WP8: Data segment for EISCAT 3D Radar

Stage 2 Report – Submitted Version

2 February 2006, IWM, DJM, *et al.*

Introduction

This document constitutes the first external progress report on Work Package 8 (WP8) of the EISCAT-3D design study. The remit of the work package is to design a system to facilitate the efficient collection and distribution of large volumes of data across the geographically extended network of future EISCAT facilities. It includes the data transfer between the instruments themselves, and from the instruments to the users and consumers of their data products. In particular, the responsibilities of Work Package 8 comprise the design of the Data Distribution System and the Secure Data Archive, as well as the design of data visualization and assimilation tools which will allow the results from the new radars to be utilised efficiently and effectively.

The following sections summarise the progress that has been made in this Work Package since the start of the project in May 2005. Although it is inappropriate to make a detailed design for the data system at this stage of the project, we can demonstrate that, even with existing technologies, potential solutions are available to deal with the high data rates which we expect to encounter in the EISCAT-3D system. In a number of areas, we have identified “interface issues” where the design of the data flow system will be influenced by considerations relevant to other work packages. These issues are under discussion with the work packages concerned, and the solutions will be clarified in future progress reports.

Terminology

Since we anticipate high data volumes in some parts of the data flow, choices must be made for the type and retention duration of stored data. To clarify the different levels of data generated by the 3D system, the following nomenclature is used:

a) Sample-level data – these consist of the complex-amplitude data coming from each element of the receiver array. These data are produced by the first-stage digitisation of the received signal. They are time-sequenced and are without integration. It is anticipated that these data will not be stored. The handling of sample-level data falls within the remit of Work Package 4; data will only enter into the data system developed under Work Package 8 after beam-forming.

b) Beam-formed data – these consist of the weighted sum of a number of streams of sample level data, each of which is multiplied by a different complex-valued weighting coefficient in order to form a “beam”. Like sample-level data, these data are time-sequenced, are without integration, remain uncorrelated, and still have a time resolution at the sub-IPP level. At the receive-only sites a number of such beam-formed data streams will be generated concurrently. Discussions of the precise interface and format specifications under which beam-formed data will enter the WP8 data system are underway elsewhere within the design study. The data rate will depend on several aspects of the data acquisition strategy. These include whether the data are split into ion-line and plasma-line streams (with different bandwidths). Thus the design of the data handling system at this level will depend on the specifications which come forth from WP4 and from WP9, which will be handling the signal processing tasks including beam-forming, data decimation, filtering and correlation. However, the data rate from a complete beam, will not exceed 40 Msamples/s (16-bit samples), thus resulting in rates of 288 Gbytes/hr/beam. This assumes that allowance has been made for the fact that the radar will not be transmitting with 100% duty cycle.

c) Correlated data – these data constitute the outputs from the signal processing tasks carried out by WP9, comprising the auto-correlated power domain data from each beam. A correlated data set will generally be appreciably larger than a beam-formed data set, because correlation will involve generating the whole lag profile matrix. However, this inflation in size may be mitigated by post-integrating in time (typically in the order of 1 second). There is no requirement that correlated data from all experiments should have the same time resolution – the time resolution will depend on the properties of the experiment (as at present).

d) Analysed data – these data are formed by appropriately post-integrating the correlated data and then passing them through an analysis program such as the GUISDAP code presently used by EISCAT. Automatic detection (and excision) of spurious signals, and the automatic recognition of scientifically interesting data may occur before post-integration. At this level, the data volumes will be roughly equivalent to those which we deal with in the current EISCAT system, except that several equivalent data sets can exist simultaneously (corresponding to a number of simultaneous beams).

e) Metadata – these are time-stamped data, containing comprehensive information about the system state and radar operating mode, as well as data from co-located and/or remote auxiliary instrumentation. They are automatically generated whenever the system is operating, such that every data-set of type (a-d) automatically gets a metadata set associated with it. Analysis and visualisation software processing any of the other data types can reference the metadata to retrieve the data context and to configure itself appropriately. Metadata can also be accessed and visualised independently for e.g. event-search or maintenance purposes.

Summary of Data Products

EISCAT will offer standard data products from the 3D system at each of the following levels:

- Beam-formed data at pulse-to-pulse time resolution will be stored in ring buffers, the duration of which will be as long as practicable, but initially at least several hours, in a manner which allows for users to access and copy selected intervals. Because of bandwidth and data volume considerations, the beam-formed data may have to be stored in an IO-optimised format.

- At least one set of time-integrated correlated data will be calculated from each set of beam-formed data, and permanently stored in a Grid/Web accessible master archive.

- At least one, and possibly several, analysed data sets will be permanently stored corresponding to each set of correlated data. These will be processed with welldocumented analysis strategies, at least one of which will be a standard strategy making the data suitable for use in “quick-look” applications and statistical studies (provision of the data analysis software is however not within the scope of the design study). Production of the analysed data will occur in real-time, as a matter of course, both for radar operators and other users wishing to monitor and access the real-time data.

- The primary data products will also be used to provide value-added parameters (examples may include velocities, conductivities, currents, heating rates).

In addition to the above standard data products, the following data products will also be available from some (but not necessarily all) experiments using the EISCAT-3D radars.

- Data on the properties of measured spectra (both ion-line and plasma line) comprising spectral power, numbers and locations of spectral peaks, asymmetries, spectral moments, etc. The extent to which such data will be available will depend on the details of the particular experiment, and the level at which data are already processed by WP9.

- Data produced by interferometry applications (cross-phase, coherence and reconstructed visibility functions) will be stored in a ring buffer of sufficient duration to allow groups interested in interferometry applications to transfer them elsewhere. Threshold logic will be used to determine which data need to be stored. Sufficient metadata will be generated to describe which data sets were used and how the interferometry calculations were carried out. The detailed requirements on the data and metadata for interferometry applications will be established in collaboration with the University of Tromsø, who are leading the development of Work Package 5.

User Control & Scheduling

In addition to the handling of the actual data streams, there is additional data provided to the systems of WP8 from the control and monitoring systems of EISCAT-3D (part of WP7). These data are described as follows:

- Provision of control data to WP8, in the sense that there is a need to control where to channel data, when, etc. These are essentially the control data to WP8 systems, both in terms of scheduled data acquisition, and as ad hoc commands (such as stop, abort, reset, etc.).

- Provision of control data that will become metadata. These are the same as the engineering data that is sent on from the control and monitoring system, but which are to be incorporated into the data-stream. In principle, this could be all data, especially if they are stored in a relational database. These data would include both the commands and the actual measured values. As an example, the demanded and actual azimuth of the transmitter beam may be stored, both of which may be incorporated into the data requested by the user. These data need to come from WP7.

- Provision of monitoring data from WP8 to WP7. These data are the engineering data which comes from the WP8 systems, but which should be made available to the entire EISCAT-3D monitoring system so as to provide a unified system. Examples of such monitoring data include system resources (CPU usage, storage capacity, storage use, data rates) and environment monitoring which are handled as part of the computing systems (computer room temperature). The data dictionary and transfer mechanism remain to be defined as part of the WP7/WP8 interface.

EISCAT-3D Data Storage, Access and Visualisation Baseline

Data Storage

File-based and/or relational storage philosophies will be used for data storage as appropriate. Both storage systems will provide secure access for users and automated, secure remote backup, be Grid-compatible and allow easy association between data and metadata.

The data rate is one of the key issues here, as it will affect the amount of data that can be handled, and the duration for which it can be stored (in the case of temporarily-stored data). However, even at the maximum average data input rate (40 Msamples/s/beam), there are identified hardware solutions that are currently available – albeit expensive – which will be able to handle this data rate.

The number of beams operated at a receiver-only site is likely to be of the order of 5 or more. With 16-bit samples, this results in 400 Mbytes/s per site of raw data. A terabyte store would thus be filled in approximately 40 minutes. However, further data decimation may be used to reduce this. The most useful method here may be to reduce the data spectrally, thus sampling only the frequencies around the ion and plasma lines. Any such spectral extraction would be performed by the signal processing aspect of EISCAT-3D (WP9). Additionally, time integration dramatically reduces the data throughput, but limits the flexibility of the data processing. Ideally the trade-off between the two extremes of maximising data preservation and minimising data throughput would be configurable at run-time. The allocation of storage, as a consumable resource, is an issue which will need to be addressed as part of the operational policy of the EISCAT-3D system. An attractive option, however, might be to allocate an independent data acquisition and storage system to each received beam, since this would minimise the risk to the entire array from the failure of such a unit, and offer useful redundancy in experiments where only a small number of beams was needed. Otherwise, data acquisition would be multiplexed onto a single storage device.

One example of a currently available data acquisition system which can handle data rates of the kind envisaged here is that used by the VLBI (astronomy/ geodesy) community [1]. Currently deployed on dozens of antenna-receiver sites across the world is the Mark 5A VLBI recorder. This is a disk-based storage system, based around a “StreamStor™” disk-interface card, manufactured by Conduant Corporation [2]. The StreamStor™ card can move data, in either direction, between any two of the three ports that it has:

- Hard-disk Array (two “packs” of 8 IDE hard disks each)
- PCI Bus (providing access to the host PC and onwards to the local network)
- Data Port (a 32-bit bus supporting the FPDP interface specification [4])

Thus, as an example, during “recording” mode, the StreamStor™ card directs data from the Data Port to the Hard-disk Array; during analysis, it directs it from the Hard-disk Array to the PCI Bus (and hence the CPU/Network); and during playback, it directs data from the Hard-disk back to the Data Port. The maximum data rate is 512 Mbps.

The Mark 5A occupies a 5U chassis and the two disk packs may be “hot-swapped”. Software allows the data recording to operating in “pendulum mode”, so that data acquisition is seamlessly transferred from one disk-pack to the other. After switching, the first unit may be removed thus, with operator attendance, the system can acquire data uninterrupted indefinitely.

The Mark 5B VLBI recorder is the logical extension of the Mark 5A, and makes use of improved data acquisition card technology, namely the StreamStor™–Amazon disk-interface card. This system is currently under development, but is scheduled for commercial availability mid-2006. The Mark 5B is a VSI-compliant system with capability to handle data rates up to 1024 Mbps; no external formatter (the extra hardware required for the Mark 5A) is necessary.

Data Access

The EISCAT-3D data system will encompass several different types of data, from sample-level beam-formed data through to science products containing estimates of derived plasma parameters. There will be several orders of magnitude difference between the data rate close to the antenna array and that at which summary science products based on derived plasma parameters are supplied to users. While some data will be freely available to all EISCAT users, access to certain data sets or data products will be regulated according to procedures similar to those presently applied by the EISCAT Scientific Association. For these reasons, different access systems will be required for different parts of the data flow.

At the user end of the chain, a web-accessible archive system will contain the final science products of the data reduction algorithms. Such an architecture could also be used for holding an archive version of the integrated correlated and analysed data set, together with the appropriate metadata. Preserving user access to data at the level of autocorrelation functions (or lag profile matrices) is important, since it will enable users to conduct their own data analyses, possibly with assumptions different to those employed in standard data reduction.

Nearer to the array end of the data flow, where the data rate will be much higher, a short-term data storage system will be needed for users requiring access to high time-resolution data, which are too large to be permanently archived (e.g. for applications such as interferometry). In this part of the data access system, the design philosophy will be based on the use of cyclic buffers, which will be over-written at intervals dictated by the data rate, but long enough for users to download data of interest. The whole data storage system will be designed to allow future expansion, including the potential for users to deploy their own data storage devices and/or processing tools at the radar sites to reduce the requirement for high-volume data transfer.

Science Products and Data formats

Because the data system will comprise various kinds of radar data and storage requirements, different data formats are appropriate to different stages of the system.

For those parts of the chain where data flow rates may be challenging, e.g. in the beam-formed or correlated data, bespoke formats optimised for I/O performance will be necessary. However at the user end of the data system, final products should conform to requirements for associating different data (e.g. correlated data, analysed data and system meta-data) software compatibility for data access and dissemination, and the need for self-describing extensible data formats.

The detailed specification of the science products from the EISCAT-3D radars will involve a future management decision, taking account of the resources available. However, in general terms, these data will comprise files of atmospheric parameters derived from the radar data. These data will be the outputs from specialised data reduction programs, whose input will be the correlated data, together with various a-priori models of the ionosphere and thermosphere. The number and type of these parameters will depend on the kind of experiment and the chosen analysis strategy. For ionospheric applications, they will include electron density, electron and ion temperature and plasma velocity. These basic parameters may be accompanied by others, either derived from a combination of measurements and models, or measured over a limited fraction of the range extent. Examples include ion-neutral collision frequency and ion composition. Other useful measurements, not generally available in current Incoherent Scatter Radar data, might include additional parameters of the neutral atmosphere, and an automated indicator of whether the target plasma is in thermal equilibrium. The data system must carefully distinguish between parameters derived from simple measurements, those calculated with the help of models, and any parameters whose derivation involves non-EISCAT measurements. In all cases, the associated metadata should be sufficient for the user to identify all the data sets and assumptions used in the derivation of these parameters.

An important requirement on the EISCAT-3D system, especially at the science end of the data chain, will be to make these data available in a format which facilitates their assimilation into software used by the scientific community. This includes combination with data from other related instruments (radars, optical imagers, riometers, etc.) to produce composite images or multi-diagnostic data products. The issue of data compatibility between data sets in Solar Terrestrial Physics will be very important in the timeframe of the EISCAT-3D radar, as new Grid and Virtual Observatory initiatives will provide environments enabling users to work with data from multiple instruments. The need to support data assimilation and combination does not necessarily constrain the choice of EISCAT data format, but imposes a requirement to interpret between the chosen format and the data containers that such software systems will use. The area of data assimilation and inter-operability is one in which the state-of-the-art is evolving quickly, hence it may be unwise to make a final decision about data formats too early in the project.

Data Visualisation

Precise specification of the data visualisation section is not appropriate at this stage of the design study. However, the corresponding section from the EISCAT 3D Design Specification Document [3] has been reviewed, and it has been confirmed that there are no changes to the strategy set out in our original proposal, other than the explicit clarification that data would be displayable in some form across the Internet for remote monitoring of experiments.

Conclusion

The principle aim at this stage in the design study is to outline the different types of data products that may be generated by the EISCAT-3D system. In addition to identifying and describing them, it is important to describe their position in the data processing sequence and to indicate the duration for which such data would be kept for reanalysis or archiving. This document has accomplished this exposition. It is apparent that clarification is required on a number of issues concerning the interfaces between WP8 and other work packages. These are as follows:

- Further clarification from WP4 and WP9 is required to establish the data rates, formats and transfer protocols used to be used in generating the beam-formed data.

- Discussion with WP5 is needed to agree the data types, data volumes and access protocols required for data and metadata in support of interferometry applications.

- We need to establish what control and monitoring data WP7 will make available, and in what form, in order to incorporate such parameters into the data-stream as metadata. In addition, WP7 needs to specify what monitoring parameters may be accepted from WP8 as part of the real-time control/monitoring system.

- Clarification is needed from WP9 on what processed data-streams will be provided to WP8. The level to which processing (spectral, decimation, integration, auto-correlation, etc.) has been performed needs to be indicated, as this will have a major impact on the data rates, and hence the storage duration, that the system will be designed to handle. The interaction between WP8 and WP9 is a major issue, whose resolution will be an important goal for the next phase of the design study.

- Owing to the potentially large data throughput, and the finite, albeit large, storage that can be provided, the issue of data retention and storage allocation is an issue that will impose constraints on the operation of the EISCAT-3D system, and will require policy decisions as part of the radar management plan.

References

[1] *"WP8 Report : High-speed data acquisition"*, D.J.McKay, 5-Jan-2006.
[2] Conduant Corporation, http://www.conduant.com/
[3] *"EISCAT 3D Design Specification document"*; FP6-2003-Infrastructures-4: EISCAT 3D Proposal #-11920 v1.0 01-Nov-2005.

— End of document —