

EISCAT_3D

The next generation European Incoherent Scatter radar system

FINAL DESIGN STUDY REPORT

DELIVERABLE D11.1



**EISCAT
Scientific Association**



Universitetet i Tromsø
University of Tromsø



Institutet för rymdfysik
Swedish Institute of Space Physics



Space Science & Technology
Rutherford Appleton Laboratory

EISCAT_3D
Work Package 11
Implementation Blueprint
Deliverable 11.1

Version 1.00, 05-Jun-2009

Editors:

Derek McKay (D.J.McKay@rl.ac.uk)
Ian McCrea (I.W.McCrea@rl.ac.uk)

Rutherford Appleton Laboratory
Chilton, Didcot, OX11 0QX
UNITED KINGDOM

On behalf of:

EISCAT_3D Design Study

The EISCAT Scientific Association,
PO Box 812,
SE-981 28 KIRUNA
Sweden

Tel: +46-980-79150
Fax: +46-980-79159

Document history (YYMMDD):

090114 – Original outline.
080224 – Implementation of the outline into the initial document draft
090508 – Re-merged all component sections into this master document
090509 – Editing of references and citations
090510 – Editing and cross-checking. Incorporating the final sections
090511 – Consistency checks, references, TLAs, title graphics, spelling & grammar, captions
090521 – Incorporated comments from internal reviewers
090521 – Added Executive Summary
090526 – Added Appendix with list of deliverables and editorial comments
090531 – Additional editorial work
090601 – Added comments from Steering Group
090602 – Additional editorial work
090603 – Inclusion of comments from internal reviewers
090604 – Editorial work, internal formatting, corrections and consistency checks
090605 – Public release

List of Participants

Tony van Eyken, EISCAT Scientific Association, Director (2003-2008)

Esa Turunen, EISCAT Scientific Association, Director (2009-)

Gudmund Wannberg, EISCAT Scientific Association / IRF Kiruna, Technical Project Leader

Henrik Andersson, EISCAT Scientific Association

Rico Behlke, University of Tromsø

Vasyl Belyey University of Tromsø

Peter Bergqvist, EISCAT Scientific Association

Johan Borg, Luleå University of Technology

Asgeir Brekke, University of Tromsø

Jerker Delsing, Luleå University of Technology

Marcus Edwall, Luleå University of Technology

Ivan Finch, Rutherford Appleton Laboratory

Anders Gabert, Luleå University of Technology

Tom Grydeland University of Tromsø (now at Discover Petroleum)

Martin Gustavsson, Luleå University of Technology

Björn Gustavsson, University of Tromsø

Ingemar Häggström, EISCAT Scientific Association

Toivo Iinatti, EISCAT Scientific Association

Jan Johansson, IRF Kiruna

Gustav Johansson, Luleå University of Technology

Jonny Johansson, Luleå University of Technology

Cesar La Hoz University of Tromsø

Tarmo Laakso, EISCAT Scientific Association

Ralf Larsen, EISCAT Scientific Association

Mikael Larsmark, Luleå University of Technology

Tore Lindgren, Luleå University of Technology

Jonas Lindström, Luleå University of Technology

Magnus Lundberg, Luleå University of Technology

Ulrica Lång, Luleå University of Technology

Jussi Markkanen, EISCAT Scientific Association

Inge Marttala, EISCAT Scientific Association

Ian McCrea, Rutherford Appleton Laboratory

Derek McKay, Rutherford Appleton Laboratory

Severin Oeschger, Luleå University of Technology

Markku Postila, EISCAT Scientific Association

Walter Puccio, IRF Uppsala

Toralf Renkowitz, University of Rostock / Institut für Atmosphärenphysik

Joel Ståbis, Luleå University of Technology

Krister Söderström, Luleå University of Technology

Lars-Göran Vanhainen, EISCAT Scientific Association

Assar Westmann, EISCAT Scientific Association

Ingemar Wolf, IRF Kiruna

Acknowledgements

The editors of this document wish to thank the following people for their contributions: Steve Crothers, Thomas Ulich, Markku Lehtinen, Anja Strømme, Ken Freeman and Izabela Bukowska.

The Technical Project Leader wishes to extend special thanks to Universität Rostock and the Institut für Atmosphärenphysik (IAP), Kühlungsborn, for their willingness to accept our proposal for a study of the EISCAT_3D array element antenna problem as the topic of Mr. Renkwitz's M.Sc. thesis, and for their kind and generous support of his work.

The members of the design study team would like to express their particular thanks to a number of people who have made it possible for us to reach this point. During the period of the study, EISCAT has had two directors, Prof. Tony van Eyken and Dr. Esa Turunen, both of whom have strongly supported this study. Prof. van Eyken did much to lay the groundwork for EISCAT_3D, including the development of the initial proposal which led to this FP6 funding. Dr. Turunen has enthusiastically picked up the baton from his predecessor, and supported the study through to completion, while at the same time beginning the planning for the next phase of the EISCAT_3D development.

Our thanks should also be expressed to those EISCAT staff members who were not directly involved in the design study, several of whom have been called upon to make extra efforts to keep the existing systems running while their colleagues were busy with the design exercise. We would like to thank each of our institutions for the support we have received from them, including their help with the monitoring and reporting of the financial aspects of this exercise. All of us owe a considerable debt of gratitude to Henrik Andersson, the EISCAT Head of Administration, who has ensured that the study kept on track from a financial and administrative point of view, and has very efficiently taken care of the many official requirements associated with reporting to the commission.

Staff from the University of Tromsø wish to thank Dr Jorge Chau, Director of the Jicamarca Radar Observatory, for his generosity in providing access to the Jicamarca Radar to conduct imaging experiments and tests and for very useful support and insights on the radar imaging technique. They also thank Dr Dave Hysell for no less very useful support and insights on the technique and for generously hosting one of them.

The EISCAT_3D team from Rutherford Appleton Laboratory also thanks the engineers, executives and sales representatives from the various companies with whom they have had considerable discussions. Thanks also go to other RAL staff: Steve Crothers, Dave Corney, Matthew Wild and Paul Gallop.

The images of the EISCAT_3D radar used on the title pages of this report were produced by Dr. Phil Moran and his colleagues at Ffab Productions (www.ffab.co.uk) and we wish to thank them for their work in support of this project.

EXECUTIVE SUMMARY



Executive Summary

- This report summarises the results from the four-year “EISCAT_3D Design Study”, a project supported by the Sixth Framework Programme of the European Union and running from May 1, 2005 to April 30, 2009.
- The aims of the Design Study have been essentially twofold, viz.
 - to study the feasibility of constructing a third-generation incoherent-scatter research radar, using cutting-edge technology throughout and providing an order-of-magnitude improvement in temporal and spatial resolution, to replace the existing, aging EISCAT VHF and UHF systems, and
 - to produce a detailed, costed design for such a system.
- The design study did not include a task to develop the science case; this was carried out as a parallel activity by EISCAT's Science Advisory Committee (SAC), later renamed as the Science Oversight Committee (SOC) from 1 January 2007. Members of the design study maintained close links with the science case development, and many of the ideas from SAC fed into the preparation of the EISCAT-3D Performance Specification Document (PSD).
- Following extensive consultation with the scientific user community in 2004 and 2005, it was determined that only *a multi-static phased array system* could reach or approach the performance demanded by present and future users.
- Accordingly, the target system specified in the EISCAT_3D Performance Specification Document (PSD) comprises a central active (transmit-receive) site (the “core”) and four receive-only sites, located on two approximately 250 km long baselines oriented N-S and E-W respectively.
- To achieve the desired performance, the proposed system design incorporates a number of innovative, ground-breaking concepts, e.g.
 - Direct-sampling receivers
 - Digital time-delay beam-forming
 - Multiple simultaneous beams from each receiving array
 - Adaptive polarisation matching and Faraday rotation compensation
 - Digital arbitrary-waveform transmitter exciter system
 - Full interferometry and imaging capabilities
 - Amplitude-domain data recording at full sampling rate
- During the four-year study, all mission-critical technical concepts have been modelled, investigated by simulations, in critical cases also by full-scale tests, and found to be realisable. Array sizes, transmitter power levels and receiver noise performance required to reach the desired time and space resolutions have also been established. Based on this work, we propose that the target system should have the following technical characteristics:

- The core will comprise 1) a 120-m diameter filled circular aperture array with ≈ 16000 elements, laid out on an equilateral triangular grid, and 2) a number (6...9) of smaller outlier receive-only arrays. It will provide:
 - a half-power beamwidth of $\approx 0.75^\circ$, i.e. comparable to that of the EISCAT UHF,
 - a power-aperture product exceeding 100 GW m^2 , i.e. an order of magnitude greater than that of the EISCAT VHF,
 - grating-lobe free pattern out to 40° zenith angle,
 - graceful degradation in case of single-point equipment failure.
- Each core array element will be made up from a radiator, a dual 300+300 watt linear RF power amplifier, a high performance direct-digitising receiver and support electronics. The recommended radiator is a crossed Yagi antenna with a minimum directivity of about 7 dBi.
- Two filled 8000-element receive-only arrays will be installed on each baseline at distances of respectively ≈ 110 and ≈ 250 km from the core. Their radiating elements will be 3- or 4-element X Yagis, essentially identical to those used in the core. The Yagis will be directed towards the core field-of-view and elevated to 45° . Outlier arrays for interferometry will also be installed.
- Advanced digital beam-forming systems will allow the generation of a large number of simultaneous beams from each array, thus eliminating the time/space ambiguity plaguing all present incoherent scatter systems and making true volumetric imaging of vector quantities possible for the first time.
- We have verified that a system meeting the performance requirements put forth in the EISCAT_3D Performance Specification Document could be built today, using existing technology, if cost were not an issue. Advances in semiconductor technology, signal processing and data storage between now and the time of placing a contract are expected to reduce component and subsystem costs to the point where a full-size core would cost about 60 MEUR and each receive-only site 20 MEUR. *We recommend that the member institutions of the EISCAT Scientific Association commit to funding and constructing such a radar system according to the results and guidelines given in our technical reports (EISCAT_3D Deliverables) within the next 5-7 years.*
- As designed, the system is highly modular and lends itself excellently to gradual expansion if funding should only be forthcoming in instalments. *In this case, we recommend that in a first phase, a 5000-element, 70-m diameter core array and at least two, 1500-element receiver sites should be constructed to replace the multi-static capabilities of the present UHF radar, which are going to be lost within the next twelve months.* This configuration would already exceed the performance of the current VHF system, providing a 1.3° half-power beamwidth, a power-aperture product of $\approx 10 \text{ GW m}^2$, and full beam steerability at the transmitter site. There are also stages beyond the outlined system involving multiple core elements arrays that could then be expanded as additional funding became available or as part of a more comprehensive initial build.
- The design study team notes with pleasure that, as a result of our efforts and the hard work of the EISCAT executives, EISCAT-3D was added to the European Strategy for Research Infrastructure (ESFRI) roadmap, at its last revision announced in December 2008.

This page is intentionally blank.

PART A

INTRODUCTION AND OVERVIEW



Table of contents

1	Introduction.....	6
1.1	Background	6
1.2	Scope of this document	7
1.3	Legal notices	8
1.4	Abbreviations and terms used	9
2	Original Objectives	14
2.1	Science drivers	14
2.2	Original objectives	19
3	Original Science Case	21
3.1	Evolutions of current EISCAT science	22
3.2	New EISCAT science.....	28
4	State of the Art in Scientific Radars	32
4.1	EISCAT	33
4.2	AMISR	35
4.3	Jicamarca.....	36
4.4	Sondrestrom	37
4.5	Arecibo	37
4.6	Millstone Hill	38
4.7	Kharkov	38
4.8	Irkutsk.....	39
4.9	MU and EAR.....	39
4.10	MST radars	40
4.11	Future developments.....	40
4.11.1	The Chinese ISR	41
4.11.2	An Antarctic ISR.....	41
4.11.3	PANSY	41
4.12	Summary.....	41
5	Design Challenges	43
5.1	Introduction	43
5.2	Transmitter system	44
5.3	The beam-steering problem.....	45
5.4	Array systems and element antennas	47
5.5	Adaptive Faraday rotation compensation.....	47
6	Methodology	50
7	System Overview	58
8	Principle Design Areas	61
8.1	Antenna design.....	61
8.1.1	Element antenna performance requirements.....	61
8.1.2	Environmental considerations.....	62
8.1.3	Antennas for the core array.....	63
8.1.4	Antennas for the Demonstrator Array.....	68
8.2	Array design	69
8.2.1	Introduction.....	69
8.2.2	Array design procedure.....	69
8.2.3	Core performance requirements and array parameters	71
8.2.4	Serviceability and survivability aspects.....	74
8.2.5	Constructional aspects	74
8.2.6	Pointing calibration and adaptive pointing corrections	76
8.2.7	The Demonstrator Array	77

8.3	Beam-forming concepts	79
8.3.1	Introduction.....	79
8.3.2	Beamforming layout	80
8.3.3	Digital time-delay beamforming.....	81
8.3.4	Finite impulse response filters	82
8.3.5	Filter design results	83
8.4	Time-keeping	85
8.4.1	Introduction.....	85
8.4.2	Cable calibration timing system.....	85
8.4.3	Simulation results.....	86
8.4.4	Measurement setup	87
8.4.5	Measurement results	88
8.4.6	GNSS timing system.....	88
8.4.6.1	GNSS timing system concept.....	88
8.4.6.2	GNSS simplified receiver.....	89
8.4.6.3	Test setup for concept evaluation	90
8.4.7	Results and conclusions	90
8.5	Receiver front-end.....	92
8.5.1	The low-noise amplifier	93
8.5.1.1	RF component characterisation	93
8.5.1.2	Temperature stabilisation	94
8.5.1.3	Directional coupler design.....	94
8.5.1.4	LNA design.....	94
8.5.1.5	LNA results.....	95
8.5.2	The A/D converter board	95
8.5.2.1	Sample clock.....	96
8.5.2.2	Anti-aliasing filters	97
8.5.2.3	Amplifier stages.....	97
8.5.2.4	ADC selection.....	97
8.5.2.5	A/D converter discussion.....	98
8.5.3	Calibration system implementation	98
8.6	Signal processing chain.....	99
8.6.1	Time domain beam-forming / beam-pointing.....	101
8.6.2	Band-limiting and down-sampling	106
8.6.3	Coherency processor / trigger for interferometry	107
8.6.4	Quick-look analysis	108
8.6.5	Polarisation matching.....	108
8.7	Polarisation estimation	109
8.7.1	Introduction.....	109
8.7.2	Polarisation estimation using eigenspace analysis.....	111
8.7.3	Direct approach for polarisation estimation.....	112
8.7.4	Subspace tracking	112
8.7.5	Evaluation	113
8.7.6	Spectrum estimation.....	115
8.7.7	Discussion and conclusions	117
8.8	Interferometry and radar imaging	118
8.8.1	Fundamentals of interferometric imaging.....	118
8.8.2	Single-baseline interferometry.....	120
8.8.3	Array configurations	120
8.8.4	Phase calibration	122
8.8.5	Image inversion for interferometry	124

8.8.6	Multi-dimensional radar imaging	124
8.9	Data system	127
8.9.1	Introduction.....	127
8.9.2	Data rate challenges	128
8.9.2.1	Incoherent scatter data	128
8.9.2.2	Interferometry data	129
8.9.2.3	Data from supporting instruments	130
8.9.3	System requirements	131
8.9.4	Hardware specification	132
8.9.5	Suggested solutions.....	133
8.9.5.1	Temporary ring buffers.....	133
8.9.5.2	The permanent archive	134
8.9.5.3	Network issues.....	134
8.9.6	Summary	134
9	Additional Design Areas.....	135
9.1	Transmitter	135
9.2	Internal data communication.....	139
9.3	Control and monitoring	142
9.3.1	Introduction.....	142
9.3.2	Benefits of an EROS-type system.....	143
9.3.3	User support	143
9.3.4	Maintenance and development	143
9.3.5	Unified interface	144
9.3.6	Assessment.....	144
9.3.7	Future development	145
9.4	Visualisation and data analysis	145
9.4.1	Technical visualisation.....	145
9.4.2	Scientific visualisation.....	147
9.5	Site infrastructure requirements	148
9.5.1	Introduction.....	148
9.5.2	Summary of the site survey 2005.....	149
9.5.3	Infrastructure.....	150
9.5.3.1	Data communications	150
9.5.3.2	Power supply	150
9.5.4	RF environment	151
9.6	Novel science applications	151
9.6.1	Introduction.....	151
9.6.2	Identification of novel uses.....	151
9.6.3	Improving ionospheric correction models	152
9.6.4	Use of long time series data.....	153
9.6.5	Ionospheric signatures of global change.....	154
10	The Demonstrator Array	156
10.1	Introduction	156
10.2	Technical overview of the Demonstrator	156
10.2.1	Antenna array – details	157
10.2.2	Equipment cabinets.....	158
10.2.3	Interface block	160
10.3	Receiver structure	161
10.3.1	Front-end subsystem	163
10.3.2	Digital down-converter	163
10.3.3	Data communications link	164

10.3.4	Interface to the EISCAT UHF digital back-end	165
10.4	Signal flow through the Demonstrator	167
10.5	Demonstrator-specific software.....	168
10.5.1	Rack-software	168
10.5.2	DDC-software	169
10.5.3	Radar controller	169
10.6	Installation period 2006–2008	169
10.6.1	List of the main milestones:.....	170
10.7	First Demonstrator incoherent scatter data.....	171
10.8	Reception of test signals	173
10.9	Environmental effects on the Demonstrator Array.....	174
10.9.1	Measurements	174
10.9.2	Proposed measurement system	175
11	Design Blueprints for Key System Elements.....	180
12	List of Potential Suppliers.....	192
13	Costs and Timescales for System Construction.....	196
13.1	Estimated costs of EISCAT_3D	196
13.2	Timescales for EISCAT_3D.....	198
14	Summary and Conclusions	202
15	References.....	204
16	Appendix: Complete List of Reports & Deliverables.....	212
17	Appendix: Publications in open literature.....	221
17.1	Journal publications.....	221
17.2	Conference publications	221
17.3	Doctoral theses	221
17.4	Licentiate thesis	222
17.5	M.Sc. theses.....	222
17.6	EISCAT workshops and user meetings.....	222

1 Introduction

This document is the only deliverable from Work Package 11 of the EISCAT_3D design study (“Implementation Blueprint”) and as such it comprises our final technical report, summarising the findings of the entire four-year design study. In this report, we present an overall design for the new EISCAT_3D radar system, integrating the results of all the work packages and the work of all the project partners. We show in detail how it is possible to construct a next-generation, multistatic phased array radar system, incorporating a large number of technical innovations which have been researched during the course of our study. By building this radar system in northern Scandinavia, European scientists will be able to realise a facility which restores Europe’s lead in incoherent scatter radar technology, and provides an assured long-term upper atmospheric science capability to replace the existing EISCAT UHF and VHF radar systems in Norway, Sweden and Finland.

1.1 Background

The current EISCAT facilities used to study the high-latitude upper atmosphere and its connection to geospace, are developed and operated by the EISCAT Scientific Association, a multi-national organisation funded by the relevant national research councils of China, Finland, Germany, Japan, Norway, Sweden and the UK, and also supported by financial contributions from France, Russia and Ukraine. The EISCAT Scientific Association was formed in 1975, when the construction of the Mainland UHF radars began. The first observations on the UHF system began in 1981, and the VHF radar first became available in 1984. In 1996, the 500 MHz radar facility on Svalbard (the EISCAT Svalbard Radar) was added to EISCAT’s suite of instruments. Further details of the existing EISCAT radars can be found in Section 4.1 of this report. For many years, the EISCAT radars have been world leaders in terms of their experimental capability, and this lead has been maintained by constant innovations within the user community which has continuously developed new experiments and data analysis techniques, aided by a major upgrade of the mainland receiver hardware in 2001.

Over the last few years, a number of trends have appeared which have inevitably made it necessary for us to renew our current hardware. Most importantly, EISCAT science has moved into new areas, including more detailed studies of the energy coupling between the upper and lower atmosphere, the linkages between the ionosphere and magnetosphere, investigations of the importance of turbulence and small-scale structure, and sensitive detection of weak-coherence targets such as micrometeoroids, cm-scale space debris, as well as into planetary radar applications. These studies are not purely scientific, but also have practical importance for applications such as global positioning, communications and space situational awareness. In turn this has presented a need for observations beyond the capabilities of the present systems, requiring capabilities such as greater sensitivity, faster scanning, multiple beams, volumetric imaging and interferometry, which no incoherent scatter radar is currently capable of providing. At the same time, EISCAT’s existing radars are becoming older and harder to maintain, and the frequency bands which have been used for EISCAT operations in the past are coming under increasing pressure from UMTS 900 mobile telephones (in the case of our UHF

frequencies) and digital audio broadcasting (in the case of our VHF frequencies). Another factor has been the recent inception of new incoherent scatter radars using phased array technology, such as the AMISR systems in the United States. While these systems are admirable in many respects, the EISCAT community collectively felt that it was possible to go considerably further in developing phased array radar which was really unique, combining high sensitivity, volumetric imaging, interferometry and multistatic observations.

This is the background to the proposal of the EISCAT_3D design study, an attempt to specify a new generation of incoherent scatter facilities for the 21st century to provide European scientists with secure access to the world's leading research instrument for upper atmospheric science. As we will show in the following pages, we have convincingly demonstrated that our ideas are achievable in practice, not just by producing a paper design for a new radar system, but also by testing the proposed technical solutions at prototype level, which has allowed us to verify that many of our design ideas work in practice.

1.2 Scope of this document

This report is organised into seven Parts (A to G). In Part A, we briefly review the original objectives and context for the design study, specifying in greater detail the scientific rationale for developing a new radar and the implications that these new science goals have for the performance of such a facility. We summarise the current state of the art by looking at the existing incoherent scatter radar facilities, finding that none are fully capable of meeting our scientific objectives. We describe the factors which have made the EISCAT_3D design particularly challenging from a technical viewpoint, and summarise the methodology which we have used in carrying out the study, including a survey of the user community to determine design requirements, not all of which could be fully realised because some were impossible or mutually exclusive. In these cases, Part A describes the compromises which were necessary to accommodate as many as possible of the desired capabilities.

In Part B, we review in detail the individual design areas which have been needed in order to realise the EISCAT_3D concept. We divide these into two categories – critical and ancillary design issues. This reflects an evolution which has gone on during the course of the study. At the time of our proposal to the EU (spring 2004), all the design areas in the study were given a roughly equal weighting. Once we began the detailed work, however, it quickly became apparent that certain aspects were particularly challenging and required comparatively more effort. These are the critical design areas which are reported in Sections 8.1 to 8.9, and the descriptions of these make up the bulk of this report. On the other hand, it was realised that some design issues were not as challenging as had originally been thought, for example because the necessary specification could be achieved by modifying a system used in the existing EISCAT radars. Comparatively less effort was therefore spent on these areas, which are described in Sections 9.1 to 9.6, although they should still be regarded as essential elements of the overall EISCAT_3D design.

In Part C, we have compiled the detailed blueprints for all areas of the study where such details exist. In general these correspond to the critical design issues, a highly-

specified design was essential to be certain that the required capability could be achieved. It was agreed early in the study that such an assurance required more than a paper design. Accordingly, the EISCAT Scientific Association has used its own resources to implement a “demonstrator array” – effectively a small working prototype of a passive EISCAT_3D array, which we have used to test critical technology aspects such as our antenna design, array geometry, beam-forming, time-keeping and signal processing. The construction and performance of the demonstrator array are summarised in Part C (Section 10). A sample of the design blueprints for the Demonstrator Array can be found in Part D.

Although we do not foresee that the EISCAT_3D system will begin to be constructed for at least two years, we have carried out the design study with the aim of demonstrating that the system can be achieved using technology which already exists. An important consideration in moving beyond the design study will be the ability to demonstrate that there are no technical obstacles, and we have shown that the system could be constructed immediately if factors such as funding and site selection were not an issue. In many cases, therefore, we have identified actual components and suppliers capable of producing key EISCAT_3D sub-systems, and these are listed in Part E. In other cases, however, novel hardware has been designed and assembled in-house, and a necessity before any move to construction is that ways should be found of fabricating such elements, which because of the distributed nature of the EISCAT_3D system will often need to be produced in large quantities. These issues are considered in Part F, where we review the factors relevant to the costs and timescales of system construction. In that part we also discuss briefly some of the political and funding developments which have occurred during the course of the study, and the implications which these have for the implementation of EISCAT_3D. In Part G, we end with a short summary and some conclusions.

1.3 Legal notices

This document has been prepared as part of the EISCAT_3D design study. Although it is publicly available, the views and opinions expressed within it are only pertinent in the context of the EISCAT_3D project and have no implications for any other project. In particular, comments made on commercially available products and services are neither endorsements nor condemnations. No part of any original material in this document should be used for any purpose without acknowledgement of the source. Requests for use of such material should be addressed to the EISCAT Scientific Association (as given in the front matter of this report).

Wherever possible, the authors have sought to use original material or material with no copyright restrictions. Where this has not been the case, permission has been obtained either verbally or in writing. Copyright in some documents and other material used herein belongs to third parties; please check copyright notices in respect of such material. References are given in a separate section at the end of the document.

1.4 Abbreviations and terms used

The following abbreviations and terms are used in this document:

A/D	Analogue to Digital
AB	Class AB, circuit design
ADC	Analogue to Digital Converter
AEU	Antenna Element Unit
AKR	Auroral Kilometric Radiation
AMISR	Advanced Modular Incoherent Scatter Radar
ARMA	Autoregressive Moving-Average (process)
ASIC	Application Specific Integrated Circuit
ATRAD	Atmospheric Radar Systems Pty Ltd, Australia
AUT	Antenna Under Test
B	Byte(s)
BM	Buffer Memory
BW	Bandwidth
Cas-A	An astronomical radio source
CCLRC	Council of the Central Laboratory of the Research Councils (UK)
C&M	Control and Monitoring
CIR	Corotating Interaction Region
CME	Coronal Mass Ejection
CMOS	Complementary Metal Oxide Semiconductor
CPU	Central Processing Unit
CRIRP	Chinese Research Institute of Radio Propagation
CW	Continuous Wave
Cyg-A	An astronomical radio source
DAB	Digital Audio Broadcasting
DAC	Digital to Analogue Conversion
DC	Direct Current
DDC	Digital Down-Conversion
DDS	Digital Direct Synthesis
DP/A	Delay-Phase shift and Add
DSP	Digital Signal Processing
E3D	EISCAT_3D
EAR	Equatorial Atmospheric Radar, Indonesia
EASI	ESR Aperture Synthesis Imaging (Interferometry system on Svalbard)
EI	EISCAT
EISCAT	European Incoherent SCATter (the scientific organisation)
EISCAT_3D	(Also EISCAT_3D), Next generation EISCAT radar project
EISLAB	Embedded Internet Systems Laboratory, LTU
ELAN	Experiment Language
ELF	Extremely Low Frequency (3 to 30 Hz)
EM	Electromagnetic
EROS	EISCAT Real-time Operating System
ESA	European Space Agency
ESR	EISCAT Svalbard Radar
ESRANGE	European Space Range, rocket range, Kriuna, Sweden
ESSA	Environmental Science Service Administration
EU	European Union

EUR	Euro (currency)
FET	Field Effect Transistor
FIR	Finite Impulse Response
FOV	Field of View
FP6	Framework Programme (EU funding programme)
FP7	Framework Programme (EU funding programme)
FP8	Framework Programme (EU funding programme)
FPGA	Field Programmable Gate Array
FSD	Fractional Sample Delay
FWHM	Full-Width Half-Maximum
GB	Gigabyte
GHz	Gigahertz
GNSS	Global Navigation Satellite System
GPS	Global Positioning System
GSM	<i>Groupe Spécial Mobile</i> , also Global System for Mobile communications
GUISDAP	Grand Unified Incoherent Scatter Data Analysis Program
HEMT	High-Electron Mobility Transistor (solid-state process)
HF	High Frequency
IAP	Institut für Atmosphärenphysik, Germany
IPS	Inter-Planetary Scintillation
IRF	Institutet för Rymdfysik (Swedish Institute of Space Physics)
IS	Incoherent Scatter
ISR	Incoherent Scatter Radar
IT	Information Technology
JASMET	Jicamarca All-sky Specular MEteor Radar
JULIA	Jicamarca Unattended Long-term Investigations of the Ionosphere and Atmosphere
kB	Kilobyte
kEUR	Kilo-Euro (currency abbreviation)
KuaFu	Proposed magnetospheric mission
LAAR	Large Antenna Aperture Radar
LNA	Low-Noise Amplifier
LOFAR	Low-Frequency Array (Dutch/EU radio telescope)
LSI	IT company
LTU	Luleå Tekniska Universitet (Luleå Technical University, Sweden)
LVDS	Low-voltage differential signalling
MAID	Massive Array of Idle Disks
MB	Megabyte
MEM	Maximum Entropy Method
MERLIN	Multi-Element Radio-Linked Interferometer (UK radio telescope)
MHz	Megahertz
ML	Maximum Likelihood
MPI	Max-Planck Institut, Germany
MSc.	Master of Science, university degree
MST	Mesosphere Stratosphere Troposphere
MU	Middle and Upper atmosphere radar, Japan
MVUE	Minimum Variance Unbiased Estimator
MW	Megawatt
NBS	National Bureau of Standards (USA)

NCO	Numerically Controlled Oscillator
NEIAL	Naturally Enhanced Ion Acoustic Lines
NOAA	National Oceanic and Atmospheric Administration (USA)
N-S	North-South
NXP	Semiconductor manufacturer (formerly part of Philips Electronics)
OCF	A UK-based IT company
OM2	Multimode (optical) Fibre specification, ISO 11801 standard
PANSY	Program of the Antarctic Syowa MST/IS radar
PB	Petabyte
PC	Personal Computer
PCB	Printed Circuit Board
PFISR	Poker Flat Incoherent Scatter Radar
PLL	Phase-Locked Loop
PMSE	Polar Mesospheric Summer Echoes
PMWE	Polar Mesospheric Winter Echoes
POSIX	Portable Operating System Interface (standardisation)
PPARC	Partical Physics and Astronomy Research Council (UK, now STFC)
PRS10	Time standard
PSD	Performance (also Primary) Specification Document
PSU	Power Supply Unit
PT100	Model of thermal sensor
PhD.	Doctor of Philosophy, university degree
Q1	First quarter (January to March)
Q2	Second quarter (April to June)
Q3	Third quarter (July to September)
Q4	Fourth quarter (October to December)
RAL	Rutherford Appleton Laboratory (STFC, UK)
RC	Resistor-Capacitor
REX	Receiver signal processing card
RF	Radio Frequency (in the E3D context, it is typically 200-250 MHz)
RFI	Radio Frequency Interference
RFMD	RF Micro Devices, company
RFQ	Request For Quotations
RT	Real-Time
RTI	Range Time Intensity
S&H	Sample & Hold
SAC	Scientific Advisory Committee
SAR	Synthetic Aperture Radar
SAS	Serial Attached SCSI (computer disk drive technology)
SCSI	Small Computer System Interface (peripheral bus technology)
SEK	Swedish Kroner (currency)
SERDES	Serialiser/De-serialiser
SETI	Search for Extraterrestrial Intelligence
SFDR	Spurious-Free Dynamic Range
SG	Sub-array designation
SIOS	Svalbard Integrated Observing System
SMHI	<i>Sveriges Meteorologiska och Hydrologiska Institut</i> , Swedish Meteorological and Hydrological Institute
SMP	Shared Memory Processor
SNR	Signal to Noise Ratio

SOC	Science Oversight Committee (formerly SAC)
SOUSY	SOUnding System (instrument)
SPEAR	Space Plasma Exploration by Active Radar (instrument)
SRI	Stanford Research International
STEREO	Two-spacecraft solar wind mission
SSTD	Space Science & Technology Department (RAL, UK)
STFC	Science and Technology Facilities Council (replacing PPARC and CCLRC, UK)
TARLAN	Transmitter and Receiver Language
TB	Terabyte
TEC	Total Electron Content
Te/Ti	Ratio of the electron to ion temperatures
THEMIS	Time History of Events and Macroscale Interactions during Substorms (Spacecraft)
TLK	Electronic component series, produced by Texas Instruments
TPL	Technical Project Leader
TTD	True Time Delay
TV	Television
UHF	Ultra High Frequency (300 MHz to 3 GHz)
UiT	Universitetet i Tromsø, (University of Tromsø, Norway)
UMTS	Universal Mobile Telecommunications System
UNIX	An operating system
UK	United Kingdom
URSI	<i>Union Radio-Scientifique Internationale</i> (International Union of Radio Science)
US	USA
USA	United States of America
VAT	Value Added Tax (UK)
VCO	Voltage Controlled Oscillator
VCXO	Voltage Controlled crystal Oscillator
VHF	Very High Frequency (30 to 300 MHz)
VLBI	Very Long Baseline Interferometry
VLF	Very Low Frequency (3 to 30 kHz)
VME	VERSAmodule Eurocard bus, computer bus standard
VSWR	Voltage Standing Wave Ratio
WP	Work Package (specifically, one of the areas of the EISCAT_3D design study)
WP1	Work Package 1 Management of Design Study and Spectrum Allocation, EISCAT
WP2	Work Package 2 Evaluation of Design Performance Goals, EISCAT
WP3	Work Package 3 Evaluation of the Options for the Active Element, EISCAT, IRF
WP4	Work Package 4 Phased Array Receivers, LTU
WP5	Work Package 5 Interferometry, UiT
WP6	Work Package 6 Active Element, IRF
WP7	Work Package 7 Distributed Control and Monitoring, EISCAT
WP8	Work Package 8 Data Archiving and Distribution, RAL / CCLRC
WP9	Work Package 9 Signal Processing, LTU, EISCAT & IRF
WP10	Work Package 10 New Ways of Exploiting Incoherent Scatter Radars, UiT

WP11	Work Package 11 Implementation Blueprint, EISCAT & RAL
WP12	Work Package 12 Networking and Reference Time and Frequency, EISCAT
WP13	Work Package 13 Enabling actions (Spectrum etc.), EISCAT & IRF
WPs	Work Packages
WWW	World-Wide Web

Please note that references to external reports, books and papers are specified using square brackets, [], and are listed in full in the References Section 15, starting on page 204.

2 Original Objectives

The EISCAT_3D radar system is envisaged as providing a significant forward step in the design and operation of incoherent scatter radars for studies of the atmosphere and near-Earth space. The design exercise has therefore required that the radar should incorporate a number of new concepts in experiment design and data handling, the need for which has become evident during more than 20 years of operation with the existing EISCAT systems, but which have not been possible to implement given the constraints of existing EISCAT hardware. In addition, the award of EU funds for the system design stipulated that we should design a system whose capabilities go beyond those of the current generation of US incoherent scatter radars, to restore Europe's technological lead in this area. While the new US AMISR (Advanced Modular Incoherent Scatter Radar) systems are undoubtedly highly capable, their design and initial implementation have not included many of the concepts proposed here, including the ability for routine interferometric operation and for the temporary storage and advanced re-processing of the lowest level data products. As such, we believe that the EISCAT_3D concept represents a significant advance on existing incoherent scatter radar designs; however, it also implies significant challenges, especially for the data handling aspects of the system given the very high volumes of data which need to be dealt with and the different (and dynamically configurable) timescales on which different data products need to be processed and stored.

An important factor in designing a new radar system, especially in fast-developing research areas such as atmospheric and space physics, is to “future-proof” the system as much as possible. This means that the system needs to be designed with sufficient capacity and flexibility to handle experimental concepts which have not yet been fully developed, such as the possibility for multi-phase (e.g. quad phase) codes and the potential to fully exploit the possibilities for multiple re-processing of the voltage-level data. This places significant demands on various areas of the design study, among them the design of the data handling system, since much of the novelty in the EISCAT_3D system design relates to methods for processing and re-processing the data.

2.1 Science drivers

The facilities of the EISCAT Scientific Association have defined the state of the art in incoherent scatter ionospheric radars at least since the commissioning of the EISCAT Svalbard Radar (ESR) in 1996. However, recognising that the capabilities of the existing EISCAT VHF and UHF mainland systems were being pushed to their limits by ever-increasing scientific demands, in June 2003 a questionnaire prepared by the SAC chairman and the EISCAT director was distributed to the SAC and Council members of all the existing EISCAT countries. The questionnaire asked EISCAT users to specify their current scientific priorities and to predict what the major research issues would be during the ten-year period of the next EISCAT agreement (2007-2016). Users were also asked to suggest what new facilities the Association would need to provide in order to meet their scientific requirements, as well as to rank the existing facilities in priority order.

Many of the EISCAT countries had held meetings of their user communities during the spring and summer of 2003 to discuss the future of EISCAT. The suggestions of these meetings were incorporated into the responses returned by the national SAC representatives.

When the responses were analysed, a vast range of new science opportunities for the ISR method was indicated, provided that EISCAT would invest in developing the technology. Some specific areas of research were repeatedly highlighted as being of fundamental importance but as yet insufficiently understood and meriting continued investigation:

- A better understanding of the interface between the ionosphere and middle atmosphere and the nature and effectiveness of the coupling mechanisms between these regions is required. Major questions remain concerning the processes which mediate energy transfer in the mesosphere and lower thermosphere, and the extent to which these are important in governing the behaviour of the lower atmosphere. One key area was likely to be the study of the impact which energetic solar and magnetospheric particles have on the chemistry of the mesosphere and upper stratosphere (e.g. effects on odd nitrogen, odd hydrogen, negative ions and ozone abundance) using data from multiple diagnostics in conjunction with detailed chemical models, such as the Sodankylä Ion Chemistry Model. The effect of auroral precipitation on the chemistry at higher altitudes (e.g. via the production of odd nitrogen) also needed much further study.
- Many of the unresolved questions in auroral physics, relating to fundamental processes in plasma physics, are still open. EISCAT had recently provided dramatic demonstrations of the importance of small-scale structure in auroral arcs, and of the processes such as plasma wave coupling occurring in connection with these. It had become increasingly clear that a full understanding of current flow in auroral structures would require an understanding of “black aurora”, the regions of no emission which may contain large electric fields and/or substantial field-aligned currents; because of the large energy densities and small spatial scales involved, these highly structured regions were expected to be rich in plasma physics phenomena. However, having scale sizes much smaller than the radar beam width, they would remain unresolved unless the emerging new technique of radar interferometry, pioneered at EISCAT, could be applied.
- Continued studies of the electrodynamic coupling between the ionised and neutral atmospheres at ionospheric heights are vitally needed. New results suggested that meso-scale structure in the neutral atmosphere might play a much greater role in determining the overall amount of energy dissipation than had previously been acknowledged. In order to fully model how the energy of the solar wind is ultimately dissipated in the upper atmosphere, a better understanding of the action of coupling processes such as Joule heating and Lorentz forcing is needed. A prerequisite for real progress in this area would be the availability of high time resolution full 3-D convection electric field data.

- The study of phenomena related to “dusty plasmas” and the Polar Mesospheric Summer Echoes (PMSE) phenomenon was expected to remain a major theme of EISCAT research. PMSE is of great interest because it is a sensitive indicator of conditions at the mesopause, and can also be used in examining the dynamics of this important region. Since the mid-1980s, EISCAT had been at the forefront of experiments to establish the nature and behaviour of PMSE. Innovative experiments using the Tromsø heater to modulate the intensity of PMSE layers had provided new theoretical understanding of the microphysical processes responsible for forming these layers. There was also growing interest in the existence and significance of Polar Mesosphere Winter Echoes (PMWE), which may not only be observable at EISCAT VHF frequencies, but also studied by heating experiments, similarly to active PMSE studies.

A wide diversity of other science topics were also mentioned, including several new fields of research that in some cases had grown out of discoveries made at EISCAT, such as:

- transient coherent echoes from compact targets in the auroral ionosphere
- extremely narrow natural layers of ionisation in the E region
- radar determination of meteor orbital parameters
- studies of micrometeor UHF head echoes and meteor-atmosphere interactions

Further investigations into all these fields are currently restricted by the limited performance of the existing UHF and VHF radars in the 50-150 km altitude range. This is particularly true of the new research topics, where the processes being studied are transient and spatially compact, so much so that improvements in sensitivity and temporal and spatial resolutions by a full order of magnitude are necessary in order to progress the understanding.

A number of compelling scientific reasons had also been advanced why EISCAT should move to continuous operations, or at least to a much higher level of operation than undertaken at the time. Ionosphere-atmosphere coupling research would benefit greatly from the availability of long, uninterrupted data series, which would make it possible to study seasonal and solar cycle changes in the mesosphere and lower ionosphere, to monitor long-period phenomena such as planetary waves and tides and to develop climatologies of these and other wind and wave modes. Auroral physicists would enjoy a much increased likelihood of finding intervals with an optimum combination of geomagnetic activity level and the appearance of interesting types of auroral structure (e.g. discrete arcs or active rays, as opposed to diffuse precipitation). The possibility of continuous operations would also make it much easier for EISCAT to provide observational support for satellites, rocket campaigns and other types of diagnostic instruments and allow the EISCAT community to play a full role in the ground-based support of future satellite missions.

Through its response to the questionnaire, the EISCAT user community had thus outlined a very strong scientific programme for the future, addressing many important unresolved issues while building on EISCAT’s traditional areas of strength and expertise. At the same time, however, it was recognised that the existing mainland EISCAT VHF and UHF radars would not be capable of delivering the performance

demanded by this ambitious program, nor could they be relied on to serve as the platforms on which to construct a new system.

Designed in the late 1970s, the VHF and UHF were never planned for continuous operations and would stand up poorly to attempts in this direction. While the tri-static UHF system had been recently upgraded with two new klystrons and was otherwise in good technical shape, it was severely Debye-length limited below 80 km and above 800 km as a consequence of the high operating frequency (928 MHz) and therefore fundamentally ill-suited to mesosphere and ionosphere topside studies. It was also under constant threat of losing its operating spectrum. Back in the 1970s, the Norwegian frequency administration had awarded EISCAT an allocation at 933.5 MHz. This frequency was already at that time internationally assigned to the so-called “fixed-to-mobile” radio service (later used for mobile phones) on a first priority basis, but still unused. However, the explosive development of mobile telephone systems during the 1990s had now led to an accelerating demand for spectrum space in the 900 MHz range. This was expected to lead to a situation where at least the Finnish and Swedish receiving sites would be forced off the air at some point during the period of the next agreement, thus effectively terminating EISCAT's unique vector velocity measurement capability.

The VHF system frequency, 224 MHz, had long ago been selected to allow successful radar operation at the very low electron densities encountered in the mesosphere and ionospheric topside without running into Debye cut-off problems. This part of the VHF spectrum was still used for terrestrial analogue TV broadcasting in the three Nordic host countries, but the TV transmitters were at the point of being closed down as part of the transition to all-digital UHF TV. While digital audio broadcasting (DAB) would certainly claim part of the vacated spectrum, it was expected that the authorities would also be favourably inclined towards creating a protected allocation for EISCAT use in this band, thus helping to safeguard the long-term future of the system.

However, the VHF system had been found to be underpowered in practice for serious work at really low electron densities. It was recognised also that the VHF system might have a very limited future, in view of the well-known pointing limitations of the VHF antenna, and the fact that the transmitter klystrons had been built to a custom design not used by any other facility, were getting fairly old, and were unlikely to be repairable in the event of failure.

Against this background, it was clear that an entirely new radar system was needed if EISCAT wanted to maintain its position at the forefront of incoherent-scatter-radar based research. In order to represent a scientifically worthwhile and future-proof investment, this new system should be substantially more powerful, sensitive and flexible than the existing mainland EISCAT VHF and UHF systems, while at the same time retaining and extending the unique, extremely powerful multi-static capability of the latter. Interferometric capabilities should be designed-in from the beginning in order to fully exploit the world-leading position which EISCAT had acquired in the area of radar interferometry. An operating frequency in the high VHF band (225-240 MHz) would ensure optimum performance in low electron density conditions (i.e. both in the middle atmosphere and in the topside ionosphere). Exactly how such a new radar system should be realised in practice would of course become

the subject for a comprehensive feasibility study, but phased-array technology was seen as offering a number of appealing properties.

Phased arrays are inherently modular. This removes any dependence on a single reflector antenna structure and allows the system to be expanded incrementally as resources become available. Thanks to the steady progress in high power semiconductor RF technology, it was now technically and financially feasible to generate the required multi-megawatt RF power level by fitting a 500-watt class transmitter module to each array element. The modularity of such a distributed system allows a graceful degradation in the event of a module failing and is thus ideally suited to continuous operations. Employing electronic beam-forming and beam-steering, phased array systems are not only capable of steerability comparable with dish systems over a wide field of view, but also of pulse-to-pulse beam steering in arbitrary directions for applications such as large-scale imaging. Provided that linear power amplifiers are used, aperture tapering (e.g. as required for sidelobe suppression) can be introduced almost trivially.

Phased array technology had been chosen by the US incoherent scatter radar community to be the basis of its new AMISR system at Poker Flat, Alaska, USA, and the Polar Cap Observatory radar, approved for funding and deployment at Resolute Bay, Canada. This radar would have technical abilities beyond those which could be provided by any of the existing EISCAT radars, either in their present form or through reasonable upgrades. Although the Resolute Bay radar would be operating at about 440 MHz, the possibility would certainly exist for some of the techniques used in its design to be applied also at ≈ 235 MHz. There might also be a possibility to share some technology with the Dutch LOFAR very large phased-array radio-astronomy project.

It was recognised that a future EISCAT system could capitalise on many phased-array technology features. For instance:

- By applying a suitable degree of aperture tapering, a VHF phased array at the Tromsø/Ramfjordmoen site would offer the possibility of making field-aligned observations in the topside ionosphere, which had previously been essentially impossible to achieve,
- At the receive-only (“remote”) sites, the introduction of phased-array technology would make it possible to form a large number of simultaneous beams. By directing these to intersect the transmitter beam at many points separated by about a scale-height, full-vector ionospheric drift velocities over a range of altitudes and the true instantaneous 3-d convection electric field would become accessible to study for the first time, as desired by e.g. scientists studying ion-neutral coupling,
- It would be possible to use subsets of all arrays as interferometer elements in addition to their regular role as parts of the main arrays.

2.2 Original objectives

As the deliberations of the governing bodies of the Association had reached this point, the call for FP6 support for Design Studies was issued by the European Commission. This prompted the EISCAT Executive to submit the application that eventually led to the FP6 EISCAT_3D project.

To quote from the proposal, the ambitious global objective of the Design Study was to produce a design for “a facility offering a level of performance surpassing that of all other facilities, both existing and under construction” and “meeting or exceeding the requirements of European scientists and data consumers over the next cycle of solar activity”. It was stated that the study would “investigate the technical feasibility of a new, next-generation VHF incoherent scatter radar with distributed power amplifiers and an upgraded antenna array for both transmission and reception, together with at least two further, remote reception facilities, using phased arrays with multiple distributed receivers”. In the first approximation, deployment at the existing mainland sites should be considered, but serious consideration would also be given to the possibility to deploy arrays at other locations such as the ESRANGE and Andøya rocket launch sites, as well as to the extent to which interferometry, MST capabilities and fully remote operation could be factored into the design. The new facility would provide high-quality ionospheric and atmospheric parameters on an essentially continuous basis for academic researchers and practical service consumers, as well as near-instantaneous response capabilities for scientists and users studying unusual and unpredicted disturbances and phenomena in the high-latitude ionosphere and atmosphere. A much greater variety of data products than that available from the existing systems would be generated, in particular time-resolved complex amplitude data.

The principal components of the envisaged system were listed as:

- A powerful VHF (225 – 240 MHz) transmitter system
- At least three receiver systems, of which two or more would use remote phased arrays capable of simultaneously receiving scattered signals from ionospheric volumes along the radar beam with sample separations better than the scale-height of the ionosphere (i.e. the altitude interval over which the ionosphere behaves as if constant) at each altitude,
- Two or more interferometry receiver systems located close to the transmitter site to support measurements of spatially restricted phenomena at resolutions less than the dimensions of the transmitted beam
- Appropriate networks and data distribution systems to allow both remote control of the facility and timely and effective distribution of both raw and processed data and results

Some of the explicitly stated performance goals were:

- Improvements in the temporal and spatial resolution (both parallel and perpendicular to the radar line-of-sight) by about an order of magnitude relative to the present EISCAT systems,
- The extension of full vector velocity measurements of from a single point to the entire altitude range of the radar, and
- An increase of the operational time by a factor of at least four (from 12 to 50%), possibly even to full-time.

In order to achieve these goals, the design study would need to encompass essential developments in areas such as e.g.

- Improved radar pulse compression techniques
- Advanced signal processing,
- Electronic beam forming,
- Adaptive polarisation control,
- Data collection, distribution, and analysis, and (possibly)
- Design and production of application-specific subsystems on silicon.

The principal deliverable resulting from the Design Study was expected to be “...a fully developed and costed design for the next generation ionospheric radar...”, including among other things:

- A construction project and time plan,
- The identification of suitable components, systems, and (European) vendors, suppliers, and developers, and
- The identification of funding options and profiles.

3 Original Science Case

As stated in the Executive Summary, the development of the science case for EISCAT-3D was not a design study responsibility, but was handled by the EISCAT Science Advisory Committee (SAC), which became the Science Oversight Committee (SOC) as from January 2007. Although it is not therefore strictly necessary for us to include a summary of the science case in the final report, we have chosen to do so in the interest of completeness, since many of the ideas discussed by SAC and SOC have fed into parts of this study. In particular we concentrate on the science case ideas as they stood in 2005, since these fed directly into the Performance Specification Document (see Section 6). Development of the science case has since continued steadily, most recently in the form of an international EISCAT-3D community meeting, which took place in May 2009, a summary of which can be found at www.eiscat.se/groups/EISCAT_3D_info/uppsala-may09/.

The Scientific Advisory Committee of EISCAT (SAC) first considered the science case for a next-generation incoherent scatter radar design in the autumn of 2003. Their subsequent report on the future scientific direction of EISCAT, presented to the EISCAT Council in November of that year, can be found on the design study project website at www.eiscat.se/groups/EISCAT_3D_info/. At the same time a number of EISCAT member countries produced their own “white papers” summarising their plans for the future of EISCAT science, which can be found at www.eiscat.se/groups/EISCAT_3D_info/inputs. The immediate driver for these documents was not actually to produce a new radar, but to re-focus the scientific mission of EISCAT ahead of the renewal of its international treaty, due at the beginning of 2006. However, it transpired that many of the science requirements specified in these documents, particularly the more ambitious plans, would either be unachievable or severely sub-optimal if they had to be realised using EISCAT’s existing radar systems. The resulting discussions lead to the formulation of the EISCAT_3D concept and the submission of the design study proposal in 2004.

Once the EISCAT_3D design study proposal had been accepted, it was realised that a further forward look exercise was needed, since the existence of EISCAT_3D would imply a significant expansion of the range of science which EISCAT was able to investigate. A panel was therefore set up under the chairmanship of Professor Mike Lockwood (a former chair of EISCAT Council) with a remit to develop a highly innovative blue-skies vision for the future of EISCAT, with no disciplinary boundaries and unconstrained by any assumptions about the available hardware. In this way it was expected that the scientific vision would lead the hardware design, rather than the scientific possibilities being constrained by the limitations of the proposed system. This committee has worked largely through email and internet contacts, and produced its most recent report in April 2008. The science planning for EISCAT_3D has therefore been a mixture of evolutionary developments, based on the near-term development of EISCAT’s existing radars, and more speculative thinking based on radical new radar concepts, some of which will not even be achievable on the timescales envisaged here. The section below presents those scientific aims which could feasibly be addressed by an advanced phased array radar, with interferometric and advanced data processing capabilities, of the kind presented in this report, assuming that such a system can begin operations within the next five years.

3.1 Evolutions of current EISCAT science

In this sub-section we consider those science areas which are within the existing remit of EISCAT, but not easily achievable with the present hardware configuration. Because they have fairly well-specified requirements, these areas provide most of the basic rationales for the EISCAT_3D system design. They will also produce a guaranteed scientific return, in that the questions and the observations needed to resolve them are relatively well understood. By contrast, the next sub-section (3.2) considers more speculative areas, arguably with higher scientific potential, but with less firm guarantees of success since the understanding of these areas is less well-developed, making it more difficult to frame detailed observational requirements. The discussion of evolutions to existing EISCAT science focused on seven main topic areas:

- (a) The structure and coupling of the ionosphere and thermosphere.
- (b) The coupling between the upper, middle and lower atmosphere
- (c) The coupling between the ionosphere and the magnetosphere
- (d) Long-term monitoring of global atmospheric change
- (e) The development of the solar wind, and its interaction with Earth's environment
- (f) The physics of the aurora
- (g) Small-scale plasma physics
- (h) Studies of meteors, space debris and Near-Earth Objects

We will briefly consider each of these areas in turn.

- (a) **The structure and coupling of the ionosphere and thermosphere.** This is the core traditional study area for radars of this type, and contains a whole host of important science issues.

The ionosphere comprises only a small part of the upper atmosphere, being embedded in the neutral thermosphere, which can be hundreds of times denser. Because of its charged nature, and (for instance) its ability to respond to electric fields applied in the geospace environment, the ionosphere provides an important means of communicating solar energy to the atmosphere and dissipating it, either directly by particle precipitation, or less directly by the acceleration of charged particles which then collide with and heat particles of the neutral thermosphere, giving rise to bulk heating of the atmosphere which can in turn generate winds, waves, upflows and various types of non-thermal situations and plasma instabilities.

One key issue, currently out of range of EISCAT's radars, is to study how this energy input is structured, which can only be achieved by having access to quasi-simultaneous wide-scale maps of the main ionospheric parameters, especially ion and neutral densities and velocities, but also ion and electron temperatures and ion compositions. Such information cannot currently be accessed by EISCAT because all of its present radars are based on dishes, which are limited to single beams, directed to particular positions. To move the subject forwards, we require a different technology, such as phased arrays, capable of multiple quasi-simultaneous beams, providing instantaneous wide

spatial coverage. While this can already be achieved by the new generation of US radars, a further major advance would be to create the capability for a multistatic phased array system, which does not exist anywhere at the present time.

This would provide the ability to measure multiple simultaneous vectors of plasma velocity, allowing an experimenter to probe the instantaneous latitude and longitude structure of the plasma velocity field and furthermore to examine how the velocity field changes as a function of height, which would help to reveal the structure of the underlying thermosphere. Particularly advanced data sets could be created by combining data from such radar systems with data from instruments such as Fabry-Perot interferometers, specifically intended to study the neutral atmosphere. This imposes a requirement for a multistatic system, with baselines of some hundreds of kilometres, but ideally more than one site per baseline, for reasons of geometry. The phased arrays at each site should be capable of generating multiple quasi-simultaneous beams. At least one site should have a transmitter capability, but the optimum flexibility would be achieved if all sites were capable of both transmission and reception, with each site able to receive backscattered signals from any of the others.

- (b) **The coupling between the upper, middle and lower atmosphere.** This is probably the single most exciting emerging field in our subject area, and poses the challenge of understanding how the upper atmosphere is changing in response to long-term natural and man-made forcing factors.

This not only includes the investigation of whether solar and upper atmosphere variability might be able to influence climate, but also the obverse study of how climatic forcing in the lower atmosphere is affecting the upper atmosphere and whether there are any means to observe and quantify such effects. The requirement here is for long-term continuous observations, both to provide an extensive synoptic database and in order to get guaranteed coverage of unusual processes such as Solar Energetic Particle Events and stratospheric warmings. The study of dynamical processes such as atmospheric tides, planetary waves and middle atmosphere turbulence is especially important, since these are of key importance in re-distributing energy from the lower to the upper atmosphere. Because EISCAT's current radars were not intended for unattended long-term operations, there is a relatively little information on the climatology of long-period atmospheric waves and tides in the polar regions, though it is known that they behave differently from the predictions of current models.

There is also an important requirement to establish the spatial and temporal variability of mesospheric phenomena such as PMSE and PMWE, and of long-term trends in the height of atmospheric layers, which may contain signals of climate change. The required observations will comprise long continuous data sequences, on timescales of months to years, covering the height range from the F-region downwards, to as low an altitude as possible. These should be supplemented with data from instruments dedicated to middle

and lower atmosphere observations, such as MST radars, lidars and wind profilers.

In addition, the combination of such EISCAT_3D data with data from comparable radars at other longitudes (e.g. northern Canada and Alaska) would be of major importance in understanding the longitudinal structure and variability of tides and planetary waves. Comparison with Antarctic data would also be of great significance in understanding differences between the middle atmospheres of the northern and southern hemispheres.

- (c) **The coupling between the ionosphere and magnetosphere.** EISCAT has made many pioneering discoveries in the area of ionosphere-magnetosphere coupling, mainly by identifying the ionospheric proxies of processes occurring in the cusp and dayside magnetopause, and currently carries out an extensive programme of observations in support of space missions such as THEMIS and Cluster. While these missions will be complete by the time that EISCAT_3D is operational, support will still be needed for the next generation of European and international magnetospheric missions, including CrossScale, KuaFu, Radiation Belt Storm Probes and Solar Sentinels, leading toward the proposed Magnetospheric Multi-Scale mission.

Such missions would be significantly compromised if high-quality ground-based radar support for their observations was lacking. Again, the possibility to carry out observations with a wide simultaneous coverage would represent a major advance on current single-beam studies, while the potential to carry out such observations with similar radars, distributed in longitude, will produce very major improvements in the observational coverage of the high-latitude ionosphere. While EISCAT has made a number of discoveries in the magnetosphere via observations of the ionosphere, a major frustration has arisen from the fact that the current radars do not have sufficient power and gain to observe the near-Earth magnetosphere directly.

If EISCAT_3D were designed with significantly higher transmitter power and/or receiver aperture than the existing radars, the potential would exist to observe altitudes of 1000 km and above, and to access the important altitude regime in which ion acceleration occurs and amplification mechanisms such as the Ionospheric Alfvén Resonator might be active. These heights cannot currently be reliably studied by ground-based radars. This imposes a need for higher transmitter power, large receiving area and a frequency sufficiently low to avoid the worst effects of the Debye length cutoff.

- (d) **Long-term monitoring of global atmospheric change.** Perhaps the main basic requirement for EISCAT_3D will be the ability to make continuous long-term routine measurements, particularly of those regions of the upper atmosphere which display the signals of long-term global change.

These synoptic observations would include the long-term monitoring of ionospheric layer heights, convection patterns, wave and tide climatologies, thickness and duration of mesospheric echo regions and topside plasma outflow. The same transmission data stream, processed differently, would also

yield continuous information on space debris, meteors, and the presence/absence of sub-beamwidth structure. In this mode, the requirement would be for a data set probably based on a number of different modulation schemes, i.e. with separate modulations optimised for different height ranges, but producing a homogeneous ionospheric data set in a well understood format, together with summary information on space debris, meteors and small structures. This requirement is likely to define the basic mode of operation for EISCAT_3D.

- (e) **The development of the solar wind and its interaction with the Earth's environment.** Uniquely among the world's incoherent scatter radars, EISCAT has carried out a strong programme of passive solar wind observations using the technique of interplanetary scintillation (IPS), in which the scintillating signal from a compact astronomical source, such as a quasar, is observed through the phase screen of plasma irregularities propagating with the solar wind, which creates the scintillation. In this way, the EISCAT radars are used in pairs or groups of three, and in conjunction with other European radio telescopes, such as the MERLIN array in the UK.

The technique is becoming even more relevant with the advent of missions such as STEREO, which allows observations of CMEs observed by optical methods and those observed by the radars. More relevant still will be future missions such as Solar Orbiter, which will allow CME and CIR events tracked using the IPS technique to be sampled in situ. Interplanetary scintillation work is one aspect which is actually well suited to single-beam, steerable dishes and the use of high frequencies. Not only is the resulting beam highly directed, but low elevations can be reached without grating lobes (which would be created using phased arrays) and high frequencies facilitate probing of the solar wind at distances closer to the solar limb than are possible at low frequencies.

One important aim of these observations is to find the processes responsible for accelerating the solar wind as it streams away from the Sun, a process in which the most important interactions occur at radial distances less than 5 solar-radii. Unlike the other requirements, therefore, the IPS capability requires the retention of the existing steerable dishes, though for passive use only, allowing them to be used for interplanetary scintillation studies whenever the Sun is at a sufficiently high elevation and a suitable radio source is within a few degrees of the solar limb.

- (f) **The physics of the aurora.** The aurora is the best known of all phenomena in solar-terrestrial physics, but several important gaps still remain in our understanding of auroral physics. A key aim is to understand the factors which determine the structure of the aurora, in other words to find out what magnetospheric and ionospheric mechanisms dictate the production of rays, curls and streamers, and to determine the electrodynamic processes which correspond, for example, to the dark lanes between aurora.

The accurate measurement of electric fields and particle fluxes in very limited spatial volumes is important here, especially because auroral arcs have scale sizes much smaller than the beam widths of the current generation of

incoherent scatter radars. Also important is the determination of the nature and origin of the precipitating particles, e.g. to measure the relative contributions of electron and proton precipitation, and to discover whether the currents can be carried by precipitating particles alone or whether they require the acceleration and outflow of thermal ionospheric plasma in order to preserve current continuity.

The conventional problem that auroral phenomena have smaller scale sizes than radar beamwidths can be mitigated to some extent by creating larger radars, e.g. using extended arrays to minimise beamwidth. However, the sampling of really small spatial scales requires the application of radar interferometry, a technique which has been developed for ISRs at EISCAT. In this application, multiple spaced receivers are used, with a variety of different baseline lengths and orientations, and their data are auto-correlated and cross-correlated to derive a visibility function which allows probing down to scale sizes of a few tens of metres.

This approach has been used on Svalbard to investigate auroral spectra with Naturally Enhanced Ion Acoustic Lines (NEIALs) which occur in small regions of rayed auroral arcs, and appear to represent regions of driven plasma turbulence. At the ESR, a purpose-built interferometer is being used to implement this technique. However, the use of a phased array has the advantage that the array can be divided into modules (or sub-arrays) whose spacings can be altered such that they provide the elements of the interferometer. No incoherent scatter radar has yet been built like this, and this application provides a strong motivation for the novel design of the central site array of the EISCAT_3D radar.

- (g) **Small-scale plasma physics.** EISCAT has carried out a world-leading programme in small-scale plasma physics, based on controlled experiments to modify the ambient ionospheric plasma using high-powered HF transmissions from the EISCAT HF Heating Facility in Tromsø, and the SPEAR facility in Svalbard.

In typical “heating” experiments, strong HF transmissions are used to create field-aligned plasma irregularities, which are then allowed to decay into smaller-scale waves and turbulence. The plasma can also be heated to high temperatures, giving rise to enhanced airglow emissions (“artificial aurora”) and changing the chemistry and composition of localised regions of the upper atmosphere. The heaters can also be used to generate electrostatic waves, which can be ducted along field lines for magnetospheric probing, including the investigation of density structure and magnetic conjugacy between radars and satellites. Experiments like this have a fundamental limitation in that they are trying to study small-scale phenomena, which often act within very spatially localised regions, with radar beams whose widths are on the order of kilometres at the altitudes in question.

An important advance would be possible if the limitations of the radar beamwidth could be removed, and EISCAT has led the way in attempts to develop experiments using small-scale interferometer arrays to achieve sub-

beamwidth resolution. The first demonstrations of the technique have used a single low-gain receiver with baselines of order 100m, in conjunction with the two large dishes of the EISCAT Svalbard Radar. Although this initial configuration is far from perfect, it has demonstrated that the plasma physical phenomena known as NEIALs (Naturally Enhanced Ion Acoustic Lines), probably caused by the collapse of Langmuir turbulence into ion-acoustic waves, have a spatial scale of the order of tens to hundreds of metres (i.e. much smaller than the transverse beamwidth of the radar). The plan is to upgrade this system to five small receiver dishes, making seven receivers in all, which would allow the complete size and shape of these regions to be categorised and might provide unprecedented information on the scale sizes of plasma density enhancements generated by heating experiments, or of structure in PMSE layers.

The development of the radar interferometer technique is the subject of a distinct Work package in this study (WP5). An important objective of the EISCAT_3D design is to integrate the capability for interferometry as a standard part of the new radar system, so that interferometric measurements will implicitly be part of every EISCAT_3D experiment. This has very significant implications, not only for the design of the EISCAT_3D antenna array, but also for the design of the data system, as we discuss in later sections of this report.

- (h) **Studies of meteors, space debris, and Near-Earth Objects.** Although not strictly speaking an application of incoherent scatter, the EISCAT radars have very powerful capabilities in studying hard targets in the Earth's space environment.

This is an important capability since, for example, the meteor input function at high latitudes is not very well-known, and characterising the statistics of such meteors can help to define the importance of populations which have been produced within the solar system, compared to those which have come from elsewhere. This in turn is important for understanding some aspects of the formation of the early solar system.

The monitoring of man-made space debris represents an important applications area for EISCAT, since many agencies which operate satellites and human space missions require a knowledge of the debris environment and how it is changing over time. EISCAT has already carried out a programme of space debris monitoring for ESA, in particular monitoring the effect of the spreading debris clouds following the Chinese anti-satellite missile test of 2004, and the Iridium-Cosmos collision of 2009.

While the necessary observing techniques are already well-developed, the ability of the present EISCAT radars to obtain a good statistical data set is limited by the fact that the radars only support a single transmitter beam, which cannot be quickly scanned. If we were able to sweep rapidly through a range of positions, we would be able to build up a more comprehensive database on the meteor and space debris environment, as well as improving our speed and position information, which would also be facilitated by the use of multiple remote antenna beams and multiple common volumes.

3.2 New EISCAT science

While the previous sub-section has addressed science topics which are essentially extensions of EISCAT's present activities, the Science Oversight Committee has also been considering some more novel applications of the EISCAT_3D radar (see also Section 9.6). These applications are not the main science drivers for the new radar, but they have the potential to constitute major new areas of activity for the association over the coming decade. A few examples are reviewed here.

- (a) **GPS correction:** Satellite global positioning has become a vital capability in many areas of human activity, from navigation to road pricing to mineral exploration. However, the standard GPS technique is inherently imprecise, due to phase delays in the satellite beacon signals imposed by the Total Electron Content of the ionosphere and plasmasphere and by tropospheric water vapour. Removing one of these sources of uncertainty, even over a very limited geographic region, would have a number of interesting implications.

If the ionosphere/plasmasphere uncertainty could be removed, e.g. by using a radar to map the electron density above Northern Scandinavia very accurately, then not only would GPS positioning potentially become much more accurate, but the remaining unknown could be studied in isolation, allowing studies of tropospheric water content using GPS techniques. An obvious application of more accurate GPS in the Arctic region would be potential uses for studying land or ocean surface height, for applications such as the monitoring of glacial thinning and ice cover. From this perspective, high-quality electron density maps of the region above Northern Scandinavia could become an applications product of the EISCAT_3D system, which could be supplied to the GPS, geodesy and meteorology communities.

- (b) **Heterodyne Reflectometry:** This is a technique used in laboratory plasma physics, in which two signal sources at slightly different frequencies are used simultaneously to produce an additional "beat" frequency, which effectively becomes another probing frequency for the radar. This capability can allow a radar to operate at frequencies very different from those at which it was designed: for example, the simultaneous use of two VHF frequencies separated by a few MHz effectively provides the radar with an HF transmission capability. The application of this technique using the current EISCAT hardware has been discussed with the idea of creating a beat frequency at the second ion gyroharmonic, a very interesting frequency for plasma physics experiments. The current generation of radars, however, cannot implement such heterodyne techniques because of limitations in their transmitter design, making it desirable to include such flexibility into the EISCAT_3D transmitter.
- (c) **Magnetotellurics:** Magnetotelluric surveying is a technique already used in geology and mineral exploration, and involves multi-point measurement of electric potential at the Earth's surface in order to calculate the magnitudes and directions of the telluric currents which flow in the surface layers of the Earth's crust. From this the conductance of the crust can be calculated, allowing inferences to be drawn about its mineralogy.

Normally, magnetotellurics exploits natural variability in the Earth's magnetic field to induce the telluric currents in the crust. In the last decade, however, it has been realised that a powerful HF radar, or ionospheric heater could emulate or supersede the requirement for natural variability by artificially modulating the auroral electrojet, creating a very efficient means for inducing telluric currents over a wide area. Although the technique would not work at the VHF frequencies for which the EISCAT_3D system is intended, it could be used in conjunction with a co-located heater system, or by exploiting the potential of heterodyne reflectometry (see above) to create artificial magnetic fluctuations which might then be exploited by the geological community. This has no strong implications for the EISCAT_3D data system, other than necessitating that the system metadata could adequately describe when the technique was in use and the transmitter frequencies involved.

- (d) **Assimilative Modelling and Virtual Observatories:** Until recent years, ionospheric science largely consisted of the simple collection and interpretation of measurements, either from single instruments or combinations of different diagnostics. However the science has now evolved to a state in which the diverse set of measurements need to be interpreted in the context of a single theoretical framework. This means that the emphasis is now increasingly turning to assimilative modelling – coupling state-of-the-art models of the upper and lower atmosphere with real-time or retrospective observations, in order to constrain and improve the modelling capability, and move towards real predictability. The techniques of assimilative modelling (where real-time data is routinely assimilated into running models) and of re-analysis (where historical data is assimilated into models of previous system behaviour) are now well-developed in meteorology, but are only in their infancy for upper atmosphere science. The fact that the EISCAT_3D radar will have such advanced capabilities, allowing very accurate multi-point, multi-parameter measurements of the ionosphere over the full range of heights, would make the EISCAT_3D data set an ideal candidate for use in assimilative modelling. This will make some demands on the data, especially in terms of the real-time availability of analysed data. One method of facilitating the use of EISCAT_3D data in such assimilative models would be to make use of the emerging generation of Virtual Observatory projects, one of whose functions is precisely to identify an assembly of heterogeneous data resources in a given science area and to bring them together in a homogeneous user interface.
- (e) **Mesosphere and Stratosphere Monitoring:** The mesosphere, at an altitude of around 80km, is the coldest region of the Earth's atmosphere, and may be very important to understanding the effects of long-term global change since it contains phenomena such as Polar Mesospheric Clouds and Polar Mesosphere Summer Echoes, which correspond to the presence of charged dust/ice complexes, and whose occurrence is believed to be increasing with time. The dynamics of the mesosphere are also very interesting since dynamical coupling across the mesopause provides one way for lower and upper atmospheric processes to link together, while the mesosphere is also

characterised by complex chemistry whose effects can extend downward into the lower atmosphere.

The stratosphere, lying below the mesosphere at about 40km, represents a local peak in the atmospheric temperature profile, corresponding to the absorption of heat by ozone, making this layer very important for understanding the atmospheric energy balance. A cooling stratosphere is predicted to give rise to a descent in the overlying atmospheric layers and there is some indication from long-term ionospheric data sets, including EISCAT data, that such a descent is occurring as the lower atmosphere warms.

Monitoring of the mesosphere and stratosphere will be an important aim of EISCAT_3D and its co-located instruments, but requires a mix of techniques, some of which are most appropriate to other types of radars. Two factors would have implications for the EISCAT_3D data system: firstly the decay times of autocorrelation functions from radar backscatter at these altitudes are very long, while vertical scale lengths of mesospheric structure are short, necessitating the use of long codes (or even piece-wise codes) for probing these regions. Secondly the density of free electrons available for scattering drops sharply below 80km, requiring the use of other techniques such as the generation of Artificial Periodic Irregularities, using the interference between an HF pump wave and its ionospheric reflection to create irregularities which can then be probed using radar Bragg scatter.

- (f) **Planetary and Lunar Radar:** In addition to scattering from ion-acoustic waves in the upper atmosphere, or from hard targets in the Earth's immediate environment, such as meteors and satellites, the potential exists to carry out lunar and planetary studies using EISCAT_3D by using active radar beams to scatter from a lunar or planetary surface.

Radio waves can penetrate clouds, dust, regolith and even bedrock; and measurements of the phase and polarisation of the returned signal from planetary surfaces might provide interesting information on surface composition, especially because the proposed EISCAT_3D frequency is well below the frequencies used in standard radio astronomy. If any radar system using a similar frequency were constructed elsewhere in the world, the possibility would exist to carry out simultaneous long baseline probing of solar system objects as is presently undertaken at a much higher frequency by radio astronomy facilities such as Arecibo and Green Bank.

A further application of EISCAT_3D for planetary studies would be to use the radar to undertake studies of Auroral Kilometric Radiation, since AKR is the strongest radiation signature produced by the Earth, and might be key to the future identification of potentially habitable extra-solar planets, since the presence of AKR indicates an internally generated magnetic field, which may be a prerequisite for life.

The AKR generated by the Earth cannot be measured at the Earth's surface, since it is reflected back into space by the ionosphere. However, a radar such as EISCAT_3D could directly observe the causal mechanisms of AKR,

believed to arise from the interaction between electron beams (e.g. from aurora) and populations of non-thermal “horseshoe distribution” electrons in a converging magnetic field. The combination of EISCAT_3D and a suitable satellite would be a powerful tool for AKR studies.

In addition, the frequency of AKR depends inversely on the size of the emitting planet, so that Jupiter-sized planets produce kilometric radiation at frequencies above the ionospheric cut-off, opening the possibility that, depending on its frequency range, EISCAT_3D could measure such emissions directly.

While many of the above factors have some implication for the type of data which might be obtained from the EISCAT_3D system, probably none of them makes a more stringent demand than the already very exacting requirement imposed by the desire to measure the ion and plasma lines arising from normal incoherent scatter in the Earth’s ionosphere.

4 State of the Art in Scientific Radars

Before describing in detail the proposed design of the EISCAT_3D radar, it is worthwhile to note the characteristics of the comparable radars which exist around the world. In total, there are 11 separate incoherent scatter radar systems currently operational, three of which are operated by EISCAT. In addition, a few radar systems designed for the study of the middle atmosphere are actually powerful enough to undertake limited experiments using incoherent scatter.

Of the purpose-built incoherent scatter radars, eight are dish-based systems while the three remaining systems (Jicamarca, Poker Flat and Resolute) are based on phased arrays. Except for the EISCAT Svalbard Radar, the dish-based radars are all at least 20 years old, although all of them have been extensively refurbished during that period. The use of dishes can result in systems with high gain and good directivity - however dish-based radars are inherently limited by their ability to look only in a single direction with a slow scan, limited by the speed at which the dish can be moved. Not all of the dishes are fully steerable; the ESR 42m dish, the Arecibo antenna and the Millstone Hill zenithal antenna are fixed systems, essentially limited to a single look direction. Where steerable dishes do exist, however, they have the added advantage of being able to reach low elevations without deforming the beam pattern or producing large grating lobes.

EISCAT has operated dish-based radars since the early 1980s, since when they have yielded a very high scientific return. Further extending EISCAT science, however, relies on exploiting new techniques which are not available to the existing radars. These include the possibility of creating multiple quasi-simultaneous beams, carrying out volumetric imaging, and realising the potential of interferometry as applied to ionospheric radars. Within the past few years, some of these techniques have begun to be developed on the three phased array radars which currently exist. The Jicamarca radar in Peru has existed since the early 1960s, but because of its location at the magnetic equator, its scientific mission has been very different to that of EISCAT. Recently, Jicamarca has seen a number of developments in interferometry and radar imaging, including some of the development work done as part of the present study. Having these capabilities available in the auroral ionosphere would be very powerful, but we also require the rapid beam-steering which is not currently available at Jicamarca.

The AMISR radars, currently being developed in the US, have a number of desirable features, including rapid beam-steering. However, these systems are not really designed for interferometry and radar imaging, their gain and timing accuracy are rather lower than we envisage for EISCAT_3D and, in the European sector, their frequency range is reserved for fixed radio links and base stations for mobile communications systems. Hence we require the creation of something entirely new – a phased array radar having the steerability of an AMISR or MU/EAR, with the directivity and gain of the best dish-based systems, and the capacity for interferometry and radar imaging which has to be factored into the array design from the beginning. This explains why EISCAT_3D will be different from any other radar currently in existence.

4.1 EISCAT

The European Incoherent Scatter Scientific Association, EISCAT, operates the world's largest system of incoherent scatter radar installations and other radio diagnostics. Two ISR transmit-receive systems, operating on respectively 928 MHz (the UHF) and 224 MHz (the VHF), as well as one of the world's most powerful HF heaters, are located close to Tromsø, Norway. The UHF system also comprises two receiver stations, located at Kiruna, Sweden and Sodankylä, Finland. A third ISR system, the EISCAT Svalbard Radar (ESR), is located close to the Longyearbyen settlement on the Arctic island of Spitsbergen.

The first-generation EISCAT system (aka the "mainland system") stands out from all other incoherent scatter facilities because it was designed as a unified system with unique features from the beginning. Its origin goes back to the 1969 URSI General Assembly in Ottawa, when a recommendation was passed urging European research groups to investigate the possibility of establishing an incoherent scatter observatory in the European auroral zone. Several meetings of interested scientists were held during 1970 and 1971 and a proposal for the facility was prepared and presented to European research councils. In 1973, a Steering Committee was eventually formed and charged with the drafting of an international agreement to be entered into by the various research councils, and with the preparation of detailed technical and administrative plans for the observatory.

To ensure optimum performance both in the high electron density conditions prevailing during active times and also in the low electron densities found in the mesosphere and ionosphere topside, it was decided to construct two separate ISRs, operating on respectively about one-half (UHF) and twice (VHF) the 70-cm wavelength used by many other ISR systems. Construction of the UHF system started in 1975-1976; the system was officially inaugurated in August 1981. Its tri-static geometry with two receive-only stations in Kiruna and Sodankylä is a unique feature among the world's ISR facilities, giving it the capability to perform unambiguous point measurements of the full vector ionospheric plasma velocity throughout the ionosphere – albeit only at one point at a time. The UHF antennas are fully steerable, 32-meter Cassegrain reflectors on wheel-on-track mounts, derived from a design originally developed for the Intelsat 1 C-band ground segment. The gain at 930 MHz is 48 dBi and the one-way beam-width is 0.6° . Typical system temperatures are 90-110 K at Tromsø and 50 K at the receiver sites.

The VHF radar came on-line about two years later. Designed for maximum sensitivity, this system is equipped with a massive 120 x 40 meter parabolic cylinder antenna, mechanically steerable in a N-S plane from a lowest elevation of 30° towards geographic north to well beyond the local magnetic field line. Limited steerability in a plane through the elevation axis and the boresight direction is provided by a system of interchangeable phasing cables. The antenna can be operated either as a single reflector with about 46 dB gain, generating a single $0.6^\circ \times 1.8^\circ$ beam, or as two independent 60 x 40 m antennas with about 42.5 dB gain, each generating a $1.2^\circ \times 1.8^\circ$ degree beam. This dual-beam mode is routinely employed for long-range mapping of the ionospheric E field to well north of Spitsbergen by phase-steering one of the beams off the boresight direction and operating at the lowest elevation of 30° .

Both systems use pulsed high-power klystrons in their transmitter power amplifiers. The VHF is currently operating off a single 1.5-MW five-cavity klystron, while the UHF runs two 1.3-MW four-cavity klystrons with a power bandwidth of about 6 MHz, allowing the transmission of pulses down to 1 μ s pulse length. The transmitted waveforms are freely programmable within the limits set by the available duty cycle (12.5 %) and pulse length (max. 2 ms), thus offering almost total flexibility in adapting the system for specific observational needs.

Towards the end of the 1980s, the Association began to draw up plans for an ISR on Spitsbergen to support the ESA Cluster multi-satellite mission with ground-truth data from the magnetospheric cusp region. For logistical and financial reasons, the new system (the ESR) would have to be a single-frequency installation operating on approximately 500 MHz, and so great efforts went into designing low-noise antennas and receiver systems to ensure good performance even in the low electron density conditions expected that far north.

Construction started in 1991 and proceeded in three phases. The first build phase, comprising a fully steerable, made-to-order moving-head 32-m Cassegrain dish antenna with 42.5 dBi gain and a 500-kW transmitter was completed on schedule in August 1996; the transmitter system was upgraded to 1 MW a year later and in 1998 a 42-m fixed dish antenna pointing up along the local magnetic field line was added. Using the two dish antennas as the two endpoints of a single baseline, radar interferometry has been successfully implemented. The recent addition of a small Yagi array located between the dishes now allows three different baselines to be formed. Recently, the Chinese Research Institute of Radio Propagation (CRIRP) has submitted a proposal to construct a third large antenna at the ESR site. The new antenna would be a fully steerable dish, with the possibility of dual use for incoherent scatter and deep space communications.

In the ESR transmitter system, eight modified UHF-TV transmitters, each running two 64-kW external-cavity klystrons in parallel, are power-combined for a total output power of 1 MW. The klystrons are normally tuned for maximum output power, resulting in a power bandwidth of about 4 MHz, but can be re-tuned for up to 10-MHz bandwidth if required at the expense of some power loss. The two antennas are connected to the two output ports of the four-port hybrid forming the last stage of the power combiner; by flipping the phase of the low-level RF drive signal going to four of the transmitters, the full 1-MW output power can be sent to either antenna on a pulse-by-pulse basis. The disadvantage of a relatively low peak power is offset to a large degree by an extremely high RF duty cycle of 25 %. Using this efficiently requires regular use of multiple frequencies and unconventional approaches to experiment design and signal processing, however.

To handle the multi-frequency receiving task, the ESR receiver employs an EISCAT-designed, fully digital multi-channel back-end with a 15-MHz, 14-bit ADC feeding into a number of individually programmable digital down-converter / low-pass filter / decimation channels, being probably the first ISR receiver in the world to use this technology. The band-limited, decimated data is normally processed into time-averaged autocorrelation estimates using multi-threaded software running in a Solaris environment on multi-CPU server platforms, but direct recording to disk of the unprocessed complex amplitude samples is also possible when special post-processing

is required. Plasma-line reception is a routine part of most standard experiments and used extensively for system calibration purposes. Starting in 2000, the receivers and signal-processors of the mainland system have been upgraded to use the same technology, such that more or less identical back-ends are now in use at all sites.

All raw data are securely archived in-house and can be provided to users on a variety of media to suit individual requirements. The autocorrelated measurement data is also fitted for ionospheric parameters (electron density, electron and ion temperatures, Doppler velocity and, when possible, mean ion mass) in near-real time using the GUISDAP analysis software suite. Analysis results are available on-line from the Association's WWW servers and also regularly uploaded to international ionospheric science databases.

4.2 AMISR

The Advanced Modular Incoherent Scatter Radar (AMISR) is an incoherent scatter radar project constructed by SRI International, under a grant from the National Science Foundation. It is a modular, mobile facility using a phased-array antenna system.

The most basic element in the AMISR design is the Antenna Element Unit (AEU), which consists of a power amplifier, polarisation and transmit/receive circuitry, an antenna, a low noise amplifier, signal processing and control and monitoring electronics. Thirty-two identical AEU's are placed in fixed locations on a "panel", with the locations determined by the desired grating lobe-free scanning of the final antenna. A panel thus represents the minimum-size phased array antenna, with 128 panels combined to form a "face", roughly 30 × 30 metres. For scientific objectives requiring maximum sensitivity, all the available antenna panels can be installed to operate as a single face. Where broad spatial coverage is crucial, the panels can be distributed over multiple faces that act as multiple separate ISRs. Separate faces can also be deployed to different locations to provide support for multiple measurement goals, to provide even greater spatial coverage or to provide capabilities not available with a monostatic radar system (e.g., tristatic drift velocity measurements).

Most presently operating ISRs rely on a few unique components (such as steerable parabolic dishes and klystron power amplifiers) that have large individual price tags both for purchase and maintenance. This is not the case for the AMISR, where the cost is due to the number of components, not their individual costs. The system development can thus proceed from relatively modest design efforts, including testing and optimisation of the AEU's, to a full deployment of an operational ISR at any pace necessary. Although there are cost factors associated with the size of each development phase of a fully functional ISR (e.g., quantity price breaks and the need to sustain interest from multiple vendors), the rate at which the system is developed can readily conform to the realities of funding cycles.

The AMISR is being constructed in two stages. The first face in Poker Flat, Alaska, has been completed and is already being used for scientific investigations. The remaining two faces are under construction in Resolute Bay, Nunavut, Canada. A

feature of the AMISR design is that it can be easily repositioned and operate in a number of different configurations, with future AMISR locations determined by a scientific advisory panel. Since each face of the AMISR functions independently, AMISR can be deployed in up to three separate locations at the same time. It is able to electronically steer the beam on a pulse-to-pulse basis, thus allowing experiments that were hitherto impossible. Remote operation and electronic beam steering let researchers operate and position the radar beam instantaneously to accurately measure rapidly changing space weather events.

Centre frequency	449 MHz	Antenna type	3 × Phased arrays
Peak power	2 MW	Element type	Crossed dipoles
Max. duty cycle	10%	Antenna size	128 × (1.5 × 3.5 m)
Pulse length	1 – 2000 μs	Antenna gain	43 dBi
Polarisation	Circular	Beamwidth (FWHM)	1°
System temperature	120 K	Field of View	±25°

Table 4.1: AMISR Technical Details

4.3 Jicamarca

The Jicamarca Radar Observatory operates a scientific facility for studying the equatorial ionosphere. Located almost at the magnetic equator, the radar can be pointed perpendicular to the nearly horizontal geomagnetic field throughout the ionosphere. The 49.92 MHz incoherent scatter radar is based on a large square array of 300 metre per side containing 18,432 half-wave dipoles arranged into 64 separate modules of 12 x 12 crossed half-wave dipoles. Each linear polarisation of each module can be separately phased, and the modules can be fed separately or connected in almost any desired fashion. This allows great flexibility, but changes cannot be made rapidly. The half power beam width of the full array is about 0.8°. The frequency bandwidth is about 1 MHz. The transmitter consists of three newly designed modules each delivering a peak power of 1.5 MW, with a maximum duty cycle of 6%, and pulses as short as 0.8-1.0 μs and as long as 2 ms. There are currently 16 receiving channels.

Owing to its long wavelength, the Jicamarca signal is not affected by Debye length problems at low electron densities, so that the radar obtains signals from altitudes of 5000 km and higher. It is also the only radar that makes routine absolute F-region measurements of electron density using Faraday rotation. The Jicamarca radar is the most sensitive MST radar in the world; in fact, it is the only true MST radar, capable of probing even the “gap” region near 50 km, partly because of its long wavelength. Jicamarca interferometric measurements, have allowed the identification of the main sources of sporadic meteor echoes and the detection and identification of meteor showers. Careful studies of meteor mass are currently underway which may have important astronomical implications. The utilisation of the radar interferometric and imaging technique has produced spectacular 2- and 3-D visualisations of the turbulent dynamics of spread-F, the electrojet and the 150-km echoes which has led to a number of insights into the physical processes of these phenomena.

Jicamarca has played an important role in the EISCAT_3D design study, since many of the experiments connected with the radar imaging development for EISCAT_3D

were performed in collaboration with JRO. Staff at Jicamarca have pioneered the radar imaging technique by exploiting the flexibility provided by the 64 antenna modules. Part or all of the antenna is used to illuminate the ionosphere, and the echo signals are received by a number of modules or sets of modules which are combined employing interferometric techniques to synthesise a 2- or 3-Dimensional image.

Jicamarca also hosts a number of other smaller radars including JULIA (Jicamarca Unattended Long-term investigations of the Ionosphere and Atmosphere) which makes nearly continuous observations of ionospheric and atmospheric irregularities for long-term trend and climatological studies, SOUSY, employed for studies of the stratosphere and mesosphere, whose capabilities approach those of the MU radar, and JASMET (Jicamarca All-sky Specular MEteor Radar) used to measure night time winds in the mesosphere. An AMISR-7 prototype is also currently on loan from SRI International. In addition the site is equipped with ionosondes; chains of magnetometers; Fabry-Perot interferometers and an all-sky optical imager.

4.4 Sondrestrom

The Sondrestrom radar is an L-band system, located near Kangerlussuaq (formerly Søndre Strømfjord) on the west coast of Greenland, and operating at frequencies around 1290 MHz. The radar incorporates a 32 metre fully steerable dish, which was previously located at Chatanika, Alaska, and before that in Stanford, California. The peak transmitter power is 5 MW, though the maximum duty cycle is only around 3%. The antenna beamwidth is 0.6° and the antenna has 50 dB gain. Pulse lengths from 2 to 500 μs can be used, with coding down to a resolution of 1 μs . While some of the radar hardware is several decades old, the receiver system has been refurbished, making it capable of running and decoding the same types of advanced coding schemes currently in use at EISCAT. A wide range of experiments is run, including wide-area scans for the mapping of convection and high-latitude morphology and experiments in support of rocket flights from the co-located rocket range.

The Sondrestrom radar forms the centrepiece of a comprehensive array of supporting equipment including single-beam and imaging riometers, a digisonde, a lidar facility, magnetometers, all-sky TV and auroral camera systems, as well as interferometers, photometers and spectrometers, ELF and VLF receivers. This collection of instrumentation means that the ionosphere and thermosphere above western Greenland are among the most intensively diagnosed regions of the Earth's upper atmosphere. In principle, the multi-instrument observatory capability which exists at the Sondrestrom facility is a template for the kind of geospace observatory which we hope will develop around the EISCAT_3D radar facilities.

4.5 Arecibo

The incoherent scatter facility at Arecibo, inaugurated in 1963, exploits the world's largest single-dish radio telescope, a 305-metre spherical reflector suspended over a natural limestone sinkhole. The feed system is suspended 150m above the reflecting surface, and can be steered up to 20° from the vertical by moving the tertiary and secondary feed assemblies along a curved circular track. The facility is devoted

mainly to radio astronomy, but a small fraction of the operating time is reserved for incoherent scatter use, utilising a 430 MHz two-klystron transmitter with a 1 MHz bandwidth and total peak power of 2.5 MW. The large collecting area of the dish means that the Arecibo system is the most sensitive of the worldwide ISRs, capable of reaching the ionospheric topside, where hydrogen ions constitute an appreciable percentage of the composition. As well as the radar facility, Arecibo also hosts a number of atmospheric lidars, spectrometers and interferometers, and an HF heater is currently nearing completion. Scientists at Arecibo have made a number of contributions to developing the ISR technique, including the use of random codes and the development of constrained analysis techniques to fit data in cases of multiple ion populations having different characteristics.

4.6 Millstone Hill

The Millstone Hill radar is located near Westford, Massachusetts, in the continental US. Despite its mid-latitude location, its high antenna gain allows studies from just short of the Arctic circle to the north, to the Caribbean region in the south. The facility consists of two dishes, a 68m zenithal antenna and a 46m steerable dish. The operating frequency is within a bandwidth of 1 MHz around 440 MHz, and the signal is generated using two klystrons, each capable of 2.5 MW, with a 6% duty cycle. Pulse lengths from 1 μ s to 2ms are possible, and the radar uses a variety of modes including binary polarisation and phase codes and chirped pulses. Twelve receiver channels can be used over a receiver bandwidth which can be up to 52 MHz.

The zenith-pointing dish began to be used for incoherent scatter studies in 1974, with the steerable dish being added in 1978. The radar operates as a key element of the meridional chain extending from Søndre Strømfjord in the auroral region, through Millstone Hill and Arecibo, down to Jicamarca close to the equator. As well as the radar, Millstone Hill also hosts a range of optical instrumentation and an ionospheric sounder. Staff at the facility have made a number of developments in signal processing, data analysis and data distribution, including the development of the Madrigal archive, which has become the standard means of disseminating ISR data.

4.7 Kharkov

The Kharkov ISR system is operated by the Institute of Ionosphere, a department of the Kharkov Polytechnics Institute (Kharkov, Ukraine). The radar system is situated at the radio physical observatory approximately 50 km from Kharkov (geographical coordinates 49.6° N, 36.3° E). Since 1972 the observatory has conducted systematic measurements of ionospheric parameters in the altitude range 100 to 2000 km.

The VHF radar system (operating frequency 158 MHz) consists of two parabolic antennas, one of them, 100 m in diameter, is directed vertically, and the other, 25 m in diameter, is fully steerable. The VHF transmitter can be operated in two modes providing peak transmitted power 2.6 or 3.6 MW. The radar is able to transmit pulses 65 to 840 μ s in length at repetition frequencies between 25 and 60 Hz. Transmitted signals can be either linearly or circularly polarised. The receiving system is of super-heterodyne type with triple frequency conversion. The receiver's bandwidth is

300 kHz. The receiving system temperature is between 100 and 250 K depending on the type of antenna amplifiers used.

4.8 Irkutsk

The Irkutsk radar, actually located 120 km northwest of Irkutsk in southern Siberia, was realised through the conversion of a former military system, known as Dnepr, for civil scientific use by the Siberian Branch of the Russian Academy of Sciences. This conversion involved the complete replacement of the receiver chain, signal processing hardware and computing systems. The resulting facility is a monostatic ISR, working at frequencies from 154 to 162 MHz, with two transmitters jointly capable of 3 MW transmitter power. The radar uses a sectorial horn antenna of dimension 246 m × 12 m, with an antenna gain of 35 dB and a noise temperature of 450 K. The beam width is typically around 0.5° in the north-south plane and 10° in the east-west plane, and can be scanned in the north-south direction between limits of 30° either side of the zenith. The radar currently uses pulse lengths from 70 to 900 μs and a pulse repetition rate up to around 25 Hz.

A major upgrade during the period 2003 to 2005 included the renovation of the transmitter, receiver and signal processing capabilities. The ability to handle multi-pulse and phase coded pulse schemes was added, together with pulse-to-pulse beam steering, multi-channel reception and real-time signal processing. Some degree of scanning in the east-west plane is now possible and the possibility also exists to carry out interferometric measurements. While still not as advanced as facilities such as EISCAT, the geographic location of the Irkutsk radar is such that it forms an important link in the chain of incoherent scatter radars including Millstone Hill, Kharkov and the Japanese MU facility. The Irkutsk radar has been contributing regular ionospheric data as part of the co-ordinated World Day programme since 1993, and also plays an active role in the monitoring of space debris.

4.9 MU and EAR

In addition to the incoherent scatter radars discussed above, there are two further facilities capable of performing incoherent scatter experiments, although they are really optimised for the study of the middle atmosphere. The MU (Middle and Upper atmosphere) radar was constructed near Kyoto, Japan in 1984, in order to investigate atmospheric and plasma dynamics from the troposphere to the magnetosphere. MU is a monostatic pulsed Doppler radar, with a frequency of 46.5 MHz, based on an active phased array of 475 crossed yagis and the same number of solid-state transmitter/receiver modules. The system is capable of very fast beam-steering, on a pulse-to-pulse basis, within 30° of the zenith. The antennas are distributed in a circular array of 103m diameter (8330 m²), the peak power is 1 MW and the average power is 50 kW. The antenna beam is conical, with a half-power beamwidth of 2.6°.

In 2001, Japanese and Indonesian scientists constructed a very similar system, known as the Equatorial Atmospheric Radar (EAR), close to the equator in Western Sumatra. The EAR frequency (47 MHz) is almost the same as that of MU, though the antenna array is slightly larger (560 dipoles, 110m diameter). The transmitter power is the

same as MU and the beamwidth is comparable (3.4°) with the same rapid steerability. The MU/EAR design is seen as being highly successful, and a plan now exists to site a similar radar facility, known as PANSY at the Japanese Syowa station in the Antarctic.

4.10 MST radars

In the troposphere and lower stratosphere (up to about 30 km) radar signals are returned from refractive index fluctuations produced by turbulence in the neutral atmosphere; these are sometimes referred to as clear-air echoes. In the upper stratosphere and lower mesosphere (between about 60 and 100 km) refractive index variations are strengthened by the strong vertical gradient in electron density. Because the scale size of the refractive index fluctuations must be of the order of one-half the radar wavelength (Bragg scattering), and the minimum scale size of turbulence increases with height from a few centimetres in the lower troposphere to several meters in the upper stratosphere and lower mesosphere, most MST radars operate in the VHF band (typically 30–60 MHz, 5–10 m wavelengths). MST radars are characterized by high- powered transmitters and large antennas. (VHF antennas range from 100 to 300 m across.) They typically make continuous measurements of 3-d wind vectors and radar reflectivity in an altitude range from 1 to 20 km, and can then also measure strong echoes from the mesosphere (above 65km) where ionisation, and phenomena such as Polar Mesosphere Summer (and Winter) echoes become important. Smaller wind profilers that lack the transmitter power and antenna area to detect returns from the upper stratosphere and the mesosphere have often been called ST (stratosphere–troposphere) radars.

Modern MST Radars are state of the art instruments capable of providing estimates of atmospheric parameters with very high resolution on a continuous basis which are essential in the study of different dynamical processes in the atmosphere. The MST technique is an important research tool in the investigation of prevailing winds, waves (including gravity waves) turbulence, atmospheric stability and other mesoscale phenomena. MST radar data provide a key input to reliable three dimensional modelling of the atmosphere which in turn improves our understanding of climatic and weather variations. A large number of MST radar and wind profiler systems are operating worldwide. They include the Morro MST radar on the EISCAT site at Tromso, the MU and EAR radars (described above), the ALWIN radar in Andenes, the UK MST radar at Capel Dewi, and large systems in India and China, among many others.

The combination of MST and ISR capabilities is very powerful for investigating a range of physical processes, including transport between different atmospheric layers, energetic precipitation, atmospheric turbulence and the physics of dusty plasmas.

4.11 Future developments

Incoherent scatter is a continually evolving field, and in addition to EISCAT_3D, there are also two other future ISR facilities currently in the planning stage. Neither of these is in any way a competitor to EISCAT_3D, indeed the proposed Antarctic

radar is likely to have a key synergy with EISCAT_3D in that the two radars together will allow the northern and southern high-latitude ionospheres to be observed simultaneously in a way which has never previously been possible.

4.11.1 The Chinese ISR

For several years, the Chinese Research Institute of Radio Propagation (CRIRP) has been in possession of a former military radar station around 100 km from Kunming, in Yunnan province of southern China. The installation consists of a fully-steerable 40m parabolic dish with a Cassegrain feed, designed to operate at UHF frequencies. The old transmitters and receivers were removed when the site was vacated by the Chinese military and, until recently, no funding has been available to convert the radar for scientific use. In the last year, however, funding has been allocated to this project and contracts have been signed to develop the necessary hardware. Much work still needs to be done, but it is hoped that the facility will be operational as an ISR facility within five years.

4.11.2 An Antarctic ISR

In August 2008, a workshop was held in Chicago to explore the possibility of deploying an ISR in the Antarctic. This concept originated with the group at Stanford Research International (SRI), in Menlo Park, California, which is responsible for developing and operating the AMISR instruments. A proposal is currently under development for submission to the US National Science Foundation. This would envisage the operation of an AMISR-type radar, probably at either the US McMurdo base, the Chinese Zhongshan base or the Australian Davis station. The workshop was attended by scientists from the USA, Australia, Japan, Peru, South Africa and the UK with the idea of establishing an international consortium to push this project forward. Locating the radar at Davis or Zhongshan would have clear benefits for EISCAT, which is close to the magnetically conjugate point for these locations. If the radar were located at McMurdo it would be nominally conjugate to the Resolute Bay AMISR, though the high magnetic latitude of this station makes discussion of conjugacy relatively meaningless. It is also likely that an AMISR will be deployed in Argentina for studies of the South Atlantic magnetic anomaly.

4.11.3 PANSY

The PANSY project is a Japanese initiative to introduce an MST (Mesosphere-Stratosphere-Troposphere) /IS (Incoherent Scatter) radar in the Antarctic Syowa Station. The centre frequency will be 47 MHz and the peak power will be 500 kW (with an average of 25 kW). The antenna array will be a quasi-circular array consisting of 1045 3-element crossed Yagi antennas with sensitivity within 30 degrees of zenith. The total collecting area will be approximately 16000 m².

4.12 Summary

From the descriptions given above, it is obvious that there exists a wide range of radar types, with a considerable variety of capabilities. Although many of these radars

satisfy certain parts of the performance requirements for EISCAT-3D, no existing radar gives a complete match to the full range of capabilities needed by the research community for a next-generation EISCAT system. Therefore, although we can use aspects of existing radar designs, the design study team has been compelled to seek some novel solutions for realising the new radar. In the next section we review the major design challenges in specifying the EISCAT-3D system, and in Section 6 we discuss the methodology which was selected in order to address these challenges.

5 Design Challenges

5.1 Introduction

The construction of a large scientific sensor like the EISCAT_3D system is a complicated, lengthy and costly process. To make the required effort and investment worthwhile in the first place, the system must provide its user community with major performance improvements (preferably “quantum jumps”) in a number of mission-critical areas right from the start. But thereafter, it must be kept in operation and maintained at the forefront for a long time, often well beyond the point at which its book value has been depreciated to zero, in order to deliver maximum return on the public funding that made the construction possible. A case in point is the EISCAT UHF system, which was first commissioned in 1981 and is still in operation 28 years later; the projected lifetime of the 3D system is 30 years from start-up.

Accordingly, the overriding challenge for the EISCAT_3D Design Study has been to produce a system design that is at the same time both cutting-edge and future-proof. That is, the system should deliver the best possible performance commensurate with the current state of the art, but it must also be able to accommodate future breakthroughs in radar theory and practice with little or no modification to the basic hardware architecture. To accomplish this will require transmitter, receiver, signal-processing and data storage subsystems with a previously undreamed-of degree of flexibility.

In 2005, all this was seen to actually be within reach, provided that state-of-the-art developments in high-speed digital electronics, digital signal processing and software radio could all be brought to bear on the design. The confluence of analogue and digital techniques in the GHz range had now reached a point where multi-bit analogue/digital converters running at more than 100 MHz sampling rate were available off-the-shelf at affordable prices. For the first time it would therefore be technically and financially realistic to consider a distributed direct-sampling receiver design. This would offer a number of important advantages – in particular, several prerequisites for the implementation of freely programmable multi-beaming, a capability long since featuring at the top of the users’ wish lists, would be met automatically. While the amount of digital data produced by an array using this technique might appear staggering at first sight (at least 2.4 Gbit/s/array element), industry-standard optical fibre systems and field-programmable gate array technology could be used to transport the thousands of gigabits per second from the individual receivers and process them into a manageably small number of beam-formed data streams. Also, thanks to the recent breakthroughs in magnetic mass storage technology, massive amounts (hundreds of terabytes) of data could now be stored on off-the-shelf hard disk arrays, thus making it possible to buffer the complex amplitude data from each beam for a significant length of time. This would allow users with special requirements to analyse interesting intervals with non-standard methods.

Similar advances had taken place also in the signal generation area. Silicon subsystems performing digital up-conversion of base-band signals to carrier frequencies in the 100-MHz range, followed by high-precision digital-to-analogue conversion, were already available. 250-MHz devices were expected to be available

by the end of the study in early 2009, thus making it realistic to consider providing each array element with its own arbitrary-waveform generator / exciter subsystem. Field-effect power transistor technology had advanced to a point where single devices could reliably provide several hundred watts of continuous-carrier RF power in the 250-MHz range at power gains exceeding 10 dB, so greatly simplifying the design of compact high-performance modular power amplifiers.

Given these extremely promising initial conditions, one might be led to assume that designing the envisaged radar system would be a fundamentally straightforward task. As the project got underway, however, it quickly became clear that as always, the devil was in the details: In order to make the favoured system architecture meet the global performance requirements put forth by the prospective 3D user community, the performance envelopes of many subsystems would have to be stretched in partially unexpected directions; in certain cases, it was even not clear if the required performance could be achieved reliably or at all. To further complicate matters, all this would have to be achieved subject to the frequency allocations awarded and restrictions imposed by the telecom administrations of Norway, Sweden and Finland. To exemplify the challenges encountered, short summaries of their effects on transmitter, beamsteering and polarisation matching subsystems and the array and element antenna designs are given below.

5.2 Transmitter system

Ever since the start of EISCAT operations in 1981, the development and implementation of radar codes for different measurement applications has been an integral part of the EISCAT scene, actively pursued by users and staff. The best-known example of this work is probably the alternating-codes technique developed by Lehtinen and Hågström [Leh87], which almost immediately became a *de facto* world standard.

Alternating codes, as well as most other coding schemes previously and currently used for incoherent-scatter measurements, are typically based on binary phase-shift ($0^\circ/180^\circ$) sequences. Recent theoretical studies by Lehtinen's group have however demonstrated that there exist many codes employing combinations of arbitrary phase and amplitude modulation that are completely free from range ambiguities. These codes show great promise and may actually be superior to constant-amplitude codes under many conditions. Unfortunately, it has not yet been possible to test them in practice, because the EISCAT transmitter systems cannot accommodate amplitude-modulated signals.

For decades, standard practice in incoherent-scatter transmitter design has been to improve the signal-to-noise ratio by maximising the transmitter RF output power within the budgetary, technical and other constraints in force. Power amplifiers have therefore been designed to run the active devices (klystrons, power grid tubes or transistors) into saturation in order to improve the DC-to-RF conversion efficiency, the transmitter exciters have traditionally only included provisions for binary phase-shift modulation, and amplitude modulation has not normally been implemented for the simple reason that it does not use all the available peak power all the time. All EISCAT transmitters, including the ESR one, are designed in this way.

In order to remove these restrictions and provide users with total flexibility up to the bandwidth limitation imposed by the frequency management authorities, *the 3D transmitter system would have to be designed as a giant arbitrary-waveform generator, combining a programmable and totally flexible exciter subsystem with a very large number of class-AB solid-state linear power amplifier modules with at least 20 dB of dynamic range.* While representing a substantial departure from current trends in RF power amplifier technology, which focus heavily on switching (class-E) designs, the envisaged system would handle any and all currently known codes and would also be able to accommodate any future, currently unforeseen developments in coding and modulation theory without any modifications, so making the EISCAT_3D transmitter system essentially future-proof. Carrying this design philosophy to its logical conclusion, i.e. providing each array element with its own self-contained digital exciter/amplifier system, it would become possible to fully control the signals emitted by each individual element antenna. This would be an extremely valuable asset for several reasons, e.g.:

- The time-delays required to steer the transmitted beam in the desired direction could be introduced at baseband and loaded into the individual exciters, so effectively eliminating the need for dedicated beam-steering hardware,
- Individual timing corrections could be applied to each element as required, thus optimising the focussing of the beam,
- It would be possible to apply any arbitrary degree of amplitude tapering across the array aperture when transmitting. This could be used e.g. to suppress sidelobes along the horizontal or to widen the transmitted beam,
- Phase offsets tailored to e.g. generate an $L \neq 0$ beam could be introduced, thus opening the door to radar experiments with signals carrying orbital angular momentum.

The frequency band being considered for the active array is squeezed in between two digital audio broadcasting channels, and the operating permit will therefore probably be issued on a non-interference basis. For this reason, *all possible measures to suppress out-of-band emissions must be taken already in the design stage.* Using the arbitrary-waveform capability to apply tailored spectrum masks to all transmissions would go a long way towards achieving this goal, but in addition both the system reference clocks and the power amplifier chains must be designed for the lowest possible phase noise. *Provisions for running the transmitter system at substantially reduced power (10% of full power) should also be designed in from the beginning,* both to reduce the operating costs when running in “ionospheric monitor mode” and to reduce the field strength close to the radar site.

5.3 The beam-steering problem

In order to recover all the information about the scattering medium contained in the incoherent-scatter signal, the instantaneous bandwidth of the radar receiving antenna must be larger than that of the scattered signal, whose spectrum is the convolution of the spectrum of the transmitted probing signal and the plasma wave spectrum at the

scattering point. Under conditions of high electron density, the bandwidth of the latter, including both the up- and downshifted plasma line complexes can be more than 25 MHz, such that the total bandwidth of the scatter signal can exceed 30 MHz. In the case of an electrically steerable phased-array antenna operating at 240 MHz, the resulting relative bandwidth of more than $\pm 6\%$ constitutes a very non-trivial problem, the more so if the simultaneous generation of several beams is contemplated.

Phased arrays employ either of two different beam steering techniques, viz. *phase steering* and *time-delay steering*. In phase steering, the excitation phases of the individual array elements are adjusted modulo- 2π to make the radiated signals phase-coherent along a wave-front normal to the desired pointing direction. In most currently operating systems using this technique, the phase-shifts are introduced through programmable analogue phase shifters, constructed from short sections of transmission line that are switched in or out of the RF signal paths as required.

In time-delay steering, delays equal to the total differences in wave propagation time are inserted in the signal paths to the different array elements, causing the radiated signals to become both time- and phase-coherent along a wave-front normal to the pointing direction. This can be accomplished e.g. by a system of switchable delay lines, which are in principle no different from the phasing lines used in a phase-steered system, only longer.

While the actual implementation of the two steering techniques can thus be very similar, there are dramatic differences in performance (see Section 8.3.1). Modulo- 2π phase steering has a couple of very serious drawbacks. Firstly, whenever the beam is steered off-boresight, its direction becomes frequency dependent. Secondly, the time-domain shape of the driving signal is not preserved in the emitted EM wavefront. In other words, phase steering is *dispersive* both in the frequency domain and in the time domain; it can be shown that attempting to use short pulses on a large modulo- 2π phase-steered array will simply not work.

By contrast, in a time-delay-steered array is *non-dispersive* both in the time domain and in the frequency domain; the beam direction is frequency-independent and the time-domain shape of the transmitted signal is preserved. It follows from the reciprocity theorem that a time-delay-steered array is non-dispersive also when used for receiving.

In the 3D Performance Specification document, the active array the field-of-view is specified as a cone centred on the vertical and extending to 40° zenith angle. At the same time, the array bandwidth must be $> 5\text{MHz}$ to match the transmitter bandwidth. The only viable way to satisfy these requirements simultaneously is to use true time-delay-steering. Assuming that the received signals are digitised at the individual array elements, each element data stream is first delayed by a specific amount. All delayed streams are then added; the result is a single data stream representing a focussed beam. By applying different sets of delays to the same data, simultaneous beams in different directions can be formed; the number of beams will be limited only by the available processing power.

However, to meet the specified beam pointing resolution requirement of 0.625° in each of two orthogonal planes, the time delays would have to be programmable with a

resolution and stability of better than 12 picoseconds (less than $3 \cdot 10^{-3}$ of a period of the approximately 240-MHz radar carrier frequency) over a range of more than 350 nanoseconds! How to best implement a freely programmable delay with this resolution in practice is a fundamental and extremely challenging design problem. But more importantly, the 12-picosecond ultimate resolution forces very stringent demands on the purity and stability of the system reference time and frequency. It also has direct implications for the choice of receiver design and architecture, such that all these areas must be considered together.

5.4 Array systems and element antennas

To the best of our knowledge, no VHF phased arrays as large as those of the projected 3D system have ever been constructed in the sub-Arctic before, and so there exists very little data on what to expect in terms of weather and environment effects on array performance and survivability. Protecting the arrays from precipitation and other external disturbances by enclosing them in radomes is out of the question due to the sheer size of the arrays (up to 150 m diameter), so both the element antennas and the array structure must be designed to be so intrinsically tolerant to climatic influences that radar operation can continue essentially 24/7 all around the year with minimal need for manual intervention, preferably even under the worst imaginable weather conditions likely to be encountered in the deployment area.

In this respect, the array element antennas pose a number of particularly difficult design challenges. They must be sufficiently broadband to allow full-bandwidth transmission and reception of both up- and down-shifted plasma lines from a full bandwidth transmitter pulse, they must exhibit the maximum directivity commensurate with the specified field-of-view of each array and they must be capable of generating any arbitrary polarisation, *but they must also maintain all these characteristics even when loaded with large amounts of snow and ice*. Since the arrays will be remotely located, they must be designed to continue to operate even in the absence of regular snow clearing; the array structures must therefore accommodate large amounts of snow (ideally a whole winter's precipitation) without becoming adversely affected.

5.5 Adaptive Faraday rotation compensation

On its way from the scattering region to a remote receiver, the scattered signal is modified by the phenomenon known as *Faraday rotation*. If not properly compensated for at the receiver end, this can cause a loss of statistics by a very large factor and is therefore a serious issue.

In a multi-static incoherent-scatter radar system like e.g. the EISCAT UHF and the planned EISCAT_3D system, a radar transmitter illuminates the ionospheric plasma with strong RF pulses. Signals scattered off the plasma electrons are then received both by a receiver co-located with the transmitter and also at a number of outlying receiving sites located some 150 – 250 km from the transmitter site. Normally, the polarisation of the transmitted wave is *circular*, i.e. as the wave passes a fixed point in space, its **E** vector is seen to rotate in the positive sense around the propagation

direction with angular frequency $= 2\pi f$, while its magnitude remains constant. Signals backscattered towards the transmitter will therefore also be circularly polarised, but with the opposite handedness – their \mathbf{k} vector is exactly reversed in the scattering process.

Scatter signals picked up at the receiving sites will however have had their \mathbf{k} vectors turned through an angle $< 180^\circ$ and are therefore *elliptically* polarised, i.e. their \mathbf{E} vectors trace out ellipses in planes normal to the respective \mathbf{k} vectors. Obviously, the obliquity and orientation of these polarisation ellipses can be computed *a priori* if the scattering geometry is known.

In order for a remote receiver to recover all power contained in the scattered signal, its polarisation response must match the polarisation of the wave intercepted by the receiving antenna aperture (cf. the standard matched-filter case!). To accomplish this in practice, the signal is received on two co-located sets of linearly polarised antennas with their polarisation planes oriented at right angles to each other. In this way the incoming wave-front is decomposed into two orthogonal vector voltage components. These are then added with appropriate amplitude ratio and phase so as to maximise the signal-to-noise ratio.

If the propagation from the scattering point to the receiver were in vacuum throughout, the scattered wave would retain its original polarisation and the solution to the matching problem would be trivial. However, since the actual propagation path is in the ionosphere and therefore contains a time-varying amount of plasma, the scattered wave is subject to time-variable Faraday rotation, i.e. its polarisation ellipse will rotate through an angle approximately proportional to the total number of electrons along the path.

Faraday rotation is strongly frequency dependent; the rotation per unit electron content suffered by a propagating wave of frequency f varies as f^{-2} . Under typical ionospheric conditions, the total rotation of a wave with angular frequency ω , propagating from the ground up to a scattering point at 300 km altitude in the ionosphere above Tromsø, Norway, and from there to a receiver at Kiruna, Sweden, is approx. $(0.3 \dots 3) \cdot 10^{18} \omega^{-2}$ radians. At the EISCAT UHF frequency, 928 MHz, this corresponds to only a few degrees and the phenomenon is therefore essentially a non-issue for a UHF system.

However, at the EISCAT_3D frequency, 235-240 MHz, Faraday rotation constitutes a major challenge; the total rotation can be anything between 90 and 270 degrees, such that more than half of the available signal power and a factor of four or more in time resolution may be lost if the receiver response is not matched to the actual orientation of the polarisation ellipse!

Since the total number of electrons along the path from the scattering point to the receiver is an unknown quantity and the receiving system may introduce time-dependent phase and amplitude biases, the polarisation state of the received signal must be regarded as unknown. The only workable way to configure a remote receiver for maximum signal recovery is then to apply an adaptive scheme, where the optimum polarisation setting is continually estimated from the received signal itself. Nothing like this has ever been attempted at any of the world's other incoherent-scatter

installations, and a successful EISCAT_3D implementation would therefore represent a real breakthrough, technically but also scientifically, as estimates of the total electron content along the un-illuminated propagation paths from the scattering regions to the receivers would be continually generated as a by-product of the optimisation scheme, so providing users with one additional observable per beam.

At least five simultaneous beams will be generated at each receiving array. Because these will all be pointing in different directions, Faraday rotation compensation must be applied to each beam separately. Fortunately, the polarisation time variation is orders of magnitude slower than the sampling rate – changes of more than a degree per second are rare. It should therefore be possible to employ even complex and computationally intensive optimisers.

The most challenging aspect of the Faraday rotation compensation problem is however that the scatter signals can be extremely weak; at the low end of the signal dynamic range the effective signal-to-noise ratio may be below 1%. Polarisation estimates determined from the actual samples may then be so noisy as to be practically worthless for correction purposes. But as it is precisely under those conditions of almost vanishing signal that every last bit of signal power must be recovered, *robust and stable state-of-the art statistical prediction and optimisation techniques would have to be employed*, combining time averages of the sample-by-sample data with appropriate *a priori* estimates of the most probable amount of rotation as a function of time of year, time of day, solar activity etc.

6 Methodology

When the opportunity to apply for EU FP6 support for a design study presented itself in 2005, it was realised that even though the ideas for a new incoherent scatter radar system had originated in the EISCAT community, the Association's resources were not sufficient to take on a project of this magnitude single-handedly. While EISCAT could provide world-class expertise in the important areas of high power RF, antennas and system control and monitoring, the Association was weak in other areas, notably interferometry, digital signal processing and data management. Contacts were therefore taken to the EISCAT group at the University of Tromsø (UiT), Norway and the EISCAT support group at Rutherford-Appleton Laboratories (RAL) in the U.K. The Tromsø group had successfully implemented radar interferometry techniques at the EISCAT Svalbard radar and was expected to play a key role in the effort to include interferometry as a routine capability, while the RAL group could bring its extensive experience of maintaining a large database of EISCAT data and supporting analysis and visualisation software into the project.

This still left the digital signal processing area uncovered. Informal contact was now made with the "Embedded Internet Systems" group (EISLAB) at the Department of Computer Science and Electrical Engineering, Luleå University of Technology (LTU), Sweden. This team, which had no previous connections to EISCAT or any other ground-based space research activity, but was known to be highly competent in both component and systems design as well as in signal processing matters, eagerly accepted the challenge, so enabling a consortium of four Parties to be formed, with EISCAT as the Coordinator.

Because the application for FP6 support had to be accompanied by a Project Plan, it was necessary to make a first cut at structuring the work already in the course of preparing the application, before all relevant user input was available. Starting from a block diagram of a generic radar system, clearly distinguishable (and supposedly independent) sub-systems like e.g. the transmitter, the receiver, the antenna and the data processor were identified. A Work Package (WP) was then constructed for each sub-system and added to the plan. Insofar as possible, WPs were set up to be self-contained and with a scope matching the expertise and skills available to a specific Party, which was then assigned the primary or full responsibility for their execution. The EISCAT Technical Director was appointed Technical Project Leader (TPL) with global responsibility for all technical aspects of the study, and a Steering Group, composed of the group leaders from the different Parties and chaired by the TPL, was set up.

While a credible framework for the project had thus been established, it was recognised that this would only go so far; if the FP6 application should succeed, the tasks of the individual WPs and the interfaces between different WPs would need to be specified in quite some detail and aligned against a global, detailed target performance specification to enable each Party to proceed more or less on its own. Only in this way would the project team stand a realistic chance to deliver a coordinated, workable design by the end of the Project.

Logically, because the Design Study was being undertaken for the benefit of the scientific user community, the target system performance specification should reflect

its requirements and wishes insofar as possible. Since the design and development work could only start in earnest once these were known, a separate Work Package, WP 2, was defined with the single task of establishing current and prospective future users' requirements for a next-generation research radar and condensing these into a technical performance specification document, which would then be used as the baseline for the work in all other WPs.

Finally, a Management Work Package, WP 1, was set up. An EISCAT Headquarters responsibility, it was to operate for the duration of the Design Study and handle top-level project management, budgetary and cash-flow control and administrative contacts to Brussels.

In this manner, the project was initially subdivided between the following twelve WPs:

(a) WP1	Management	EISCAT
(b) WP2	Evaluation of design performance goals	ALL
(c) WP3	Active element (transmitter) options	IRF
(d) WP4	Phased array receiver design	LTU, EISCAT
(e) WP5	Interferometry	UIT, EISCAT
(f) WP6	Active element design	IRF
(g) WP7	Distributed control and monitoring	EISCAT
(h) WP8	Data archiving and distribution	RAL, EISCAT
(i) WP9	Signal processing	EISCAT, IRF, UIT
(j) WP10	New developments	UIT
(k) WP11	Implementation blueprint	EISCAT
(l) WP12	Time and frequency	EISCAT

Formal definitions of the scope and tasks of each WP can be found in Annex 1 to the EISCAT-EU FP6 contract [DSD].

Upon the start of the Design Study on May 1, 2005, the TPL and the WP 2 project group set about to prepare a Questionnaire which was then circulated to a large number of scientists and research groups known to have an interest in the plans for the 3D system. The Questionnaire was organised in three sections. "General Questions" were included to give an idea of the potential size and makeup of the user community for the new radar. In the "System Specification" section, recipients were asked to quantify what level of spatial and temporal resolution and what field-of-view would be required at different altitudes to enable them to advance their research in a decisive way. Finally, in the "Data Products" section, they were asked to rate the relative importance of having different types of data available in real-time. In early June 2005, the Questionnaire was e-mailed to some 40 individual scientists known to be past, present or potential new EISCAT users. A second mailing to an extended group of ionospheric scientists was made in early August 2005. When the process was closed in mid-September, 2005, a total of 22 replies had been received. Several of these were collective replies from entire research groups, so the response reflects the views of about 40 individuals.

Preliminary results from the Questionnaire were presented to the scientific community in an EISCAT_3D Performance Specification meeting on September 12-13, 2005 at

the Swedish Institute of Space Physics, Kiruna, Sweden. Immediately thereafter, the 3D Steering Group worked out and agreed on the technical baseline for the remainder of the project, based on the user response, and on the detailed definition of the resulting work and its distribution among the individual WPs. A formal Baseline Document, “D2.1 Radar Performance Specification” [D2.1], was finally published on November 1, 2005. This document contains a statistical, de-personalised breakdown of the user responses as an appendix.

Ever since the first tentative ideas for the 3D system were being formulated, it had been understood that several of its novel performance features, like e.g. the digital beam-forming, the multi-beaming and the polarisation tracking, were strongly dependent on the successful implementation of practical solutions to many of the issues highlighted in the “Design challenges” section, e.g. the generation and distribution of extremely low jitter, high-precision clock signals throughout the array and the development of robust statistical optimisation algorithms. In the process of working out the Baseline Document, the Steering Group noted that in these respects the Project was about to start navigating uncharted waters. In particular, it was not known if there existed stable and reliable beam-steering and beam-forming solutions that could be expected to operate in the climatic conditions encountered at the planned sites, or indeed if any useable approaches existed in the first place.

Also, it was obvious that no funding agency would be prepared to put up the amount of money required for the construction of a 3D system unless it were convinced that not only was there a sound science case for the system, but also that a system delivering the required performance could really be built. The Steering Group therefore saw the primary task of the Design Study as proving that the cornerstone concepts of the envisaged system really held and that a system with the required performance could be constructed using existing technology. For these reasons, it was felt that all performance estimates and demonstrations of array performance based on computer simulations, no matter how promising, should be validated by actual measurements performed before being accepted as workable.

For this purpose, a small 224-MHz receive-only array (the “Demonstrator”, see Sections 10 and 11) was constructed at the Kiruna site. Financed entirely by EISCAT funding and operated in coordination with the VHF (224-MHz) and HF Heating systems in Tromsø, this “proof-of-concept” array would serve as a test-bed on which mission-critical concepts, hardware and software, including but not limited to the direct-sampling receivers, the digital beam-forming system, the adaptive polarisation tracker and the coherency trigger software, could be tested, debugged and validated under realistic conditions using real radar returns from the ionosphere. While a future target system receiver array would probably not be a one-to-one copy of the Demonstrator, the receiver, the beam-former, the data transmission systems and the real-time software developed for it would be directly re-usable if desired; some results could even be expected to be of direct relevance to the design of the active (“Core”) array.

Having decided on this course of action, the Demonstrator became the *de facto* focal point for much of the subsequent work. All effort in WP 4 and major parts of WP7, WP 9 and WP 12 were re-focussed onto developing the actual receiver, beam-former and optical data communications hardware and real-time software required to make

the Demonstrator into a fully working receiving array with the capability to generate at least three simultaneous beams with freely programmable elevations and transfer the raw data into the EISCAT data processing environment. The intention was that once debugged and commissioned, the array would be maintained in regular operation for a considerable time to gain experience of how stable the beam-forming really was in practice. Being able to access the raw data in real-time, as well as being able to retrieve long sequences of recorded raw data after the fact, would also enable members of the WP5 and WP9 groups to validate and optimise the polarisation tracking algorithms and the coherency trigger software on real ionosphere echoes.

In early 2008, the Study had progressed to a point where the need to consider not only technical details but also logistical, environmental and legal aspects of the 3D system had become pressing. It was realised that this was an area that had clearly been overlooked when the initial Project Plan was formulated, such that there was no logical place for the required actions and no resources set aside for their completion. The Steering Group therefore proposed that the Project Plan should be expanded to include a new Work Package:

- WP13 Enabling actions (Spectrum etc.) EISCAT/IRF

One of the most important tasks for the new WP would be to continue the dialogue with the frequency management authorities in the three Nordic countries, with the goal of obtaining a legally binding and firmly protected spectrum allocation before or by the end of the Study. Resources would be made available through re-distributing some of the unused effort and funding allocated to other Work Packages.



Figure 6.1: The 48-element Demonstrator array at the Kiruna EISCAT site in a winter setting, November 2007.

At the outset, some 25 individuals were associated with the Design Study to various degrees. Changes to the composition of a group this size will inevitably take place over a four-year period. Some individuals may retire; others may find new employment and resign from their positions. In both cases, the lost manpower must be replaced quickly lest the work be held up or stand still; this may turn out to be very difficult if the individuals leaving the team happen to be key persons possessing unique expertise and skills.

The Design Study has been faced with precisely this latter problem on a couple of occasions. In the summer of 2007, a key individual on the UiT/WP5 team announced that he was leaving the science field for a position in industry. The UiT team nevertheless managed to fully meet its commitments in WP5 through a combination of great efforts by the other group members and highly appreciated external help from fellow scientists working at the Jicamarca Radar Observatory in Peru.

In November 2007, EISCAT's most experienced RF engineer left the Association to take up a position with the Norwegian civil air traffic administration. This was a potentially even more serious loss, since this person had the expertise needed for designing and constructing a test bed for the evaluation of a range of RF power semiconductor devices with special regard to their suitability for long-pulse operation (known to be an especially very difficult operating regime). He was also supposed to carry out the actual test program, including a long-term, full-power survivability test stretching over many months. The successful completion of these tasks was a key component in the overall transmitter design, which required that this expertise should be replaced promptly – but none of the four parties could offer an individual with the required skills.

Fortunately, the Kiruna branch of the Swedish Institute of Space Physics (IRF) operates several atmospheric radars and had among its staff several very senior and experienced radar engineers with long experience of both EISCAT and other radar systems. Furthermore, the EISCAT_3D TPL, Dr. Wannberg, had for some years held a fall-back position as an IRF staff scientist, from which he had been on open-ended leave of absence during his tenure as EISCAT Technical Director. This combination of circumstances was now exploited. Following negotiations between EISCAT and IRF, an IRF EISCAT_3D project group was formed at the institute, the full responsibility for WP3 and WP6 and large parts of WP9 was transferred from EISCAT to the IRF, Dr. Wannberg returned to his IRF position on March 1, 2008 to head up the institute's EISCAT_3D activities and a very experienced IRF staff member was appointed as the leading engineer and charged with completing the transmitter work. These changes were later formalised through a decision by the highest governing body of the EISCAT_3D project, the General Assembly, and duly approved by the FP6 Project office, making the IRF the fifth full party to the EISCAT_3D project. From this point onwards, progress in WP3 and WP6 has been excellent.

The accession of IRF to the project also allowed the EISCAT_3D team to have access to a highly skilled FPGA designer. His services were enlisted in WP9 and he duly came up with a practical design for the Demonstrator beamformer, something which had eluded the project team up until this time.

Many of the most important contributions to the design work have been made by individuals who have had different parts of the project assigned to them as their Ph.D. or M.Sc. thesis subjects. In this respect, the extensive NEC-based design of element antennas for the Core array and the computer simulations of overall array performance performed by Toralf Renkwitz of the University of Rostock deserve special recognition. These results, which will be of fundamental importance to the final design, were obtained by Mr. Renkwitz as his M.Sc. thesis task without any FP6 support but with full support from the Institut für Atmosphärenphysik (IAP), Kühungsborn, which is hereby gratefully acknowledged.

As far as the project team is concerned, all key issues and major design challenges identified at the outset or in the course of the work have thus been addressed in a satisfactory manner. On the other hand, some of the things included in the initial project plan have not been addressed in the manner foreseen when the application for funding was being formulated.

One of these is the promised comprehensive system design, complete with blueprints and worked out to a level of detail that would have allowed it to be used as the basis for an RFQ procedure. There is a logical reason for not pursuing this goal more actively. No funding for the target EISCAT_3D system is yet available and none is likely to appear in substantial amounts for at least another two to three years. Since many of the electronics components to be used in the construction of the system are in a state of rapid, continual development, any design completed today in this area would be overtaken by events and probably even obsolete before construction became possible. Much of the effort spent on preparing a detailed design at the present time would therefore be wasted. What we have done at this stage is to verify that the required performance goals are achievable, while leaving the precise selection of components to be specified.

Considering the size and scope of the full 3D project, the project team is of the opinion that, if a turn-key solution to the procurement of the new radar is required, then partnerships with major players in the telecom arena (Nokia, Ericsson or similar) will be needed. There may also be scope for participation by consortia of smaller, local industries in the production of certain subsystems. Since the large telecom companies already possess vast amounts of experience in the area of constructing base stations for mobile telephone systems, and understand the practical issues relating to weatherproofing and installation of complicated electronics systems in the field, a pragmatic way to approach the purchasing exercise might be to follow up the present study with a second step, where a team of skilled antenna, electronics and signal processing engineers would work out a set of highly detailed performance requirements and complement these by constructing “performance prototypes” of the radar hardware, i.e. reference units exhibiting or exceeding the minimum required performance in all critical respects. An RFQ could be then issued to the telecom industry throughout Europe, asking for quotes on a system meeting or exceeding the technical specifications but prescribing little or nothing about how bidders should go about meeting them; this would leave the door open for each company to leverage its particular strengths.

This page is intentionally blank.

PART B
SYSTEM DESCRIPTION



7 System Overview

The early view of the EISCAT_3D system was essentially conceptual. It was to a large extent a natural outgrowth on earlier proposals for the direction of the Association after 2006 (when the Agreement would be up for renewal) generated by the Futures Committee and the Director. These had identified a number of areas where improvement or renewal of the mainland radar systems would be required in order for EISCAT to remain at the forefront of incoherent-scatter research for at least another decade. In particular, the reinstating of VHF field-aligned operation capabilities was regarded as mandatory and a much greater beam-steering flexibility at VHF than that offered by the existing reflector antenna was needed. Multi-static, multi-beam receiving capabilities and greatly improved spatial resolution were also very high on the list.

The system configuration envisaged at the time of the FP6 application (Figure 7.1) was one of several possible approaches to addressing these needs. It would comprise at least the following components:

- (a) One or several, as yet unspecified high-power VHF transmission system(s), possibly re-using the existing parabolic-cylinder VHF antenna and klystron transmitter at Tromsø but also including a phased array. Both high-power klystrons and distributed transmitter solutions would be considered,
- (b) At least three receiver systems, one of which would be co-located with the transmitter facility and the other two being receive-only (“remote”) sites. Phased arrays, capable of simultaneously receiving scattered signals from a number of ionospheric volumes along the radar beam, would be employed at least at the remote sites,
- (c) Two or more interferometry receiver systems, located close to the transmitter to support measurements of spatially restricted phenomena at resolutions better than the dimensions of the transmitted beam, and
- (d) Data networks and distribution systems allowing remote control of the facility and effective distribution of both raw and processed data and results.

It was however recognised that this configuration was unlikely to remain unchanged throughout the study. As the Project got underway, this quickly proved to be true. The current view of the EISCAT_3D system configuration, shown in Figure 7.2, is the result of four years of painstaking detail-level analysis of the pros and cons of different approaches to realising the individual subsystems and integrating these into a unified system.

Some of the most important differences between the initial and final concepts are:

- (a) The active phased array will be circular rather than square,
- (b) Only a distributed, solid-state transmitter system will be considered,
- (c) The Tromsø VHF reflector antenna will be fitted with solid-state transmitter modules and initially used in a low-power mode for test purposes as well as for middle atmosphere work; it will however not be an integral part of the final 3D system,
- (d) To improve the remote site array aperture usage, two remote sites per baseline will be deployed; one of these (the “close-in” one) will be dedicated to low-altitude work and the other one to F2 / topside work,
- (e) Each active array element will be provided with its own digital signal generator/exciter. In this way, both the transmitter beam-steering function and a possible aperture tapering function can be absorbed into the exciter, thus eliminating the need for a separate modulator and phase control system.




Figure 7.1: Conceptual-level block diagram of the EISCAT_3D radar system as envisaged at the time of application for EU FP6 support for the Design Study in 2005.




Figure 7.2: Revised EISCAT_3D concept, April 2009.

8 Principle Design Areas

In the next two sections of this report, we discuss in detail the various design considerations which are needed to specify the EISCAT_3D radar facility. Rather than breaking down the description according to the different work packages of the design study, we take the approach of describing the various sub-systems of the eventual radar. Although these broadly follow the division of work packages, there are instances where the sub-system design has resulted from a number of work packages collaborating together. In each section, we identify the work packages which have contributed to a particular area and the reports and deliverables where more details can be found. Section 16 provides a complete list of these design study outputs, together with the URLs at which they can be found on the project web site.

We have divided the consideration of the EISCAT_3D sub-systems into two parts, reflecting the way that the study has turned out in practice. As noted earlier, it was found as the study proceeded that some areas turned out to be more demanding than had originally been foreseen and, as a result, most of our effort has been concentrated on those areas. These are the “principle design areas” discussed in Section 8. In Section 9, we discuss the other design areas which, while still very important to the design of the integrated radar system, have occupied less time in the design study phase, either because they turned out to be less demanding than originally envisaged, or because it was recognised from the outset that these areas were less challenging.

Both sections are ordered approximately according to signal path, i.e. reflecting the journey taken by a signal from the time that it enters the system through the antenna array, through the various detection and signal processing stages to its eventual storage in the long-term archive.

8.1 Antenna design

8.1.1 Element antenna performance requirements

The 3D array element antennas must satisfy a number of partially conflicting performance requirements. They must be sufficiently broadband to allow full-bandwidth transmission and reception of both up- and down-shifted plasma lines from a full bandwidth transmitter pulse, they must exhibit the maximum directivity commensurate with the specified field-of-view of each array but at the same time have as little gain as possible towards the horizon, they must be capable of generating any arbitrary polarisation and they must be sufficiently insensitive to climatic effects to maintain their characteristics even when loaded with snow and ice.

We define antenna bandwidth as the frequency range over which the magnitude of the antenna feedpoint complex reflection coefficient $|s_{11}| < -20$ dB. This definition is appropriate for a transmitting antenna, since for the transmitter to deliver its rated output power the antenna impedance must be maintained inside a fairly narrow range. Obviously, the transmitted spectrum can be no wider than the frequency band allocated to the system by the authorities (229.928 – 236.632 MHz in the present case). This is therefore also a lower limit on the antenna bandwidth, but more will be

helpful in ensuring that the antenna maintains its basic performance also when loaded with snow and ice. When used for reception, the impedance-matching requirement is less stringent, which is a great help for the designer.

In the EISCAT_3D Performance Specification Document [D2.1], the system field-of-view (FOV) is specified as follows:

“The beam generated by the central core transmit/receive antenna array will be steerable out to a maximum zenith angle of 40° in all azimuth directions. At 300 km altitude, the radius of the resulting field-of-view is approximately 200 km. In the N-S plane this corresponds to a latitudinal coverage of $\pm 1.80^\circ$ relative to the transmitter site. The antenna arrays at the 3D receiving facilities will be arranged to permit tri-static observations to be made throughout the central core FOV at all altitudes up to 800 km.”

This requires that the full opening angle of the Core antennas is at least equal to the maximum zenith angle, i.e. $\pm 40^\circ$. At the receive-only sites, antennas with an opening angle of $\pm 30^\circ$ can be used. As nothing is said in the PSD about how much gain drop can be accepted at the FOV limits, we have – somewhat arbitrarily – fixed this at the -1 dB level.

The basic electrical performance specifications can now be summarised as follows:

- Centre frequency 235.0 MHz
- $S_{11} < -20$ dB bandwidth $> \pm 6$ MHz
- -1 dB opening angle $\pm 40^\circ$ (core array elements), $\pm 30^\circ$ (receive-only array elements)
- Arbitrary polarisation
- Gain at angles $> 75^\circ$ off-boresight as low as possible, ideally at least -16 dB

8.1.2 Environmental considerations

Weather conditions in the 3D deployment area are a challenge to the antenna designer. In the wintertime, snow and/or ice collecting on the antenna elements will lower the feedpoint impedance and cause a downward shift and narrowing of the passband. The directional pattern will also be affected to some degree. To ensure that the antennas remain within specifications over the nominal design bandwidth, they must therefore be designed with a substantial amount of extra bandwidth on the high-frequency side of the passband.

Maintenance and survivability aspects are extremely important. At the envisaged 3D Core site in northern Norway, close to the present EISCAT Ramfjordmoen installation, the weather can change from sub-Arctic to maritime in a matter of hours. In midwinter, a thick layer of hoarfrost build up on exposed objects during long cold periods; later in the season, blizzards depositing up to a meter of wet, heavy snow in a day frequently hit the area. The Core radiator elements must be mechanically sturdy so as not to deform or collapse even under such weather conditions.

The receive-only sites, being located farther inland, will be much less exposed to severe icing and wet snowfall than the Core site, allowing also physically large medium gain antennas, e.g. (6+6) element X Yagis, to be considered from the survivability point of view. On the other hand, at these locations long periods of very cold weather with attendant accumulation of large amounts of hoarfrost will be quite frequent. If the antennas are too narrowband to begin with, the resulting de-tuning can be so severe as to render the array useless, as demonstrated by measurements performed on the thin-element Demonstrator Yagis and the Demonstrator array [Lin08]. Again, a successful design must provide a substantial amount of extra bandwidth on the high-frequency side.

8.1.3 Antennas for the core array

Each Core element antenna must provide a minimum of 7 dBi gain in order to match the array inter-element distance of 0.7λ . It must also be capable of generating any arbitrary polarisation. Two antenna types that meet these criteria are an X dipole-reflector combination and a short X Yagi antenna.

After carefully considering both electrical performance criteria and practical aspects, the X dipole/groundplane alternative has been eliminated in favour of the Yagi. This is primarily because a groundplane is perceived as a potential source of mechanical design complications and serious operational problems. By constructing it from wire netting it can be made lightweight and nominally transparent to precipitation, but for it to approximate an electrically continuous groundplane it must cover essentially the whole array aperture and the mesh size can be no larger than approximately 0.06λ . At 235 MHz, this is ≈ 70 mm. With a mesh size this small, snow and hoarfrost accumulating on the mesh wires will soon clog the mesh and then the whole reflector area will start to collect precipitation. To prevent the surface from becoming permanently deformed by the accumulated snow, sturdy, thick-wire netting must be used, offsetting much of the weight and cost gain obtained by using netting in the first place.

Eventually, the accumulated snow must be removed. This can probably only be done manually, which conflicts with the 3D design goal of essentially unattended operation. To enable personnel to get at the upper groundplane surface, access hatches must be installed at intervals of a few meters; these will add to the cost and complexity of the array structure.

Short-boom, thick-element Yagi antennas should handle the winter conditions much better. Employing large-diameter parasitic elements, a three-element Yagi can produce the required 7 dBi gain over up to 10 % relative bandwidth. Under conditions where heavy snowfall would clog a mesh groundplane completely, an array of such Yagis remains virtually transparent to the precipitation. If extraordinarily bad weather should occasionally require manual intervention in order to protect the structural integrity of the array, snow clearing can be performed from the ground with simple tools (e.g. long brooms).

Having decided in favour of an X Yagi element, the next step was to investigate the electrical and optical characteristics of some Yagi configurations by numerical simulation, in order to find one or more designs meeting or exceeding the PSD gain,

bandwidth and polarisation purity requirements, while at the same time exhibiting relative insensitivity to weather effects. Several Yagis would then be built according to the most promising design and validated by actual measurement. Their performance as element radiators in very large arrays would also be simulated and analysed using a combination of NEC and Matlab.

The simulation work was performed by Mr. Toralf Renkwitz of Universität Rostock, Germany as his M. Sc. thesis project. Mr. Renkwitz had substantial previous hands-on experience of using the well-known NEC electromagnetics software suite to simulate and analyse Yagi antennas. Recognising that his results would be of direct interest to the instrument development conducted at the Institute of Atmospheric Physics (IAP) Kühlungsborn, the IAP kindly put its computing facilities, including its NEC software installation, at his disposal free of charge. Mr. Renkwitz successfully defended his M. Sc. thesis in June 2008.

The end product of the simulation exercise is a (3+3) element X Yagi (the “Renkwitz Yagi”) meeting nearly all the performance requirements listed for the Core array radiator. This design uses straight 20 mm diameter tubular elements, a folded dipole radiator and a half-wave transmission line balun [Bal82]. Its mechanical layout and dimensions and its predicted electrical and optical characteristics are listed in Table 8.1.1 and Table 8.1.2.

To validate the simulation results, a full-size Renkwitz Yagi employing tuneable dipoles has been constructed to the dimensions given in Table 8.1.1 (Figure 8.1.1). The elements are passed through the 30-mm square boom and welded to it on both sides to create a sturdy structure and ensure positive electrical contact. In this configuration, the boom exerts a shortening effect on the elements; these must therefore be lengthened by a few mm [Ren07].

Element	Mechanical length	Position
Reflector	657 mm	0 mm
Folded dipole	550 mm	+ 228 mm (centre)
Director	484 mm	+ 410 mm

Table 8.1.1: Renkwitz Yagi element lengths and positions

Impedance	200 + j0 Ω
Gain	7.11 dBi
Main lobe FWHM	89°
Attenuation @ 40°	2.52 dB
Attenuation @ 75°	7.79 dB
Front/back ratio	20.1 dB
Bandwidth ($s_{11} = -20$ dB)	(-12.75 / +11.25) MHz
Bandwidth ($\Phi_z = 5^\circ$)	(-9.00 / + 6.75) MHz

Table 8.1.2: Predicted electrical and optical characteristics of the Renkwitz Yagi.

Complex impedance measurement data obtained after fine-tuning the dipole shorting bar positions (*Figure 8.1.* and *Figure 8.1.*) are in very good agreement with the simulated performance. It is particularly gratifying to note that 24-MHz bandwidth can be achieved in practice. A 19-antenna test bed for validation of the mutual

coupling behaviour [Ren07, pp. 49-62] is under construction. Measurement data from this setup will be published separately.



Figure 8.1.1: *The first 235-MHz Renkwitz Yagi, constructed to the mechanical specifications given in Table 8.1.1 under typical midwinter conditions. In this beta version, the folded dipoles are made with movable shorting bars to allow a limited amount of fine-tuning. The production run antennas will have their dipoles made from formed continuous lengths of aluminium tubing. The blue cables that are strapped down onto the boom are half-wave baluns transforming the nominal 200 ohm balanced feedpoint impedance to 50 ohms unbalanced.*




Figure 8.1.2: Smith diagram plot of the measured terminal impedance Z of the front half of the Renkwitz Yagi evaluation model. The predicted reactance sign reversal at mid-band (the “cusp”) is clearly visible next to marker # 3. The reference plane for this measurement is at the end of a piece of 50-ohm coax, extending backwards from the feedpoint to a point approx. 30 cm behind the reflector element, so the impedance vs. frequency trace appears rotated relative to the predicted one.




Figure 8.1.3: Reflection coefficient s_{11} magnitude vs. frequency for the front half of the evaluation model Renkwitz Yagi. The measured -20 dB s_{11} bandwidth, (- 11.4 / + 14.0) MHz slightly exceeds the predicted one but is shifted upwards in frequency by about 2.5 MHz relative to mid-band. The reason for this is being investigated.

The effects of ice on the feedpoint impedance and passband of the Renkwitz Yagi have been investigated in a separate series of simulations [Ren07, pp. 44-46]. The expected passband shape when the antenna is covered by a smooth 3 mm thick ice layer with $\epsilon_r = 5$ and conductance $G = 4 \cdot 10^{-4}$ S/m is shown in **Figure 8.1.** A downward shift of the passband is clearly visible; the $s_{11} = -20$ dB / VSWR = 1.22 points are shifted from 222.5 MHz to 220 MHz and from 246.5 MHz to 241.5 MHz, respectively.




Figure 8.1.4: Predicted voltage standing-wave ratio (VSWR) vs. frequency of the Renkwitz Yagi when covered by a 3 mm thick layer of ice with $\epsilon_r = 5$. The $s_{11} \leq -20$ dB / VSWR ≤ 1.22 passband is shifted downwards by close to 4 MHz relative to that of an ice-free antenna, but still meets the design target of $s_{11} \leq -20$ dB over (235 ± 6) MHz.

The effects of ice have also been investigated experimentally. A single-polarisation test antenna, sitting on its outdoor test stand at an ambient temperature of -15° C, was repeatedly sprayed with tap water to allow an ice layer to form on all elements, and the feedpoint impedance was measured for every 0.5 mm of additional ice thickness.

Figure 8.1.1 shows what happens when the whole antenna is covered by a 2.5 mm thick ice layer. The measured $|s_{11}|$ pass-band shape is very similar to the simulated results of Figure 8.1.1. The -20 dB s_{11} points have shifted downwards in frequency to approx. 216 and 239 MHz. Thus, while the ice-loaded antenna still exhibits a bandwidth far exceeding the required 12 MHz, it no longer meets the -20 dB s_{11} requirement at 241 MHz, the nominal upper band edge.

Since the actual -20 dB s_{11} bandwidth of a Renkwitz Yagi is almost 25 MHz, the series-production design could be re-optimised for a centre frequency about 2.5 MHz above the nominal band centre, i.e. 237.5 MHz. This should ensure that the upper passband edge stays well above 241 MHz even when the antenna is covered by a 2.5 mm ice layer. The associated gain loss at 235 MHz should be almost unnoticeable under normal conditions.

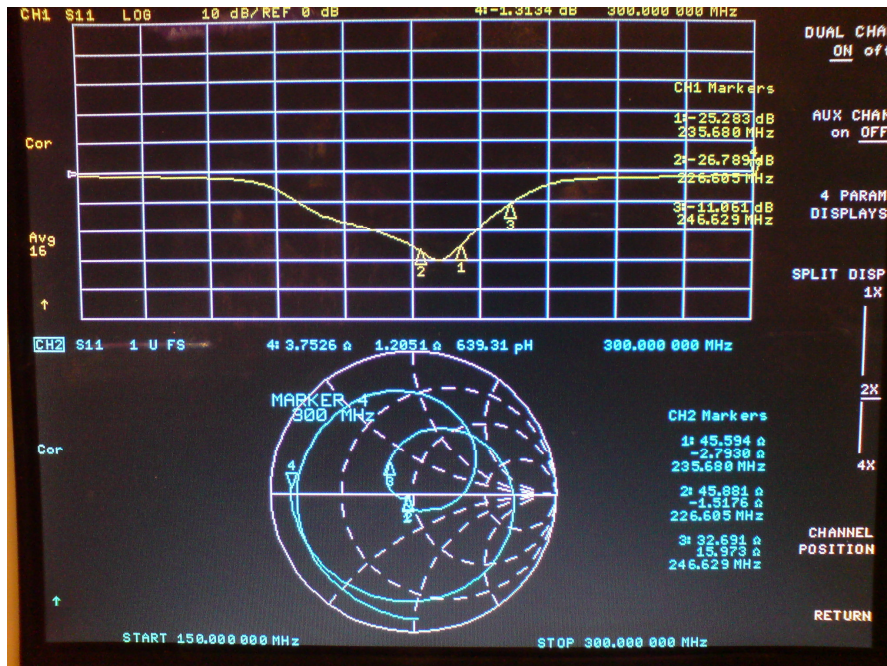


Figure 8.1.5: Reflection coefficient s_{11} vs. frequency for a Renkwitz Yagi covered by an approx. 2.5 mm thick ice layer, applied by repeatedly spraying water onto the antenna elements with a spray bottle. For this measurement the antenna was kept at an ambient temperature of -15°C . The $-20\text{ dB } s_{11}$ pass-band has now shifted to approx. (216 – 239) MHz and no longer meets the (235 ± 6) MHz band-width requirement on the high-frequency side.

8.1.4 Antennas for the Demonstrator Array

All proof-of-concept tests planned to be carried out with the Demonstrator Array (e.g. time-delay beam-steering, multi-beaming and coherency triggering) use the existing EISCAT 224-MHz VHF transmitter in Tromsø, Norway as the illuminator. Since the Demonstrator-VHF common field-of-view is very restricted (it is defined by the length and direction of the Tromsø – Kiruna baseline and the essentially non-existing E-W steerability of the VHF antenna), medium gain (11 – 12 dBi) antennas with fairly narrow main lobes can be employed to advantage, leading to the following antenna parameters.

- Centre frequency 224.0 MHz
- Bandwidth (defined as $s_{11} < -20\text{ dB}$) $> \pm 3\text{ MHz}$
- Relative gain -1 dB at 10° off-boresight, in all azimuth directions and over full bandwidth

In the U.S. there is an amateur radio frequency band at 222 MHz. Several U.S. antenna manufacturers carry Yagi antennas for this frequency, with gain and bandwidth characteristics matching our requirements, in their product lines. At our request, a California-based company, M2, Inc., re-optimised their 222 MHz seven-element Yagi design for 224 MHz. The resulting design, a (6+6) element X Yagi on a 130 cm long boom, labelled 224-6RM, has 11.6 dBi gain at 224 MHz, a -1 dB beam-width of $\pm 13^{\circ}$ and a bandwidth far exceeding our requirements. At the recommended stacking distance of 1.68 m or 1.25λ , only 48 of these antennas are required to fill the

planned 150 m² array. This antenna was therefore selected as the element radiator for the Demonstrator.

During the 2007/2008 winter season, the completed Demonstrator array exhibited severe de-tuning whenever covered by more than a few millimetres of frost. Since the 224-6RM is a relatively high-gain, narrowband design, this was not unexpected. To get a better understanding of the problem, extensive numerical simulations of the 224-6RM under various amounts of dielectric loading have been performed; systematic measurements on a single antenna in severe icing conditions have also been carried out [D3.2]. These demonstrate beyond doubt that under weather conditions expected to occur several times every winter season, the 224-6RM suffers a passband shift of more than 6 MHz. This negative but extremely valuable result underlines the importance of ensuring that the antennas eventually selected as receive-only array elements for the full 3D are sufficiently broadband and tolerant of weather effects.

8.2 Array design

8.2.1 Introduction

A very important part of the remit of Work Packages 3 and 4 is to establish the main design features of both the active element (“Core”) array and the receive-only arrays and to offer guidelines for how to realise them in practice. This work has been pursued on several fronts in parallel. Relevant parts of the system performance requirements (field-of-view, altitude coverage, time and altitude resolution, angular resolution etc.) have been translated into technical specifications of e.g. required array size and structure, radiator performance, transmitter power and practical arrangements, comprehensive analytical and numerical analyses of phased array characteristics and performance have been performed [Ren07; Mar06] and several important design, maintenance and cost-driver aspects of the Core array have been considered and reported on in Deliverable 3.2 [D3.2]. While the present summary is necessarily brief, the results and conclusions presented [D3.2], [D4.2], [D6.2] and in Section 10 of this document, are sufficiently detailed to enable a follow-on project to quickly proceed to a detail level design and costing. An executive summary can be found in [D3.2].

8.2.2 Array design procedure

Two of the fundamental performance specifications of any radar system, namely its field-of-view and its spatial resolution transverse to the beam, are entirely governed by the mechanical and optical characteristics of the used antenna. The transverse resolution is essentially determined by the antenna beamwidth, which to first order is inversely proportional to the transverse extent of the antenna aperture; this holds whether the aperture is a filled array or a reflector. Accordingly, if the angular resolution and operating wavelength are specified, it is straightforward to derive a zeroth-order estimate of the required array size.

The useable field-of-view of a plane phased-array is at most equal to the grating-lobe-free solid angle region centred on a normal to the array (assuming that it is immobile).

The size and shape of this region are determined by the inter-element distance and the element placement geometry. When designing to a set of performance specifications, both these parameters are manipulated until the grating-lobe-free region totally encompasses the required field-of-view; the design is then iterated until the desired result is achieved with a minimum number of elements.

A performance requirement of specific relevance to the incoherent-scatter radar case is the radar time resolution, i.e. the integration time required to estimate some observed quantity to a given statistical accuracy. This is a system-level parameter, which in addition to being a function of the antenna aperture size also depends on the available peak transmitter power, the modulation pattern used and the receiver noise temperature. A phased-array system is a special case in this regard: Since the transmitter is distributed across the array, each array element being equipped with its own power amplifier module, the total available peak power is the individual module peak power times the number of elements. The system power-aperture product, an often used figure-of-merit [D3.2, p.26] is therefore approximately proportional to the fourth power of the array diameter! This very strong dependence implies that for any given time resolution performance requirement and any given transmitter module power level, there exists a narrow “threshold range” of array sizes; an array smaller than the lower end of this range is essentially incapable of providing the desired time resolution, whereas one only slightly larger than the upper end of the threshold range will deliver substantially better performance than requested.

To start the design process, the initial set of fundamental performance requirements is first analysed according to the flowchart depicted in Figure 8.2.1. The designer may then find that some of them are unrealistic and/or mutually contradictory and need to be modified. He is also likely to find that the design is restricted by boundary conditions like e.g. limited real-estate, construction time deadlines and a funding ceiling. At some stage – and even before the detailed design can begin – the design process therefore almost inevitably mutates into a linear-programming exercise. This has proven to be true also of the EISCAT_3D array design.




Figure 8.2.1: Flowchart of the array design procedure used to establish zeroth-order estimates of several important EISCAT_3D antenna array parameters.

8.2.3 Core performance requirements and array parameters

Design target values for the 3D system field-of-view, transverse resolution, power-aperture product and coupled height/time resolution under different ionosphere conditions have been laid down at the beginning of the design study [D2.1]. They are repeated below for convenience:

- Centre frequency between 220 – 250 MHz, subject to allocation
- Peak output power ≥ 2 MW
- Power-aperture product > 100 GW m²
- Transverse resolution at 100 km altitude better than 150 m
- Grating-lobe free radiation pattern out to 40° zenith angle in all azimuth directions
- Electron density estimated to $< 10\%$ statistical uncertainty under the following conditions:

Altitude [km]	Electron density [m ⁻³]	T _e /T _i	Ion composition	Height resolution [m]	Integration time [seconds]
80	1 × 10 ⁸	1.0		≤100	30
100	3 × 10 ⁹	1.0		100	1
150	1 × 10 ¹⁰	1.0	50% NO ⁺ , 50% O ⁺	100	1
300	3 × 10 ¹⁰	2.0	100% O ⁺	300	1
800	3 × 10 ¹⁰	3.0	5% H ⁺ , 95% O ⁺	1000	10
1500	1 × 10 ¹⁰	4.0	10% H ⁺ , 90% O ⁺		60

Table 8.2.1: Assumed conditions for evaluation of statistical uncertainty

- Evaluating these specifications according to Figure 8.2.1 procedure, it was promptly established that an array extent of more than 1000 metres would be needed to satisfy the (150-meter@ 100 km) transverse resolution requirement [D3.2, pp.10-11]. Since a filled 1-km diameter array is a gigantic project (cf. the Square Kilometre Array), far exceeding any practical and financial boundary conditions, it was concluded that this particular design target must be met through an interferometric arrangement.

Next, a comprehensive numerical analysis of a number of different array geometries was performed [Ren07; D3.2 pp.11-14]. This demonstrated that a filled circular aperture, employing equilateral grid geometry with an inter-element distance of 0.7λ , can provide essentially grating-lobe free performance out to 40° zenith angle; full NEC analyses of arrays of up to 500 elements showed a very well-behaved sidelobe pattern. Based on these results, it was decided to adopt the circular aperture or equilateral-triangular grid geometry as the preferred geometry and apply it to the next step in the design process, i.e. determining the system time resolution as a function of array size. This analysis proceeded in two steps: First the beam-widths and power-aperture products of arrays were computed; in the next step, signal-to-noise ratios and required integration times were evaluated [D3.2 pp.26-30].

A $100 \lambda/120$ -m diameter, 16000-element array has a half-power beamwidth of $\approx 0.75^\circ$, i.e. comparable to that of the EISCAT UHF. Assuming a power level of 600 W per array element, in accordance with interim results from the Work Package 6 RF power device testing program, its power-aperture product will be 106 GW m^2 , i.e. right on the 3D design specification target and about one order of magnitude greater than that of the EISCAT VHF when operated in single-beam, dual klystron mode (Mode 1). It is therefore recommended as a suitable design target for the 3D Core.

Comparing the joint time/height resolution requirements laid down in [D2.1] to the level of performance expected from the 16000-element array was disappointing; they were found to be extremely unrealistic and it was established that even a 36000-element array would fail to meet them at all altitudes by a large factor (25...350)! However, *provided that the height resolution can be relaxed by a factor of 4...10 (from 100 m to 1 km at 150-km altitude and from 1 km to 4 km at 300-km altitude)*, the 16000-element array would easily meet the 1-second/10-second time resolution requirements. This is a consequence of the fact that in an incoherent-scatter

measurement in the ionosphere, the signal-to-noise ratio is proportional to the square of the vertical extent of each range gate, other things being equal.

If the 16000-element array cannot be constructed right away due to a shortage of funds, a 5000-element, 70-m diameter array would be a logical first step. Providing a 1.3° half-power beam-width, a power-aperture product of $\approx 10 \text{ GW m}^2$, and full steerability, it would exceed the performance of the current VHF system from the start and could then be expanded as additional funding became available. In any case, a basic set of receive-only outlier arrays for interferometry should be installed from the start to meet the horizontal resolution performance requirement.

The above considerations can be summarised as follows:

- The active (“Core”) element shall be a 120-m diameter, 16000-element filled circular aperture array with the elements laid out on an equilateral triangular grid and an inter-element distance of 0.7λ . This will have a half-power beamwidth of $\approx 0.75^\circ$, i.e. comparable to that of the EISCAT UHF. Its power-aperture product will be $\approx 100 \text{ GW m}^2$, i.e. about one order of magnitude greater than that of the EISCAT VHF.
- To get started with limited funding, a 5000-element, 70-m diameter array would be a logical first step. It would exceed the performance of the current VHF system, providing a 1.3° half-power beam-width, a power-aperture product of $\approx 10 \text{ GW m}^2$, and full steerability, and could be expanded as funding became available.
- In either case, a basic set of receive-only outlier arrays for interferometry should be installed from the start to meet the 150-m transverse resolution performance requirement.
- The receive-only arrays will be of approximately the same size as the Core and use the same basic geometry; due to the more restricted field of view required at these sites, the inter-element distances may be chosen slightly larger.
- The extreme joint time/height resolution requirements laid down in PSD Section 2.12 are unrealistic, but relaxing the height resolution by a factor of 4...10 (from 100 m to 1 km at 150-km altitude and from 1 km to 4 km at 300-km altitude), puts the 16000-element array in a position to meet the PSD time resolution requirements.
- The user community must reconsider whether incoherent-scatter altitude resolutions better than 1 km are really scientifically meaningful and required at altitudes above 150 km before the 3D project proceeds to the RFQ phase – any improvements in this regard will have to be bought at very capital high cost and may be of marginal use for most applications.

8.2.4 Serviceability and survivability aspects

The 3D radar system will be required to operate essentially unattended, around-the-clock and in all seasons. However, to the best of our knowledge no phased-array as large as the projected 3D arrays has ever been constructed in the sub-Arctic before and so there exist very little data on what to expect in terms of weather and environment effects on array performance and survivability. It is therefore prudent to design for the worst imaginable environmental and weather conditions and other challenges likely to be encountered in the deployment area.

In the summertime, reindeer and elk can stray into the antenna fields, potentially damaging radiators and feedlines. As already discussed in connection with the element antennas, winter weather conditions at the Core site can be extremely variable. It is not uncommon for the snow layer to eventually reach a thickness of more than a metre just before the spring thaw sets in; the EISCAT Heating system air-insulated aluminium coax antenna feedlines, which run on low supports close to the ground, have frequently been deformed by the sheer weight of the snow.

To minimise the risk of any part of the array becoming damaged by heavy snow in this way, the array support structure must be designed to be as transparent to precipitation as possible. Since the available space between the vertical structural members is likely to be too small to allow regular mechanised snow removal, it is also recommended that antennas, feedline systems and support structure are elevated at least some 2.5 – 3.0 m above ground level, which should provide plenty of room for a whole winter's accumulated precipitation underneath. If at all possible, all arrays should also be fenced-in to keep game and the public out. The fenced-in area should preferably be made substantially larger than the actual array, such that the fence can do double duty as a snow-screen.

8.2.5 Constructional aspects

In addition to the actual antenna, each Core array element comprises an RF power amplifier, a transmitter exciter, a receiver and a data processing unit. These sub-systems must all be securely weatherproofed and supplied with electrical power, full-duplex high-speed data communication channels and reference time and frequency signals. Distribution grids for these services must be put in place before, or concurrently with, the construction of the actual array. Since more than 16,000 array elements are required to approach the full target performance specification, an array architecture where each element is a self-contained unit is clearly impractical from the utility distribution point of view.

However, at the 3D operating frequency, 235 MHz, several factors conspire to allow a very substantial degree of array sub-grouping. At this frequency, cable losses are much lower than at UHF and in addition the sky noise temperature is relatively high (>100K), such that a few Kelvin of additional noise impacts the signal-to-noise ratio only marginally. This allows the designer to dramatically reduce the equipment housing and distribution grid complexity by housing all electronics subsystems for a group of some tens of elements in a single equipment container, connecting to the actual antennas through low-loss coaxial cable runs of some metres. Substantial cost savings can be expected to result; assuming N_R radiators per group, the number of

equipment containers, power feeds, data network nodes and time and frequency distribution cables are all reduced by a factor $1/N_R$.

The proposed equilateral triangular grid structure lends itself naturally to a self-similar grouping scheme. Any array element and its six nearest neighbours can be circumscribed by a regular hexagon, seven such hexagons can be circumscribed by a larger “hexagon” with seven times the area and 49 elements, and so on in hierarchical order (7, 49, 343, 2401, 16807) elements per sub-group, where the last group is essentially equal to the fully expanded array. Following a detailed analysis of the effects of cable losses [D3.2 pp.31-33], the 49-element group is recommended as offering a favourable trade-off between cost and construction gains and loss of statistics at low signal-to-noise ratios. With an inter-cell distance of 4.9λ or 6.2 m, the 49-element cell is also a very suitable size for the interferometry application.




Figure 8.2.2: Illustrating how the 3D Core array can be built up from close-packed 49-element sub-groups, each of which can be regarded as composed of seven 7-element hexagonal cells. The Figure shows a top view of a 343-element, approximately 18-m diameter array group, formed from seven sub-groups (outlined in red). Each sub-group is served by a common, approx. 2-m by 2-m equipment container (indicated by a blue square at the centre of each sub-group) containing all RF, signal processing and control and monitoring electronics.

Figure 8.2.2 shows how seven 49-element sub-groups, each with its own approx. $2.0 \times 2.0 \times 2.8$ m equipment container, are close-packed to form a 343-element group. In this arrangement, the shortest distance between any two containers is about 3.8 m, leaving ample room to drive around in the array with a 4x4 motorcycle and a trailer to bring in equipment and/or remove failed units. In principle, whole containers could be brought in and out of the array in this way, suggesting that regular maintenance and major repairs could be done on a unit-exchange basis, with the actual work being done

at a service centre. Assuming realistic system parameters [D3.2 p.34], a standard 3-phase 400 V, 35 A feed per container should be enough to handle all electrical power demands; a multi-core optical fibre cable should easily handle the 24 Gbit/s maximum expected data rate out of each container.

The same basic array layout and organisation are recommended also for the receive-only arrays. Since these have no transmitters, the power demands are much lower (a 400 V, 16 A three-phase feed per container should be plenty), but because multiple beams will be synthesised for nearly 100% of the time, the data bandwidth requirements are doubled.




Figure 8.2.3: A side view of the 343-element array group of Figure 8.2.2. Each hexagon represents a seven-element cell, comprising six element radiators at its corners and one at its centre. The array is assumed to be elevated at least 3 m above average ground; the actual element radiators and the array support structure are suppressed for clarity. Seven 2 x 2 x 2.8 m equipment containers, each serving 49 radiators, are situated under the array.

8.2.6 Pointing calibration and adaptive pointing corrections

Beam-forming and beam-steering in the EISCAT_3D system is performed by summing the several thousand digital data streams from the individual array elements with appropriate time delays to generate one or more beams in the desired direction(s). The delay applied to a particular data-stream is the sum of a pointing-

direction-dependent term, an offset term representing the best current estimate of all system-dependent time delays, and a correction term that adaptively corrects for environmental and other time-dependent effects. Calibration of the timing system in the array will be performed using the techniques described in section 8.4 of this report. Correction term values will then be determined by scanning a beam over the nominal positions of the strong celestial calibrator sources Cas-A and Cyg-A and fitting a correction matrix to the observed residuals. A variant of this time-proven technique has been used with great success for more than 25 years to calibrate the EISCAT UHF system antenna pointing.

Cas-A ($\delta = + 58^\circ 33'$) will be in the grating-lobe-free region of the Core array for nearly 90% of the time. At any given azimuth, Cyg-A ($\delta = + 40^\circ 36'$) is about 18 degrees lower in the sky and therefore ideal for mapping out the Core beam at large zenith angles towards the south. With the flux densities of both sources exceeding 6000 Jy at 235 MHz [Baa77], a properly focussed 49-element Core array sub-group should see them at better than unity SNR. A possible way to set up the arrays initially is therefore to first optimise the focussing of each sub-group separately on one or both sources (the system geometry is such that the receive-only arrays will only be able to observe Cas-A). Signals from seven neighbouring cells are then combined; the relative time offsets between cells are adjusted for maximum SNR, and so on in a hierarchical manner until the whole array has been focussed.

In order to maintain accurate pointing at all times, one of the Core beams will be dedicated to continually scanning a region of the sky centred on Cas-A; data from this beam will be used to adaptively update the pointing correction matrix. The receive-only array pointing corrections will be similarly updated whenever Cas-A is inside their fields-of-view. At other times, they will be checked against the Core transmitter beam; this will allow the accuracy of the lateral component of the respective directional vector to be verified and corrected if necessary.

8.2.7 The Demonstrator Array

In an effort to make the EISCAT_3D system as flexible and future-proof as possible, many novel design concepts, e.g. direct-sampling receiver systems, fractional-sample time-delay beam-steering, the simultaneous forming of multiple beams, adaptive pointing (self-) calibration, adaptive polarisation matching, coherency detection and generation of interferometry processor trigger signals, were made part of the system design from the beginning. The purpose of the Demonstrator is to serve as a test-bed on which these mission-critical concepts can be tested in practice and under realistic climatic conditions. As some of them (e.g. adaptive polarisation matching and coherency detection) can only be validated through the actual reception of incoherent-scatter signals from the ionosphere, the Demonstrator array at the EISCAT Kiruna site has been designed for a centre frequency of (224 ± 3) MHz, which enables it to receive transmissions from the existing EISCAT VHF transmitter at Ramfjordmoen, Norway.

The 48-element array is laid out as a rectangular structure with the long axis in the Ramfjordmoen-Kiruna plane. The elements are arranged in twelve equispaced rows of four antennas each (Figure 8.2.4). The element radiators are (6+6) element X Yagis with $\pm 25^\circ$ half-power opening angle, directed towards the north and elevated to 55°

In each row, signals from the four antennas are combined in two 4→1 coaxial in-phase power-combiners, one for each of the two polarisations. This produces a broadside (E-W plane) array pattern with a half-power opening angle of about $\pm 6^\circ$. For each of the two polarisations separately, the output signals from each array row can then be combined to form either a single beam or several beams at different elevation angles in the Kiruna-Tromsø plane. This arrangement (beam-steering only in a vertical plane) is sufficient because the VHF transmitter beam is also essentially only steerable in the N-S plane. It has the added advantage that the beam-steering function is reduced to a one-dimensional problem – simpler and faster to implement than full 3-D steerability, but sufficient to validate concepts and algorithms.




Figure 8.2.4: Schematic side and top views of the Demonstrator array.




Figure 8.2.5: Array E-plane radiation pattern.

The combined signals are processed in 24 direct-digitising receivers. After down-conversion and decimation, the output data are converted to optical signals and routed into the site building via optical fibre for digital beam-forming and analysis.

To enable a systematic checking-out of the beam-focussing algorithms and the time-keeping system, a near-field test source has been installed. This is a horizontal dipole – plane reflector combination, mounted on a telescoping mast located approximately 4 m in front of the array and looking down on it at 55° depression. By changing the mast height over the (9 – 21) m range, the relative propagation time delays to the different array rows can be varied in a controlled manner. A R&S vector signal generator generates and injects different types of RF signals, including sub-microsecond pulses and spread-spectrum signals.

Tests and experiments carried out so far have been done on one of the two polarisations at a time. In this mode, fully analysable incoherent-scatter signals from 300 km altitude have been received at about 4% signal-to-noise ratio. It is estimated that when the adaptive polarisation system has been brought fully on-line, the SNR will exceed 6% at $1.0 \cdot 10^{11}$ m⁻³ electron density.

Additional information can be found in Section 10.

8.3 Beam-forming concepts

8.3.1 Introduction

Since the EISCAT_3D radar will be a Large Aperture Array Radar (LAAR), it is necessary to use time-delay beamforming instead of the more commonly used phase-delay beamforming (see also Section 5.3). This is necessary since the length of the shortest radar pulse that the radar need to be able to receive is smaller than the size of the array aperture. Therefore, if phase-delay beamforming were to be used, signal energy would be lost whenever the illuminated part of the array is less than the array itself. This is because the phase-delay method only cares about the phase of the signal, i.e. it can at most delay one whole wavelength of the incoming signal. At 224 MHz (the Demonstrator Array frequency), the wavelength is roughly 75 cm which is a lot smaller than the maximum length inside the array aperture which can be up to 300 m. By using time-delay beamforming, this problem is removed since any delay can be compensated for by the system.

To achieve a versatile array radar, pointing angles, number of simultaneous beams and rapid redirection of beams are essential. In a time-delay beamforming array, this translates into a high number of different delays for each antenna element in the array that have a short switching time. This can be done in the analogue domain by introducing different delay-lines in the signal path for each antenna element, or it can be done in the digital domain by delaying the signal through filters. The benefits of the digital solution easily outrank the analogue counterpart; it is a lot faster in switching, it requires less hardware per beam and per antenna element, and it provides an easy way to increase the number of beams after the system has been built. The main drawback of the digital system is the difficulty to achieve the necessary timing

accuracy, see Section 8.4, to achieve the beamforming. For the EISCAT_3D project, digital time-delay beamforming was chosen to be the most appropriate solution.

8.3.2 Beamforming layout

To have an expandable and reliable simulation tool, each part of the physical system has been designed separately so that parts can be added, updated and removed without affecting the rest of the system. The incoming signal to the array is created as highly over-sampled data and then fed through each step of the receiver system, i.e., anti-aliasing filters, S&H, ADC and beamformer. White noise is added to simulate background and temperature noise, and timing jitter is added to provide means to evaluate the effect of non-ideal timing. The positioning of the EISCAT_3D system receive sites are shown in Figures 8.3.1, yielding two different sets of sweep variables and mounting directions for the Yagi antennas used in the array. The values in are corrected for the curvature of Earth, which is $0.009^\circ/\text{km}$. This also affects the azimuth width of the receiver arrays, since the sending array has a 30° wide beam pointing straight up. The maximum beam-steering limits necessary to see the whole beam in Tromsø are larger than the -3 dB beam-width of the Yagi antennas, see Figure 8.3.2. Thus, the Yagi beam-width is a limit in the setup of the test array, and this limitation is used in the simulations done below. The antennas have been permanently mounted as the beam-steering limits listed in Table 8.3.1 suggest, but with a maximum steering limit of 30° .

Limit	Site 1	Site 2
Elevation	$46.09^\circ \pm 30.61^\circ$	$65.22^\circ \pm 42.45^\circ$
Azimuth	$0^\circ \pm 38.03^\circ$	$0^\circ \pm 68.55^\circ$

Table 8.3.1: Beam-steering limits for the different receive sites.




Figure 8.3.1: Geographical locations of the transmitter and receiver sites of the existing EISCAT UHF system. The EISCAT_3D sites will be at similar locations.




Figure 8.3.2: Gain pattern of the Yagi antennas used in the EISCAT_3D test array, as applied in the simulations. The pattern is a sinc curve fitted to the actual data points from the Yagi antenna data sheet. Note: There are no back- or sidelobes as they are not needed for the simulation.

8.3.3 Digital time-delay beamforming

As stated in Section 8.4.1, the overall demand on the timing error of the EISCAT_3D radar is to have a standard deviation of less than 160 ps. To reduce the timing error introduced by the hardware, the Analogue-to-Digital Converter (ADC) is placed as close to each antenna element as possible as described in Section 8.5. After the incoming signal has been digitised through under-sampling at 80 MHz in the ADCs, each signal needs to be digitally delayed so that the incoming signal from each antenna element is matched to every other antenna element in the array. After the antenna elements are lined up correctly, the sole remaining step is to sum the signals to create a beam.

The delaying of the signal can be divided into two parts, whole sample delay which delays by multiples of the sampling frequency, and Fractional Sample Delay (FSD) which delays the signal parts of a sample. This first part is easily achieved in the digital domain by dropping samples, and the second part can be implemented by digital filters. For the EISCAT_3D project, Finite Impulse Response (FIR) filters were chosen.

8.3.4 Finite impulse response filters

To achieve the requested beam-steering granularity for EISCAT_3D, the minimum step size of the delay for each element must be below 13.1 ps [Ste06]. This granularity is not affected by the standard deviation of the timing error, but instead contributes to the error. There are five main design criteria to consider when designing a time-delay FIR filter; group delay error, phase error, amplitude error, filter length and coefficient resolution.

- The group delay error is important since creating a group delay is what we are trying to achieve. Thus any error in the group delay reduces the delay accuracy from each antenna element, causing less than optimal beam-forming. It also causes our broadband signal to be distorted since different frequencies in the band will have different delays. The design criteria for the beam-forming granularity sets the minimum fractional delay; 1/1024th of 12.5 ns yields delay granularity of ~12.2 ps, and thus a quantisation error of ± 6.1 ps. This leaves only ± 0.55 ps of error to the filter design since the total error is set to be ≤ 13.1 ps, which has proved very difficult to achieve. By increasing the number of filters to 8192, (for a memory cost of 81 kB for a set of 1024 filters), the beam granularity is reduced to ~1.5 ps, leaving ± 5.8 ps for the filter design.
- Since group delay is the derivative of phase with respect to frequency, it should not be necessary to optimise on phase delay as well. However, this is only true for a single filter by itself. The phase shift of any one filter is not important, but the difference between filters is, i.e. if the phase is shifted 30° by some filters and -150° by others, they would cancel out and thereby degrade the beam. To prevent this from happening, the phase error of each filter must be small, i.e. within the previously stated ± 5.8 ps, thus the difference between the filters will be equally small. Since both the group delay error and the phase error manifest themselves as time errors they can be added together. Thus, it is the sum of the two errors that must meet the ± 5.8 ps accuracy demand.
- The amplitude error is important since it will affect the beamwidth, side-lobes and nulls. This is because beam-forming works by constructive and destructive interference between antenna elements through summing of signals, so if some signals are larger than others they can affect the summing adversely. The main effect is beam-widening, but even a few percent of amplitude error will cause only limited beam-widening. Awaiting more specific simulation results, we have set a conservative limit of $\leq 1\%$ for the amplitude error.
- The length of the filters, or the number of taps, directly affects the amount of hardware necessary to realise the filters. Assuming an FPGA with the same calculation clock speed as the sampling rate, a 36 tap long filter will require 36 multipliers to be realised.
- Similarly, the coefficient resolution also affects the necessary hardware as the resolution of the multipliers must match the filters.

The demands on the FIR-filter design are summarised in Table 8.3.2.

Criteria	Limit
Group Delay Error	} ≤ 5.8 ps
Phase Delay Error	
Amplitude Error	$\leq 1\%$
Filter Length	36 taps
Filter Resolution	18 bits

Table 8.3.2: Summary of the FIR-filter design criteria

8.3.5 Filter design results

After optimizing the design of the FIR-filters created for the EISCAT_3D beamformers and studying data sheets of different hardware, it was concluded that FIR-filters with 18-bit coefficient resolution and a length of 36-taps would fulfil the set demands, see Figures 8.3.3 – 8.3.5. As can be seen, the criterion of less than 5.8 ps timing error is met for all delays over all frequencies.

While to the date of writing this document, there has been no testing of the beamforming filters in hardware, extensive simulations have been performed to evaluate the accuracy of the filters [Joh08].




Figure 8.3.3: Phase errors for each filter (fractional delay) over the EISCAT_3D frequency band for a 36-tap long filter with a coefficient resolution of 18 bits.




Figure 8.3.4: Group errors for each filter (fractional delay) over the EISCAT_3D frequency band for a 36-tap long filter with a coefficient resolution of 18 bits.




Figure 8.3.5: Amplitude errors for each filter (fractional delay) over the EISCAT_3D frequency band for a 36-tap long filter with a coefficient resolution of 18 bits.

8.4 Time-keeping

8.4.1 Introduction

The necessary accuracy of the timing system in the EISCAT_3D radar has been thoroughly investigated during the course of the project [Ste06, Joh08]. The results show that for the receiver arrays to function within the defined parameters, the total error in the timing between any antenna elements in the array may not have a standard deviation of more than 160 ps. This error consists of three main contributors; the receiver hardware, the timing system, and antenna phase-centre movement due to external factors such as weather. Thus, it is reasonable to have a target of a third of this error, ~50 ps, for the timing system by itself.

There are two different approaches to solving the timing distribution system for EISCAT_3D. The first system is a cable-based system that continuously evaluates the delay between different parts of the timing system, thereby keeping track of any shifts in delay through the system. The second one is based on Global Navigation Satellite Systems (GNSS), which tracks the difference in phase of the signals from satellites to the timing units placed in the centre of each sub-array [Ste08].

There are benefits and drawbacks of both systems: The cable-based system is less prone to disturbances but needs a lot of cables in the array and can only run when the radar is off. The GNSS-based system could be sensitive to disturbances since it is based on radio signals, but if disturbances are significant to the GNSS signal, chances are that the radar itself also is unusable. The GNSS system has a lot less hardware requirements, i.e. almost no cables, and thus could be a lot cheaper in a large array.

Both systems have been evaluated and are described in the following sections but at different levels. The cable calibration timing system has been simulated, built and tested in the test array in Kiruna, whereas the GNSS based system has been simulated and tested on an even smaller scale (single base-line test) at Luleå University of Technology.

8.4.2 Cable calibration timing system

One way of achieving phase synchronisation was first suggested by Grover [Gro94]. The idea is to send a pulse down a line and let it reflect at the end point, measuring the average time between the outgoing and reflected pulses. By doing so, one actually measures the time of the reflection, regardless of where on the line you are positioned when doing the measurement. Thus, using only one line for the whole array, each antenna element measures the same time with great accuracy. By synchronizing another distributed clock with this measurement, one would achieve phase-synchronised clocks in the whole array.

This method would prove difficult to achieve in the EISCAT_3D project, mainly because of difficulties in keeping the signal levels large enough through the array and to create a sensor with good enough accuracy to measure the time of each pulse. However, a version of the system has been developed where each LNA of the system is connected to every other LNA, see Figure 8.4.1. By injecting a sinusoid signal in

the LNAs in different order and with different frequency, the actual relationship between the length of the cables used for the calibration can be decided. The direction of the injected signal can also be controlled, so that the length of signal path can be deduced not only between the LNAs, but also from each LNA to the phase centre of the connected antenna and also through the normal signal chain to the ADC.

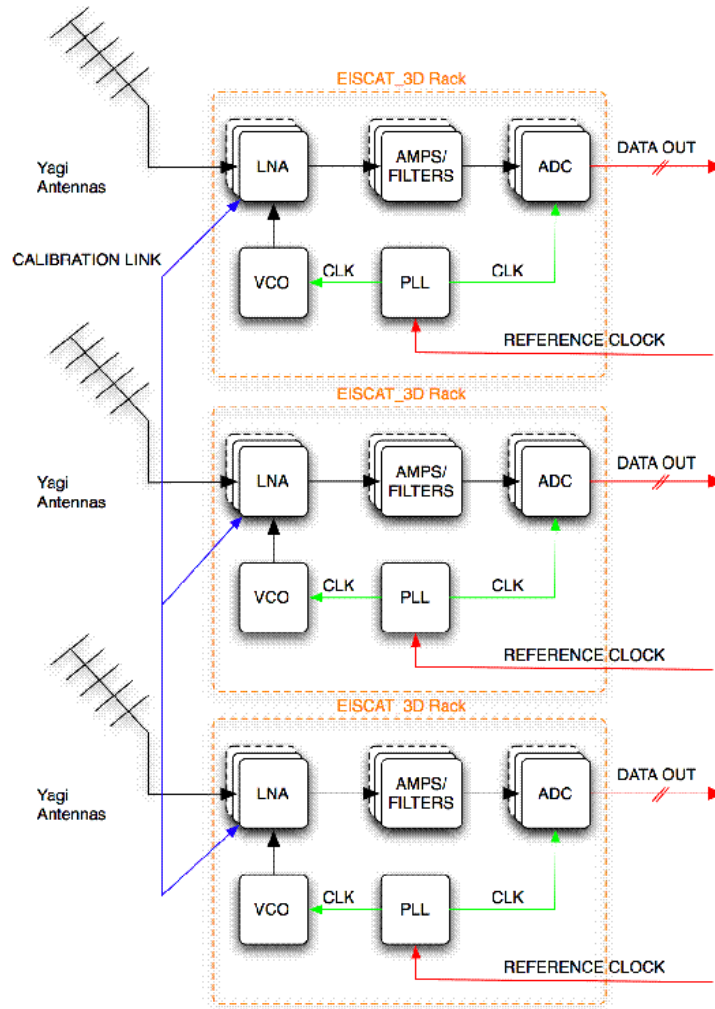


Figure 8.4.1: Diagram of the EISCAT_3D LAAR racks containing the cable calibration system. Each LNA is connected to every other LNA in the array via the CALIBRATION LINK to allow calibration of the different units to each other.

8.4.3 Simulation results

The proposed calibration system has been simulated in Matlab by computing the resulting n-port when all the components of the system are interconnected by cables of random length. The directional coupler is modelled based on data extracted from four coupler boards, both in terms of capacitance/inductance matrix and the board to board variations of all the parameters. The coupler board switches are modelled by a 3Ω resistance and 1.25 pF terminal capacitances, with deviations 0.1 common and 0.02 difference between switches. The resistors are configured for a 20 dB local signal injection attenuation. If we set the antenna and front-end reflections to $-(10..13) \text{ dB}$ and $-(15..18) \text{ dB}$ (chosen randomly), respectively, we get delay and amplitude

standard variations of 32 ps and 6.0%, which can still be considered acceptable. Increasing LNA reflections to $-(10..13)$ dB increases the error to 45 ps and 7.7% respectively. On the other hand, eliminating LNA reflections completely only improves the performance to 20 ps and 5%. Eliminating both LNA and antenna reflections reduces the errors to 14 ps and 4.5%, respectively. Thus, the $-(15..18)$ dB range of return loss seem like a reasonable target for the LNA design.

The results above are of course specific for a specific set of variations, and should by no means be seen as the ultimate achievable accuracy, as better coupler and component tolerances, combined with a good method for compensating for known offsets (based on measurements during manufacturing and field calibrations), could be expected to yield significant improvements. Note that the current front-end calculation do not attempt to make any corrections for known errors, e.g if we make production measurements on the coupler/LNA boards, nor does it compensate for the limited directivity of the coupler which results in antenna impedance measurements being way off.

8.4.4 Measurement setup

To evaluate the actual performance of the cable-based calibration system, a C program has been developed which can be started through the server software installed in the test array in Kiruna. The calibration program is connected to the channel boards of the original EISCAT system, capable of capturing 256 samples of I- and Q-data for each row and polarisation of the test array at 1.25 MHz. This limits the bandwidth of the system to 625 kHz, and with a sampling frequency of 80 MHz and a further digital down-mixing by 16.3 MHz, this places the signal band for the calibration system in the test array from 223.075 MHz to 224.325 MHz. In a future system, this limitation would not necessarily apply as it is currently limited by the maximum frequency of the channel boards.

For each calibration measurement, the channel boards capture 50 μ s of data every 50 ms for a period of 5 seconds. At every second 50 μ s data block, the couplers and VCO outputs in each LNA are set by the calibration program. This is necessary to map all the different signal paths between each ADC of the whole array. When the lengths of the signal paths are known, the time difference between each ADC can be calculated. At the start of each 5 second data block, the EISCAT radar controller sends a synchronisation pulse to the calibration program so that each data block contains well defined settings of the couplers and VCOs.

The calibration program saves 3 different data blocks which each uses a different VCO frequency to allow resolving phase ambiguities since the signal paths between some rows in the test array are greater than the wavelength of the injected frequency.

After analysis of the captured data in Matlab, a matrix of the correct time- and amplitude differences between each ADC in the array is stored and later used for correct determination of the beamforming filter settings.

8.4.5 Measurement results

At the time of writing, no measurement results from the test array in Kiruna were available. While large sets of data have been collected, the analysis is not yet complete. The results will be published in Deliverable 4.2.

8.4.6 GNSS timing system

In the system described here, low-end L1-only GNSS receivers are used in a highly application specific environment, to provide the picosecond accuracy timing. This is done by dividing the LAAR into small sub-arrays of e.g. 16 elements each. The maximum length of the cables distributing the clock is then reduced to 6 m which is short enough to be calibrated by length approximation only, assuming that the clock distributed to each sub-array is known. By inserting a Global Navigation Satellite System (GNSS) receiver at each of these sub-arrays, to provide a clock reference that is unaffected by changing conditions over the array, the antennas are now timed to the specified accuracy.

Other benefits of building a GNSS timing system include lower cost due to reduced amount of coaxial cable throughout the array, and that a continuous cable length calibration system that ensures the timing accuracy of the distributed clock system is no longer necessary.

8.4.6.1 GNSS timing system concept

Each of the sub-arrays will contain a Phase Locked Loop (PLL) in which the distributed frequency reference is reproduced and distributed to the local GNSS receiver, the radar ADCs, and a signal injection system located as close to the radar antenna elements as possible, to calibrate the analogue signal path of the system, as shown in Figure 8.4.2.




Figure 8.4.2: Diagram of the EISCAT_3D LAAR receiver front end and GNSS timing unit

The main purpose of the PLL is to adjust the phase of the reference clock to be equal throughout the array. This is achieved by creating a closed loop feedback from the GNSS receiver to the PLL and adjusting the phase according to the phase differences in the received satellite signals in respect to a reference GNSS receiver. The reference receiver is a high-end receiver which is used in conjunction with application specific software to produce the information sent to each of the sub-array receivers that is necessary to calculate the phase difference of the local clock compared to the reference clock. The PLL can now be used to make the needed adjustments so that all local clocks have the same phase throughout the array.

The information sent from the reference GNSS receiver to each of the GNSS timing units is; which satellites to use, Doppler-shift, tracking chip, and expected phase and time. This information will allow the sub-array receivers to only be capable of tracking a low number of satellites, no more than six, and using the tracked phase differences to calculate the expected phase of the local PLL. Thus, full capability receivers are not needed, but instead a Field-Programmable Gate Array (FPGA) will be used with a GNSS RF-frontend to control the PLL.

8.4.6.2 GNSS simplified receiver

In general, an off-the-shelf GNSS L1-receiver is rated to produce a clock with an error of less than 50 ns, which is about 1000 times higher than the necessary 50 ps accuracy. However, specific conditions apply to this GNSS timing system that improves the accuracy and will relax the requirements on the receiver, such as:

- A short base-line system, i.e. the maximum distance between two GNSS antennas is 300 m which infer all significant external errors, such as atmospheric, ionospheric, and ephemeris errors, in this application to be common over the array.
- A common high accuracy reference clock is available throughout the array to all receivers, which removes a significant part of the clock drift errors between the receivers.
- Externally based selection of which satellites are to be used for the position & time solution to exclude any timing errors arising from the use of different geometry matrices in the position & time calculations.
- All receivers are stationary and the time constant of the change of cable length in the reference clock distribution can be expected to be in the order of 30 min. This enables the use of a very long integration time, up to 30 min, of the timing solution which will reduce thermal noise significantly.
- Phase measurements from one satellite only are sufficient to calculate the timing error between the sub-arrays since the relative position of each receiver is known with good accuracy. However, more satellites will increase accuracy and also reduce any positional error of the GNSS antenna.
- No integer ambiguity solution is necessary, since the relative position of the receivers is known with good accuracy and the absolute time difference between the receivers is insignificant, only the phase of the distributed clock is important.

These conditions all relax the requirements on each of the GNSS timing units to the point where each timing unit only need to be capable of tracking a low number of satellites, measuring the phase of each satellite and then calculate the timing and position error of its own location relative to the reference antenna. With the use of an FPGA, the correlators and tracking can be built in hardware and the calculation part can be built in software with the use of existing processing cores that are available for many FPGAs, all within a single chip.

8.4.6.3 Test setup for concept evaluation

Test measurements have been performed in an outdoor environment during windy winter conditions, clear weather at -10° C and wind speeds up to 20 m/s in gusts, with two antennas (Novatel GPS-702-GG) placed randomly, but precisely surveyed, at approximately 5 m distance from each other placed on a rooftop to simulate the conditions in the EISCAT_3D LAAR. Intermediate Frequency (IF) data from the antennas were collected during a one hour measurement with a NordNav Multi-FrontEnd receiver and then post-processed in a Matlab script to calculate time and position difference between the antennas. As clock reference a rubidium frequency standard (PRS10, Stanford Research Systems, Inc.) was used.

8.4.7 Results and conclusions

After running the post-processing on the collected data, the expected and measured phase differences between the test- and reference antennas were plotted for each tracked satellite see Figure 8.4.3. The expected and measured values are almost identical except for a bias that differs from satellite to satellite. The bias is an

indication of a timing difference between the two antennas, or rather the receivers connected to the antennas. The sign of the bias in this plot differs simply because the plot ignores the integer ambiguity for clarity.

Without averaging, the standard deviation of the clock difference is $\sigma = 50.11$ ps, which is within the usability range for the timing of the EISCAT_3D LAAR. With the 5 min averaging filter applied, this is reduced to $\sigma = 20.52$ ps.

The use of a GNSS based picosecond accuracy timing system has thus been shown to be feasible to achieve using relatively simple calculations. The need for a timing error of less than 50 ps for the EISCAT_3D LAAR has been met with margin using only 5 min of integration time, yielding a timing error of approximately 21 ps. Due to the simplifications of the GNSS position & time calculations possible for the highly application specific conditions that apply for the EISCAT_3D LAAR, a GNSS timing unit can be created in an FPGA at a relatively low cost compared to other timing distribution solutions.




Figure 8.4.2: Measured and expected phase difference between test- and reference antennas. Dashed lines indicates the expected phase differences calculated from the antenna position difference and current satellite geometry, and the solid lines indicates the actually measured phase differences. Integer ambiguity is not resolved in this plot and the data has been unwrapped to simplify the plot. No averaging has been applied to the data.

8.5 Receiver front-end

This section gives an overview of the antenna front end, as shown in Figure 8.5.1 below. The two primary active parts in the receiver chain are the low noise amplifier (LNA) and the A/D converter board.

The **LNA** is built around two amplifier stages and a band-pass filter. The input gate is protected by anti-parallel Schottky diodes. The first stage is followed by a band-pass filter and then by a monolithic gain block. To facilitate system calibration the LNA includes a directional coupler and four analogue switches. This allows the calibration of phase and amplitude variations between channels. Signal injection is performed locally at one LNA at a time, this signal is also distributed over coaxial cable to all other channels in the array. Doing this at all LNAs sequentially makes the calibration independent on the network of coaxial cables.

The **A/D converter board** contains additional filters and two additional gain stages. After amplification and filtering, the signal is digitised by a 16 bit ADC running at 80 MHz. The clock signal is generated by an on-board VCXO which is phase locked to a 10 MHz external clock reference. The last gain stages are fully capable of producing signal levels in excess of the tolerance of the A/D converter. Thus, a protection circuit consisting of a Schottky diode limiter and a steerable attenuator is preceding the ADC. The A/D converter board delivers a 16 bit + clock LVDS output data stream.

The signal chain described above is for one single channel. Each A/D converter board does hold two identical units as described. This can be used to perform the A/D conversion of the signals from both polarisations on one antenna using the same A/D converter board. The following sections give an in depth description of the LNA, the A/D converter boards, and of the front end hardware implementation of the calibration system.




Figure 8.5.1: Block schematic for one channel.

8.5.1 The low-noise amplifier

The design of the low-noise amplifier (LNA) is aimed at fulfilling the specifications listed in Deliverable D2.1 [D2.1]:

- a bandwidth of 30 MHz ,
- a system noise temperature of 50 K and,
- a spurious free dynamic range of 70 dB.

In addition, it was decided that the LNA should have a return loss of no less than 17dB. A high return loss is desirable when building antenna arrays, as this will help mitigate some of the problems with coupling between the antennas in the array. Furthermore it was deemed desirable to filter out unwanted signals as early as possible in the system, as these signals may otherwise intermodulate and produce in-band signals which, once created, are impossible to get rid of. In case the final system uses a smaller bandwidth, both noise figure and return loss can be improved by minor redesign.

The manufactured LNA boards as shown in Figure 8.5.2 include input protection, two amplification stages, a passive filter, calibration signal injection circuitry and various voltage regulation and bias regulation circuits.




Figure 8.5.2: Photograph of one of the completed LNAs

8.5.1.1 RF component characterisation

As the applications for very low Noise Figure (NF) amplifiers operating at VHF frequencies are limited, so is the number of available semiconductor devices. The device with the best specified performance at a frequency not too different from the intended frequency range was the ATF-541M4 E-HEMT from Avago, with a specified noise temperature at optimal noise match of 0.1dB (7 K) at 500MHz.

At RF and microwave frequencies the noise figure is calculated from the measured output noise when two or more different but well known noise levels are present at the input of the device. In order to fully model the noise in an active device (both voltage

and current noise, and the correlation between these), it is necessary to measure the noise figure when several different input impedances are present at the input of the device. This work was performed as a separate M.Sc thesis work within the WP4. The work resulted in a model for the transistor operating at 224 MHz. This model was used in the further optimisation of the LNA.

Simpler (one or two-port network analyser) measurements were used for characterizing various passive components. In particular, a number of wire-wound surface mount inductors from different manufacturers were measured and evaluated, as it became apparent early on that the performance (i.e losses/Q-factor) of currently available inductors suitable for mass-produced electronics would limit the noise figure of the LNA. The electrical models derived from these measurements were used when designing the LNA.

8.5.1.2 Temperature stabilisation

Due to the strict timing requirements for the final system a study on temperature characterisation was performed. It was early concluded that it was not feasible to fully temperature stabilise the whole system. On the other hand, the only parts critical to the timing are certain components in the calibration system. Because these components have been placed on the LNA boards, we can temperature stabilise the whole LNAs in the current system. Further study will be required to determine if even this stabilisation is necessary.

8.5.1.3 Directional coupler design

Each LNA board includes one directional coupler and a number of semiconductor switches for routing calibration signals into the receiver path. Due to PCB size and insertion loss constraints a coupler much shorter than the normal $1/4 \lambda$ (~20cm at 224MHz) was used. The main implications are that the coupling will change significantly over the frequency band, and that the coupling per unit length will have to be fairly large. In order to minimise the insertion loss due to losses in the input signal path a coupler with transmission lines on opposite sides of a ground plane (with a coupling slot) was designed with a wide (for low loss and thus low noise figure) conductor for the incoming signal at one side of the board, a ground plane close to the other side of the board and a suitably narrower conductor for the calibration signal.

8.5.1.4 LNA design

A Matlab tool for simulating interconnected noisy two-ports was implemented. This tool was then used with a particle swarm optimiser to find an "optimal" LNA given a number of constraints on noise figure, return loss, bandwidth, frequency selectivity and stability margin. The coupler, input protection diodes, the filter following the transistor, and a model of the second gain stage (an RFMD RF3376 monolithic gain block) were all included in the simulations used by the optimiser. In order to improve the yield of manufactured amplifiers it was required that the solution should remain stable when components deviated a certain amount from their optimal values.

It is worth noting that the requirement of a high return loss (17 dB) over a large bandwidth (30MHz) leads to a significant trade off in noise temperature from the sub-10 K noise temperature (at optimal noise match) of the transistor to a noise

temperature in the range of 40 to 50 K, due to both suboptimal (from a noise perspective) input matching, and losses in said matching network. After exploring the different possible trade offs and some different topologies the design shown in Figure 8.5.3 was selected for further testing.




Figure 8.5.3: LNA topology selected after optimisation.

8.5.1.5 LNA results

Overall, the design performs as intended. The 24 working units assembled so far are currently undergoing a long-term test at the EISCAT site in Kiruna. Some variations in the frequency characteristics and gain of the LNAs have been observed, attributed mainly to the loose tolerances on available high quality factor surface mount inductors.

In preliminary tests (a single LNA board) 16 kV human body model (20 pulses of each polarity) applied to input of the unpowered LNA failed to induce measure performance degradation. The SFDR of the LNAs has been measured only when connected to the ADCs in the test receiver, but as this combination shows it shows SFDR of 70 dB or better.

8.5.2 The A/D converter board

The advantage of constructive under-sampling is that no intermediate frequency is used, which in turn means no additional local oscillator or mixer. Instead the sampling frequency of the analogue to digital converter (ADC) is chosen so that the frequency band of interest falls between the harmonics of the Nyquist frequency, and thus is sampled without loss of information. As with any sampling process, it is vital that suitable anti-aliasing filters are used, or signal and noise in other bands will fold into the signal.

In order to test the feasibility of using constructive under-sampling and to be able to test beamforming and calibration on array level an A/D converter board was designed. The A/D converter board as pictured in Figure 8.5.4 holds two channels which allow the use of one board per antenna with two polarisations. As shown in the block schematic in Figure 8.5.5, both channels are clocked by one on board VCXO which is

phase locked to an array wide distributed reference clock. Main components and features of the A/D converter board are further discussed in the sections below.

8.5.2.1 Sample clock

The system uses a locally generated sample clock which is phase locked to some other reference signal that is distributed to all receivers. The clock generation is done with one VCXO (Crystek CVHD-950) and one PLL (Analog Devices ADF4002) for each pair of ADCs. This means that the noise due to sampling clock noise will only be correlated between the two polarisations of each antenna, not between antennas. These PLLs are locked to a 10MHz reference signal that is distributed to all receivers. Even though this is the best commercially available VCXO we have been able to find, it limits the maximum SNR of the ADCs to little more than 65dB, as explained further below.




Figure 8.5.4: Photograph of one of the A/D boards.

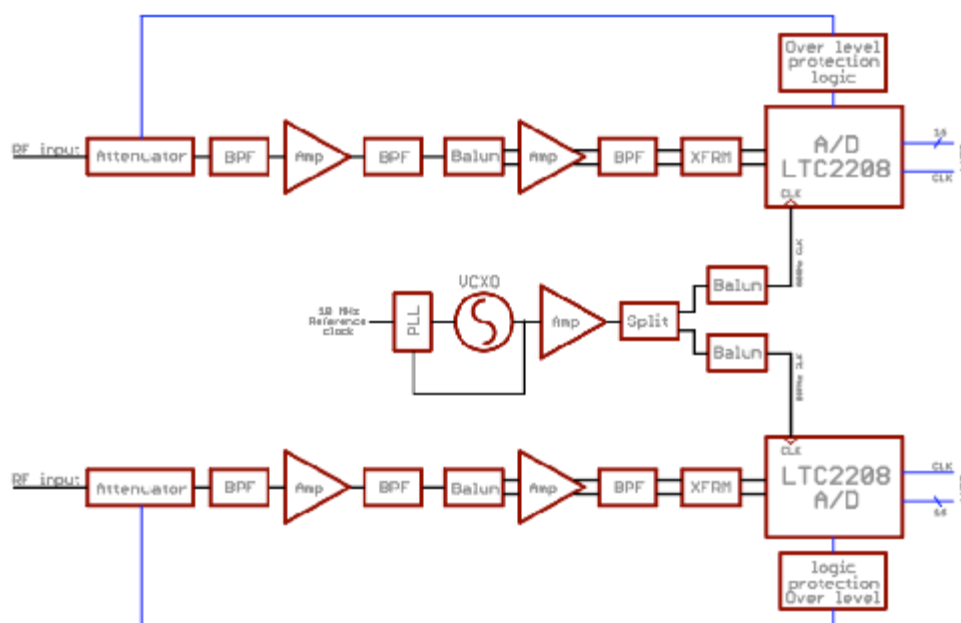


Figure 8.5.5: Block schematic of one A/D card.

8.5.2.2 Anti-aliasing filters

The anti alias filtering is split between three filters. The first filter on the ADC board is a monolithic filter with a sharp characteristic but fairly high insertion loss, which implements the bulk of the out of band signal rejection, especially at frequencies close to the 200-240 MHz band. The two later filters are built from discrete surface mount components, have a much less sharp pass-band, and are used to attenuate out of band noise generated by the the amplifiers after the first filter and strong signals at frequencies far from the band of interest (e.g 100MHz, 450 MHz, 900MHz). These filters have a significantly lower insertion loss, lowering required output level from the preceding amplifier stage, and thus the distortion generated in that amplifier.

8.5.2.3 Amplifier stages

One problem with directly sampling the RF signal is that of amplifying high frequency signals to the high levels required by an ADC ($1-3 V_{pp}$), without introducing significant amounts of distortion. This doesn't affect the first amplifier stage on the ADC board, where a generic high-linearity monolithic amplifier is used, but for the final stage a fairly powerful amplifier (HELA-10D) as the last stage was the only available solution. In order to protect the ADCs from being destroyed by the high output power this amplifier is capable of producing, e.g in case of abnormally high input signal to the system, two types of input protection were implemented: diodes attempt to limit the voltage from the amplifier to safe levels, and when significant current is flowing through these diodes an attenuator at the input of the ADC board is activated, attenuating the input signal by about 40dB, thus keeping the diodes from overheating. Another drawback of this amplifier is the rather high power consumption of about 3W, to be compared to the ADC which consumes about 1.3 W. If more distortion can be tolerated a lower powered buffer could be used, with significantly lower power consumption. This would also make the input protection circuit obsolete.

8.5.2.4 ADC selection

The main figures of merit of any ADC operating at RF frequencies are the signal to noise ratio (SNR) and the spurious free dynamic range (SFDR), typically measured at -1dB full scale (dBFS). Because both of these decrease with increasing signal frequency it becomes necessary to use a more expensive ADC when using direct undersampling than would have been the case if the signal had been converted to an intermediate frequency first. Furthermore, noise in the sampling clock will be scaled by $20 \cdot \log_{10}(F_{sig}/F_{clk})$, which even for the best available VCXOs will result in a larger noise contribution than from the ADC itself when sampling a 224MHz signal at 80MHz. In practice we have found that this limits the maximum achievable SNR over a 40 MHz bandwidth at 224 MHz to be little better than 65 dB, which is 8 dB or more lower than the theoretically achievable SNR when using a good ADC with a noise-free sampling clock. In our implementation we use LTC2208 ADCs from linear technology, selected mainly based on low distortion when undersampling a 200-240 MHz signal combined with moderate power consumption.

8.5.2.5 A/D converter discussion

A total of 13 ADC boards have been manufactured. Due to the issue with amplification of noise in the sampling clock when undersampling the SNR of all ADCs is limited to about 65dB. The SFDR of all boards and channels is better than 70 dB.

As the quantisation noise of the ADC will add directly to the noise temperature of the system, the noise contribution from the ADC must be balanced against the maximum SNR by placing an appropriate amount of gain before the ADC. In the current system the gain is somewhat too high, which places the noise from the ADC 20 dB below the other noise sources, with a contribution to the total system noise temperature of 2.9 K. This is of course very low, but comes at the cost of reducing the maximum system SNR to only about 45 dB. In our current tests this is of no consequence, but for a final system it would probably be desirable to reduce the gain preceding the ADC by about 6dB, thus increasing the system SNR to 51dB at the cost of increasing the noise contribution to 12 K. Generally speaking the use of undersampling as a means of simplifying the system has worked well. The higher cost of the ADC will have to be compared to the cost of the otherwise required additional local oscillator and mixer. The use of an intermediate frequency would also reduce the complexity of the anti-aliasing filters.

8.5.3 Calibration system implementation

The LNA PCBs include the electronics required for implementing a calibration system that will theoretically be capable of mitigating any timing/phase and amplitude errors in a multi-channel receiver system, without being dependant on any external precision components (such as long cables of accurately known length). This system is based on using a separate passive signal distribution network through which test signals can be sent. The key is to have a passive signal distribution network connecting all LNAs to each other and/or to a reference station that is reciprocal, which means that the transfer function (both phase and magnitude) is exactly equal in both directions between any two ports of the network. As it turns out, almost any passive component (including cables and power splitters) are reciprocal, and so is any network built from only passive components.

The signal is injected at the originating LNA after some attenuation (in addition to the nominal -20dB coupling of the directional couplers), and is distributed to the other LNAs where it is injected without additional attenuation. It is possible to select at which of the ports of the directional coupler the signal is applied thus injecting the signal either towards the LNA or towards the antenna, or not to inject the signal at all. These options allow us to measure and compensate for the coupling between the antennas, and to measure the reflection coefficient of each antenna (for diagnostic purposes). The topology of the calibration signal network remains an open question, in a large system it could be beneficial to use an asymmetric network with one or more reference systems where the signal is injected without coupler directly into an LNA, thus gaining about 20dB signal level. While not necessary for small arrays, this could be important when building larger systems. For our tests we have used a fairly

symmetrical network as shown in Figure 8.5.6 to avoid the need for specialised hardware.




Figure 8.5.6: Schematic of the signal injection design.

8.6 Signal processing chain

In the early days of incoherent scatter, the processing of the radar receiver output voltage into range-gated spectral parameters, and the storage of those at a physically acceptable time resolution, were always major system bottlenecks. When the first EISCAT systems were designed in the 1970s, megawatt-class klystrons, big reflector antennas and parametric amplifiers were readily available, but high-speed A/D conversion technology was still in its infancy, computers were very expensive, painfully slow and totally incapable of handling the receiver data directly, and hard-disk storage technology was only being introduced. The only feasible way to handle the signal processing problem was to restrict the receiver pass-band to the maximum extent possible without losing part of the ion-line spectrum, digitise the band-limited signal, process and time-average the resulting (50 – 100) ksamples/s data stream in application-specific hardware – the famous “Alker Correlator” [Alk78] – and finally store the Correlator output on nine-track magnetic tape at an effective recording rate of a few tens of kilobytes per second. In fact, achieving the computational throughput required to do full justice to the transmitter and receiver capabilities was still very much non-trivial even when the EISCAT Svalbard Radar was being designed in the early 1990s.

However, the enormous progress in magnetic data storage technology during the last decade has now made it possible to record raw data on affordable off-the-shelf hard-disk arrays at rates of hundreds of megabytes per second. Similarly, advances in integrated circuit performance have made it possible for the average scientist to have an almost incredible amount of data processing power sitting on her office desk – her

PC or Mac is probably fitted with at least 1 Gbyte of RAM and running at a dual-core CPU clocking at 2 GHz or better. It is also likely to sport in a Terabyte or more hard disk storage. If that is still not enough, the local university data centre can probably offer fast, on-line-accessible disk storage in the hundreds of terabytes or more, capable of recording streaming data at rates exceeding several gigabytes per second.

An important consequence of these developments is that EISCAT users now have ready access to tools that allow them to analyse, in a relatively short time, vast quantities of raw data to which little or no pre-processing has been applied. This has triggered a demand for full-bandwidth amplitude-domain data from the new radar, and research into the development of analysis algorithms operating directly on these.

Accordingly, the current tentative EISCAT_3D receiver / signal processing system architecture is laid out to allow digitised voltage data to be extracted from A/D converters installed as close to the antennas as possible and stored essentially unprocessed, leaving the application-specific extraction of physical parameters from the data to the end user.

Storing all the data produced by the individual array elements separately is a practical impossibility, but by applying a few steps of pre-processing, the data rates can be brought down to manageable levels without losing essential information in the process. For the purpose of the Design Study and this summary report, “signal processing” is defined as precisely that low-level pre-processing.

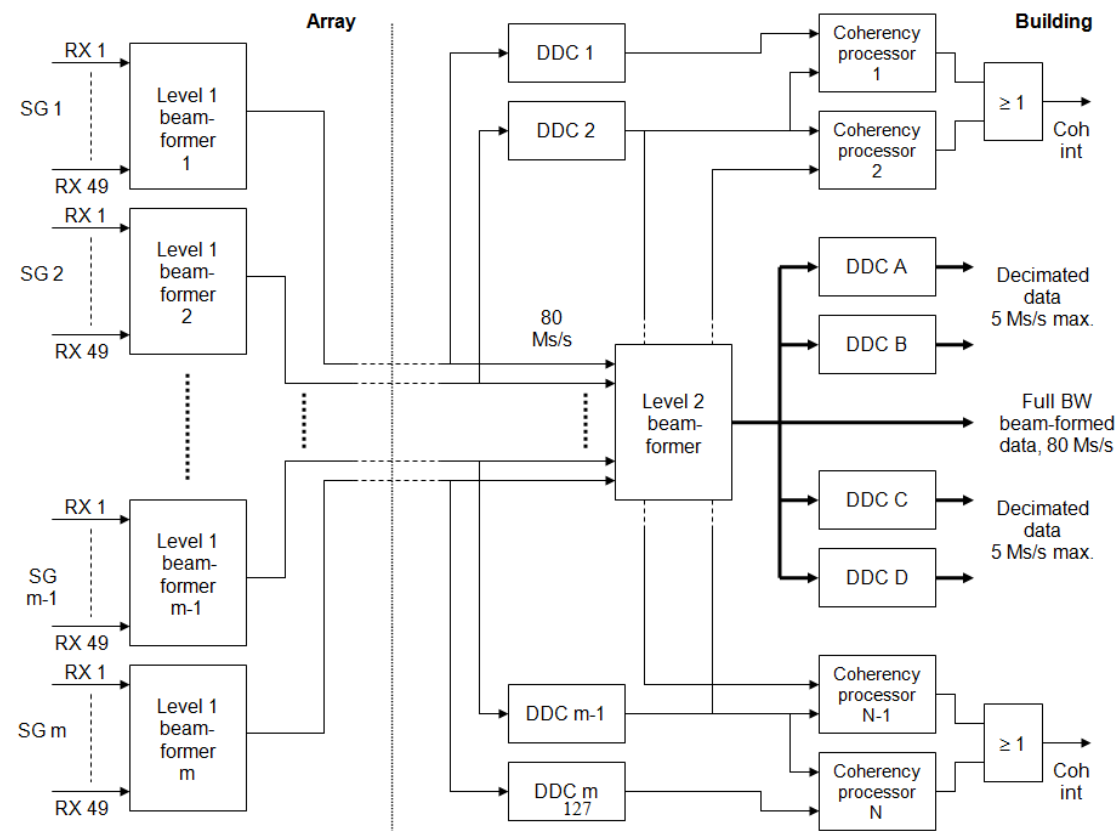


Figure 8.6.1: Schematic diagram of a subset of the data flow through the EISCAT_3D signal processing system, from the individual ADCs to the storage/network units.

Figure 8.6.1 (on the previous page) shows a simplified view of the planned signal flow through the signal processing system, from the output of the individual ADC to the access point into an optical backbone network (disregarding for the moment the buffering of data on hard media). The abbreviations SG 1...SG m refer to the m first 49-element sub-arrays, each of which generates $(49 + 49)$ 80-Ms/s data streams that are delivered to the signal processing system.

Three clearly distinguishable pre-processing functional blocks can be identified in the flow diagram:

- 1) Time-domain beam-forming / beam-pointing,
- 2) Band-limiting and down-sampling,
- 3) Coherency detection and interferometry triggering

We will now consider each of these stages. An actual FPGA-implementation of a beam-former will be described in some detail, while the other stages will be discussed more at the conceptual level, to demonstrate that the hardware needed to build a full specification 3D signal-processing system is available today.

It must be stressed in this context that substantial progress in FPGA technology is to be expected before the design of the full 3D system is finalised. The final signal-processing system architecture is therefore likely to end up looking rather different from the one proposed here – but it would still have to provide at least the basic processing functions listed above.

8.6.1 Time domain beam-forming / beam-pointing

The fully populated Core array will contain about 16000 X-Yagi element antennas. Each receive-only array will contain at least half this number, such that the total system will comprise about 48000 elements! Each element will be equipped with a dual channel (one channel for each of the two polarisations), direct-sampling receiver system, covering a bandwidth of at least 30 MHz and outputting two continuous 16-bit sample streams at a rate of approximately 80 – 90 MHz. The total sample rate generated in the Core is thus in the order of $2.9 \cdot 10^{12} \text{ s}^{-1}$ – an impossible number, even given today's data storage facilities.

One way to transform this into a tractable problem is to process the receiver outputs in a *beam-former*, where the data streams from the individual elements are combined into a very few (at most 6-10) composite data streams, each running at the same 80-90 MHz rate but including contributions from all array elements and representing a fully focussed beam pointing in a desired direction. Apart from making the system data rate manageable, this approach is also appropriate from an information-theory point of view; in a radar system, meaningful signals will only be received from those regions in space that are illuminated by the radar transmitter and so only those data representing beams that look into those regions need to be stored.

After a thorough analysis of the system performance requirements, *true time-delay* (TTD) emerged as the only viable approach to the beam-forming task. TTD is non-dispersive, i.e. the direction of the formed beam is independent of frequency. This quality is an absolute necessity for the 3D system to function as intended, since large bandwidth modulations (up to 5 MHz) will be routinely employed and plasma line reception at frequency offsets of up to ± 15 MHz ($\pm 7\%$ of the centre frequency) will also be a routine feature.

TTD is conceptually very simple: the data streams from the individual array elements are delayed by different amounts and then summed. The delays are selected such that only those components representing signals entering the array from the desired direction add constructively, while all others are suppressed. For a given array, the maximum delay is determined by the physical extent of the array and the maximum beam off-boresight angle. In the case of the 3D Core array with a diameter $\cong 160$ m and a maximum zenith angle of 40° , the maximum delay is about 370 ns, i.e. about 30 12.5-ns sampling intervals. At the same time, in order to achieve the beam pointing resolution of 0.06° specified in the 3D Performance Specification Document, it must be possible to set the individual time-delays with a resolution better than 15 ps, i.e. about one-thousandth of a sampling interval!

To handle this task “on the fly” at 80-90 MHz, we propose to use a system of *fractional-sample delay (FSD) filters*, i.e. all-pass FIR filters set up to work as interpolators. The FSD technique makes use of a consequence of the sampling theorem: *when a band-limited signal has been sampled at a rate exceeding twice the highest frequency present in the signal, the resulting sample series contains all information present in the original signal, so it is possible to reconstruct its instantaneous value at any past time by interpolation of a sufficiently long sequence of samples.* This technique has earlier been used in e.g. sonar, but as far as we have been able to determine, this is the first time that it is being proposed for use in a large research radar system.

As shown in Figure 8.6.2, a basic TTD beam-former consists of two main parts, viz. a set of FSD filter units and a full-adder (Σ). The FSD function is realised as a generic FIR filter. One FSD unit is required per element antenna and beam. The coefficients c^m , $m = 1 \dots n$, determine the filter group delay and must be determined independently for each element antenna and beam direction. By adding more taps to the filter, the integer part of the required delay can also be realised in the same structure.

The data streams output from the FIR filters are all delayed by the correct amounts, such that summing them in the Σ block immediately generates the desired beam-formed data stream. Multi-beaming can now be realised almost trivially by letting several beam-formers with different settings run in parallel on the same input data. This feature will be particularly valuable at the receiver sites, where it will make it possible to observe a number of volumes along the transmitter beam truly simultaneously.

During the Design Study, several different approaches to synthesising FIR delay filters exhibiting almost perfect phase-linearity over a 30-MHz bandwidth have been tried. The LTU team has performed extensive Matlab simulations to verify the validity of each approach; 48-tap filters have been shown to provide the required fractional delay accuracy over a 30 MHz wide band while introducing very little amplitude ripple (less than a fraction of a dB).

The basic arithmetic operation required to realise an FIR filter is the Multiply-Accumulate (MAC). One MAC is required per filter tap and clock cycle, so a single 48-tap FSD filter running at 90 MHz requires $4.32 \cdot 10^9$ MACs per second; the whole array (32000 individual receivers) requires close to $1.4 \cdot 10^{13}$ MACs per second per beam.

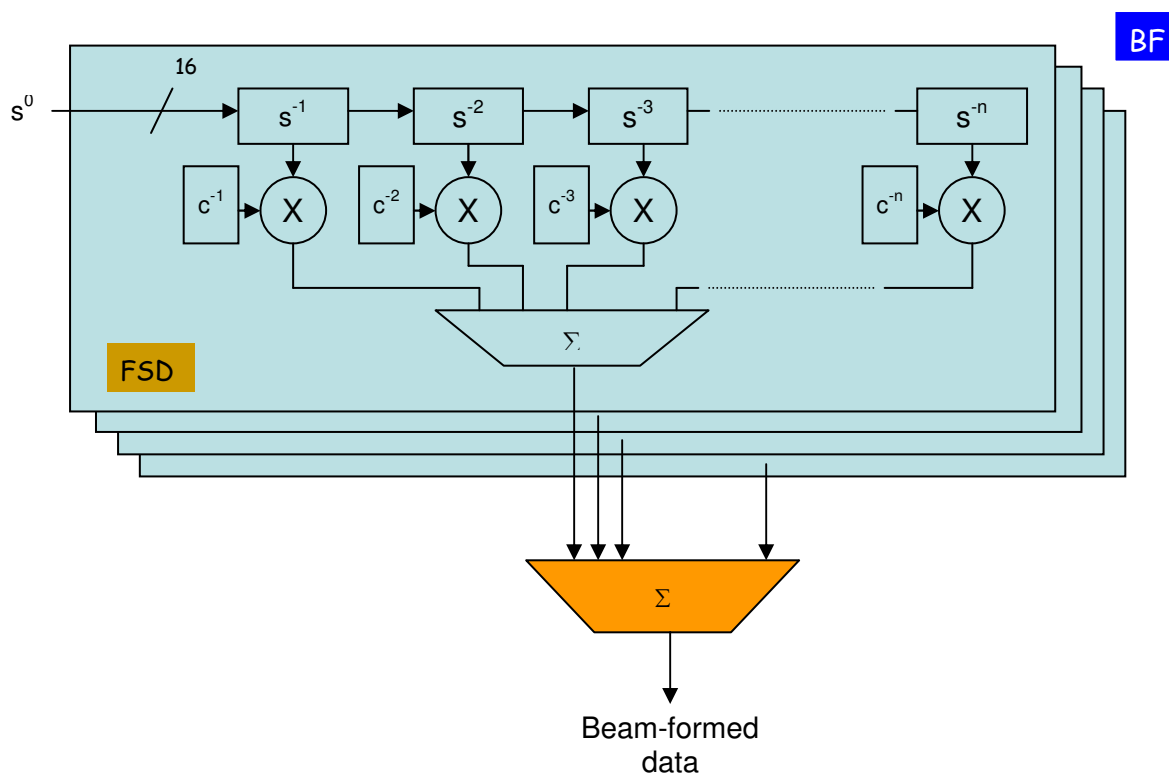


Figure 8.6.2: Block diagram of a basic TTD beam-former, constructed from a number of FIR filters and a full-adder.

This massively parallel, computationally intensive task is an ideal application for *field-programmable gate array* (FPGA) technology. With the possible exception of application-specific integrated circuits (ASICs), no other technology can presently match the computational throughput / electrical power consumption ratio of FPGA. FPGAs can also be reconfigured in the field by reprogramming whenever improved algorithms are developed. This technology is therefore the logical choice for any application that is expected to remain in continual development over its entire lifetime, like the EISCAT_3D system.

To verify the TTD/FSD concept in practice, an FPGA-based beam-former has been developed for the Demonstrator array. A single Xilinx SX95T Virtex-5™ FPGA, implemented in 65-nm CMOS technology, was found to have sufficient raw parallel computing power to generate three simultaneous beams from the Demonstrator in a single chip. To speed up the design and implementation work, an evaluation board carrying the FPGA, a number of optical transceivers and most of the other required interface logic has been purchased. Figure 8.6.3 is a photograph of the evaluation board; a block diagram of the features available on the board is shown in Figure 8.6.4.




Figure 8.6.3: Photograph of the SX95T FPGA evaluation board used to implement the digital beam-former for the 3D Demonstrator array. Two optical transceivers can be seen on the left; an add-on board housing another four optical transceivers has been designed and is currently being manufactured at the EISCAT Sodankylä site.

The SX95T contains a total of 640 MAC blocks that can be clocked at up to 550 MHz. The multipliers are 25- by 18-bit 2^s complement units and the accumulators can handle 48-bit precision. The chip also contains several application-specific hardware blocks (“sub-cores”) that can be used to unpack and de-serialise the stream of high-speed (1.6 Gb/s) data directly from the optical links to the computational core and to off-load the results onto other serial links. In an optimal configuration, the chip can theoretically deliver $3.5 \cdot 10^{11}$ MACs per second, so something in the order of 50 SX95Ts could handle all computations required to form a single beam from the Core! In practice, a substantially larger number will be required, since the real bottleneck now turns out to be the I/O capabilities of the chip rather than the MAC rate.

In the Demonstrator application, the FPGA computes three beams with X and Y polarisation from an “apparent” array of twelve antennas (i.e. the 24 signals from the twelve array rows). Down-sampled I / Q signals from each antenna row are read-in at a data rate of 5 MHz. 48 FIR filters per beam are computed; with three beams, a total of 144 FIR 36-tap filters are thus running in parallel. To do this, the chip only needs to be clocked at 400 MHz.

To fully configure the three beams with correct delays, a total of 72 FIR coefficient filter sets have to be uploaded to the FPGA. Beam-forming sums are computed separately for the I and the Q data streams belonging to each beam, but since these use the same delay in their path from each antenna, identical FIR filter coefficients are used for the FIR pair. Loading, control and read-back are handled over a dedicated Ethernet connection to the host computer. Figure 8.6.5 shows a simplified diagram of the signal flow and resource usage on the SX95T evaluation board.

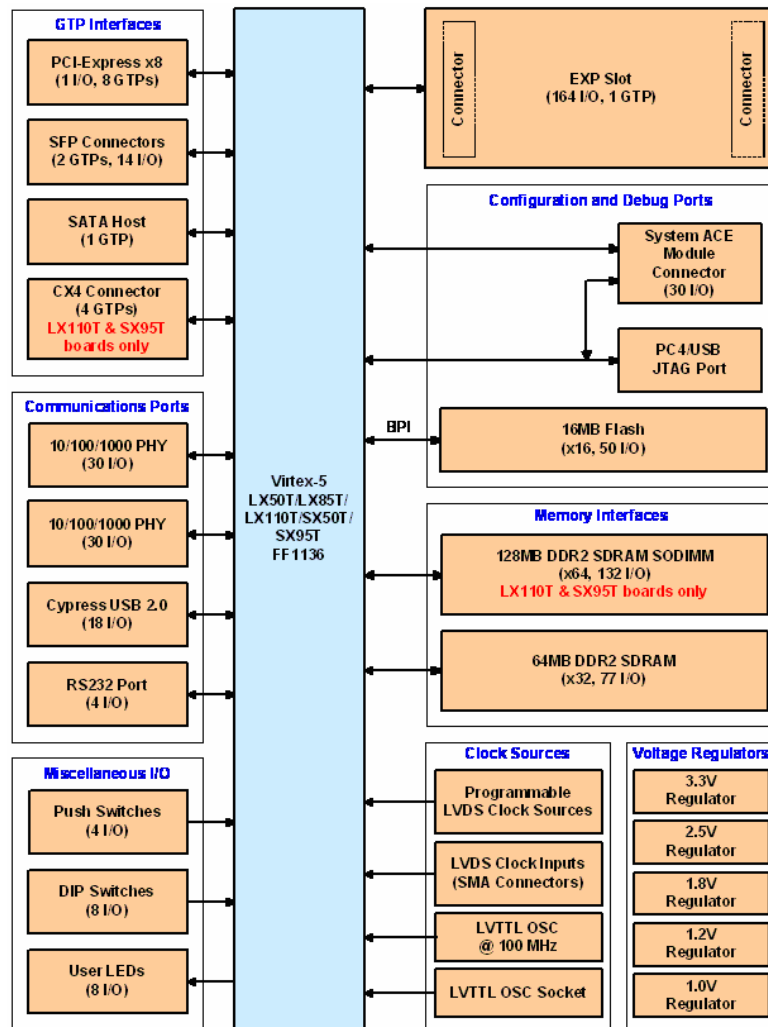


Figure 8.6.4: Block diagram of the resources available on the SX95T evaluation board

As shown in Figure 8.6.1, a convenient and practical way to arrange the beam-former system in the full-size arrays will be to build it as a hierarchical structure, where all data streams from the 49 X-Yagis connected to each equipment container are first fully beam-formed at the container level. This would probably require at least ten SX95T chips. The outputs from these “Level 1 beamformers” are directly useable as input to the next level, Level 2, which would essentially only consist of a set of full-adders, a small number of vector multipliers and the cross-correlators required for the coherency detection and triggering system. An advantage of this arrangement is that the data streams from the individual containers are conveniently available to the coherency detector and also available for storage whenever the coherency threshold is exceeded.

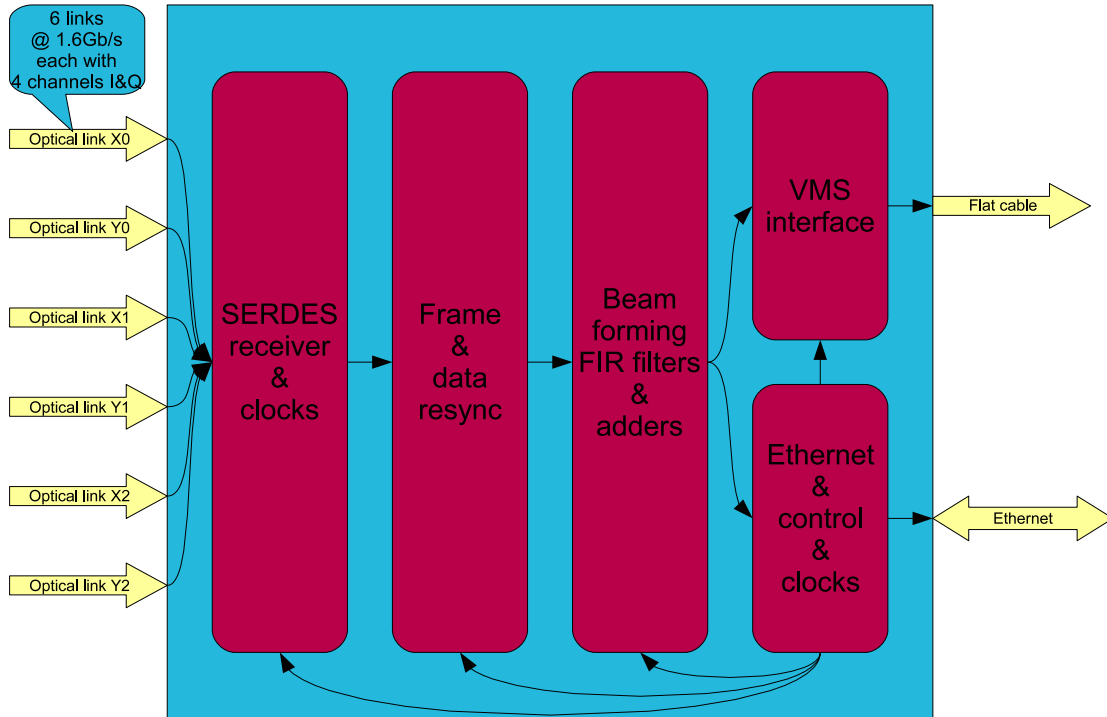


Figure 8.6.5: Schematic diagram of the Demonstrator beam-former signal flow

8.6.2 Band-limiting and down-sampling

The full bandwidth of the EISCAT_3D receiver system is specified at 30 MHz in order for the system to be able to simultaneously intercept signals at offsets of up to ± 15 MHz from the transmitter frequency; this corresponds to plasma lines from a plasma with an electron density in the order of $1.5 \cdot 10^{12} \text{ m}^{-2}$. However, most of the information contained in the scatter signal is concentrated in the *ion line*, which is narrow (only a few kHz wide) and always located at most only a few kHz away from the transmitter centre frequency. The actual bandwidth of the received ion line is of course larger; it is equal to the extent of the convolution of the ion line spectrum with the spectrum of the transmitted probing signal and can therefore be as much as 5 MHz, if a high-resolution modulation is used. In a more typical case, it will be in the order of 200 kHz or less than one percent of the full receiver bandwidth.

Many of the near-real-time data processing modules in the system, including the radar interferometry coherency detectors, the polarisation optimiser routines running at the receiving sites and the “standard quick-look” use only the ion line data. To reduce the computational load posed by these applications, the signal processing chain includes a number of digital down-converter units (DDCs).

The functions performed by the DDCs are: 1) down-converting the band-segment of interest to baseband in a digital mixer, 2) band-limiting the resulting data stream as required by passing it through a programmable FIR filter, and 3) decimating the data stream from the FIR filter to a low rate, commensurate with the filter bandwidth.

There will be two levels of DDCs. In the first level, the data streams from the individual sub-arrays are processed; the outputs from these DDCs (DDC 1...DDC m) will be used as the input signal for

the coherency processors. The second DDC level (DDC A...DDC D) uses the fully beam-formed signal output from the beam-former Level 2 as its input and supplies band-limited data to the polarisation optimiser and the quick-look analysis routines discussed in Section 8.6.4.

DDC subsystems are available as application-specific integrated circuits (ASICs) from a number of manufacturers. One example is the ISL5416 four-channel ASIC used in the Demonstrator receiver system. Many FPGA realisations of DDCs also exist and can be had in the form of so-called “IP cores”, most of which are commercial products that are licensed to the end user for a fee. Either alternative would do the job; no recommendation as to which one to choose is made here as the market situation is likely to change substantially before the design must be finalised.

8.6.3 Coherency processor / trigger for interferometry

Radar interferometry will be a standard feature of the EISCAT_3D radar. It will be possible to record the full-bandwidth raw data from a large number of previously selected sub-arrays (≈ 20) in order to enable the subsequent processing of the data into 3-d volumetric images. But because the aggregate data rate to disk in this mode is some 10 - 20 times higher than when the system is running in its normal, beam-forming mode, sustained interferometry operation is not possible.

On the other hand, since running the system in interferometry mode is only meaningful when there are targets exhibiting a certain degree of coherency in the beam, an automatic *coherency processor* is included in the signal-processing system. Its job is to continually generate estimates of the degree of coherency between the signals received by pairs of separated sub-arrays at some time delay τ . Whenever the coherency is found to meet or exceed some set threshold, the coherency processor issues an interrupt to the radar control system, requesting that the whole array should switch over into interferometry recording mode.

The coherency processor comprises three main parts, viz. the *cross-correlator/auto-correlator*, the *decision logic* and the *buffer memory*. The cross-correlator computes the product $r_{jk}(n,t,\tau)$ of a pair of samples from the complex-valued, band-limited and down-sampled data streams from two sub-arrays j and k located the endpoints of an interferometry baseline:

$$r_{jk}(n,t,\tau) = s_j(n,t) \cdot s_k^*(n,t+\tau)$$

The products are averaged point-for-point over a number of consecutive radar cycles, n_{\max} , to generate an estimate of the cross-correlation or coherency between the two signals received by the two sub-arrays at a time delay τ :

$$\langle r_{jk}(t,\tau) \rangle = \sum_{n=1}^{n_{\max}} s_j(n,t) \cdot s_k^*(n,t+\tau)$$

Before being input to the decision system, the $\langle r_{jk}(t,\tau) \rangle$ estimates are normalised by dividing them by the square root of the product of the respective autocorrelations $\langle r_{jj}(t,\tau) \rangle$ and $\langle r_{kk}(t,\tau) \rangle$, computed as above, in order to eliminate effects of variations in the absolute scattering cross sections.

The decision system may initially be realised as a set of hard logic comparators that generate an interrupt, instructing the array to switch over into full interferometry mode, whenever the coherency exceeds some preset threshold level. More elaborate algorithms, implemented in

software, are also being considered. For the triggering system to work as intended, several tens of coherency estimates must be computed simultaneously in near real-time; to handle the resulting computational demand in a compact and efficient manner, the correlators and comparators will be implemented in FPGA.

In order to make it possible to study the development leading up to the formation and detection of a coherent target, all data used to compute the $\langle r_{jk}(t, \tau) \rangle$ estimates must be buffered for at least the number of radar cycles over which the averaging is performed, such that the data can be retrieved and post-processed whenever a coherency interrupt is generated. There are potentially at least two different ways to implement the buffering; either in a local RAM buffer or FIFO, or the data can be transferred into the central data store, where a dedicated ring buffer would have to be implemented on hard disk. The cleanest and least bandwidth-demanding alternative is probably to include a RAM buffer in the coherency processor.

8.6.4 Quick-look analysis

The band-limited data from each fully formed beam in the EISCAT_3D system is fundamentally equivalent to the data output from the “channel boards” in the existing EISCAT VHF, UHF and ESR receiver back ends. The data processing software following the digital down-converters will therefore need to perform essentially the same tasks as the *lag_wrap* and *decodump* processes running in the current EISCAT systems, that is, it should process the channel board data into complex autocorrelation coefficients, decode and re-shuffle these, and save the resulting autocorrelation *lag profiles* to disk.

The full lag profile set is one of the primary final data products from the system and will be securely archived. For control and monitoring purposes, selected parts of the lag profile data will be analysed in near-real time for physical parameter values (electron density, temperatures, Doppler velocity and polarisation) by e.g. the GUISDAP incoherent-scatter analysis software suite.

The *lag_wrap* and *decodump* routines are easily integrated into any UNIX/Linux environment. They are initialised from simple scripts, describing how the raw data are arranged and how the data manipulation primitives should be applied. They are written in multi-threaded fashion, using POSIX pthread, and can thus immediately make full use of the computing power of bigger shared memory processor (SMP) machines if and when required; their performance will scale in proportion. Since these routines and the other current EISCAT data processing software already meet the 3D post-DDC processing requirements and can be readily re-used, we conclude that no major development of new data processing software is needed to implement a quick-look analysis function into the 3D system.

8.6.5 Polarisation matching

Although not included in the data flow diagram in Figure 8.6.1, the real-time merging of orthogonal polarisation components offers a possibility to reduce the data rate from the arrays by a factor of two and should therefore be briefly mentioned here:

To recover all energy and information present in the scattered signal from a particular direction on the sky, two co-linear beams must be formed, one for each of the two orthogonal polarisations. When the polarisation of the received wave is known, the two resulting 80-MHz data streams can be combined into a single data stream containing all the signal energy without losing any information about the target state. The merging operation, which immediately reduces the data rate to disk by a factor of two, can be performed by a vector rotation and subsequent addition, two

operations requiring only a handful of MAC blocks and therefore representing a totally negligible computational load in comparison to the beam-forming task; the merging technique is therefore very appealing from a storage and bandwidth conservation point of view.

At the Core array, it will be possible to arrange matters such that the received polarisation is always close to circular, making beam merging a straightforward process. The biggest gains can however be reaped at the remote sites, where five or more simultaneous beams will be in operation. Here, the received polarisations are time-varying functions of the amount of Faraday rotation suffered by the scattered signals on their way from the scattering region to the receiver, and therefore not known *a priori*; the actual polarisations must be determined from the respective received signals and the corresponding scaling factors must be computed and fed back into the FPGA system. These operations require a fair amount of statistical processing of the incoming data in near-real time, using the so-called *subspace tracking* technique. For this purpose, a high performance, general-purpose computing platform must be attached to the signal processor. The required procedure for polarisation estimation is described in the following section.

8.7 Polarisation estimation

8.7.1 Introduction

The purpose of this part of the study has been to derive a simple adaptive solution for matching the receiver to the unknown polarisation state of a transmitted signal affected by variable Faraday rotation during propagation. The intention is naturally to minimise any noise enhancement due to polarisation mismatch, thereby improving the Signal to Noise Ratio (SNR). The proposed solution has been obtained by recognising that the polarisation vector can be viewed as the dominant eigenvector of the measurement covariance matrix. With this insight, an approach based on subspace tracking has been used to determine the underlying polarisation from measured data. The abilities of the derived solution have been investigated and it has been shown that that performance similar to perfect polarisation matching can be achieved. However, these abilities were evaluated using simulations in a controlled environment, and the transfer into a real data setting has to be performed with great care.

The proposed approach tracks the polarisation by seeking the dominant eigenvector of the measurement covariance matrix. In other words, the scheme locates the subspace in two-dimensional complex space that inherits the most energy. In order for such a strategy to be successful, especially in low SNR scenarios, the noise contribution has to be white (in polarisation) with good accuracy, i.e., having equal power in all subspaces. If this assumption does not hold, the adaptive scheme will tune to strong noise components rather than the desired polarisation state. Sources for non-white noise components can include, for example, non-equal amplification of the two channels. Various methods can be used to guard against this. One preferable approach, which has become a *de facto* standard in array signal processing, is to pre-whiten the data using noise only measurements, collected for instance before or after the signal of interest.

The purpose of this section is to document and discuss the chosen strategy to track the polarisation state of the received field. The initial task of the receiver is to demodulate the received signal to base-band, sample at an appropriate rate, and beamform to an appropriate location in space. These steps are performed individually on both the horizontal and vertical polarisation channels,

producing two series of discrete-time signals $x_H[k]$ and $x_V[k]$, respectively. k is here a sample index. By writing $x_H[k]$ and $x_V[k]$ in vector notation, the received data can be modelled as¹

$$\begin{bmatrix} u_H[k] \\ u_V[k] \end{bmatrix} = \begin{bmatrix} \cos(\beta) \\ \sin(\beta)e^{j\delta} \end{bmatrix} u[k] + \begin{bmatrix} n_H[k] \\ n_V[k] \end{bmatrix}, \quad k = 1, \dots, K, \quad (1)$$

see [3]. In (1), β and δ are the polarisation parameters, $u[k]$ is the backscattered signal, and $n_H[k]$ and $n_V[k]$ are independent complex Gaussian noise source components in the horizontal and vertical channels respectively. In order to simplify the notation we introduce,

$$\mathbf{x}[k] = \begin{bmatrix} u_H[k] \\ u_V[k] \end{bmatrix}, \mathbf{a}(\beta, \delta) = \begin{bmatrix} \cos(\beta) \\ \sin(\beta)e^{j\delta} \end{bmatrix}, \mathbf{n}[k] = \begin{bmatrix} n_H[k] \\ n_V[k] \end{bmatrix},$$

where $\mathbf{x}[k]$ is the measurement vector, $\mathbf{a}(\beta, \delta)$ the polarisation vector, and $\mathbf{n}[k]$ the noise vector. Using this shorthand, (1) can be written as

$$\mathbf{x}[k] = \mathbf{a}(\beta, \delta)u[k] + \mathbf{n}[k], \quad k = 1, \dots, K. \quad (2)$$

Our ultimate desire is to measure properties of the backscattered signal, $u[k]$. To accomplish this it is often convenient, as an intermediate step, to first extract the backscattered signal $u[k]$ from the measured signal $\mathbf{x}[k]$. In doing so we would like to find an estimate of $u[k]$ while suppressing the noise as much as possible. Within the EISCAT_3D system, this can be achieved using the Minimum Variance Unbiased Estimator (MVUE). For the problem in (2), the MVUE (also the Maximum Likelihood (ML) estimator) is given by

$$\hat{u}[k] = (\mathbf{a}(\beta, \delta)^H \mathbf{a}(\beta, \delta))^{-1} \mathbf{a}(\beta, \delta)^H \mathbf{x}[k] = \mathbf{a}(\beta, \delta)^H \mathbf{x}[k] = \cos(\beta)u_H[k] + \sin(\beta)e^{-j\delta}u_V[k], \quad (3)$$

see [Kay93]. In order to implement (3), $\mathbf{a}(\beta, \delta)$, i.e., β and δ have to be known. Propagation phenomena, such as Faraday rotation and signal path delays, however, make these partially unknown at the receiving site. We therefore require a method to extract the polarisation vector $\mathbf{a}(\beta, \delta)$ from the data, minimising noise enhancement due to polarisation mismatch, and thereby improving the Signal to Noise Ratio (SNR).

¹ In reality we also have an unknown and constant phase offset; however, all succeeding processing is based on spectral characteristics for which this off-set vanishes. Furthermore, in the prescribed model, the polarisation parameters are assumed to be fixed in time. Although this is not exactly true in our application, the polarisation parameters change very slowly, so that over a finite interval they can be considered fixed.

8.7.2 Polarisation estimation using eigenspace analysis

To formalise our analysis we assume that both $u[k]$ and $\mathbf{n}[k]$ are distributed according to zero mean complex Gaussian distributions with (co)variances σ_s^2 and $\sigma_n^2 \mathbf{I}$ respectively². Furthermore it is assumed that the desired signal $u[k]$ is independent of the noise $\mathbf{n}[k]$. The Gaussian assumption regarding $u[k]$ and $\mathbf{n}[k]$ in turn means that $\mathbf{x}[k]$ is Gaussian distributed, so that it is completely characterised by its first and second moments, i.e., its mean vector and covariance matrix. If we study the random vector $\mathbf{x}[k]$ we note that since both $u[k]$ and $\mathbf{n}[k]$ have a zero mean, so will $\mathbf{x}[k]$ ³. The conclusion is that, since the mean is known, all statistical information regarding $\mathbf{x}[k]$ can be found in its covariance matrix,

$$\mathbf{R}_x = E\{\mathbf{x}[k]\mathbf{x}[k]^H\}.$$

Since the desired signal is assumed to be independent of the noise, we can decompose the measurement covariance matrix as

$$\mathbf{R}_x = E\{\mathbf{x}\mathbf{x}^H\} = \mathbf{a}E\{uu^*\}\mathbf{a}^H + E\{\mathbf{n}\mathbf{n}^H\} \quad (4)$$

$$= \sigma_u^2 \mathbf{a}\mathbf{a}^H + \sigma_n^2 \mathbf{I} \quad (5)$$

$$= \begin{bmatrix} \mathbf{a} & \mathbf{a}_\perp \end{bmatrix} \begin{bmatrix} \sigma_u^2 + \sigma_n^2 & 0 \\ 0 & \sigma_n^2 \end{bmatrix} \begin{bmatrix} \mathbf{a} & \mathbf{a}_\perp \end{bmatrix}^H, \quad (6)$$

where we have dropped the dependence on β, δ and k for notational convenience. \mathbf{a}_\perp is here the unit normal vector that is orthogonal to \mathbf{a} , i.e.,

$$\mathbf{a}_\perp = \begin{bmatrix} -\sin(\beta)e^{-j\delta} \\ \cos \beta \end{bmatrix}.$$

We note that (6) is the eigen-decomposition of \mathbf{R}_x with eigenvalues $\lambda_1 = \sigma_u^2 + \sigma_n^2, \lambda_2 = \sigma_n^2$ and corresponding eigenvectors \mathbf{a} and \mathbf{a}_\perp . In the statistical signal processing community, \mathbf{a} is usually called the signal subspace while \mathbf{a}_\perp is denoted the noise subspace, see [Kri96]. From the subspace decomposition in (6) we make the following observations :

- The largest (dominant) eigenvalue of \mathbf{R}_x (λ_1) equals the power in the estimated backscatter

$$\sigma_u^2 = \lambda_1. \quad (7)$$

- The smallest eigenvalue (λ_2) equals the noise power

$$\sigma_n^2 = \lambda_2. \quad (8)$$

- The signal power is given as

$$\sigma_u^2 = \lambda_1 - \lambda_2. \quad (9)$$

- The polarisation vector \mathbf{a} is the dominant eigenvector.

The last observation is of primary interest since our aim is to estimate the polarisation vector $\mathbf{a}(\beta, \delta)$.

² The Gaussian assumption is made for analytical tractability. Even if the underlying signals do not follow such an assumption, the methods will still work, but only exploit second order properties.

³ The zero mean assumption is natural since the considered signals are down-converted versions of the corresponding bandpass signal. Furthermore, valuable information is rarely included in the bias and most systems remove the mean value, implicitly or explicitly, prior to further processing.

8.7.3 Direct approach for polarisation estimation

Following the observations made above, we note that a natural strategy in estimating the polarisation vector \mathbf{a} is as follows:

1. Form an estimate $\hat{\mathbf{R}}_x$ of \mathbf{R}_x based on measured samples, $\mathbf{x}[k], k = 1, \dots, K$.
2. Perform the eigenvalue decomposition of $\hat{\mathbf{R}}_x$

$$\hat{\mathbf{R}}_x = [\mathbf{v}_1 \quad \mathbf{v}_2] \begin{bmatrix} d_1 & 0 \\ 0 & d_2 \end{bmatrix} [\mathbf{v}_1 \quad \mathbf{v}_2]^H, \quad \text{with } d_1 > d_2.$$

3. Take the dominant eigenvector of \mathbf{R}_x as the estimator for \mathbf{a}

$$\hat{\mathbf{a}} = \mathbf{v}_1.$$

What remains is to decide upon an estimator for \mathbf{R}_x , in order to carry out step 1 above. Covariance matrix estimation is a difficult task that has inspired much research effort over the last decades, see for instance [Yan94], [Sve05] and references therein. By far the most frequently encountered approach is to use the sample covariance matrix

$$\hat{\mathbf{R}}_x = \frac{1}{K} \mathbf{S}, \quad (10)$$

where \mathbf{S} is the multivariate S-statistic,

$$\mathbf{S} = \sum_{k=1}^K \mathbf{x}[k] \mathbf{x}[k]^H. \quad (11)$$

Although \mathbf{S} is a sufficient statistic for \mathbf{R}_x and (10) is the Maximum Likelihood estimator, (10) has well known weaknesses, see [Sve05]. These weaknesses, however, primarily lie with respect to the eigenvalues and since we are primarily concerned with estimating the eigenvectors, the use of (10) is appropriate. Moreover, since our primary interest is to estimate the eigenvectors of \mathbf{R}_x we can use either (10) or (11) since these share the same eigenspace. In conclusion our polarisation estimator can be given as

$$\hat{\mathbf{a}} = \text{ev}_1\{\mathbf{S}\}, \quad (12)$$

where ev_1 denotes the dominant eigenvector.

8.7.4 Subspace tracking

In the application considered here, the polarisation parameters are not completely constant but vary slowly over time. Hence, in the EISCAT_3D system, an adaptive technique that tracks the polarisation state is required. Using the the framework discussed above, a straightforward approach to accomplish such continuous adaptation is to use a sliding window and at time m apply the following estimate

$$\hat{\mathbf{a}}_m = \text{ev}_1\{\mathbf{S}_m\},$$

with

$$\mathbf{S}_m = \sum_{k=m-K+1}^m \mathbf{x}[k] \mathbf{x}[k]^H. \quad (13)$$

Such an implementation, however, requires the formation of \mathbf{S}_m in (13) and a 2x2 eigenvalue decomposition for each new sample, which requires fairly large computing power. This is not a new challenge and in order to overcome this difficulty, a number of adaptive algorithms for signal subspace tracking have been developed in the signal processing community. In the general case, the data are of size N (in our case $N=2$) and the rank of the signal subspace equals L (in our case $L = 1$). The literature on these techniques is extremely rich due to its importance. Algorithms can

basically be classified into three groups, depending on complexity. An excellent review of the development up to 1990 can be found in [Com90] which treats solutions with complexity of order $O(N^2)$ and $O(NL^2)$. Since then, a variety of fast subspace tracking solutions (complexity in the order of $O(NL)$) have been proposed. The most promising techniques are found in [Yan95], [Bad05] and [Dou08]. For the one-dimensional signal subspace case ($L=1$), all methods show similar characteristics. We utilise the approach given in [Dou08] which for our particular ($L=1$) problem yields the solution given below

- Initialise with a_0
- for $m=1$ to
 - $r_m = \mathbf{a}_{m-1}^H \mathbf{x}[m]$
 - $\mathbf{T}_m = \mathbf{a}_{m-1} + \alpha \mathbf{x}[m] r_m^*$
 - $\mathbf{a}_m = \frac{\mathbf{T}_m}{\|\mathbf{T}_m\|}$
- end

The motivation for choosing this solution is that it shows good adaptation speed and offers a simple implementation using only a single design parameter α . The value of α controls the sensitivity of the applied algorithm. Large α yields very sensitive solutions that are able to track fast changes, however, yielding noisy estimates. In some sense a large α corresponds to choosing a small K in (13) above. In contrast, small values of α render accurate estimates that cannot track fast changes in the data (similar to having a large K in (13) above). Setting a value for α is hence a trade-off between adaptation speed and estimation error, a trade-off that depends on the underlying characteristics, defined through R_x and its time-variability.

8.7.5 Evaluation

To study the performance of the derived estimators we will study a case in which the polarisation changes throughout the data set. We will also study the effect of proper polarisation estimation when estimating the spectrum of the underlying signal. We consider a data set with 100,000 samples. The true underlying values of β and δ start off by being constant, $\beta = \pi/4$ and $\delta = -\pi/3$. After a certain number of samples (in the case of β , 30,000 samples, and in the case of δ , 10,000 samples), the parameters start to change in a linear fashion. After a while, the two parameters return to being constant, β stabilises at $\pi/16$ for samples greater than 80,000 ($m > 80,000$) and δ remains fixed at $-\pi/8$ for $m > 90,000$. The evolution of β and δ is illustrated by the dashed lines in Figure 8.7.1.




Figure 8.7.1: Performance in terms of estimates for β and δ for the batch-type processor in (13). From left to right (a)-(c). (a) Sample support equals $K = 500$. (b) Sample support equals $K = 3000$. (c) Sample support equals $K = 10000$.

In Figure 8.7.1 the performance of the batch method, described in (13), is also shown for a $-3dB$ scenario. The performance using three different values of K , describing the number of signal samples used to form the eigenvector estimate, is shown. In Figure 8.7.1(a), showing the performance for $K = 500$, we note that the resulting eigenvector estimate is fairly noisy. In contrast, Figure 8.7.1(c) shows the performance using $K = 10.000$. Here the length of the data support suppresses the variance; however, the method is not able to track the channel variation in an appropriate fashion. For this specific scenario it appears that using $K = 3000$ offers a good trade-off between suppressing variance and being able to track the polarisation changes. The performance of this estimator is shown in Figure 8.7.1(b).

For the subspace tracking approach given above, results corresponding to those in Figure 8.7.1 are given in Figure 8.7.2.




Figure 8.7.2: Performance in terms of estimates for β and δ for the subspace tracker described in Section 8.7.4. From left to right (a)-(c). (a) Sensitivity parameter $\alpha = 0.004$. (b) Sensitivity parameter $\alpha = 0.0006$. (c) Sensitivity parameter $\alpha = 0.0002$.

Again we consider three different settings. In Figure 8.7.2(a), the performance using a comparatively large value of α , here 0.004 , is shown. Comparing this to the results in Figure 8.7.1(a) we note that the subspace tracking approach yields results similar to batch processing. These results are also verified by comparing Figures 8.7.2 (b)-(c) to Figures 8.7.1 (b)-(c). In conclusion, it appears that subspace tracking is a highly accurate and relatively simple way of obtaining the polarisation vector from data.

In all of the estimates considered above, the variance of the δ -estimate increases tremendously as β approaches zero. This should not come as a surprise, since the polarisation vector $\mathbf{a}(\beta, \delta)$ becomes independent of δ for $\beta = 0$, see (1). This behaviour in terms of β and δ is, however, not an artifact, rather it is a property of the chosen parametrisation. Hence, these properties will not carry over to $\mathbf{a}(\beta, \delta)$ or affect its ability to project onto the proper subspace. To verify this observation, Figure 8.7.3 shows the time-varying performance in estimating signal power where, throughout the data set, the previous 10.000 estimates of $u[k]$ are used to estimate signal power. The same data set as in Figures 8.7.1 and 8.7.2 are used. In addition to the batch processor using $K = 3000$ samples, and the subspace tracker using $\alpha = 0.0006$, Figure 8.7.3 also shows the performance obtained using a fixed polarisation vector. In order to get a fair comparison, the fixed values of β and δ are here chosen as the average ones, i.e., $\beta = (\pi/4 + \pi/16)/2 = 9\pi/64$, and $\delta = (-\pi/3 - \pi/8)/2 = -11\pi/48$. As a reference, the performance of the estimator that exploits perfect polarisation knowledge is also shown.

By studying the performances in Figure 8.7.3, we note that the accuracy in estimating the signal power does not decrease for the adaptive polarisation estimators as β approaches zero and the

variance of the δ estimate increases. This verifies our observations above. In Figure 8.7.3 we also note the substantial advantage of tracking the polarisation parameters, as the squared error in estimating signal power is up to 100 times smaller for the adaptive schemes compared to the case that utilises fixed average values. Only for samples between $m = 40,000$ and $60,000$ are they comparable; but it is exactly in this region where the average values of the polarisation parameters are close to matching the true underlying ones.

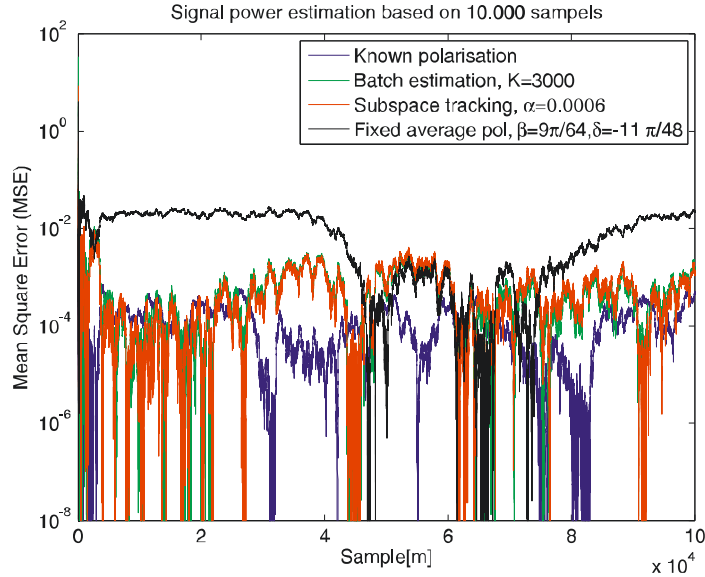


Figure 8.7.3: Time-varying performance in estimating signal power. The underlying polarisation parameters change in accordance to the scenario in Figures 8.7.1 and 8.7.2. The SNR is -3dB.

8.7.6 Spectrum estimation

To further validate the characteristics of the subspace tracking based signal estimator we will also carry out a simulation that investigates the improvements in terms of spectrum estimation quality. In order to simulate a case that resembles the EISCAT scenario, the underlying signal $u[k]$ is in this case not white noise but instead an ARMA-process whose spectral characteristics resemble those of signals encountered in ionospheric measurements. Specifically, we use a 5th order ARMA process described through the following difference equation:

$$u[k] = 1.9759u[k-1] - 2.0135u[k-2] + 1.1026u[k-3] - 0.3276u[k-4] + 0.0407u[k-5] + \\ + 0.0196v[n] + 0.0882v[n-1] + 0.1471v[n-2] + 0.0981v[n-3] - 0.0294v[n-5] \\ - 0.0098v[n-6],$$

where $v[n]$ is a zero-mean white complex valued Gaussian process with variance $\sigma_v^2 = 1$. The underlying filter is a fifth order Butterworth low-pass filter (cut-off frequency $\pi/3$) with an additional zero at $z = 0.5$. The spectrum of this process is given in Figure 8.7.4(a) and its power is close to one. To enable comparison to previous findings we consider the same time-varying polarisation behaviour as studied in Figures 8.7.1, 8.7.2 and 8.7.3; we assume a dataset of 100,000 samples, and the noise variance, σ_n^2 , is set to two, rendering an SNR of approximately -3dB.

To estimate the spectrum, the Bartlett technique using 200 frequency bins is chosen [Sto97]. Given that the data length is 100,000 points, the power spectrum is averaged using 500 measurements. As a reference, the Bartlett estimate of the desired signal $u[k]$, i.e., $\hat{S}_u(e^{j\omega})$ is shown in Figure

8.7.4(b). As we cannot measure $u[k]$ directly, but only consider an estimate, see (3), our ability to estimate the spectrum depends on the accuracy in retrieving $u[k]$. We here study two such estimators which were both studied above. The first is obtained by tracking the polarisation vector using subspace tracking. The sensitivity parameter is set to $\alpha = 0.0006$. The corresponding spectrum estimate is denoted $\hat{S}_{\hat{u}_{sst}}(e^{j\omega})$. The second approach, also considered above, utilises a fixed polarisation assumption. As in Figure 8.7.3, the polarisation parameters are here given by the average over the data set, $\beta = 9\pi/64$, $\delta = -11\pi/48$. The resulting spectrum estimate is denoted $\hat{S}_{\hat{u}_{av}}(e^{j\omega})$. We deliberately leave out the batch approach since the results of this technique are so close to the subspace tracking approach, see Figure 8.7.3.

In Figure 8.7.4(c) these two spectral estimates are shown along with the estimate that is obtained using perfect knowledge of the polarisation state throughout the data set. This estimate is denoted $\hat{S}_{\hat{u}_0}(e^{j\omega})$. At first sight, the performances of all three spectral estimates seem very similar, although we notice the bias of two compared to Figure 8.7.4(b), which is due to the noise component in the estimates of $u[k]$ ($\sigma_n^2 = 2$). However, we observe that the spectrum obtained using subspace tracking is very close to the one obtained using perfect polarisation knowledge. This again confirms our findings that estimating the polarisation parameters yields similar performance to a clairvoyant receiver that knows the true state of the polarisation throughout.

Studying the spectra in Figure 8.7.4 more carefully, we also notice that the spectrum of the fixed approach, using an average polarisation setting, has slightly lower energy than the one obtained through subspace tracking. This inevitable property comes from the polarisation mismatch and the fact that the receiver is not able to capture all of the signal energy. We can study this property in more detail by looking at the two quantities,

$$L_{sst}(e^{j\omega}) = \frac{\hat{S}_{\hat{u}_{av}}(e^{j\omega}) - \sigma_n^2}{\hat{S}_{\hat{u}_0}(e^{j\omega}) - \sigma_n^2} \quad L_{av}(e^{j\omega}) = \frac{\hat{S}_{\hat{u}_{av}}(e^{j\omega}) - \sigma_n^2}{\hat{S}_{\hat{u}_0}(e^{j\omega}) - \sigma_n^2}.$$

We note that $\hat{S}_x(e^{j\omega}) - \sigma_n^2$ corresponds to the portion of the power (as a function of frequency) that is due to the desired signal (using any given method). Since $\hat{S}_{\hat{u}_0}(e^{j\omega}) - \sigma_n^2$ is the estimate obtained using perfect polarisation knowledge, $L_x(e^{j\omega})$ can be interpreted as the fractional power (as a function of frequency) obtained using that same method, compared to having perfect polarisation knowledge. In Figure 8.7.4(d), $L_{sst}(e^{j\omega})$ and $L_{av}(e^{j\omega})$ are shown over the interval $-\pi/3 < \omega < \pi/3$, constituting the frequencies of relevant signal power. Studying Figure 8.7.4(d) we note that the subspace tracking approach does not seem to lose signal power, compared to having perfect polarisation knowledge. The approach using a fixed polarisation vector, on the other hand, appears to lose as much as 10% of its signal power throughout the band of consideration. Again this stresses the importance of treating the polarisation in a proper way.

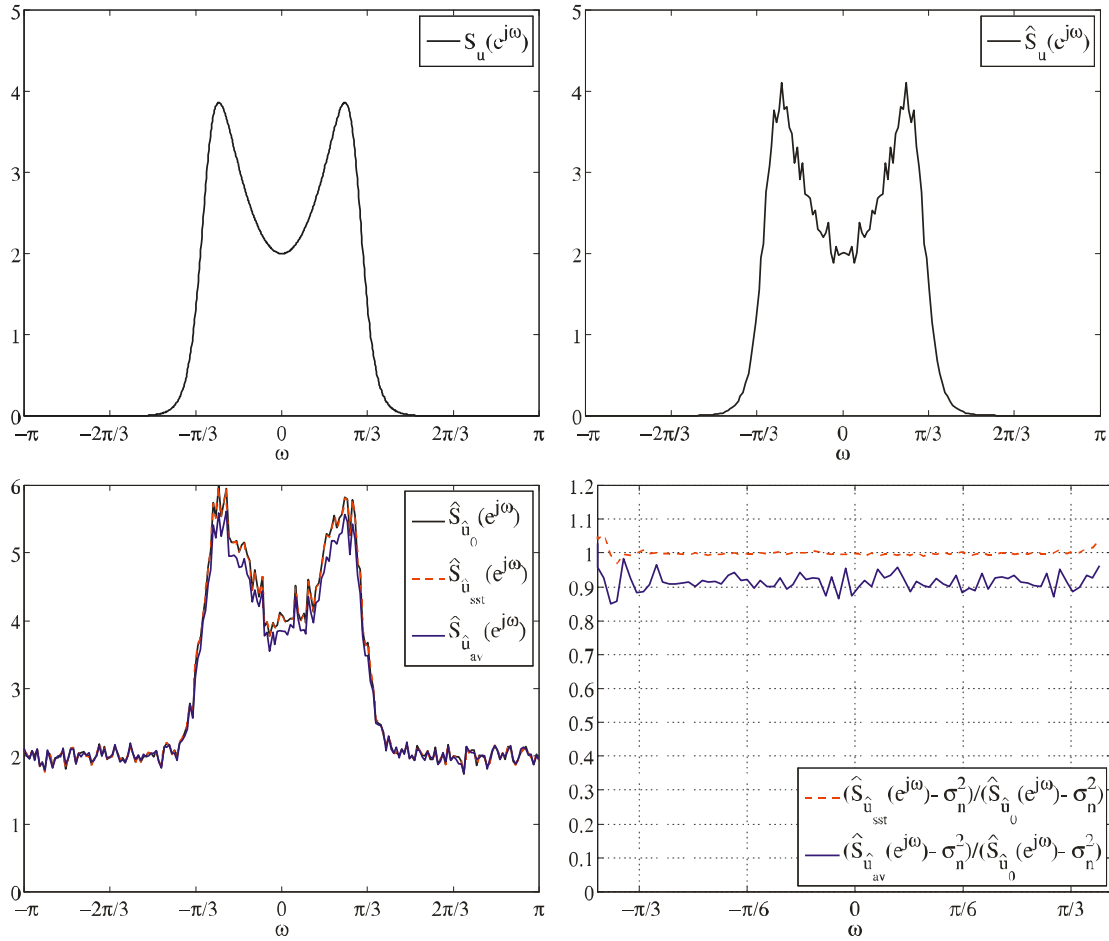


Figure 8.7.4: Performance of spectrum estimation using the extracted signal. From left to right and top to bottom (a)-(d). (a) True spectrum of the desired signal. (b) Spectrum estimation using the true underlying signal. (c) Spectrum estimation using the measured signal. (d) Fractional power loss.

8.7.7 Discussion and conclusions

The utilisation of subspace tracking seems like a very good strategy in obtaining adaptive polarisation matching. Using low complexity adaptation, performance similar to perfect polarisation matching can be achieved. However, our discoveries have been evaluated using simulations in a controlled environment, and the transfer into a real data setting has to be performed with great care. Our approach tracks the polarisation by seeking the dominant eigenvector of the measurement covariance matrix. Hence, our scheme locates the subspace in two-dimensional complex space that inherits the most energy. In order for such a strategy to be successful, especially in low SNR scenarios, the noise contribution has to be white (in polarisation) with good accuracy, i.e., equal power in all subspaces. If such an assumption is not accurate, the adaptive scheme will tune to strong noise components rather than the desired polarisation state. Sources for non-white noise components include non-equal amplification of the two channels, and mutual coupling.

In real applications robustness against non-white noise contributions can be approached using different strategies. One approach is to pre-whiten the data using noise only measurements, collected for instance before or after the signal of interest, see [Kri96]. This is a standard procedure in array signal processing. A second approach is to utilise the Bayesian framework and consider a-priori knowledge regarding the polarisation parameters. Although the polarisation is not exactly known, fairly good knowledge of the polarisation uncertainty can be expressed by using the known bistatic geometry and previous history of for instance annual variations. In a Bayesian setting, the

posterior distribution is expressed (up to a scaling factor) as the product of the likelihood and the prior. If the prior has limited support, the solution of a Bayesian approach is then guaranteed to lie within these bounds. In principle, the Bayesian framework could also help us, through the exploration of subjective prior knowledge, to achieve better estimates for cases in which the noise is truly white. However, at the SNR levels that are targeted here, the number of required measurements, K , is very large, whereby the likelihood almost certainly dominates any reasonable prior, i.e., if the range of the likelihood is within the that of the prior any reasonable prior will be constant over the range of the likelihood. Hence, it does not matter whether the likelihood or the posterior distribution is utilised, since they are equal up to an unknown scale.

8.8 Interferometry and radar imaging

8.8.1 Fundamentals of interferometric imaging

The built-in interferometric capabilities of the EISCAT_3D system - complemented with multiple beams and rapid beam scanning - is what will make the new radar truly three-dimensional and justify its name. The basic foundations upon which the imaging capabilities will be implemented have been provided by Work Package 5 of the EISCAT_3D design study, led by the University of Tromsø.

The technology to be employed by EISCAT_3D is Aperture Synthesis Imaging Radar (ASIR) and is closer to the technology used by radio astronomers (VLBI, Very Long Baseline Interferometry) to image stellar objects (Thompson et al., 2001) than to the SAR (Synthetic Aperture Radar) technique used onboard airplanes and satellites to map the Earth's surface and other planetary surfaces. In the radio astronomy case the source itself spontaneously emits radiation that is collected by a number of passive antennas. In the case of a radar, the transmitter – acting exactly like a camera flash - illuminates the target (the ionosphere or atmosphere) and a number of antennas collect the reflected radiation - exactly as in the radio astronomy case (or like the lens of a camera).

From this point on, the two cases are essentially identical (though Earth's motion is an important difference). The procedure to construct the image of the target consists in calculating the cross-correlation between the signals from all different pairs of receivers. These values represent samples of a function in two dimensions (the positions of the pair of receivers, known as baselines; specialised to the usual case in which the baselines are contained in a plane) called the visibility. The imaging inversion problem consists in obtaining the image from the visibility, which, in principle, should be a 2-dimensional Fourier transform. However, in virtually all cases, the samples of the visibility are sparse, uneven and truncated, which leads to a highly singular inversion problem requiring carefully crafted algorithms. The image obtained from the inversion algorithm is called the brightness distribution, or just brightness, and it is a function of the angle of arrival of each point in the target, that is, the brightness is the angular distribution of the target's intensity. In the radar application it is often advantageous to first decompose the receiver signals into their spectral (frequency) components and apply the imaging inversion to each spectral component separately to obtain one image per spectral component.

It is a non-trivial task to express visually the image as a function of three space variables, one frequency variable, and time - in all, five independent variables. The time variable can be taken care of by displaying images (frames) in rapid succession, that is, in the form of animation or movie. There still remain four independent variables of which only two can be employed using conventional plotting techniques. Since the mathematical relationship between the visibility and the image is a simple (2-D) Fourier transform, all the accumulated applied knowledge on Fourier

transforms in other domains can be applied to imaging. For instance, the resolution of the image is determined by the longest distance between receivers (longest baseline); the larger structures (bandwidth) are determined by the shortest baseline; resolution and bandwidth are related by the Nyquist theorem; windowing effects and their mitigation are the same; and so on.

In the EISCAT_3D design study we have established all the conditions and parameters relating to the description just given above, for the new radar to have a 3-dimensional capability. The core antenna could be composed of 20 or more modules (sets of around $7^3 = 343$ antenna elements) accompanied by a few outlying modules, the latter to comply with the resolution requirements. This is an optimum and flexible antenna layout from which favourable configurations can be quickly implemented to obtain the angular resolution and coverage of the required image; several simulated configurations have been suggested. The phase accuracy of the timing system has been specified to fulfil the desired image resolution.

A novel way to calibrate the imaging system has been proposed using the phases obtained from measurements of the usual incoherent scattering signals for which these radars are designed. The inversion algorithm, based on the Maximum Entropy Method (MEM) has been implemented and tested on simulated and real world data, the latter obtained with the imaging-capable radar of Jicamarca. Methods to represent visually a function of five independent variables - with various degrees of completeness and compression - have been investigated and tested with simulated data.

The deliverables of Work Package 5 summarise the recommendations regarding general and elementary conditions which have to be fulfilled by a phased array radar in order to have built-in imaging capabilities, touching upon the lengths and number of the baselines, the placement of the antenna elements, the active elements, the receiving elements, and signal processing. The properties of the Fourier transform imply that the shortest baseline determines the angular size of the image (bandwidth in Fourier parlance). Thus, the distance between neighbouring antenna modules is what will determine the maximum angular extent of the image. Assuming that the length of a module is 16λ (WP3 Report) the angular coverage will be $1/16$ radians or about 3.6λ which translates into a horizontal size of about 6 km at 100 km range and 18 km at 300 km range. However the image size will be determined by the size of the illuminated volume which is considerably smaller. In the case of strong signals it will be advantageous to utilise a compact subset of the entire array for transmission to cover a wider illuminated beam, consequently resulting in a wider image. Since the transmitted power is distributed, there will be a proportional loss of power. A more effective and elegant alternative is to use aperture illumination coding that employs the entire array [Woo01].

The longest baseline will determine the angular resolution of the image. It is expected that the smallest scattering structures will be encountered in the auroral E-region, as optical measurements indicate the existence of structures of some tens of metres in visible aurora [Tro01]. A useful target resolution is therefore a structure of 20 metres width at 100 km range that subtends an angle of 2×10^{-4} radians which implies a baseline length of 5000λ or 6250 m for a radar frequency of 240 MHz. Since the dimension of the core array will be at most a few 100 metres, a few outlier antennas have to be built.

The number of samples of the visibility function that can be measured is equal to the number of different baselines, which in turn is equal to the number of different receiving antenna pairs, or $n(n - 2)/2$, where n is the number of receiving antennas. Allowing for 4-10 outlier antennas, the remaining will be taken from modules in the core array. The number of the latter will depend on the total number of modules and on their size. Images of reasonable quality have been obtained by a relatively small number of antennas at the Jicamarca Radar, of the order of 6-8 under high signal to noise ratio conditions. At least as many will be supported by the EISCAT_3D radar. For instance

the combination of 8 core modules and 4 outliers will give 66 baselines (somewhat less if there are redundant baselines). Various optimum baseline configurations and signal processing for real-time monitoring are the subject matter of the Configurations Report of September 2006 [D5.1] that is summarised below.

8.8.2 Single-baseline interferometry

The Stage 2 report from Work Package 5 [WP5c] contains the full derivation of basic expressions for the observable complex cross-correlation and complex coherence in a two-element radar interferometer. The expressions are two-dimensional generalisations of the one-baseline interferometer with arbitrary orientation employing separate antennas for transmitting and receiving. The WP5 Report [WP5c] contains a simplified derivation, the results of which have been employed in the processing and interpretation of interferometric measurements using the EISCAT Svalbard Radar [Gry03] and in the early experiments made at the Jicamarca Radar [Far81]. The report includes the definition of the terminology employed in the field and a preliminary discussion of signal processing requirements to monitor the ionospheric targets in real time searching for high coherence events.

In the Stage 3 report [WP5d] it is shown that the phase of the measured visibility function depends linearly on the time delay between the signals arriving at the two antennas of a baseline. Furthermore, one phase period maps to one fringe period, and each of these to one radio-wave period. Thus, a useful condition on the accuracy of the phase - which determines the quality of the measured image - is that the random variations (time jitter) of the timing system, δ , should be a small fraction of the period of the radiation (the radar), $\delta \ll T_o$, T_o being the period of the radiation. A practical condition can be that the total phase inaccuracy be about 10° , or 1/40th of a period, which gives 100 ps for a 250 MHz radar. In fact, the random time variation in this condition should be the accumulated random time variation. Since the accumulated random time variation decreases in proportion $1/\sqrt{n}$, where n is the number of accumulations (when the jitter is statistically independent) the actual allowed time jitter is decreased by this factor. For instance, beam forming is performed by the addition of simultaneous samples (multiplied by phase factors) taken from a number of antenna elements, usually the elements of a module. For 343 elements to a module, the allowed time jitter of the individual samples will be reduced by a factor of $\sqrt{343} \sim 20$, or 2 ns instead of 100 ps.

8.8.3 Array configurations

A meeting between WPs 4, 5, 8 and 9, held in Kiruna in early July of 2006, led to the adoption of a two-level beam forming architecture [McK06]. This architecture, when compared to others that were under consideration, is simpler to implement and produces the desired results for the full array beams, while it eliminates the need for separate beam formers for interferometry, increases the configuration flexibility considerably, and immediately provides interferometric capability for all receiver beams. It is an optimum global solution that satisfies the requirements of interferometric imaging, as the module antennas of the second level constitute the antenna units upon which antenna baselines can be constructed with great flexibility under software control, very possibly at the experiment design/implementation level. The configurations that can be obtained are limited only by the size of the modules in the first level and by the longest linear dimension of the entire array. An important task of the design process was then to survey the parameters of the limiting configurations on the basis of this architecture. Advantageous geometric patterns into which the module antennas can be organised into baselines are described for various sizes and geometries of the core array.

The two-level beam forming architecture consists in combining the signals into beams in two stages. In the first stage, module beam formers will apply a Delay-Phase shift and Add (DP/A) operation to antenna element amplitudes and accumulate the results for groups of elements to form module beams. On entering the second level there is a bifurcation: the module beams from all modules are added to form the full beam; then the same module beam signals are routed to the interferometric sub-system, where they are downsampled to a bandwidth matched to ion line observations. The module beam signals are the starting points for interferometric signal processing and storage.

The purpose of the real-time processing is to decide when to store amplitude data from the interferometry sub-system. The production of images from the amplitude data that is stored will be made offline. It is sufficient to monitor a small number of baselines distributed in orthogonal directions in real-time in order to detect interesting events and trigger the storage of the data streams by thresholding (see Section 8.6.3). Profiles of cross- and auto-correlation and complex coherence functions of a selected set of modules will be calculated in real-time. Any one or all of these functions will be thresholded in order to decide whether to store the data streams or not. The selection criteria may use low threshold values to begin with, to allow for some false positives (data stored without interesting features) and minimise the number of false negatives (interesting events lost). With experience, threshold adaption will be utilised. Temporary storage of amplitude data in ring buffers will allow the permanent storage of data for some seconds prior to a triggered event and to continue for as long after the end of an event (See Section 8.9.2.2). The records will incorporate the selection criteria that originated the trigger and the corresponding range. Generally, the smallest scattering structures will produce higher coherence regardless of baseline, while larger structures will only produce high coherence from shorter baselines. Thus, an initial strategy on the choice of which baselines to monitor in order to catch as many events as possible would be to monitor a couple of short and a couple of intermediate baselines oriented N-S and E-W respectively.

Data for interferometry will be stored as amplitude data. This is a consequence of the fact that with an integration time of 0.1s, the storage requirements are still smaller for amplitude data than for the quadratic quantities such as correlated or spectral data - the reason being that the number of baselines increases quadratically with the number of modules used for interferometry. In addition storing the amplitude data gives added flexibility during the offline processing of the data.

The design of module configurations suitable for imaging has been divided in two parts: configurations for the core (involving a subset of the core modules), and configurations of outlier modules. Randomly placed modules, as well as optimally selected subsets of various regular (dense) arrangements have been analysed for the core configurations. An optimal configuration is one that fills maximally the baseline space (the u-v space of the visibility function). That means generally the most even distribution of baselines with a minimum of redundant baselines. The specifications are given by the number of baselines and their placement on a two-dimensional coordinate system. This is done best graphically as in Figure 8.8.1 which includes the antenna module configuration itself and the distribution of baselines (the sampled points of the visibility) side by side. A few configurations have been also illustrated with their synthesised beams. A nice example is described in this summary. This is a core array using 37 hexagonal modules in a hexagonal grid of 7 lines shown on the left side of the figure. A configuration of 10 modules – shown shaded on the left side of the figure – has been found which achieves 44 non-redundant baselines with only one repeated baseline indicated by the shaded modules on the right side of the figure which shows the baseline coverage. The modules and baselines to be used for real-time monitoring are marked with an “R”.




Figure 8.8.1: A core antenna implemented with 37 hexagonal modules showing a 10-module configuration (shaded hexagons on the left side) that achieves 44 non-redundant baselines and one repeated. The baselines are shown on the right side. The modules marked with an “R” are the suggested for real time monitoring. The “X” indicates the single-module, zero-spacing.

For sizes of the core array under consideration of the order of one hundred metres (80λ) it will be necessary to deploy outlier modules, which are assumed to be identical as the core antenna modules. The core modules will provide the short baselines and their particular imaging configuration will not have a decisive impact on the configuration of the outlier modules. Thus, a 21-module core with an 8-module configuration for imaging was employed for the simulations carried out for this report. The guiding principle is that the synthesised beam should have a minimum of artefacts, e.g. sidelobes. Therefore, the synthesised beam should be as clean as possible in order to give the image inversion algorithm a good initial estimate. Module configurations were evaluated by the maximum sidelobe level in the synthesised beam inside a region in the vicinity of the on-axis position. Based on experiences from radio astronomy, nominal module positions were sought along log-spiral arms. To take advantage of the property of random distributions to have a tendency to produce low sidelobes, the starting point was a random distribution of 100 module antennas spread along a log-spiral path of appropriate size. A sample of 9 modules with favourable sidelobe properties were taken from this distribution based on 9 nominal positions.

We refer to the reader to the original report [WP5d] for the detailed results. It is clear that the particular configurations described are among many others also with favourable properties that could be considered. The main point is that, at this stage of the design activity, it is more important to establish the methods and selection criteria rather than to specify a definitive configuration, which should be done at a later stage of the design and implementation of the project.

8.8.4 Phase calibration

During some of the early interferometric experiments carried out using the two antennas of the EISCAT Svalbard radar, Grydeland [Gry07] discovered that the phases of the complex coherence function of incoherent scatter signals were very well ordered – that is, they were approximately constant within the signal bandwidth of incoherent scattering – even when the amplitude of the coherence was rather low. It was promptly realised that what was going on is as follows.

Under normal and quiet situations, the ionosphere is homogeneous within scale lengths much larger than a typical horizontal extent of the transmitter antenna beam over much of the ionosphere, certainly up to and well above the peak of the F-layer. In this situation, and under the most general conditions regarding the geometry of a radar interferometer, it turns out that the visibility function (the complex cross-correlation between two separated receivers) has a constant phase equal to zero, so that the visibility is a real number. This affords a convenient method for global relative calibration of the baseline phases of an aperture synthesis radar used for ionospheric imaging, as Grydeland had suspected.




Figure 8.8.2: Spectra and complex coherence measured by the ESR. The phase in the lower panel has been calibrated by subtracting the measured phase shown in the first panel. A slight improvement of the coherence amplitude can be appreciated.

Figure 8.8.2 shows an illustrative measurement and its calibration. The upper panel shows from left to right two frames with spectrograms of IS signals measured by the 32 and 42 metre antennas of the EISCAT Svalbard Radar, respectively, while the third and fourth frames show the amplitude and phase of the coherence function of the two signals as a function of Doppler frequency. One can appreciate that the phases are constant within the bandwidth of the signals as defined by the power spectra shown in the first two frames. Likewise, the coherence amplitude of the signal is clearly delineated and stands out from the noise even when its value is very low. The second row shows the results after the calibration was performed. The procedure is to find the phase correction that maximises the amplitude of the coherence and apply the correction to the complex signal amplitudes. One can verify that after the correction, the coherence amplitude has improved. One should not be misled by the slight improvement seen in this example, which is not the object of the present exercise. In an actual imaging experiment, the visibility function is sampled at many points (baselines) and the quality of the resulting image after inversion/restoration depends decisively on the accuracy of the visibility phases.

Many other calibration methods have been reported in the radio astronomy literature and recently also by investigators that employ the Jicamarca radar for ionospheric imaging [Cha08]. A relative simple but effective method of absolute calibration is to measure the radiation of a radio star of known position. The radio star appears to the radar receiver as a point source which results in clean and accurate phase and time delay measurements from which correction phases can be extracted.

8.8.5 Image inversion for interferometry

Amongst the many image inversion/restoration algorithms employed by the much bigger radio astronomy community, two stand out, namely the CLEAN procedure and the Maximum Entropy Method (MEM) [Cor99]. After some considerations, it was decided to adopt the MEM algorithm for development and tests in this project. The CLEAN procedure assumes that the true image is composed of a small number of point sources – which is the case in many astrophysical situations, but not generally in the radar applications. The MEM has a more mature mathematical foundation and, in the version we have chosen, is effective and robust as shown by the Jicamarca implementation [Hys02]. Briefly, the numerical problem is to find an extremum of the following functional, using Hysell's notation and the Einstein summation convention:

$$E[f(e_j, \lambda_j, \Lambda, L)] = S + \lambda_j(g_j + e_j - f_i h_{ij}) + \Lambda(e_j^2 \sigma_j - 2 - \Sigma) + L(I_i f_i - F)$$

where f is the sought after brightness distribution, $S = -f_i \ln(f_i/F)$ is the entropy, I_i is a vector of ones, $F = I_i f_i$ is the integrated (total) brightness, g_j is the measured visibility, h_{ij} is the point spread function that contains the Fourier kernel, e_j are the random errors, σ_j^2 are the (theoretical) expected error variances, and Σ parameterises the error norm, effectively constraining it. The remaining quantities are Lagrange multipliers as follows. λ_j defines the fundamental constraint on the entropy functional by relating the measured visibility (including the random errors) to the sought after brightness that makes the entropy functional an extremum. The other Lagrange multipliers put additional constraints that typically would ensure an improvement of the quality of the final solution: Λ puts a bound on the error norm equal to a preset value equal to Σ ; and L constrains the total brightness effectively ensuring that the solution will be non-negative (positive semi-definite).

The software implementing the algorithm has been written as a Mathematica package. The Mathematica application, with its powerful symbolic and numerical facilities, was considered advantageous to use in order to produce a first prototype in a short period of time. The algorithm has been tested on simulated data and on real-world data taken with the Jicamarca Radar. An interesting configuration of 7 core modules and an outlier module was found that produces an evenly distribution of baselines coordinates with a minimum of gaps and was utilised in the tests. Figure 8.8.3 shows the configuration and the results of inverting an image composed of 5 localised gaussian blobs as described in the figure. The source code of the Mathematica package is available for download at the EISCAT_3D website.

8.8.6 Multi-dimensional radar imaging

The problem of presenting the multi-dimensional brightness distribution data obtained from imaging radar in a way that its physical significance can be appreciated in a visually intuitive manner is very challenging. Since the brightness distribution function is a function of five independent variables: three spatial coordinates (two bearing angles and range), Doppler frequency, and time, $B(\alpha; \beta; r; f; t)$, the problem is how to represent points of a 6-dimensional space on a two dimensional surface - usually on a computer monitor or on a sheet of paper. In fact, it is feasible to implement this representation by using the two spatial coordinates of the displaying surface augmented by the three values that define a colour, and animation to represent time. However, colour as representing three quantities (mathematically/physically correct), does not lend itself to an intuitive appreciation by the human brain (whether this skill can be gained by training is beside the point). A mechanism of dimension reduction is necessary.




Figure 8.8.3: Five-blob image reconstructed from the visibility measured with the 8-antenna configuration with 7 core antennas of the Jicamarca array and 1 outlier shown in the upper left panel. The second upper left panel shows the baselines generated by the configuration. The rightmost upper panel shows the core antennas used in the configuration in red and their coordinates alongside. It is a remarkably optimum configuration that produces an excellent reconstruction.

A very powerful mechanism of dimension reduction (Thomas and Cook, [Tho]) is to use projections, the simplest and most versatile technique being just to collapse one or more dimensions by integrating the dimensions left out from the graphical representations. This technique is automatically multi-level and provides unambiguous data visualisation at different levels of displayed details. At the first level, all but one of the independent variables are integrated, and the result can be plotted as a curve. An example that has physical relevance is the total brightness as a function of time, the total brightness being synonymous with total power or intensity; in the case of pure incoherent scattering, the total brightness would be proportional to the total electron content (TEC). At the next level three independent variables are integrated, leaving a function of two independent variables. A common case is when the surviving variables are range and time and the graphical representations are known as Range-Time-Intensity plots or RTIs, which are usually represented as colour or gray scale maps or as contour plots. These two levels result in static plots. At the third level, keeping time and integrating any two of the remaining variables, one obtains plots as in the previous levels that can be animated by displaying them one at a time successively as a function of time, namely one produces a movie: for instance, dynamic spectrograms as a function of range and time. The fourth level exploits more fully the advantages brought about by imaging by integrating only one independent variable - which is not usually the time - and encoding the value of the brightness distribution and the Doppler frequency with two of the three variables of colour

space. It is advantageous to use the Hue-Saturation-Value or HSV colour space with hue to encode the Doppler frequency and the saturation of the brightness distribution keeping the third colour parameter - the value - constant, usually 100%, that is, only the brightest colour combinations are employed. The Doppler is conveniently divided in three regions of hue with green tones to encode frequencies around zero and blue and red tones to encode the blue and red shifted frequencies, naturally. As mentioned earlier, employing the three variables of the HSV colour space in order to make a dynamic (animated) graphical representation of the brightness distribution as a function of all its independent variables is not advisable since it appears difficult for a viewer of such representation to “decode” the colours into its three quantitative components in an intuitive manner. As things are, even the two-colour variable representation of the 4th level described above requires a little effort to “get used to”.

Cuts come to the rescue, in a way, since their use avoids integrating one variable, and instead, replaces it with a cut. This is accomplished by setting up a 3-dimensional coordinate system of the three physical dimensions [(x; y; z) or (α ; β ; r)] on the 2-dimensional displaying medium (the computer monitor) using perspective and displaying in it a cut through any of the three axes at a specified value. The cut is displayed as a colour map where brightness distribution is encoded by colour hue and Doppler by colour saturation of the map's pixels. Time is “encoded” by animation. Alternatively the cut can be displayed by itself as an animated colour map. Figure 8.8.4 shows an illustration in which both alternatives have been included and the cut has been made on the range axis at $R = 145$ km for simulated data. The EISCAT_3D web site contains several illustrative examples of the visualisation package for imaging developed in this project using simulated data and real data obtained with the Jicamarca radar.

Prototype software to visualise the brightness distribution as described above has been developed as a Mathematica package *E3Dvisualize.m*. The package contains two basic functions for simulating (*SimulateBrightnessData*) and plotting the brightness data (*PlotImagingData*). To generate simulated multi-blob brightness data we use a Gaussian model of brightness distribution for each of a specified number of blobs with specified trajectories as a function of time. The programs and the examples can be downloaded from the EISCAT_3D web site.




Figure 8.8.4. Cross section of simulated brightness data at a given range of $R = 145$ km.

8.9 Data system

8.9.1 Introduction

In designing a data system for EISCAT_3D, the goal is to preserve the maximum degree of flexibility, in order to allow the user to work with the various different levels of data product. This includes the capability to interact with beam-formed data at the sample level, allowing users the possibility to control the signal processing (Section 8.3), interferometric analysis (Section 8.8) and even to facilitate data analysis straight from the sample-level data stream (see the discussion in Section 8.6).

In order to satisfy the dual requirements of totally flexible data handling, together with the need for an archived data set which is of a manageable size for long-term storage, we have designed a two-stage system for handling data within EISCAT_3D. The beam-formed sample-level data, together with data from the interferometry system, and some high-volume data from supporting instruments, are streamed to a large ring buffer designed to hold up to a few days worth of low-level data, after which these data will be over-written. The ring buffer allows the low-level data to be stored for long enough to allow it to be optimally processed, in terms of subsequent auto-correlation and integration in time and range. The latency time of the buffer must therefore be long enough to allow multiple processing strategies to be applied before the low-level data are over-written. The final optimally-derived data products, which are typically at least an order of magnitude smaller, are then transferred to the permanent data archive. At the same time, a second copy of the incoherent scatter data is separately passed through a default signal processing strategy in order to produce the quicklook data needed for control of experiments. Figure 8.9.1 shows an overview of the data system topology.

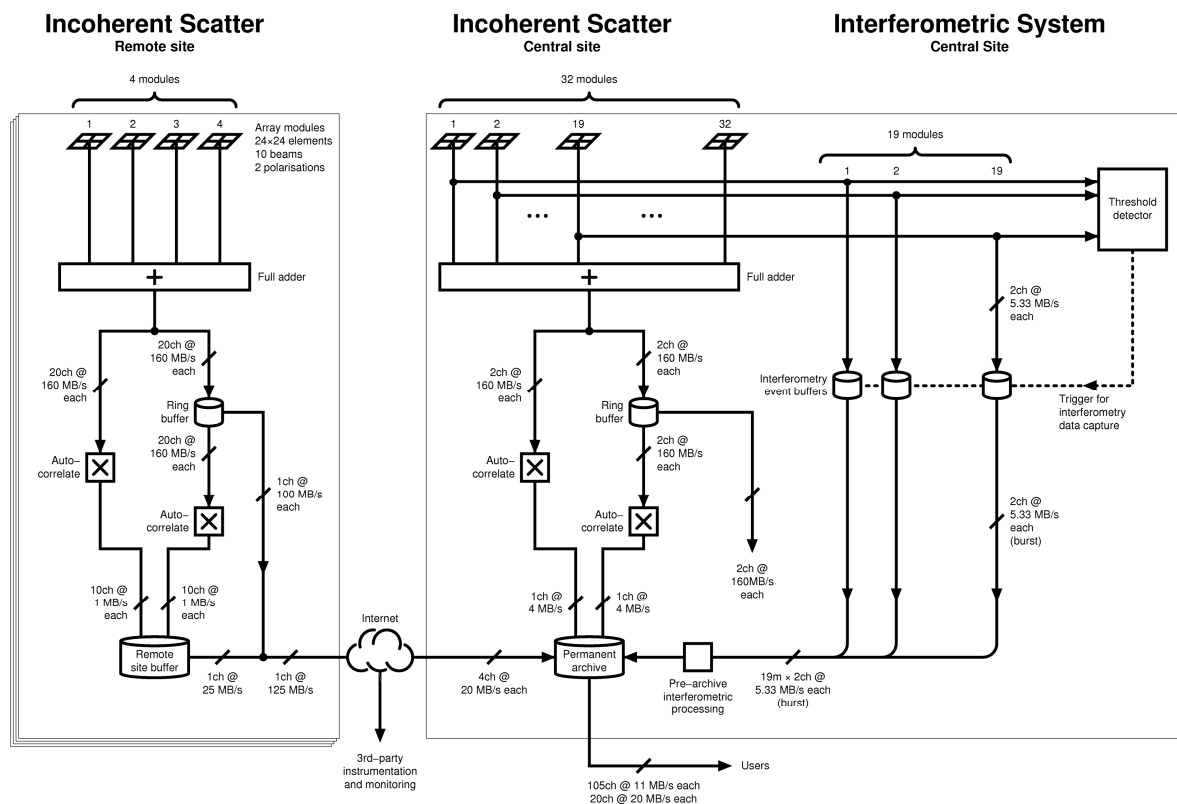


Figure 8.9.1: Topology of the data system for the EISCAT_3D radars. Note that in this context, a module is a collection of antennas in the phased array that are grouped together for both technical and scientific reasons. Unlike the interferometric part of the study (Section 8.8) which used modules of $7^3 = 343$ elements, the data processing was considered using square modules of $24 \times 24 = 576$.

8.9.2 Data rate challenges

The major challenge is that the lowest-level data products of EISCAT_3D will be so large that it is simply impossible to store them permanently. In current incoherent scatter radars such as the existing EISCAT system, a substantial amount of pre-processing is done before the data become available to science users. This reduces the data volume to an acceptable level for distribution and archiving, but also takes away much of the flexibility for subsequent data handling – e.g. because some level of auto-correlation and spatial/temporal integration has already been performed, making it impossible to construct the original sample-level data streams. In EISCAT_3D, we wish to retain this flexibility, as well as offering a number of new products, such as interferometric data, which are not available in standard experiments on other radars, and where considerable flexibility is needed to process the data optimally. Finally, the vision for EISCAT_3D is that the radar sites should become “solar-terrestrial observatories”, acting as magnets for large amounts of supporting instrumentation, some of which would also have high data rates and would benefit from using the same data system infrastructure.

8.9.2.1 Incoherent scatter data

As outlined in Sections 8.1 and 8.2, the system comprises a large central site of some 16,000 elements, all of which are capable of being sampled continuously on two orthogonal polarisation channels at 80 MHz with at least 14-bit sampling. In addition, there are four passive remote sites, each of around 10,000 elements, again sampled on two polarisation channels at 80 MHz with at least 14 bits. This means that each array element would be producing data at a rate of around 320 MB/s, giving a data rate of 9.6 TB/s for the whole of the central site.

While in principle the beam-forming for EISCAT_3D could be done entirely in software by operating on all this sample data, this idea was rejected at a very early stage in the design study because of the impossibility of manipulating such large quantities of data in real time. Instead, the beam-forming will be done by FPGAs (Section 8.6.1), so that, from the data system perspective, the lowest-level data will be a stream of sample-level beam-formed data at a rate of 320 MB/s per beam, or 27.6 TB per beam per day for the central site. At the remote sites, the useful data rate per beam is significantly lower, because it is limited by the pulse length and the intersection between the transmitter and receiver beams. It is assumed, however, that each remote site will use multiple fixed simultaneous beams, which means that the sample-level data rate from the remote sites might well be higher than that from the central site. We have assumed that the data rate for the remote sites will be around 800 MB/s per site. Integrating across all five EISCAT_3D sites, this gives a total of around 300 TB per day of sample-level data from the full EISCAT_3D system, of which around 30 TB comes from the central site and 70 TB from each of the remotes (see Deliverable 8.3, Section 3.2, [D8.3]).

Even at this level, the data rates are challengingly large for short-term storage; however storage of the full 80 MHz sampled data is unlikely to be necessary for most experiments. The assumption of 80 MHz sampling corresponds to a bandwidth of 30 MHz, sampled 33% above the Nyquist frequency. Although this bandwidth allows EISCAT_3D to cover the whole of the possible range of frequencies occupied by the plasma lines under all electron density conditions, the spectral features due to incoherent scatter only occupy a small percentage of this region. For example the ion line data, providing most of the information, could be bandpassed and sampled at around 200 kHz, as long as this passband was less than the modulation bandwidth of the sounding pulse.

Assuming a duty cycle of order 75%, this would reduce the data rate at the central site to around 140 GB/day (1.6 MB/s) for ion line data.

In experiments where the sampling resolution is dictated by the ion line spectral bandwidth, the data rate would be around 500 GB/day (5.7 MB/s) for each remote site (assuming ten simultaneous beams). The detection of plasma lines entails an additional amount of bandpassed data, which might be greater than that required for the ion line, since plasma line bandwidths can be of the order of hundreds of kHz. An additional complicating factor is that, because the location of the plasma lines varies with electron density, search algorithms will need to be established to provide a search range for the plasma lines. The design of such algorithms has not yet been considered in detail, and should be pursued at a later phase of the project. Nonetheless, in principle, the required bandwidth needed to search for the ion lines and plasma lines might be only a relatively small fraction of the full 30 MHz, as long as the modulation bandwidth of the sounding pulses is sufficiently small.

If we assume that the incoherent scatter data product temporarily stored in the ring buffers will consist of (conservatively) band-passed sample-level ion line and plasma line data from each site, the input data rates reduce to the order of 10 MB/s from each site, meaning that incoherent scatter data accumulates in each cyclic buffer at each site at a rate of around 0.8 TB/day.

In practice, however, it is very likely that the modulation bandwidth will set the required sampling rate in many EISCAT_3D applications, since the greatly increased figure-of-merit of the new radar will allow the possibility to use much shorter pulses than are possible with the current EISCAT system. For example, a coding scheme based on 1 μ s bauds would require 1 MHz sampling to cover the full modulation bandwidth, and even shorter pulses are likely to be possible. This would immediately increase the ion line data rate requirement to around 700 GB/day (8 MB/s) and the requirement for wider bandwidth coverage will be further increased by the use of multiple simultaneous frequencies.

The use of high-bandwidth modulations does not require any change in the architecture of the data system, but does imply a scaling in the necessary investment. As we show below, the kind of system we have been examining for use in the ring buffer is capable of holding almost a month's production of continuous EISCAT_3D data produced in a simple single frequency, spectral bandwidth-limited experiment. However as the baud lengths reduce and the number of frequencies increases, the latency time of a fixed-size ring buffer reduces. The ultimate ring buffer specification selected for EISCAT_3D is likely to reflect a trade-off between the desired resolution of the experiment and the length of time for which it is desirable to buffer the data, in order to be certain that they have been optimally analysed.

8.9.2.2 Interferometry data

In collaboration with Work Package 5, we have made some initial estimates of the bandwidth and storage requirements for handling the data from the interferometry system. In normal operation, it is assumed that interferometry will only be carried out the central site, where 19 baseline pairs will be used to make continuous interferometric measurements. These will be continuously tested for coherency, with sub-beamwidth structures only being present when the coherency threshold is exceeded. However, it is not sufficient only to store the interferometry data during periods when the coherence threshold is exceeded – ideally we would like to retain the data from periods immediately before and after this, so that EISCAT_3D can be used to investigate the rise and decay of sub-beamwidth structures. The philosophy of using temporary buffer storage, already outlined above, is well suited to such a goal.

In principle, we have calculated the input data rate from the interferometry system to be 202.9 MB/s, or 730.5 GB/hour (see Deliverable 8.3, Section 3.2.7 [D8.3]). This assumes that long pulses can be used for interferometry, so that modulation bandwidth issues do not arise. In addition, current interferometry work using the ESR suggests that sub-beamwidth structures will only be present for a very small proportion of the time. If we assume that only 5% of the total interferometry data will be useful (which may well be an over-estimate) the average interferometry data rate falls to 36.5 GB/hour. This would mean that interferometry data would be transferred to the cyclic buffer of the central site at an average rate of around 0.8 TB per day for subsequent processing into visibility and brightness functions, a similar rate to that needed for the spectral-bandwidth limited incoherent scatter data.

In our data system design, we have assumed that a storage area of order 60 GB capacity will be reserved for the continuous writing of interferometry data (enough to hold about 5 minutes of data from 19 baselines). This area will not be part of the ring buffer, but will be a separate interferometry buffer, perhaps even using RAM. These data will be continuously tested against the coherency threshold criteria outlined in Section 8.6.3. If the threshold is not exceeded the data will be over-written. However once the coherences on a given number of baselines begin to exceed the threshold, the content of the interferometry buffer will be continuously transferred to the central site ring buffer. Interferometry data will continue to be transferred to the ring buffer until the coherences have fallen persistently below threshold values for more than a user-specified period, providing ample time to examine the decay of the scatterers. In this way, the raw data from the interferometry system can be stored in the central site ring buffer for long enough for it to be optimally processed into visibility patterns and brightness functions, which will then become part of the permanent archive.

8.9.2.3 Data from supporting instruments

Part of the vision of EISCAT_3D is that its sites should become “solar-terrestrial observatories”, acting as magnets for a variety of supporting instrumentation which would share the same data system infrastructure. Hence we have attempted to evaluate the temporary storage and long-term archive requirements for a comprehensive collection of instruments distributed around the EISCAT_3D sites (see Deliverable 8.4, [D8.4]). In doing so we have assumed that complex instruments requiring manual intervention, or with particularly high data rates, will be restricted to the central site only, whereas instruments designed for long-period unattended operations, and with typically much lower data rates will be accommodated at the remote sites.

As might be expected, the most stringent requirements come from instruments such as high-resolution auroral imagers, where data rates of order 250 GB per day are possible during continuous operation. Such an operation is, however, relatively unlikely, since such images would only be of interest during dark periods with visible aurora. Nonetheless, during geomagnetically active winter periods, the data rate from such an instrument might exceed that from a single beam of band-passed ion line data from EISCAT_3D. In principle, the data from such an imager can be handled in a similar way to incoherent scatter data – i.e. they can be temporarily stored in the central site ring buffer in order to allow some pre-processing such as image selection or compression which might enable the data volume to be reduced significantly for storage in a permanent archive (either the EISCAT_3D central archive or an independent data store).

For all other instruments at the central site, the problems are significantly less severe. In some cases, this is because the instruments produce much less raw data, while in other cases the data set becomes smaller due to pre-processing carried out within the instrument’s own dedicated computer system. At the remote sites, the most complex supporting instruments (which tend to have the highest data volumes) have been deliberately omitted, and the contribution made to any one of the

remote site ring buffers from the remaining supporting instruments is therefore appreciably smaller. We have estimated (Deliverable 8.4, Section 5.4 [D8.4]) that the data rate arriving at the ring buffers due to supporting instruments will be around 350 GB/day (4.2 MB/s) at the central site, and around 100 GB/day (1.2 MB/s) at each of the remote sites. This compares to around 1.6 TB per day at the central site from the other EISCAT_3D data sources (0.8 TB each for incoherent scatter and interferometry) and around 0.8 TB per day (due to incoherent scatter only) at each remote site. A reasonable estimate is that, in total, the ring buffers will receive data from all sources at a rate of ~2 TB/day at the central site, and ~1 TB/day at each remote.

The proportionate impact of the supporting instrument data on the permanent archive is more significant, however. This is because the high-level data products derived from the supporting instruments are more similar in size to the low-level data. We have calculated (Deliverable 8.4, Section 5 [D8.4]) that the volume of permanently archived data from supporting instruments is likely to be of order 1 TB per week across all five EISCAT_3D sites, compared with around 4 TB per week for the data products from EISCAT_3D. In our judgement, this 25% increase in the growth rate of the permanent archive is acceptable given the potential increase in scientific return.

8.9.3 System requirements

As can be seen from the previous sections, both parts of the EISCAT_3D data system (i.e. the ring buffers and the permanent archive) need to be flexible enough to handle input data from a variety of sources (incoherent scatter, interferometry and supporting instruments), while at the same time being able to support a sufficient output data rate to cope with the demands of specialised processing algorithms developed by users, data copying or archive backup, and data transfer both within the data system itself and to external users.

Table 8.9.1 summarises possible input and output requirements for the various parts of the archive system, listing the sources of the data and its potential users/applications at each stage.

Data Type	Archive Type	Input Data Rate (TB/day)	Output Data Rate (TB/day)
Incoherent Scatter	Central Site Buffer	0.8	up to 126
	Remote Site Buffer	0.8	up to 126
	Permanent Archive	0.5	up to 76
Interferometry	Central Site Buffer	0.8	up to 126
	Permanent Archive	0.1	up to 76
Support Instruments	Central Site Buffer	0.4	up to 126
	Remote Site Buffer	0.1	up to 126
	Permanent Archive	0.2	up to 76

Table 8.9.1: Input and output data rate requirements for the EISCAT_3D data system

The input data rates for incoherent scatter and interferometry assume that the sampling is dictated by spectral bandwidth. For the output data rates, the specified requirement would allow the whole of the ring buffer to be read out in around 18 hours – this might be necessary following an extended network outage at a remote site. It is also assumed that third-party users will require the capability to copy data from the ring buffer into their own storage systems. The output requirement for the central archive reflects what is likely to be a generous assumption on the use of the archive by multiple simultaneous users, combined with regular backup and replication of the archived data.

As well as housing the permanent archive, we have assumed that the central site will host a substantial central computing capability, which will also be accessible to the users. This is essential

because processes such as data analysis (the reduction of the data to the main plasma parameters) involve going from very high volume auto-correlated or sample-level data, to very low volume files containing a relatively small number of parameters. While the analysis can be run on almost any platform, it is far more sensible from a network point of view to keep the analysis capability close to the archive, avoiding the need to transfer high-volume data sets to multiple user sites.

8.9.4 Hardware specification

In order to satisfy the requirements of the system specified above, Work Package 8 developed a list of hardware requirements for the EISCAT_3D data system, and then embarked upon a market survey to ascertain whether any existing vendors had systems capable of meeting our storage and data rate requirements. Our specifications were as follows.

For the short-term “cyclic buffer” storage:

- A minimum of 56 TB of short-term storage, available at all times.
- A minimum of four input channels and eight output channels, each capable of 160 MB/s or better on any one input or output channel.
- A minimum of 38 input channels and 38 output channels, each capable of 6 MB/s or better on any one input or output channel.

In principle, this configuration would allow the temporary storage of around 20 days of spectrally-limited data at the central site, and 40 days data at each remote. This ought to be sufficiently long to allow each data set to be optimally processed before it is deleted and also to guarantee data security at the remote sites in the event of an extended network outage.

For the long-term permanent archive system, the specified requirement was:

- (a) A minimum of 1 PB of usable storage.
- (b) A minimum of 400 TB of usable on-line storage.
- (c) A minimum of 600 TB of usable “near-on-line” storage (e.g provided by an automated tape system) or additional on-line storage.
- (d) Potential to extend at ~200 TB per year
- (e) Capability to accept input data of at least 311 MB/s over all input channels, and reading and outputting 928 MB/s over all output channels.
- (f) Capability to expand to 105 output channels, providing an output bandwidth of 1555 MB/s.
- (g) Capability to sustain data rates of 20 MB/s on 4 input and 20 output channels.

This configuration would allow the storage of around 5 years worth of EISCAT_3D data from all sources.

For both of the above systems it was required that the power drawn by the data system should be less than 300 kW, that system management software and monitoring tools should be available, and that at least 36 months of hardware and software support could be provided.

A significant requirement for the data systems at the remote sites is that they should be capable of long-period unattended operation. In the design study, we have assumed that the EISCAT_3D central site should be an active site in all senses, i.e. that there should be a local staff with engineering support close at hand, if not permanently on site. For the remote sites, however, it is assumed that there will be no local staff, meaning that the sites should be capable of long-period unattended operations. It is assumed that there will be regular visits for technical maintenance and

that, in the event of an equipment failure, an engineer would visit the site. Because the sites might be inaccessible for long periods in winter (e.g. due to snow conditions) the data system will need the capability to handle individual disk failures by “graceful degradation”, so that the system continues to operate until some critical number of failures has been accumulated, by which time it should have been possible to restore the system to a nominal state.

A further significant threat to the operation of the remote sites is the loss of internet connectivity for an extended interval. In this case, it will be impossible to transmit the archive data to the central site, meaning that such data will have to be stored locally until the connection is restored. This is possible by re-configuring the ring buffer storage so that, as the internet downtime increases, a progressively smaller amount of ring buffer space is allocated to temporary storage of low-level data, with a progressively larger proportion being allocated to the long-term storage of archival data. Only when the amount of space remaining in the ring buffer is insufficient for the real-time processing of the low-level data would it be impossible for operations to continue, so an important consideration was to make this period sufficiently long that network outage of such a length would be unlikely even in harsh winter conditions.

It is not seen as being desirable that the permanent archive system should all be acquired at the start of the project, since the archive is only expected to grow at around 200 TB per year, and a phased deployment would allow us to take advantage of future advances in storage technology. However it was required that an initial archive system should offer at least 300 TB of archive storage.

8.9.5 Suggested solutions

As noted in Section 8.4, the temporary “ring buffer” storage and the long-term “permanent archive” represent two separate sets of requirements, each of which has a different solution. However it is essential that the systems suggested for each are capable of inter-operating effectively. The easiest way of accomplishing this would be to contract the construction of the data system entirely to a single vendor; however our study has established that this is not strictly necessary, since sufficient standardisation exists to make us confident about the inter-operation of separate solutions from different suppliers, although a single supplier may still be desirable for ease of integration and maintenance.

8.9.5.1 Temporary ring buffers

In the EISCAT_3D system design, there will be one ring buffer data store at each remote site and one at the central site. Each of the ring buffers will in principle have a rather large number of input channels, since they must support data inputs corresponding to the different beams of the EISCAT_3D radar, the different baseline pairs of the interferometry system and the various other types of observing instruments deployed in support of EISCAT_3D which have a high-volume data rate. However, not all of these input channels would have the same data rate (see Section 8.1). In our study, we have examined two types of solution: one in which each input channel is kept separate, the other in which the number of input channels can be reduced by multiplexing the input data into fewer channels with a higher data rate. In principle, we believe this second option to be superior. One of the potential vendors (Union/Data Direct) has specified an example system with 20 initial input channels with double redundant links, each capable of 160 MB/s. These are then multiplexed into 8 links each with a bandwidth of 6 GB/s, and written to an array of 10 drive enclosures, each 50% populated with 300 SAS drives, each of 300 GB capacity. The resulting storage system has a capacity of 90 TB, of which 72 TB is actually usable storage capacity. The fact that the system uses a very large number of multiple disks offers a good configuration for graceful degradation, while the use of 50% populated drive enclosures allows the storage capacity

of the system to be doubled if a future need arises for additional capacity. The power consumption of this configuration is of order 10 kW and the non-operational temperature range of the components is between 85°C and -40°C, likely to be acceptable in all but the most prolonged extreme conditions.

8.9.5.2 The permanent archive

The permanent archive will exist at the central site only, and needs to comprise a mixture of on-line and near-on-line storage since it is assumed that, beyond a certain age, the data will no longer be in daily use. However, our goal has been to ensure that any user can have near-instantaneous access to any of the archive data for applications such as data mining and event searching. In our study we quickly identified one technology which seems very suitable for the central archive, namely a MAID system (Massive Array of Idle Disks). MAID systems comprise large disk arrays with a “sleep mode” feature which allows individual tiers of disks to spin down after a user-defined inactivity period, between 5 and 330 minutes, to save on power and cooling. One potential vendor specified a system very close to our requirements, with 1.2 PB of raw capacity (960 TB of formatted capacity) with 24 GB/s of backend bandwidth across 20 channels, RAID6 protection to allow mirroring and graceful degradation, 20 TB/hour backup speed and modular, redundant, hot-swappable components. The system is also highly scalable with the capability to support 1200 disk drives behind a single system, giving a capacity up to four times larger than traditional storage platforms. It is an open architecture, with the capacity to inter-operate with almost any commercial file system.

8.9.5.3 Network issues

As can be seen from the requirements specified above, the data system requires inter-site connections with speeds of order 10 Gb/s. It may be unrealistic to assume that these speeds will be available from the commercial internet, so that proprietary provision may be needed. As part of our study, we made a survey of current capabilities for data transfer, finding that internet hardware and protocols supporting 10Gb/s links were already available, though fairly new and not yet fully standardised (Deliverable 8.3, Section 6.2 [D8.3]). In addition, it is recommended that a second low-bandwidth link should be used as a backup to transmit house-keeping information in the event that the primary network fails. This would at least provide sufficient information to verify that the remote radars had been put into a safe state if the primary network fails. A range of options, based on technologies such as microwave links, satellite links and mobile telephony, already exist which could fulfil this requirement.

8.9.6 Summary

The data system Work Package has developed a two tier-design for the EISCAT_3D data system based on a combination of high-volume short-term storage at each of the EISCAT_3D sites, and a large permanent archive at the central site. We have shown that the kind of technology needed to implement these systems already exists, and is sufficiently scalable to accommodate the likely requirements of the new radar. At current prices this technology is expensive, but probably not the largest cost element of the EISCAT_3D system.

9 Additional Design Areas

As stated at the beginning of Section 8, we have broken our detailed description of the EISCAT_3D system down into two sections. In Section 8, we reviewed those parts of the system which the design study revealed to be the most challenging, and on which most time was correspondingly spent. In this section we consider the remaining design areas which, while still essential for the specification of a fully functioning radar, proved to be a less severe technical challenge. In some cases, such as visualisation and control systems, this was because it turned out that software already used by EISCAT, or existing elsewhere in the incoherent scatter community, could be adapted to do the required job. In other cases, such as the site infrastructure surveys, the requirement was always for a time-limited exercise to establish the conditions which pertained at the various possible remote site locations and whether these would be suitable for further consideration as the site for potential EISCAT_3D arrays.

9.1 Transmitter

The basic 3D transmitter performance requirements, as listed in the “EISCAT_3D Performance Specification Document”, are:

- Centre frequency: between 220 – 250 MHz, subject to allocation
- Peak output power: ≥ 2 MW
- Instantaneous -1 dB power bandwidth: ≥ 5 MHz
- Pulse length: 0.5–2000 μ s
- Pulse repetition frequency: 0–3000 Hz
- Modulation: Arbitrary waveforms, limited only by power bandwidth

Within this document, several centre frequencies in the range 220–250 MHz will be mentioned, owing to the availability of off-the-shelf components for testing and the pre-availability of the existing VHF facilities for Demonstrator Array trials.

As outlined in “Design Challenges” (Section 5), in an ideal 3D system these requirements should be met by fitting a self-contained digital arbitrary waveform/linear RF power amplifier unit (Figure 9.1.1) to each array element. The resulting system configuration would be capable of handling any arbitrary combination of phase and amplitude modulation, while at the same time providing the hardware functionality required for implementing both the time-delay beam-steering function and, if required, array aperture tapering on transmit. The arbitrary-waveform capability could also be used to suppress out-of-band emissions by applying tailored spectrum masks to all transmissions, something which might turn out to be required at all times in order to prevent interference to the active DAB channels immediately below and above the 3D spectrum slot. WP 6 has focussed its efforts on demonstrating that this “ideal” transmitter system can be implemented now, using components available off-the-shelf.

The arbitrary-waveform generator building block is clearly a key component. Since tens of thousands of these will be required, they must be available at low cost for the scheme to be practically and economically feasible. While the required functionality can be provided by a laboratory vector signal generator/random waveform generator attachment, such as is available from e.g. Rohde & Schwarz, this comes with a price tag in the order of 50 kEUR and is therefore out of the question.

However, a limited offering of affordable digital up-converter ASICs incorporating the desired functionality was in fact available already at the beginning of the Design Study in 2005. At that

time, no device was yet fast enough to be directly useable as the exciter for a 3D transmitter, but steady progress toward higher clock frequencies was being made. Since it was expected that devices capable of directly generating a modulated carrier at 230 MHz would be available by the end of 2008, the WP6 team has kept the development in this area under observation. A silicon subsystem offering all the required performance is now indeed available: Analog Devices, Inc. recently released the AD9957, a 1-GHz universal I/Q modulator and agile upconverter. This ASIC contains a direct digital synthesiser (DDS), an integrated 14-bit digital/analogue converter clock multiplier circuitry, digital filters, multi-chip synchronisation logic and other DSP functions. Operating at 1 GHz internal clock rate and 250 MHz I/Q data throughput rate, this device can directly generate analogue output signals at up to 400 MHz. Best of all, the chip is available off-the-shelf and it is also relatively inexpensive; current pricing is around USD 25 in 1000+ quantities! A practical solution to the transmitter exciter design problem thus appears to be within reach.

Unfortunately, because the AD9957 became commercially available at such a late stage of the 3D project, there has so far been no possibility to evaluate its capabilities in full. However, two AD9957 evaluation kits have been ordered and delivered and preliminary evaluation results will be included in Deliverables D6.2 and D9.2. Further work will concentrate on mating the AD9957 to the clock distribution and data transmission strategies developed in WPs 4 and 12 and to incorporate it in the test bed (see below) to enable full-bandwidth, random-waveform testing of the power amplifier chain.

Deliverable D 3.2 “Options for the Active Element” recommends that if the construction of the active element (the “Core”) must proceed in stages, at least 5000 elements should be deployed in a first phase. Since each element radiator has two independent radiating systems at right angles to each other, this minimum system will comprise at least 10000 transmitter modules. To meet the peak output power performance specification of ≥ 2 MW, the RF power amplifier in each module must therefore deliver at least 200 W peak output power. At 225 MHz, a range of current RF power semiconductor devices can do this with ease – at least on paper – and so it has been decided to aim for at least 300 W per module, preferably generated by a single device, to ensure that the 2 MW peak power target will be met by a comfortable margin.

However, almost without exception, RF power semiconductors rated for operation in the 220 – 250 MHz range are specified only for Band III analogue TV transmitter service. In this application, the devices are operated in class-AB and biased to run at a constant dissipation of about 50 % of their rated output. The RF power actually generated during one scan line is controlled by the video signal content and is typically very low; maybe 10 – 20 % of rated power. The output reaches its peak value only during the 4.7 μ s line sync pulse at the end of a scan line (every 64 μ s) and also during the 160 μ s frame sync pulse at the end of a half-frame (every 20 ms), corresponding to an effective RF duty cycle of only 7.3 %.

By contrast, a typical IS radar transmit-receive cycle is between 1 and 10 ms, that is, in the same order as (6...60) TV scan lines – but the RF pulse lengths are typically in the 0.5 – 2 milliseconds range, corresponding to an effective RF duty cycle of up to 20 %. Transmission is almost always at full saturated output power and the device is completely cut off between pulses. In this operating mode the device is subject to a cyclic thermal load, whose period is in the same range as the thermal time constant of the bonding wire – chip – substrate system, thus exposing the bonding wires to cyclic stresses that have been known to induce material fatigue and premature device failure. It is a matter of serious concern that no manufacturer data is available regarding the ability of devices to stand up to these operating parameters and thermal stresses.

The WP 6 team eventually concluded that the only way to gain full confidence in the suitability of a particular RF power device for IS radar operation at 230 MHz would be to run extended full power stress tests (several thousand hours) on a statistically significant sample (hundreds of devices), using realistic waveforms. An exercise of this magnitude is clearly outside the scope of the feasibility study, but regarded as essential in the procurement phase. At that point, prospective bidders for the target system contract would have to be requested to carry out such tests or present relevant statistics.

On the other hand, in line with the general policy established for the feasibility study, it was regarded as vital to demonstrate that at least one current device, available off-the-shelf, could meet the basic power output and power bandwidth requirements reliably with no degradation over a long period (ideally several thousands of hours).

For this purpose, contacts were established to the well-known semiconductor manufacturer Philips/NXP. Following some preliminary inquiries and clarifications, the 3D project received several free samples of the Philips BLF 248 VHF power FET for evaluation. This device has been in full production for several years and can be regarded as mature as far as its intended application goes. It is a push-pull silicon N - channel FET, rated at 300 W nominal CW output power up to 225 MHz when operated in class AB. According to the manufacturer's data sheets, a typical device can be expected to produce 350 W CW at 70° C operating temperature.

An RF test bed, containing three BLF 248-based single-stage amplifier modules based on a 225-MHz design provided by the device manufacturer, was constructed for the purpose of determining device performance and characteristics when operated under typical EISCAT_3D conditions. Two of the three BLF 248-stages can be run at full saturated RF output power, thus simulating the configuration to be used for an actual array element. A block diagram and a brief description of the technical details of the test bed are given in Figure 9.1.2. Full details will be found in Deliverable D6.2.

A comprehensive test program was started immediately once the test bed was ready. Unfortunately, because of the late start of this part of the study (it only got seriously underway once WP6 had been transferred to the IRF) the test program is still running at this time and will have to continue for some time after the end of the formal FP6 project in order to derive significant statistics.

Some of the most important items under investigation are:

- To establish whether the BLF 248 might be prone to thermal-stress-induced failures when operated with long pulses at full saturated output power; if intrinsic to the design, such “infant-mortality” problems would start to show up already in an extremely small number of devices,
- To analyse the heat-management aspects of the amplifier design; this is done with the aid of a number of PT100 thermal sensors embedded in the heat-sinks,
- Detailed phase noise tests of each individual amplifier. To identify whether switching noise on the drain supply affects the phase noise, tests will be run both with a switchmode PSU and with the amplifiers powered from storage batteries,
- Operation at reduced drain voltage will be tried to verify that the amplifiers can be operated at reduced output ($\approx 10\%$ of full power) with close to optimum DC-to-RF conversion efficiency without re-tuning.

After fine-tuning at 28 V U_{DS} and constant-carrier saturation drive, all BLF 248 amplifier units produce over 300 W output at 235 MHz.

So far, two long-duration simulated radar runs, employing 2-ms pulses and 25 % duty factor, have been performed. The first run, started on February 6, was made at slightly reduced power (≈ 250 W). When this had to be terminated on March 11 due to a scheduled mains power break, the amplifier system had been running with no breaks for almost 800 hours. A second run at the 300-W level is currently in progress.

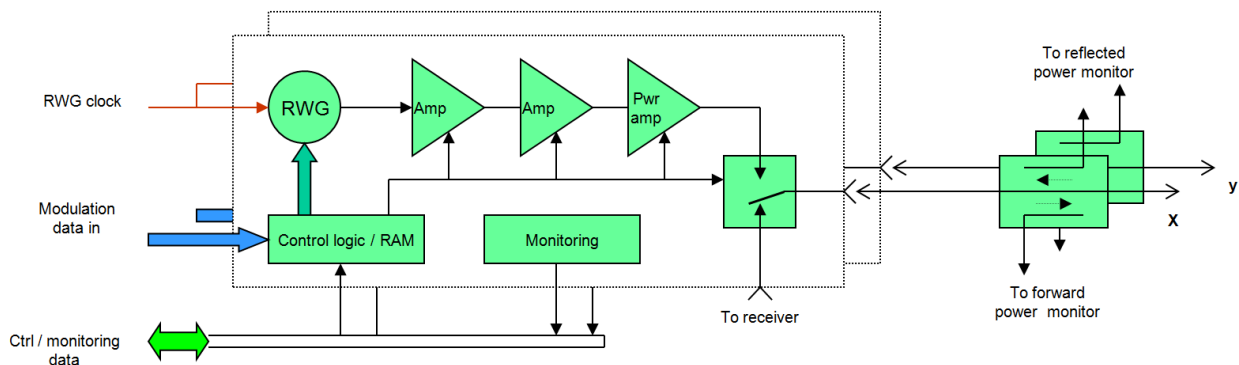


Figure 9.1.1: Schematic diagram of an element transmitter, comprising two identical units, one for each of the two polarisations. Each unit contains a digital random waveform generator/upconverter (RWG), two intermediate power amplifier stages, a 300+ watt power amplifier, a transmit-receive switch, a RAM bank and control and monitoring logic. Depending on the degree of mutual coupling between neighbouring element radiators, circulators and waster loads may have to be fitted between the power amplifier and the transmit/receive switch.

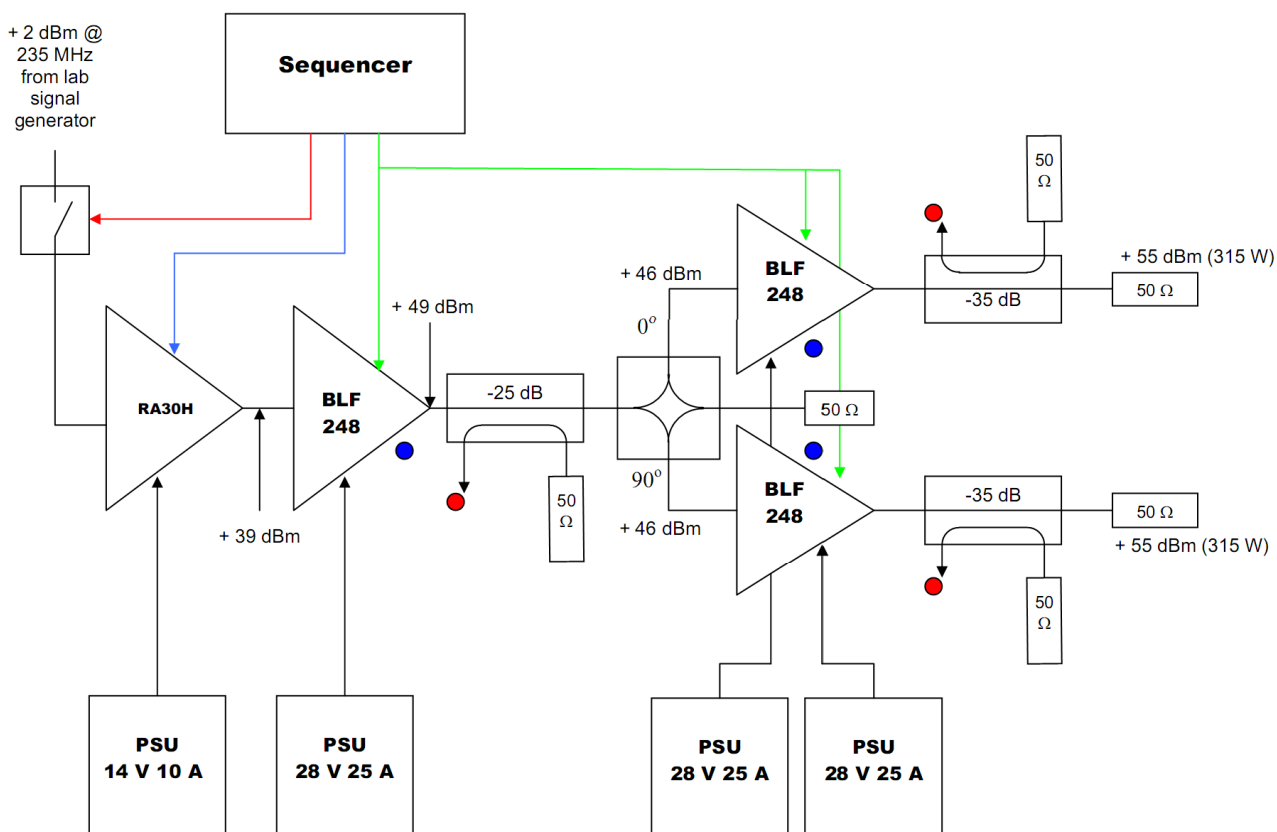


Figure 9.1.2: Block diagram of the EISCAT_3D 235-MHz power FET test bed. A +2 dBm carrier signal, generated by a standard laboratory signal generator, is on/off modulated by a PIN diode switch and amplified in a Mitsubishi RA30H2127M VHF power module to about +39 dBm. This drives a BLF 248 amplifier to about 80 W (+49 dBm) output, which is split into two 40-watt signals, driving two BLF 248 amplifiers to a peak output of 315 watts each at 1 dB gain compression. An EPROM-based sequencer generates the logic signals controlling the PIN diode RF switch and the gate biases of all amplifiers. Red (RF) and blue (temperature) dots indicate monitoring points.

Preliminary data from the phase-noise tests indicate that the amplifier chain adds only an insignificant amount of noise sideband power to that already present in the driving signal. Happily, this appears to be the case also when the amplifiers are powered from switchmode power supplies.

An unexpected problem was discovered during the long duration tests: The time constant of the voltage control feedback loop in the particular switchmode PSU used throughout the test bed is in the order of a millisecond. During a 2-millisecond pulse, this causes the drain voltage on the BLF 248 transistors to ramp first down by about 1 V and then, once the feedback cuts in, back up again, thus amplitude-modulating the RF output by about 7 %. While this in no way affects the serviceability of the RF amplifiers themselves, it indicates that when drafting system specifications to be used as the basis for an RFQ procedure, the required transmitter amplitude stability must be specified explicitly.

In summary:

- Integrated circuits combining an arbitrary-waveform generator with a direct digital synthesiser capable of generating a carrier signal directly at the EISCAT_3D operating frequency (235 MHz), thus meeting all performance requirements for the 3D transmitter exciter, are now available off-the-shelf at affordable prices. *The most difficult transmitter design problem (coming up with an exciter design that can be produced by the thousands at low cost) can thus be considered solved.*
- A basic RF power amplifier module, employing a mature, VHF power FET, has been operated for many hundreds of hours under long-pulse (2 ms) conditions at the 300-watt level. No device degradation or infant mortality has been observed. It has thus been demonstrated that *a workable 235-MHz power amplifier design, based on an off-the-shelf device and capable of handling the most demanding incoherent-scatter radar waveforms, exists and could be duplicated today if required.*

9.2 Internal data communication

More than 16,000 array elements will be required to approach the full 3D Core array target performance specification. In order to resolve the polarisation state of the return signal, which contains important information, each array element will be equipped with two direct-sampling receivers. Thus more than 32,000 individual high-speed data streams will be generated in the Core alone. The receive-only arrays will probably be smaller, but taken together, they will add up to a similar number of elements and data streams. Collecting, pre-processing and merging all the resulting raw data into a small number of beam-formed data streams is a formidable task whose magnitude is made painfully obvious by the following calculation:

An incoherent-scatter signal returning from the ionosphere contains several line complexes (ion line, gyro lines and plasma lines). Under high electron density conditions, these may be spread over as much as ± 15 MHz relative to the transmitter frequency. Irrespective of the receiver technology used, a real-valued sample rate of at least 60 Msamples/s is required to properly represent the information contained in this bandwidth. Each array element will therefore generate at least 120 Msamples/s (60 Msamples/s for each of the two polarisations). The actual sampling rate must be some 30 % higher than this minimum value to account for the finite pass-band-to-stop-band transition region of the analogue anti-aliasing filters preceding the A/D-converters. Accordingly, an effective sample rate of 80 Msamples/s/polarisation has been assumed throughout this study.

To maintain the required receiver spurious-free dynamic range of 70 dB, the effective sample word-length must be at least 12 bits, which implies that least 14-bit converters should be employed. To be on the safe side, 16-bit converters and a sample length of 16 bits have been assumed throughout the study.

The number of bits generated per second by each array element is therefore equal or greater than

$$16 \cdot 2 \cdot 80 \cdot 10^6 = 2.56 \cdot 10^9, \quad \text{i.e. } 2.56 \text{ Gbit/s}$$

While this may not appear unmanageable at first sight, the real magnitude of the data management problem becomes apparent when the number of array elements is factored in; the total number of bits produced per second then becomes

$$2.56 \text{ Gbit/s/element} \cdot 1.6 \cdot 10^4 \text{ elements} = 4.096 \cdot 10^{13} \text{ bits/s}$$

It would of course be possible to handle this huge raw data rate by brute force, sending the data from each element over a dedicated line to a central processing unit. But having considered the practical implications of going this route, the project team concluded that a much simpler and more cost-effective approach would be to reduce the data rate by a very substantial factor immediately following the individual receivers but before sending the data out onto the network. This insight contributed to the development of the hierarchical array architecture described in Section 8.2. In this scheme, the receivers for a hexagonal group of 49 dual-polarisation element antennas are housed in a common equipment container together with a sub-part of the beam-forming system that will combine the data streams from all 49 antennas into a small number (5 – 10) of partially beam-formed streams, thus effecting a data rate reduction by a factor of (0.1 – 0.2). The array-internal data communication problem is thereby effectively reduced to finding off-the-shelf solutions to two main problems, i.e. on the one hand how to solve the short-distance communications requirements inside each equipment container and on the other hand how to network the ≈ 300 equipment containers.

Since the signal from each antenna element is essentially noise-like and uncorrelated to the signals from all the other 48 antennas, the summing of all signals will cause the effective word-length to grow only by some three bits. Our initial estimate of 16 effective bits per sample thus holds also for the partially beam-formed streams. The total output data rate from each container will therefore be of the order of (10 – 20) Gbit/s, which is in a range that can be handled by a pair of standard multi-mode optical fibres.

Optical fibre technology thus emerges as the natural choice for the internal networking of the array. However, which protocol to use is not clear at first. A bewildering range of different industry-standard protocols with widely varying complexity, performance and overhead exist. Also, having settled on a protocol, the electrical output signals from the beam-formers in each container must be converted from parallel to serial format, media converted into optical signals and put onto the fibre; the process must then be reversed at the receiving end.

Work Package 12 has therefore focussed on a thorough survey of currently available industry standards in the areas of format conversion, serial protocols, media conversion and optical fibre cable. A comprehensive report is now available as Deliverable D12.2 “Array Networking and Communication Report” [D12.2]. A brief summary of some of the most important findings is given below:

- Today's electrical serial communications standards are designed around the *lane* concept. A *lane* is a set of differential pairs, one pair for transmit and one pair for receive. A *link* is a communication path between two devices. An xN link is composed of N lanes i.e. x4 link is composed of 4 lanes. Three current industry-standard protocols for box-to-box serial communication, viz.
 - 10 Gb Ethernet,
 - Fibre Channel and
 - InfiniBand

can provide sufficient bandwidth to handle the (10 – 20) Gbit/s rates out of each equipment container. *Serialiser/Deserialiser* (SerDes) units are available off-the shelf for all of these protocols. The SerDes converts data from parallel to differential-serial and from differential-serial to parallel at rates exceeding 10 Gbit/s (full duplex). SerDes units were initially used in box-to-box communication, but are nowadays also used in chip-to-chip communication. The basic building blocks of a SerDes are a differential transmitter/receiver, a serialiser/deserialiser, a parallel 8B/10B or 64B/66B encoder/decoder and transmitter and receiver buffers (FIFOs).

- All three protocols employ encoding of the primary data to guarantee enough transitions in the serial datastream for the transmitter/receiver to create/extract the embedded clock and to DC-balance the serial datastream; this serves to reduce inter-symbol interference. The actually achievable serial link transfer rate depends on which encoding method is used, but all three protocols can achieve at least 8 Gbit/s link bandwidths by employing multiple lanes.
- Since links are bi-directional by definition, implementing a link to an equipment container automatically establishes a high-bandwidth uplink data path from the control centre to the antenna elements, thus allowing the use of a fully digital transmitter exciter system where the desired transmitter waveform is uploaded to a digital upconverter/DAC unit as a digital sequence. As the bandwidth required for this operation is considerably less than that required for the off-loading of the received data, no special precautions should be required in the design phase.
- Optical transceivers for both multi-mode and single-mode fibre are available off-the shelf and can be interfaced to essentially any standard SerDes chip. It may become necessary to use single-mode fibre in certain parts of the system because of the distances involved (multi-mode fibre does not work reliably at distances exceeding 600 m).
- Because of the very high clock speeds involved, only a minute amount of mutual clock jitter (< 40 ps rms) can be tolerated and clock jitter cleaner circuitry may be needed at both ends of the fibre circuits. Alternatively, the receiver ADC clocks can be used.
- *All hard- and software required to construct a fully working array-internal data communications system exists now and can be had off the shelf; there are in fact several almost equally good options.*
- *It is however recommended that the progress in this area is carefully monitored and that no firm decisions are taken until the last minute.*

To gain some practical experience with setting up and operating a gigabit/s serial network, the Demonstrator array system has been wired up for optical data transfer. 58.0 meters of standard 50/125 μm multimode fibre (OM2) is used to establish six 1.5 Gbit/s full duplex optical serial links, three for each polarisations, between the array and the EISCAT site computer room. Texas Instruments TLK 1501 SerDes units are used to serialise the output data from the digital down-converters housed in the three equipment containers in the array field and re-convert it back to parallel electrical signals at the receiving end. The receiver frontend and the digital down-converter are controlled and configured via a USB interface running over a dedicated fibre. A USB extender is employed and supports full speed 12 Mbit/s USB 1.1 over the 50/125 μm multimode fibre. This setup is described in Section 10.3.3.

9.3 Control and monitoring

9.3.1 Introduction

At the existing EISCAT radars, the user-level Control and Monitoring (C&M) software system is called EROS (“EISCAT Real-Time Operating System”). It will be argued here that the present-day EROS could, and indeed preferably should, be used as the starting point also for EISCAT_3D C&M. At this stage of the project, it is premature to attempt to actually define the C&M commands for the subsystems of EISCAT_3D. For one thing, there are not yet sufficient details available to allow such issues to be realistically addressed. More importantly, from the point view of developing the user-level C&M system, it does not matter much what the low-level control programs of devices and subsystems actually do in detail, as long as the radar subsystems are “EROS-conforming”. An EROS-conforming subsystem or device has these properties:

- Engineering-level C&M software for the subsystem exists.
- That software is able to access all the necessary aspects of the hardware.
- That software provides programmable C&M in a form that allows interfacing to EROS in some “standard” way.

Long experience has shown that it is normally both fast and safe to integrate a conforming subsystem into EROS. Therefore, it is unnecessary, and even wasteful, to go into details about EROS commands for a subsystem too early. Even at the time when the project actually starts placing tenders for the various subsystems to the industry, it will not be necessary to have a detailed EROS-level interface ready. The requests for quotations should merely request that the engineering subsystems be EROS-conforming. The approach here is that the hardware comes first; the software will be adapted as necessary. This marching order has worked well for the existing EISCAT radars.

The work description of WP 7 stipulates a literature study to inspect *the current state of the art in real-time control systems architecture and software* in order to find if an existing control system *meets the specific requirements of the EISCAT_3D project*. Such a study was not made, for early on it became clear that the in-house EROS system can be enhanced to meet also the needs of the EISCAT_3D radar. Once it is accepted that EROS is able to do the job, the argument for actually letting EROS do the job becomes overwhelming, and any further search for a suitable C&M becomes moot.

9.3.2 Benefits of an EROS-type system

The basic assumption is that EISCAT_3D C&M system ultimately needs to be maintained and operated by EISCAT, even in the case that much of the hardware would be built and maintained by an external (possibly commercial) consortium. That is, it is assumed that it is EISCAT staff, together with the scientific user community, that is mainly responsible for the actual EISCAT_3D operation. Another assumption is that at the time of EISCAT_3D, there will be other, older EISCAT radars like the EISCAT Svalbard Radar, still in operation. With these premises, there is strong motivation to develop EISCAT_3D control and monitoring system gradually, starting from the existing EROS.

9.3.3 User support

The user-level C&M software is crucial for the usability of a radar system. The C&M software is the radar's user interface. With scientific radar, it is important that users with different levels of technical radar expertise can make use of the system. In the future, we expect many EISCAT users having less technical experience than is typical in the present day. This will make a well-tested and well-known user interface even more important in the future. It must be possible to offer simple ways of operating the radar for some users, by providing ready-made packages of operation, while also making it possible for experienced users to have full access to all features of the radar. The way this essential flexibility is achieved in EROS is by the use of a scriptable command language, called ELAN, which can control all aspects of the radar. It is possible for experienced users, including EISCAT staff members, to write ELAN command scripts for their technically less experienced colleagues. The expanding library of such scripts can be either used as such, or can be used as starting point for smaller or larger modifications. This flexibility should be provided by any EISCAT_3D C&M system. EROS has been built from the ground up to facilitate full programmatic control of the EISCAT radars.

9.3.4 Maintenance and development

Software developers responsible for the mission-critical C&M must be able to respond quickly. In a scientific context, the user requirements for the software can change rapidly. A new piece of hardware may become available, and must be integrated into the system, or an existing piece of hardware can develop a fault that still allows meaningful operation to continue, if only the control software can be suitably hacked. It must be possible to adapt the software quickly, cheaply, and without fuss, to whatever will be required by the hardware situation on the one side, and by the user wishes on the other side.

It is necessary for EISCAT to have good in-house understanding and complete in-house control of the C&M software. The idea of starting to fill web-forms of a commercial software service contract in order to get some obscure bug fixed, or some curious user request to fulfilled, while the paying scientist with her rocket on the launch-pad is biting her nails, is too nightmarish to contemplate. The EROS system has been built entirely in-house, and is very familiar to EISCAT software staff. EROS is implemented as a lightweight, loosely coupled system of interacting processes, where most faults will be well isolated and relatively easy to locate and correct. Even substantial problems can be repaired in-flight when an experiment is going on, without taking the system down. To preserve the flexibility of software service and development also for the EISCAT_3D, it seems desirable that all essential parts of the overall C&M software not only be designed but actually implemented by EISCAT staff, based on the existing EROS core.

A related advantage in adapting EROS as the C&M system also for the future EISCAT radars is that the regular software development and maintenance work going on at the existing EISCAT radars, will then automatically benefit the new system also. The present radars will act as a good testing environment for the software of the new system.

9.3.5 Unified interface

A key aim in EROS design has been to provide a uniform interface for the radar hardware across all EISCAT radar systems and sites. The purpose is that from the user's point of view, all EISCAT radars look and feel as much the same as possible, and feel part of a single system. This must be seen beneficial also in the future, only, to reign-in on the increasing hardware complexity brought by the EISCAT_3D, more so. As a relatively general-purpose, command-based radar control environment, EROS is quite good in supporting this kind of unifying approach. The devices and subsystems may be different from site to site and from system to system, but they all are still controlled by just some commands. Often both the command name and command parameters can be kept intact between systems, and only the internal implementation will need to be made system-dependent. In the very least, the command name can be kept intact, and only some of the command parameters need to be adapted for specific system. This kind of unification would be lost if one would use EROS at the old radars but something else at EISCAT_3D.

9.3.6 Assessment

EROS has proved to be a capable and rather convenient system for controlling and monitoring of all current EISCAT radars, but it is still also under continuous development. Over the 25+ years lifetime of EROS, there has been one complete pre-implementation from scratch, and in addition several major re-writes. And, as a smaller modification, support for some new EISCAT facilities has been recently added to the EROS system, such as the "E3D" module for the EISCAT_3D demonstrator array (WP 7.1). It is probably fair to say that the current version of EROS, EROS 5, is a reasonably mature, but in no way frozen, C&M system.

The envisioned EISCAT_3D radar, though advanced in terms of antenna and transmission hardware and receiver signal-processing performance, should not be so much qualitatively different from the present systems that it could no more be handled by EROS, by adding new modules. Things that *would* make life difficult for EROS would be anything that would necessitate the EROS core itself to handle "hard" real-time operating system, but there are no such needs in the present EISCAT systems, nor do there appear to be in the EISCAT_3D.

Nevertheless, in addition to being able itself to take care of the soft real-time needs, the EROS core is also capable, via special hardware assistance, to serve some crucial, hard real-time needs where events *must* happen at the specified time to be useful at all. This arrangement makes it possible for EROS to control those events in the transmission and reception that need to occur at a specified fraction of a microsecond. It appears that this combination of soft and hard real-time control is still sufficient also for the EISCAT_3D.

Thus, there seems to be no fundamental reason why the present, well-tried, in-house EROS could not be scaled up to meet the quantitatively increased, but qualitatively still pretty much unchanged, control and monitoring needs of the EISCAT_3D system as well.

9.3.7 Future development

During the design study, with the needs of EISCAT_3D in mind, EROS was updated from version 4 to version 5, by solving two critical, long-standing problems; after this update, all the currently known unknowns in the EROS implementation are now under control. One more major EROS update, EROS 6, is needed *before* the system is powerful enough to start incorporating EISCAT_3D subsystems. The update will make EROS console, which so far has been available only locally, available over the network. This is essential for EISCAT_3D that should be able to operate mostly unmanned. EROS would also benefit of stronger use of several new enhancements that have been taking place in the underlying open-source programming language system in recent years. A graphical user interface for radar monitoring is also being planned together with the EROS 6 update. Provided that the EISCAT_3D subsystems are delivered EROS-conforming as is suggested, they can be brought under EROS control as soon as they become available. It may be appropriate to upgrade the EROS version name to EROS 7 or EROS 3D.

9.4 Visualisation and data analysis

The visualisation needs for a complex scientific instrument like a phased array incoherent scatter radar can be divided in a number of ways. In the design study, this area was addressed as part of Work Package 8, and summarised in the final deliverable from this work package [D8.6]. The visualisation needs were first separated into technical and scientific needs, and then each of those into real time and post experiment visualisation. The next section describes some of the technical visualisation needed for monitoring a phased array radar, while the following section addresses the requirements for scientific visualisation.

9.4.1 Technical visualisation

To ensure proper operation and, more importantly, to notify the user of problems and errors, efficient visualisation of the system status is needed. This is obviously a most important real time need. Post experiment visualisation is also of importance for documentation of proper radar operation and detection of possible deviations from the intended experiment.

Provided the real time visualisation need is met, the post experiment visualisation needs should be satisfied, as long as the system status data is stored. Since a modular computer controlled phased array radar is more complex than radars with transmission from a single parabolic antenna, the visualisation need for system monitoring and status surveillance is also more complex. For single-dish radars it might be sufficient system monitoring to record and display a few scalar parameters such as antenna pointing direction, transmitted power system temperature/noise level and control transmission pattern and data flow. For a phased array radar, however, the system monitoring must be done at the transmitter/receiver level. With as large a number of transmitters and receivers as will be used in EISCAT 3-D the system monitoring data have to be visualised in a clear and understandable way in order to be useful. Here we can draw heavily from the experience gained from the Poker Flat Incoherent Scatter Radar (PFISR), the first of the Advanced Modular Incoherent Scatter Radars built and operated by SRI International. Even though the technical design of EISCAT 3-D will differ from PFISR, both systems are computer controlled phased array radars, and thus share most system characteristics relevant for visualisation of system status. Figure 9.4.1 shows one of the system monitoring visualisation used at PFISR, displaying the transmitted power. Further examples, showing the radar beams available and used by the current experiment and the ambient temperatures can be found in [D8.6].




Figure 9.4.1: The PFISR real time monitoring display shows transmitted power of the individual transmitters. This is an efficient mode to display the system status of a phased array radar. Figure courtesy T. Valentic, used with permission.

The AMISR/PFISR visualisation software is open source and available and appears to be a good building block to adjust and modify to fit the design of EISCAT 3-D. Since the beam shape and pointing of a phased array are obtained by controlling the phase and power of the individual transmitters it would be preferable to log the transmitted antenna patterns. At the HAARP ionospheric heating facility, the antenna pattern is logged and displayed in real time, one such beam pattern is shown in Figure 9.4.2. This is feasible for a high power HF radio wave transmitter where there are one or two simultaneous beams that typically change direction on time-scales of 10s of seconds. For an incoherent scatter radar switching between multiple beams at time-scales of 10 ms it is obviously unfeasible to do in real time. However, it might be worthwhile at a reduced cadence. One could then use a time averaged antenna pattern in combination with its variance to show the stability of the radar beams.




Figure 9.4.2: Two HAARP antenna patterns are displayed. The one to the right has a minimum in the centre and is used to transmit a beam with angular momentum. Figure courtesy of T. B. Leyser, used with permission.

9.4.2 Scientific visualisation

In the case of scientific visualisation, there might be a more distinctive difference between the real time and post experiment visualisation needs. The most important objective with real time visualisation is to support decisions, for example how to run external instrumentation, and possibly if or when to launch rockets. The natural evolution from the current real time visualisation of analysed data, generally a figure with updating panels showing electron density, temperatures and ion velocity, is to replicate those data for each beam. This could be done with one figure for each parameter with panels for each beam, as is currently used at PFISR, or with one figure per beam and panels for electron density, electron and ion temperature and ion velocity. The preferred layout will of course depend on the experiment at hand and will also depend on personal preferences and should if possible be left to the user. For the current EISCAT “RT-graph” displaying back-scattered power, ion and plasma line spectra, the same extension could be made, again with either one figure per beam displaying all different parameters or one figure per parameter displaying for example the ion line spectra for all beams etc, with similar layout considerations.

Another design feature of the real time visualisation system should be to allow the users to link or plug in tailor-made data visualisation functions, but more importantly functions to signal for special observational conditions. The reason this functionality will be needed is maybe best explained by an illustrative example: Occasionally the EISCAT radars observe Naturally Enhanced Ion Acoustic Lines (NEIALs) that are orders of magnitude above the thermal back-scatter. At times these echoes give spectra with large cross correlations when observed with the two EISCAT Svalbard Radars. This makes it possible to use interferometric techniques to determine angle-of-arrival. With EISCAT 3-D it will be possible to use interferometric imaging to determine the horizontal structure of these echoes. A signal monitoring the total back-scattered power at altitudes above 500 km could be used as a trigger for NEIAL events after which data from separate sub-arrays could be stored.

This simple signal, and others we might design, could of course be implemented during the development of the visualisation system. However, since research will always be the study of what is at best poorly understood, it will always be impossible to accurately predict what might become signals of future interest. Two requirements for this to work are obviously that the visualisation system is built modularly, and that the data formats are self-contained.

The post experiment visualisation aims at presenting the physical parameters of interest in a clear and understandable manner. The most straightforward visualisation is to use figures with a similar format to the current EISCAT standard plots. However, plots with panels displaying the parameters of interest as function of time and altitude stacked together beam by beam might not be sufficient to properly display events where there is horizontal motion of spatial structures. For such conditions, or when several radar beams are used to build a three-dimensional image of the ionosphere, the best approach would be stacking the beams together in a three-dimensional block and show selected slices and cuts, as in Figure 9.4.3. Such displays nicely show the spatial variation in the selected cuts, and if extended to animations with respect to time, one can also see the temporal variation.

These stacked plots might, however, miss interesting evolution that happened between cuts, so the user must take proper care when selecting cuts. For plots like these, one can also use cuts that follow some characteristic parameter. For example, one might also chose a cut at the altitude of peak E-region electron density. As time progresses, the altitude of the cut will then vary but one is assured to see the spatial variation of the E region peak. When one wants to combine the information on several spectral features colour coding can be used. This makes for displays that are very rich in information. For EISCAT 3-D where the time variation of the back-scattered spectra

can be measured in the full three spatial dimensions we will get a data product that is 5-dimensional (x; y; z; f; t) and we will have to resort to displays which reduce that dimensionality.




Figure 9.4.3: Electron densities at three horizontal cuts and one vertical, displaying the E region electron density in and around an auroral arc. Figure courtesy of J. Semeter, used with permission.

Volume rendering techniques have largely been left outside this brief overview of visualisation methods. The main reason for this is that volume visualisation with transparency/ translucency is good for qualitative overviews of an event. However, experience of working with three-dimensional imaging of aurora and radio induced optical emissions (also known as artificial aurora) has shown that the human visual perception is not well developed to interpret three-dimensional intensity variations. We have adapted to see motion, contrast, and the three-dimensional location of surfaces of nearby objects with stereoscopic vision, but not to see the three-dimensional density variation of conditions such as fog (Private communication: Bill Hibbard, creator of *vis5d*, a leading visualisation package for high-dimensional data: <http://vis5d.sourceforge.net/>). Volume rendering methods do produce nice looking overview animations. For example, see www.eiscat.se/groups/EISCAT_3D/WP8/d86/arc_movie.gif (showing an animated volume rendering of a section of an auroral arc as the observer views the arc from different angles) and www.eiscat.se/groups/EISCAT_3D/WP8/d86/emissions_movie.gif (showing the temporal evolution of radio induced optical emissions above EISCAT Ramfjord, with the viewer located at a fixed point looking eastward.).

9.5 Site infrastructure requirements

9.5.1 Introduction

Work Package 13 has addressed issues related to the preparations for the next phase of the EISCAT_3D project in so far as they are relevant to the Design Study.

In addition, a team of engineers from the Kiruna and Sodankylä EISCAT sites performed a site survey during the autumn 2005 and spring 2006. The work was done as part of WP2: Evaluation of design performance goals.

The site survey resulted in a tentative model of the EISCAT_3D geometry. The central core (transmitter site) is assumed to be located near the present EISCAT site at Ramfjordmoen, Norway, and the receiving sites near Porjus and Abisko (Sweden), Säytsjärvi (Finland) and Masi (Norway). However, recent information indicates that the city of Tromsø have plans to allow industrial facilities to locate in the Ramfjordmoen valley. The establishment of an EISCAT_3D site is dependent on an interference-free surrounding over a long period of time (min. 20 years). An alternative location for the central core is near Överbygd, 80 km south of Tromsø. The exact locations of the EISCAT_3D sites will not be decided within this Design Study.




Figure 9.5.1: One possible layout of the EISCAT_3D system.

9.5.2 Summary of the site survey 2005

The objective of the Site survey was to find potential construction sites for the EISCAT_3D radar system, on the basis of the EISCAT_3D Design Specification Document:

- A central transmitting/receiving core, located close to the EISCAT Tromsø radar site at Ramfjordmoen, Norway.
- Two receiving facilities, located at ground distances of 90-120 km roughly south and east of the transmitting facility, respectively.
- Two receiving facilities, located at ground distances of 220-280 km roughly south and east of the transmitting facility, respectively.

The areas near Abisko (Sweden), Masi (Norway), Inari (Finland), Porjus (Sweden) and Tromsø (Norway) were selected as interesting for the survey. A hand-held GPS unit and digital maps were bought in order to map the potential sites.

At each location, the search criteria for a potential antenna site were:

- For the central core — a 300 by 300 meter area, relatively flat and dry, with possibilities to place smaller interferometry antenna arrays at roughly 120 degrees angular separation and extending out to a distance of about 1 km from the system midpoint.
- For the receiving sites - a 300 by 300 meter area, relatively flat and dry, and with an incline of 2-20 meters, in the direction towards Tromsø.
- The absence of TV/radio transmitters and cell phone base stations in the neighbourhood.
- Infrastructure – data communications, power supply

In parallel to this, the field team made RFI measurements at the potential EISCAT_3D sites. One objective with the measurements was to investigate if there are interference signals in the frequency range 230 – 240 MHz. The other objective was to measure the field strength from the TV/Radio transmitters in the area, and collect basic data for the radar receiver design engineers.

The measurements were done using EISCAT-owned equipment:

- Spectrum analyser: Rohde& Schwarz FSP7
- Omnidirectional antenna: Discone Diamond D-130J, 25-1300 MHz
- Log-Periodic antenna: Create 5130-1N, 50-1300 MHz

The results of the measurements are presented in the EISCAT_3D web site, which includes a presentation of the four potential EISCAT_3D sites that were investigated by the site survey team [D13.1].

9.5.3 Infrastructure

To keep the construction costs down, it is important to have a good infrastructure in place in the area of a construction site. Access to a backbone fibre link and a power line, not too far away from the site is an advantage.

9.5.3.1 Data communications

Very high speed data communication links are essential to the operation of the new radar. Initially, a minimum of 5 Gbit/s bandwidth will be required, this will rise to >10 Gbit/s as the system is expanded into the full configuration.

9.5.3.2 Power supply

The average power consumption of the transmitter site can be estimated at 5-7 MW (10 MW peak power output, 50% efficiency, 20% RF duty cycle + some reserve capacity). Access to a power line close to the site is an obvious advantage, but we must accept that a branch line may have to be constructed.

9.5.4 RF environment

It is essential that the RF (Radio Frequency) environments at the potential construction sites are free from interference. The radar receivers are extremely sensitive to interference from external sources. The external sources that can interfere with Incoherent Scatter Radar measurements are numerous and include:

- TV/Radio broadcast emissions.
- GSM/UMTS 900 base stations.
- Aeronautical and satellite communication.
- Radio equipment.
- Consumer electronics: computers, wireless systems, telephone equipment, etc.

Besides these main concerns, secondary concerns as sources of RFI (Radio Frequency Interference) can be high-voltage power lines and motor traffic: cars, snowmobiles, four wheelers, etc. The experience from EISCAT UHF (at 930 MHz) operation shows that a snowmobile can interfere with the radar measurements if it is within a distance of 1 km.

To conclude, one must avoid contact with human settlements. A good construction site should be located away from any populated area, preferably, outside the field-of-view from any town or village.

Extensive RFI measurements must be carried out at the potential construction sites. This includes measurements of the field strength and in-band interference. The field strength measurements aim at locating the strong signals that could potentially cause saturation in the EISCAT_3D receivers. Strong signals in the antenna side-lobes can also cause intermodulation generated products in the receiver front-end. The in-band interference measurements aim at locating weaker signals, inside the frequency band of interest. The RFI measurements require high sensitivity and well-calibrated instrumentation.

9.6 Novel science applications

9.6.1 Introduction

Work Package 10 of the EISCAT_3D design study was set up to identify and evaluate any non-traditional or new uses that potentially could be implemented with the EISCAT_3D radar. This work package was originally intended to produce three deliverables. However an extra deliverable was added during the course of the study, to examine the contribution which EISCAT_3D could make to the study of global atmospheric change.

9.6.2 Identification of novel uses

In the initial phase of this work package, a search was undertaken for recent scientific papers dealing with non-traditional uses of geospace research radars. Based on the results of the search, a list of possible new uses was constructed, with a special focus on global climatic change. In total over 200 documents were processed of which 150 pertained to climatic change and 50 to non-traditional uses of incoherent scatter radars. As a result of this search, the evaluation of the extent to which incoherent scattering radars, or similar, have been employed for climatic studies became an important aspect of the work package.

Apart from global climate change, the following major subjects were identified and their prospects analysed. The results are summarised in Table 9.6.1, below.

An assessment of the desirable technical specifications of the new radar was then undertaken to identify what characteristics the EISCAT_3D radar would need to possess to make contributions in the above areas. The resulting specifications did not deviate considerably from those produced in the survey conducted by Work Package 2.

In general, this part of the study demonstrated how difficult it is to find entirely new uses. Most imaginable new applications of IS radars have been tested in one or another form, however the potential for climatic studies provided by long-term EISCAT data appears very high and promises new important scientific results, in particular after the new radar is put into operation with its superior specifications and the prospects of continuous operations.

Subject	Future Prospects
Improving ionospheric correction models	Feasible
Artificial ionospheric targets	Feasible, warrants further investigations
Space debris	Feasible
Planetary radar	Feasible
Magnetospheric radar	Seems feasible – requires high sensitivity and longer integration intervals. Requires further investigation.
General relativity theory	Requires exceedingly high accuracy. Seems unfeasible but requires further consideration.
Meteor radar	Feasible
SETI (search of extra-terrestrial intelligence)	Feasible, requires specialised software and processing

Table 9.6.1: Potential new uses of EISCAT_3D and an assessment of their prospects.

9.6.3 Improving ionospheric correction models

Plasma turbulence in the ionosphere has significant importance for communication, navigation and surveillance systems based on trans-ionospheric radio links, since radio scintillation caused by electron density irregularities in the ionosphere may result in signal degradation and outage. In addition, some remote sensing techniques may also experience severe disturbances. With the help of ionospheric models it is possible to correct for these disturbances up to certain levels. However, these models do not represent very well the real ionosphere at high latitudes, since models have been built with the use of empirical data obtained at mid- and low-latitudes.

Because of this, some effort was spent in investigating the feasibility and utility of employing long time-series incoherent scatter data to improve the integrity of trans-ionospheric radio communication signals, especially at high latitudes. A critical area in which this contribution can be important is in improving the ionospheric models used by the GPS and the Galileo global navigation satellite systems. The second report produced by WP10, [D10.2a] gives an overview of satellite navigation systems and remote sensing satellite systems. In addition, the physical basis of the Total Electron Content (TEC) is introduced, as well as different ways of retrieving other ionospheric parameters that play a role in the degradation of the radio signals. Finally, a review is presented of previous work concerning the measurement of electron density by means of incoherent

scatter radar, in particular by EISCAT, including a comparison with electron density measurements obtained by other instruments and existing models. The overviews and reviews presented in this report justify why this should be an important area of work for EISCAT_3D, and lay the groundwork for future studies.

9.6.4 Use of long time series data

Work Package 10 identified a number of important uses of EISCAT data which would be made possible by the analysis of high-quality long-term data sets. These included studies of global atmospheric change, radio wave propagation, satellite communications and global navigation and positioning systems. In order to facilitate this type of study, some time was spent in developing a software tool capable of retrieving and processing data from the long-term EISCAT database. Although the tool was developed using existing EISCAT data, its use will be equally applicable to data from EISCAT_3D.

The aim of the software package was to identify long-term trends in the data and to facilitate the development of an empirical model of high-latitude electron density in order to improve on current models, which are poorly characterised at high latitude. Models of this kind have potentially important applications for:

- Climatic studies in the upper parts of the atmosphere
- Corrections for satellite data especially Galileo and GPS
- Space weather studies
- Correction of SAR remote sensing data

A software tool was developed to operate on the existing Madrigal database (<http://www.openmadrigal.org>) which holds a wide variety of data from EISCAT and other incoherent scatter radars. The python programming language was chosen for the development, because it is internet compatible, object oriented, open source- and standards based, platform independent, and non-proprietary. The resulting program is highly modular and can be easily scripted to execute many common necessary tasks such as acquiring particular data-sets, and integrating the acquired data with given strategies. Initially, data acquisition was made via internet connections, but this proved to be too slow given the sheer size of the data-sets. The solution was to create a local database, and this proved to be fast and effective. Using the software, the entire dataset of electron density was acquired and processed in various ways in order to have synoptic overviews over one (Svalbard) and two (mainland) solar cycles, with different binnings to emphasise different trends of hourly, daily, seasonal, and solar cycle behaviour.

An example is shown in Figure 9.6.1, which shows the seasonal variability of electron density as a function of altitude over one solar cycle covering from 1997 to 2006 over Svalbard [Bel07]. The figure shows in a striking way the variability over height, season and phase of the solar cycle. Subsequently these data were compared to the predictions obtained from the IRI model (International Reference Ionosphere), showing that the model deviated significantly from reality. Further examples, together with more information about the development of the software and the potential uses of long-time series EISCAT data, are given in the third deliverable from this work package [D10.2b].




Figure 9.6.1. Electron density over Svalbard integrated over 3 hours and binned by season. The whole EISCAT Svalbard Radar data until the end of 2006 was used in compiling this figure.


9.6.5 Ionospheric signatures of global change

The increase of anthropogenic greenhouse gas emissions in the atmosphere is expected to cause an increase of the global temperature in the lower atmosphere and the oceans, while a decrease is expected in the middle atmosphere and ionosphere [UI100]. For instance, Basseur and Hitchman [Bra88] anticipate a cooling of the stratosphere by 8-15 K, whereas Roble and Dickinson [Rob89] predict the mesosphere and thermosphere to cool by 10 and 50 K respectively, for a doubling of concentrations of carbon dioxide and methane.

These temperature changes are much larger than those in the lower atmosphere, and some of their consequences should therefore be easier to detect. Among others, changes in the electron concentration are expected [Ris90] that can be parameterised by the values of maximum plasma frequency in the E- and F-regions (foE and foF2) and the height of the F-region peak, hmF2. Earlier work on this topic was reviewed as part of this Work Package, and an assessment has been made as to what and how EISCAT_3D can contribute to studies of Global Change [D10.6].

The long-term time series of ionospheric parameters already organised as synoptic data (see the previous section) over one and two solar cycles can be utilised for global change studies, showing that it is feasible to employ incoherent scatter radar data for this purpose. A few important issues have to be taking into account, however. To ensure data quality, it is advisable that changes in equipment and experimental setups (such as modulation codes) should be kept to a minimum. Long-term planning would minimise the impact of changes made to software and hardware on the small residuals that represent the trends that can be attributed to global change.

PART C
PROOF OF CONCEPT
THE
DEMONSTRATOR
ARRAY



10 The Demonstrator Array

10.1 Introduction

The EISCAT_3D Demonstrator system, located at the EISCAT Kiruna site in Sweden, has been entirely funded by EISCAT. The main purpose of the Demonstrator is to serve (in conjunction with the existing VHF system in Tromsø) as a test bed which can be used to validate a number of techniques central to the overall EISCAT_3D design concept (see also Section 9.2). These techniques include, among others,

- array electrical and optical characteristics,
- digital beam steering,
- multi-beaming,
- adaptive pointing correction, and
- adaptive polarisation matching.

In the EISCAT receiver, the 24 streams are initially stored into six so-called buffer memories. From there, they can be processed with the standard EISCAT signal processing machinery.

The Demonstrator system has proved invaluable, allowing the design teams to refine and validate their work in many parts of the project, giving confidence in both the appropriateness and the practicality of the evolving design solutions. Much of the project description reported in this document is therefore constructed around interactions with this system.

10.2 Technical overview of the Demonstrator

The Demonstrator system consists of two parts: (1) a small phased-array antenna of 48 crossed-Yagi elementary antennas, arranged into a 12 rows by four columns rectangular grid, and (2) a receiver system connected to the array. This allows prototypes of individual parts of the complete EISCAT_3D receiver subsystem, such as the sampling receiver and optical fibre interconnections, to be evaluated in an environment approximating the target system. A block diagram of the system is shown in Figure 10.1.

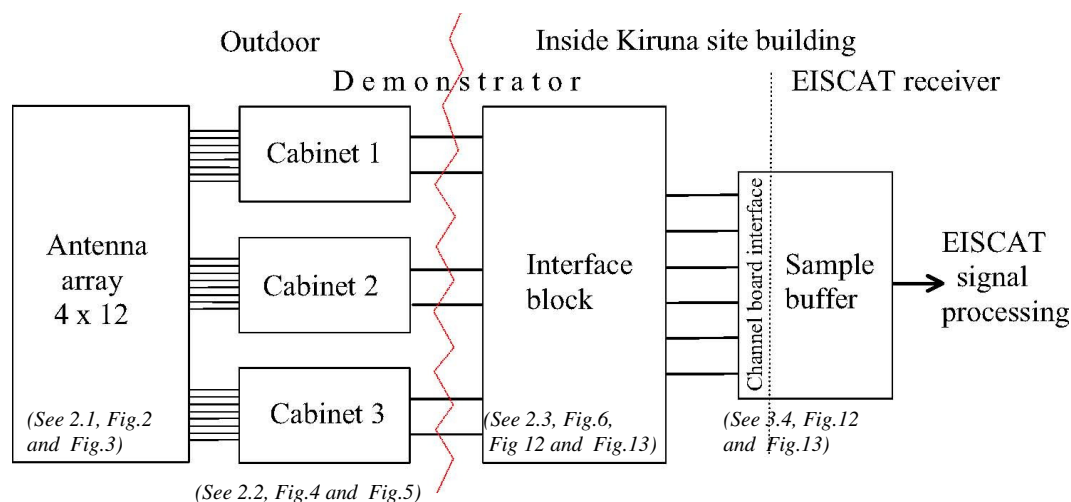


Figure 10.1. Simplified block diagram of the Demonstrator system.

The antenna array has been constructed on the west side of the EISCAT Kiruna site building. It is a 12×4 rectangular grid of (6+6) element X-Yagi antennas, measuring 24.6×6.6 metres. The long axis is oriented towards the Tromsø EISCAT site. The Yagis are grouped into three 4×4 element sub-arrays. A small equipment cabinet is located in the middle of each sub-array. Each cabinet houses hardware for eight signal paths. Each path comprises a front-end/receiver, an analogue-to-digital converter, a digital down-converter, and data transfer equipment. Once inside the site building, the sample streams from all cabinets enter an interface block, where they are converted from optical to electrical signals, de-serialised and transferred into a sample buffer. From there, the data are read into the EISCAT UHF receiver back-end and processed using standard EISCAT signal processing and visualisation software. All time-critical hardware and data transfer functions, and most set-up functions, are handled by the EISCAT UHF Radar Controller and the EROS operating system, thus making the Demonstrator system appear to the operator as just an additional antenna or signal path, equivalent to the existing UHF signal path. A set of Demonstrator-specific commands has been implemented into EROS for this purpose.

10.2.1 Antenna array – details

Figure 10.2 gives a schematic overview of the antenna array (see also Sections 8.1.4 and 8.2.7). The element antennas have a gain of 11.5 dBi and a -3 dB beam-width of about $\pm 25^\circ$. The four antennas in each row are separated by about 1.6 meters, as recommended for maximum E plane gain by the manufacturer. The row-to-row distance is about 2.0 meters. All Yagis are pointed northwards and elevated by approximately 55° from the horizontal, with the two sets of linear elements oriented at $\pm 45^\circ$. Signals from all four antennas in each row are combined in two analogue power combiners, one for each linear polarisation. There are thus a total of 24 output channels from the array. Obviously, with this arrangement digital beam-steering is only possible in the vertical plane.




Figure 10.2. Schematic side and top views of the 224-MHz Demonstrator array. One set of 4-to-1 power combiners is highlighted in blue and red.



Figure 10.3. The Demonstrator array in a winter setting.

10.2.2 Equipment cabinets

The Demonstrator array is divided into three 4×4 element sub-arrays. An insulated and ventilated instrumentation cabinet is located in the centre of each sub-array (Figure 10.3). This is a small hut, about one metre high, with base area of about 0.7 square metres. Each cabinet houses a standard 19-inch instrumentation rack. This has simplified the assembly and integration of the receiver system greatly, since most of the receiver and signal processing subsystems are built to the 19-inch standard. A close-up of a fully instrumented cabinet is shown in Figure 10.5. In order to simplify the initial laying and eventual upgrading of the cabling connecting the array to the EISCAT receiver, timing and control systems, underground cable ducts have been installed between the cabinets and the site building.

The hardware in each cabinet serves four antenna rows and comprises:

- two sets of four front-end receivers each,
- a serialising systems and optical fibre transceivers,
- control and timing devices,
- power supplies, and
- a heater.

Each front-end receiver set comprises four (low-noise amplifier → ADC → digital down-converter) signal chains, such that each cabinet can process the altogether eight signals generated by four rows of element antennas (see Section 8.5). Since the two signals from each row represent two polarisation components of the received signal that are not parallel to the local vertical or horizontal, they are labelled X and Y for convenience.

The system is set up such that one of the receivers handles all X polarisation data and the other receiver all the Y polarisation data. In the last stage of the digital down-converter, four decimated data streams, representing the band-limited X or Y data from four antenna rows, are multiplexed

into a single serial data stream, which is then media-converted and transferred into the site building over an optical fibre link. The arrangement is illustrated in Figure 10.4.




Figure 10.4. Cabinet hardware block diagram

Loading and control of the front-end receivers is done via the control and timing unit. This provides a USB link to a Linux computer, which retrieves status information from the LNA and also controls the LNA. Since the vendor-provided software suite for configuring the digital down-converters unfortunately only runs under Windows 2000, a dedicated PC, connected through a second USB link, is provided for this purpose. Time-critical signals from the EISCAT radar controller are transmitted to the receiver equipment on dedicated fibres. The reference clock signal is distributed via a coaxial cable.



Figure 10.5. Equipment cabinet with the full complement of receiver hardware installed.

10.2.3 Interface block

Figure 10.6 shows the structure of the interface hardware. The data interface block receives the six serial data streams representing the 24 data channels and divides the streams into six groups of four serial data streams each, which are then converted back into a parallel format matching the interface into the EISCAT receiver.

The reference clock for the front-end receivers is a GPS-stabilised 10 MHz signal, taken from the frequency distribution amplifier in the site's instrument hall. It is sent on a coaxial cable to equipment cabinet # 2 in the array, and distributed from there to the other two cabinets. Figure 10.13 shows part of the Demonstrator system's interface hardware, in the receiver room of the Kiruna site.

The interface hardware also media-converts all electrical control signals generated by the EISCAT radar controller, like e.g. the calibration injection signal and the DDC synchronisation command, into optical signals before transmission to the array.

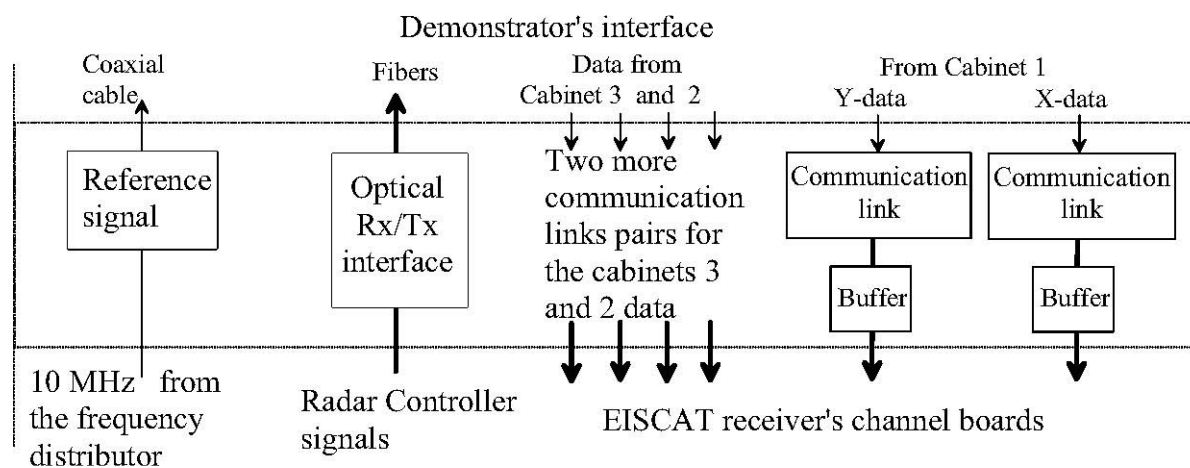


Figure 10.6. Interface hardware block diagram.

10.3 Receiver structure

Figure 10.7 shows a block diagram of the Demonstrator receiver and signal processing hardware.




Figure 10.7. Main functional blocks of one Demonstrator signal channel. The complete system comprises 96 such channels, two for each of the 48 dual-polarisation antennas. All equipment above and to the left of the dashed red lines is installed in one of the equipment cabinets out in the array; all equipment in the lower right-hand corner is installed in the control room of the Kiruna EISCAT site.

The system architecture shown here was selected as a compromise between two partially conflicting aims: on the one hand, to include and validate as many as possible of the target 3D receiver components and functions, and on the other hand to interface the Demonstrator to the existing EISCAT UHF receiver hardware and software, thereby making the array fully operational as quickly as possible:

- The (low-noise amplifier - 80 Msamples/s ADC) direct-digitising front-end system has been designed and constructed by a team from Luleå University of Technology (LTU). It can process a 28 MHz wide frequency slot with very good dynamic range and is functionally up to the full EISCAT_3D receiver specifications. If required, the electrical design could, to a large extent, be copied as designed and used in a much bigger array. In order to minimise cable losses, the front-end systems are located as close as possible to the antennas, i.e. in the equipment cabinets,
- The digital down-converter units (DDCs) are specific to the Demonstrator and will not appear in exactly the same context in the target receiver design. They are set up to low-pass filter the 80 Msamples/s front-end data streams and decimate them to ≈ 5 Msamples/s, which is a sample rate that the UHF receiver can cope with (just barely...). The filters are currently set up to pass a total bandwidth of 1 MHz per data stream, enough for ion line work,

- The decimated data from all rows is serialised by application-specific circuits (ASICs) and transferred to the interface block in the site control room, where it is de-serialised and media converted back to electrical signals,
- These are fed into the existing EISCAT UHF receiver channel boards, transferred from those to the VME single-board computer that controls the UHF receiver back-end and processed into normal lag-profiles by the standard lag_wrap software,
- All time-critical control signals are generated by the radar controller (RC) of the EISCAT receiver, and are distributed via the interface block to the front-end, the DDC, and the communication link. The DDC, the optical communications link and the control and interface hard- and software have been realised by the EISCAT project team.
- The interface block will soon be expanded by the addition of an FPGA-based beam-former system, programmed to combine signals from all rows into beam-formed data streams. When this is done, the FPGA logic will also replace the de-serialiser hardware currently used. The FPGA development has been done at the Swedish Institute of Space Physics (IRF). Detailed descriptions of each functional block are given in Sections 10.3.1- 10.3.4.

The receiver hardware is assembled in three sub-racks, which are set on top of each other. A fully configured sub-rack, ready for mounting in its equipment cabinet, is shown in Figure 10.8.




Figure 10.8. Photograph of a complete front-end / DDC / SerDes subsystem.

10.3.1 Front-end subsystem

In the front-end subsystem, the (224 ± 14) MHz signal from the antenna is first amplified in a low-noise amplifier (LNA). The amplified signal is de-aliased by passing it through a 28 MHz wide band-pass filter, further amplified in a second amplifier stage and finally sampled by the ADC and the samples are sent into the DDC. A block schematic of a front-end is shown in Figure 10.9.




Figure 10.9. Structure of the front-end subsystem.

The LNA is designed around a low-noise E-HEMT (Avago ATF-541M4) followed by a band-pass filter and a monolithic gain block (RFMD RF3376). To facilitate system calibration, the LNA includes a directional coupler and four analogue switches. This allows calibration of phase and amplitude variations between channels. Calibration signals are distributed over coaxial cable to all channels in the array simultaneously, but injected sequentially into one LNA at a time.

The ADC board contains additional filters and two gain stages and delivers a 16-bit (plus clock bits) output data stream to the DDC. The analogue input signal is amplified, filtered and digitised by a Linear Technology LTC2208 16-bit ADC running at 80 MHz. An on-board crystal oscillator, phase-locked to a 10 MHz external clock reference, generates a low-jitter sampling clock signal.

The last gain stage is capable of producing an output signal level that could damage the ADC. To minimise the risk of this happening, a protection circuit consisting of a Schottky diode limiter and a controllable attenuator is inserted ahead of the ADC.

10.3.2 Digital down-converter

The digital down-converters (Section 8.6) are used to band-limit and decimate the 80 Msamples/s front-end data stream. These functions are realised with the Intersil ISL5416, which is a wideband, four-channel, programmable digital down-converter chip. Each channel comprises a numerically controlled oscillator, a digital mixer, digital filters, automatic gain control, and a decimating filter. A block diagram of the ISL5416 is shown in Figure 10.10.




Figure 10.10: Block diagram of the Intersil ISL5416 DDC chip

All four channels are independently programmable and may be updated in real time. Each channel can take its input data from any one of the four digital input buses. Input data rates up to 95 Msamples/s are supported. The input data stream is mixed with a complex-valued local oscillator signal, generated by a 32-bit programmable numerically controlled oscillator (NCO), and the mixer output is then filtered. Filtering functions include a multi-stage cascaded-integrator-comb (CIC) filter, two programmable FIR filters (the first up to 32-taps, the second up to 64-taps), a half-band interpolation filter, and a re-sampling/decimating FIR filter. The overall decimation ratio is from 1 to 4096. The internal data path is 20 bits wide. Channels may be cascaded for increased bandwidth. Outputs from the part are available over the parallel, serial or microprocessor interfaces. Selectable outputs include I samples, Q samples, and automatic gain control (AGC) settings. Data from all four channels can be time-multiplexed onto any one of the available output ports; the Demonstrator DDCs are all programmed to operate in this mode.

10.3.3 Data communications link

A serialiser-deserialiser (SERDES) is a pair of functional blocks commonly used in high-speed communications. The two SERDES blocks, one at either end of a communication link, convert data between serial data and parallel formats. The basic building blocks of a SERDES are:

- a differential transmitter-receiver,
- a serialiser-deserialiser,
- an encoder-decoder (e.g. 8b/10b, 64b/66b line codes) and
- transmitter and receiver buffers (FIFOs).

In the Demonstrator system, the SERDES boards installed in the equipment cabinets convert the time-multiplexed 16-bit DDC output data streams into serial form and output them on optical links, running at 0.8 Gbit/s. These terminate in the interface block in the EISCAT control room, where a set of optical receivers convert the optical signals back to electrical signals, which are then re-parallelised by a battery of SERDES boards.

The communication links are realised with Texas Instruments TLK1501 evaluation module boards and industry-standard optical transceivers. The TLK1501 SERDES chip uses high-speed phase lock loop (PLL) technology, serialises, encodes (8b/10b), transmits data on a differential pair and supports an effective serial interface speed of (0.6...1.5) Gbps.

Figure 10.11 shows how X polarisation complex data from array rows 1-4 are transferred through the serial link. Eight 16-bit complex data points, representing the I and Q parts of the four data words, are packed into a 128-bit frame, consisting of eight 16-bit slots.

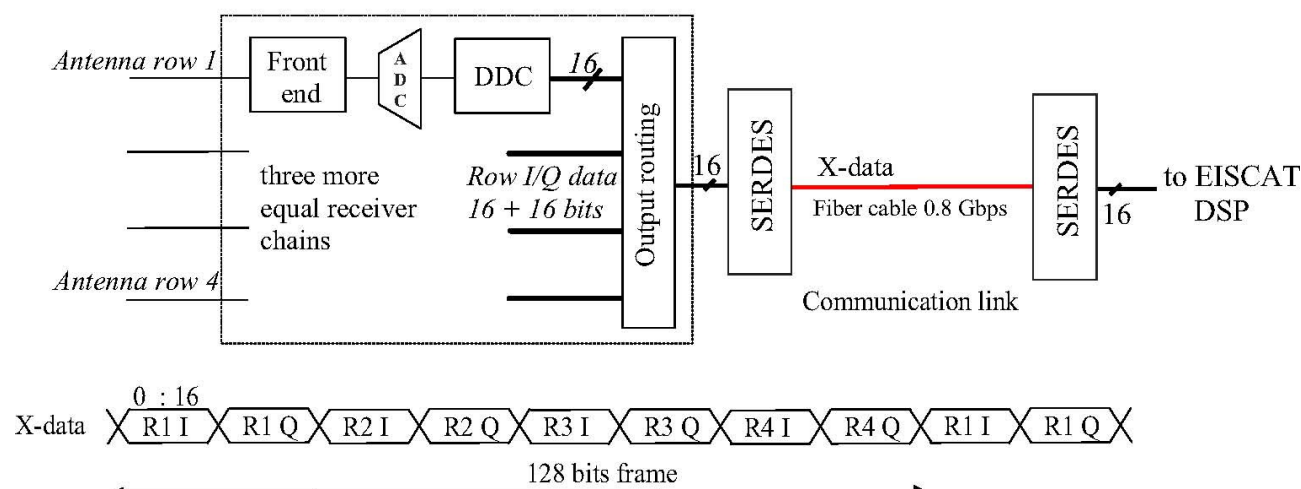


Figure 10.11. Serial fibre link between the Demonstrator receiver and the EISCAT UHF DSP.

10.3.4 Interface to the EISCAT UHF digital back-end

The re-parallelised data streams out of the receiving SERDES circuits are dumped into the EISCAT UHF receiver via a set of specially modified “channel boards” (Figure 10.12).

A standard EISCAT receiver *channel board* is essentially a single-channel digital receiver, where, in the standard EISCAT application, high-speed sample data are dumped into the channel board via the Front Panel Data Port (FPDP) (an industry-standard interface for high-speed parallel data transfer), converted to base-band, low-pass filtered, decimated, gated and finally stored in a double-bank buffer memory (BM), from which the data are then read out by the lag_wrap software correlator program. A special SERDES/FPDP data adapter board has been designed and constructed. It accepts the output from the SERDES system and adapts it both electrically and mechanically to the FPDP input. This makes the Demonstrator data stream available to the existing EISCAT receiver back-end, so allowing standard EISCAT software to be used to process the data from the array. The interface arrangement is shown in Figure 10.12.

Because the DDC boards in the Demonstrator receiver already provide the down-conversion, filtering, decimation and gating functions, most of the hardware on a standard channel board can be bypassed; the only operation required is to buffer the Demonstrator data in the double-bank memories. To accomplish this, a special piggyback module that bypasses the NCO and FIR filter has been designed. This module also generates proper timing signals for the dual-page buffer memory.




Figure 10.12. Demonstrator/EISCAT receiver data interface

The digital signal processing VME crate at the Kiruna EISCAT site can accommodate six channel boards. All those are needed when the data are to be taken from the whole Demonstrator array. Accordingly, six NCO/FIR bypass modules and six FPDP data buffer boards were fabricated. Figure 10.13 shows how all six channel-boards are connected with FPDP flat cable to the output connectors of a SERDES/FPDP buffer rack.




Figure 10.13. Demonstrator receiver-EISCAT receiver interface equipment. The rack for the six SERDES/FPDP buffers can be seen at the top, the EISCAT DSP rack is in the middle and in the lowest rack there is the optical RX/TX hardware for the radar controller.

10.4 Signal flow through the Demonstrator

The transformations suffered by a signal propagating through one of the 24 receiver paths in the Demonstrator system are shown in Figure 10.14.

The 224-MHz signal from the antenna is amplified in the low noise amplifier, band-pass filtered, and further amplified in a high performance amplifier, such that the average noise level at the output is in the order of two ADC bits; this is to suppress ADC quantisation noise at the expense of

a 12-dB reduction in dynamic range, which has been deemed to be acceptable. The signal is then digitised to 16-bit of resolution by an ADC running at 80 MHz sampling rate.

Next, the digital data stream is input to a DDC where it is converted to a complex-valued base-band stream, low-pass filtered and decimated. Finally, the decimated data stream is transferred into the EISCAT UHF receiver sample buffer for further signal processing, such as correlation computations.




Figure 10.14. Signal flow through a receiver chain.

10.5 Demonstrator-specific software

The engineering-level software for the Demonstrator system consists of modules to set up the front-end, the A/D, and the digital down-converter, for each of the 24 signal chains. When the FPGA-based “beam-former” hardware modules are installed, these will require the addition of a corresponding software module to initialise them and load configuration parameters.

10.5.1 Rack-software

Both the receiver front-end and the ADCs are accessed via the "rack" software module, which provides a comprehensive suite of control and monitoring functions addressing e.g. the protection diodes, the LNAs, the ADCs, the switching on and off the power supplies etc. Temperatures, system status and other engineering data can be read back.

10.5.2 DDC-software

In the Demonstrator setup, Intersil ISL5416 four channel digital down-converter (DDC) ASICs handle the digital frequency translation, filtering, and decimation. The DDC chip needs to be loaded with proper operational parameters like decimation rates, centre frequencies and filter parameters. To speed up the system construction, a number of off-the-shelf ISL5416 evaluation boards have been used together with a vendor-supplied utility program that only runs under Windows 2000. As the Demonstrator system evolves, the onboard microprocessor interface, with custom-made software, will be used; this will greatly simplify the integration of the Demonstrator into the EISCAT standard control-and-monitoring system, EROS.

10.5.3 Radar controller

The radar controller (RC) is a state machine with a time resolution of 100 ns, which initiates and controls all time-critical actions on the radar hardware. This device is programmed through a high-level language, TARLAN, where all actions are set to happen at certain time delays from the start. The source file is compiled, the resulting binary file is downloaded into the RC and starts to execute at the specified experiment start time. To handle the Demonstrator array, the TARLAN syntax and compiler have been expanded with a few Demonstrator-hardware-specific commands.

10.6 Installation period 2006–2008

The original EISCAT_3D design study plan included the construction of a proof-of-concept phased-array Demonstrator antenna, with sufficient receiving area to detect an ionospheric signal transmitted from the EISCAT VHF system in Tromsø. After our application for FP6 support for the Demonstrator had been declined, it was decided to nevertheless construct the planned system, which would now be financed entirely by EISCAT. In the spring of 2006, a 10 x 25 m piece of land near the Kiruna EISCAT site building was prepared to receive the Demonstrator array. The support structure was assembled in July. An order for 48 crossed-Yagi antennas, tuned to 224 MHz, was placed with a California-based company, M2 Antennas, Inc.. The antennas were received, assembled, tested and installed in the array in November 2006–January 2007. The required power combiners were designed and prototyped in-house, but the production was left to the industry.

Next, all power combiners and the associated cabling were installed and a set of coaxial delay lines, proportioned to generate a single beam at 55° elevation was fitted. Finally, the overall impedance characteristics of the phased array were verified by network analyser measurements and found to be very good.

Conclusive proof of correct operation was first obtained on October 17, 2007, when the array successfully recorded VHF incoherent scatter signals from 280 km altitude in the ionosphere above the Tromsø EISCAT VHF radar. As far as we have been able to determine, this is the first documented bistatic reception of the EISCAT VHF signal. While one of the two orthogonal polarisations was very noisy for some still undetermined reason, the noise temperature of the other polarisation channel was determined to be approximately 130 K — a very good value at 224 MHz. The SNR on this polarisation approached 4 % and the data could easily be fitted for all standard ionospheric parameters. All in all, the first reception of the VHF signal was a very important milestone in the overall project.

During 2008, a great deal of effort went into the manufacturing of the front-ends, the integration of the DDCs and the communication link and various other hardware-related tasks. In the 8th meeting

of the Steering Group on September 17, 2008, it was decided that the Demonstrator system should be operational at the beginning of November 2008. Around mid-September, all three equipment cabinets were still empty except for the basic infrastructure: mains, heating and ventilation system, USB hub, fibre optic cables to the site building (22 fibres), and coaxial cables to the site building and between the cabinets.

Careful coordination of the very intensive installation period was now necessary, because the hardware was being assembled at three separate locations: LTU, the Kiruna EISCAT site and the Sodankylä EISCAT site. Therefore, regular teleconference meetings were arranged to plan and coordinate the installation work. In the first meeting, 22 September 2008, a detailed timetable for the installation period was worked out. This indicated what was planned for each week, how many persons would be involved, and who would actually be working at the Kiruna site. The progress was monitored in weekly teleconferences, and the timetable was updated when needed.

In mid-October 2008, a test signal was sent through the whole signal path, from the LNA into the EISCAT receiver, for the first time - though yet only in a lab environment. The data were then processed with the standard EISCAT signal processing software and the results graphically displayed on a computer screen. This verified that the basic design concept worked.

The finalising of the receiver hardware, and designing test software, continued intensively from mid of October to the second full week of November. The hardware was bench-tested unit-by-unit in the lab. When the basic performance of all the 24 signal chains had been thoroughly verified, the hardware was finally installed in the three outdoor cabinets and the required signal and control cables were connected.

A simple demonstration measurement was now set up. Short pulses ($1\mu\text{s}$ @ 223.660 MHz frequency, 34.0dBm) were radiated from a test transmitter mast (see Figure 10.2). All rows in the Demonstrator array received these signals successfully on 13 November.

10.6.1 List of the main milestones:

Year 2006

- *May* — Ground preparation work on a 10 by 25 meter area at the Kiruna site.
- *June* — Support structure for the antenna elements erected. Twelve rows with 4 antennas in each row.

Year 2007

- *January* — 48 Yagi antennas assembled and installed in the array.
- *June* — RF combiners and RF cables installed.
- *August* — Three outdoor cabinets placed inside the array to house the front-end receiver system.
- *October* — First radar tests. A temporary test receiver was used together with analogue beam forming (phase-delay cables and an analogue combiner). A signal transmitted with the VHF transmitter in Tromsø and scattered from the ionospheric F layer at 280 km altitude was successfully received. The signal-to-noise ratio was about 4 %.

Year 2008

- *October* — Test signal (228.970 MHz) sent through one chain of the Demonstrator receiver into the EISCAT receiver in a lab environment.
- *November* — Installation of the Demonstrator receiver. Low noise amplifiers and ADCs, digital down-converters and SERDES modules installed in the cabinets. Signal from the test transmitter received successfully.

10.7 First Demonstrator incoherent scatter data

The Demonstrator antenna array is designed for a centre frequency of 224 MHz to make it possible to conduct a range of proof-of-concept tests and experiments where the EISCAT VHF transmitter in Tromsø is used as an illuminator.

The first step in the commissioning of the Demonstrator was to verify proper operation of the basic antenna array. Several tests were carried out for this purpose. To reduce the number of unknowns insofar as possible, the digital receiver system was not used. Instead, the signals from the twelve array rows were combined into a single signal per polarisation after first having been passed through passive delay lines made from low-loss coax cable and cut so as to create a single beam at 55° elevation, i.e. aligned with the boresight direction of the individual Yagi antennas.

At this elevation, the array beam intersects the vertically pointing VHF transmitter beam at an altitude of about 280 km above Tromsø. The two beam-formed signals (representing the X and Y polarisations) are amplified by low-noise GaAs FET amplifiers ($T_n \approx 30$ K), down-converted and finally introduced into the EISCAT UHF receiver at the first intermediate-frequency level, 112 MHz. With this arrangement, the tests can be run as standard EISCAT experiments under EROS control and using the standard EISCAT signal processing software.

The test experiment, a derivative of the tau-6 code, has been run on four occasions (Oct. 16, 2007, Oct. 17, 2007, Nov. 2, 2007 and Feb. 5, 2008) for a total of 10 hours. The transmission consists of a set of (16 x 84- μ s) alternating-code pulses for a total pulse length of 1344 μ s. The reception was done on two orthogonal linear polarisations (called X and Y), to be able to study the effects of Faraday rotation of the scattered signal.

In all tests, one of the two polarisation channels delivers a signal-to-noise ratio approaching 4%, while the signal power in the other polarisation channel is very low. This could possibly be partly due to a more or less constant Faraday rotation during the experiments, which were all carried out in similar ionospheric conditions. However, the difference in SNR between the two channels is greater than what can be explained by an unfavourable orientation of the polarisation ellipse. It has been noted that there exists a poorly understood gain difference between the polarisation channels; the missing signal might also be related to this problem, which is still being investigated.

The data from the “good” polarisation channel have been analysed with the GUISDAP incoherent-scatter data analysis software suite. Figure 10.15 shows electron density and electron and ion temperature values from the 2 November 2007 run. After post-integrating the raw data for 2 minutes, the error bars on the output parameters are comparable to those obtained in a typical UHF experiment under comparable ionospheric conditions. From the known distance between Kiruna and Tromsø and the actual time of the signal-carrying samples, the effective Demonstrator beam elevation was calculated to be 52.4°. This is slightly different from the geometrical elevation but demonstrates that the array essentially works as expected. The difference might be partially due to

some snow cover on the elements, as all experiments were carried out in wintertime. The side-lobes of the array are rather strong, as several echoes from hard targets were seen at very different heights, such as meteors from about 100 km and space debris from about 800 km.




Figure 10.15: Analysed data from the 2 November 2007 Demonstrator test experiment. From top to bottom: Electron density, electron and ion temperatures and ion drift velocity along scattering bisector. 2 minute averaging of the input (auto-correlation function) data has been employed.

10.8 Reception of test signals

A simple test has been set up to check the whole antenna / digital receiver chain with a known signal shape. Short ($8\ \mu\text{s}$) pulses are transmitted through a mast-mounted transmitter antenna located in front of the Demonstrator array and picked up by the individual rows. Each transmit/receive cycle is $204.8\ \mu\text{s}$ long and subdivided into 256 samples of $0.8\ \mu\text{s}$ each. The $8\text{-}\mu\text{s}$ pulse is transmitted $50\ \mu\text{s}$ into the cycle.

Figure 10.16 shows the pulses received on four different array rows. The red and blue colours show the real and the imaginary part of the complex signal amplitude respectively. The test shows that the system works as expected. No attempts to phase the signals from the individual rows have been made yet; such tests are planned as soon as the FPGA-based beam-former has been completed and installed.




Figure 10.16: $0.8\text{-}\mu\text{s}$ pulses received on four different array rows and processed through the complete signal chain.

Antennas operating in an arctic environment may have their properties degraded significantly due to snow falling on the antennas. This could cause severe problems in applications where knowledge about the gain and pointing direction of the main beam of the antenna array is needed in absolute numbers. In Section 8.1.4 the performance of the antenna elements during snowfall is assessed both individually and in an array configuration. Some of the material presented in this section have previously been published in [Lin08]. In Section 10.9.2 a measurement system for the EISCAT_3D antenna array is proposed, this section is based on material from [Lin09].

10.9 Environmental effects on the Demonstrator Array

In this Section, the behaviour of an individual Demonstrator element antenna when influenced by snow is studied in some detail and a near-field measurement system for monitoring the phase behaviour of the target arrays is proposed.

The 224-MHz element antenna selected for the Demonstrator is a (6 + 6) element X Yagi antenna, designed and manufactured by M2, Inc. of Fresno, California (USA). It is a derivative of an existing (7 + 7) element X Yagi tuned to 222 MHz and was selected mainly because it was inexpensive, available more or less off-the-shelf and could be delivered on very short notice.

However, this antenna is constructed with relatively thin (diameter 6 mm) elements and as a consequence it is relatively narrowband. It was therefore expected to be relatively sensitive to snow and ice coverage, prompting the present study. Here, only one of the polarisations has been studied.

10.9.1 Measurements

The antenna return loss was measured on January 18-21, 2008, at the LTU facilities in Luleå, Sweden. During this period, a considerable amount of snow fell and the temperature oscillated around 0° C, resulting in a very heavy and wet snow/rain mix. In Figure 10.17, the return loss of the antenna at 224 MHz is shown as a function of time (top right). It is also shown as a function of frequency (left) for two selected occasions marked in the plots to the right. In addition, the centre frequency and bandwidth of the antenna as well as the measured precipitation at a meteorological station in Luleå (data was provided by the Swedish Meteorological and Hydrological Institute, SMHI) are shown.

Before the snowfall started, the antenna bandwidth was ~ 6 MHz, and the passband was centred on 224 MHz. As the snowfall began and continued, the whole antenna passband was shifted steadily downward in frequency. As this process continued, the reflection coefficient increased until the passband was entirely below the desired frequency, at which point the antenna essentially became non-operational. This behaviour can be seen in the plots of the centre frequency and bandwidth.

Measurements performed on the test-array outside Kiruna, Sweden, showed a mutual coupling of about -40 dB between adjacent antenna elements when the antennas were free from snow. The mutual coupling has also been simulated for a four-by-four antenna array with the same dimensions as the test array. The simulations were done without snow on the antennas, with 0.25 mm of snow, and 0.5 mm of snow. The snow was assumed to be distributed uniformly over the whole array. In all cases the coupling coefficient of less than -35 dB. Thus, the risk of scan blindness does not seem to increase due to snowfall. It should however be noted that these results are valid for the antenna elements used in the test array only. The antenna elements used in the full-size array will be designed for larger bandwidth and thus different. The coupling effects under different conditions (i.e. snowfall) should therefore be considered during the design process of both the antenna elements and the antenna array.




Figure 10.17: The measured return loss (s_{11}), bandwidth (BW), centre frequency (f_c), and precipitation during snowfall. To the left the return loss is shown as a function of frequency for two selected occasions. These are marked in the plots to the right with dashed and dash-dotted lines, respectively.

10.9.2 Proposed measurement system

Since the EISCAT_3D radar will operate continuously, the far-field of the antenna elements will need to be measured and verified regularly. In particular, the measurement system must be able to detect any changes in the phase characteristics. In this section a measurement system using the equivalent electric current method for the antenna array is proposed. To reduce the size of the matrices in the system of equations to be solved, the near-field of each antenna element is measured separately, with the probes used as test transmitters (the term probe will still be used throughout the text for simplicity). This is possible since all elements will have individual and independent receiver front-ends. The presented method has been simulated in order to illustrate the performance of the approach.

The measurement system is described in detail in [Lin09]. It can be summarised as:

1. The field at a given point in space (e.g. at the location of a probe) depends on the current distribution on the AUT and can be calculated for the general case using the near-field integral [Kil00]. The current distribution is here described using a sum of trigonometric functions.

2. To enable the use of both near-field and far-field sources, the difference between the current distribution on two different antennas is used.
3. A linear system of equations is obtained which is solved using either the least-squares method or a Kalman filter [Kay93].

The performance of an implementation of the method described above is assessed using numerical data from simulations using the Numerical Electromagnetics Code. In the simulations, all effects due to mutual coupling between nearby antenna elements (e.g. scan blindness) have been neglected. These effects can be taken into account when performing the measurements by using additional basis functions and careful selection of the probes used at a given time.

The antenna array considered here is similar to the EISCAT_3D Demonstrator array and the antenna elements are assumed to be M2 Demonstrator X Yagis. The measurement system proposed here is designed with the final radar system in mind, where the size of the antenna array could be on the order of 125-by-125 m. The antenna elements are mounted at an elevation angle of 55° and also rotated 45° with respect to the local horizontal. Again, it should be noted that this system setup is only considered to define realistic locations of the probes. It is thus not designed with any array parameters in mind.

The probes are assumed to be mounted on 150 m high towers, which are spread out around the antenna array. Although there may be some interference due to reflections from the towers this effect should be small since there are no towers in the main beam direction of the antenna array. There are three probes on each tower, at 50, 100, and 150 m height. This setup is shown in Figure 10.18. Note that in this figure, there are no probes in the main beam of the array, which was one of the main requirements of the measurement system.




Figure 10.18: The proposed measurement system setup. The number of antenna elements plotted is in the figure for clarity lower than in the actual system. The probes are denoted by circles.

In the simulations considered here, the signals transmitted by the probes are assumed to be noisy. "Noise" here refers to all possible sources of error in the observations. It is assumed that the noise is Gaussian distributed with zero mean. Since the amplitude of the signal received by the antenna elements will depend on both the distance to the probes and the antenna gain in the direction of the probes the signal-to-noise ratio (SNR) will vary significantly between different probes. The SNR for the AUT considered in this section varies from less than -7 dB to over 32 dB for different probes.

To better evaluate the performance, the complex error in the estimated far-field was calculated. This can be seen in Figure 10.19. A total of 100 simulations were used to create this plot. It can be seen that the Kalman filter gives more accurate results than the least-squares solution. Also, the systematic error is small for both the least squares solution and the Kalman filter solution.




Figure 10.19: The complex error induced due to the noise on the test signal (ϵ denotes the error). The units are V/m.

This page intentionally left blank.

PART D

DESIGN BLUEPRINTS FOR KEY SYSTEM ELEMENTS



11 Design Blueprints for Key System Elements

In the following pages, we reproduce a number of the engineering drawings used in the design of the demonstrator array. As we discuss in Parts E and F (Sections 12 and 13) of this document, it is not yet appropriate for us to attempt to draw the detailed system designs for the final EISCAT_3D system – however the following diagrams represent the means by which many of the critical system elements needed for EISCAT_3D can be implemented using present-day technology.

The designs included here represent outputs from Work Package 4 (receiver design), Work Package 9 (signal processing) and Work Package 12 (time and frequency) and represent the efforts of design team members in the EISCAT Scientific Association, IRF Kiruna and the University of Lulea. The actual construction and testing of the demonstrator array was led by EISCAT, and financed using money from EISCAT's own hardware development resources. A very comprehensive description of the demonstrator array can be found in the previous section (Section 10), and further diagrams illustrating various aspects of the design can be found in Sections 8 to 10.

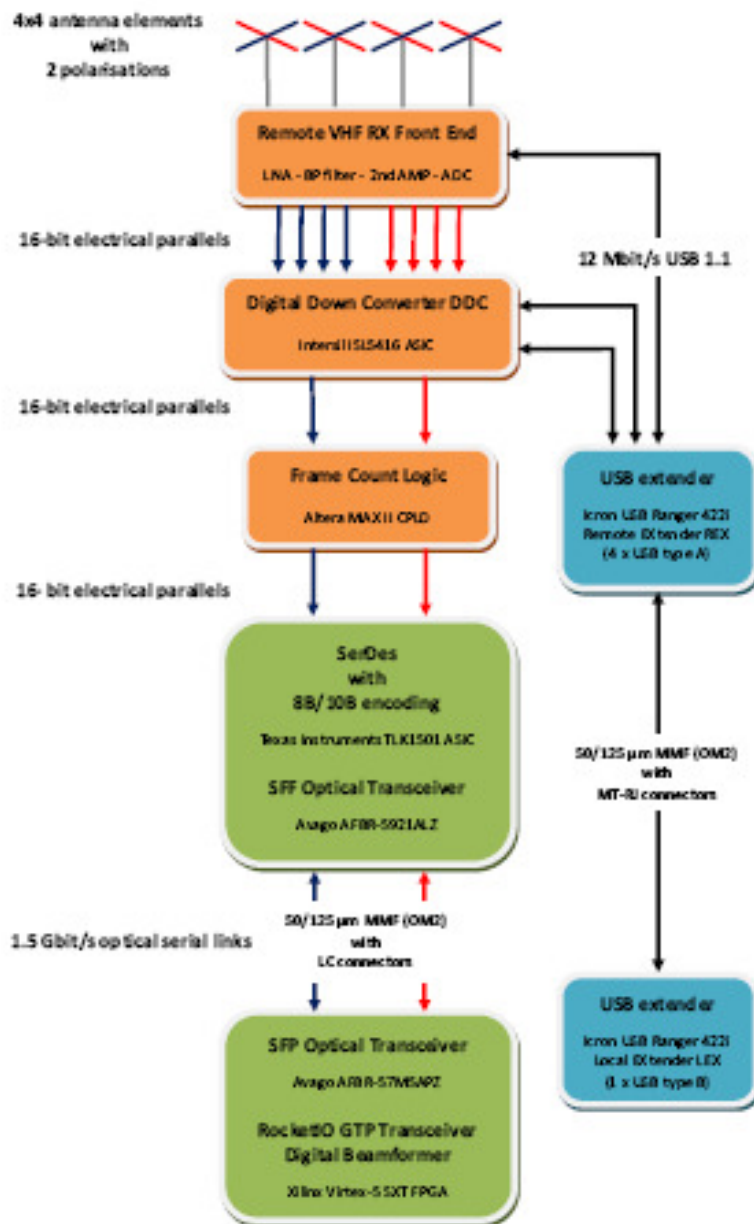


Figure 11.2: Networking, reference time and frequency system.

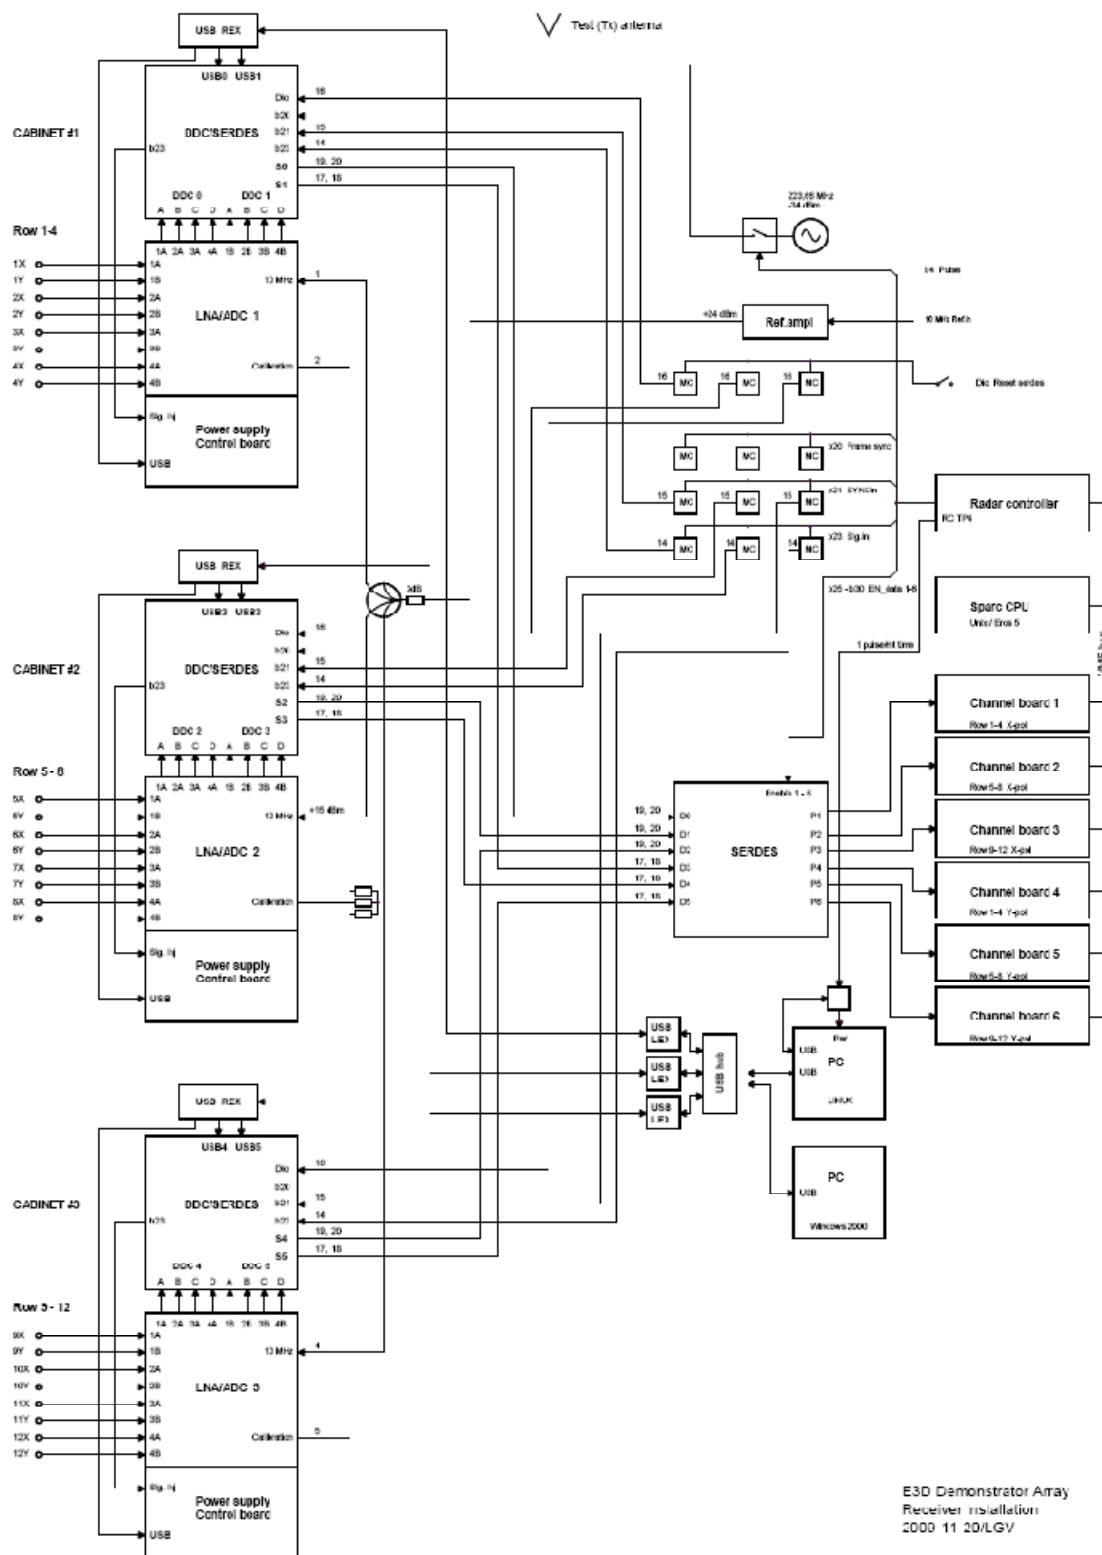


Figure 11.3: Demonstrator Array Receiver Installation.

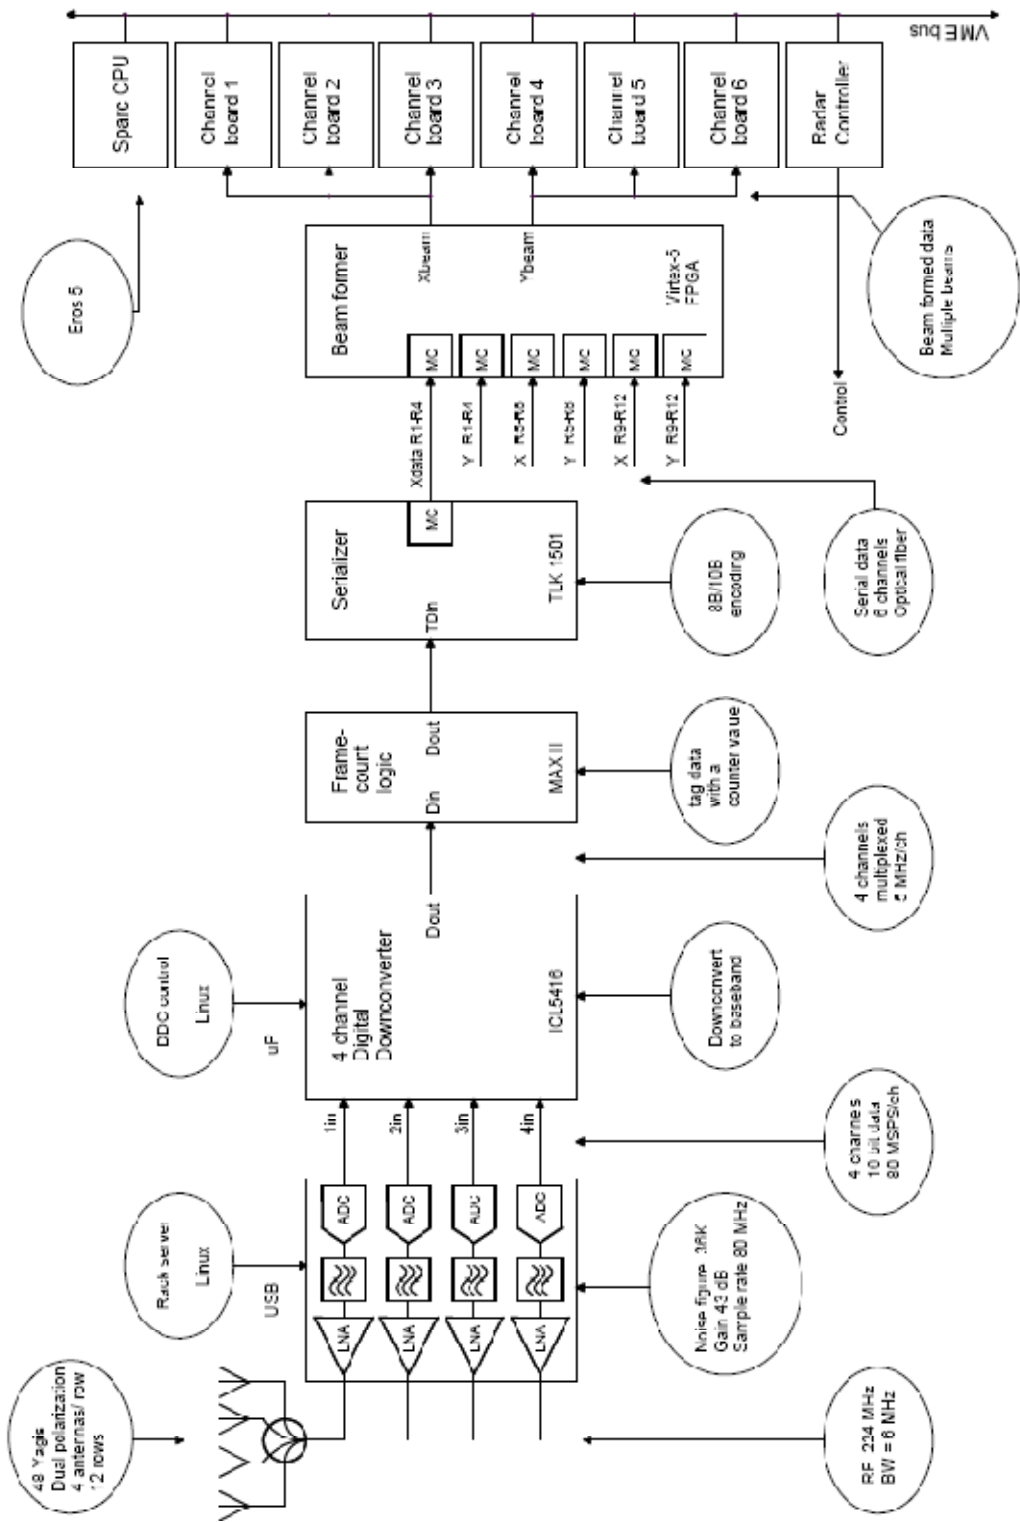


Figure 11.4: Receiver block diagram for the EISCAT_3D Demonstrator Array.

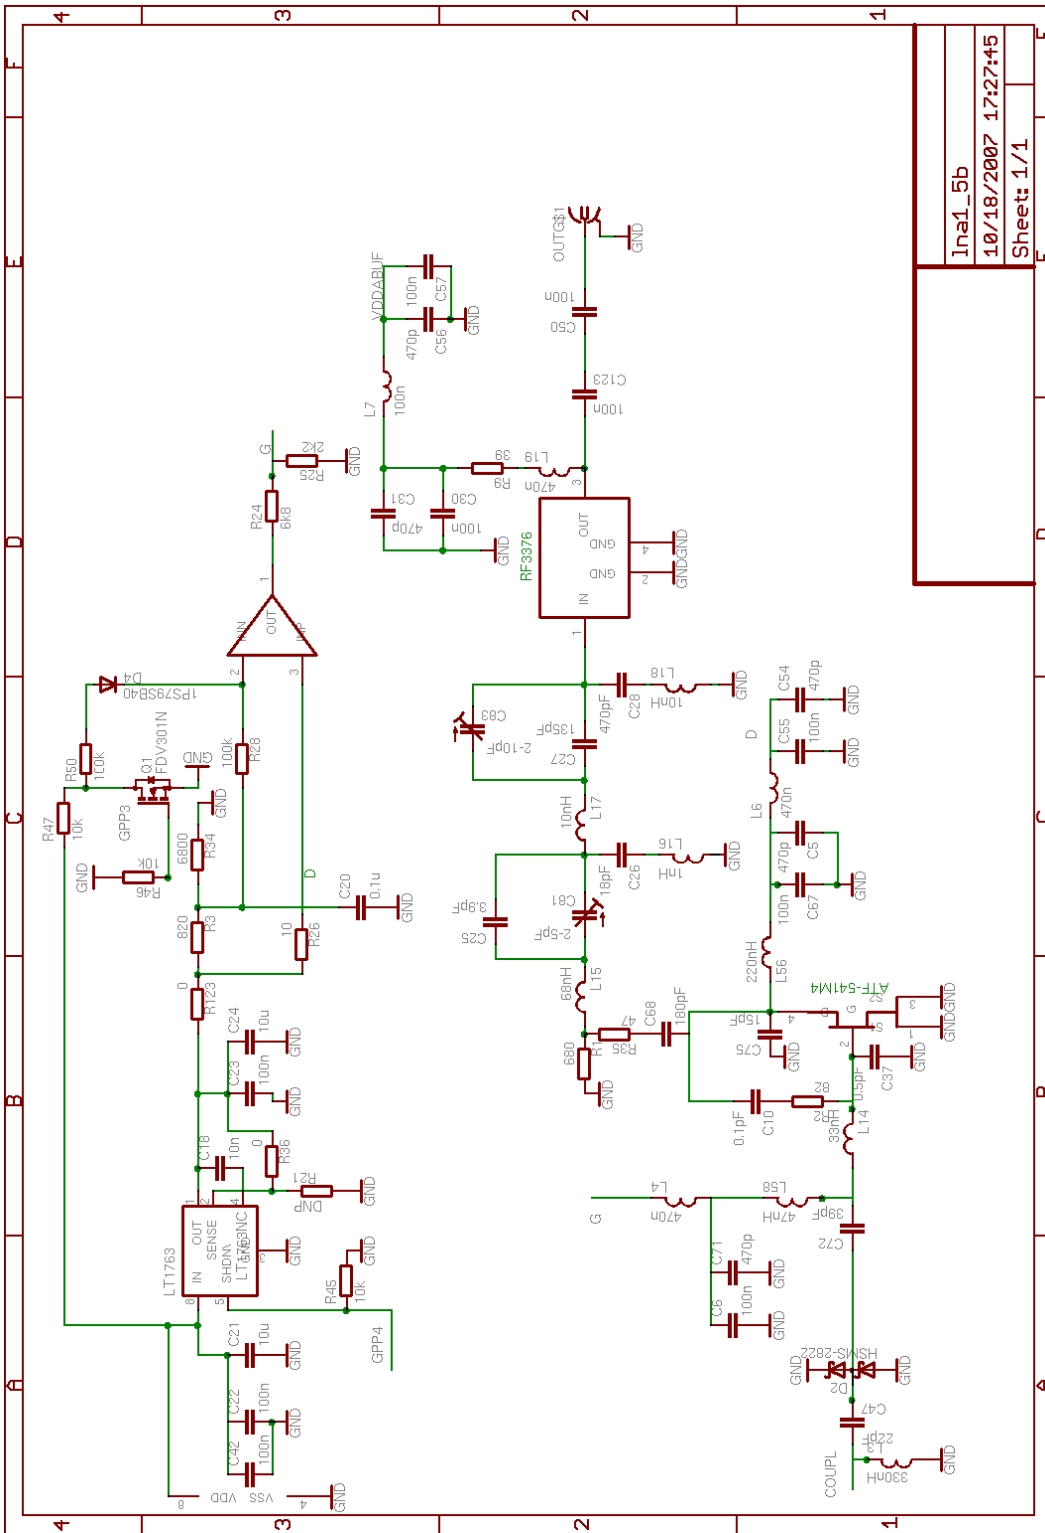


Figure 11.5: Low-Noise Amplifier design for the EISCAT_3D Demonstrator Array.




Figure 11.7: Framecount Logic System for the EISCAT_3D Demonstrator Array.




Figure 11.8: Block schematic of FPGA evaluation board.




Figure 11.9: Block diagram of the signal path inside the beam-forming FPGA.

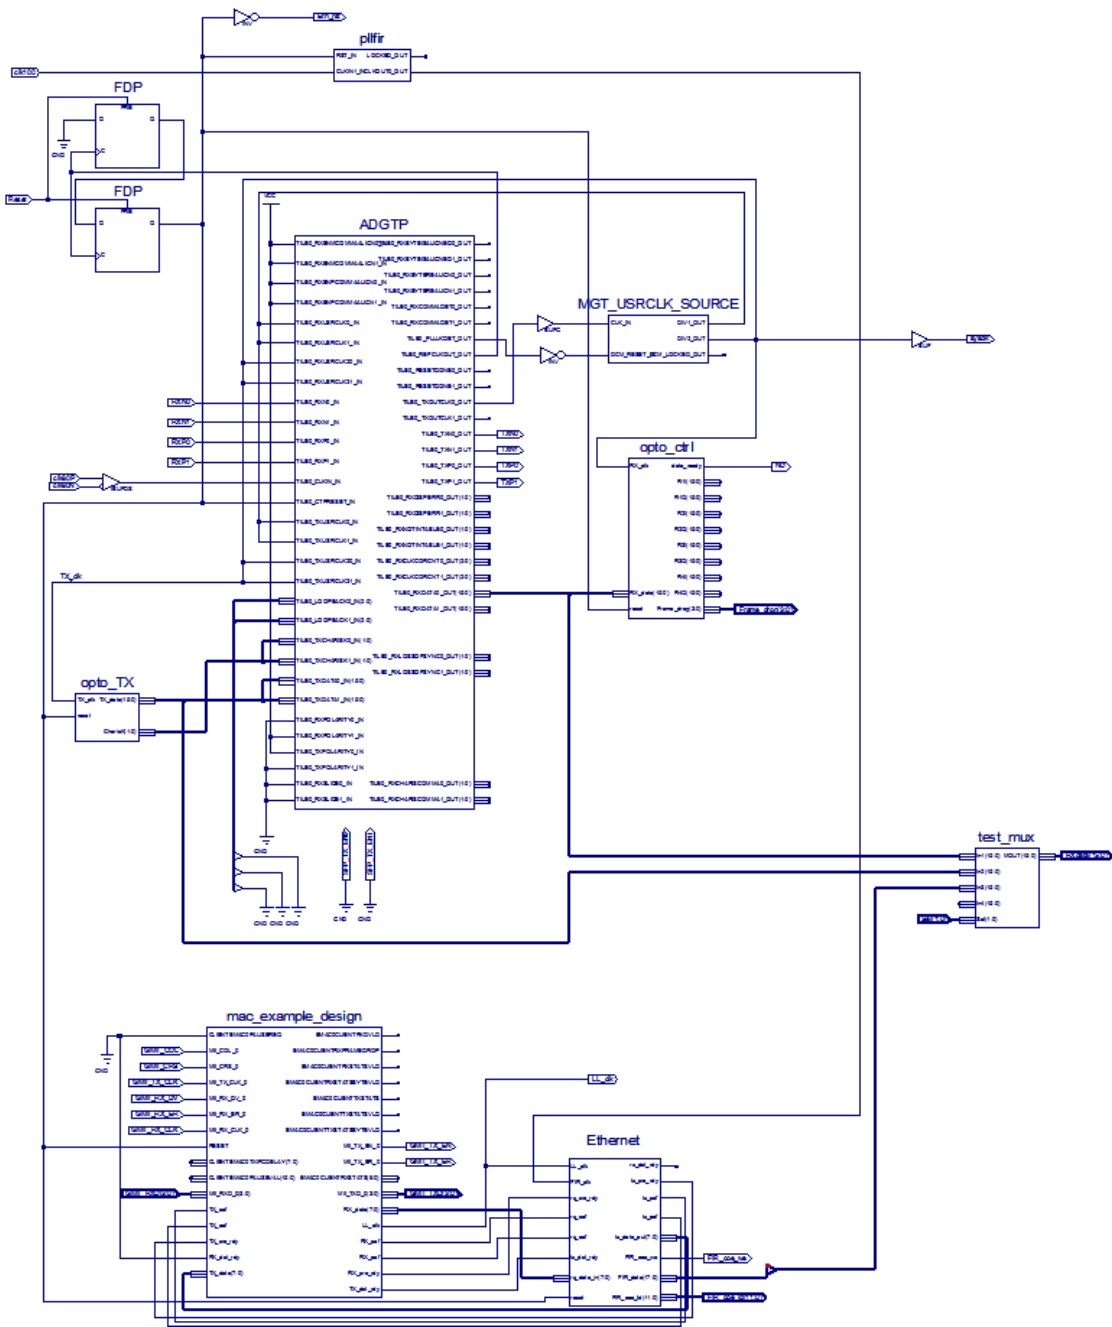


Figure 11.10: Top level schematic of the FPGA-based beamformer.




Figure 11.11: FIR filters and adders for each optolink.

PART E

IDENTIFICATION OF POTENTIAL SUPPLIERS



12 List of Potential Suppliers

When this study began in mid-2005, the expectation was that it would be finished with something resembling a complete list of the parts and suppliers needed for the construction of EISCAT_3D, enabling the organisation to move straight into a procurement exercise. This has proven to be unrealistic, for a number of reasons.

Firstly, EISCAT_3D will be a huge system, comprising at least 48,000 antenna elements, 96,000 receivers and 32,000 transmitter modules, all capable of performing to a very high level specification. It was obvious from the outset that this volume of components could not simply be ordered “off the shelf”, but that EISCAT would need to identify industrial partners (e.g. from the electronics or telecommunications industries) with the capacity to mass-produce these components. Furthermore, although we have constructed the demonstrator array using a number of commercially-sourced components (such as the antennas from M2 Inc. of California) it became apparent that some of these components were simply not adequate to provide the basis of the new radar, e.g. because of the de-tuning of the antennas when they were affected by frost and snow (see Section 8.1.3). This meant that it has been necessary to go back to basics, designing and fabricating our own antennas, such as the Renkwitz Yagi described in Sections 6 and 8.1. This design is very application-specific, and we have not yet reached the stage of identifying potential suppliers capable of fabricating such units on industrial scales, nor is it reasonable for us to do so until we are further into the implementation phase of the project (see Section 13).

The same is true of a number of other design study areas. For example, the low-noise amplifiers and ADCs were fabricated to our own design (though using lower-level commercial components), the beam-former was developed by creating filters designed specifically for EISCAT_3D, and the electronics needed to provide the basis of the transmitter system have only just become available because of the very high timing precision that is needed, leaving us very little time for system testing. This illustrates that in many cases, commercial solutions for the EISCAT_3D hardware do not exist. This does not mean that they cannot be constructed, but many of the sub-systems we require would need to be specifically manufactured for use in EISCAT_3D.

The case of the transmitter hardware illustrates another important point – the development of electronics in some of these areas is currently moving so fast that it would be extremely unwise to commit to a particular manufacturer at this time. What we have proven in the design study is that the concepts behind the EISCAT_3D system are achievable with present-day technology. While we have constructed a demonstrator array which verifies that an EISCAT_3D radar can work in practice, the system we have built cannot be scaled up to 48,000 elements without an unreasonable expense, which can only be mitigated by finding manufacturing partners able to produce the requisite purpose-built sub-systems in bulk. The speed of technology development in this area makes it desirable to leave the identification of such partners until as late as possible in the implementation phase.

There was, however, one area of the design study in which specific vendors were deliberately approached to provide quotes, or at least suggest hardware solutions, against the design study specifications. This was in the design of the data system (Section 8.9) where a key objective was to establish whether the data system design could be realised with current hardware, and obtain some idea of how much such a system might cost. A total of 23 possible vendors were identified and approached, and some provided highly detailed breakdowns of costs and suggested hardware – the details are reported in the relevant project deliverable [D8.3, part E].

In this report, the following suppliers have been mentioned as having provided hardware which has been used in the course of the design study:

- Analog Devices (receiver and transmitter systems)
- Avago (receiver system)
- Create (equipment used in the site surveys)
- Crystek (receiver system)
- Intersil (down-converters for demonstrator)
- M2 Inc. (antennas for the demonstrator array)
- Mitsubishi (transmitter system)
- NordNav (receiver used for testing array timings)
- Philips/NXP (transmitter system)
- Rohde and Schwarz (transmitter system and equipment used in the site surveys)
- Texas Instruments (internal data communication)

In addition, the following suppliers were identified and approached as being potential vendors for the data archive and distribution system:

- Bridgehead Software – www.bridgeheadsoftware.com
- Brocade Communications – www.brocade.com
- Copan Systems – www.copansys.com
- Data Direct Networks – www.datadirectnet.com
- EVault – www.evault.com
- Fujitsu-Siemens – www.fujitsu-siemens.com
- Headway Technology Group – www.headway.co.uk
- Hewlett-Packard – www.hp.com
- Hitachi Data Systems – www.hds.com
- IBM – www.ibm.com
- Isilon – www.isilon.com
- Latitude UK – www.latitudeuk.com
- LeftHand Networks – www.lefthandnetworks.com
- Logicalis – www.uk.logicalis.com
- LSI – www.lsi.com
- NetApp – www.netapp.com
- OCF – www.ocf.co.uk
- Pillar Data Systems – www.pillardata.com
- Quantum – www.quantum.com
- Silicon Graphics Ltd – www.sgi.com
- Solid State Solutions Ltd – www.s3.co.uk
- Union – www.unionsolutions.co.uk
- 3PAR – www.3par.com

As stated above, this should not be taken to imply that any of the companies mentioned here would necessarily have any advantage by the time that funding for EISCAT_3D is in place and the project is ready to issue requests for quotations. Although we are now in a position to say that the project is achievable with current technologies, the choice of actual suppliers will depend on a range of factors including price, ability to mass-produce and willingness to manufacture very application-specific hardware. The results of such an assessment simply cannot be predicted this far in advance.

This page is intentionally blank.

PART F

COSTS AND TIMESCALES
FOR
SYSTEM CONSTRUCTION



13 Costs and Timescales for System Construction

In this section, we briefly summarise our current view of the costs and timescales for the future evolution of the EISCAT-3D project. We discuss the extent to which these have been clarified by the design study exercise, and identify those areas where significant uncertainties remain. We also describe our views of the next steps in taking the project forward, exploiting both future EU funding opportunities and the national infrastructure programmes of interested member countries.

13.1 Estimated costs of EISCAT_3D

Our current cost estimates for the EISCAT_3D system have been made by bringing together information from a variety of sources.

In the first place, we have been fortunate to be able to draw on nearly 30 years' worth of collected experience from the EISCAT Scientific Association. In the 1970s the Association planned and constructed the old mainland VHF and UHF systems; in the 1990s the exercise was repeated with the design and construction of the EISCAT Svalbard Radar (ESR). The Association's records thus constitute a good reference, particularly as regards the fraction of the total construction costs going into civil works and supporting infrastructure.

Secondly, in the course of the present Design Study we have built a Demonstrator array, which provides a working illustration of the key technologies needed for the implementation of the new radar. The construction of the Demonstrator has involved the procurement, evaluation and further development of antennas, beam-formers, signal processing hardware etc., thus giving us valuable insights into the electronics marketplace.

Thirdly, in some areas of the study, such as the design of the data system, we have entered into detailed discussions with commercial vendors to establish whether our performance specifications were achievable with present-day technology and to get an idea of the financial implications.

Arriving at the final cost of EISCAT_3D is not, however, simply a matter of adding all these component costs together. As stated in Sections 8.1, 8.2 and 8.8.3, the target central site configuration comprises approximately 16,000 dual-polarisation antenna elements, together with additional "outrigger" arrays for interferometry. To this should be added up to four remote receiver sites, each comprising up to 8,000 elements each. A fully instrumented EISCAT_3D system will thus comprise at least 35,000 element antennas, 32,000 transmitter modules and 70,000 receivers.

The following tables show estimates made in 2007 of the costs of such a system, separating the costs of the active central site from the passive remote site. Since the EISCAT Scientific Association is legally a Swedish corporation using the Swedish crown (SEK) as its primary currency, the estimates were original done in SEK. They are here converted to Euros for the sake of convenience, assuming an exchange rate of 1 EUR = 10.6 SEK.

On the basis of these figures, an EISCAT_3D system consisting of one active site and four passive remotes would cost around €280M. Since this calculation was carried out, the value of the Swedish Crown against the Euro has declined by around 30%. Fortunately, however, the price estimates for most of the above sub-systems were based on US components, meaning that the most important exchange rate is that of the Euro against the US dollar, which is similar now (April 2009) to the equivalent exchange rate in 2007, so that this estimate of the total cost in Euros is still approximately valid.

Functional Unit	No. of Units	Cost in M€
Antenna systems	16000	15
Receivers/ADCs	32000	25
Transmitter System	32000	40
Timing System		4
Climatised Packaging	330	8
Internal data comm.		3
Beamformers	3	1
Data store		1
External data comm.		0.5
Building and infrastructure		0.5
Swedish VAT at 25%		25
Subtotal Active Array		123

Table 11.1: Possible construction costs for the EISCAT_3D active site

Functional Unit	No. of Units	Cost in M€
Antenna systems	8000	9
Receivers/ADCs	16000	12
Timing System		5
Climatised Packaging	170	3
Internal data comm.		1.5
Beamformers	3	0.5
Data store		0.5
External data comm.		0.5
Building and infrastructure		0.5
Swedish VAT at 25%		8
Subtotal Active Array		40.5

Table 11.2: Possible construction costs for one EISCAT_3D remote site

It is now necessary to introduce a number of caveats. The above estimates were largely made on the basis of the known costs of the various subsystems built for the Demonstrator and take only a nominal account of the savings that can be expected to result from industrial mass-production of receivers, transmitter modules and other electronics subsystems. We believe that it should be possible to obtain a factor 2 cost saving on these, equating to almost a factor two reduction in the total costs of the EISCAT_3D system, by contracting out the bulk production to an industrial partner with previous documented experience from this kind of work and the facilities required for fully automated production (such as e.g. any of the big European telecom corporations).

Taking account of possible savings achievable through mass production, the cost of constructing a 16000-element active array is therefore estimated to be of order €60M, with each passive 8000-element array costing 65 % less. In the present design study, we have envisaged a target system with one active and four passive sites, giving a total cost of €140M. This is essentially the minimum system required to achieve our initial performance goals. One must however be very cautious here: there is a large uncertainty in the mass production rates of the various system elements, especially since no actual tenders have yet been issued, or vendor quotes obtained for most of the relevant components.

At this point, it is important to recognise that EISCAT_3D is designed to be modular in construction, making it possible to build a smaller system originally and expand it as additional funding becomes available. It is shown in Project Deliverable 3.2 that already a 5000-element active site would provide raw performance equal to the best that the present EISCAT VHF system has ever delivered, but with dramatically increased flexibility and low-altitude capabilities. The construction of an array of this size would therefore be a logical first phase in a staged development of the project. Extrapolating from the above estimates, the cost would be

in the order of € 25 – 30 M. The receiving sites could also be initially constructed with half-size arrays, bringing the first phase cost to something in the order of € 70M.

Obviously it is not sufficient merely to construct the EISCAT_3D system – resources will also be needed for operations. Here it is much easier to extrapolate costs from the operations of the existing EISCAT radars, with operating costs estimated to be around €5M per year for a basic EISCAT_3D system, as described in this report.

13.2 Timescales for EISCAT_3D

The original EISCAT_3D FP6 proposal foresaw three stages of development following on from the initial design study; namely a preparatory phase (2009-2011), a construction phase (2011-2015) and an operational phase (2015-2045). In fact, we are currently somewhat ahead of schedule by virtue of the fact that the EISCAT_3D project is now included on the ESFRI roadmap. This is seen as being a critical step in the building of the political consensus needed to realise EISCAT_3D. We also have a working 4x12 element Demonstrator array, whose development was not formally part of the design study, but was constructed by EISCAT using its own resources. Added to this, we now have promises of long-term allocation of a 6.7 MHz wide slot in a highly suitable portion of the VHF spectrum in Norway, and we are in active discussion to secure protection for the same band in Sweden and Finland. All of these developments put us in an excellent position to begin the next phase of the project.

The preparatory phase, which is beginning now, has a number of critical objectives:

- a) Securing European political support for the EISCAT_3D facility
- b) Obtaining strong government and research council backing in the three Nordic countries, where the facility will be constructed.
- c) Obtaining support from the relevant authorities of the non-European members and associates of EISCAT (China, Japan and Russia) and other international partners.
- d) Obtaining the remaining frequency allocations required for active operation of the EISCAT_3D radars in northern Scandinavia.
- e) Selecting suitable sites and obtaining permission to deploy arrays in those locations
- f) Establishing a formal project structure
- g) Recruiting additional staff to carry out the increasing amount of technical work
- h) Clarifying those remaining technical issues that have not been resolved by the design study.
- i) Specifying the elements of the system hardware which require mass production,
- j) Identifying industrial partners capable of fabricating such systems
- k) Identifying suppliers for the remaining elements of the EISCAT_3D system
- l) Prototyping hardware elements,
- m) Drafting and issuing RFQs for the various system elements
- n) Obtaining firm tenders for same
- o) Securing sufficient funding guarantees to enable construction of an initial system

The preparatory phase will necessarily involve a strong element of salesmanship. We are confident that EISCAT_3D provides a superb opportunity for the European and international scientific community; not only to those groups involved in atmospheric and space physics, but also for institutions which are active in developing the relevant technological aspects such as ASIC design and advanced signal processing. Engaging such groups is absolutely necessary in order to provide the strong and competent engineering community needed to support and develop the project; without the support of such a community, EISCAT_3D will not survive in the long term.

In the construction phase, the goals will be as follows

- a) Conversion of the existing EISCAT VHF system to operate at the EISCAT_3D frequency (233 MHz)
- b) Construction of one or more passive remote sites
- c) Construction of an active site, initially at the minimum level for scientific operations
- d) Expansion of the active site, and addition of further remote sites, as funds become available
- e) Addition of further active sites (possibly by conversion of some passive remotes)

It is clear from the above that the construction phase will overlap with the operational phase of the system, since it would be possible, indeed even advisable to begin operations with a small system whose size can be increased incrementally as more funding becomes available. This type of staged development is actually a good guarantee of viability, since it ensures that development and testing are done on the basic system, minimising the technical risk of adding more modules.

It should also be mentioned that the timeframe for the construction of EISCAT_3D is likely to coincide with a number of allied initiatives, with which we expect certain synergies. The Nordic governments are committed to the development of the “Northern Scandinavian window to Geospace”, the large collection of experiment platform infrastructure already existing in northern Norway, Sweden and Finland. This includes one of the world's best networks of ground-based observatories, combining optical and radio methods, two rocket launch facilities (at Andøya and ESRANGE) and the world's most powerful HF heater run by EISCAT itself. Some of these facilities have ambitious plans for further development, and the potential for high-latitude geophysics will be further augmented by the SIOS (Svalbard Integrated Observing System) programme, which promises significant development in the observational capabilities based on Svalbard, where EISCAT already has a strong presence. The EISCAT_3D development will add additional value to these related initiatives, making this period a very suitable time to be developing Europe's next-generation incoherent scatter radar.

The present EISCAT systems, which began operations in 1981, have played a major role in unifying European scientists in this area and giving them a common focus. Hence we would be looking for member nations to provide long-term guarantees of commitment to the facility, probably on a model similar to the existing five-year rolling membership. Experience with the existing systems has demonstrated that this commitment brings ample rewards in terms of scientific return – in terms of cost per publication, EISCAT must rank as one of the most cost-effective of Europe's major science facilities. In designing EISCAT_3D we have attempted to specify a system which will be capable of meeting the requirements of Europe's upper atmospheric and space physics community for an operational lifetime of at least the next 30 years and guaranteeing Europe's continuing competitiveness in this vital field of research

This page is intentionally blank.

PART G

SUMMARY AND CONCLUSIONS



14 Summary and Conclusions

In this report, we have summarised the results from the four-year “EISCAT_3D Design Study”, which has been supported by the European Union under its Sixth Framework Programme, and carried out during the period from May 2005 to April 2009. This report has constituted a high-level summary of the outputs from the thirteen work packages which made up the study, identifying how the constituent activities have related to each other and to the final system design, and briefly describing the main conclusions of each. During the course of the study, we have developed the technical specification of a third-generation incoherent scatter radar for the European research community, exploiting the latest technology to provide an order of magnitude improvement in performance and sensitivity compared to the current EISCAT radars.

The EISCAT_3D concept, as set out in this document, is based on a large central site, consisting of a 120m diameter filled circular core array with around 16000 antenna elements, supplemented by a number of small outlier receive-only arrays for interferometric imaging. The design also includes four remote sites, each consisting of around 8000 antenna elements, two on a north-south baseline and two on an east-west baseline, with the baselines intersecting at the central site. The inner remote sites will be located around 110 km away, and the outer remotes will be around 250 km away from the central site. This design offers significant improvements in sensitivity and functionality compared to conventional radars used for incoherent scatter up to now, including the ability to perform interferometric imaging as a standard part of system operations and allowing the volumetric imaging of the auroral ionosphere.

In addition to the theoretical exercise of analysing and simulating all the mission-critical aspects of the new radar system, we have constructed a working demonstrator array, funded from internal EISCAT resources, in order to verify that our design concept is actually feasible. The demonstrator array has been a major benefit to the project, enabling us to make important refinements to our initial ideas on array structure, antenna design and signal processing strategies. On this basis, we have established that an advanced radar system, fully meeting the performance targets set by our user community, could be constructed today using existing technology, if funding were not an issue.

Even in a report of this length, it has not really been possible to do justice to all of the technical work which has been done in connection with this study. Many further details can be found in the published deliverables from the various work packages of the design exercise, all of which are available on the project web site at www.eiscat.se/groups/EISCAT_3D_info, and in the publications in the open scientific literature which are listed in the appendices.

Conducting this study has demonstrated to many of us just how fast the key technologies are changing, in aspects of radar systems such as beam-forming, coding, signal processing and analysis. During the last four years, we have several times become aware of emerging new technologies with the potential to offer significant improvements in performance compared to our initial ideas, which reflected the capabilities available when we submitted the proposal. This is an important reason why we should not yet freeze the design; since we are still some time from assembling the funding to build the radar, it is certain that even better approaches to some of the problems will become available. Indeed we are already aware of promising new approaches in coding, software beamforming, inversion-based signal processing and data analysis which we need to explore in the next phase of the project.

The completion of the design study is only the start of the process of realising the EISCAT_3D facility. The project will now move into an implementation phase, in which further technical

aspects, including the issues connected with mass-producing the large number of required system elements, will now need to be addressed. In addition, there are still the considerable number of political, financial and scientific elements which need further work. We anticipate that these issues will be taken forward over the coming months, partly through nationally-funded activities and also via a further application for EU funding under the 7th framework. Several of the FP6 study team will remain engaged in these further activities, and those who will be moving on to other things will do so with our profound thanks and good wishes.

15 References

- [Alk78] H.J. Alker, *Measurement Principles in the EISCAT system*, EISCAT Technical Note 78/5
- [Baa77] J. W. M. Baars, R. Genzel, I. I. K. Pauliny-Toth and A. Witzel, *The Absolute Spectrum of Cas-A; An Accurate Flux Density Scale and a Set of Secondary Calibrators*, *Astron. Astrophysics* 61, 99-106 (1977)
- [Bad05] R. Badeau, B. David, and G. Richard, *Fast approximated power iteration subspace tracking*. *IEEE Trans. Signal Process.*, 53, 2005.
- [Bal82] C. A. Balanis, *Antenna Theory; Analysis and Design*, John Wiley & Sons (1982), ISBN 0-471-60352-X, page 368.
- [Bel07] Belyey, V., and C. La Hoz, *First results of an open-source based program for acquiring long time series of ionospheric parameters from the EISCAT madrigal database*, Abstract (2007), 13th EISCAT International Workshop August 6-10 2007, Åland, Finland.
- [Bra01] Bracewell, Ronald, *Fourier Analysis and Imaging*, 2004, Springer ISBN: 978-0-306-48187-1.
- [Bra88] Brasseur, G. and M.H. Hitchman, *Stratospheric Response to Trace Gas Perturbations: Changes in Ozone and Temperature Distributions* (1988), *Science*, 240, 634-637.
- [Cha08] Chau, J. L., D.L. Hysell, K.M. Kuyeng, and F.R. Galindo (2008), *Phase calibration approaches for radar interferometry and imaging configurations: equatorial spread F results*, *Ann. Geophys.*, 26, 2333-2343.
- [Cla99] Clark, B.G, *Coherence in Radio Astronomy*, Chapter 1 of Taylor et al., 1999. [Tay99]
- [Com90] P. Comon and G. H. Golub. *Tracking a few extreme singular values and vectors in signal processing*. *Proc. IEEE*, 78(8), Aug. 1990.
- [Cor99] Cornwell, T., *Deconvolution*, Chapter 8 of Taylor et al., 1999 [Tay99].
- [D1.1] EISCAT 3D Deliverable D1.1, 1st Annual Report of the EISCAT_3D Design Study Period. 2005-05-01 - 2006-04-30
- [D1.2] EISCAT 3D Deliverable D1.2, 2nd Annual Report of the EISCAT_3D Design Study Period: 2006-05-01 - 2007-04-30, published 2007-06-12
- [D1.3] EISCAT 3D Deliverable D1.3, 3rd Annual Report of the EISCAT_3D Design Study Period: 2007-05-01 - 2008-04-30, published 2008-06-10
- [D2.1] EISCAT 3D Deliverable D2.1, *EISCAT_3D Radar Performance Specification Document* Ver 1.1 November 7, 2005
- [D3.2] EISCAT 3D Deliverable D3.2, *Options for the Active Element*, 2009-02-19.
- [D5.1] EISCAT 3D Deliverable D5.1, *EISCAT_3D Radar Imaging Arrays Configurations Report*, submitted September 27, 2006

- [D5.2] EISCAT 3D Deliverable D5.2, *Radar Imaging Algorithms Report*, submitted January 30, 2008
- [D5.3] EISCAT 3D Deliverable D5.3, *EISCAT_3D Radar Multidimensional Imaging Radar Data Visualisation Report*, Published 2009-03-12
- [D6.1] EISCAT 3D Deliverable D6.1, Interim Report *A report on WP6 progress as of Feb 1, 2009*
- [D7.1] EISCAT 3D Deliverable D7.1, *Basic set of control and monitoring software primitives for the WP4 demonstrator*, published 2009-03-12
- [D7.2] EISCAT 3D Deliverable D7.2, *Control and Monitoring Subsystem*, published 2009-03-15
- [D8.1] EISCAT 3D Deliverable D8.1, Interim report on progress in Design Study Work Package 8 (EISCAT 3-D Data System) submitted February 3rd 2006
- [D8.2] EISCAT 3D Deliverable D8.2, *Design of the data archiving and distribution system*, 2007-04-12.
- [D8.3] EISCAT 3D Deliverable D8.3, *Data archiving and distribution technical study*, 10th September 2008
- [D8.4] EISCAT 3D Deliverable D8.4, *EISCAT-3D as a geophysical observatory supporting instruments and the effect of their data on the design of the archive systems*. 2009-03-09.
- [D8.5] EISCAT 3D Deliverable D8.5, *Required interfaces to the data and archive system*, 2009-02-04.
- [D9.1] EISCAT 3D Deliverable D9.1, *Basic set of processing primitives for operation of the Demonstrator*, published 2009-04-17
- [D9.3] EISCAT 3D Deliverable D9.3, *Polarisation Tracking*. 2009-05-17.
- [D10.1] EISCAT 3D Deliverable D10.1, *New techniques, Progress Report*, 19 March 2006
- [D10.2a] EISCAT 3D Deliverable D10.2a, *Use of Long Time Series ISR Data for Improving Ionospheric Correction Models at High-Latitudes*, submitted 2007-04-27
- [D10.2b] EISCAT 3D Deliverable D10.2b, *First results of an open-source based program for acquiring long time series of ionospheric parameters from the EISCAT Madrigal database*, submitted 2007-04-27
- [D10.6] EISCAT 3D Deliverable D10.6, *How EISCAT_3D can contribute to studies on Global Change*, published 2009-05-02
- [D12.2] EISCAT 3D Deliverable D12.2, *Final report on network and signal transport*, 2009-04-07.
- [D13.1] EISCAT 3D Deliverable D13.1, *Additional input material related to the preparation of D1.5/D11.1*, published 2009-4-16

[D13.3] EISCAT 3D Deliverable D13.3, *Review of steps to procure use of potential sites*, published 2009-4-16

[D13.4] EISCAT 3D Deliverable D13.4 *Review of other issues affecting the possibility to move forward*. 2009-04-07.

[DSD] *EISCAT 3D Design Specification Document*.
http://www.eiscat.se/groups/EISCAT_3D_info/P_S_D_7.pdf.

[Dou08] X. Doukopoulos and G.V. Moustakides. *Fast and Stable Subspace Tracking* IEEE Trans. on Signal Processing, 56(4), 2008.

[EAHM2] EISCAT 3D Notes from the 2nd EISCAT_3D All-Hands Meeting, IRF-Kiruna, Sept 16-17, 2008

[EAR05a] EISCAT 3D 1st Activity Report - 1st Activity Report of the EISCAT_3D Design Study Period. 2005-05-01 - 2006-04-30

[EAR06a] EISCAT 3D 2nd Activity Report - 2nd Activity Report of the EISCAT_3D Design Study Period. 2006-05-01 - 2007-04-30, published 2007-06-12

[EAR07a] EISCAT 3D 3rd Activity Report - 3rd Activity Report of the EISCAT_3D Design Study Period. 2007-05-01 - 2008-04-30, published 2008-06-10

[EIB] *EISCAT 3D Introduction and background*

[ENG] EISCAT 3D : The next generation ISR system "EISCAT_3D" - the next generation European Incoherent SCATter radar system

[ENI]EISCAT 3D National Inputs

[Far81] Farley, D. T., H. M. Ierkeic, and B. G. Fejer (1981), *Radar Interferometry: A New Technique for Studying Plasma Turbulence in the Ionosphere*, J. Geophys. Res., 86(A3), 14671472.

[Gro94] W. Grover, *A new method for clock distribution*,. IEEE Trans. Circuits Syst. I, Fundam. Theory Appl. (USA), vol. 41, no. 2, pp. 149 - 60, Feb 1994.

[Gry03] Grydeland, T., C. La Hoz, T. Hagfors, E.M. Blixt, S. Saito, A. Strømme, and A. Brekke (2003), *Interferometric observations of filamentary structures associated with plasma instability in the auroral ionosphere*, Geophys. Res. Lett., 30, 1338, doi:10.1029/2002-GL016362.

[Gry07] Grydeland, T., C. La Hoz and V. Belyey, *Phase calibration of an aperture synthesis imaging array by means of incoherent scattering*, Abstract (2007), 13th EISCAT International Workshop August 6-10 2007, Aland, Finland.

[Hys03] Hysell, D. L., and J. L. Chau (2006), *Optimal aperture synthesis radar imaging*, Radio Sci., 41, RS2003, doi:10.1029/2005RS003383.

[Hys97] Hysell, D. L., and R. F. Woodman (1997), *Imaging coherent backscatter radar observations of topside equatorial spread F*, Radio Sci., 32(6), 23092320.

- [Jac99] J.D. Jackson. *Classical Electrodynamics*. Wiley, 3rd edition, 1999.
- [Joh08] G. Johansson, J. Borg, J. Johansson, M. L. Nordenvaad, and G. Wannberg, *Simulation of Post-ADC Digital Beam-Forming for Large Aperture Array Radars*, submitted to Radio Science, 2008
- [Kay93] S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory*, Prentice Hall, New Jersey, USA, 1993.
- [Kil00] P-S. Kildal, *Foundations of Antennas, a Unified Approach*, Studentlitteratur, Lund, Sweden, 2000.
- [Kri96] H. Krim and M. Viberg. *Two Decades of Array Signal Processing Research: The Parametric Approach*. IEEE Signal Processing Magazine, 13(4):67-94, 1996.
- [Kud91] Kudeki, E., and F. uuu (1991), *Radar Interferometric Imaging of Field-aligned Plasma Irregularities in the Equatorial Electrojet*, Geophys. Res. Lett., 18(1), 4144.
- [Leh87] Lehtinen, M. And Häggström, I. (1987): *A new modulation principle for incoherent scatter measurements*. Radio Sci. **22** 625-634.
- [Lin08] T. Lindgren and J. Ekman, *Performance of a Yagi antenna during snowfall*, International Symposium on Antennas and Propagation, Taipei, Taiwan, October 2008.
- [Lin09] T. Lindgren and J. Ekman, *A Measurement System for the Complex Far-Field of Physically Large Antenna Arrays under Noisy Conditions Utilizing the Equivalent Electric Current Method*, submitted to IEEE Trans. Antennas Propagat., 2009.
- [Mar06] J. Markkanen, *A Matlab package for Phased Array Beam Shape Inspection*, www.eiscat.se/groups/EISCAT_3D_info/e3dant.pdf (2006)
- [McK06] McKay, D.J., *EISCAT-3D WP4/5/8/9 signal processing meeting*. Summary of discussions from the meeting at 2006-07-03, 05.
- [Mes95] X. Mestre. *Improved Estimation of Eigenvalues and Eigenvectors of Covariance Matrices Using Their Sample Estimates*. IEEE Transactions on Information Theory, 54(11), 2008.
- [Mis06] P. Misra and P. Enge, *Global Positioning System, Signals, Measurements, and Performance*, 2nd ed. P.O. Box 692, Lincoln, Massachusetts 01773: Ganga-Jamuna Press, 2006.
- [MTR] EISCAT_3D Mid-term Review Presentations, published 2007-10-23
- [Ren07] T. Renkwitz, M. Sc. thesis: “*Analysis and Optimisation of Medium Gain X-Yagi Antennas for the EISCAT 3D 237.5MHz Incoherent Scatter Radar Active Array*”, www.eiscat.se/groups/EISCAT_3D_info/MScWorkToralfRenkwitzEISCAT3D (2007)
- [Ris90] Rishbeth, H., *A greenhouse effect in the ionosphere?* (1990), Planet. Space Sci., 38, 945-948, doi:10.1016/0032-0633(90)90061-T.

[Rob89] Roble, R.G., and R.E. Dickinson, *How will changes in carbon dioxide and methane modify the mean structure of the mesosphere and thermosphere?* (1989), *Geophys. Res. Lett.*, 16,1441-1444.

[SAC03] EISCAT 3D SAC Report October 2003

[SAC71] EISCAT 3D Documents from SAC71

[Ste06] G. Stenberg, J. Borg, J. Johansson, and G. Wannberg, *Simulation of Post-ADC Digital Beam-Forming for Large Area Radar Receiver Arrays*, in *International RF and Microwave Conference*, 2006.

[Ste07] G. Stenberg, *Advancement of Atmospheric Research Tools*, Licentiate Thesis, Luleå University of Technology, 2007.

[Ste08] G. Stenberg, T. Lindgren, and J. Johansson, *A Picosecond Accuracy Timing System Based on L1-only GNSS Receivers for a Large Aperture Array Radar*, in *ION GNSS 2008*. Institute of Navigation, 2008.

[Sto97] P. Stoica and R.L. Moses. *Introduction to Spectral Analysis*. Prentice Hall, 1997.

[Sve05] L. Svensson and M. Lundberg. *On Posterior Distributions for Signals in Gaussian Noise with Unknown Covariance Matrix*. *IEEE Transactions on Signal Processing*, 53:3554--3571, 2005.

[Tay99] Taylor, G.B. (Ed.), C. L. Carilli (Ed.) and R. A. Perley (Ed.) (1999), *Synthesis Imaging in Radio Astronomy II*, ASP Conference Series, Vol. 180; xxxiii, 704p; California, USA; ISBN 1583810056.

[Tho] Thomas, J.J, and K.A. Cook, *Illuminating the Path: The Research and Development Agenda for Visual Analytics*. National Visualisation and Analytics Centre (NVAC). Online http://nvac.pnl.gov/docs/RD_Agenda_VisualAnalytics.pdf

[Tro01] Trondsen, T. S., and L. L. Cogger (2001), *Fine-scale optical observations of aurora*, *Phys. Chem. Earth (C)*, 26 (1-3), 179188, doi:10.1016/S1464-1917(00)00105-7.

[Ull00] Ullich, T., *Solar variability and long-term trends in the mesosphere* (2000), Ph. D. thesis, Sodankylä Observatory, Sodankylä, Finland.

[WP5a] EISCAT 3D WP5 Report - *Considerations and requirements regarding WP5.Interferometry*, September 10, 2005

[WP5b] EISCAT 3D WP5 Report - *Final Radar Imaging Algorithms Report*, submitted August, 2008

[WP5c] EISCAT 3D WP5 Report - *Fundamentals of radar interferometry I. One baseline Stage 2 report*, April 2006

[WP5d] EISCAT 3D WP5 Report - *Radar Interferometer Phase Calibration using the Visibility function of Incoherent Scattering Report*, May 2, 2007

[WP5e] EISCAT 3D WP5 Report - *Report on Synchronisation for interferometry*, September 2006

[WP8a] EISCAT 3D WP8 Report - February 3 2006 *Progress report on Work Package 8 of the EISCAT-3D Design Study*, submitted February 3rd 2006

[Wan93] G. Wannberg. *Wave Polarisation and Software Polariser for the EISCAT UHF System*. Technical report, 09/57, EISCAT, www.eiscat.se/groups/Documentation/BlueBooks/Polariser, 2009.

[Woo71] Woodman, R. F. (1971), *Inclination of the Geomagnetic Field Measured by an Incoherent Scatter Technique*, J. Geophys. Res., 76(1), 178184.

[Woo97] Woodman, R. F. (1997), *Coherent radar imaging: Signal processing and statistical properties*, Radio Sci., 32(6), 23732391.

[Woo01] Woodman, R. F., and J. L. Chau (2001), *Antenna compression using binary phase coding*, Radio Sci., 36 (1), 4551, doi:10.1029/2000RS002388.

[Yan94] R. Yang and J. Berger. *Estimation of a Covariance Matrix Using the Reference Prior*. The Annals of Statistics, 22(3):1195--1211, 1994.

[Yan95] B. Yang. *Projection Approximation Subspace Tracking*. IEEE Trans. on Signal Processing, 43(1), 1995.

This page is intentionally blank.

APPENDICES



16 Appendix: Complete List of Reports & Deliverables

The following is a list of the publicly available reports and deliverables produced during the EISCAT_3D design study, mainly comprising the deliverables and reports from the 13 work packages. These documents can be downloaded from the open access pages of the EISCAT web site, and have been collected together at: www.eiscat.se/groups/EISCAT_3D_info/

A: Background and Planning Documents

A brief high-level summary of the EISCAT_3D project can be found at: [www.eiscat.se/groups/EISCAT_3D_info/Introduction and Brief Background.pdf](http://www.eiscat.se/groups/EISCAT_3D_info/Introduction%20and%20Brief%20Background.pdf)

The EISCAT Scientific Advisory Committee first considered the scientific case for a next-generation EISCAT radar in October 2003. Their report can be found at: www.eiscat.se/groups/EISCAT_3D_info/overall_summary.doc

In addition, a number of EISCAT member countries produced “white papers” setting out their requirements for a new radar. These can be found at: www.eiscat.se/groups/EISCAT_3D_info/inputs

B: Original Application and Description of Work

The current version of “Annex 1” to the above contract, specifying the work to be done in each work package, can be found at: www.eiscat.se/groups/EISCAT_3D_info/Annex_I_1May2008

C: Work Package Outputs

The majority of the available material consists of the reports and deliverables of the thirteen work packages of the design study, as specified in the agreed description of work (Annex 1). The available documents are as follows:

Work Package 1: Management of Design Study (EISCAT)

This was the work package dedicated to the management of the design study and the reporting of the financial and administrative status to the EU. Two documents were published in each year of the study, a deliverable and an activity report, summarising the year's activities. In total, five deliverables were planned from this work package – the four annual reports, plus a final report summarising the overall achievements of the study.

Deliverable 1.1: 1st Annual Report of the EISCAT_3D Design Study. Period: 2005-05-01 - 2006-04-30

www.eiscat.se/groups/EISCAT_3D_info/1st_Annual_Report_EISCAT_3D.pdf

Year One Activity Report: 1st Activity Report of the EISCAT_3D Design Study. Period: 2005-05-01 - 2006-04-30

www.eiscat.se/groups/EISCAT_3D_info/1st_Activity_Report_EISCAT_3D.pdf

Deliverable 1.2: 2nd Annual Report of the EISCAT_3D Design Study. Period: 2006-05-01 - 2007-04-30, published 2007-06-12

www.eiscat.se/groups/EISCAT_3D_info/2nd_Annual_Report_EISCAT_3D.pdf

Year Two Activity Report: 2nd Activity Report of the EISCAT_3D Design Study. Period: 2006-05-01 - 2007-04-30, published 2007-06-12

www.eiscat.se/groups/EISCAT_3D_info/2nd_Activity_Report_EISCAT_3D.pdf

Deliverable 1.3: 3rd Annual Report of the EISCAT_3D Design Study. Period: 2007-05-01 - 2008-04-30, published 2008-06-10

www.eiscat.se/groups/EISCAT_3D_info/3rd_Annual_Report_EISCAT_3D.pdf

Year Three Activity Report: 3rd Activity Report of the EISCAT_3D Design Study. Period: 2007-05-01 - 2008-04-30, published 2008-06-10

www.eiscat.se/groups/EISCAT_3D_info/3rd_Activity_Report_EISCAT_3D.pdf

Deliverable 1.4: 4th Annual Report of the EISCAT_3D Design Study. Period: 2008-05-01 - 2009-04-30, published 2009-06-10

www.eiscat.se/groups/EISCAT_3D_info/4th_Annual_Report_EISCAT_3D.pdf

Year Four Activity Report: 4th Activity Report of the EISCAT_3D Design Study. Period: 2008-05-01 - 2009-04-30, published 2009-06-10

www.eiscat.se/groups/EISCAT_3D_info/4th_Activity_Report_EISCAT_3D.pdf

Deliverable 1.5: Summary of design study activities

www.eiscat.se/groups/EISCAT_3D_info/Deliverable_D1_5.pdf

Work Package 2: Evaluation of Design Performance Goals (EISCAT)

This work package was charged with surveying the users' requirements for the performance of the new radar, and synthesising these into a specification document, which formed the only deliverable from this work package.

Deliverable 2.1: EISCAT_3D Radar Performance Specification Document. November 7, 2005
www.eiscat.se/groups/EISCAT_3D_info/P_S_D_7.pdf

Work Package 3: Evaluation of the Options for the Active Element (EISCAT, IRF)

The aim of this work package was to survey the possible different approaches for the transmission capabilities of the new radar and evaluate the effectiveness of each in terms of scientific return, performance and ease of manufacture. Two deliverables were originally planned, however the work plan was later modified so that all of the outputs of Work Package 3 were combined into a single deliverable.

Deliverable 3.2: Options for the active element: February 19, 2009
[www.eiscat.se/groups/EISCAT_3D_info/D3.2 Options for the Active Element](http://www.eiscat.se/groups/EISCAT_3D_info/D3.2%20Options%20for%20the%20Active%20Element)

Work Package 4: Phased Array Receivers (LTU)

This work package was tasked with developing the concept for the receiver and antenna systems capable of realising a multistatic multiple-beam phased array system. Two deliverables were originally planned, however the work plan was later modified so that all of the outputs of Work Package 4 were combined into a single document.

Deliverable 4.2: EISCAT_3D radar receiver/antenna subsystem report: May 30, 2009.
www.eiscat.se/groups/EISCAT_3D_info/d42

Work Package 5: Interferometry (UiT)

This work package was responsible for studying the requirements needed to realise an interferometric imaging capability for the EISCAT_3D radar system, so that observations enabling the resolution of sub-beamwidth scale scattering structures could be made as a standard part of the radar operations. Three deliverables were planned for this work package, and all were produced. In addition the work package produced regular reports summarising its activities. The final report was an overview of the progress of the entire work package.

WP5 first report: Considerations and requirements regarding WP5, September 10, 2005
www.eiscat.se/groups/EISCAT_3D_info/WP5_Considerations_and_requirements_regarding_Interferometry2005-09-10.pdf

WP5 second report: Fundamentals of radar interferometry I: One baseline Stage 2 report, April 2006
www.eiscat.se/groups/EISCAT_3D_info/WP5_Interferometry_Stage2_Report_April06.pdf

WP5 third report: Report on Synchronisation for interferometry, September 2006
www.eiscat.se/groups/EISCAT_3D_info/WP5_Report_on_Synchronisation_for_interferometry_September2006.pdf

Deliverable 5.1: EISCAT_3D Radar Imaging Arrays Configurations Report, September 27, 2006
www.eiscat.se/groups/EISCAT_3D_info/D5_1_EISCAT_3D_Radar_Imaging_Arrays_Configurations_Report.pdf

WP5 fourth report: Radar Interferometer Phase Calibration using the Visibility function of Incoherent Scattering Report, May 2, 2007
www.eiscat.se/groups/EISCAT_3D_info/WP5_Report3_Radar_Interferometer_Phase_Calibration_using_the_Visibility_function_of_Incoherent_Scattering_2007-05-02.pdf

Deliverable 5.2: Radar Imaging Algorithms Report, January 30, 2008
www.eiscat.se/groups/EISCAT_3D_info/D5.2_WP5_Report_on_Image_Inversion_for_Interferometry_January2008.pdf

WP5 fifth report: Final Radar Imaging Algorithms Report, August, 2008
www.eiscat.se/groups/EISCAT_3D_info/WP5_FinRep_Image_Inversion_August2008.pdf

Deliverable 5.3: EISCAT_3D Radar Multidimensional Imaging Radar Data Visualisation Report, March 12 2009
www.eiscat.se/groups/EISCAT_3D_info/D5.3_WP5_Visualization_report.pdf

WP5 Final Report: Final Report on Work Package 5. May 8 2009.
www.eiscat.se/groups/EISCAT_3D_info/wp5_final_report.pdf

Work Package 6: Active Element (IRF)

This work package proceeded in parallel with the later stages of WP3, designing the active component of the EISCAT_3D radar including the generation and distribution of the RF signal, the transmitting antennas and the necessary control and monitoring systems. Two deliverables were planned, both of which were achieved.

Deliverable 6.1: Interim Report on Work Package 6, February 1 2009
[www.eiscat.se/groups/EISCAT_3D_info/D6.1 Interim Report](http://www.eiscat.se/groups/EISCAT_3D_info/D6.1_Interim_Report)

Deliverable 6.2: Active Element Subsystem Design, May 30, 2009.
[www.eiscat.se/groups/EISCAT_3D_info/D6.2 draft 1.doc](http://www.eiscat.se/groups/EISCAT_3D_info/D6.2_draft_1.doc)

Work Package 7: Distributed Control and Monitoring (EISCAT)

This work package was tasked with evaluating the various possible concepts for control and monitoring of the various functions of a highly distributed multi-element radar system, and making a recommendation for how these should be handled. Two deliverables were planned and achieved.

Deliverable 7.1: Basic set of control and monitoring software primitives for the WP4 demonstrator, March 12 2009
[www.eiscat.se/groups/EISCAT_3D_info/D7_1 Basic set of control and monitoring primitives for the WP4 demonstrator array 2009-03-12.pdf](http://www.eiscat.se/groups/EISCAT_3D_info/D7_1_Basic_set_of_control_and_monitoring_primitives_for_the_WP4_demonstrator_array_2009-03-12.pdf)

Deliverable 7.2: EISCAT_3D Radar Control and Monitoring Subsystem Report, March 15 2009
[www.eiscat.se/groups/EISCAT_3D_info/D7_2 Radar Control and Monitoring Subsystem Report 2009-03-15.pdf](http://www.eiscat.se/groups/EISCAT_3D_info/D7_2_Radar_Control_and_Monitoring_Subsystem_Report_2009-03-15.pdf)

Work Package 8: Data Archiving and Distribution (RAL, UiT)

This work package was assigned the task of designing a data-handling system capable of coping with the very large data volumes expected from EISCAT_3D. It included both short-term storage and long-term archiving, as well as the consideration of data visualisation and user interactions with the various archiving systems. Six deliverables were planned and achieved.

Deliverable 8.1: Interim report on progress in Design Study Work Package 8 (EISCAT 3-D Data System) February 3rd 2006

www.eiscat.se/groups/EISCAT_3D_info/stage_2_report_final.pdf

Deliverable 8.2: High-level design of the data archiving and distribution system. April 12 2007.

www.eiscat.se/groups/EISCAT_3D_info/deliverable_8_2_final.pdf

Deliverable 8.3: Data archiving and distribution system for EISCAT-3D. February 12 2009.

www.eiscat.se/groups/EISCAT_3D_info/d83/

Deliverable 8.4: EISCAT-3D as a geophysical observatory. March 9 2009.

www.eiscat.se/groups/EISCAT_3D_info/d84_20090309.pdf

Deliverable 8.5: Required interfaces to the data and archive system. February 4 2009.

www.eiscat.se/groups/EISCAT_3D_info/deliverable_8_5_090130.pdf

Deliverable 8.6: Visualisation for EISCAT-3D. March 18 2009.

www.eiscat.se/groups/EISCAT_3D_info/visualisation.pdf

Work Package 9: Signal Processing (LTU, EISCAT, IRF)

This work package focused on the remaining signal processing aspects not already covered by other work packages. These included the issues of polarisation control and signal processing optimisation in low SNR situations. Five deliverables were planned, all of which were achieved.

Deliverable 9.1: Basic set of processing primitives for operation of Demonstrator. April 17 2009.
www.eiscat.se/groups/EISCAT_3D_info/Deliverable_9_1

Deliverable 9.2: Signal Processing Subsystem Report, June 5 2009
http://e7.eiscat.se/groups/EISCAT_3D_info/Deliverable_D9_2.pdf

Deliverable 9.3: Polarisation Estimation using Eigenspace Analysis. May 17 2009.
www.eiscat.se/groups/EISCAT_3D_info/D9.3/

Deliverable 9.4: Complex Covariance Matrices with Efficient Implementation Strategies
www.eiscat.se/groups/EISCAT_3D_info/D9.4/

Deliverable 9.5: Polarisation Estimation using Subspace Tracking
www.eiscat.se/groups/EISCAT_3D_info/D9.5/

Work Package 10: New Ways of Exploiting Incoherent Scatter Radars (UiT)

The aim of this work package was to conduct a survey of the possible “non-traditional” uses of incoherent scatter radars and explore the potential of implementing such uses on the EISCAT_3D systems. Six deliverables were originally planned from this study, however modifications in the work plan meant that only four deliverables were produced (though one of these was in two parts).

Deliverable 10.1: First progress report on WP10. June 9 2006.

www.eiscat.se/groups/EISCAT_3D_info/D10_1_Progress_Report2006-03-19.pdf

Deliverable 10.2 (part 1): Use of Long Time Series ISR Data for Improving Ionospheric Correction Models at High-Latitudes. April 27 2007.

www.eiscat.se/groups/EISCAT_3D_info/D10.2_Progress_Report_Part1_2007-04-27.pdf

Deliverable 10.2 (part 2): First results of an open-source based program for acquiring long time series of ionospheric parameters from the EISCAT Madrigal database, April 27 2007.

www.eiscat.se/groups/EISCAT_3D_info/D10.2_Progress_Report_Part2_2007-04-27.pdf

Deliverable 10.3: How EISCAT_3D can contribute to studies on Global Change. May 5 2009.

www.eiscat.se/groups/EISCAT_3D_info/D10.6_How_EISCAT_3D_can_contribute_to_studies_on_Global_Change.pdf

Deliverable 10.4: Final report on Work Package 10. May 1 2009.

www.eiscat.se/groups/EISCAT_3D_info/wp10_final_report.pdf

Work Package 11: Implementation Blueprint (EISCAT, RAL)

This work package was tasked with producing the final technical summary of the design study, combining the outputs from all of the other work packages into a high-level system description. The present document, containing a final technical overview of the whole design study, is the only deliverable of this work package.

Deliverable 11.1: Final Report of the EISCAT_3D design study (this document!). June 7 2009

www.eiscat.se/groups/EISCAT_3D_info/D11.1/

Work Package 12: Networking and Reference Time and Frequency (EISCAT)

The need to maintain time synchronisation across the antenna arrays at an accuracy less than a few tens of picoseconds constitutes a significant challenge for the development of EISCAT_3D. This work package was tasked with examining possible methods for guaranteeing time synchronisation, and developing a technical solution to ensure that the required accuracy could be maintained. Two deliverables were proposed from this work package, both of which were achieved.

Deliverable 12.1: Time, Frequency and Synchronisation Subsystem Report. December 18 2008
www.eiscat.se/groups/EISCAT_3D_info/deliverable-12-1

Deliverable 12.2: EISCAT-3D radar array: networking and communication report. April 7 2009.
www.eiscat.se/groups/EISCAT_3D_info/DeliverableWP12.2

Work Package 13: Enabling Actions (EISCAT, IRF)

This work package was set up to explore the issues related to the identification of actual radar sites, the securing of operating permissions to use the selected radar frequency and other similar work concerning the logistical details of developing EISCAT_3D. Five deliverables were proposed for this work package, and four were realised. The planned review of actual or potential funding sources (D13.2) was deferred until a later phase of the study.

Deliverable 13.1: Additional WP13 input material related to the preparation of D1.5/D11.1. April 16 2009.
www.eiscat.se/groups/EISCAT_3D_info/Deliverable_13.1/

Deliverable 13.3: Review of steps to procure use of potential sites. April 16 2009.
www.eiscat.se/groups/EISCAT_3D_info/Deliverable_13.3

Deliverable 13.4: Risks in site selection. April 7 2009.
www.eiscat.se/groups/EISCAT_3D_info/Deliverable_WP13.4

Deliverable 13.5: Frequency Allocations: Status Report, May 30 2009.
http://e7.eiscat.se/groups/EISCAT_3D_info/del135

17 Appendix: Publications in open literature

The performed work has resulted in a number of publications in international conferences and journals, as well as academic publications for M.Sc., licentiate and Ph.D. degrees. Additionally, contributions relevant to the design study have been made to EISCAT workshops. Below is a complete list of this material.

17.1 Journal publications

A Pico-Second Level Cable-Based Calibration System for Large Aperture Array Radars, G. Stenberg, J. Borg, and J. Johansson, To be submitted.

Simulation of Post-ADC Digital Beam-forming for Large Aperture Array Radars, G. Stenberg, J. Borg, J. Johansson, M. Lundberg-Nordenvaad, and G.Wannberg, under revision for Radio Science

The Reference Prior for Complex Covariance Matrices with Efficient Implementation Strategies, L. Svensson and M. Lundberg-Nordenvaad, IEEE Transactions on Signal Processing, 2009

A Measurement System for the Complex Far-Field of Physically Large Antenna Arrays under Noisy Conditions Utilizing the Equivalent Electric Current Method, T. Lindgren and J. Ekman, submitted to IEEE Transactions on Antennas and Propagation, 2009.

17.2 Conference publications

Picosecond Level Error Detection using PCA in the Hardware Timing Systems for the EISCAT 3D LAAR, G. Johansson, F. Häggglund, and J. Carlsson, To be published

Simulation of Post-ADC Digital Beam-Forming for Large Area Radar Receiver Arrays, G. Stenberg, J. Borg, J. Johansson, and G. Wannberg, IEEE International RF and Microwave Conference 2006

Temperature stabilization of electronics module, A. Gabert, J. Borg, J. Johansson, IMAPS Nordic Annual Conference 2006

A Picosecond Accuracy Timing System Based on L1-only GNSS Receivers for a Large Aperture Array Radar, G. Stenberg, T. Lindgren, and J. Johansson, ION GNSS 2008.

Low-complex maximum likelihood estimator for the linear model, J. Ståbis and M. Lundberg-Nordenvaad, To be published

Performance of a Yagi antenna during snowfall, T. Lindgren and J. Ekman, International Symposium on Antennas and Propagation, Taipei, Taiwan, October 2008.

17.3 Doctoral theses

Doctoral thesis of G. Stenberg, 2009 (in writing).

Doctoral thesis of T. Lindgren, 2009 (in writing).

17.4 Licentiate thesis

Advancement of atmospheric research tools, G. Stenberg, 2006

17.5 M.Sc. theses

Low-Noise Amplifier Design and Optimization, M. Edvall, 2009

Temperature stabilization of electronics module, A. Gabert, 2006

Analysis and Optimisation of Medium Gain X-Yagi Antennas for the EISCAT 3D 237.5MHz Incoherent Scatter Radar Active Array, T. Renkwitz, 2007

RF front end for atmospheric radar, K. Söderström, (in preparation 2009)

17.6 EISCAT workshops and user meetings

A large number of presentations have been specifically made on the EISCAT_3D project at the following workshops and user meetings:

- Kiruna 2005
- Åland / Sodankylä 2007
- Uppsala 2009

— End of Document —