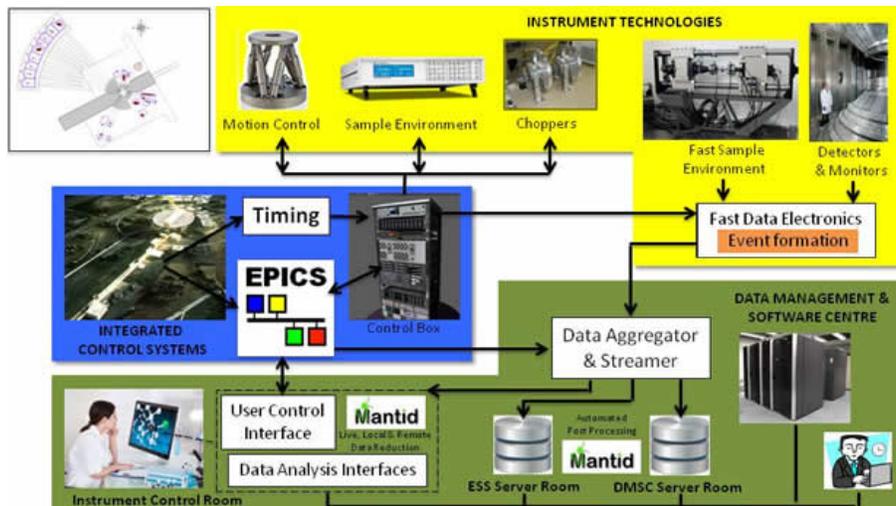


Pan-European ESS Data Management and Software Centre Takes Form in Copenhagen

SEPTEMBER 15, 2015

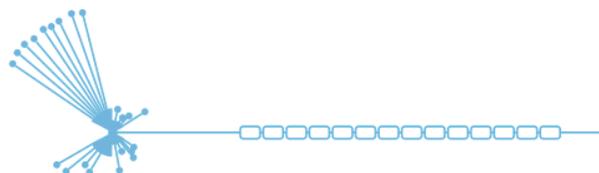
The Experiment Data. The ESS Data Management and Software Centre is an integrated part of the design and construction of the ESS instrument suite and a key driver in the development of the facility's user program.

COPENHAGEN and LUND—From the moment the first neutrons produced by the European Spallation Source (ESS) register their existence on a detector, the raw experiment data will flow from Lund, Sweden, to Copenhagen, Denmark, and then on to the ESS scientific user community throughout Europe. Constructing and operating this comprehensive data workflow is the core function of the ESS Data Management and Software Centre (DMSC), which is a vital part of the ESS scientific user program.



A schematic showing the domain of the ESS Data Management & Software Centre relative to Instruments, Instrument Technologies, and ICS. IMAGE: ESS

'DMSC is an integral part of the ESS Science Directorate. Our work is completely focused towards supporting the scientific research programme of ESS,' states Mark Hagen, the Centre's director, who previously worked at the ISIS neutron source in the UK and the Spallation Neutron Source (SNS) at Oak Ridge National Laboratory in the USA during their respective construction phases.



'In the near term—the next few years—we are very much focused on enabling the instruments to operate by developing the control software, the data reduction software, and the initial data analysis software," explains Hagan. 'In the longer term we are dedicated to ensuring that the data analysis and modelling software is available to extract the maximum amount of scientific information from the experimental data.'

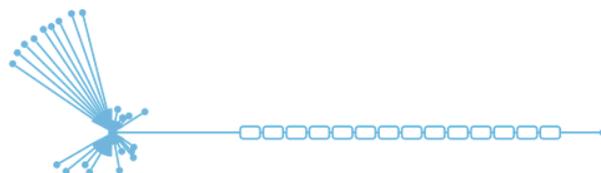
Processed in real time, the continuous stream of experiment data from the ESS instruments will be converted to scientifically meaningful information that can be displayed on a user interface, on-demand and worldwide. Through an iterative, software-assisted process of visualisation and analysis, researchers will be able to view the output of a neutron scattering experiment as it is taking place and begin to interpret results.



The DMSC staff in Copenhagen. From left, Mark Hagen, Sune Rastad Bahn, Thomas Holm Rod, Brian Lindgren Jensen, Simon Heybrock, Danielle Adonis, Jesper Rude Selknaes, Petra Aulin, Tobias Richter, Thomas Larsen, Jonathan Taylor, Torben Nielsen, Michael Wedel, and Céline Durniak. Absent from the photograph are Daniel Fulla Marsa and Peter Willendrup. The group is standing in front of an apple tree planted on the UCPH campus that has been grown from a sapling originating from the tree under which Newton supposedly first conceived his theory of universal gravitation. PHOTO: ESS

Network of European Collaboration

The success of the ESS scientific user program, which is expected to begin in 2023, will rest on a strong network of global research collaborations and in-kind partnerships diligently cultivated across the project for more than three decades. [The ESS In-Kind Contributions \(IKC\) model](#) is of a scale never attempted before in a research infrastructure project. In order to achieve its goals, DMSC will likewise take a community-driven approach to software development that fits neatly within the ESS IKC model.



To develop the data reduction and analysis software required by ESS, DMSC expects to draw upon the experience and resources of in-kind partners that include ISIS, the Paul Scherrer Institute (PSI) in Switzerland, and the Jülich Centre for Neutron Science (JCNS) in Germany. Other neutron facilities and university groups in Europe and globally will also participate, such as the USA's SNS.

Additionally DMSC is part of two recently established European Union programmes funded through Horizon 2020—[SINE2020](#) and [BrightnESS](#). Each will provide additional focus in a software area critical to ESS.

In SINE2020—world-class Science and Innovation with Neutrons in Europe for 2020, managed by Institut Laue Langevin (ILL)—DMSC will lead a Joint Research Activity (JRA) on Data Analysis and Software to develop a pan-European, modular, interoperable analysis suite. This work will be done in collaboration with ILL and Laboratoire Léon Brillouin (LLB) in France, JCNS/MLZ, ISIS, and the Swiss Spallation Source (SINQ) at PSI.

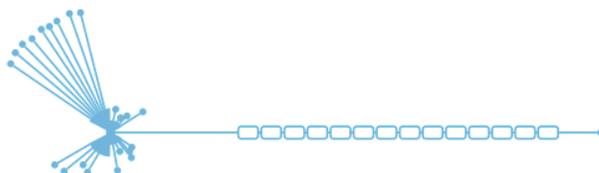
BrightnESS is a €20 million grant that will support ESS through a broad range of activities, including work packages in the area of detector event processing software. Detector event processing is a task traditionally done via electronics, but ESS's software-based processing will allow scientists to be more flexible and selective, and potentially more precise, in discriminating and locating neutron events on detectors. This project will add to DMSC's expertise that of scientists and software professionals from PSI, the University of Copenhagen (UCPH) and the synchrotron source Elettra in Italy.

'One of the things I really find great about ESS with DMSC is that it's an international collaboration,' says Sune Rastad Bahn, a UCPH and Technical University of Denmark (DTU) alumnus who joined DMSC in May 2015. 'I think we should always remind ourselves it's about all of the partner countries. It goes all the way down to the individual scientists who will visit us. They should all feel they're part of ESS.'

Ramping Up

The DMSC is located on the campus of UCPH and has itself emerged from an early partnership between ESS and the University. Originally known as the Interim Data Centre at UCPH's Niels Bohr Institute, the group began its work as part of the Danish pre-construction contribution to the ESS project. On January 1, 2014, the Interim Data Centre staff members officially joined ESS as DMSC, and were organized into four groups: Data Management, Instrument Data, Data Analysis & Modelling, and Data Systems & Technologies.

The scheduled ramp-up for activities within DMSC necessarily aligns closely with that of the engineering and construction of the ESS instrument suite and the development of the facility-wide [Integrated Control System \(ICS\)](#). As the data centre prepares for hot commissioning of the first ESS instruments, the next two years will see a rapid growth in collaborative software development involving many partners. Currently, an extensive



operation to gather requirements from neutron science users and facilities across Europe is underway.

DMSC currently has a staff of 15, which will soon grow quickly as in-kind partners begin their work. A decade from now, when ESS is fully operational with thousands of experiments performed annually, the staff is expected to peak between 60-70 employees working as part of the larger support structure of the ESS scientific user program.

'Historically in neutron scattering the instrument scientists had to do everything for an instrument,' explains Hagen. 'They were the people who had to fill the cryostat, tighten bolts on the instrument, and work on the data analysis. Nowadays, it's clear that the amount of work to do for a neutron scattering experiment has gotten larger and larger and it really takes a team to support the experiments. A team with a variety of different skills, and the people at DMSC will be part of that team.'

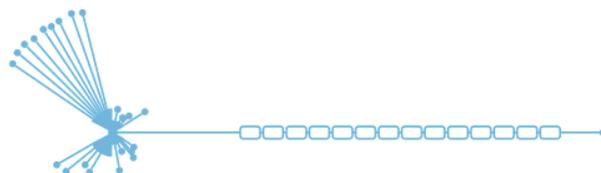


DMSC group leaders, clockwise from top left: Jonathan Taylor, Sune Rastad Bahn, Thomas Holm Rod, and Tobias Richter. PHOTOS: ESS

The Hardware

One of DMSC's first priorities is to establish reliable and easy access to experiment data. To this end, a system of software servers is currently being designed for placement in Lund and Copenhagen, with the focal computational power housed at the DMSC's building in Copenhagen.

'As soon as the infrastructure design is ready, we will do a lot of preliminary set-up and testing,' says Sune Rastad Bahn, group leader for the Data Systems & Technologies group that has responsibility for DMSC's hardware. 'As we get closer to the instruments being



commissioned, we will start to scale up. The infrastructure will be rock solid. Instrument users should rightfully have full trust their data is kept safe.”

A dedicated optical fibre will connect the two sites for data transfer across the Öresund Bridge and tunnel that span the waterway between Sweden and Denmark. The cable is scheduled to be in place by 2018.

The Instrument Interface

To begin an experiment on one of the neutron instruments in the ESS Instrument Hall in Lund, researchers will rely on the instrument’s experiment control interface. DMSC’s Instrument Data group and in-kind partners PSI and the Science & Technology Facilities Council (STFC) in the UK are developing the interface software.

The experiment control interface will be used to control a wide range of user-definable instrument settings. Depending on the instrument, these may include chopper configurations, adjustments to beam-defining components like collimation and aperture size, and sample environment parameters such as temperature and pressure. Naturally, this information is also critical to calculations that determine the significance of the raw data generated at the detectors, and is recorded as the experiment’s metadata.

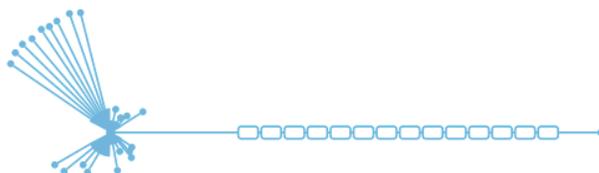
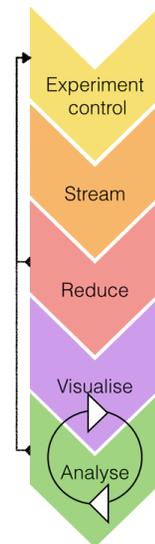
‘The control interface will be a modular software system,” explains Jonathan Taylor, leader of the Instrument Data group. ‘It will have the same look and feel across the entire instrument suite of ESS with specific customisations for each of the instruments, depending on the differing science areas served by the instrument and the different methods the instrument scientist wants to use.”

This expandable and customisable set of discrete software modules is being designed to form the basis of an adaptable standard, applicable throughout Europe. This software will direct mechanical adjustments to the instrument and sample environment through its interface with EPICS, the industry-standard Experimental Physics & Industrial Control System, which is implemented and managed by the ESS ICS Division.

‘Our work will be done with in-kind partners, and essentially we’re going to take an existing experiment control framework and develop on top of it,” explains Taylor. ‘The aim is then not to just have commonality across ESS but also commonality with our partner neutron labs in Europe.”

Stream, Reduce, Visualize, Repeat

Once the parameters for the experiment have been set and the instrument and sample environment adjusted, the experiment is run. Data collection begins when the neutrons scattered from a sample are recorded in the instrument’s detectors. For each of these



detected neutrons the positions and arrival times are recorded as event data organized in 'frames'. A frame is the time interval between neutron beam pulses, which at ESS will be 71 milliseconds. Each frame of neutron data, along with its metadata, is then transmitted to data-aggregating computing stations in Lund and Copenhagen.

From here, the goal is to achieve an optimal processing speed that allows ESS to stream the raw data, process it, and return meaningful and scientifically valid data back to its users in as close to real time as possible. The visual data will be made available, in Lund or elsewhere, as a continuously streamed channel, available on demand. Individual scientists monitoring the experiment will be able to subscribe and 'tune in' frame by frame.

'We plan to have a prototype data streaming system running by the end of 2016,' says Tobias Richter, leader of the Data Management group. 'It may not have the performance at that time for the full flux of ESS, and some of the processing may not be in place, but we will have a pipeline for continuous improvement to reach our targets. It is critical to transport the data correctly, and this really motivates us to do a good job and write reliable software.'

The data-aggregating software and systems managed by Richter's group will also write a backup copy of the data both in Lund and Copenhagen for the rare case that the stream is interrupted. The Data Management group, together with its European in-kind partners, will also supply data writing and data cataloguing services for ESS users.

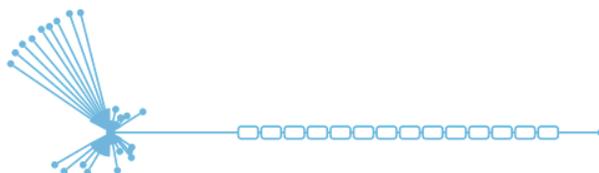
Once aggregated, the raw data is then reduced into scientific units—such as momentum transfer (Q) from time-of-flight data—using Mantid, an Open Source software framework originally developed at ISIS. This large-scale live data reduction—to be developed by Taylor's Instrument Data group in collaboration with STFC—will allow the data to be modelled into meaningful representations and returned to the instrument users as dynamic one-, two-, three- and four-dimensional visualisations of the particle interactions measured by the experiment.

From Reduced Data to Scientific Understanding

The process of converting the reduced data into scientific models that can be visualized and interpreted is an iterative one that typically requires human interaction with the data. Delivering the software to enable this is the role of the Data Analysis & Modelling group, headed by Thomas Holm Rod.

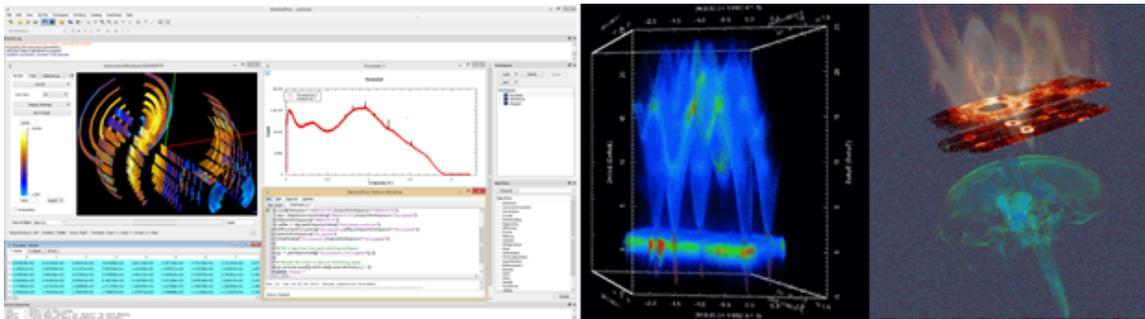
'It is crucial for us to get our in-kind contributions in place soon so that we can start on providing and developing the data analysis software that will be required to ensure the scientific success of the first ESS instruments when they come online,' says Holm Rod. 'These in-kind contributions are essential for us because they are a means to leverage the experience and knowledge of the European neutron scattering community.'

To develop and test their data analysis and modelling software, the group will use virtual



data generated by a neutron instrument simulation program called McStas.

‘Peter Willendrup, from DTU, is a lead developer of McStas and spends a third of his time in my group to assist us with using it,’ says Holm Rod. ‘Torben Nielsen in my group has already used McStas to test aspects of the data flow and analysis for [the instrument LOKI](#), which is a small-angle neutron scattering instrument and one of the first instruments to be endorsed for construction at ESS.’



Mantid data reduction and visualisation in action, using data from a simulated neutron scattering experiment. IMAGE: ESS

A Foundation for Future Development Across Europe

In addition to delivering the software for data analysis for the first days of instrument operations at ESS, it is also the goal of Holm Rod’s group to deliver a platform that can be used for future development of more sophisticated analysis tools. This platform will be designed to enable developers in the neutron scattering community to contribute in a controlled fashion.

For these same reasons, the Data Analysis & Modelling group will coordinate the SINE2020 JRA on Data Analysis and Software—in partnership with ILL, JCNS/MLZ, LLB, ISIS, and PSI—with the aim of creating the foundation for a common software suite at European neutron sources. This highlights another important aspect of IKCs, which is that partner facilities in many cases will be able to use the software that they contribute to ESS at their own facilities. This provides ESS with the opportunity to test the software on real data from real instruments during the ESS Construction Phase, when it will not be possible to test it on ESS instruments.

###

