\mathrm{MO}}(f)$, if the STALO is derived from frequency multiplication.

However, care should be taken in predicting the actual phase noise spectrum of the STALO based purely on the MO phase noise, as modern radars typically use frequency synthesis techniques to generate the carrier. STALO phase noise is then a composition of the multiplied MO phase noise within its loop bandwidth, but independent voltage controlled oscillator (VCO) phase noise outside this bandwidth. Nevertheless, using the simple frequency multiplication is often sufficient if a quick approximation is required.

In a bistatic receiver, the transmit and receiver STALOs are matched, but mutually independent, and therefore there is no cancellation of phase noise in the bistatic receiver. In fact, the phase noise adds, leading to even higher levels of measured phase noise, reducing the SCV performance in comparison to the monostatic scenario. The receiver phase noise in the bistatic case is computed as,

$$S_{\phi-\mathrm{IF-Bi}}(f) = 2S_{\phi-\mathrm{STALO}} = 4M^2\mathcal{L}_{\mathrm{MO}}(f) \quad [\mathrm{dBc/Hz}] \quad (2.12)$$

In both cases, most modern receivers directly sample the IF signal, which adds additional phase noise to the digital representation of the analog signal. The phase noise contribution from the ADC and thermal noise processes is generally smaller than the noise of the STALO. Nevertheless, a more complete description includes these terms. The total phase noise of the sampled IF, S_ϕ, therefore includes a contribution from the ADC phase noise, $S_{\phi-\text{ADC}}$, and thermal phase noise, $S_{\phi-\text{th}}$. The total phase noise is then,

$$S_\phi = \sqrt{S^2_{\phi-\text{IF}} + S^2_{\phi-\text{ADC}} + S^2_{\phi-\text{th}}} \qquad [\text{dBc/Hz}] \qquad (2.13)$$

Where the ADC phase noise is a combination of the ADC sample jitter and reference clock jitter converted to phase noise by integrating over the baseband bandwidth, $S_{\phi-\text{sample}} = S_{\text{jitter}} + S_{\text{clk}}$ [63]. The ADC phase noise is further increased by the ratio of the IF, f_{IF}, to sample clock frequency, f_s.

$$S_{\phi-\text{ADC}} = S_{\phi-\text{sample}} + 20\log\left(\frac{f_{\text{IF}}}{f_s}\right) \qquad [\text{dBc/Hz}] \qquad (2.14)$$

Doppler Processing

For modern Doppler processing, the sampled data are reshaped into a data matrix, where the slow-time dimension represents subsequent pulses in the same range bin. Doppler processing therefore effectively resamples the data at the much slower PRF, which causes significant aliasing. Aliasing folds the phase noise spectrum from the PRF to the much higher receiver bandwidth, where the number of folds, n, is the ratio of the receiver bandwidth to the PRF, or $n = \text{BW/PRF}$ [64]. The phase noise, in dBc, seen in the Doppler spectrum, $S_D(f)$, is then,

$$S_D(f) = S_\phi(f_d) + \sum_{i=1}^{n} S_\phi(f_d + i \cdot \text{PRF}) + \sum_{\text{i=1}}^{\text{n}} \text{S}_\phi(\text{f}_\text{d} - \text{i} \cdot \text{PRF}) \qquad (2.15)$$

Where $S(f_d)$ is the Doppler bin of interest, and we include spectral folds above and below the PRF.

Coherent Processing Phase Stability

Coherent pulse integration, Doppler processing, and synthesis imaging require phase stability over the CPI.

In a monostatic radar, phase errors during the CPI are not limited by phase drift due to the single STALO. Instead, any loss in coherence is due to non-cancelled phase noise components, which are assumed to be white[1]. Integrating multiple coherent samples, leads to the white noise averaging away to some fixed value. The reduction in coherent gain, known as the coherent integration loss, L, can be evaluated analytically as the sum of pulses in Gaussian white noise. However, it can also be more simply approximated using Richards' approximation formula, given that there are tens of pulses to be integrated [65]. The approximated coherent loss formula is,

$$L = -4.343\sigma^2, \; N >> 1 \qquad [\text{dB}] \qquad (2.16)$$

where σ^2 is the white phase noise variance.

In bistatic radar, the non-target induced phase drifts as a function of the relative phase stability of the MOs. A first order requirement for the RMS phase change during the CPI ranges from $1/10^{th}$ of the RF wavelength, λ or 3.6°, to a maximum of $\pi/2$ for modern processing schemes [14]. The choice of maxumum phase drift depends on the application, where BiSAR applications typically have the most stringent requirements. The required maximum frequency error during the CPI, $\sigma_y(T_{int})$, is then,

$$\sigma_y(T_{int}) = \frac{\Delta\Phi}{2\pi f_C T_{\text{int}}} \qquad [\;] \qquad (2.17)$$

where $\Delta\Phi$ is the maximum RMS phase error, f_C is the RF carrier frequency and T_{int} is the CPI. However, this formula assumes a constant drift, or fixed offset in frequency, which is not typically a realistic scenario, given frequency instability.

---

[1] This assumption is valid under normal radar integration periods.

Sandenbergh therefore models the maximum allowable phase drift in units of time for a bistatic radar with RF frequency f_c, of the MOs as,

$$\Delta\tau_\Phi = \frac{\Delta\Phi}{2\pi f_c} \qquad [\mathrm{s}] \qquad (2.18)$$

Where $\Delta\tau_\Phi$ is the maximum allowable phase drift of the MOs, based on the predicted time-domain phase deviation over the integration time, $x(\tau_{\text{int}})$.

$$\Delta\tau_\Phi \leq x(\tau_{\text{int}}) \qquad [\mathrm{s}] \qquad (2.19)$$

Where, similar to equation 2.9, the phase drift of the MOs in time is,

$$x(\tau_{\text{int}}) = y_0\tau_{\text{int}} + \Delta T K_T \tau_{\text{int}} + \frac{1}{2}D\tau_{\text{int}}^2 + \sigma_x(\tau_{\text{int}}) \qquad [\mathrm{s}] \qquad (2.20)$$

In practice however, assuming a constant offset during the short integration times of a radar may be acceptable. However, CPIs that last longer than a few seconds for applications such as passive coherent location (PCL) radar, radio astronomy, or multistatic radars operating with sub-Hz Doppler resolution, should consider equation 2.20.

2.2 Time and Frequency Transfer

Accurate and stable synchronisation and syntonisation is not only crucial to multistatic radars, but is also a requirement in various economic, industrial, and scientific fields. For instance, fifth generation (5G) communication standards have more stringent time interval and jitter requirements than current standards [66]; whilst, the current European Securities and Markets Authority (ESMA) standard requires stock markets to be synchronised to UTC with an accuracy of less than 1 us [67]. In line with the goals of distributed radar, TFT is fundamental to the operation of other microwave sensors like radio astronomy arrays and long-baseline interferometers, such as the Very Large Array (VLA) and European VLBI Network (EVN). In particular, the upcoming Square Kilometre Array (SKA) requires ultra-stable frequency distribution from active hydrogen masers to hundreds of antennas, as well as nanosecond accurate timestamping on timescales of decades for long-term pulsar timing [68].

TFT describes the process of transferring a reference value of TF from a system master source to a remote receiving site, over some connecting medium. After a time transfer, if the two sites are closely aligned in time they are said to be synchronised, and if their frequencies are closely aligned, they are said to be syntonised. Phase synchronisation of the microwave carrier will also be needed for fully coherent systems. The *accuracy* of time or frequency transfer is related to the mean offset in time or frequency between the two sites after the transfer. The *stability* is related to the uncertainty, or variance, in the data as a function of time.

2.2.1 TFT Methods

There are three commonly used methods for transferring a reference value time and/or frequency from a master site to a remote slave location. This section briefly summarises the methods, to assist in section 2.2.2. For a detailed review of these methods, including uncertainty estimates, see Levine [69].

Syntonising the frequency of a remote clock is somewhat easier than time transfer, as one can simply transfer the frequency directly for the remote, if neither time nor phase information is of importance. However, the channel includes sources of noise and attenuation, that will reduce the signal SNR, and increase the the level of phase noise seen at the slave site. This issue can be mitigated with the use of a narrowband PLL, locked to a sufficiently stable oscillator. The PLL bandwidth can be chosen, such that the remote end averages the short-term channel fluctuations, and effectively tracks the long-term frequency stability of the master reference. Time transfer, which is used in synchronisation and clock comparisons, is more involved, and requires that one can transmit a reference value of time from the master site to the remote, slave site.

One-way methods use knowledge of the geometric and propagation delays from a master time source to the receiver. In one-way time transfer, the master transmits its local value of time, which then propagates through a noisy channel to the slave site. This requires that the signal's propagation can be accurately modelled, taking into account the propagation velocity, effective path length, and dynamics of the sites. The receiver can then simply add the calculated delay to the received time value. One-way methods are effective given a comprehensive model of the system and environmental delays, making them ideal candidates for GNSS TFT. In theory, the phase of the remote site could be synchronised to the master, as the time value is sent at the start of each cycle, and one has an accurate model of the effective path length, which equates to an integer number of cycles and a fractional cycle, which is the estimated phase. However, one-way methods do not deal with unpredictable path length changes, which limits the stability of the TFT process, and makes it impractical to synchronise phase with microwave carriers.

Two-way methods exploit the reciprocity of the transfer channel, by assuming that the path length from the master to the slave is the same path from the slave to the master. Two-way methods only require that each node can transmit and receive time packets, and do not require an exact model of the path delays. In a two-way time transfer[2], a node, typically the master, initiates a time transfer by broadcasting its local time. The slave records this time and responds with its own local time. The path induced delays are assumed to be the same in both directions, thus the master clock can calculate the time offset and one-way delay, which is assumed to be half the round-trip delay. The master can then send back a time update with the new time value, compensating for the one-way delay. Continuous two-way exchanges allow the system to calculate and track changes in the path delay, which reduces instability. The network time protocol (NTP) and PTP are examples of two-way time transfer.

---

[2] Here we are assuming a half-duplex system for brevity, but full-duplex two-way time transfer is achievable.

Common-view methods are used when there is a TF reference that shares common coverage with the the master and slave sites. The one-way path length induced errors are common to both sites, and therefore cancel out in a time comparison. A means of communication between the sites is still required for clock comparisons. Common-view methods using GNSS satellites are typically used by metrological institutes for UTC comparisons, and can be accurate to an order of a few nanoseconds for a well calibrated system.

2.2.2 TFT Techniques

Fixed infrastructure, specifically in the case of low-loss coaxial and fibre optics, typically offer simple and cost effective means of time and frequency transfer. Cable based systems are less prone to EMI and multipath, and are often useful during prototyping due to the availability of low-cost coaxial cable, specifically for standard MO frequencies in the 10-100 MHz range. As stated previously, a data network is also usually required for control and monitoring, and thus can benefit from the fixed cable installations.

However, the reliance on a cable based medium for TFT, does require additional infrastructure to be considered in the overall design, which may prove impossible for mobile systems or non-stationary radars. Additionally, other issues such as the increased cost of laying and housing the cable in a suitable environment should be considered. Nevertheless, such fixed systems can be installed for projects where the radar nodes are required to remain in place, such as a netted coastal defence radar.

LVDS and Coaxial Cable

Of the several electric cable types, the lowest loss for radar applications are typically low-loss coaxial (specifically for carriers higher than very-high frequency (VHF)). However twisted-pair cables, as found in standard Ethernet applications, are feasible for TFT over short baselines for the lower frequency MO and data transfer. If the nodes are placed on larger baselines (several kilometres), low loss coaxial cables are more applicable.

Frequency transfer is easily achieved with a direct coaxial cable between nodes, some low-noise gain, and a narrowband PLL to reject the thermally induced noise in the cable. Simple one-way time transfer can be used if the cable length and propagation velocity is known. Experiments can be initiated if the transmitter sends a pre-pulse to the receiver.

An example of multistatic synchronisation using low-voltage differential signalling (LVDS), is the NetRAD radar system [52, 55], designed and built by UCL and the UCT. Its initial design was configured for short baselines of several metres, using 50 m unshielded twisted pair (UTP) CAT6 cables. A low phase noise 100 MHz OCXO MO was distributed to each receiver from the transmitter using an LVDS line driver. The MO was multiplied up remotely using PLLs to generate the 2.4 GHz LOs. The CAT6 cable also carried a 100 MHz discrete synchronisation clock for each node's FPGA timing and control system, as well as an LVDS pair for the main bang epoch pulse. Whilst the phase was stable during operation, the authors found significant phase changes when the cable was displaced.

Fibre-optics

Fibre-optic cabling has several notable improvements over coax, specifically in terms of significantly reduced signal attenuation, ultra-wide bandwidth and immunity to EMI. Optical fibre cables also weigh appreciably less, and are more flexible than low-loss coax, making them applicable to weight sensitive applications and difficult terrain.

Fibre-optic techniques have been employed in telecommunication systems for more than 40 years and have a long history of transferring both digital and analog signals over thousands of kilometres. However, digital transmissions are more common, and thus a wealth of standards have been created for data applications.

If one would like to replace the RF cabling with fibre-optics, microwave photonics can be employed. Microwave photonics [70] and RF-over-fibre (RFoF) have developed tremendously over the past decade, meaning that it is feasible to replace receiver RF cables, allowing digitisation to occur further from the receiving nodes. RFoF techniques are already extensively used in Astronomy for RF or LO distribution, particularly for phased arrays [71]. Unlike copper, fibre does not suffer severe attenuation with increasing frequency, but is susceptible to high noise figures due to the modulation and detection process [72]. Commercial RFoF systems are available allowing the RF or LOs to be transmitted over several kilometres. Several RFoF techniques exist, such as detecting the beat frequency of two phase coherent laser sources spaced by the microwave frequency, or modulation of the laser's intensity using an internal or external RF modulator.

Unfortunately, fibre is also significantly susceptible to temperature instability and shock and vibrations, where expanding fibre induces phase instability due to path length variations and increased chromatic dispersion [73]. Distribution of an LO or RF frequency along the fibre would therefore incur significant additive phase noise and distortion if no compensation scheme were in place. Whilst temperature drift is generally a slow process, and can be heavily reduced by burying the cables approximately a metre underground, acoustic noise is a serious concern. acousto-optical modulator (AOM)s in a feedback system are therefore practically needed to reduce the additive acoustic phase noise [74].

With proper link stabilisation, fibre-optic systems have been demonstrated to provide the ultimate in performance for accurate TF dissemination. Using two-way optical time and frequency transfer (TWOTFT) techniques [75,76], one can disseminate accurate and stable TF over thousands of kilometres.

For instance, Liwczyski et all installed a 420 km TWOTFT system in Poland (based on a modulated 100 MHz Caesium reference) for UTC comparisons between Warsaw and Pozna. The dissemination was better than 1.5 ns for a one day average, which exceeds the calibration uncertainty of the GPS precise point positioning (PPP) time comparison method used between the two sites [76]. No published work on multistatic radar using TWOTFT is available, likely due to the current prohibitive cost and fixed infrastructure requirements. Nevertheless, the extreme distances that can be reached using fibre TWOTFT, and the performance that one could achieve, may be useful for future large baseline bistatic radars.

Direct Breakthrough or Scattered RF Signal

If an adequate LoS exists between the transmitter and receiver, one could use the direct breakthrough to synchronise the receivers [14]. However, these signals are often significantly attenuated and may suffer from various interfering effects such as multipath, dispersion or signal jamming. This technique can be used in both CW and Pulsed systems [77].

If the receivers are out of LoS, and the bistatic pair survey a common coverage area, the scattered RF signal from a clutter cell in common coverage can be used. Willis [14] briefly mentions using tropospheric scattering for radars out of direct LoS. However, this technique is likely only suited for over-the-horizon (OTH) radar [78] where the bandwidth is on the order of 10 kHz, and thus the stability requirements are significantly relaxed.

Atomic Clocks

Whilst not a TFT technique per say, atomic clocks are well suited to multistatic radar frequency syntonisation. However, atomic standards, like the venerable HP5071a, are notoriously expensive, and generally constrained to laboratories. The relative drift between atomic clocks will be considerably slower than even the best quartz reference, allowing the clocks to be resynchronised on longer timescales.

Miniaturised, chip-scale Rubidium and Caesium clocks are also now commercially available. Chip-scale versions are also commonly integrated into modern GNSSDOs for improved holdover during GNSS denial.

GNSS Systems

GNSS[3] based TFT has several advantages over alternate TFT methods, and are typically utilised in low-cost and mobile systems. All GNSS satellites are equipped with an ensemble of high performance atomic clocks, which are required in order to provide the time and frequency stability needed for terrestrial navigation. These atomic clocks are periodically synchronised and corrected by their respective timebase, which is further traceable to UTC [79].

A specialised GNSS reference specifically aimed at high-end time and frequency applications, known as a GNSSDO, locks onto a GNSS receiver's PPS or carrier phase signal using a narrowband PLL. The PLL, which may be hardware or software defined, steers the output oscillator using the phase error between its own VCO, which is typically an OCXO, and an averaged version of the receiver PPS. This allows the GNSSDO to exploit the inherent long-term time and frequency stability of GNSS, by steering the frequency drift the OCXO to the long term frequency accuracy of GNSS. GNSSDOs therefore exhibit the short-term stability of the OCXO within the PLL bandwidth, but have their long-term frequency instability (from ageing and random walk effects) improved by tracing the GNSS constellation's inherent Caesium stability. For these reasons, GNSSDOs have become integral devices in applications that require low-cost, stable time and frequency dissemination, ranging from telecommunication base-stations, laboratory primary references to smart power grids and financial timeservers.

GNSSDOs are therefore highly attractive solutions to distributed radars, simplifying the design of the time and frequency system. Firstly, they are inherently required to provide positional information, thereby automatically determining the position of the nodes with respect to a global geo-spatial reference frame; secondly, low-cost GNSS receivers are readily available from commercial vendors; and more importantly for TFT, the GNSS signal allows a simple and cost effective means of highly accurate and stable time transfer, traceable to UTC, that does not require any line-of-sight between the sensors. Additionally, the

---

[3] GNSS is an umbrella term for all satellite based navigation providers, specifically GPS, Global Navigation Satellite System (GLONASS), Galileo, Beidou.

output signals of commercial GNSSDOs are also standard crystal frequencies (1, 5 or 10 MHz) and a PPS, which integrates into most modern radar hardware. However, it should be noted that the PPS of standard GNSSDOs are typically generated by the GNSS receiver, and are not averaged versions generated by the GNSSDO. This is an important merit to note when deciding on a commercial system, as standard GNSS receivers have PPS timing jitter on the order of 100 ns [80].

Premium GNSSDOs use chip-scale atomic clocks for improved hold-over performance. If the GNSS signal is compromised, the frequency will drift slower with respect to that of an OCXO based GPSDO. Johnson [81] compares a number of commercial rubidium and crystal references for a multistatic system and quantifies the effects of GNSS signal loss. In a rubidium GPSDO, the author reported a frequency offset of 15 mHz after 24 hours of GNSS loss, which will not considerably affect Doppler performance. Range drift of a approximately 450 m over 2 days was found.

Operating a GNSSDO at every node in a multistatic radar can also allow the radar to exploit the common-view TFT principle, where the common sources of atmospheric instability between each satellite and the receivers are effectively cancelled by the common GNSS signal path to each device [69].

Wurman et al were likely the first to prototype a fully functional pulse-Doppler bistatic radar using a retrofitted WS88 radar as the transmitter and a a passive receiver [82]. Time synchronisation was carried out using a pair of GPSs and a start signal carried over telephone link.

One of NetRAD's goals, was to collect bistatic sea clutter, in order to develop a clutter database under various bistatic angles. However, the limitation of a 50 m baseline, impairs the maximum bistatic angle to near quasi-monostatic levels at observation ranges of only a few 100 m. Therefore NetRAD was upgraded with three independent GPSDOs in order to extend the baselines to several kilometres. The work was carried out by Sandenbergh, a doctoral student from UCT, and is extensively covered in his thesis and publications [1, 17, 56, 57].

TF synchronisation was achieved using multi-channel GPSDOs. In a co-located laboratory setting, the GPSDOs were shown to be synchronised to a peak-to-peak time error of approximately 6 ns [57]. On larger baselines, this value is expected to deteriorate as the receivers move out of the common transmission path, limiting the maximum baseline to several kilometres. NetRAD however, did not require baselines larger than a few kilometres, and thus the above result is a decent approximation of the performance of the synchronisation accuracy.

During field trials, Sandenbergh found that the GPSDOs would take too long to settle to their optimum time constant, τ, of 1000 s, and thus developed a quick-lock software defined filter (SDF) using Matlab [57]. The reduction in settling time proved invaluable for changing the node positions in a reasonable time for daily experiments, as well as allowing the filter to quickly dampen transients that cause the GPSDO PLL to unlock. Sandenbergh was also able to reduce a portion of the correlated phase in the bistatic node using the direct breakthrough as a phase reference.

Microwave and Optical Links

Microwave and free-space optical links can be used where sufficient line of sight is available. However, these systems are prone to intentional or accidental jamming and multipath. Terrestrial links in particular must have clear LoS, and typically be mounted high enough, where multipath can be mitigated. In space-borne systems, like the Tandem-X BiSAR satellites, one can make effective use of synchronisation pulses over a microwave link [83], allowing phase stability to 1° of RF phase.

Free-space optical links have improved in recent years, allowing synchronisation over turbulent air paths to femtosecond level accuracy. Bergeron et al developed a 4 km free-space TWOTFT link that was able to synchronise a microwave clock to an optical clock to within 13 fs. The link stability was on the order of 10×10^{-14} for 1 s averages.

2.3 The White Rabbit Project

The LHC, which is housed in the main complex of CERN, is one of the largest scientific instruments ever built. The LHC is part of a massive collider network, composed of thousands of nodes, with several detectors, accelerators, and control centres. The functioning of the network requires that each device in the network shares a highly precise and accurate notion of time, as well as a means to communicate and distribute data and control triggers. A network based approach suits the LHC's time and data requirements. However, standard network based time transfer solutions like NTP and PTP, are limited to time accuracies ranging from milliseconds to microseconds. The WR project was therefore proposed to overhaul CERN's old GMT network. This section will provide a brief overview of the WR project and relevant information for its use in multistatic radar. For a comprehensive discussion of the protocol see Wlotowski's dissertation [84].

2.3.1 White Rabbit Overview

WR, or WR-PTP, is a protocol that extends the standard PTP protocol, by utilising L1 syntonisation (similar to SyncE) and a DDMTD phase measurement, which enhances time transfer accuracy to the sub-nanosecond regime. Figures 2.3 from Rizzi et al [86] and 2.4 from Lipinski et al graphically demonstrate the WR synchronisation procedure between a master and slave WR devices. The protocol uses PTP to timestamp Ethernet frames at the physical (PHY) level of the Open Systems Interconnection (OSI) model, and L1 syntonisation using PLLs to lock each slave to the local master's downstream frequency, effectively creating a common notion of frequency for the network (125 MHz in the standard case of 1 Gb Ethernet).

A coarse PTP time synchronisation with a resolution of 8 ns is achievable using the nominal period of the 1 Gb Ethernet clock. The L1 link syntonisation between pairs of master and slave nodes is achieved by encoding the master clock with 8b/10b line encoding, and using a clock and data recovery (CDR) circuit to recover the link clock for the slave PLL. Unlike in a standard PTP network,

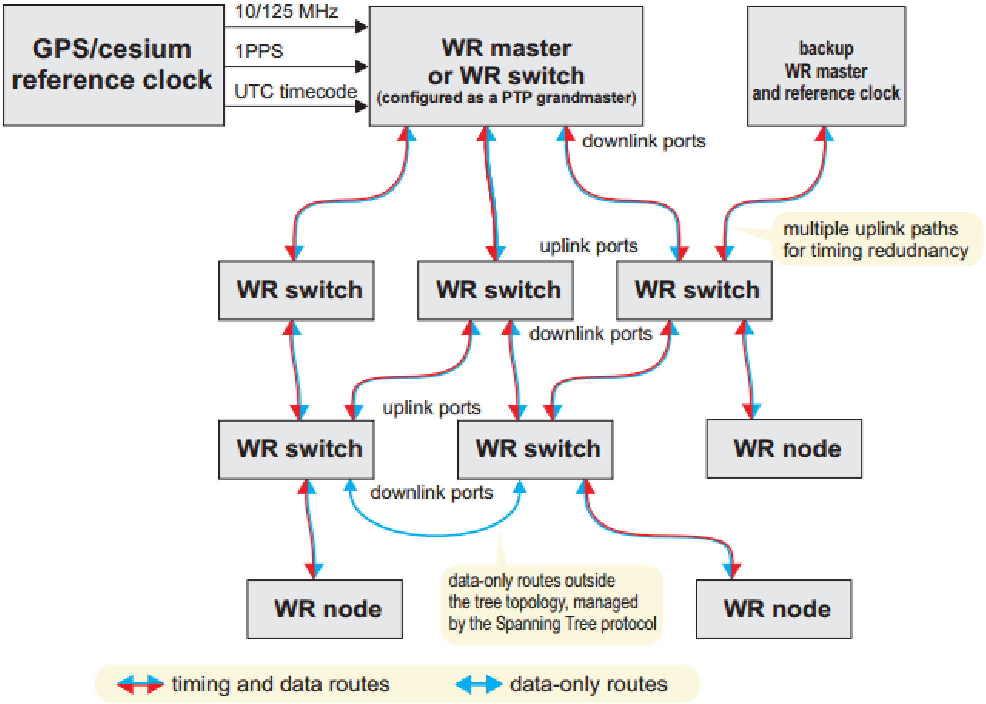

Figure 2.2: A simple diagram of a WR network showing the hierarchical clock tree (from [85])

where the rate of sync messages must be high in order to compensate for the drift induced by the free-running oscillators, the link syntonisation means that the rate of WR sync messages can be significantly reduced, increasing non-sync related network bandwidth.

Whilst WR guarantees sub-nanosecond synchronisation, the coarse gigabit clock resolution is almost an order of magnitude higher, meaning that the clock output would be ambiguously limited to a precision of 8 ns. WR improves upon the PTP resolution with the use of an FPGA DDMTD implementation [87]. In short, the DDMTD uses a phase comparator to measure the phase between the beat notes of the master and recovered slave clocks, each mixed against a common transfer oscillator. The phase is converted to units of seconds and added as a fine resolution component to the PTP calculation. The remaining time uncertainty is due to the clock jitter out of the slave's PLL, which is typically below 50 ps.

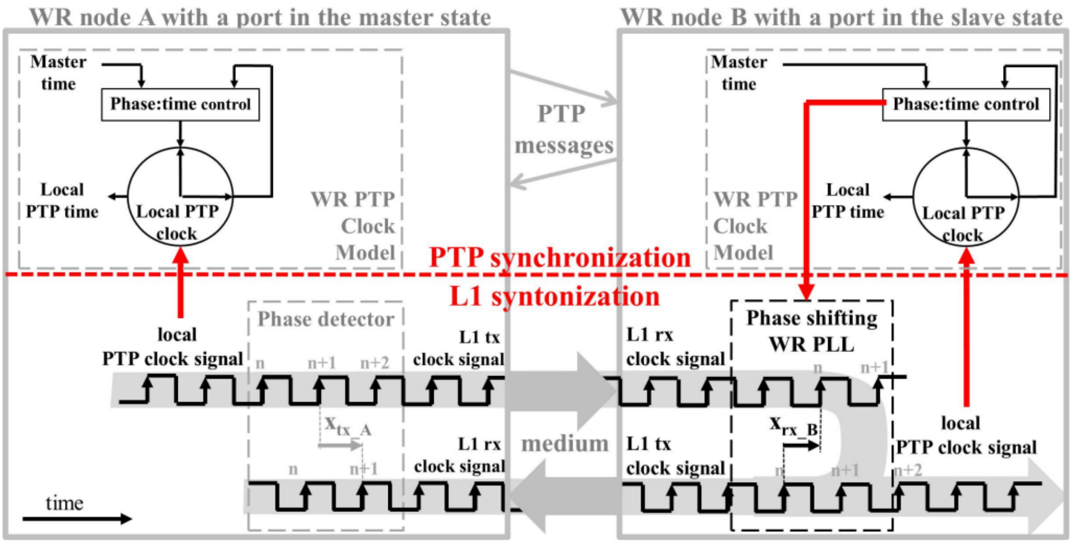

Figure 2.3: Synchronisation and syntonisation in the WR protocol (from [86])

Additional slaved nodes, or hops, are afflicted by higher levels of jitter due to PLL peeking [84].

On hardware, WR-PTP is made possible by an FPGA core called the WR PTP core (WRPC) [88], which is required in any device that wants to implement WR. A White Rabbit network (WRN) inherits PTP's sense of clock hierarchy, as seen in figure 2.2, where all nodes in the network can be thought of as clocks. The clocks' time accuracy is relative to the first node at the top of the tree, designated as the grandmaster (GM) clock. The clocks in the network are specialised WR switches and end-nodes (see section 2.3.2), known as boundary clock (BC)s and ordinary clock (OC)s respectively (as in a standard PTP network). OCs are clocks that are at the end of a network branch, and are used to output timing and control signals for the specific application. BCs are responsible for routing data traffic to the respective OCs, but adjust the transmitted timestamps to compensate for the routing packet delay.

The network is arranged in a tree topology to avoid time loops, and, as in PTP, uses the best master clock algorithm (BMCA) algorithm to select an optimal clock source, and allow for redundant sources of TF in case of failure of a timing source.

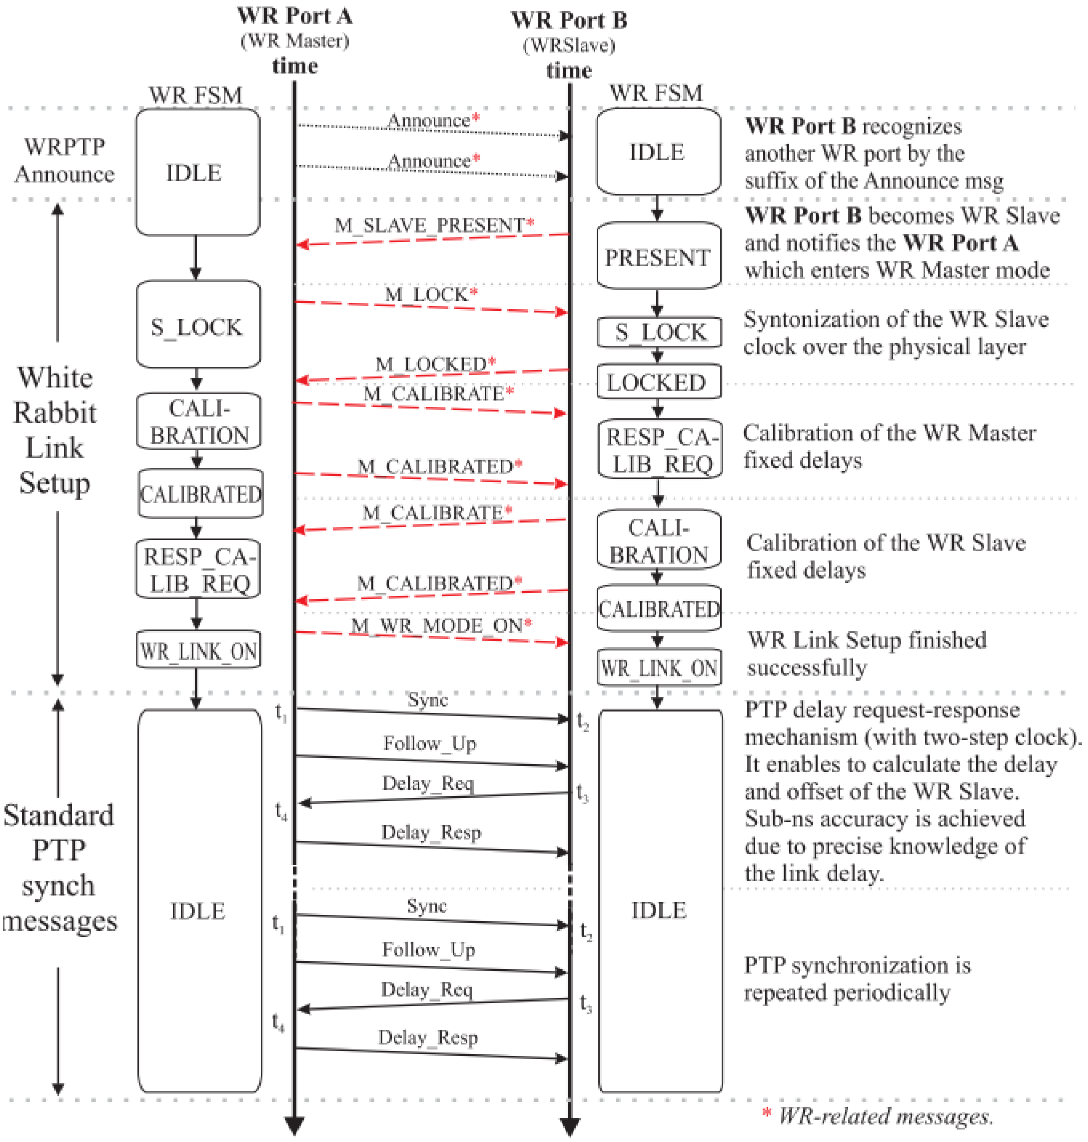

Figure 2.4: Synchronisation flow diagram for WR (from [85])

2.3.2 Hardware

Although WR is a protocol, it does require specialised hardware and software to be effectively implemented. Fortunately, the protocol, the hardware, and the software, are under CERN's open-hardware licence (OHL) policy [89] (currently v2.0). Hardware and software designs of the current generation of the White

Rabbit switch (WRS), simple PCIe carrier card (SPEC), and various FMCs are freely available in their respective repositories on CERN's open-hardware repository (OHWR) [35]. However, commercially available products based on the OHWR designs are available from several vendors, speeding up implementation.

A standard network is composed of WR compliant equipment, and in general contains the following components:

1. *White Rabbit Switch*: The WRS is an 18-port switch that implements WR synchronisation and facilitates Gigabit switch routing. The current switch makes use of a Xilinx Virtex 6 FPGA for the WR gateware, which includes the WRPC FPGA core for WR timing. The switch can act in grandmaster mode, where the switch becomes the primary time source for the network, or it can be slaved to other switches using the WR protocol. When in grandmaster mode, the switch requires a 10 MHz frequency reference and a PPS input. The reference clock's frequency stability and accuracy directly impact the TF performance of further WR nodes, and therefore low phase noise GNSSDO or atomic clocks are recommended. In most distributed hardware applications, only the relative synchronisation accuracy to the grandmaster switch or other nodes is important, thus traceability to UTC is of minor importance. However, in applications where the network must align to some timebase, the switch can utilise time of day information through an National Marine Electronics Association (NMEA) interface (included by most GNSSDOs), or using a local or internet NTP timeserver.

2. *Simple PCIe Carrier Card*: A WR SPEC [90] allows the use of specialised FMCs designed using OHWR guidelines to be used within a WRN. Communication to the SPEC is handled through a peripheral component interconnect express (PCIe) port, and can therefore be used in a standard personal computer. The SPEC contains a Xilinx Spartan 6 for WR, along with the necessary clocking resources, which allows it to include the WRPC core, and thus provide WR timing to the node.

3. *Fast Mezzanine Cards*: The OHWR repository hosts designs for several FMCs [91], which are responsible for interfacing with a user's hardware.

There are several WR FMC designs, ranging from ADCs and digital-to-analog converter (DAC)s, to general purpose input and output (GPIO) and programmable delays. This thesis uses the fine delay FMC, which is able to delay an input pulse to up to 4 output channels in the range of 600 ns to 120 s, with sub-nanosecond accuracy. However, it is also possible to generate programmable pulse trains, with frequencies up to 10 MHz [92].

4. *Single-mode Fibre and Gigabit Transceivers*: Single-mode fibre is used in a standard WRN as it provides the lowest attenuation and dispersion. For links up to 10 km using single-mode fibre, WR encourages the use of 1000BASE-BX10 bidirectional transceivers, operating with 1490 nm and 1310 nm wavelengths [93].

2.3.3 Applications

WR is currently being used in several projects, ranging from cosmic ray detectors to metrology institutes. An application study by Lipinksi [48] highlights the major use cases, as well as enhancements to the original system. Lipinski separates the applications of WR into five distinct categories:

1. *Time and Frequency Transfer*: The primary goal of any WR network is to distribute an sub-nanosecond accurate representation of the network time. Several national time institutes, such as MIKES are using WR to compare atomic clocks and disseminate accurate representation of UTC to clients [39].

2. *Time Triggered Control*: Time critical triggers need a tight degree of synchronisation to trigger simultaneously, or with accurately programmed delays. Time triggered control allows an ordered set of triggers to be pre-programmed in each node, which allows nodes to trigger within a nanosecond. Users of time triggered control generally include dispersed hardware, as is the case in complex collider networks. For instance the GSI replaced their GMT timing network with WR, and have been operating the WRN

since 2015. Lipinksi additionally includes trigger distribution as a separate class to discriminate between control and distribution, however for brevity they are similar in operation. Time triggered control may be of special interest in multistatic radar, where synchronised operations such as starting the PRF and digitisers, or priming high powered amplifier (HPA)s is required.

3. *Precise Time Stamping*: Accurately, and precisely determining the time interval between two spatially dispersed events is a common challenge in engineering, physics and economics. Therefore the most common use of WR is for its sub-nanosecond timestamping. Essentially, precise time stamping of incoming pulses or analog samples, allows events in the network to be related with sub-nanosecond accuracy. This is particularly important in dispersed senors, which attempted to localise the spatial position of an event. Financial institutions also require the timestamping accuracy for traceability to UTC.

4. *Fixed-Latency Data Transfer*: Data is streamed from each node in a synchronous and deterministic manner, reducing data collisions when multiple nodes are streaming data. The WRS uses a WR-streamer core in the gateware, which acts as a first-in, first-out (FIFO) buffer, and allows data to be sent to its destination with fixed and known delay. Fixed-latency data transfer also allows dispersed systems to use closed control loops, as feedback values are guaranteed to have constant latency. The fixed data latency may be a powerful feature for real-time coherent data fusion schemes in future multistatic radars.

5. *Radio-Frequency Transfer*: In a RF-over-WR application, an analog RF input is digitized by a specialised RF FMC. The digitised samples are shared across the network, and are used by a direct digital synthesis (DDS) in each WR node to re-synthesise the RF input with a fixed delay and known phase. Synthesised RF signals are therefore coherent, and avoid the difficulties in phase alignment that plagues analog RF distribution. The verified jitter for such a system is less than 10 ps for frequencies up

to 355 MHz. Improvements to the jitter and phase accuracy are possible using the newer low-jitter WRS.

Dispersed RF distribution is therefore likely the most exciting feature for multistatic radars. The LOs for each radar could be synthesised using this concept, ensuring phase accuracy and stability on the order of picoseconds, with negligible phase drift through slow-time due to the WR protocol.

2.3.4 Limits of White Rabbit Performance

Frequency Stability

The ultimate limits on WR performance, in terms of TF stability, is tied to the implementation of WR on existing FPGA hardware. A 2018 study by Rizzi et al [94] quantifies the short-term and close-in phase noise performance of the WRS, and characterises the sources of additive phase noise due to the DDMTD and Gigabit transceiver. Additional, non-limiting phase noise issues are primarily caused by poor hardware decisions in early versions of the WRS. In particular, the original WRS used a mixed-mode clock manager (MMCM) PLL to lock the switch to an external 10 MHz frequency reference in *grandmaster mode*, which adds significant phase noise to the otherwise low-phase noise reference. However, this issue can be retroactively solved on old switches by adding a low-jitter daughterboard, which uses an analog PLL instead. Switches fitted with the low-jitter daughterboard show an order of magnitude improvement in short-term frequency stability.

Temperature Stability

As part of the link budget, WR compensates for variable delays of the medium due to changing temperature. However, there are some residual and uncompensated phase shifts due to changes in the delay of the fixed hardware delays from the printed circuit board (PCB) traces and internal FPGA delays. These effects are not always minimal, specifically if the nodes have a large temperature differential, and will impact the synchronisation accuracy. Li et al [95] tested the effect of temperature changes using a WR node placed in climatic chamber. The group found a variation in the synchronisation of 300 ps over a temperature change of 50° Celsius. Luckily, the phase changes are linear with temperature, so a compensation scheme was implemented by the group, reducing the temperature induced phase change to 50 ps over the same temperature range.

Scalability

A study in 2016 was conducted by Diaz et al [96] to determine the effects of daisy-chaining several WR nodes together. A total hop of up to 15 WR light embedded node (WR-LEN)s was achieved, with the 15^{th} node being slightly over the guaranteed 1 ns accuracy (1.0731 ns) with a standard deviation in jitter of 174.5 ps. The team then tested up to 18 hops, where the 18^{th} hop had degraded in time accuracy to 2.6487 ns with 490 ps jitter.

Transfer Mediums

WR synchronisation does not necessarily require fibre as the transfer medium, and can be achieved using other mediums such as coax or microwave links, given that correct calibration for the coax SFPs is used. That being said, single-mode fibre provides the stipulated TFT performance, whereas other mediums will have deteriorated performance. Microwave links may be preferable in multistatic sensing if fixed infrastructure is unavailable, but there is LoS.

A 2015 study by Rico et al was likely the first investigation into using WR synchronisation based on sub-6 GHz microwave links [97]. However, the study only achieved synchronisation below 50 ns using SyncE and PTP, and was unable to fully utilise WR due to microseconds of jitter on the estimated path length delay. Recently (2020), a study by Gilligan et al, used millimetre-wave P2P links operating in the 71- 71 GHz carrier range, with a link separation of of 500 m [49]. The group used 2 WR-LENs for the WR nodes, which were connected to the millimetre-wave links via corresponding duplex SFPs, with custom calibration values of the hardware delays. Their results, which are directly compared to a fibre based WRN, show that millimetre-wave links are capable of sub-nanosecond time transfer, and only slightly deteriorate the frequency stability (ADEV of 7.1×10^{-12} vs 3.7×10^{-12} at 1 s for millimetre link and fibre, respectively). The group attribute the deterioration in the ADEV to the frequency stability of the free-running dielectric resonator oscillator (DRO)s used by the links for carrier synthesis.

2.4 Conclusion

Coherent multistatic radars require TFT for their synchronisation and syntonisation requirements. Broadly, there are three requirements for TF in a multistatic radar. Time synchronisation is required for accurate bistatic ranging, with the time error typically being on the order of a tenth of the compressed pulse width. The clocks must maintain this accuracy during the holdover interval, which depends on the rate of synchronisation of the clocks and the frequency stability of the oscillators. Frequency accuracy of the STALOs is required for bistatic Doppler accuracy. Coherent integration and Doppler processing require phase stability during the CPI. The phase drift during this interval is dominated by the frequency offset between the STALOs. Meeting these requirements requires TFT from one of the many available techniques.

Fixed infrastructure based TFT is typically more robust, with the simplest techniques using a simple one-way TFT on coaxial cables. Fibre based systems offer a premium on long range TFT, as fibre is resilient to EMI, has very low attenuation, and is significantly cheaper than coaxial cable. TWOTFT has been proven to provide the most stable form of TFT, allowing optical clocks hundred of kilometres apart to be accurately compared.

If fixed infrastructure is unavailable, which is typically the case in mobile systems, wireless techniques are more appropriate. There are numerous wireless methods, ranging from using atomic clocks that are periodically re-synchronised, to microwave and optical links, and utilising the direct breakthrough signal. In particular, TFT utilising GNSSDOs is highly appealing. GNSSDOs provide a low-cost TF reference, which has the short-term frequency stability of an OCXO and the long-term stability of UTC, they are accurately synchronised in time to an absolute standard (typically UTC) within several nanoseconds, and they provide an accurate global spatial reference. However, GNSS-denial is becoming more of a concern, and therefore other systems should be evaluated.

One intriguing option is WR, which provides sub-nanosecond synchronisation to any node within a WRN. WR is an Ethernet protocol based on PTP and L1 syntonisation, allowing it to be used as a standard deterministic data network and timing and synchronisation network. Several updates to the hardware and protocol, such as the inclusion of a low-jitter daughterboard for the WRS, have improved the performance of the time transfer, and reduced residual clock jitter. WR is being used in hundreds of projects, ranging from telescope arrays to financial institutions. However it has had little interest in the radar community, despite having highly attractive characteristics, like the commercial availability and open hardware, an included data network, phase synchronisation and frequency syntonisation, and numerous plug-and-play FMCs that offer extended capability. The next chapter will describe the design of a WR system for use in coherent multistatic radar.

Chapter 3

Design of a White Rabbit Disciplined Oscillator Network

In this chapter we discuss the design of a WRN for use in the NeXtRAD radar system. We first establish the requirements for the WRN, both based on NeXtRAD's interoperability requirements, and the synchronisation and timing requirements established in chapter 2. An overview of NeXtRAD is available in Appendix A.

WR was designed as a sub-nanosecond synchronisation network, indicating that the relative timing accuracy between nodes, and the relative long-term frequency stability, should improve on the performance of the current multi-channel GPSDO network. However, based on available literature, it was hypothesised that the commercially sourced fine delay FMCs were likely to inject higher levels of phase noise into the radar.

We therefore considered two designs, the first being a system using only commercially available WR equipment, specifically a fine delay FMC, and the second being a hybrid WR and GPSDO design, called a WRDO.

It was hypothesised that the short-term frequency stability of the WR fine delay design was inadequate for the phase noise requirements of NeXtRAD, and that the WRDO design would improve on the short-term stability. The short-term stability will therefore likely be similar to the GPSDOs, but will improve over the GPSDOs on longer-term averages. The relative phase between two WR nodes is also locked, thus the WRDO is expected to see reduced phase drift, in comparison to the GPSDOs.

Sandenbergh had previously found the optimum PLL time constant for the GPSDOs by comparing the GPS PPS stability to the free-running OCXOs on an ODEV plot. This method was not used for the WRDOs, as the PPS accuracy is below the resolution of the TICs. Instead, two time loop constants were empirically tested: 10 s and 100 s.

3.1 NeXtRAD Time and Frequency System

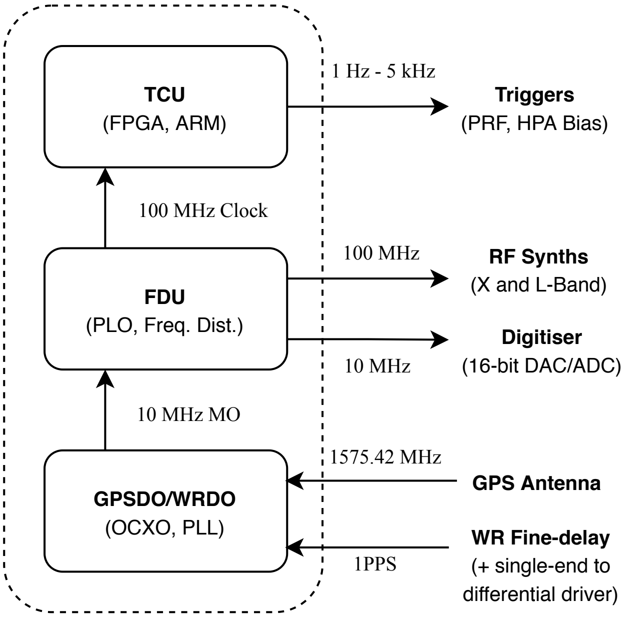

Figure 3.1: NeXtRAD node TF system.

The TF system is responsible for generating reference frequencies for each device, setting up the start time of an experiment, and triggering the PRF and control switches. The TF system, which is depicted in figure 3.1 comprises three sub-units, namely a TF reference (typically the UCT GPSDOs), which generates the MO frequency reference and synchronises the radar nodes; the FDU, which distributes frequency references and triggers for each other sub-system; and the TCU, which generates the PRF, biases the HPAs, and controls the L-band polarisation switch and receiver and exciter subsystem (REX) transmission frequency.

Each node is also equipped with a local node controller laptop, which is used for local control and coordination with the rest of the network. For the TF system, the node controllers are needed to loosely synchronise the node to a local NTP timeserver, and are required to run the GPSDO control GUI.

3.1.1 UCT GPSDOs

The UCT GPSDOs were designed by Sandenbergh, and are extensively covered in the author's doctoral thesis [1]. Essentially, each GPSDO is comprised of a Motorola M12+ Timing GPS receiver [98], phase detection circuitry, an SC-cut Oscilloquartz 8788 OCXO [99], and clock shaping circuitry. A timing-grade Motorola GPS antenna [100] with a mask angle of 20° is used to reduce multipath interference and limit geometric dilution of precision (GDOP) from satellites on the horizon.

Phase locking is achieved by comparing a digital 1 MHz clock, derived from the low-phase noise OCXO, to the PPS transition of the GPS. Figure 3.2 diagrams the PLL architecture. The 10 MHz output from the OCXO is first divided by 10, and then the time error between a rising edge of digital version of the 1 MHz and the PPS output of the receiver is measured, using an ACAM TDC-GP2 [101]. The time-to-digital converter (TDC) has a resolution of 65 ps and an RMS time precision of 50 ps.

The TDC outputs a digital 16-bit reading, which can be used within a closed loop GPSDO implementation of the Spartan 3 FPGA, or output over a serial bus to use an external SDF. Sawtooth correction and outlier removal can be enabled, which respectively reduce the GPS receiver's quantisation error and ignore outlying phase points, that would otherwise cause transient behaviour.

However, this algorithm was never implemented, thus phase transients still occur. The output of the filter is a 20-bit digital control word that tunes the DAC to steer the OCXO's frequency in the direction that minimises the phase error, effectively nulling the OCXO ageing.

The optimum PLL loop time constant for the GPSDOs is 1000 s, which was empirically found by comparing the MDEV of the free running OCXOs with the MDEV of the PPS output from the GPS receiver, and noting the point at which their sigma-tau plots intersect. The intersection defines the point at which the frequency stability of the OCXO is dominated by the frequency steering to the long term stability of the GPS PPS. The value of 1000 s is common for high quality OCXOs, but can be significantly longer with chip-scale atomic clocks which have improved holdover, but cost considerably more.

In practical applications, the software filter has been preferred, as it allows various PLL time constants to be tested before FPGA implementation. This also allows a quick-locking algorithm designed by Sandenbergh to be used. The algorithm vastly improves the acquisition time, by using an initially lower loop time constant (wider bandwidth) of 75 s, and gradually increasing the time constant once the PLL is locked to the optimum 1000 s [57]. This feature drastically assists in lock times when the nodes are relocated, and when transient signals cause the GPSDO to lose lock.

Frequency Syntonisation

Once each GPSDO has been left running long enough[1], and the 1000 s time constant has been reached, the OCXO frequency should be accurate to around 1×10^{-14} for one day averages. Measuring the absolute frequency, and therefore gaining frequency traceability to GPS, would require a careful string of calibrations. However, GPSDOs are self-calibrating, so the assumption that the frequency is syntonised to GPS with this accuracy is typically valid [67].

In NeXtRAD however, relative syntonisation between radar nodes is more im-

---

[1] Reaching constant ageing for the OCXO can take several weeks of continuous operation.

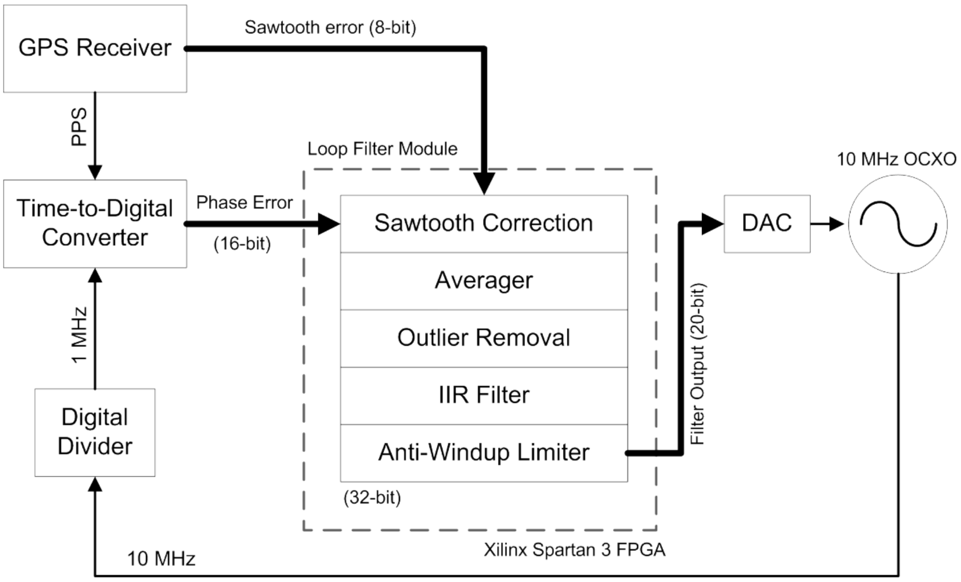

Figure 3.2: Block diagram of UCT GPSDOs (image from Sandenbergh [1]).

portant than absolute frequency accuracy. Figure 3.3 shows the typical MDEV between two GPSDO units. On short-term intervals the frequency stability is dominated by the OCXO stability. Sandenbergh measured a peak-peak frequency error of 1.3×10^{-10} for the closed-loop implementation of the 1000 s time-constant.

Time Synchronisation

During an experiment, the future start time, or *epoch* time, is sent to each node controller over the local area network (LAN) as a future UTC timestamp. The GPSDO GUI reads the epoch time, and sets the alarm time and alarm enable registers in the GPSDO.

The GPSDO alarm triggering system (depicted in figure 3.4), was designed by Sandenbergh to synchronise the start time of a radar experiment. In essence, the mechanism is based on a series of logic operations implemented in the FPGA fabric, which allows the GPSDO to fire an epoch pulse on the smoothed UTC transition.

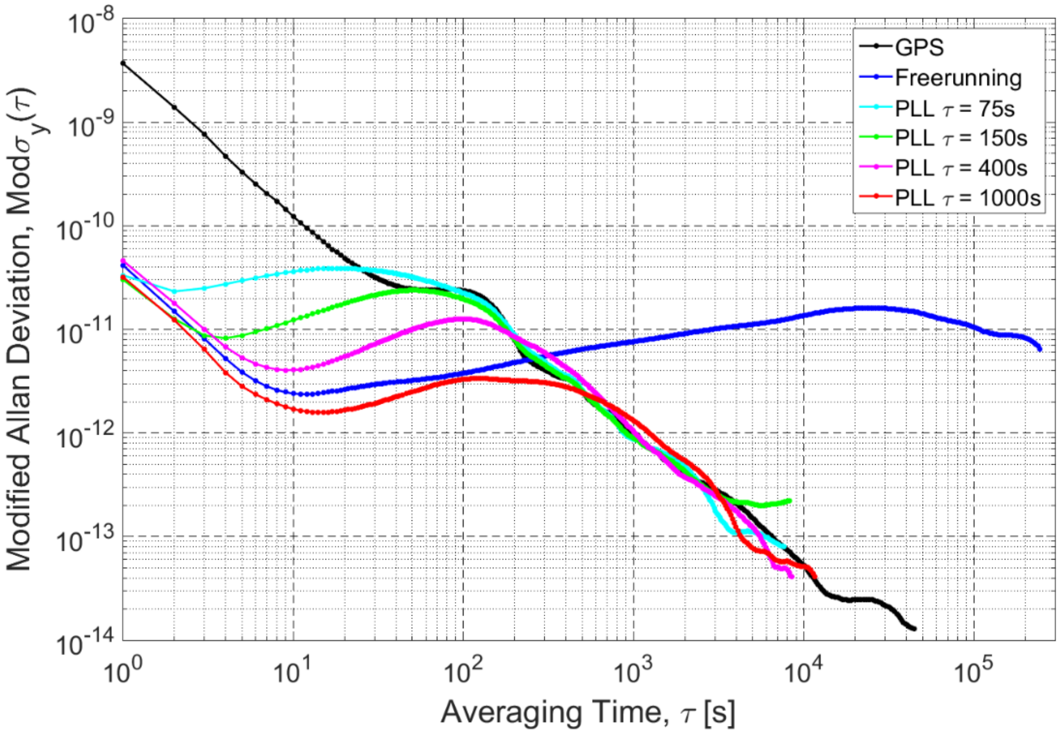

Figure 3.3: MDEV of UCT GPSDO with various time constants, GPS PPS, and free-running OCXO (image from Sandenbergh [1]).

The epoch pulse is not aligned with the GPS receiver's PPS transition, as this contains pulse jitter on the order of 20 ns, with sawtooth compensation enabled. Instead, the GPS PPS is offset by 500 ns, such that the leading edge lands roughly in the centre of two clock cycles of the PLL 1 MHz feedback clock.

The epoch pulse is therefore triggered by an AND gate on the following rising edge of the 1 MHz clock. The epoch pulse is therefore aligned with the positive zero phase transition of the MO output frequency. This is an important point, because it casts the radars time accuracy into a measure of the relative phase between the MOs. This assumption is not always valid, as commercial GNSSDOs often do not synchronise the output PPS to the OCXO.

Systematic UTC time offsets between two GPSDO units are typically seen, primarily caused by differences in cable lengths to the GPS antenna. These errors can be mostly cancelled in calibration by changing the M12's GPS receiver cable

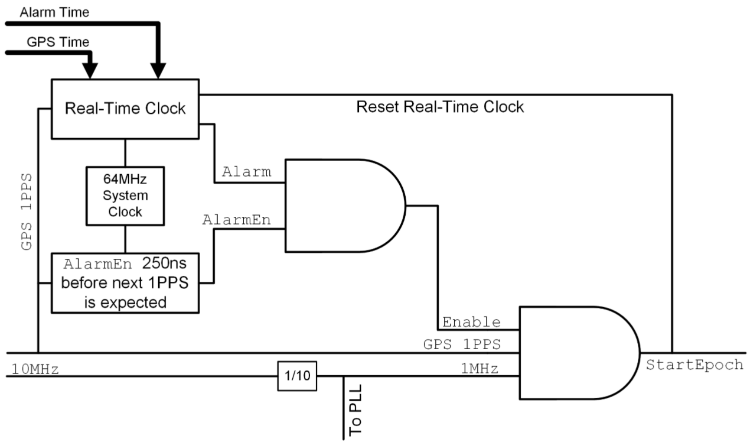

(a) GPSDO alarm mechanism

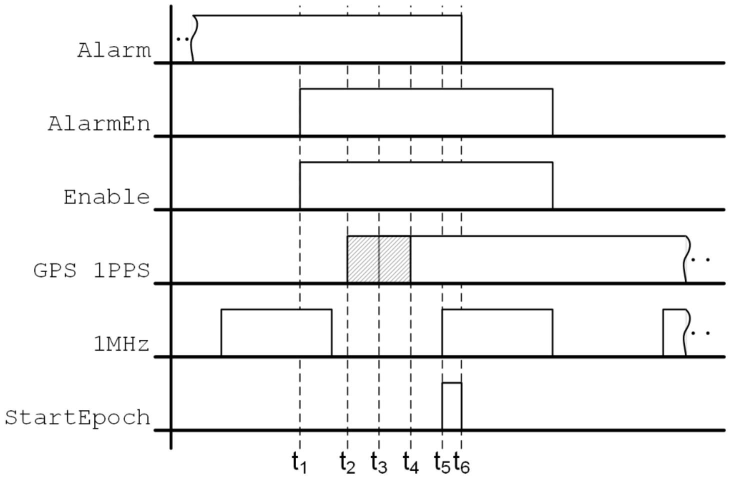

(b) GPSDO alarm timing diagram

Figure 3.4: GPSDO alarm and triggering system (images from [1]).

delay, with a resolution of 1 ns.

3.1.2 Frequency Distribution Unit (FDU)

Several of NeXtRAD's subsystems require external frequency references, which are either 10 MHz or 100 MHz, in sine and digital waveforms. The FDU is therefore responsible for the distribution of the frequency references, as well as distributing delayed pulse triggers. The frequency references are ultimately derived from a 10 MHz MO input, which is typically from the UCT GPSDOs. The 10 MHz input is split into 8 channels using a 1:8 RF splitter. Six of the channels are amplified to +7 dBm for the 10 MHz sine output section. One of the channels is used to derive the 10 MHz 3.3 V complementary metaloxidesemiconductor (CMOS) clock signal, which is fed into a 1:6 clock buffer for the 10 MHz digital output section. The remaining channel is amplified and input into a 100 MHz Wenzel phase-locked oscillator (PLO), which is also passed to a 1:8 splitter, before generating the 100 MHz sine (+7 dBm) and digital (3.3 V) sections, in the same manner as the 10 MHz.

3.1.3 Timing and Control Unit (TCU)

The TCU is responsible for creating the stable PRF, biasing the HPA during transmission, and switching between receive polarisations for the L-band carrier. The TCU is based on a Reconfigurable Hardware Interface for computatioN and radiO (RHINO) board [102], which is an FPGA based compute platform intended for use in RF applications. The RHINO contains an Advanced RISC Machines (ARM) processor to reprogram the FPGA, as well as read and change registers as needed. The TCU's clock is directly sourced from a digital 100 MHz CMOS output of the FDU.

In NeXtRAD's use-case, the pulse parameters for each experiment are uploaded to the TCU over the LAN prior to the start time. The gateware is then programmed to generate the correct timing signals such as the PRF, bias signal for the amps, polarisation switching, carrier frequency switching, and the total number of transmitted pulses.

Once programmed, the TCU enters a wait mode, which awaits an epoch trigger signal from the time reference, which is typically from the GPSDO alarm mechanism. Further calibration of the systematic time offsets is possible, by delaying the timing outputs by integer clock periods, giving a tuning resolution of 10 ns if needed.

3.2 NeXtRAD Time and Frequency Requirements

The TF requirements for NeXtRAD are determined by the requirements in chapter 2, as well as the requirements imposed by NeXtRAD's hardware, namely the phase noise requirements of the REX.

3.2.1 Time

From equation 2.5, NeXtRAD's 45 MHz pulse bandwidth requires a bistatic synchronisation accuracy of 2.22 ns RMS.

From equation 2.7, maintaining this level of synchronisation during holdover requires clock frequency stability of 1.11×10^{-9} for a 1 s update rate. However, as explained previously, this equation assumes linear phase drift, and does not take into account the mechanisms of the time synchronisation system, initial time offsets, and frequency drift characteristics of the oscillators. Using Sandenbergh's model (equation 2.9), the time requirement consists of a synchronisation error, and holdover component, which should be less than the 2.22 ns. The time drift is dominated by the frequency offset between the RTC oscillators, which in NeXtRAD's case are derived from the MO 10 MHz.

3.2.2 Frequency

Frequency Accuracy

Relative frequency accuracy is required for accurate bistatic velocity measurements. Although no current requirement exists for NeXtRAD's velocity accuracy, a value of 0.1 m/s is chosen, as the radar is used in maritime contexts where small boats in high sea states are of interest. Assuming a velocity accuracy $0.1\,\mathrm{ms}^{-1}$, from equation 2.1.3, NeXtRAD requires a maximum frequency error, $y(t)$, below 3.34×10^{-10}.

Phase Noise

Phase noise offset [Hz]	REX Requirement [dBc]	Measured FDU [dBc] (Wenzel specification [dBc])
100	-118	-129 (-138)
1k	-146	-151 (-159)
10k	-165	-160* (-177)
100k	-165	-160* (-179)

Table 3.1: Phase noise requirement for REX 100 MHz reference and measured FDU output. Wenzel specifications given in brackets, and measurement noise floor given in asterisks.

NeXtRAD's active and passive node STALOs ultimately set the requirements on transmitter and receiver phase noise, requiring as *low as possible* 100 MHz frequency references. The MO phase noise requirement was therefore set by the REX reference frequency requirements. To meet these requirements, the FDUs incorporate ultra-low phase noise Wenzel PLOs. The phase noise of the FDU channels was measured using a Rohde and Schwarz FSU26 spectrum analyser [103] , which can measure phase noise using cross-correlation of the PLL phase measurement technique. The phase noise measurements, along with measurements of several other frequency references is available in appendix B.

The FDUs use low noise op-amps to drive each channel to 7 dBm, adding negligible phase noise in the form of 1/f noise. The measured phase noise of the FDU is within 10 dBs of the PLO's specifications up to 1 kHz [104], and sufficiently below the REX requirements, see table 3.1. The noise floor at frequencies higher than 1 kHz is likely dominated by the FSUP26's phase noise floor, along with the automatically selected measurement bandwidth.

The Wenzel PLL loop bandwidth is 5 Hz and consequently, within the loop bandwidth, the PLOs frequency stability is dominated by source's quality.

Table 3.2: WR network Specifications

Device	**Specifications**	**Notes**
White Rabbit Switch	WRS [105] (S/W v5.1 and H/W v3.3)	Updated Firmware Set to Grandmaster Mode
Fibre Transceivers	Axcen bidirectional SFPs [106] 1310/1490 nm	Default calibration delays
Node	SPEC board [90] Fine-Delay FMC [92] PC OS Ubuntu 12.04 LTS	FD set to PPS on channel 1 10MHz channel 2
Fibre	Single-mode 2 m and 5 km lengths	
GPS Clock	Meinberg Lantime M400	TCXO Version [107]

Coherent Integration

Phase stability is needed for coherent integration and Doppler processing. In monostatic radar, phase stability during a CPI is typically guaranteed due to the cancellation of close-in phase noise and drift. The bistatic node's phase will drift as function of the relative carrier frequency error between the transmitter and receiver.

The maximum acceptable phase drift is up to phase quadrature during the CPI for an coherent integration loss of 3 dB. From equation 2.17 the required frequency error for NeXtRAD is 1.92×10^{-10} for 1 s integrations of the L-band carrier (1.3 GHz).

3.3 White Rabbit and GPSDO Setup

3.3.1 White Rabbit Network

The WRN in this thesis consists of a WRS, two SPECs with fine-delay (FD) FMCs, a GPSDO clock for the grandmaster reference, standard bidirectional fibre transceivers and two length of single-mode fibre cable (2 m and 5 km). A figure of the apparatus is shown in figure 3.5. Table 3.2 details the specifications

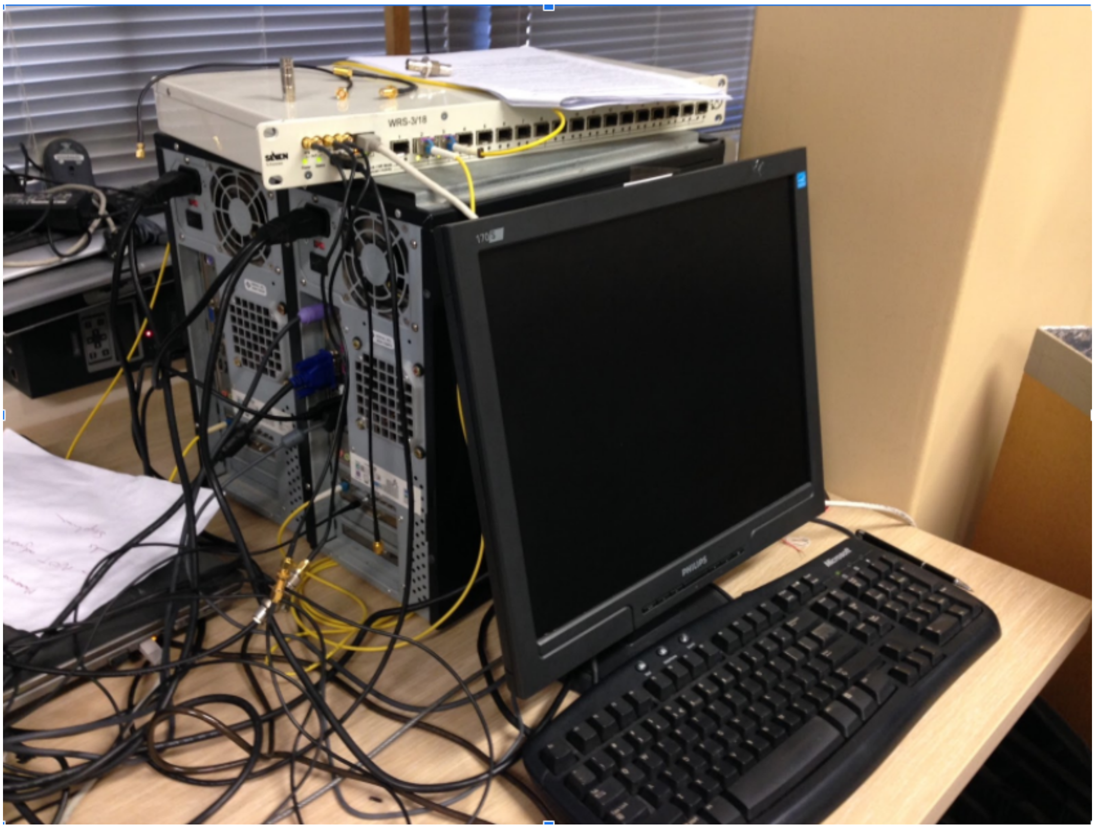

Figure 3.5: WR network setup. WR network setup for laboratory measurements. Network is composed of 2 computers with SPEC cards and fine delay FMCs, and a WRS in grandmaster mode (top).

of the equipment used in this thesis.

3.3.2 Updates and Calibration

The switch's firmware was updated to accommodate version 5.1 of the WRS software, and the network was calibrated as detailed in [108]. It was found that the default alpha values for the recommended SFPs performed well, and no further SFP calibration was required.

The switch was set to grandmaster mode, and a Meinberg M400 NTP time server was used as the TF reference for the network. Time of day information was not required in our application, however the WRS was configured to receive NTP

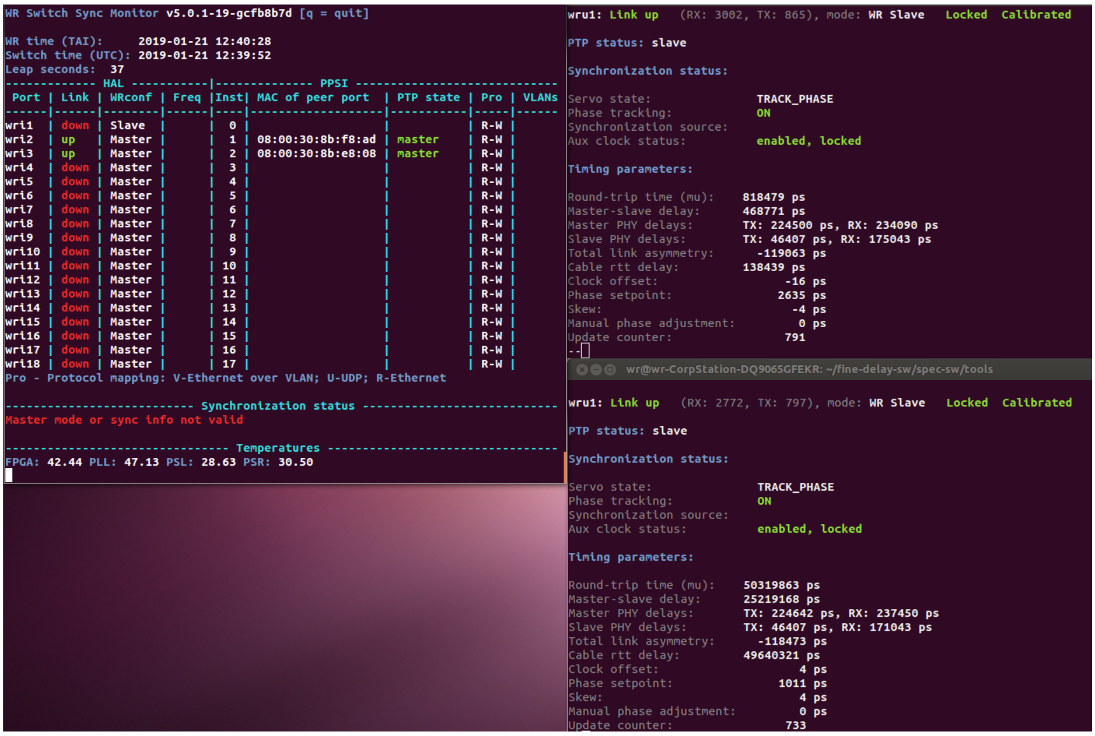

```
WR Switch Sync Monitor v5.0.1-19-gcfb8b7d [q = quit]

WR time (TAI):     2019-01-21 12:40:28
Switch time (UTC): 2019-01-21 12:39:52
Leap seconds:  37
------------- HAL -----------|------------- PPSI --------------------------------
 Port | Link | WRconf | Freq |Inst| MAC of peer port  | PTP state | Pro | VLANs
------|------|--------|------|----|-------------------|-----------|-----|------
wri1  | down | Slave  |      |  0 |                   |           | R-W |
wri2  | up   | Master |      |  1 | 08:00:30:8b:f8:ad | master    | R-W |
wri3  | up   | Master |      |  2 | 08:00:30:8b:e8:08 | master    | R-W |
wri4  | down | Master |      |  3 |                   |           | R-W |
wri5  | down | Master |      |  4 |                   |           | R-W |
wri6  | down | Master |      |  5 |                   |           | R-W |
wri7  | down | Master |      |  6 |                   |           | R-W |
wri8  | down | Master |      |  7 |                   |           | R-W |
wri9  | down | Master |      |  8 |                   |           | R-W |
wri10 | down | Master |      |  9 |                   |           | R-W |
wri11 | down | Master |      | 10 |                   |           | R-W |
wri12 | down | Master |      | 11 |                   |           | R-W |
wri13 | down | Master |      | 12 |                   |           | R-W |
wri14 | down | Master |      | 13 |                   |           | R-W |
wri15 | down | Master |      | 14 |                   |           | R-W |
wri16 | down | Master |      | 15 |                   |           | R-W |
wri17 | down | Master |      | 16 |                   |           | R-W |
wri18 | down | Master |      | 17 |                   |           | R-W |
Pro - Protocol mapping: V-Ethernet over VLAN; U-UDP; R-Ethernet

-------------------------- Synchronization status ----------------------------
Master mode or sync info not valid

------------------------------ Temperatures ----------------------------------
FPGA: 42.44 PLL: 47.13 PSL: 28.63 PSR: 30.50
```

```
wru1: Link up   (RX: 3002, TX: 865), mode: WR Slave   Locked  Calibrated

PTP status: slave

Synchronization status:

Servo state:               TRACK_PHASE
Phase tracking:            ON
Synchronization source:
Aux clock status:          enabled, locked

Timing parameters:

Round-trip time (mu):    818479 ps
Master-slave delay:      468771 ps
Master PHY delays:       TX: 224500 ps, RX: 234090 ps
Slave PHY delays:        TX: 46407 ps, RX: 175043 ps
Total link asymmetry:      -119063 ps
Cable rtt delay:         138439 ps
Clock offset:                  -16 ps
Phase setpoint:               2635 ps
Skew:                           -4 ps
Manual phase adjustment:         0 ps
Update counter:               791
--
```

```
wru1: Link up   (RX: 2772, TX: 797), mode: WR Slave   Locked  Calibrated

PTP status: slave

Synchronization status:

Servo state:               TRACK_PHASE
Phase tracking:            ON
Synchronization source:
Aux clock status:          enabled, locked

Timing parameters:

Round-trip time (mu):    50319863 ps
Master-slave delay:      25219168 ps
Master PHY delays:       TX: 224642 ps, RX: 237450 ps
Slave PHY delays:        TX: 46407 ps, RX: 171043 ps
Total link asymmetry:      -118473 ps
Cable rtt delay:         49640321 ps
Clock offset:                    4 ps
Phase setpoint:               1011 ps
Skew:                            4 ps
Manual phase adjustment:         0 ps
Update counter:               733
```

Figure 3.6: The switch (left) and SPEC (right) Monitoring GUIs showing the WRN is correctly synchronised using the WR protocol and that the delays are calibrated. The master-slave delay corresponds to the 2 m cable (above) and 5 km cable (below).

timestamps from the Meinberg timeserver if UTC traceability for the network is required at a later stage. A screenshot of the operational network is shown in figure 3.6.

3.4 A White Rabbit Disciplined Oscillator

In this section, the design of a WRDO for TF referencing in NeXtRAD is discussed. Each FDU already generates a low phase noise 100 MHz frequency reference, but requires a stable 10 MHz as a phase reference for the internal Wenzel PLO. The WR network therefore requires that each node can produce a stable 10 MHz signal, which is also synchronous with each-other node.

The first implementation that was evaluated, was to investigate if there was a native way of generating the 10 MHz PLO reference using only commercially available WR equipment, and in particular, if the fine delay FMC was adequate based on its phase noise spectrum and any additional requirements for integration.

The second design integrated WR directly into the current network by supplying the GPSDO units with an external PPS from the WR nodes. The second approach, which has been called a WRDO, was selected, as it has a few additional benefits, such as reduced close-in phase noise, GPS geo-spatial referencing, seamless integration, and clock redundancy if WR or GPSDO is needed.

3.4.1 White Rabbit Fine Delay Design

The standard WR GPIO FMC can produce a series of pulse-width controllable output pulses to create a transistor-transistor logic (TTL) square wave, however the highest frequency it readily produces is on the order of a few kilohertz. This means that for most practical radar applications, it cannot be used as a primary reference frequency, which generally require low phase noise 10 - 100 MHz references. The WR developers were consulted for a way of generating a 10 MHz reference, and it was suggested that the FD mezzanine board be used.

The primary function of the FD is to generate user programmed delays from input to output between 600 ns to 120 s with a resolution of 10 ps [92]. The ACAM TDC time tags the input of a pulse with a resolution of 28 ps and 55 ps precision. In a long-term performance test, the input trigger to delayed output pulse's jitter was found to be better than 100 ps, whilst the mean accuracy of the delayed pulse to its expected delay was below 1 ns [109].

However, for frequency requirements the FD can also be used create pulse trains with adjustable pulse widths, allowing a maximum 10 MHz TTL level signal to be generated. The outputs are driven by fast 2 v/ns AD8009 opamps and provide 50Ω source impedance. The inputs have a programmable 50Ω termination resistor for interfacing. Using the on-board TCXO, the FD has a rated time-base accuracy of ±2.5 ppm, but within a WR locked network it will exhibit long-term caesium-quality performance with respect to the Grandmaster clock [110].

In order to use the fine-delay with the FDU, the input signal must be a 4 dBm sinusoidal signal into a 50Ω system. The TTL level square wave from the WR FD cards was therefore converted into the appropriate form by attenuation of the signal level and filtering the square wave to attenuate the odd harmonics present in the signal, leaving only the fundamental 10 MHz sinusoid.

A few options were discussed for the filter, namely passive, active or crystal filters. Whilst crystal filters were ideal for notching the 10 MHz reference, they are liable to greater shifts in their cutoff frequency due to temperature fluctuations. An active Sallen-Key architecture, 5^{th} order low-pass Butterworth filter was constructed using the AD filter design wizard [111]. The cutoff frequency was 10 MHz, with the filter designed for 50Ω matching.

Fine Delay Phase Noise

The phase noise out of a free-running filtered fine delay card was measured, showing significantly higher short-term instability than the GPSDO OCXO (see Appendix B for measurements). The result is skewed by the lack of WR synchronisation, as the noise within the nominal WR PLL bandwidth (30 Hz) should be similar to that of the WRS. The switch was not available during measurements, but based on [94], the phase noise should be approximately −80 dBc within the loop bandwidth.

When set as reference to the FDU, within the 5 Hz Wenzel bandwidth, the higher short-term instability from the fine delays was postulated to deteriorate the phase noise of the FDU channels. Unfortunately, phase noise was not measured for the FDU when referenced by the WR fine delays. However the results are likely similar to the free-running FDU, albeit with the close-in phase noise of the fine delay reference up to 5 Hz. The integrated jitter will therefore be higher when referenced by the fine delay, even when considering the cards on synchronised to the WRS.

Triggering the radar nodes using the fine delay

If the fine delay was selected as the primary TF reference for NeXtRAD, alterations to the triggering mechanism would need to be made. Time triggered control is one of the main applications of the fine-delay. If this method was chosen, one could incorporate it into NeXtRAD's current trigger system by forwarding the future epoch time to each node. Each fine delay would output a trigger pulse to its local TCU, which would trigger the PRF, ADC sampling, and HPA biases.

3.4.2 White Rabbit Disciplined Oscillator

The WRDO was designed to simplify interfacing WR into the current NeXtRAD system, without completely redesigning the TF network. In this regard, the principle of the WRDO was to use the UCT GPSDOs as the reference to the FDU, but discipline the OCXO using the more stable PPS from a WRN. The GPSDOs have an external PPS input, which in this case is derived from the WR fine delay cards.

The pulse-to-pulse jitter of the PPS is in the picosecond regime, and does not suffer from the sawtooth error created by the GPS receiver. The more stable PPS therefore allows the GPSDO loop bandwidth to be relaxed, as shorter averages are required to meet the stability crossover point. Furthermore, the phase synchronous nature of WR was expected to assist in improving the coherency of the radar for long-term stability. However, the output of the fine delays is single-ended, therefore a single-ended to differential driver board was additionally required to interface the two devices.

SDF PLL Loop Constant for WR

The loop time constant of the SDF effectively determines the PLL loop bandwidth, and consequently how much of the frequency instability due to the PPS reference is seen at the output of the PLL. The optimum constant is the averaging time on a sigma-tau plot, where the time-domain stability from the reference source and OCXO intersect.

The optimal loop constant for the WR PPS was initially unknown, thus Sandenbergh's PPS stability method was attempted. However, the HP53131A TICs used in this thesis have a resolution of 0.5 ns, whilst the time error between two WR cards is below 1 ns. Determining the frequency stability using the PPS signals was therefore not feasible using the prescribed method.

Instead, an educated guess was made based on Lipinski et al's [42] results. The group measured the relative PPS stability between two fine delay units, and calculated an ODEV of approximately 10×10^{-10} for 1 s averages, with a white phase modulation slope (τ^{-1}). The estimated crossover point of the free-running OCXO and WR PPS is on the order of 50 s to 100 s. However, this estimate is predicated on the use of a standard WRS, and will likely be much lower when using a low-jitter switch which has an order of magnitude improvement in jitter. Although, this does assist the estimated value in this thesis, as no low-jitter board was available.

To accommodate the different loop constants, Sandenbergh's Matlab source code was altered, allowing various final loop constants to be programmed as needed. Two loop constants, 10 s and 100 s were selected. The experiments were run with both constants, to determine the best value given NeXtRAD's requirements. Hypothetically, the 10 s WRDO will lock to the WR source faster, and be steered to the stability of the WRN sooner. However, this will likely add more short-term instability, which is undesirable for typical radar CPIs. Conversely, the 100 s WRDO will have improved short-term instability due to more of the OCXO stability being present in the short-term region, but will take more time to reach the drift-free slope of the WR PPS.

3.4.3 GPSDOs

All three GPSDOs were placed in a laboratory with LNR400 cables running from the units to the roof of UCT, where the timing antennas were mounted. Cable delays had previously been measured, and the delays were removed from the GPS receivers' UTC offset with nanosecond level resolution.

A laptop running the Motorola WinOncore GPS program was used to set the receivers to average the position for 2 days, where-after they were set to position hold mode. The GPSDO GUI was then used in closed loop mode with sawtooth removal enabled, and the SDF set to the optimal 1000 s time constant. An average of 8 satellites were tracked throughout the measurement time.

3.5 Conclusion

In this chapter, NeXtRAD's TF system and requirements were discussed, and the design of WR network for use in multistatic radar was introduced.

NeXtRAD requires a synchronisation accuracy of 2.22 ns, and a fractional frequency accuracy below 1.11×10^{-9} for a 1 s synchronisation rate. Absolute frequency accuracy is required for Doppler accuracy, which should be below 3.34×10^{-10}. A more stringent requirement exists for the relative syntonisation between the transmitter and receiver for coherent processing, which requires a relative frequency accuracy below 1.92×10^{-10} for the L-band carrier of 1.3 GHz. The phase noise requirements are determined by the REX phase noise requirements, which are met by the FDU's Wenzel PLO.

Two WR designs were proposed for use in NeXtRAD, one using only commercially available WR equipment, and a hybrid design using the UCT GPSDOs. The WR fine delays were hypothesised to have higher levels of short-term instability due to their higher phase noise, when compared to the GPSDO OCXOs. However, both approaches will be used as the frequency reference for the FDUs, and thus the noise within the 5 Hz PLL bandwidth of the Wenzel PLOs will likely contribute the largest difference between the two designs.

Time triggering using the fine delays would require significant changes to NeXtRAD's control system, whereas the WRDO design already incorporates the current GPSDO trigger mechanism, making the transition easier, and allowing both WR and GPSDO to be used interchangeably.

Only recently (2020), has a mainstream low phase noise WR endpoint for frequency referencing become a reality. The high precision second (HPSEC) [112] is an ultra-low phase noise solution, with similar characterises to the WRDO and FDU. The HPSEC has redundant frequency references in the form of several 10 MHz, 100 MHz, and 1 GHz outputs, and is synchronised to WR by a built-in SPEC7. Any future work related to low phase noise frequency referencing could consider the HPSEC as an alternative, coupled with a low-jitter variant of the WRS.

Chapter 4

Time-Domain Measurements of Coherent Frequency References

In this chapter, the TF performance of the WRDO network is analysed under laboratory conditions. In addition, the current network of multi-channel GPSDOs are measured against the WRN, allowing a direct comparison of the performance of each TFT technique.

A high resolution, femtosecond-level measurement of the relative phase error between two TF references was conducted using the UCT DMTD. The phase error between two 100 MHz sinusoidal frequency references output by the FDUs are recorded for each for each TFT system. The DMTD reference frequency offset was set to 1 Hz. Both the phase and fractional frequency results are plotted for evaluation of the relative phase and frequency errors. The phase data are further processed to obtain metrics of TF stability, specifically the ODEV, MDEV, TDEV, and MTIE.

4.1 DMTD Test Bench

The relative TF stability between two TF references is measured using the DMTD technique, with a nominal sampling interval of 1 Hz. The UCT DMTD system is employed to measure the phase error series between two 100 MHz references, which are derived from two FDU units, locked under GPSDOs or WR. This section discusses the DMTD technique, measurement setup, and the data processing.

4.1.1 Dual Mixer Time Difference Technique

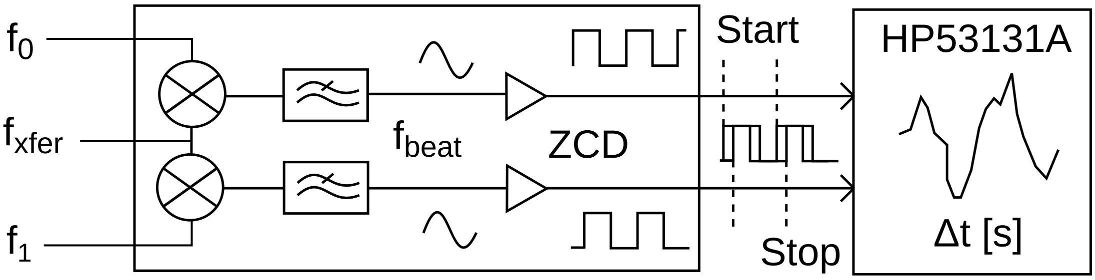

Figure 4.1: MO DMTD measurement diagram.

This thesis uses the DMTD technique to measure the phase and frequency stability between two FDUs. The DMTD technique, depicted in figure 4.1, is required to increase the resolution of phase measurements, which is typically limited to hundreds of picoseconds using simple time-interval methods. The DMTD device used in this thesis was designed and built by Sandenbergh (see [1], Appendix C), based on a design by William Riley [113].

The DMTD technique is a well established method for measuring the time-domain phase stability between two precision oscillators with femtosecond level resolution [87, 113, 114]. The DMTD technique uses heterodyning to reduce the frequency of two nominally similar oscillators, f_1 and f_2, to a much lower frequency, where their phase difference can be precisely measured.

In a DMTD device, a third slightly-offset reference oscillator with frequency f_{xfer}

(1-100 Hz higher than the oscillators under test) called the transfer oscillator, is mixed against two oscillators under test. The mixer products from each channel are low-pass filtered to remove the double-product, leaving just the low-frequency beat-note, f_{beat}, with a frequency equal to the offset of the transfer oscillator. The beat-notes are then input into zero-crossing detector (ZCD)s to increase the edge transition rate, which helps avoid aperture jitter in the next stage.

The low-pass filtered, ZCD, beat-notes are then input into a high-resolution TIC. One of the inputs is used to start the counter, whilst the other is used to stop. Noise from the reference oscillator is common to both mixers, and thus mostly cancels in the time difference[1]. Phase errors between the original oscillators are amplified by the heterodyne factor,

$$\mathrm{HF} = \frac{\mathrm{f_{nominal}}}{\mathrm{f_{beat}}} \qquad [] \quad (4.1)$$

The overall resolution of the DMTD system is then the ratio of the TIC resolution to the heterodyne factor,

$$\mathrm{DMTD_{res}} = \frac{\mathrm{TIC_{res}}}{\mathrm{HF}} \qquad [\mathrm{s}] \quad (4.2)$$

The measured time difference between the beat notes is linearly related to the phase error between the two oscillators, scaled by the heterodyne factor. Sampling the phase series (in units of seconds) generates a phase history of their relative phase error, which can then be used for calculating the respective time domain measures of frequency stability. In particular to this thesis, the primary metrics are the phase series, the ODEV, MDEV, TDEV, and MTIE.

Recently, advances in SDR have allowed high resolution measurements of oscillators to be made using low-cost SDRs [115]. The performance of SDR based oscillator metrology has been shown to improve on purely analog methods, as many sources of error due to the analog components can be mitigated. SDR methods should be considered in future projects, as they also require consider-

---

[1] Track length differences of the reference limit the higher frequency noise cancellation

ably less equipment, and will cost a fraction of the price.

4.1.2 Test Setup

This section outlines the test-bed, graphically depicted in figure 4.2, where each pair of TF references will be analysed. An image of the testbed is shown in figure 4.3.

Every TF reference generates a 10 MHz reference frequency which is used as the reference for the Wenzel PLOs. In total, there are 4 TF references that were tested, namely, the filtered fine delay FMCs, the WRDOs with 10 s and 100 s time constants, and the UCT GPSDOs with a time constant of 1000 s.

4.1.3 Choice of sampling time and recording length

The UCT DMTD device was designed to accept a maximum transfer oscillator offset of 100 Hz, and consequently, a maximum sample rate of 100 Hz. A DMTD sample rate of 1 Hz was chosen for the laboratory measurements, as the higher frequency phase noise components above 1 Hz are already described by the phase noise measurements in B. Higher data rates also incur more storage and processing, thus 1 s is a practical value for long-term measurements, specifically because we are interested in second long CPIs, and how the relative phase changes over time.

The maximum recording length was chosen to include the maximum tau in the frequency stability measurements, which is 1000 s in the case of the UCT GPSDOs. Recordings of several hours were made for each reference, but the total recordings were truncated to approximately 8 hours, which is long enough to capture the long-term steering effects.

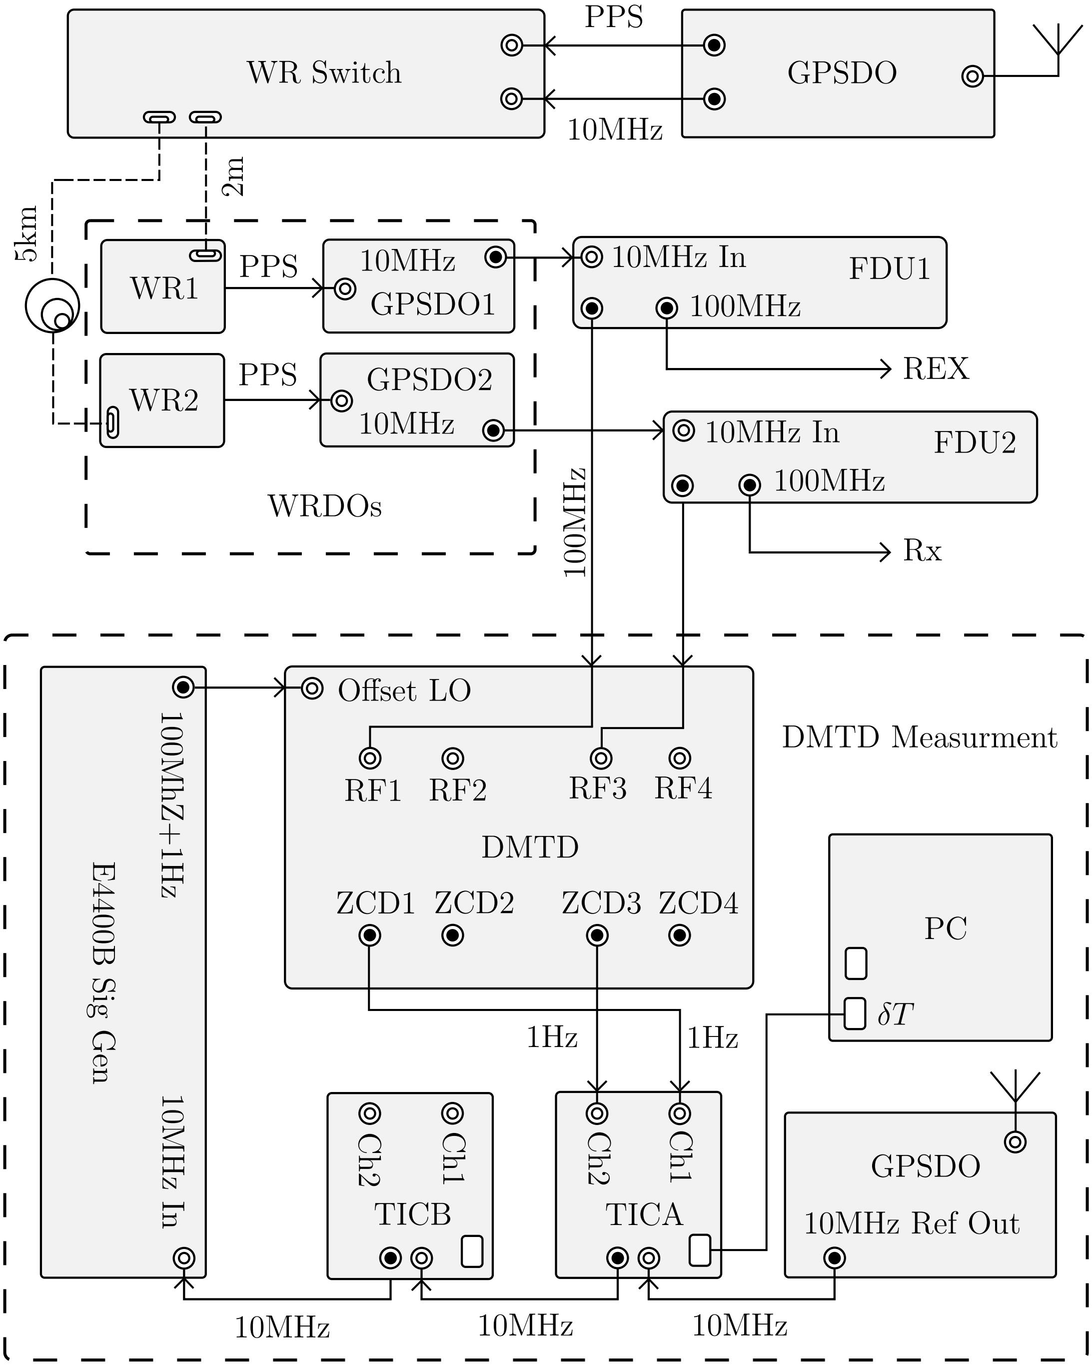

Figure 4.2: DMTD measurement testbed setup.

Figure 4.3: WR and GPSDO test bench. Left (from bottom up): E4400B synthesiser, HP53131A TICs, UCT DMTD, fine delay filters. Right: UCT GPSDOs

4.1.4 Cabling

Cable length differences to the input of the DMTD cause additional phase errors that are not cancelled in the differencing operation. Long cables also increase the presence of thermal noise, and reduce the SNR of the frequency source.

Cables were therefore kept as short as possible, and low-loss 50Ω ULC-1.5FT-SMSM+ Minicircuit SMA cables [116] were used for the most vital sections of the DMTD, which were predominantly the sections out of the device under test (DUT)s, and the sections connecting the ZCD output to the TICs. Short lengths of RG-58 cable were used to daisy-chain the Meinberg frequency reference to the synthesizer and TICs for reference frequency accuracy and coherence across the measurement equipment.

4.2 Stability Analysis

The time-domain frequency stability analysis for each experiment closely follows the analysis procedure in chapter 10 of William Riley's handbook of frequency stability [117]. The stability metrics are used to visualise the noise processes and long-term stability of the respective TFT network.

4.2.1 Data Recording

The TIC prints the phase error each second over an RS232 interface to a recording laptop. The phase series between for each measurement is recorded at a rate of 1 sample/s. The data are saved to a binary file, before being processed in Python to yield a text file with the timestamp and phase reading.

4.2.2 Preprocessing

Preprocessing the data is a vital component of any stability analysis. In cases where the phase changes by integer cycles, the phase becomes bound to $[0, 2\pi]$, in units of time. The phase data are therefore scaled by 2π and unwrapped to avoid phase wraps, before being scaled down by the heterodyne factor. The heterodyne factor for the 100 MHz frequencies and 1 Hz offset, is 1×10^8. The unwrapped TIC time error readings are therefore scaled down by 1×10^8

Outlier removal

Frequency outliers are removed prior to computing each statistic, which would otherwise bias the computed deviations. Phase jumps in the data are rare, but do happen for currently unknown reasons. Experimentation with the source of the outliers, shows that they are not due to the DMTD measurement process, but are likely related with the GPSDO electronics. A quick judgement may be that they are caused by GPS interference, but they also appear in the WRDOs, and are therefore probably related to the phase detection circuitry.

A mean absolute deviation (MAD) bound of 5 is used on the frequency data to flag outlying values. The corresponding frequency points are deleted, and the deleted values are interpolated between the nearest frequency samples on either side of the deletion. The frequency data are then reconverted to phase data. In more extreme cases, where many phase jumps are apparent, only sections of the data without the jumps are included.

4.2.3 Data Interpretation

Visual inspection of phase and frequency errors

The phase data are then used to visually analyse the phase, $x(t)$, and average fractional frequency (frequency error), $y(t)$. The frequency error plotted represents the fractional frequency using the minimum sample time of 1 s. Additionally, a 1-minute frequency average is plotted with the frequency plots, to indicate non-white excursions in the frequency data.

Sigma-Tau plots

The phase data are then used directly to compute the ODEV, MDEV, TDEV, and MTIE. The metrics are computed using Python and AllanTools [50]. AllanTools is a Python library built for TF metrology, and has been verified against the popular Stable32 written by William Riley [118]. Sigma-Tau plots for each metric are produced using the matplotlib [119] Python package.

4.3 DMTD Calibration

The DMTD is calibrated to measure the noise floor of the measurement equipment. The noise floor should be at least an order of magnitude lower than the expected frequency stability of the TF references.

4.3.1 DMTD Noise Floor

The HP53131A TICs were calibrated as per their user manuals, by using the quick-calibration technique. This involved injecting a clean 10 MHz CMOS source derived from the FDU into channel 1, and initiating the quick calibration routine, which routes the source to channel 2 and measures the static time offset between the two channels. The calibrated TIC jitter appeared to be white phase modulated (PM) noise, with a peak-peak error of 500 ps, bounded by the TIC resolution.

The noise floor of the DMTD was recorded by injecting a split 100 MHz signal from 1 GPSDO unit into 2 channels of the DMTD, where the lowest noise combination of DMTD channels was found to be channel 1 and 3. The transfer clock was synthesised using a Agilent E4400B [120], with an offset of 1 Hz. The device was left to warm-up for several hours, until the phase began to settle to a fixed offset.

Figure 4.4 shows the time error and fractional frequency series. Phase drift over the 8 hour period is minimal, with an approximate a peak-to-peak error of the mean of 0.6 ps. The long-term drift is likely due to a gradual temperature change.

The MDEV plot, which is used to analyse the frequency stability of the noise floor of the device, is given in figure 4.5. White PM appears up to a time average of about 40 s, thereafter flicker PM noise is dominant up to the flicker floor. The flicker frequency modulated (FM) floor appears near 4×10^3 s at a value of 1×10^{-17} , but a longer recording is required to find its true value.

Temperature drift is a slow process that affects the stability estimation on the order of several hours, and its effects are observed in this data. However, owing to the co-located nature of the devices under test, the air-conditioned laboratory temperature, and the equipment being subjected to the same environmental conditions, the ambient temperature and humidity were not recorded.

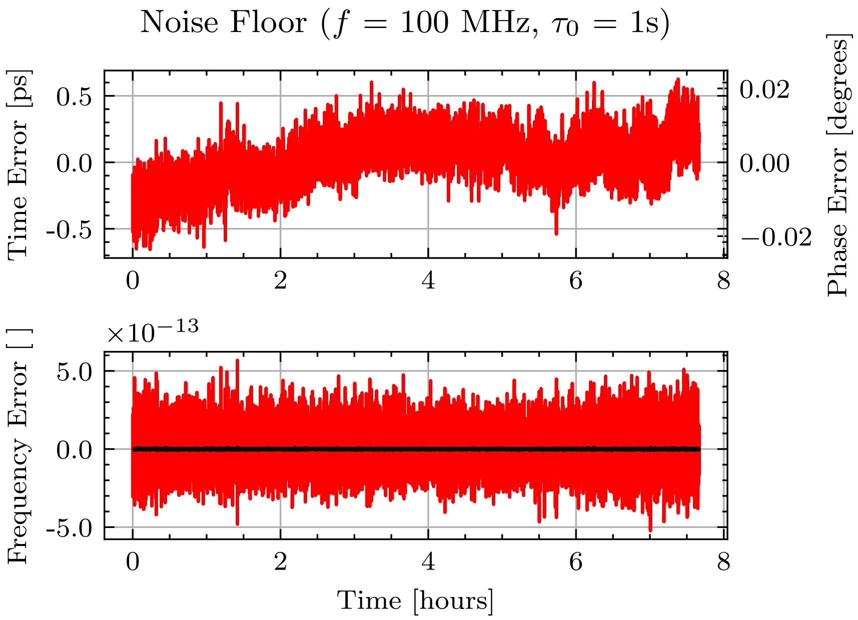

Figure 4.4: Phase of the DMTD noise floor

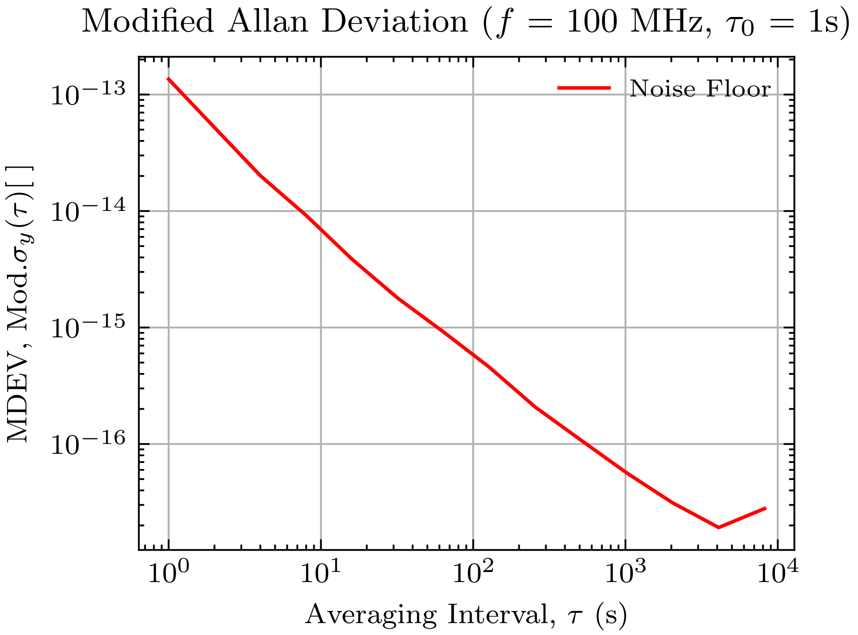

Figure 4.5: MDEV of the DMTD noise floor

4.4 Results

The results of each experiment are plotted as a phase error series using the scaled TIC readings. Initial time offsets are removed using the mean of the phase error. The average fractional frequencies are plotted alongside the phase recordings. The phase results are further processed for the selected time-domain measures of TF stability. A table of the stability metric results is given in table 4.3.

4.4.1 Phase and Fractional Frequency

The measured phase error and fractional frequency error (at $\tau = \tau_0 = 1s$) for each experiment are plotted in figures 4.6 to 4.10. The mean error of each phase recording is removed prior to processing, as to zero-centre the phases. No obvious phase drift is apparent in the data, thus the phase data are presented without frequency offset removal. A one minute moving frequency average is overlaid on each frequency plot. The estimated frequency drift was computed for an 8 hour average using a linear least-squares regression of the fractional frequencies, and the results are tabulated in table 4.2. The combined histograms of the time (phase) error and fractional frequency are plotted in figures 4.11 and 4.12 respectively.

The fine delays (figure 4.6) have an initial phase drift of approximately 401 ps over 2 hours (long-term frequency error of 5.5×10^{-14}), likely due to uncompensated temperature to phase conversion, as seen by Li et al [95]. Thereafter the frequency offset is stabilised, leaving predominantly white and pink PM. The long-term phase deviation is the lowest of all the references, bound to approximately ±5°, which would lead to a peak-peak phase error of 130° at L-band. However, the short-term frequency stability is higher than the OCXO references, as hypothesised, with a maximum average frequency error of 1.2×10^{-10}.

Both WRDOs normally have their phase bounded within ±10°, however the 100 s time constant WRDO occasionally has a large phase excursion, marginally greater than 1 ns. The transients last on the order of the damping time (3 −

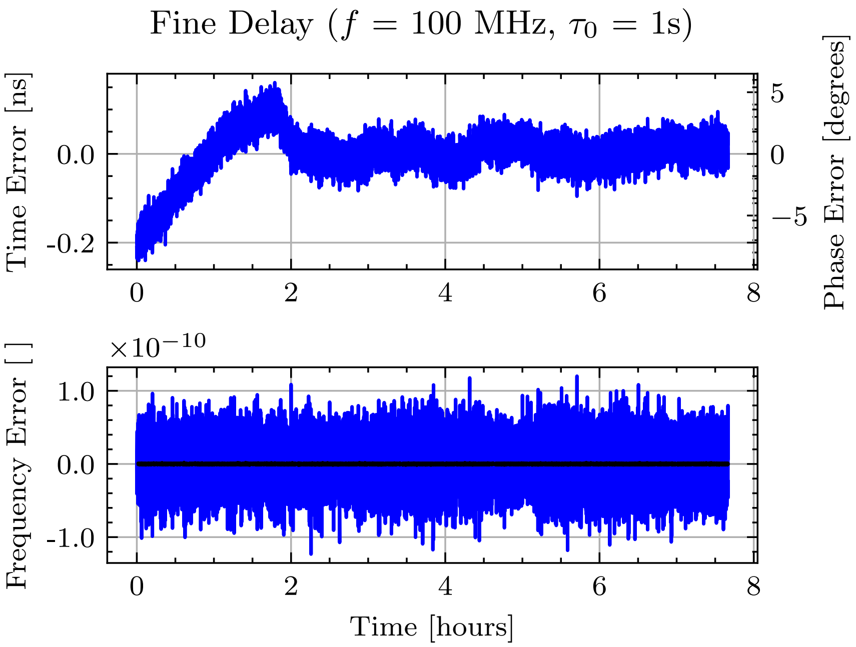

Figure 4.6: Wenzel PLOs referenced by WR fine delays.

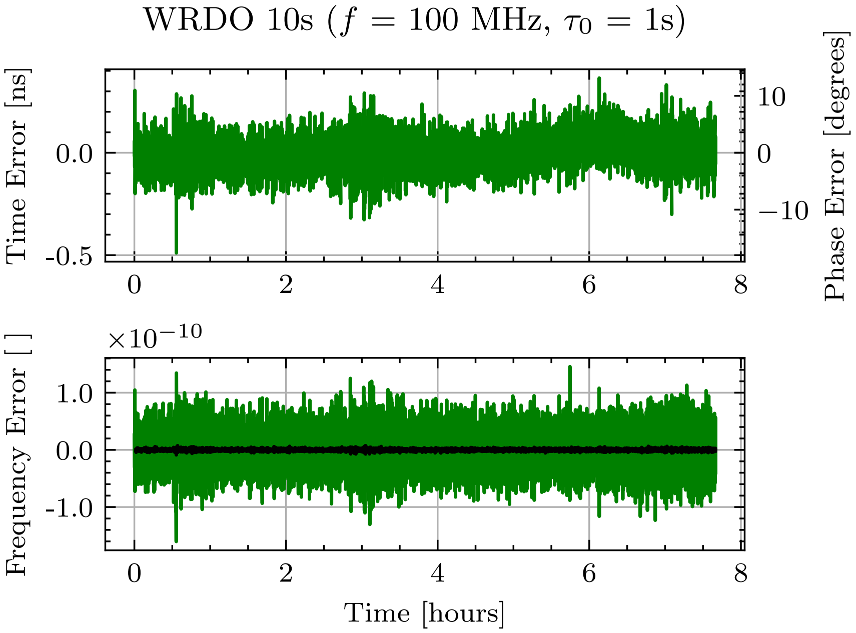

Figure 4.7: Wenzel PLOs referenced by WRDOs with with 10 s loop constant.

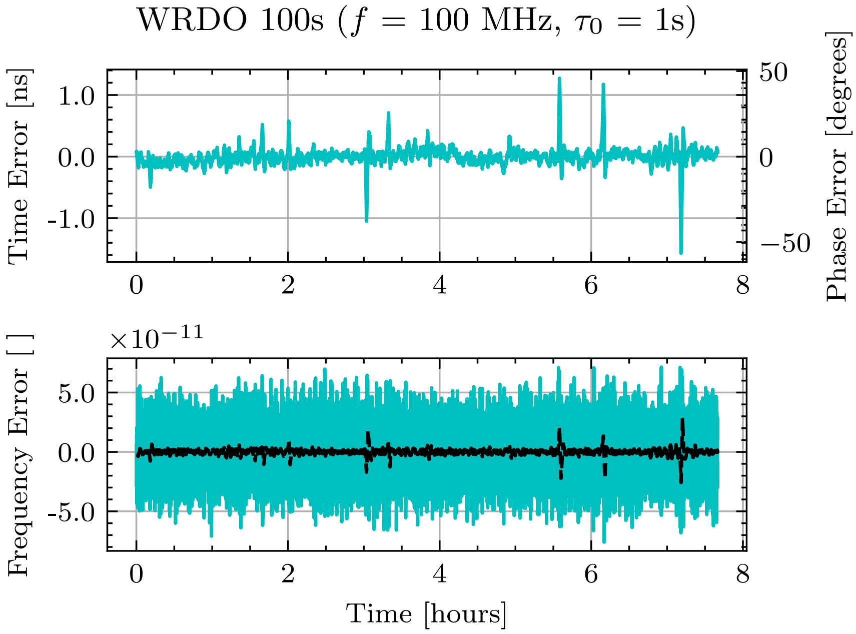

Figure 4.8: Wenzel PLOs referenced by WRDOs with with 100 s loop constant.

4τ =300-400 s). Figure 4.9 shows two 2000 s (33 minute) sections of the WRDOs with 100 s time constants, with and without transient behaviour. An overlay of a 10 s moving average for the frequency errors is plotted to better illustrate the average frequency transience. Over the 8 hour run, a total of 7 transients were recorded. A phase transient of 1.5 ns would lead to an error of 702° change in RF phase at 1.3 GHz. The origin of the transient behaviour isn't entirely understood, but is likely due to transient frequency errors from cable movement, noise on the single-ended to differential converter, or from the GPSDO phase detection. The highest fractional frequency error for 1 s averages is 6.68×10^{-11} , which is approximately 2.5 times lower than the frequency requirement.

The 10 s WRDOs have improved long-term phase synchronisation, but are afflicted by higher levels of short-term noise. There is no obvious evidence that the transients occur in the 10 s configuration, however the higher frequency noise is likely masking their appearance. For use in a multistatic radar, the more stable long-term time and phase synchronisation is tempting, as it leads the way to full coherence. However, typical NeXtRAD CPIs are on the order of a second,

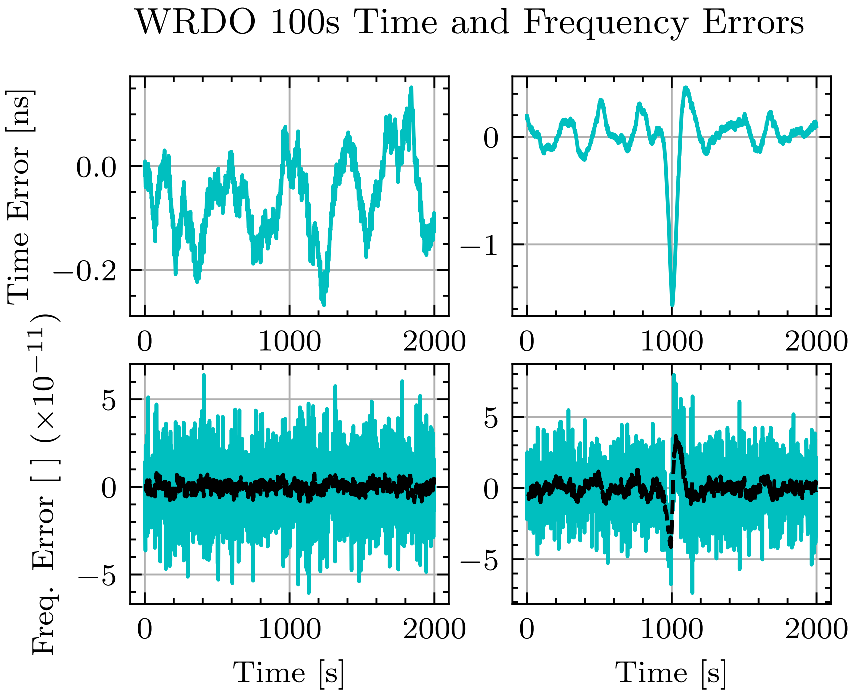

Figure 4.9: Left: Normal WRDO operation. Right: Large transient after the 7 hour timestamp.

and therefore the lower short-term noise from the 100 s is a considerably better option, even when considering the transient nature.

The GPSDO phase and fractional frequency data are plotted in figure 4.10. The GPSDOs have the best short-term noise performance, due to the longer time-constant, and consequently more time using the OCXO output. However, the phase drift is significantly higher due to the stability of the GPS PPS. The frequency transfer to the GPSDOs is also mostly independent, thus they drift independently. This is in sharp contrast with the WRDOs which are locked to the same frequency source, and phase synchronised. Over the 8 hours, the peak-peak time error between the GPSDOs reached 9.64 ns, which is several nanoseconds higher than the values measured by Sandenbergh for the same recording length. This is likely due to a shorter position average of 2 days, prior to running the experiments.

Figure 4.11 shows the combined histograms of phase. There is a large dispersion

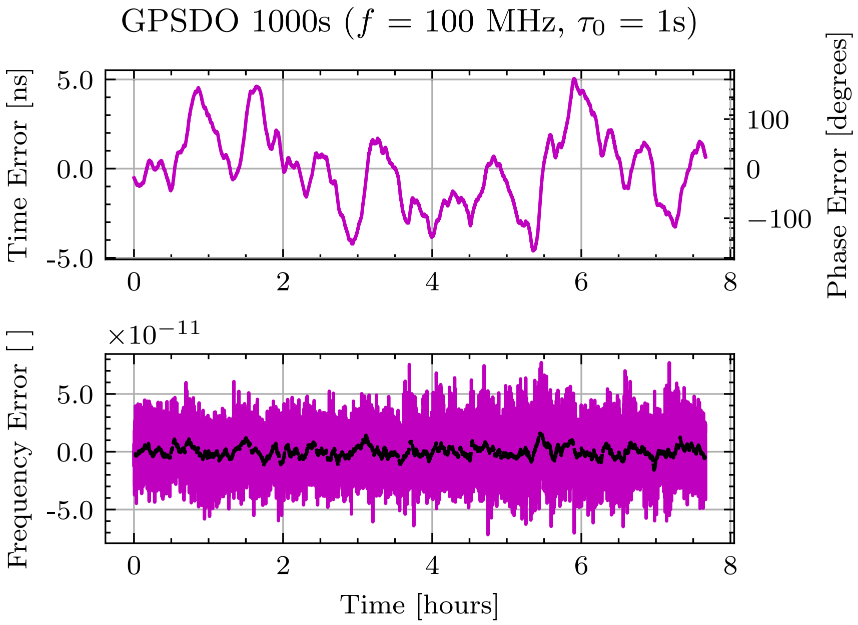

Figure 4.10: Wenzel PLOs referenced by GPSDOs with 1000 s loop constant.

of the GPSDO phase data, thus the GPSDOs are plotted on a separate scale for clarity. Tables 4.1 and 4.2 report the standard deviation, assuming normally distributed data for the phase and frequency. However, it must be stressed that these are far from normal distributions, due to coloured (long-term correlation) noise components. All of the distributions fail standard tests of normality, such as the DAgostino and Pearson test. The maximum frequency errors of all of the references are sufficiently below the coherent requirement of 1.94×10^{-10} for 1.3 GHz. However, this is assuming that there is no additional frequency error caused by mismatches in RF synthesizers.

Table 4.1: Peak-peak and standard deviation of phase and frequency errors.

Reference	Pk-Pk $x(t)$ [s]	S. Dev. $x(t)$ [s]	Pk-Pk $y(t)$	S. Dev. $y(t)$
Fine Delay	4.01e-10	5.3e-11	2.43e-10	2.86e-11
WRDO 10s	8.54e-10	7.62e-11	3.06e-10	2.96e-11
WRDO 100s	2.84e-09	1.43e-10	1.6e-10	1.85e-11
GPSDO 1000s	9.64e-09	2.16e-09	1.49e-10	1.67e-11

Note: The standard deviations assume a normal distribution, which is certainly not true in oscillators.

Table 4.2: Fractional frequency results over 8 hours.

References	Max. $\lvert y\rvert$	Std. Dev.	Est. y_0	Est. Drift (D)
Fine Delay	1.2e-10	2.86e-11	2.89e-14	-1.48e-18
WRDO 10s	1.46e-10	2.96e-11	3.81e-15	-3.37e-19
WRDO 100s	6.68e-11	1.85e-11	-1.69e-14	1.15e-18
GPSDO 1000s	7.71e-11	1.67e-11	2.78e-14	1.09e-18

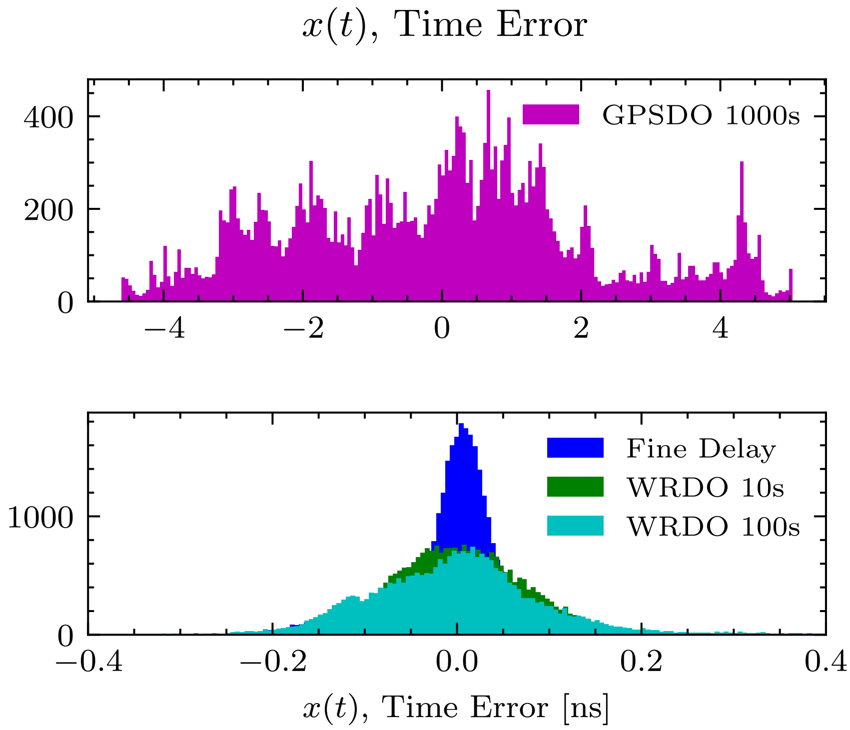

Figure 4.11: Combined time error histograms. GPSDOs separate for clarity. GPSDO bin size = 50 ps. WR bin sizes = 5 ps.

4.4.2 Time and Frequency Stability

The TF stability statistics are given to illustrate the stability of each respective TFT method. One should interpret these statistics as errors in the prediction of the frequency (ODEV and MDEV) or time (TDEV) at the respective averaging interval. For example, an MDEV of 1×10^{-11} for a 1 s average, would imply that there is uncertainty in a 100 MHz frequency source of 1 mHz. These statistics are used in contrast to standard variance, as they converge for all types of coloured

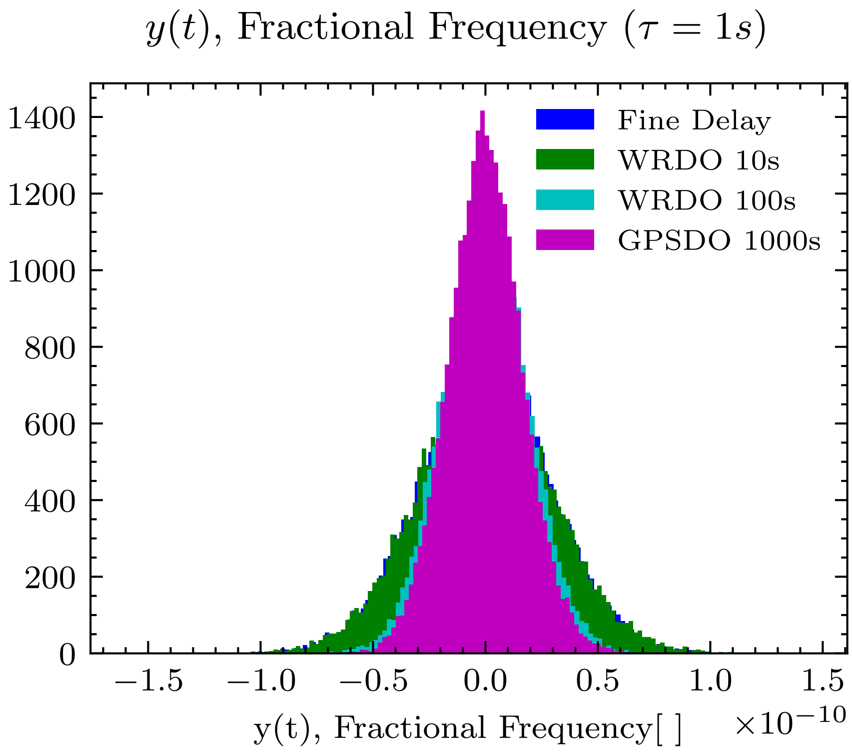

Figure 4.12: Combined fractional frequency histograms. Bin size = 2×10^{-12}

noise.

Numeric values for the TF stability results are available in table 5.1. The ODEV and MDEV of each experiment are combined, and shown in figures 4.13 and 4.14 respectively. The ODEV is depicted as customarily done in TF metrology, however, discussions of the frequency stability will focus on the MDEV. The MDEV is used as it is able to distinguish between white from pink PM at the shorter averages. Confidence intervals for all the metrics have not been depicted in the plots due to their negligible magnitudes when using overlapping statistics.

The fine delay noise has the highest short-term instability (1 s) of 3.5×10^{-11}, and averages down with predominantly white PM. Some additional instability, likely due to the temperature effects mentioned previously, occur around 1000 s averages. The WRDOs and GPSDO have similar instability for 1 s averages due to the OCXO, with values of 2.3×10^{-11} and 2.2×10^{-11} for the 10 and 100 s

Table 4.3: Time domain stability results

Reference	Metric	1.0e0	1.6e1	1.4e2	1.0e3	8.2e3
Fine Delay	ODEV	3.5e-11	2.1e-12	2.6e-13	4e-14	7.3e-15
	MDEV	3.5e-11	4.7e-13	2.6e-14	1.7e-14	7.4e-16
	TDEV	2e-11	4.4e-12	1.9e-12	1e-11	3.5e-12
	MTIE	1.2e-10	1.4e-10	1.5e-10	2.1e-10	4e-10
WRDO 10s	ODEV	2.3e-11	8.1e-12	9.3e-13	1.3e-13	1.6e-14
	MDEV	2.3e-11	4.3e-12	1.1e-13	1.3e-14	3.1e-15
	TDEV	1.4e-11	4e-11	8.4e-12	7.5e-12	1.5e-11
	MTIE	1.6e-10	7.8e-10	7.8e-10	7.8e-10	7.9e-10
WRDO 100s	ODEV	2.2e-11	2.7e-12	2.1e-12	2.5e-13	2.7e-14
	MDEV	2.2e-11	2.1e-12	1.4e-12	5.8e-14	4.3e-15
	TDEV	1.3e-11	1.9e-11	1e-10	3.4e-11	2e-11
	MTIE	8.3e-11	6.3e-10	2e-09	2e-09	2.8e-09
GPSDO 1000s	ODEV	1.9e-11	2e-12	2.6e-12	2.9e-12	4.8e-13
	MDEV	1.9e-11	1.3e-12	2.3e-12	2.3e-12	3.8e-13
	TDEV	1.1e-11	1.2e-11	1.7e-10	1.3e-09	1.8e-09
	MTIE	7.7e-11	3.3e-10	2e-09	6.5e-09	9.6e-09
Noise Floor	ODEV	1.4e-13	1e-14	1.4e-15	1.9e-16	3.4e-17
	MDEV	1.4e-13	3.9e-15	4.5e-16	5.6e-17	2.8e-17
	TDEV	7.8e-14	3.6e-14	3.4e-14	3.3e-14	1.3e-13
	MTIE	5.7e-13	8.2e-13	8.8e-13	1.1e-12	1.2e-12

WRDOs respectively, and 1.9×10^{-11} for the GPSDOs.

However, the 10 s WRDOs immediately show instability peaking to the 10 s crossover due to the loop control, before the noise averages down as white PM to 500 s. Thereafter the noise takes on a pink PM characteristic to its maximum tau.

The 100 s WRDO and GPSDO closely follow the free-running Wenzel stability up to their respective loop constants. Thereafter the frequency stability of the 100 s WRDO follows a white PM characteristic due to the frequency ageing being steered. The GPSDO should also achieve a similar result based on Sandenbergh's results, but the limit on recording time hides the long term nature.

Measures of time stability and maximum error are given by the TDEV in figure 4.15 and the MTIE in figure 4.16. The TDEV can be used in conjunction with the

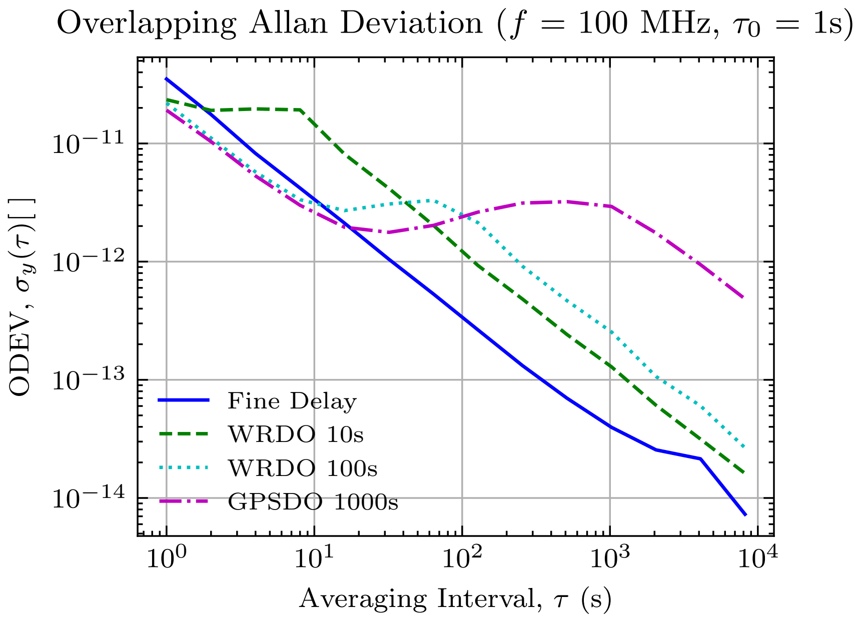

Figure 4.13: Overlapping Allan Deviation of all references.

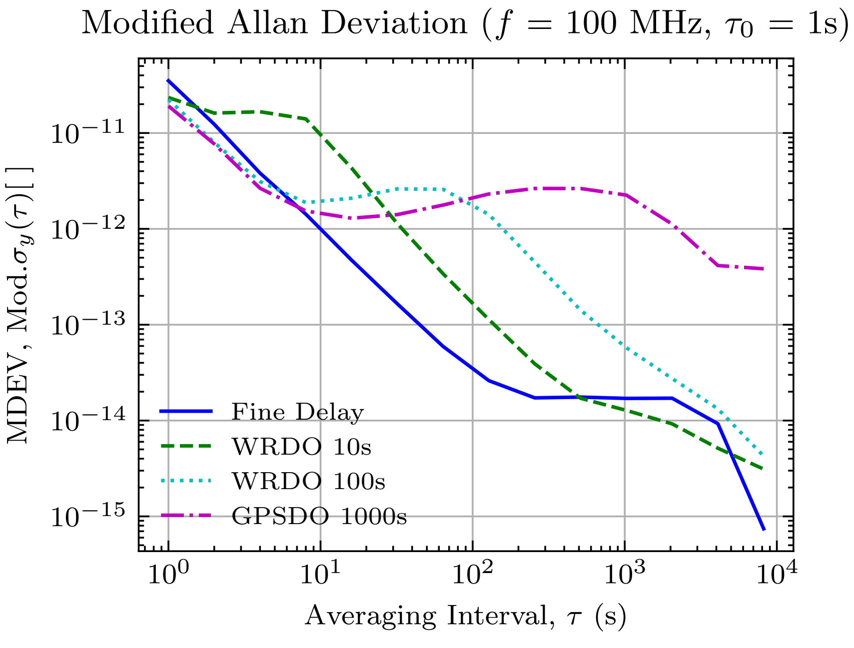

Figure 4.14: Modified Allan Deviation of all references.

initial time/phase offset, initial frequency offset, and frequency drift, to measure the estimated time/phase error using equation 2.17. This is likely only useful for estimations of the phase drift during denial, where the long-term phase can be estimated by the holdover stability of the MOs. The MTIE gives an indication of the maximum expected change in phase over the operational time. This metric is effective for determining the maximum expected range error. The GPSDOs have the largest MTIE, with a peak-peak error of 9.64 ns (4.34 times higher than the requirement), equating to a bistatic range error of 2.89 m, which is almost a change in a single range bin (3.3 m). The WR based references typically remain below the requirement, however the spurious changes in the 100 s WRDO cause the time error to surpass the requirement.

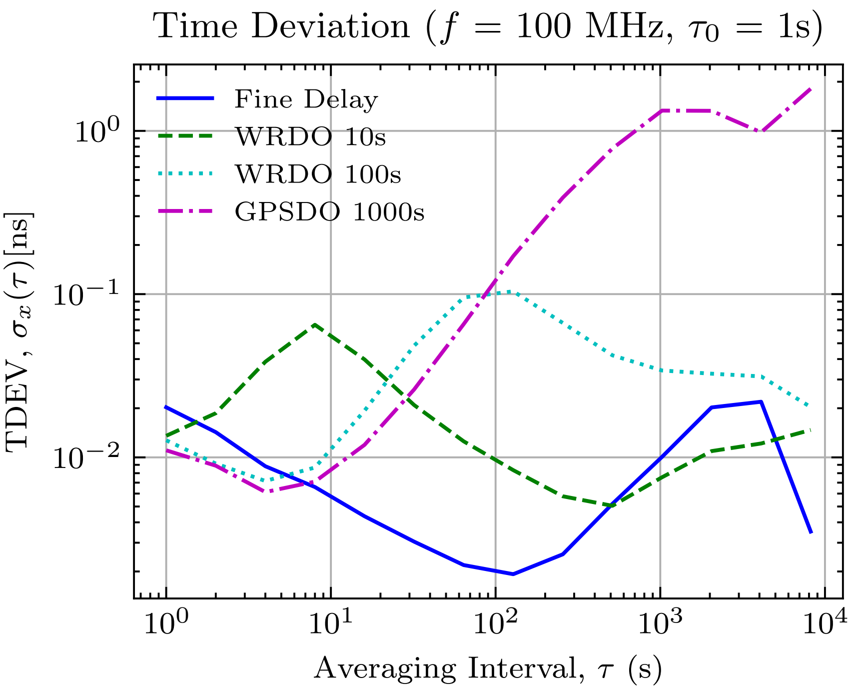

Figure 4.15: Time Deviation of all references.

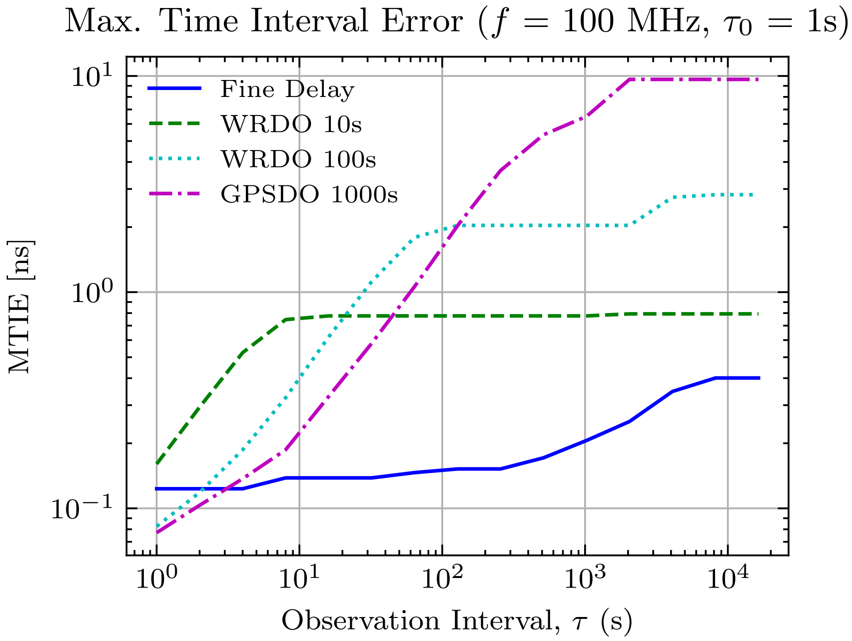

Figure 4.16: Maximum Time Interval Error of all references.

4.5 Conclusion

This chapter detailed the phase and frequency stability of both WR and multi-channel GPSDOs. A DMTD technique was used to measure the phase error between two units with femtosecond level resolution and sub-picosecond jitter. The peak-to-peak noise floor of the DMTD device was found to be approximately 1 ps for 100 MHz oscillators, with a predominantly white PM noise. The frequency stability is more than one order of magnitude below all references for all tau, thus the results of the experiments have negligible measurement noise.

The fine delay and 10 s WRDOs had the lowest long-term phase drift with an MTIE of 0.4 ns and 0.79 ns respectively, which translate to maximum phase changes of 187° and 370° during the longest observation interval. The 100 s WRDO is usually constrained to a peak-peak phase error below 0.4 ns, however, occasionally there are large transient frequency errors, which cause phase errors of approximately ±1 ns. The source of the transients aren't known for certain, but are likely due to EMI on the single-ended to differential driver board. Nevertheless, the frequency error remains sufficiently below the requirement. The GPSDOs have the largest phase changes over the maximum observation window, with an MTIE of 9.6 ns. The larger phase uncertainty in the long-term is due to the increased 1000 s loop time and the higher degree of frequency independence of the GPS receivers.

Whilst long-term phase stability is attractive for fully coherent multistatic radars, in practical systems interested in Doppler, short-term instability is more of a concern than long-term phase drift, due to CPIs typically lasting on the order of a second or less. All of the references have maximum fractional frequency errors that are below the NeXtRAD relative frequency error, however the short-term instability of the fine delays and 10 s WRDOs are higher, and will add higher levels of bistatic phase noise.

The 100 s WRDOs and GPSDOs both show exceptional short-term frequency stability of 2.2×10^{-11} and 1.9×10^{-11} respectively for 1 s averages. Note that the 10 s WRDO has a slightly higher value of 2.3×10^{-11} at 1 s, however, the instability is flicker FM and consequently does not average as white noise, which is detrimental in coherent processing.

Short-term coherence is achieved for all references, and full coherence may be possible in WRDO based systems, specifically the 100 s WRDO if the transient sources can be mitigated. All TF references should be usable within NeXtRAD based on the TF requirements in chapter 3. However, the higher levels of phase noise in the fine delays and 10 s WRDOs are not ideal, and therefore, the 100 s loop constant WRDO and the 1000 s GPSDOs will be used as the TFT system in NeXtRAD.

Chapter 5

Coherent Multistatic Radar TFT using WRDOs and Multi-channel GPSDOs

This Chapter details the use of a WRN for TFT within an operational coherent pulse-Doppler multistatic radar. Both the 100 s WRDO network and the multi-channel GPSDOs are used for the TF system of the NeXtRAD radar.

The NeXtRAD radar was set in quasi-monostatic bistatic geometry due to hardware constraints on the fibre-optic cabling and measurement equipment. Measurements in this chapter are only concerned with the time, phase and frequency stability due to the underlying TFT system, and not on typical NeXtRAD performance.

The radar was set to record a large static clutter return from Roman Rock, a lighthouse 1.83 km into False Bay, Simon's Town. Two PRF were chosen for the measurements. A medium 1 kHz PRF was used to evaluate the close-in phase noise and coherent integration loss of the radar, whilst a 100 Hz low PRF was used to capture the long-term phase stability of the radar. The long-term results additionally make use of the DMTD method presented in the previous chapter to simultaneously record the phase of the MOs.

Higher levels of phase noise are present in the bistatic data, as expected. In the medium PRF case, a larger than expected initial bistatic frequency error (2.03×10^{-10}) and drift (5.3×10^{-12}) is measured for the GPSDOs, partly due to a mismatch in the synthesised LOs, but also partly due to the shorter setup period. The phase drift of the longer measurements perform similarly to the laboratory results, and the DMTD results accurately predict the performance of the real radar results.

The simultaneously recorded DMTD data was also used to reduce the long term phase error to monostatic levels. Bistatic phase noise reduction using the breakthrough was not evaluated due to the proximity of the transmitter.

5.1 NeXtRAD Setup

5.1.1 Radar Geometry

The radar was set to a quasi-monostatic bistatic arrangement (bistatic angle, $\beta \approx 0$), with a baseline of approximately 3 m. Figure 5.1 is a diagram showing how the radar nodes are connected in a LAN network, along with the measurement equipment.

The antennas were placed on the roof of a non-disclosed location in Simon's Town, South Africa. Lengths of 3 m coaxial cables were run from each antenna into the building for the monostatic and passive nodes. Image 5.2 shows the arrangement of the antennas and the GPS antenna mount.

5.1.2 Target

Loop-back tests using a length of RF cable from the transmitter to receiver may have been sufficient to test the frequency stability of the radar. However, it was decided to run complete transmission tests with a real target, to gain insight into the real-world performance of NeXtRAD under the various TFT networks. The

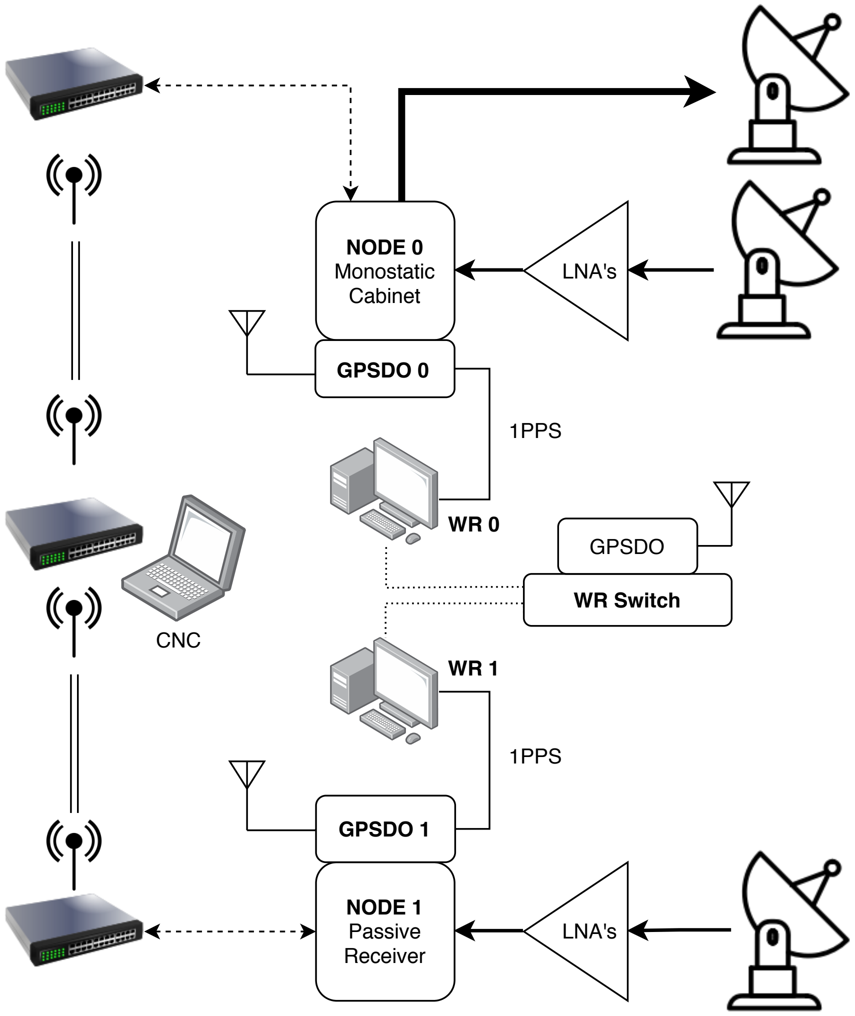

Figure 5.1: NeXtRAD Radar setup with TFT system.

Roman Rock lighthouse, as seen in image 5.3, was chosen as a target, because it has a few properties that make it ideal as a phase calibrator. Firstly, it has a significantly high RCS, thus radar echos have high SNR, and secondly, it is located 1.8 km into False Bay, meaning that it is only competing with sea-clutter. During calm to moderately windy days, the low sea-state, coupled with the low grazing angle of the antennas (< 3°), causes mostly specular reflection, separating Roman Rock significantly from competing clutter.

Figure 5.2: NeXtRAD monostatic and passive antennas (left) and GPS antenna mount (right). Set in a quasi-monostatic bistatic geometry.

5.1.3 Radar Parameters

Only the L-band carrier was used for these experiments, operating at a carrier frequency of 1.3 GHz and a peak power of 1.6 kW. The polarisation was set to vertical for transmit and receive (VV), which was assumed to further separate Roman rock from the sea-clutter due to Roman Rock's vertical extent. The pulse was chosen to be a 5 μs linear frequency modulation (LFM) pulse with 45 MHz of instantaneous bandwidth.

Figure 5.3: Roman Rock in False bay as radar target.

5.1.4 GPSDO and WRDO setup

All 3 GPSDO antennas were mounted to a pole, and placed in clear view of the sky on the same roof. The position of the GPSDOs was set to average for 24 hours, before position hold mode was engaged. The standard deviation of the 3D error was below 1 m, and all GPSDOs had more than 8 tracked satellites on average.

The WRDO network was then setup as before (chapter 3). Newer single-ended to differential driver boards were procured in an attempt to mitigate the transients in the laboratory results.

Figure 5.4 shows the inside of the temperature controlled room. The cables to the GPSs and antennas are run through the windows, which are closed as much as possible to reduce temperature fluctuations.

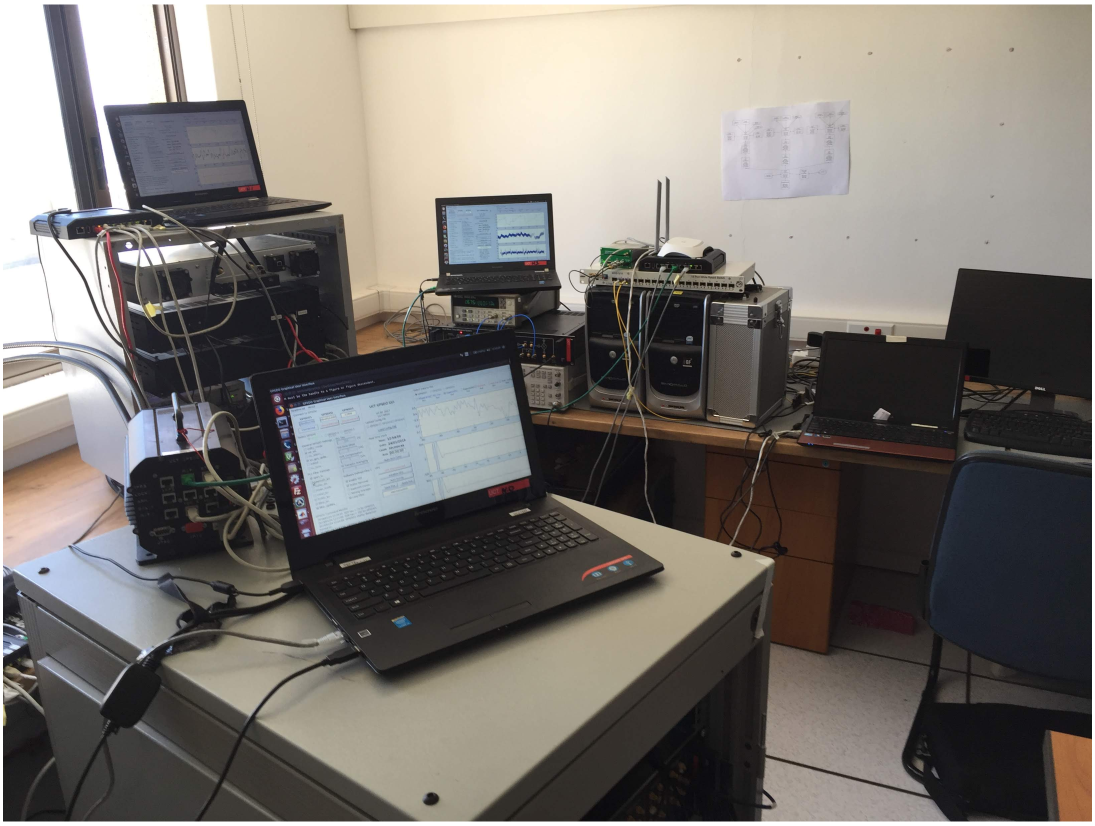

Figure 5.4: Inside control centre. Background: DMTD measurement equipment, WRN. Foreground: GPSDO, GPSDO GUI on node controllers, and active and passive nodes.

5.2 Measurement Testbed

A testbed was created to simultaneously measure the phase stability of the MOs using the DMTD technique, and the effect of the MOs on the phase of real measured radar data. Temperature drift was considered to be negligible during the length of a typical experiment, as most equipment, except the antennas and low-noise amplifier (LNA)s, were inside an air-conditioned room set to 25°C.

5.2.1 DMTD Measurement and Radar Data

The DMTD experiments were run in essentially the same manner as chapter 4, however a shorter sample time of 0.1 s was chosen to analyse shorter-term phase noise of the MOs.

In addition to the DMTD data, the multistatic radar data is simultaneously recorded for the monostatic and bistatic nodes for a low PRF of 100 Hz. The 100 Hz PRF was chosen as it was sufficiently low enough to record 20 minutes of radar data, and to include at least an averaging time of 300 s for the time-domain frequency stability metrics for the WRDO. The choice of 300 s was made to keep the recordings short enough to observe the WRDO frequency stability past the loop time constant. This does mean that the GPSDOs aren't obsevered to their maximum time constant, however these results are predominantly concerned with shorter time constants to keep the radar data file sizes manageable.

The 32-bit in-phase and quadrature (IQ) data, for the 2048 samples per pulse, equates to ADC binary file sizes of 983 MB per node. The measurement interval of the radar was further truncated in software to 15 minutes to assist with the computational time, as 900 s is long enough to observe a tau of 300 s.

In addition to the low PRF data, a medium PRF of 1 kHz was selected to observe the effects of the MOs on the sampled phase noise and coherent integration loss. Due to the magnitude higher data rate, the recordings were limited to 1 minute. Simultaneous DMTD data was not recorded for the medium PRFs due to the short recording interval.

5.2.2 Performance Metrics and Processing

In addition to direct observation of the phase data, the radar and DMTD data were evaluated using time-domain frequency stability metrics, using the same procedure as in chapter 3.

In particular, the MDEV, TDEV, and MTIE are used to evaluate TF stability of the MOs. A minimum sampling time of 0.1 s was selected to observe higher frequency noise components than those of the previous chapter. In addition, the sampled phase of the 1 kHz radar data was used to evaluate the phase noise and estimate the coherent integration losses. The time metrics were not computed for the radar data, as they are predominantly concerned with clock synchronisation, rather than carrier phase stability, where the MDEV is of more use.

The data were processed using AllanTools [50] and data are plotted with the Matplotlib library [119].

5.2.3 Radar Data Processing

The recorded IQ data were reformatted to a fast-time/slow-time matrix, range blanked to 500 range bins, windowed with a Hanning window, and then pulse compressed using a matched LFM waveform. The matched filtered data was visually inspected to determine the range bin corresponding to Roman Rock in fast-time for both the monostatic and bistatic nodes.

The argument of the complex samples, which is the phase series, was unwrapped along the slow-time dimension to remove 2π discontinuities. For comparison to the DMTD data and use in the time-domain stability metrics, the angular phase error is converted to time error by multiplying the phase by $2\pi f_c$, where f_c is 1.3 GHz for the L-band carrier.

DDS Frequency Offset Adjustment

Due to the passive receivers and REX being developed by separate manufacturers, the DDS tuning algorithms for the LOs are slightly different. This causes an issue, as there is a very slight static frequency offset at all synthesised LOs due to the tuning resolution, causing an apparent shift of the static clutter spectrum.

The tuning algorithms were unable to be changed due to their propriety nature and additional development costs. A correction to the recorded bistatic data was therefore done in post-processing by frequency shifting the IQ data along

Table 5.1: Results of DMTD and 100 Hz Bistatic Recordings

		Tau [s]				
TFT	Metric	0.1	1	10	100	Max*
GPSDO (Node 1) [Phase Corr.]	MDEV [1e-12]	21.43 (95.672) [99.32]	3.872 (20.47) [23.59]	1.828 (2.165) [2.721]	4.021 (4.667) [0.679]	0.747 (0.204) [0.1759]
	TDEV [1e-10]	0.012	0.022	0.105	2.432	1.725
	MTIE [ns]	0.113	0.116	0.214	1.512	7.463
WRDO (Node 1) [Phase Corr.]	MDEV [1e-12]	82.43 (92.31) [112.7]	8.040 (34.23) [35.81]	1.951 (2.589) [3.616]	1.997 (1.853) [0.367]	0.269 (0.176) [0.036]
	TDEV [1e-10]	0.048	0.046	0.112	1.173	0.062
	MTIE [ns]	0.046	0.069	0.177	0.601	0.644

No brackets=DMTD, ()=bistatic node, []=DMTD corrected bistatic results.
* Maximum Tau in Computation - MDEV: 400 s (300 s) [250 s], TDEV: 400 s, MTIE: 1200 s.

the slow-time dimension. This was was achieved by multiplying the complex samples by a complex exponential with the exact calculated frequency offset.

The calculated offset for 1.3 GHz, based on discussions with Reutech Radar Systems, is 94.249 844 55 Hz. Uncorrected data would see a constant Doppler shift of the same magnitude, translating to an apparent velocity offset of approximately 10.9 m/s.

5.3 Results

5.3.1 1 kHz PRF Bistatic Data

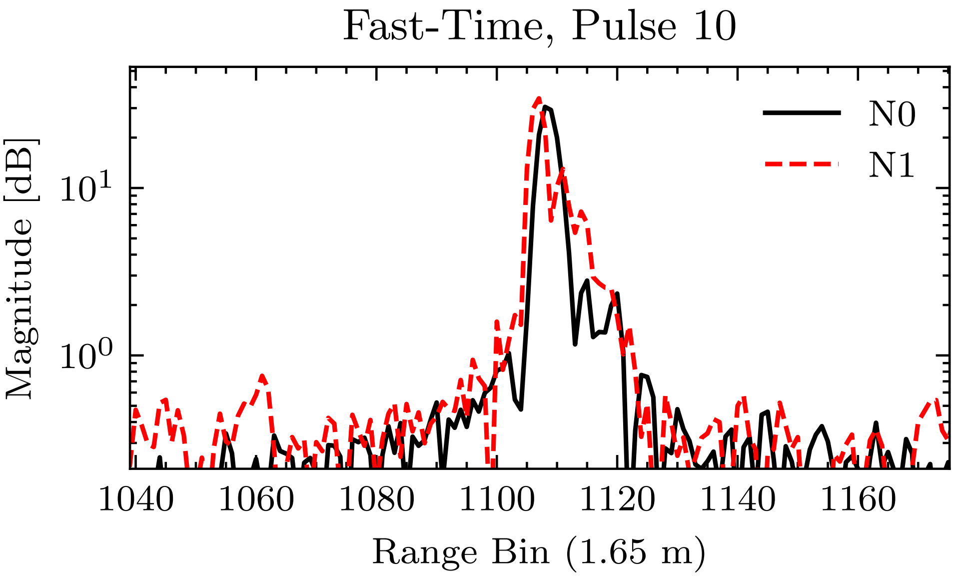

Figure 5.5: Roman Rock range bin for Monostatic (N0=1108=1828 m) and Bistatic (N1=1107=1287 m).

Figure 5.5 shows fast-time for an arbitrary pulse for the GPSDO arrangement. Roman Rock is located in range bin 1108 and 1107 for the monostatic node and bistatic node respectively. Note, that the real range resolution is 3.3 m, however NeXtRAD over-samples by a factor of 2. The WRDO range initially fell within the 1104^{th} range bin. The range was shifted by changing the bistatic receiver's GPS receiver cable delay by 4 ns.

A time series of the phase for the 1 kHz PRF data is given in figure 5.6. The periodic phase jumps in the bistatic nodes (~ 5 s) are not well understood at the moment, but are likely due to spurious mixing with the high levels of powerline modulation.

The WRDOs have a negligible frequency error (1.11×10^{-12}) and drift (8.05×10^{-15}) during the minute long recording. However, the initial frequency error (2.03×10^{-10})

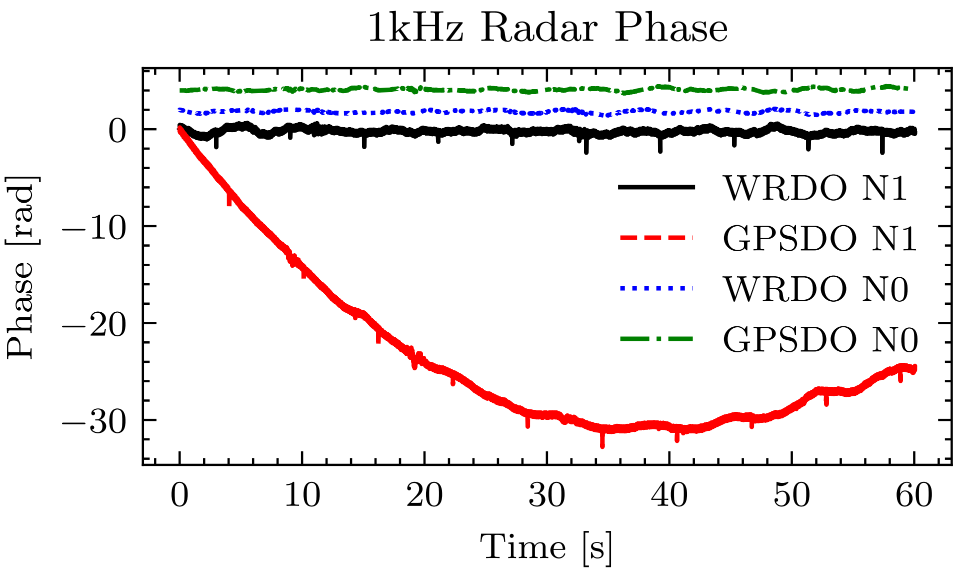

Figure 5.6: High PRF (1000 Hz) measurement of radar phase.

and drift (5.3×10^{-12}) of the bistatic node's GPSDO is significantly higher than predicted. The large offset is most likely due to the setup time of the experiment, and not leaving the GPSDOs to run sufficiently long to reduce the offset to the levels seen in chapter 4, or a transient frequency event due to GPS RFI.

The initial offset would lead to 95.2° of phase drift within 1 s, even though the frequency error is normally within specification. Figure 5.7 shows the estimated frequency and phase drift, with overlays of the estimated linear regression drift lines. The frequency data has been averaged over a period of 1 s to show the average fractional frequency for the 1 s CPI.

Figure 5.8 shows the corresponding single-sideband power spectral density of phase in $dBrad^2/Hz$. GPSDO data without drift removal is used to illustrate the effects of the frequency offset, thus there is significantly more power at 1 Hz. An increase in phase noise for the bistatic nodes is apparent, along with a large increase in the spurious powerline frequency at 50 Hz. Improved filtering of the mains supply in both nodes is required to mitigate this interference, where the powerline spurs reach $-23 dBrad^2/Hz$ in the bistatic nodes. The higher levels of phase noise in the bistatic node was expected, due to the uncorrelated phase

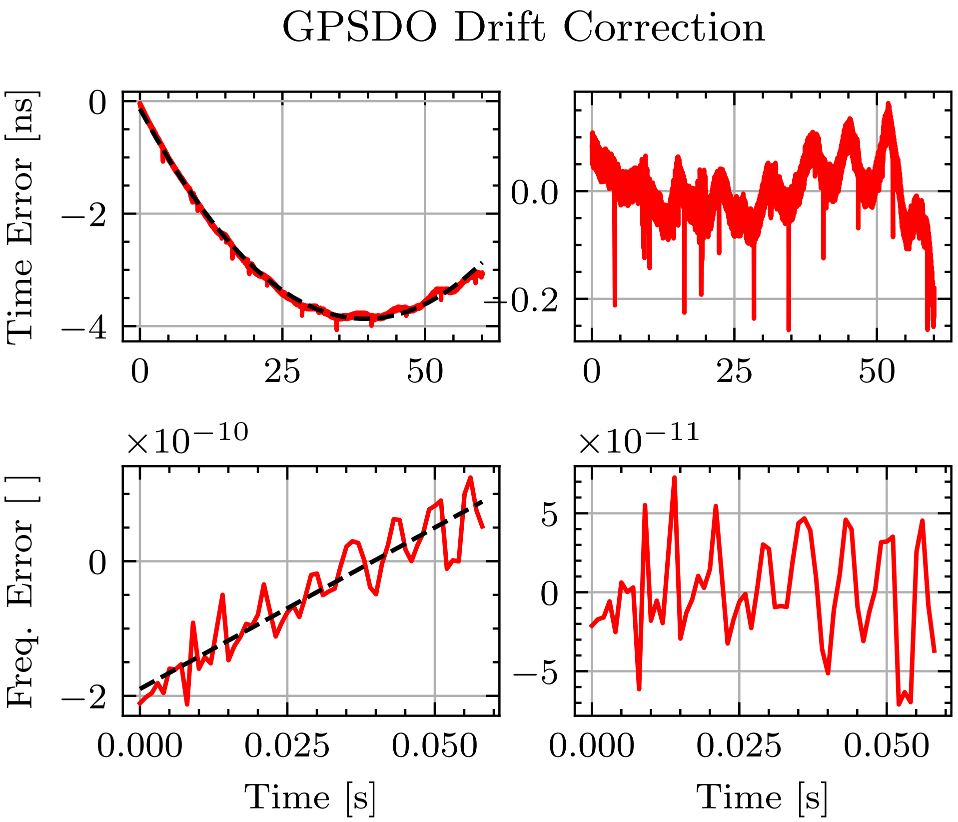

Figure 5.7: GPSDO with frequency drift and offset removed. Frequency data for 1000 s phase average. Estimated offset of 2.03×10^{-10} and frequency drift of 5.3×10^{-12}

noise of the transmitter and the effects of PRF folding of the IF from aliasing. The white PM noise floor of the monostatic node is 10 dB below the bistatic, whilst the flicker corner appears below 50 Hz. The bistatic receiver appears to have flicker PM noise from 1 Hz out to 500 Hz, whereas the flicker corner can be seen in the monostatic node at 60 Hz. Random walk of phase is seen at 1 Hz due to the phase drift and steering.

The coherent integration loss is plotted in figure 5.9 as a function of pulses integrated. Using Richards' analytic expression [65], the phase variance of the monostatic WR and GPSDO nodes (Node 0) are 0.021 and 0.019 radians, leading to negligible coherent integration losses of 0.0019 dB and 0.0015 dB, respectively. Due to the GPSDO's high frequency error at the beginning of the recording, the

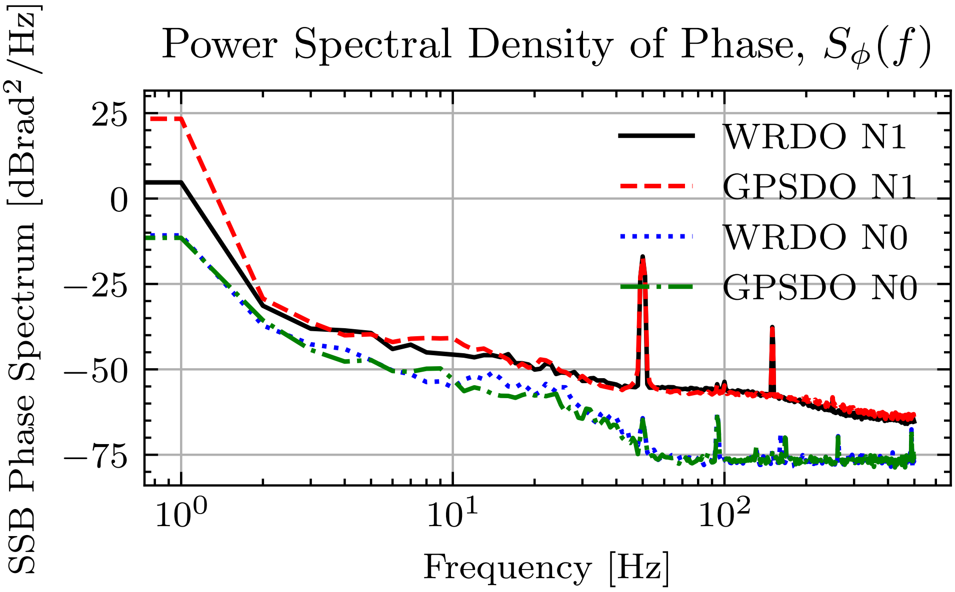

Figure 5.8: Single-sided spectral density of phase of static target.

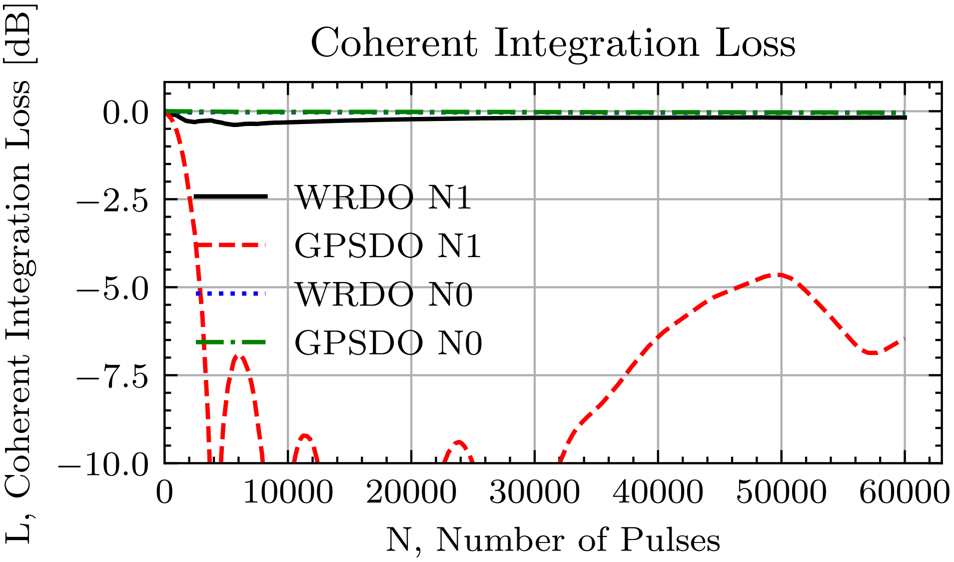

Figure 5.9: Coherent loss function vs number of pulses integrated for phase data in figure 5.6.

coherent loss reaches 3 dB within 2 s, whilst the WRDOs remain within 0.4 dB. These results indicate that both references should be suitable for the typical coherent integrations periods of L-Band multistatic radars, which are usually limited to below 1 s due to target range bin migration and Doppler smearing.

5.3.2 DMTD Phase and 100 Hz Bistatic PRF Data

Although higher PRFs are preferable for short-term instability measurements, excessive data are recorded. Using a PRF of 100 Hz, simultaneous radar and DMTD data were recorded during a the 15 minute interval.

Phase and Frequency

Figure 5.10 shows the recorded phase series of the 100 Hz experiment and the simultaneous DMTD recording.

In figure 5.10a the peak-to-peak drift of the bistatic WRDO node is bounded below 1 ns. The bistatic GPSDO node drifts to an error of 6 ns (57 rad/3266° at 1.3 GHz) in 900 s, and continues to drift past the recording length. Figure 5.10b shows the DMTD measured MO phase errors, which portray a similar picture, showing the dependence of the bistatic receiver phase on the relative MO phase error. The additional GPSDO drift is apparent, with a maximum value of 7.64 ns.

The corresponding fractional frequencies are shown in figure 5.11. The WRDOs have higher short term frequency noise than the GPSDOs, likely due to the larger loop bandwidth and single-ended to differential converter. The higher WRDO noise is more apparent in the DMTD data.

The largest frequency error for the WRDOs is 5.46×10^{-11} and 1.80×10^{-10} for the DMTD and radar data, respectively. Likewise, the largest frequency errors for the GPSDOs are 5.62×10^{-11} and 1.08×10^{-10} for the DMTD and radar data, respectively. The GPSDO DMTD maximum error appears to be due to an outlier in frequency at 200 s.

The maximum errors are still below the NeXtRAD requirements. Surprisingly, both maximum frequency errors for the DMTD data are below the results in chapter 4, which were 6.68×10^{-11} and 7.71×10^{-11} for the WRDOs and GPSDOs, respectively. Both TF systems therefore surpass the requirements for TFT for L-band.

Frequency Stability

Figure 5.12 shows the corresponding MDEV for the radar and DMTD.

The increased powerline ripple is apparent in the passive nodes' short-term stability, appearing as sinusoidal modulation below $\tau = 0.2s$. Monostatic nodes have a white PM characteristic for long-term (>1 s) recordings due to close-in phase noise cancellation. In comparison, one can see the additive close-in phase noise in the bistatic node due to the uncorrelated receiver STALO. There is surprisingly more short term instability in the WRDOs, which is likely due to increased high frequency phase noise from the single-ended to differential driver adaptation.

The common decrease in stability around 20 s for the radar data is due to the STALO PLL loop bandwidth. A similar increase is seen in the DMTD noise floor, however this is attributed to the PLL of the reference synthesiser.

5.3.3 Time Stability and Maximum Error

Figure 5.13 shows the TDEV and MTIE results of the MO's for the DMTD measurement. Peaking of the time instability for the WR and GPSDO occur at the 100 s and 1000 s time constants, respectively, thereafter the time instability traces the long-term stability of the underlying PPS sources. In the MTIE data, the WRDOs are bounded below 1 ns for all tau. The GPSDOs do not settle to a maximum error in the limited measurement interval, however, from the previous laboratory measurements, it was found that the peak-to-peak error is bounded below 10 ns for observation intervals longer than 1×10^4s.

Whilst these results are similar to the laboratory measured results, the WRDOs do not surpass a phase error of 1 ns, as they had previously done. This is likely

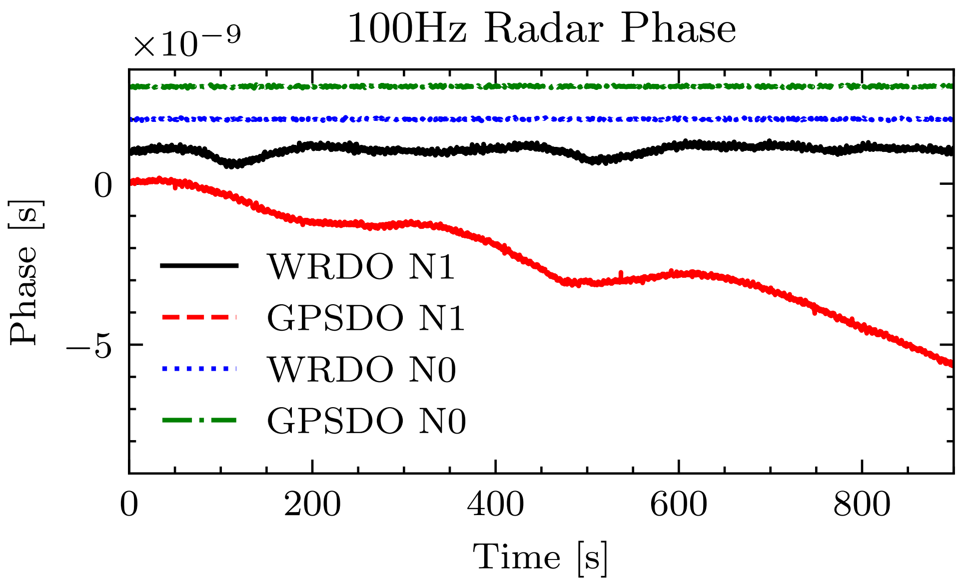

(a) 15 minutes 100 Hz radar phase series

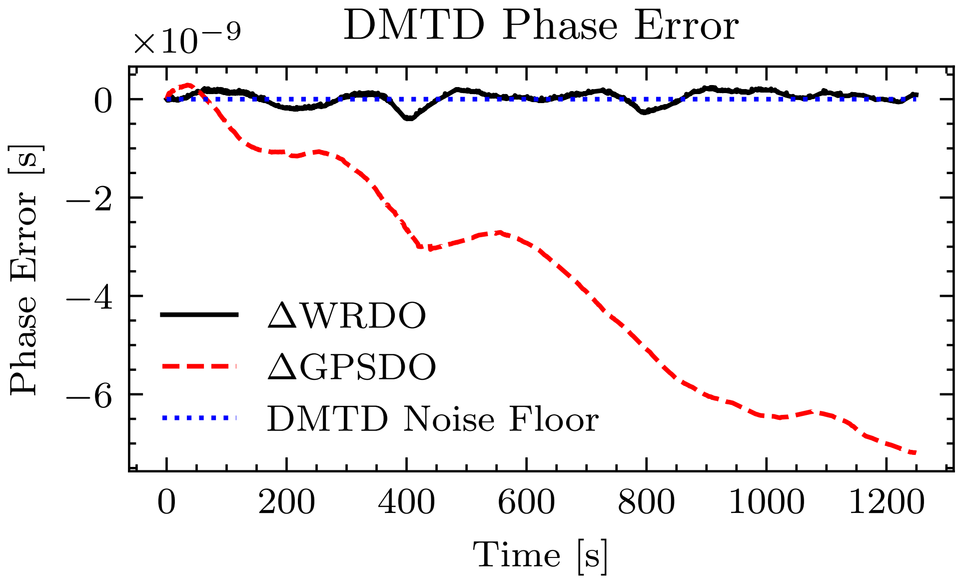

(b) 20 minutes DMTD phase series

Figure 5.10: Phase series of radar (100 Hz PRF) and DMTD (0.1 s sample spacing, τ_0). Initial phase offsets removed. In (a) the phases have been offset by +1 ns for clarity.

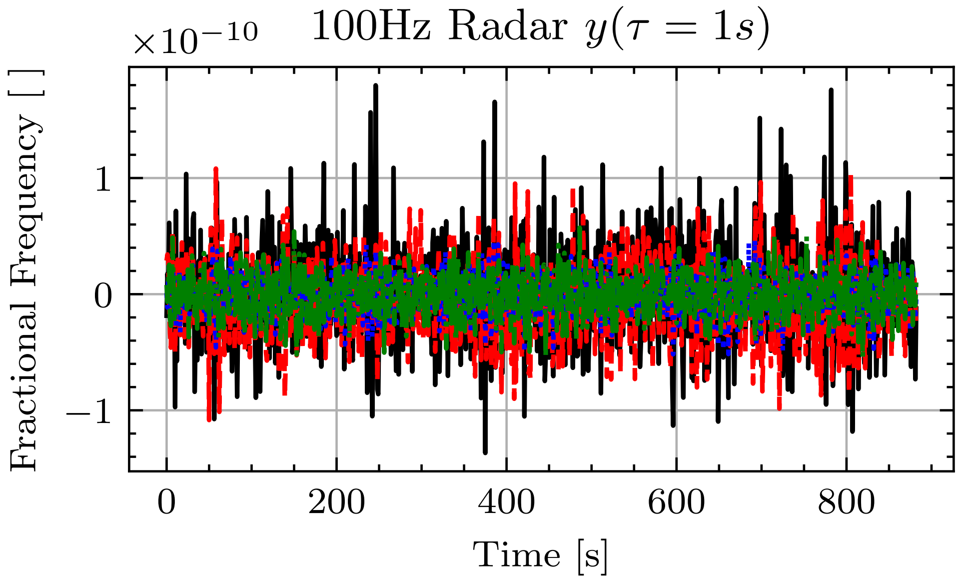

(a) 100 Hz radar fractional frequency for 1 s average ($\tau = 100\tau_0$)

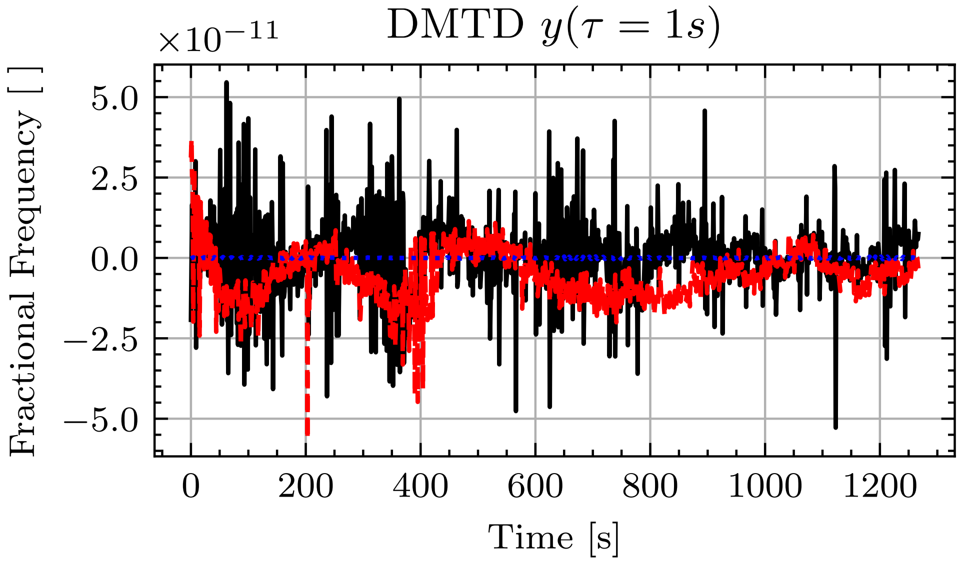

(b) DMTD fractional frequency for 1 s average ($\tau = 10\tau_0$)

Figure 5.11: Average fractional frequency series for radar (100 Hz PRF) and DMTD, averaged with $\tau = 1s$.

just a fortuitous occurrence, and a more robust picture should include these larger, albeit uncommon phase excursions.

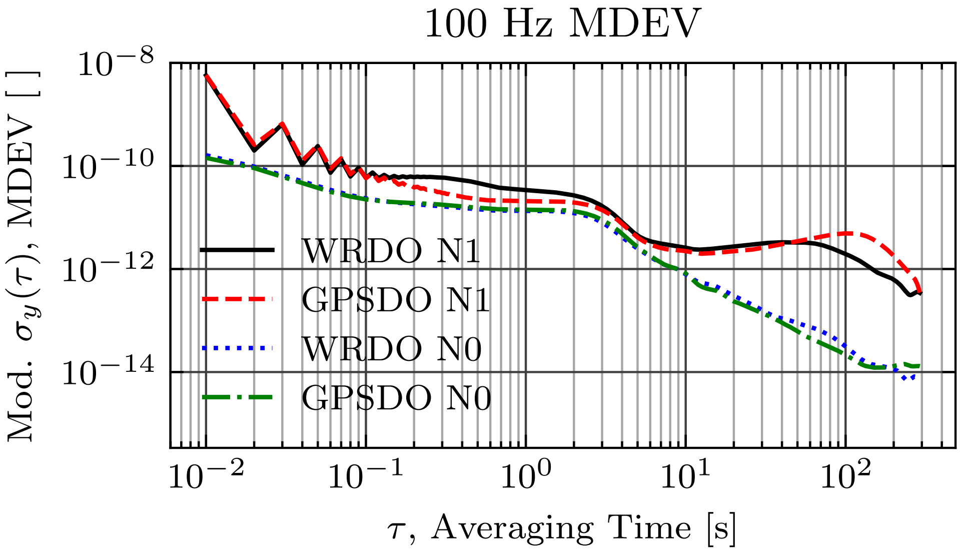

(a) 100 Hz Radar MDEV

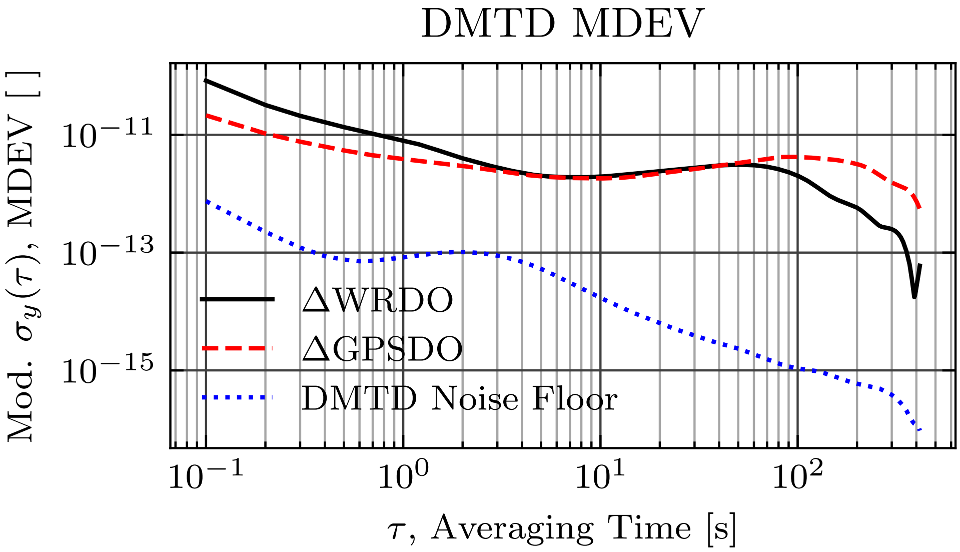

(b) DMTD MDEV

Figure 5.12: MDEV of Radar (100 Hz PRF) and DMTD (0.1 s τ_0).

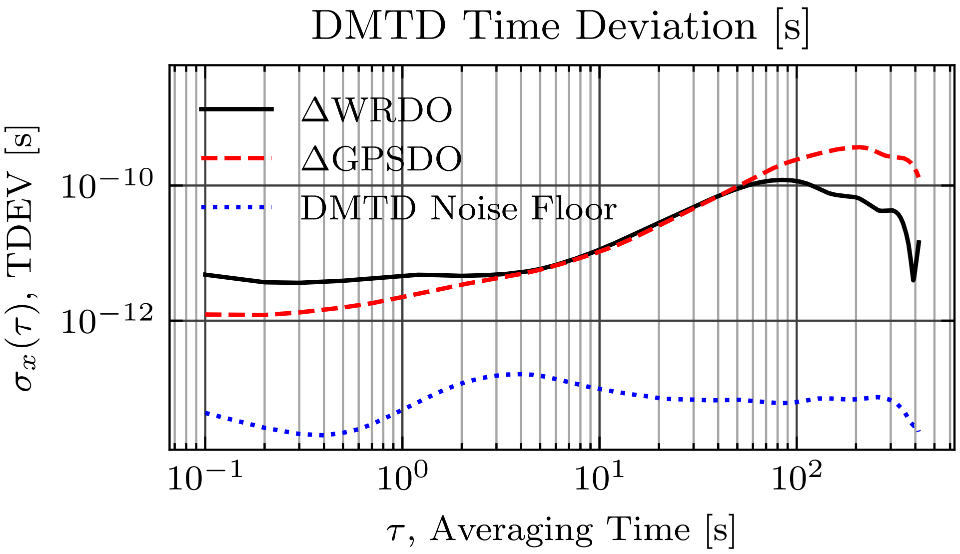

(a) DMTD TDEV

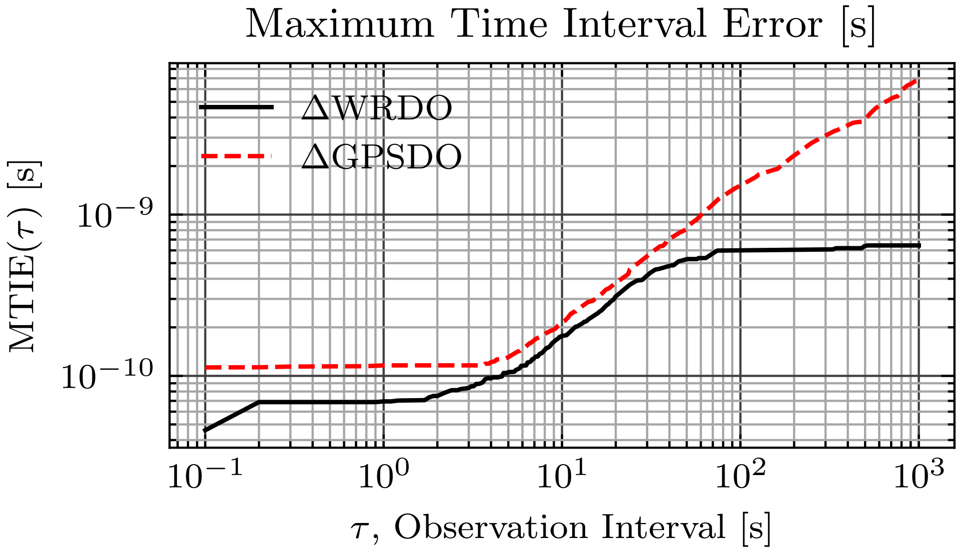

(b) DMTD MTIE

Figure 5.13: TDEV and MTIE of 100 MHz MOs.

5.3.4 Phase Correction using DMTD Phase

The DMTD phase data can be used in post-processing to reduce the influence on the GPSDOand WRDO steering dynamics, improving the long term bistatic

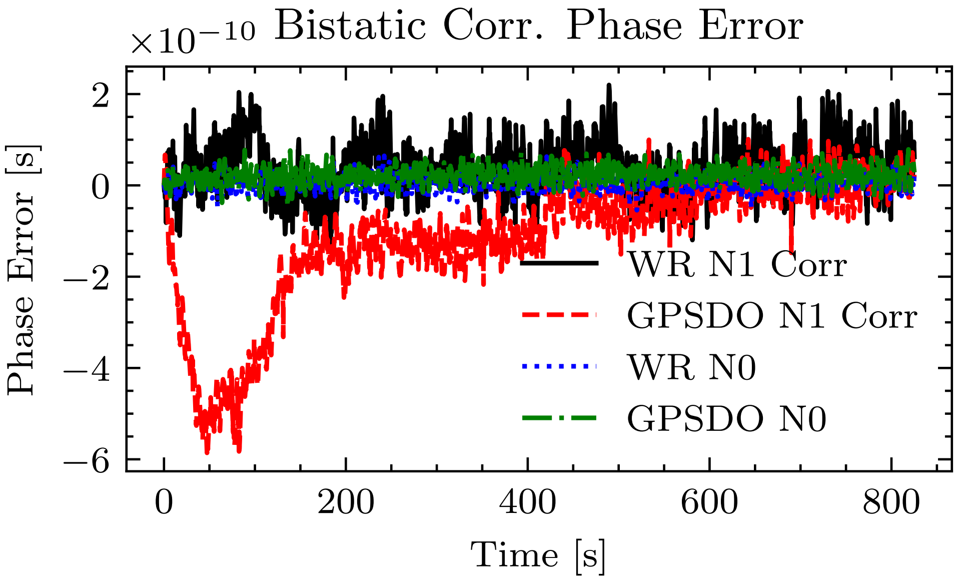

Figure 5.14: First differenced radar phase using DMTD phase data.

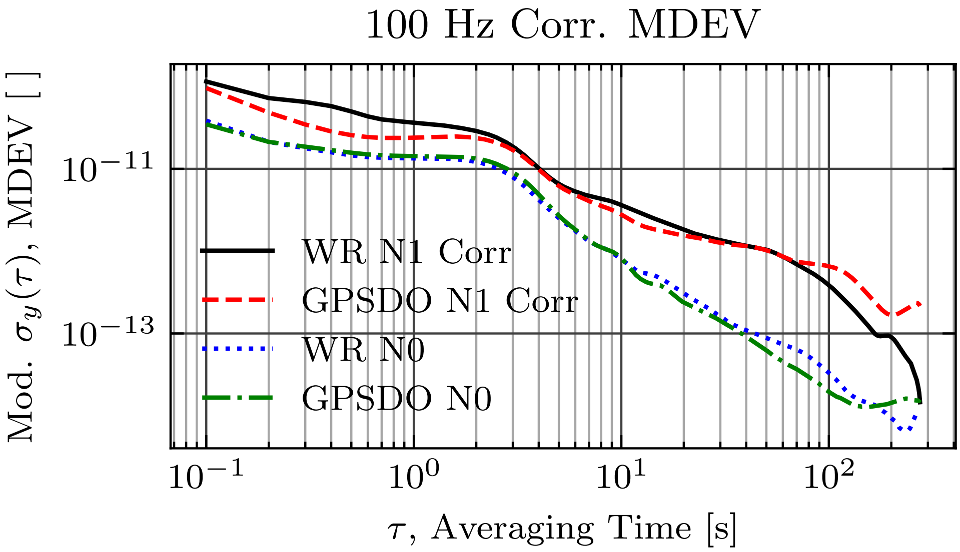

Figure 5.15: MDEV of phase corrected 100 Hz bistatic receiver. Long term stability improvement.

phase uncertainty. The DMTD generates samples at a tenth of the 100 Hz radar PRF, so downsampling or interpolation of the DMTD is required to partially remove the bistatic phase error.

As a proof of concept, downsampling by 10 was used, but effectively removing the phase error at higher PRFs would require interpolation. Figures 5.14 and 5.15 are phase and MDEV plots respectively, demonstrating the long-term phase correction. The datasets were offset by several seconds, thus a peak search of the cross-correlation function was used to align the phase data. Whilst there are still residual phase errors, there is a large reduction in bistatic phase drift, most notably for the GPSDO synchronised node.

Practically implementing this form of phase correction in a real multistatic radar is however unlikely, and a more practical solution is phase correction using the bistatic sampled phase from the direct breakthrough RF in a sidelobe as done in Sandenbergh's trial [1], or echos from a static scatterer in common coverage. This approach was not conducted in this trial due to the co-located nature of the experiment.

5.4 Conclusions

Table 5.2: Requirements vs. Results

	Requirements	GPSDO	WRDO
Clock Synch. Accuracy*	2.22 ns	7.64 ns	0.64 ns
Clock Stability (1s)	1.11×10^{-9}	3.87×10^{-12}	8.04×10^{-12}
STALO Stability (1s)	1.92×10^{-10}	2.47×10^{-11}	3.42×10^{-11}

*The MTIE results for the largest observation interval are used to demonstrate the worst case peak-to-peak clock error.

This chapter detailed the use of WR and multi-channel GPSDOs in an operational coherent pulse-Doppler radar – NeXtRAD. A lighthouse named Roman Rock 1.83 km into False Bay, Simon's Town, was used as a static target to evaluate the phase of the active monostatic node, with comparison to a passive bistatic receiver. The DMTD technique was once again used to directly compare the phase between the MOs and the real bistatic radar phase. A summary of the key results is given in table 5.2, with comparisons to the established TF requirements.

The baseline in these experiments was limited to 3 m between the nodes, therefore both the WRDOs and GPSDOs were co-located, as in the laboratory tests of chapter 4. In theory, the timing baseline for WR was still evaluated, due to the 5 km fibre-optic reel. However, the co-located nature of both references, means that the results are optimistic. In real use-cases, there would be temperature gradients and additional sources of error, such as ground vibration for the fibre-optic cables, and atmospheric turbulence for the GPS signals, which would cause additional frequency instability.

A medium 1 kHz PRF was used to evaluate the short-term frequency instability and coherent integration loss. A high initial frequency offset of 2.03×10^{-10} and drift of 5.3×10^{-12} was present in the GPSDO bistatic node. The rapid loss of coherent gain illustrates that assuming synchronisation and syntonisation is definitely not best practice. One might be tempted to remove the phase drift in processing, however this is typically contrary to the goals of radar, where the drift may be due to target motion. Bistatic breakthrough suppression should be able to reduce the phase noise to acceptable levels, however it was not done in these trials due to the proximity of the transmitter to the receiver.

In the low 100 Hz PRF trial, the frequency stability of the DMTD results are within a few decibels of the laboratory results, and are surprisingly below the laboratory trials. However, the WRDOs have higher levels of short-term instability, once again likely due to the new single-ended to differential driver. The bistatic node's time error closely tracks the DMTD phase, however there is additional noise due to the up- and down-conversion stage, and the additional phase noise due to aliasing from the PRF. The additional noise leads to an increase in the short-term frequency instability, with MDEVs of 2.47×10^{-11} and 3.23×10^{-11} at 1 s for the GPSDOs and WRDOs, respectively. The maximum fractional frequencies are

The MTIE of the WRDOs at the maximum observation interval is below the laboratory results, with a peak-peak error of 0.64 ns. These results do not have any transient frequency events, and therefore are optimistic.

The results illustrate that both GPSDOs and WR are exceptional candidates for low-cost TFT in multistatic radar, and have been successfully operated in NeXtRAD.

The X-band carrier was not used in these experiments due to high levels of HPA noise and a lack of X-band limiters for the LNAs. Any future work should focus on the higher carrier, given cleaner HPAs, as the higher frequency imposes more stringent demands on the relative frequency stability.

Chapter 6

Conclusions

6.1 Summary

This thesis evaluated the performance of a WRN for synchronisation and syntonisation of a coherent pulse-Doppler multistatic radar. WR is a protocol that enables thousands of WR compliant devices to be synchronised to a grandmaster clock with sub-nanosecond accuracy, and picoseconds of jitter. This research was inspired by the vulnerabilities of GNSS based TFT, which are at risk of hostile or accidental jamming, and therefore WR was proposed as a potential replacement to a network relying on GNSS for TFT. Both a WRN and a network of multichannel GPSDOs were used as the TFT system of an operational multistatic radar – NeXtRAD.

The first research question asked whether WR was capable of meeting the requirements for TFT in coherent multistatic radar. A comprehensive review of the TF requirements was therefore given in chapter 2, highlighting the main requirements for accurate and stable synchronisation and syntonisation.

The first requirement considered the time synchronisation of the network. Time synchronisation is required for bistatic range accuracy, and a typical requirement is on the order of a tenth of a compressed pulse width. Frequency stability of the clock is required to maintain this level of synchronisation during holdover. NeXtRAD requires a time accuracy of 2.2 ns, with a frequency stability of 1.11×10^{-9} for 1 s holdovers. Requirements are imposed by the absolute frequency accuracy for velocity estimation, and the relative frequency stability between the transmitter and receiver for coherent processing and Doppler estimation. The NeXtRAD requirements for absolute frequency accuracy are $—3.34 \times 10^{-10}\,|$ for a velocity accuracy of ±0.1 m/s, and relative frequency stability of 1.92×10^{-10} during a 1 s CPI for L-Band.

In addition to the requirements, chapter 2 also considered several other options for TFT, paying specific attention to GNSS and WR. In particular, a brief overview of WR was given, with focus on the application, limitations, and enhancements to the technology. A comprehensive evaluation on the limits of WR for TFT ability was conducted by Rizzi et al [94], where it was discovered that the MMCM was injecting additional phase noise into the WRS, and consequently reducing the phase noise performance of the network. Further analysis on the performance of WR under extreme temperature fluctuations was conducted by Li et al [95], were able to compensate the residual phase errors. There have been several enhancements of the technology, but likely the most important change has been the inclusion of a low-jitter daughterboard for the WRS, that replaces the MMCM with a hardware based PLL, improving the frequency stability by an order of magnitude. Recently, [49] et al showed that WR can be implemented using sub-millimetre wave P2P links, which may be an interesting avenue for mobile multistatic systems. RF-over-WR is also possible, where the generated RF is based on a DDS, that generates the arbitrary RF waveform using control words, distributed on the WRN. This implementation will allow the phase of the RF to be very accurately controlled, which is exciting for fully coherent multistatic and MIMO systems.

The design of a WRN for use in NeXtRAD was then described in chapter 3. A simple WRN was procured for measurements, which consisted of a WRS, 2 SPECs, 2 fine delay FMCs, 4 fibre Gigabit SFPs, and fibre optic cable in 2 m and 5 km lengths. As the low-jitter daughterboard was not available for these measurements, the switch did not have the phase noise improvement, and the hypothesis for the first research question was that the phase noise out of the fine delay cards would be too high. To address this issue, the WRDO was introduced, which is a hybrid design that uses the PPS output from a WR node as the external PPS reference for a GPSDO. The WRDO therefore combines the WRN's long-term stability, with the short-term frequency stability of a low phase noise OCXO. The additional benefits of using the WRDOs in NeXtRAD is the inclusion of GPS spatial referencing, a quick means of swapping between GPS and WR, and a radar triggering mechanism that was already well implemented into NeXtRAD. For use in NeXtRAD, all of the references were used to steer an ultra low-phase noise Wenzel PLO, which is part of the FDU used by each NeXtRAD node.

The TF performance of the WRN and multi-channel GPSDOs was then evaluated in chapter 4. The devices were measured within a temperature controlled laboratory setting to analyse the phase and frequency characteristics of the respective TFT network at zero-baseline. A DMTD technique was used, which allowed the relative phase between two references to be measured with sub-femtosecond resolution, and sub-picosecond jitter. The time error between two FDUs was measured with a sample time of 1 s, and the results were used to derive the frequency errors for a 1 s interval. The phase results were further processed to obtain several time-domain TF stability metrics, specifically the ODEV, MDEV, TDEV, and MTIE.

The maximum fractional frequency error was used to gauge the worst expected frequency offset for coherent processing. All reference combinations were below the 1.94×10^{-10} requirement, however the fine delays and WRDOs with a 10 s SDF time constant had higher levels of short-term instability, which would deteriorate the close-in phase noise performance of NeXtRAD. However, both of these combinations had the best long-term phase stability due to the WRN. The 100 s WRDO and GPSDOs had the best short-term performance, but had increased phase wander. The phase wander of the WRDO was afflicted by several transients, that caused the phase to change rapidly from a nominal phase error of ±400 ps to ±1.5 ns. The source of the transients is likely due to the single-ended to differential driver, or errors in the phase detection circuitry. Nevertheless, the short-term performance is more important, thus the 100 s WRDOs were chosen for use in NeXtRAD.

In order to answer the second research question, the final chapter, chapter 5, evaluated the TF performance of WR within the NeXtRAD radar. NeXtRAD was setup on the roof of an disclosed location in Simon's Town, South Africa, in a quasi-monostatic geometry. A lighthouse, Roman Rock, was chosen as a static target, to measure the phase of the monostatic and bistatic nodes, when synchronised using WR or GPSDO. A medium 1 kHz PRF was used to measure the phase noise and coherent loss. A larger than expected frequency offset and drift for the GPSDOs was measured, which surpassed the requirement for coherent integration. This error was likely due to the shorter setup time of the system. Both the monostatic and bistatic nodes had presence of a 50 Hz mains line modulation, however due to transmit-receive correlation in the monostatic node, the spurious frequency was heavily suppressed. This was not the case in the bistatic node, where the spurs had a magnitude on the order of -23 dbRad/Hz in the single side-band spectral density of phase. The bistatic phase noise was also appreciably higher due to the matched, but uncorrelated STALOs, coupled with the folding of the phase noise spectrum due to PRF aliasing.

A low 100 Hz PRF was then used to measure the phase over an extended period of 15 minutes, where the radar measurements were augmented with measurements of the MO phase using the DMTD. The WRDOs showed increased frequency noise, again likely due to the single-ended to differential driver picking up interference, however no large transients were present in the data, leading to a peak-peak time error of 0.64 ns. The GPSDOs showed similar performance to the laboratory results, and drifted by 7.64 ns during the 15 minute recording. The long-term phase results of the radar are highly correlated with the MO phase, and so a first difference of the radar phase with DMTD phase was performed to remove the phase wander the picosecond regime. However, using such a method is impractical in real system, thus one should use the direct breakthrough to compensate for the phase wander and close-in phase noise.

6.2 Conclusion and Recommendations

The WR project is a phenomenal feat of engineering in the TF discipline. Until its inception, time transfer over computer networks was limited to the level of microseconds using PTP. Improving on this, required specialised, in-house solutions, where standardisation and vendor support was absent. Now, with the aid of fully opensource hardware and software, and industry backing, the promise of sub-nanosecond synchronisation is a reality.

Since the beginning of this work, the WR community has grown exponentially, with hundreds of labs and institutions evaluating WR for their TFT needs. With this growth, new improvements to WR have emerged, that push the protocol to its fundamental limits. Improvements, like the low-jitter daughterboard, RF-over-WR, and HPSEC are exciting avenues that will enhance the performance of TFT to levels beyond what was achieved in this thesis. It is highly recommended that future coherent multistatic radars consider WR, along with these improvements. Additionally, this thesis did not utilise the Gigabit Ethernet functionality, which would allow not only TFT, but deterministic data transfer for real-time multistatic sensing applications.

This thesis demonstrated the power of WR in the domain of microwave sensing, and showed that it is indeed possible to synchronise and syntonise a coherent multistatic radar with limited alterations to the WRN. A novel WRDO was introduced, that suppressed the short-term instability of the WRN by steering low phase noise OCXOs. The TF performance of a WRN improves upon GPSDO based systems, and avoids the worry of GNSS-denial for critical systems.

Appendix A

NeXtRAD

Table A.1: A Comparison between NetRad and NeXtRad

Parameter	**NetRAD**	**NeXtRAD**
Frequency	2.4 GHz (S-Band)	8.5 - 9.2 GHz (X-Band) 1.235 - 1.365 GHz (L-Band)
Polarisations	VV HH Manual Adjustment	VV VH HH HV
Bandwidth	50 MHz	45 MHz
PRF	50 Hz - 3 kHz	50 Hz - 10 kHz
Pulse Width	0.1 - 10 μs	0.1 - 20 μs
Waveforms	LFM	LFM, NLFM Provision for others
Peak Tx Power	500 W	450 W (X-Band) 1.6 kW (L-Band)
Sample Rate	100 MHz	180 MHz
Synchronisation	UCT GPSDO	UCT GPSDO White Rabbit

NeXtRAD, NetRAD's successor, is a prototype multistatic radar that is being jointly developed by UCT [121] and UCL [122]. In short, it is a fully polarimetric, coherent Pulse-Doppler multistatic radar, comprised of one active transceiver and two passive receivers, and therefore a SIMO radar. The radar uses approximately 10° beam-width directive gain antennas designed to work in X- (horn antennas) and L-Band (parabolic reflectors). See Table A.1 for a full list of Nextrad's system specifications, providing a comparison to the NetRAD system.

A.0.1 Active and Passive Nodes

Nodes in NeXtRAD are nearly identical, and each consist of a transceiver or receiver for the active and passive nodes respectively, a TF subsystem, a Pentek industrial computer for SDR, LNAs, node controllers, and an antenna subsystem. The active node additionally has HPAs, and an additional antenna subsystem for transmission to reduce the cost of purchasing high power RF switches. The passive nodes also include large power supplies with inverters, which allow the nodes to be operated in remote areas with no local mains supply.

REX and Passive Receivers

The REX is a dual-band up- and downconverter designed by Reutech Radar systems [123], and generates the corresponding L- or X-band RF transmissions. The REX converts an IF of 125 MHz to an RF of 1.235-1.365 GHz for L-band or 8.5-9.2 GHz for X-band. However, in practice, the X-band frequency is limited to 8.5 GHz due to roll off of the X-band HPA frequency response. The REX is a 2 LO superheterodyne system, and uses an external 100 MHz reference frequency, which is multiplied to 1 GHz for use in the X-band and L-band LO synthesisers. The 1 GHz is also doubled to 2 GHz for the first LO. Whilst the REX is dual-band, it cannot simultaneously transmit both bands, and must electronically switch to the chosen band via a control signal available using a UTP interface.

The passive receivers were designed and commissioned by UCL, and essentially follow the same downconversion chain as the REX. However, the synthesisers use 48-bit DAC control words, as opposed to REX's 32-bit synthesisers. Whilst

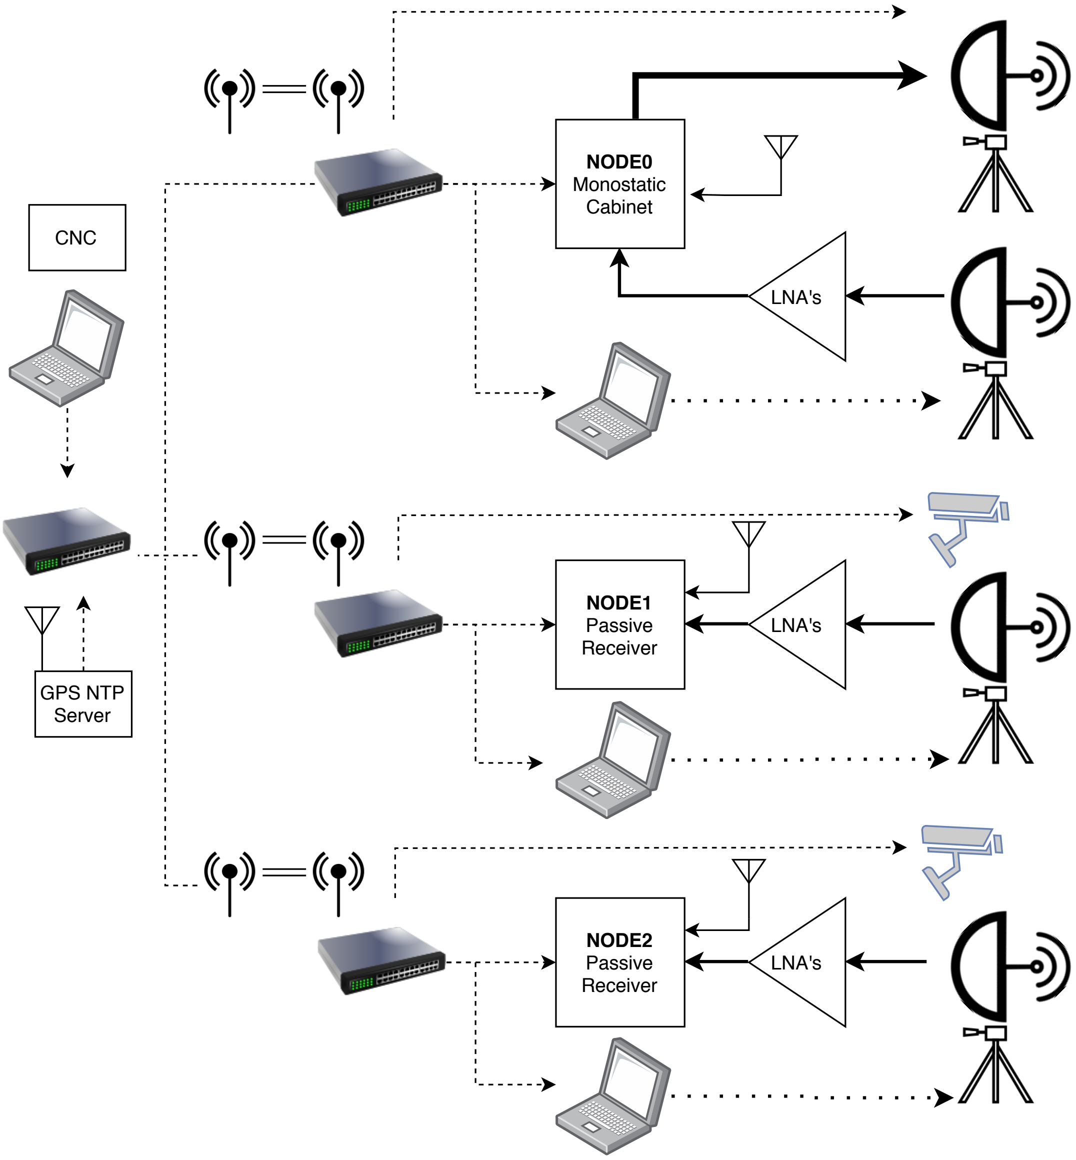

Figure A.1: NeXtRAD multistatic network

this should not typically be a large issue, the algorithms used to synthesise the correct LOs on the REX also incorporate some truncation of the 32-bits, which causes a slight frequency error of approximately 90 to 200 Hz for the L- and X-Bands respectively. The UCL receivers on the other-hand, are able to tune to the correct frequency with an accuracy in the sub-Hz range.

High Power Amplifiers

HPAs are used to increase the output RF from the REX for transmission to 1.6 kW and 450 W for L- and X-band, respectively.

Antenna Subsystem

The antenna subsystem is composed of electronically steerable Skywatcher AZ-EQ6 GT mounts [124], L- and X-band antennas, optical scopes for boresight alignment, internet protocol (IP) cameras, and an IP65 rated box housing the LNAs and L-band polarisation switches.

The parabolic L-band antennas were designed by Paine [125], whilst the X-band horn antennas were designed by Cheng [126]. Both antennas were designed to have a half-power beamwidth (HPBW) of 10°, however the L-band antennas were sub-optimally designed to save on cost and therefore have a slightly higher measured HPBW of approximately 13.9° in azimuth for H-pol. The L-band antennas also have fairly large sidelobes as a consequence of the design. Due to mounting issues, the X-band antenna is not aligned perfectly along the L-band boresight, therefore some parallax is present, although this issue is negligible for ranges larger than a few tens of metres.

The LNA boxes are fixed to the antenna frame, and are connected to the antennas using rigged low-loss coaxial cable. Limiters are placed in the first input stage to stop direct leakage during transmission which would otherwise damage the amplifier stages. The L-band stage uses both a Minicircuits ZX60-P162LN+ and ZX60-V33LN for amplification, with an electronically selectable switch to change polarisation. Both X-band polarisations uses 2 Minicircuits ZX60-153LN for amplification. The amplifiers and switch are powered custom linear power supply.

The active node antennas are connected to the HPA via high power Heliax cables. The receivers are connected to the LNA boxes with 2 m low-loss coaxial cable.

Digitisation

Digital waveform synthesis for the transmitted IF, as well as sampling of the received IF, is handled by a Pentek Cobalt 78621 , which uses a Pentek Cobalt 71261 digital upconverter (DUC) and digital downconversion (DDC) [127] mounted on a PCIe carrier for use in an industrial computer. The Cobalt 71261 has 2 800 MHz DAC channels, and 3 200 MHz ADC channels. A complete discussion of NeXtRAD's digitisation using the Pentek is given in du Plessis's dissertation [128].

The transmitted pulse waveform, centred on the 125 MHz IF, is generated by DUC. Waveforms are stored in the industrial computer's long-term memory as a list of discrete 32-bit samples, which are used for the baseband IQ representation of the chosen waveform. During an experiment, the waveform samples are placed in the Cobalt's DDR3 RAM, where a complex multiplication and interpolation by the DAC generates the IF I and Q waveforms on DAC channels 1 and 2 respectively. The REX input only requires the I component, thus the Q channel is terminated in 50 Ω. Almost any waveform can be synthesised this way, however NeXtRAD typically uses LFM chirps. Whilst the original pulse bandwidth was chosen to be 50 MHz, the radars analog bandwidth was discovered to be 45 MHz, thus the synthesised chirps were designed with this in consideration.

For downconversion from the received IF, NeXtRAD uses a sample rate of 180 MHz for each channel. The choice of sample rate folds the IF into the second Nyquist zone, to a centre frequency of 55 MHz. DDC is used to shift the spectrum to baseband, giving an effective complex sampling rate of 90 MHz. Baseband samples are temporarily stored on a solid state drive (SSD) as a binary file of interleaved 16-bit I and Q words, before being copied over to long-term storage hard disk drive (HDD)s for later data exploration. As there are only 3 receiver channels, channel 1 is used for L-band, and channels 2 and 3 are used for X-band V-pol and H-pol respectively. The electronically selectable polarisation switch is used to choose between H or V-pol for the L-band channel.

Node Controllers

Each node has a local laptop computer running Ubuntu 14.04 LTS, which is responsible for controlling the pedestals, running the GPSDO software, and housing the network file system for the TCU.

Time and Frequency System

The TF system is responsible for generating reference frequencies for each device, setting up the start time of an experiment, and triggering the PRF and control switches. The TF system is comprised of three sub-units, namely the UCT GPSDOs, which generate the MO frequency reference, radar TF synchronisation, and epoch triggering; the FDU, which distributes frequency and triggers to each device; and the TCU, which generates the PRF, biases the HPAs, and controls the L-band polarisation switch and REX transmission frequency. More detail of each subsystem can be found in chapter 3.

A.0.2 Network

The radar is controlled and operated from a single CNC laptop connected to a LAN. Each node has a local node controller laptop, which connects to the network via its own high power power-over-Ethernet (PoE) switch, which is connected to the rest of the LAN using long range unidirectional Ubiquiti Bullet 5 GHz WiFi links [129]. The links are rated up to several kilometres, given ideal conditions and clear LoS. The TCU, Pentek industrial computer, and IP cameras are also connected to the LAN via each node's local switch.

A TimeMachines GPS NTP server [130] is connected to the CNC's switch, which loosely synchronises all network devices to an accuracy of several milliseconds. This ensures that the data is synchronous to within 1 s, and that the network shares a common notion of time for event triggers.

Device orchestration is currently achieved using Ansible and a Python developed GUI. Once users have entered chosen parameters and meta-information (sea-state, weather conditions, experiment description etc.) using the GUI, a header file called 'Nextrad.ini' is generated, which contains parameters for each radar node to adjust their hardware configurations to the selected values in preparation of the next experiment. When the CNC operator ready to start an experiment, Ansible is then used to forward the header file to each node, where custom daemons monitor changes to the file, and adjust each node accordingly.

A.0.3 Operational Overview

The CNC operator sets the experimental conditions (Douglas sea-state, target GPS coordinates, pulse parameters etc.) and a future UTC start time using a GUI developed in Python. This information is stored in a 'Nextrad.ini' header file, which is then distributed over the network to each node using Ansible [131]. Several systems are primed based on the header configuration, and upon reaching the UTC start time, the PRF, HPA switching, and polarisation switching are generated from the TCU on the active and passive nodes.

The MO frequency reference and epoch triggers are created by the same GPSDOs used in NetRAD. The synchronisation accuracy is therefore similar to NetRAD. The one-way cable delays from the GPS antennas to the receivers are removed by offsetting the GPS PPS by the cable delay.

On every rising edge of the PRF, the DUC on the active node's SDR, a Pentek Cobalt 7160 [127], generates a pulse waveform, usually a 45 MHz LFM chirp, centred on an IF of 125 MHz. The IF is upconverted by the REX to the chosen RF carrier and amplified using the corresponding HPA to a peak power of 1.6 kW for L-band and 450 kW for X-band. A separate transmit-only antenna is used to avoid purchasing high power RF circulators.

The received signal is amplified by LNAs before being downconverted to an IF of 125 MHz in each node. All three channels of the Cobalt's ADC are set to sample at a nominal 180 MHz, with a decimation factor of 2, giving an effective sampling rate of 90 MHz. A Hilbert transform generates the quadrature channel for IQ demodulation. The radar uses bandpass sampling, so the 125 MHz centre frequency is folded into the second Nyquist zone to a frequency of 55 MHz. The DDC then downconverts the spectrum to baseband, which is temporarily stored on the node's SSD as three independent binary files.

The CNC controller can then run a quicklook processor called Nextlook [132,133], to rapidly verify if the files are worth saving. If the files are satisfactory, the recordings are transferred to external hardrives which are then uploaded to a long-term storage network drive.

The live footage from the video cameras can be streamed over the network and viewed by the CNC operator, but this is seldom done due to the increased data bandwidth of the network, even with low resolution settings. Instead, the video data is recorded and saved along with each experiment.

Appendix B

Reference Phase Noise Measurements

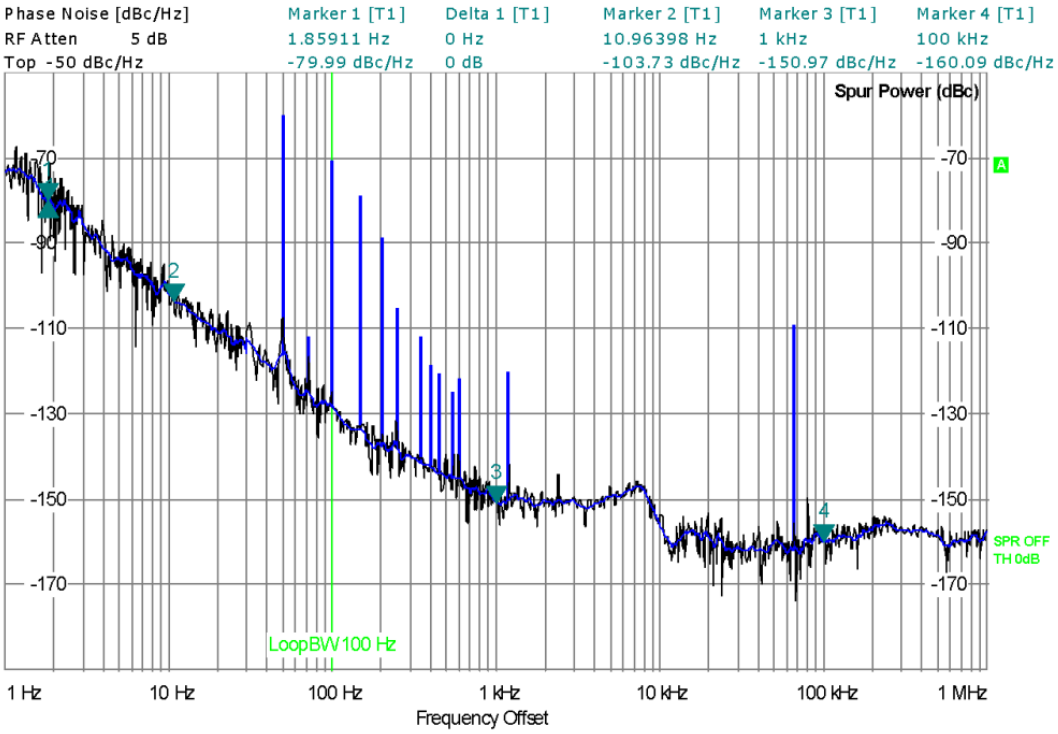

Figure B.1: FDU phase noise at 100 MHz.

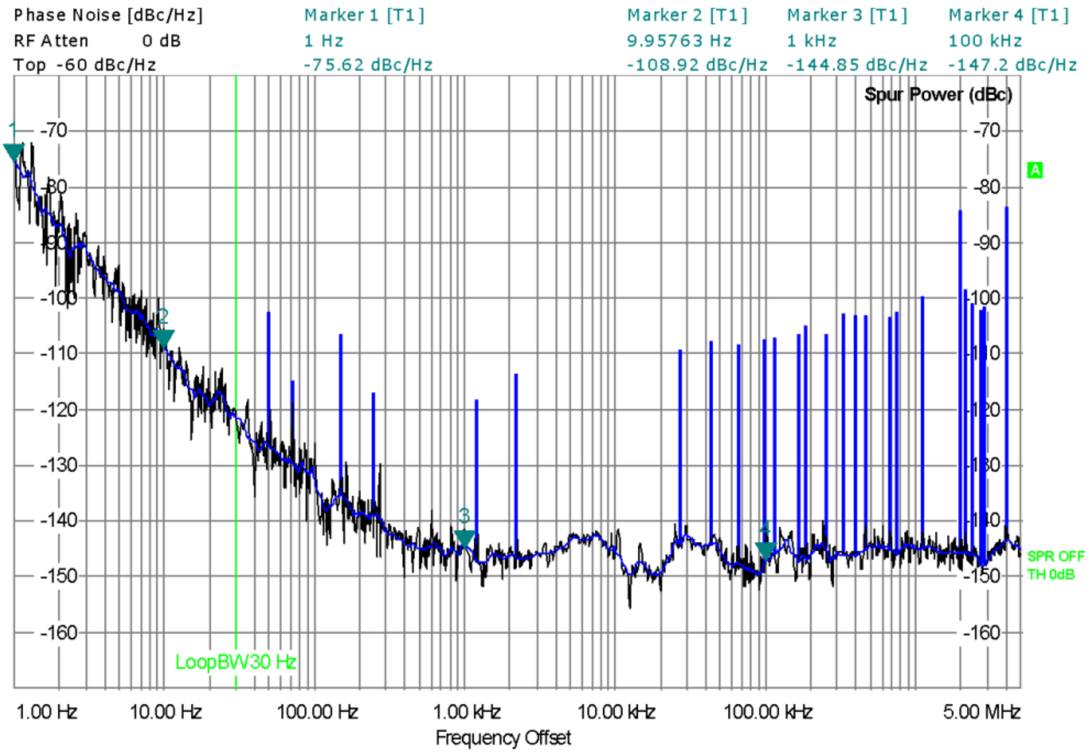

Figure B.2: UCT GPSDO phase noise at 10 MHz.

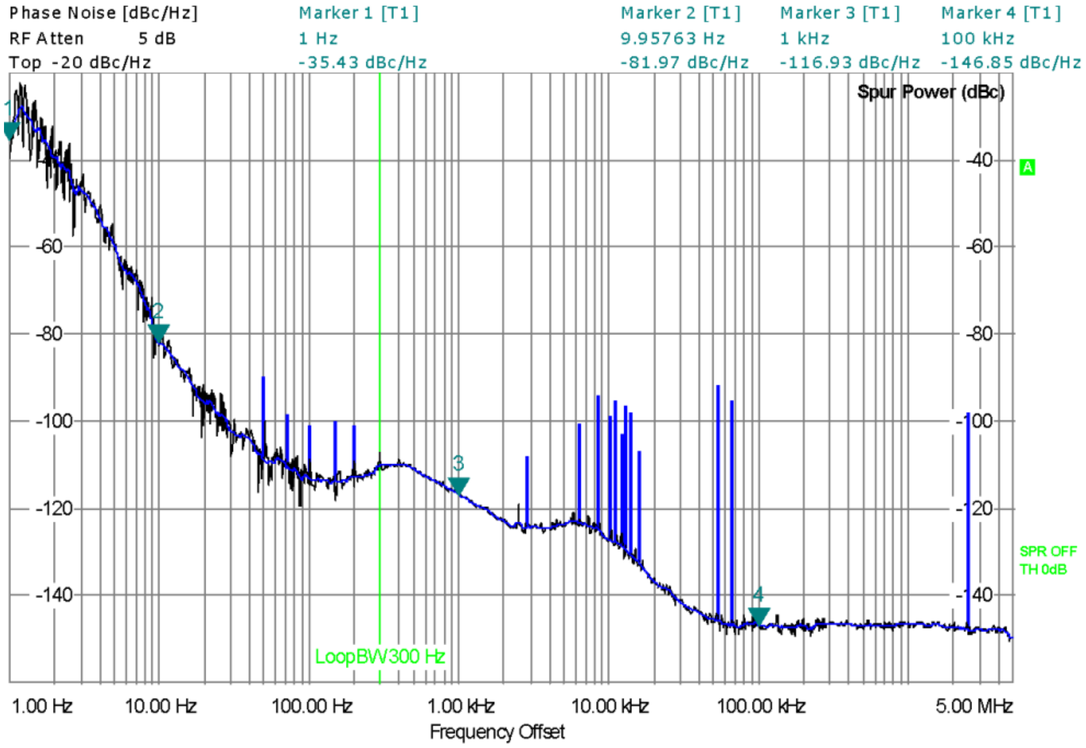

Figure B.3: WR fine delay phase noise at 10 MHz.

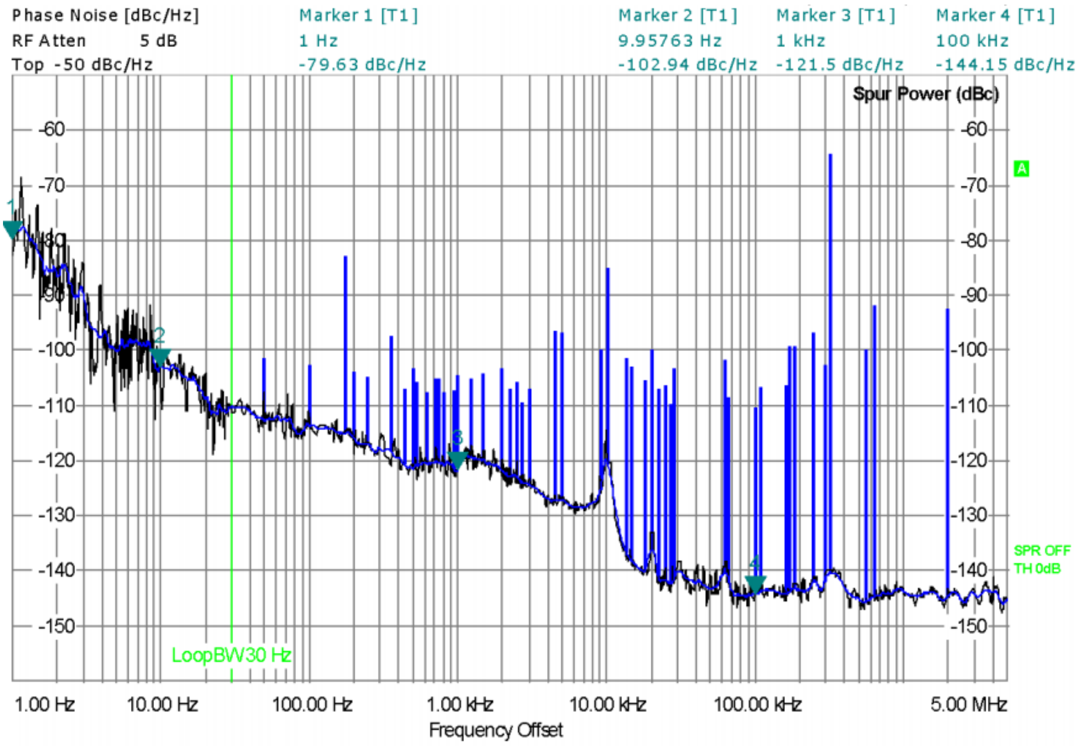

Figure B.4: Meinberg M400 NTP clock phase noise at 10 MHz.

Appendix C

Time and Frequency Stability

This appendix is intended to introduce the concept of frequency stability, and introduce time-domain metrics of frequency stability, namely the MDEV.

C.1 Phase Noise

The perfect frequency reference would be an oscillator that generates a single, unmodulated tone, modelled as a sinusoidal signal, $v(t)$, with some nominal amplitude, V_0, frequency, f_0, and initial phase, ϕ_0. However, real oscillators are affected by noise due to their environment, electronics, aging, and imperfections in the manufacturing process. Random amplitude, $\epsilon(t)$ and phase $\phi(t)$ modulation cause the signal power to spread into the tone's sidebands, whilst frequency inaccuracy, δf, causes a constant offset in the nominal frequency. A realistic frequency source therefore contains and uncertainty in the phase, $\phi(t)$ and amplitude, $\epsilon(t)$,

$$v(t) = [V_0 + \epsilon(t)] \sin(2\pi f_0 t + \phi_0 + \phi(t)) \tag{C.1}$$

Errors due to the uncertainty in the amplitude are almost always negligible[1],

---

[1] In ultra-stable oscillators, the treatment of amplitude uncertainty becomes important due

due to the inherent closed loop gain stability from the Barkhausen criteria for oscillators. Phase and frequency errors on the other-hand are not bounded by the loop gain, and are therefore not compensated. Phase and frequency uncertainty is due to several stochastic and systematic sources of noise. Systematic errors, like fixed frequency offsets, are deterministic in nature, and can therefore be mitigated by careful calibration. However, errors due to the stochastic phase and frequency uncertainty, collectively called *phase noise*, are inherently random properties of the oscillator, and therefore limit the purity of the tone.

C.2 Time-Domain Frequency Stability

In the frequency-domain, phase noise is defined as the single-sided power spectral density of the phase of the signal, at a frequency offset by the carrier frequency. Phase noise is measured in decibels below the carrier in a 1 Hz bandwidth (dBc/Hz). It is represented as a logarithmic plot, $\mathcal{L}(f)$, from frequencies close to the carrier, or *close-in* phase noise (1 Hz - 1 kHz), up to higher frequencies (1 MHz) bounded by the Nyquist sampling frequency of the measurement.

The distribution of the spectrum shows the general phase noise structure of the oscillator, where the slope of spectrum can approximated by several linear segments. The linear segments, are based on power laws, and allow one to determine the dominant phase noise processes.

White phase noise has equal power at all frequencies, and determines the noise floor of the oscillator. White noise is due to thermal agitation of electrons Nyquist-Johnson noise. Flicker phase noise, also known as pink or $1/f$ noise, is a ubiquitous noise source in active electronics and shows a strong signature in oscillators. Lower-frequency noise is due to random-walk, or Brown noise processes, which are due to environmental effects and ageing of the oscillator. Table C.1 gives the estimated power law for each noise process.

---

to amplitude-to-phase conversions, however the effects are minimal in comparison to the phase noise.

Table C.1: Noise Power Laws

Noise Process	$\mathcal{L}(f)$	**ADEV**
Random Walk FM	f^{-4}	τ^{1}
Flicker FM	f^{-3}	τ^{0}
White FM	f^{-2}	$\tau^{-1/2}$
Flicker PM	f^{-1}	τ^{-1}
White PM	f^{0}	τ^{-1}

Phase-noise and drift inherently lead to differences in the quality of each nodes reference. However, spectral measurements of the phase-noise only provide a short-term picture of the frequency stability, missing any long-term effects such as frequency drift due to ageing. A measurement of the relative frequency stability using a time-domain measure provides a richer understanding of the frequency stability. The contribution of the underlying noise processes can be analysed using the time-domain data.

Stability in the Time Domain

Time domain measures of frequency stability are useful for characterising oscillator stability on timescales longer than a second. Due to the presence of coloured phase noise, the phase series does not follow a normal distribution on longer timescales. As a consequence, the mean phase error does not converge, and therefore typical statistical measures of uncertainty, such as the sample variance, are unable to properly characterise the data. The two-sample, or Allan variance (AVAR), is therefore the Institute of Electrical and Electronics Engineers (IEEE) standard measure of frequency stability, as the uncertainty converges for most common phase noise processes.

The AVAR, and it's root ADEV, are powerful predictors of frequency stability, allowing designers to understand the frequency uncertainty of their system, and characterise the underlying phase and frequency noise processes. A plot of the ADEV, or *sigma-tau* plot, provides a useful tool in graphically determining what noise processes are present at various timescales. An estimate of the noise process can be made by looking at the slope of the line on the log-log sigma-tau plot.

However, the standard AVAR has a few limitations namely, 1) it does not distinguish between white and pink phase noise, 2) it requires extensive data to reduce the uncertainty of the estimate (improve confidence in the estimate), and 3) it does not converge for frequency drift. Therefore, newer measures based on the AVAR, such as the overlapping Allan variance (OVAR), modified Allan variance (MVAR) and Hadamard variance (HVAR) are typically used. The MVAR in particular has a phase averaging property that allows white and flicker phase noise to be separated, whilst also effectively reducing the estimate's error by using overlapping samples. The MVAR is also used in determining time variance (TVAR), which helps predict the RMS time error between two clocks given a time interval.

The MVAR can be calculated from frequency samples, or preferentially directly from the phase series,

$$\mathrm{Mod}\sigma_{\mathrm{y}}^2(\tau) = \frac{1}{2\mathrm{m}^2\tau^2(\mathrm{N}-3\mathrm{m}+1)} \sum_{\mathrm{j}=1}^{\mathrm{N}-3\mathrm{m}+1} \left[\sum_{\mathrm{i}=\mathrm{j}}^{\mathrm{j}+\mathrm{m}-1} \mathrm{x}_{2\mathrm{m}+\mathrm{i}} - 2\mathrm{x}_{\mathrm{m}+\mathrm{i}} + \mathrm{x}_{\mathrm{i}} \right]^2 \tag{C.2}$$

Where, N is the number of phase measurements averaged over the interval $\tau = m\tau_0$, m is the averaging factor, and τ_0 is the hardware sampling interval. The TVAR can then be directly calculated from the MVAR, as

$$\sigma_x^2(\tau) = \left(\frac{\tau^2}{3}\right) \mathrm{Mod}\sigma_{\mathrm{y}}^2(\tau) \tag{C.3}$$

Bibliography

[1] S. J. Sandenbergh, *Development and Performance Evaluation of a Multistatic Radar System.* PhD thesis, University College London, 2008.

[2] S. Lewis and M. Inggs, “Synchronisation of Coherent Netted Radar Using White Rabbit compared with one-way multi-channel GPSDOs,” *IEEE Transactions on Aerospace and Electronic Systems*, pp. 1–1, 2020.

[3] S. Lewis, J. S. Sandenbergh, and M. R. Inggs, “Evaluating an off-the-shelf white rabbit system to synchronise network radar via optic fibre,” *2017 IEEE Radar Conference, RadarConf 2017*, pp. 1657–1662, 2017.

[4] S. Lewis, M. Inggs, and J. Sandenbergh, “Investigation of white rabbit for synchronization and timing of netted radar,” in *2015 IEEE Radar Conference*, pp. 214–217, 2015.

[5] P. Beasley, M. Ritchie, H. Griffiths, W. Miceli, M. Inggs, S. Lewis, and B. Kahn, “Multistatic Radar Measurements of UAVs at X-band and L-band,” in *2020 IEEE Radar Conference (RadarConf20)*, pp. 1–6, 11 2020.

[6] R. Palama, F. Fioranelli, M. Ritchie, M. Inggs, and S. Lewis, “Measurements of Multistatic X and L Band Radar Signatures of UAVs,” in *2019 International Radar Conference (RADAR)*, pp. 1–5, 09 2019.

[7] M. Inggs, S. Lewis, R. Palama, M. Ritchie, and H. Griffiths, “Report on the 2018 trials of the multistatic nextrad dual band polarimetric radar,” in *2019 IEEE Radar Conference (RadarConf)*, pp. 1–6, 04 2019.

[8] H. A. Zebker, T. G. Farr, R. P. Salazar, and T. H. Dixon, “Mapping the world’s topography using radar interferometry: the TOPSAT mission,” *Proceedings of the IEEE*, vol. 82, no. 12, pp. 1774–1786, 1994.

[9] D. Mehrholz, F. W. Leushacke, R. Jehn, H. Klinkrad, and M. Landgraf, “Detecting, tracking and imaging space debris,” *ESA Bulletin*, vol. 109, pp. 128–134, Feb. 2002.

[10] Kurekin AA, Loveday BR, Clements O, Quartly GD, Miller PI, Wiafe G, Adu Agyekum K., “Operational Monitoring of Illegal Fishing in Ghana through Exploitation of Satellite Earth Observation and AIS Data,” *Remote Sensing*, 2019.

[11] A. R. Taylor, “Neutral hydrogen in the Milky Way,” in *EAS Publications Series* (L. I. Gurvits, S. Frey, and S. Rawlings, eds.), vol. 15 of *EAS Publications Series*, pp. 243–252, Jan. 2005.

[12] M. Inggs and A. Mishra, “A new taxonomy for symbiotic EM sensors,” 02 2017.

[13] B. Paul, A. R. Chiriyath, and D. W. Bliss, “Survey of RF communications and sensing convergence research,” *IEEE Access*, vol. 5, pp. 252–270, 2017.

[14] M. I. Skolnik, *Radar Handbook*. McGraw Hill, second ed., 1990.

[15] H. Griffiths, “Early history of bistatic radar,” *2016 13th European Radar Conference, EuRAD 2016*, pp. 253–257, 2016.

[16] C. J. Baker and H. Griffiths, “Bistatic and Multistatic Sensors for Homeland Security,” *Advances in Sensing with Security Applications*, pp. 1–22, 2006.

[17] M. Inggs and S. Sandenbergh, “Evaluation of coherent netted radar carrier stability while synchronised with GPS-disciplined oscillators,” in *IEEE Radar Conference*, (Kansas City), pp. 1110 – 1105, 2011.

[18] S. Haykin, “Cognitive radar: a way of the future,” *IEEE Signal Processing Magazine*, vol. 23, pp. 30–40, Jan 2006.

[19] "SKA Homepage." https://www.skatelescope.org/, 2017.

[20] A. K. Mishra and M. Inggs, "White space symbiotic radar: A new scheme for coexistence of radio communications and radar," in *2015 IEEE Radar Conference*, pp. 56–60, Oct 2015.

[21] M. Labib, V. Marojevic, A. F. Martone, J. H. Reed, and A. I. Zaghloui, "Coexistence between communications and radar systems: A survey," *URSI Radio Science Bulletin*, vol. 2017, pp. 74–82, Sep. 2017.

[22] D. Jenn, Y. Loke, M. Tong, E. C. Yeo, and R. Broadston, "Distributed phased arrays with wireless beamforming," in *2007 Conference Record of the Forty-First Asilomar Conference on Signals, Systems and Computers*, pp. 948–952, Nov 2007.

[23] C. A. Tong, *A Scalable Real-time Processing Chain for Radar Exploiting Illuminators of Opportunity*. PhD thesis, University of Cape Town, 2014.

[24] M. Inggs, G. Inggs, S. Sandenbergh, W. Al-Ashwal, K. Woodbridge, and H. Griffiths, "Multistatic Networked Radar for Sea Clutter Measurements," in *IEEE Geoscience and Remote Sensing Symposium*, (Vancouver), pp. 4449 – 4452, 2011.

[25] "Risk to u.s. critical infrastructure from global positioning disruptions." https://rntfnd.org/wp-content/uploads/DHS-National-Risk-Estimate-GPS-Disruptions.pdf.

[26] S. Bartlett, "A wide-area multi-application PNT resiliency solution." https://www.gpsworld.com/innovation-enhanced-loran/.

[27] J. T. Curran, M. Navarro, M. Bavaro, and P. Closas, "A look at the threat of systematic jamming of GNSS." http://insidegnss.com/auto/sepoct17-CURRAN.pdf.

[28] "Russia jams NATO." http://www.atlanticcouncil.org/blogs/new-atlanticist/electronic-jamming-between-russia-and-nato-is-par-for-the-course-in-the-future-but-it-has-its-risky-limits.

[29] K. Hill, "Jamming GPS is illegal, dangerous, cheap and easy." https://gizmodo.com/jamming-gps-signals-is-illegal-dangerous-cheap-and-e-1796778955.

[30] "Selective Availability." https://www.gps.gov/systems/gps/modernization/sa/.

[31] V. Sreeja, "Impact and mitigation of space weather effects on GNSS receiver performance," *Geoscience Letters*, vol. 3, no. 1, 2016.

[32] D. J. Kessler, N. L. Johnson, and M. Matney, "The Kessler Syndrome : Implications to Future Space operations," *Advances in the Astronautical Sciences*, vol. 137, pp. 1–15, 2010.

[33] "Orolia website." https://www.orolia.com.

[34] Microsemi, *Cesium Atomic Clocks to backup GNSS/GPS receivers in Communications Networks*. https://www.microsemi.com/document-portal/doc_download/134694-cesium-atomic-clocks-to-backup-gnss-gps-receivers-in-communications-networks.

[35] "CERN Open Hardware Respository." https://ohwr.org/welcome.

[36] "CERN Website." https://home.cern/.

[37] "GSI website." https://www.gsi.de/en/researchaccelerators.htm.

[38] "White Rabbit Projects Page." http://www.ohwr.org/projects/white-rabbit/wiki/WRUsers.

[39] "MIKES project page." http://www.ohwr.org/projects/white-rabbit/wiki/Mikes.

[40] E. Dierikx, A. Wallin, T. Fordell, J. Myyry, P. Koponen, M. Merimaa, T. Pinkert, J. Koelemeij, H. Peek, and R. Smets, "White rabbit precision time protocol on long distance fiber links," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 63, pp. 1–1, 01 2016.

[41] "Cherenkov Array Website - WR Timestamping." https://www.cta-observatory.org/project/technology/array-control/.

[42] M. Lipinski, T. Wlostowski, J. Serrano, P. Alvarez, J. David Gonzalez Cobas, A. Rubini, and P. Moreira, "Performance results of the first white rabbit installation for CNGS time transfer," *ISPCS 2012 - 2012 International IEEE Symposium on Precision Clock Synchronization for Measurement, Control, and Communication, Proceedings*, pp. 13–18, 2012.

[43] S. Sandenbergh and M. Inggs, "A common view GPSDO to synchronize netted radar," in *IET International Conference on Radar Systems*, (Edinburgh, UK), pp. 1 – 5, 2007.

[44] N. J. Willis, *Bistatic Radar*. SciTech Publishing Inc., second ed., 2005.

[45] M. Weiss, "Synchronisation of bistatic radar systems," in *IEEE Geoscience and Remote Sensing Symposium*, (Anchorage, USA), September 2004.

[46] E. F. Dierikx, A. E. Wallin, T. Fordell, J. Myyry, P. Koponen, M. Merimaa, T. J. Pinkert, J. C. J. Koelemeij, H. Z. Peek, and R. Smets, "White rabbit precision time protocol on long-distance fiber links," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 63, pp. 945–952, July 2016.

[47] Open Hardware Repository, *White Rabbit Users list*. http://www.ohwr.org/projects/white-rabbit/wiki/WRUsers.

[48] M. Lipinski, E. Bij, J. Serrano, T. Wlostowski, G. Daniluk, A. Wujek, M. Rizzi, and D. Lampridis, "White rabbit applications and enhancements," in *IEEE International Symposium on Precision Clock Synchronization for Measurement*, pp. 1–7, 09 2018.

[49] J. E. Gilligan, E. M. Konitzer, E. Siman-Tov, J. W. Zobel, and E. J. Adles, "White rabbit time and frequency transfer over wireless millimeter-wave carriers," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 67, no. 9, pp. 1946–1952, 2020.

[50] *Allan Tools*, 2020 (accessed May, 2020). https://github.com/aewallin/allantools.

[51] V. S. Chernyak, *Fundamentals of Multisite Radar Systems: Multistatic Radars and Multistatic Radar Systems.* CRC Press, 1 ed., 1998.

[52] T. E. Derham, "Design and evaluation of a coherent multistatic radar system," *PQDT - Global*, p. 327, 2005.

[53] T. Derham, S. Doughty, C. Baker, and K. Woodbridge, "Ambiguity functions for spatially coherent and incoherent multistatic radar," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 46, no. 1, pp. 230–245, 2010.

[54] P. Stinco, M. Greco, and F. Gini, "Data fusion in a multistatic radar system," in *Proceedings of the Institute of Acoustics*, vol. 32, pp. 146–151, 09 2010.

[55] S. Doughty, *Development and Performance Evaluation of a Multistatic Radar System.* PhD thesis, University College London, 2008.

[56] J. S. Sandenbergh and M. R. Inggs, "A common view GPSDO to synchronize netted radar," *Radar Systems, 2007 IET International Conference on*, pp. 1–5, 2007.

[57] J. S. Sandenbergh and M. R. Inggs, "Synchronizing network radar using all-in-view GPS-disciplined oscillators," *2017 IEEE Radar Conference, RadarConf 2017*, pp. 1640–1645, 2017.

[58] B. G. M. Herbert, "Modelling of phase noise from frequency references in bistatic radar," in *IMA Conference on Mathematics in Defence*, pp. 1–10, 2015.

[59] W. Wang, "Analytical modeling and simulation of phase noise in bistatic synthetic aperture radar systems," *Fluctuation and Noise Letters*, vol. 06, pp. 297–303, 01 2006.

[60] G. Krieger and M. Younis, “Impact of oscillator noise in bistatic and multistatic SAR,” *IEEE Geoscience and Remote Sensing Letters*, vol. 3, no. 3, pp. 424–428, 2006.

[61] M. A. Richards, J. A. Scheer, and W. A. Holm, *Principles of Modern Radar*, vol. One. Scitech Publishing, 2010.

[62] J. L. Auterman, “Phase stability requirements for bistatic SAR,” in *IEEE National Radar Conference*, (Atlanta), March 1984.

[63] T. Neu, *Clocking the RF ADC: Should you worry about jitter or phase noise?* Texas Instruments, 2017.

[64] J. A. Sheer, *Principles of Modern Radar*, ch. 12. Scitech Publishing, 2010.

[65] M. Richards, “Coherent integration loss due to white gaussian phase noise,” *Signal Processing Letters, IEEE*, vol. 10, pp. 208 – 210, 08 2003.

[66] L. Śliwczyński, P. Krehlik, J. Kolodziej, H. Imlau, H. Ender, H. Schnatz, D. Piester, and A. Bauch, “Fiber-optic time transfer for UTC-traceable synchronization for telecom networks,” *IEEE Communications Standards Magazine*, vol. 1, pp. 66–73, March 2017.

[67] M. Lombardi, A. N. Novick, G. Neville-Neil, and B. Cooke, “Accurate, traceable, and verifiable time synchronization for world financial markets,” *Journal of Research of the National Institute of Standards and Technology*, vol. 121, p. 436, 2016.

[68] S. Schediwy, D. Gozzard, C. Gravestock, S. Stobie, R. Whitaker, J. Malan, P. Boven, and K. Grainge, “The mid-frequency square kilometre array phase synchronisation system,” 05 2018.

[69] J. Levine, “A review of time and frequency transfer methods,” *Metrologia*, vol. 45, pp. S162–S174, December 2008.

[70] J. Yao, “Microwave photonics,” *Conference Proceedings - International Conference on Indium Phosphide and Related Materials*, vol. 27, no. 3, pp. 212–214, 2012.

[71] R. Beresford, W. Cheng, and P. Roberts, “Low cost RF over fiber distribution for radio astronomy phased arrays,” in *2017 XXXIInd General Assembly and Scientific Symposium of the International Union of Radio Science (URSI GASS)*, pp. 1–4, Aug 2017.

[72] C. H. Cox, E. I. Ackerman, G. E. Betts, and J. L. Prince, “Limits on the performance of RF-over-fiber links and their impact on device design,” *IEEE Transactions on Microwave Theory and Techniques*, vol. 54, no. 2, pp. 906–920, 2006.

[73] L. Śliwczyński, P. Krehlik, and M. Lipinski, “Optical fibres in time and frequency transfer,” *Measurement Science and Technology*, vol. 21, March 2010.

[74] L.-S. Ma, P. Jungner, J. Ye, and J. L. Hall, “Delivering the same optical frequency at two places: accurate cancellation of phase noise introduced by an optical fiber or other time-varying path,” *Optics Letters*, vol. 19, no. 21, p. 1777, 1994.

[75] Ł. Śliwczyński, P. Krehlik, and M. Lipiński, “Optical fibers in time and frequency transfer,” *Measurement Science and Technology*, vol. 21, p. 075302, may 2010.

[76] Ł. Śliwczyński, P. Krehlik, A. Czubla, Ł. Buczek, and M. Lipiński, “Dissemination of time and RF frequency via a stabilized fibre optic link over a distance of 420 km,” *Metrologia*, vol. 50, pp. 133–145, feb 2013.

[77] A. Aguasca, J. M. Plana, and A. Broquetas, “1992 ieee,” in *IEEE Frequency Control Symposium*, pp. 492–498, 1992.

[78] C. Li, X. Chen, and Z. Liu, “Over-the-horizon time and frequency synchronization for maneuverable radar system based on troposcatter,” in *2016 IEEE International Conference on Signal and Image Processing (ICSIP)*, pp. 484–487, Aug 2016.

[79] J. Levine, M. Lombardi, and D. Matsakis, “Metrological and legal traceability of time signals Metrological and legal traceability of time signals,”

in *Proceedings of the 49th Annual Precise Time and Time Interval Systems and Applications Meeting*, (Reston), pp. 59–71, 2018.

[80] L. Gasparini, O. Zadedyurina, G. Fontana, D. Macii, A. Boni, and Y. Ofek, "A digital circuit for jitter reduction of GPS-disciplined 1-PPS synchronization signals," in *IEEE International Workshop on Advanced Methods for Uncertainty Estimation in Measurement*, pp. 84 – 88, 08 2007.

[81] T. Johnsen, "Time and frequency synchronization in multistatic radar. consequences to usage of GPS disciplined references with and without GPS signals," in *Proceedings of the 2002 IEEE Radar Conference (IEEE Cat. No.02CH37322)*, pp. 141–147, 2002.

[82] J. Wurman, M. Randall, C. L. Frush, E. Loew, and C. L. Holloway, "Design of a bistatic dual-doppler radar for retrieving vector winds using one transmitter and a remote low-gain passive receiver," *Proceedings of the IEEE*, vol. 82, pp. 1861–1872, Dec 1994.

[83] H. Breit, M. Younis, U. Balss, A. Niedermeier, C. Grigorov, J. Hueso-Gonzalez, G. Krieger, M. Eineder, and T. Fritz, "Bistatic synchronization and processing of TanDEM-X data," in *2011 IEEE International Geoscience and Remote Sensing Symposium*, pp. 2424–2427, 2011.

[84] T. Wlostowski, "Precise Time and Frequency Transfer in a White Rabbit Network," Master's thesis, Warsaw University of Technology, 2011.

[85] M. Lipiński, T. Włostowski, J. Serrano, and P. Alvarez, "White rabbit: A PTP application for robust sub-nanosecond synchronization," *IEEE International Symposium on Precision Clock Synchronization for Measurement, Control, and Communication, ISPCS*, pp. 25–30, 2011.

[86] M. Rizzi, M. Lipiski, T. Wlostowski, J. Serrano, G. Daniluk, P. Ferrari, and S. Rinaldi, "White rabbit clock characteristics," in *2016 IEEE International Symposium on Precision Clock Synchronization for Measurement, Control, and Communication (ISPCS)*, pp. 1–6, Sep. 2016.

[87] P. Moreira, P. Alvarez, J. Serrano, I. Darwezeh, and T. Wlostowski, "Digital dual mixer time difference for sub-nanosecond time synchronization in ethernet," in *IEEE International Frequency Control Symposium*, pp. 449–453, 2010.

[88] "White Rabbit PTP Core project page." https://ohwr.org/project/wr-cores/wikis/Wrpc-core.

[89] CERN, "CERN open hardware licence wiki." https://ohwr.org/cernohl.

[90] "SPEC Project Web Page." http://www.ohwr.org/projects/spec.

[91] "FMC Project Web Page." https://ohwr.org/project/fmc-projects.

[92] T. Wlostowski, *Fine Delay Design Notes*, 2013.

[93] "White Rabbit compatible Fibre and SFP." https://ohwr.org/projects/white-rabbit/wiki/SFP, 2020.

[94] M. Rizzi, M. Lipinski, P. Ferrari, S. Rinaldi, and A. Flammini, "White rabbit clock synchronization: Ultimate limits on close-in phase noise and short-term stability due to FPGA implementation," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 65, pp. 1726–1737, Sep. 2018.

[95] G. Gong, H. Li, W. Pan, and J. Li, "Temperature effect and correction method for LHAASO KM2A timing synchronization node," in *2014 19th IEEE-NPSS Real Time Conference*, pp. 1–1, 2014.

[96] F. Torres-Gonzlez, J. Daz, E. Marn-Lpez, and R. Rodriguez-Gmez, "Scalability analysis of the white-rabbit technology for cascade-chain networks," in *2016 IEEE International Symposium on Precision Clock Synchronization for Measurement, Control, and Communication (ISPCS)*, pp. 1–6, September 2016.

[97] M. Rico, J. Aubry, C. Botteron, and P. Farine, "ns-Level Time Transfer over a Microwave Link using the PTP-WR Protocol," *IFCS EFTF (Joint*

Conference of the IEEE International Frequency Control Symposium & European Frequency and Time Forum), 2015.

[98] R. M. Hambly and T. A. Clark, "Critical evaluation of the motorola m12+ gps timing receiver vs. the master clock at the united states naval observatory, washington, dc," in *34th Annual Precise Time and Time Interval (PTTI) Meeting*, pp. 109–116, 2002.

[99] "Oscilloquartz 8788 datasheet." `http://www.tamadevice.co.jp/pdf/oscilloqartz/8789.pdf`.

[100] C. Systems, "Motorola T2000 Timing Antenna Specifications." `https://www.cnssys.com/cnsclock/backup/T2000.pdf`, 2017.

[101] "ACAM TDC-GP2." `https://www.pmt-fl.com/products/tdc-gp2`.

[102] S. Scott, " RHINO: reconfigurable hardware interface for computation and radio," Master's thesis, University of Cape Town, 2011.

[103] "Rohde and Schwarz FSU26 Spectrum Analyser." `https://www.rohde-schwarz.com/us/product/fsu-productstartpage_63493-7993.html`, 2017.

[104] "Wenzel PLO." `http://www.wenzel.com/wp-content/uploads/Premium-HF-PLO.pdf`.

[105] CERN, "White rabbit switch wiki." `https://www.ohwr.org/projects/white-rabbit/wiki/switch`.

[106] "Axcen SFP Website." `http://www.axcen.com.tw/page/product/p03.aspx?kind=137`.

[107] Meinberg, "Meinberg M400 product page." `https://www.meinbergglobal.com/english/products/advanced-rail-mount-ntp-server.htm`.

[108] "WR Calibration." `https://www.ohwr.org/projects/wr-calibration/wiki/wiki-white-rabbit-calibration`.

[109] T. Wlostowski, *FMC Delay 1ns 4cha (Fine Delay) Long term test report*, 2013.

[110] "Fine delay Project Web Page." http://www.ohwr.org/projects/fmc-delay-1ns-8cha/wiki.

[111] "Analog filter design wizard." https://www.analog.com/designtools/en/filterwizard/.

[112] CERN, "HPSEC wiki page." https://ohwr.org/project/hpsec/wikis/home.

[113] W. J. Riley, *A Small Dual Mixer Time Difference (DMTD) Measurement System*. www.wriley.com/A%20Small%20DMTD%20System.pdf.

[114] D. W. Allan, "Report on NBS dual mixer time difference system (DMTD) associated with phase 1 of GPS built for time-domain measurements," tech. rep., NIST Time and Frequency Division, 1976.

[115] J. Sherman and R. Jrdens, "Oscillator metrology with software defined radio," *Review of Scientific Instruments*, vol. 87, 05 2016.

[116] "ULC-1.5FT-SMSM+ Minicircuit SMA cables." https://www.minicircuits.com/WebStore/dashboard.html?model=ULC-1.5FT-SMSM%2B.

[117] W. J. Riley, *Handbook of Frequency Stability Analysis*. National Institute of Standards and Technology (NIST), July 2008.

[118] "Stable32 website." http://www.wriley.com/.

[119] Matplotlib, *Matplotlib Python Library*. https://matplotlib.org/.

[120] "Agilent E4400B RF synthesiser." https://www.keysight.com/en/pd-1000002784%3Aepsg%3Apro-pn-E4400B/esg-a-series-analog-rf-signal-generator-1-ghz?cc=ZA&lc=eng.

[121] U. of Cape Town, "UCT website." http://www.uct.ac.za/.

[122] U. C. London, “UCL website.” https://www.ucl.ac.uk/.

[123] “Reutech Radar Systems Website.” https://www.reutechradar.com/, 2020.

[124] “Skywatcher AZ-EQ6 GT Pro.” http://skywatcher.com/product/az-eq6-synscan/, 2020.

[125] S. Paine, “Design and implementation of a dual polarised L-band parabolic dish antenna for NeXtRAD ,” Master’s thesis, University of Cape Town, 2016.

[126] P.-K. Cheng, “NeXtRAD antenna design: X-Band dual polarised conical horn antenna,” Master’s thesis, University of Cape Town, 2016.

[127] “Pentek Cobalt 71620.” https://www.pentek.com/products/detail.cfm?model=71620.

[128] D. D. Plessis, “Integration and testing of a digital transceiver for a dual frequency, pulse-doppler radar,” Master’s thesis, University of Cape Town, 2016.

[129] “Ubiquiti website.” https://www.ubnt.com/airmax/bulletm/.

[130] TimeMachines, “TimeMachines GPS NTP Server.” https://timemachinescorp.com/producttimeservers/, 2020.

[131] “Ansible website.” https://www.ansible.com/.

[132] D. Jordan, M. Inggs, and D. O’Hagan, “NeXtLook a lightweight, real-time quick-look processor for NeXtRAD,” in *International Conference on Radar Systems (Radar 2017)*, pp. 1–6, Oct 2017.

[133] “Nextlook Github Repository.” https://github.com/nextrad/nextlook.