Science & Technology Facilities Council

ISIS

Project BRIEF	Document Ref: Type the Reference number here
Project: PACE: Proper Analysis of Coherent Excitations	
Authors: Toby Perring (original-issue 3) Chris Marooney (issue 4)	Date: 10/11/17 Updated: 29/11/2020 (issue 4)
Customer Agreement:	Date: --/--/--
IMC Acceptance	Date: --/--/--

Version History

Issue No.	Issue date	Summary
1	10 Nov 2017	As submitted to IMC
2	25 June 2018	Project finances, staffing and timeline now refer to Project Management Plan. Change to Section 3.1 (Objectives) objective no.1: the scope of framework now includes full Python/C++ implementation as the eventual goal
3	02 March 2020	Explicitly states that DFT calculations are outside the scope of PACE. Notes that SpinW parallelisation is now a separately funded project.
4	29 Nov 2020	Minor reformatting. The separate SpinW project is now noted at first reference. The CASTEP interface code has now been named Euphonic.

PROJECT BRIEF

1 PURPOSE: SUMMARY OF THE PACE PROJECT

PACE (Proper Analysis of Coherent Excitations) is a software project to provide an integrated visualisation, simulation and fitting environment, on massively parallel and distributed computing, which interfaces to experimental data and materials modelling codes. PACE will:

- perform virtual experiments before and during the beam time to make most effective use of neutron beam;
- enable quantitative analysis of the full set of data collected during the experiment with proper account of instrument resolution;
- lower the barrier for users to analyse their data, so increasing the number and quality of publications, and reducing the time between experiment and publication.

The major bottleneck in the productivity of the Excitations Group instruments is easy-to-use software that can effectively analyse the vast datasets that they collect. Existing packages to visualise and manipulate the data struggle to handle the data, materials modelling codes are either not implemented for high performance computing (HPC) or their output is not integrated into the analysis framework, and proper account of instrument resolution is mostly not performed as it requires parallelisation of the current codes and HPC to be feasible in most cases.

In essence, the PACE deliverables are the parallel computing implementations and integration of proven data analysis packages and modelling codes in use on the Excitations Group instruments – the HORACE data analysis and visualisation framework, the TOBYFIT fitting and resolution function convolution package (both developed at ISIS), the parallelisation of the third party modelling package SpinW (spin waves), the parallel computation of scattering functions for single crystals and powders from the output from CASTEP (lattice vibrations), and the development of a generic interface to other materials modelling codes.

The project requires funding for 12 person-years' of effort, mostly for software engineering and computational scientist effort, but this also includes a 3 year fixed term neutron scattering post for someone with instrument modelling and scientific computing experience. The effort of several members of the Excitations Group will also be assigned to the project. PACE is a capital project as it delivers a software product. In the operations phase it will be run and maintained by existing Excitations Group instrument scientists, with up to 0.5 FTE of support from Scientific Computing Department (SCD). For full exploitation of PACE in the operations phase, a new long-term data scientist post is highly desirable, but the project is not dependent on this. PACE is distinct from Mantid, which only covers data reduction. PACE relies on the succesful implementation of DAaaS (Data Analysis as a Service), an independently funded project currently being implemented in SCD, which will provide the hardware capability and operation sufficient for the needs of PACE.

PACE can be viewed as an instrument upgrade for all four Excitations Group spectrometers, at a cost of approximately £1M, somewhat less than the cost of one of the guide upgrades to MAPS or MARI.

2 BACKGROUND

2.1 Background to the project

The latest generation of spectrometers at ISIS such as LET and MERLIN routinely measure complete maps of the four dimensional wave vector and energy dependent scattering from single crystals, accounting for

1/2 - 2/3 of the beam on those instruments, and a significant fraction on MAPS. These datasets, which can be as large as 0.5 terabytes (TB) each, collected at the rate of 1TB per day, contain the full information about strength, range and symmetry of the atomic scale interactions, and therefore provide the most sensitive discrimination between different theoretical models. In LET and MERLIN, ISIS has two world leading spectrometers, and in MAPS and MARI world competitive spectrometers that are undergoing major hardware upgrades that will increase their neutron flux by as much as an order of magnitude. However the development of analysis software methods and computing infrastructure has consistently trailed behind the developments in neutron instrumentation. In practice this results in major bottlenecks in the analysis of the data that limit the extent and accuracy of the information that is extracted. Furthermore, it limits the speed to publication, and in many cases analysis of an experiment remains in abeyance with no realistic prospect of it restarting. Currently, approximately 1/3 of experiments on the chopper spectrometers result in a publication (although it should be noted that some experiments simply fail, and sometimes more than one experiment is needed to make a publication), and of those that are published the mean time to publication is 2 – 2.5 years.

The bottlenecks can be divided into technical aspects and user interaction aspects. Taking the technical aspects first:

(1) The data visualization and manipulation package HORACE, developed at ISIS and in widespread use at other neutron facilities, has much of the right functionality to visualize and manipulate the data sets, as proven over many years of use and development on the instruments. However, most of its functions do not scale to be able to handle the latest data volumes (e.g. symmetrisation of data), and none take advantage of distributed computing hardware necessary for quantitative modelling, as described in points (2) and (3) below.

(2) The modelling codes needed to quantitatively analyse data are either not implemented for high performance computing (e.g. SpinW, the *de facto* leading linear spin wave modelling code), or their output is not integrated into analysis codes (e.g. CASTEP, for which only bespoke 'homebrew' codes exist to present the phonon calculations in graphical form, and which do not calculate the scattering cross-section in the momentum-energy coordinates required for direct quantitative comparison with data).

(3) Quantitative analysis of all data, whether single crystal or powder, requires proper account to be taken of the instrumental resolution function. The long-established TOBYFIT resolution function formalism (T.G.Perring) has recently been implemented as a prototype in HORACE, but does not work on high performance computing architecture.

The pressing need for the data analysis applications to work on massively parallel computing architectures can be demonstrated by an example of SpinW being used to model 1% of a typical total four-dimensional dataset taken on LET: this would a characteristic problem in the early stages of data analysis when key subsets of the full data set are being scrutinised. On a single CPU this takes ~1 hour to simulate; supposing there are 8 parameters in the model, fitted over 10 iterations, a fit will take 80 hours; with resolution convolution the time to fit could be as much as an order of magnitude greater. To make such analysis routine will require hundreds of CPUs to be available and the code recast for distributed computing architecture, to enable a turnaround in 10 minutes to 1 hour in this example. Similar CPU requirements are expected for simulations of the scattering from CASTEP output. The above illustrate why the full quantitative analysis that the data deserve is not currently being performed. Instead, the effects of resolution are generally ignored and only the simplest analysis, for example fitting Gaussians to get peak positions, is performed.

Major bottlenecks arising from user interaction aspects of the current software and hardware are as follows:

(1) Most users are simply overwhelmed by the volume of data that is collected. It takes much longer to analyse the data compared to, for example, a reactor based triple axis spectrometer (TAS) – anecdotally, from interviews of a few users, by a factor of 4 or 5. In consequence (i) principal investigators regularly prioritise TAS data analysis over ISIS data analysis, and (ii) students often finish their PhDs before analysis of an experiment is complete, resulting in orphaned projects.
(2) Notwithstanding the fact that the analysis, modelling and resolution convolution codes do not work on multiple CPU servers or on distributed computing, in their current forms they do not in any case work in an integrated fashion, further increasing the barrier for users to analyse their data.
(3) At some point every experiment requires bespoke analysis. Except for the most highly experienced and motivated users, the close assistance of expert instrument scientist support during analysis is essential.
(4) Reliable computer hardware operating on a 24/7 basis is vital for users to have confidence in a centrally provided analysis service. This has proven to be problematic within ISIS up to now.

PACE aims to reduce these bottlenecks. Its goal is to provide an integrated data visualization, manipulation, modelling and refinement environment, utilizing the infrastructure of the DAaaS system (Data Analysis as a Service) currently being implemented and tested by Scientific Computing Department (SCD), which will offer several hundred CPU in its own right and also seamless access to the thousands of CPUs of SCARF and other resources in the SCD Cloud.

To quantify the influence PACE would have on the spectrometer programmes, we attempted to divide the experiments performed from Summer 2013 to Summer 2016 on the four chopper spectrometers (MERLIN, LET, MAPS, MARI) into three classes, according to whether: (i) PACE would have made little difference to the outcome (the experiment was a failure, the outcome ultimately uninteresting, or theory is simply not well enough developed); (ii) PACE would have likely resulted in an improved or higher impact publication, faster publication, or a publication of an experiment that has been languishing, due to PACE lowering the barrier to analysis; (iii) PACE would have made analysis possible where currently it is not, for example modelling spin waves in a complex system with SpinW and/or because resolution considerations are important. This is necessarily a rather qualitative study, and the figures vary between instruments, but the figure is typically 50% – 60% of experiments on the Excitations Group instruments would benefit from PACE.

2.2 Relationship with other projects in ISIS

The major software development program within ISIS is Mantid. PACE is distinct from Mantid, which only covers data reduction. The Excitations Group spectrometers already use Mantid to perform all their data reduction and correction. Algorithms will be written so that PACE can read the Mantid multidimensional workspace classes (matrix workspace and MDworkspace), and vice versa. This will enable PACE to directly interact with the output of the data reduction and correction applications.

The set of visualisation tools for multidimensional data in Mantid (loosely referred to as VATES) has the potential to be a powerful collection. However, visualisation has not been highlighted as a blocker in the analysis presented above. Further, the core of Mantid is about to undergo a major re-engineering (Mantid Version 4) that is planned to last for some 18 months. The requirement that Mantid compile on all platforms, its multiple ownership by several facilities, and the fact that details have still to be decided mean that it would act as an unnecessary brake on PACE to be have a dependency on Mantid, when the key requirements for PACE (massively parallel implementation, resolution convolution, and seamless interaction with third party codes) are ones that are outside the scope of Mantid.

2.3 Relationship with other projects in STFC and ouside STFC

PACE relies on the succesful implementation of the DAaaS (Data Analysis as a Service) project in SCD. This will provide virtual computers to users on a cloud of multicore servers, ultra high-speed disk servers and SCARF. It is expected to have the hardware capability sufficient for the needs of PACE.

PACE will parallelise the SpinW application that performs linear spin wave theory for arbitrary magnetically ordered structures with arbitrary magnetic interactions (Heisenberg, Ising, Dzyaloshinskii-Moriya etc.) (https://www.psi.ch/spinw/). This is being maintained at Paul-Scherrer Institute (PSI), and a staff member at the ESS is about to employed to continue the maintenance and further development of SpinW. Part of this effort will be put towards parallelisation of SpinW. The PACE team and the developers of SpinW have a good working relationship (and have together already started to explore the options for parallelization), and there is presently an informal agreement to work with ESS and PSI. The current HORACE and SpinW codes already have a prototype interface.

The resolution function sub-project within PACE will rely on McStas (http://www.mcstas.org/) to perform Monte Carlo simulations of the instruments up to the sample position, to automatically create lookup tables of neutron wave vector and time distributions on an experiment-by-experiment basis. This code is well-established and widely used at numerous neutron scattering facilities, and there is considerable experience of the use of McStas both within the Excitations Group and ISIS in general. Reliance on McStas constitutes a low risk.

2.4 Upgrade in the context of other instruments at ISIS

PACE will provide an integrated data visualization, manipulation, modelling and refinement environment driven by the needs of the Excitations Group. However, it will be designed in such a way that the key software components will accessed by a well-defined an API (Application Programming Interface) which will make it straightforward to, for example, replace the resolution function model with one for indirect geometry instruments, or to interface to other third part modelling codes. As an example, single crystal magnetism experiments are regularly performed on OSIRIS; simply replacing the resolution function model will make available all the modelling and fitting capabilities of PACE to these experiments. The analysis framework will make no assumption of direct, indirect geometry, and will accept elastic or inelastic data.

The Mushroom is a potential future inelastic instrument project that with projected count rates 50 times greater than LET. Without the software of PACE and the DAaaS hardware infrastructure, analysis of data taken on Mushroom will simply be impossible.

2.5 Upgrade in the context of developments at other facilities

At the SNS the simulation package McVine (http://www.mcvine.org/) is a ray tracing program that models instruments with sample scattering kernels. Like McStas, it is geared to simulating the outcome of experiments, and with a sufficiently detailed model of the instrument enables simulation to be performed that fully account for the effects of resolution and for multiple scattering. The approach is very different to that of PACE, however, which is geared to the realities of data analysis that involves the highly iterative process of examining key data sections, with empirical modelling of backgrounds, testing of models, parameter fitting of putative models on those data sections, in an ever increasing volume on each iteration, finally building up to full refinement of model parameters. A key element of PACE in the second half of the project is the development of analysis strategies and capturing these workflows, culminating in the building a GUI based 'workbench' that is both flexible and accessible to inexperienced or irregular users.

3 PROJECT DEFINITION

3.1 Objectives of the Implementation Phase:

PACE is a software project with several interacting but distinct components that need to be created. The key deliverables are as follows:

1) Optimisation of the current HORACE data analysis and visualisation framework by
    - Parallelisation for multiple core computers and distributed computing, operating on the Data Analysis as a Service (DAaaS) and SCARF high performance computing services at STFC;
    - Provision of file-backed operations so that datasets larger than the available memory can be processed.

2) Optimisation and extension of the current TOBYFIT resolution convolution and least-squares fitting application by
    - Production of a parallelised version of for parallel computing, operating on DAaaS and SCARF, using the framework developed above in (1).
    - Extension of the formalism to use Monte Carlo simulations of the primary spectrometer from McStas for greater accuracy of the resolution model for neutron guides.

3) Parallelisation of SpinW (spin wave modelling and scattering cross-section computation).
    - Formally this activity is now a separately funded project; however its results are integral to the use of PACE with multiple modelling codes.

4) Computation of scattering function for single crystals and powders from *ab initio* determination of phonons using CASTEP and from other codes that output force constant matrices.
    - This utility is now named Euphonic.

5) Development of an Application Programming Interface (API) to third party modelling codes so that they can be interfaced to the HORACE framework and the TOBYFIT resolution convolution application. Integration of SpinW and CASTEP into the framework using the API.

6) Construction of a GUI based 'workbench' for managing analysis of data with refinement of parameters in resolution broadened models for scattering i.e. GUI interface to TOBYFIT.

7) Mantid based manipulation and GUI based visualisation of powder data that accesses the modelling and resolution convolution/fitting capabilities of PACE.

8) PACE will be handed over for routine operation on the instruments and first formal release for users at end of the project.

In more detail:

(1) Data analysis framework. The backbone of PACE is a data analysis framework that is essentially a distributed computing implementation of the proven Matlab-based HORACE application. The minimal implementation is for a Matlab interface with key parts written as parallelised C++ to use multicore and distributed computing. Initially it will be run from the command line, as the present HORACE. In this way users will be able to immediately use the capabilities of PACE, as the current HORACE functionality will be carried across. From a developers' point of view the project risk is minimized as the literally hundreds of bespoke functions in HORACE, ranging from the graphics to highly specialised algorithms (for example, to efficiently sample points of absorption in cylindrical ^{3}He detectors), will be used essentially unchanged. The framework will read and write the standard Mantid data workspaces (matrix and MDworkspace) so that PACE and Mantid can be used in conjunction.

During the initial analysis phase the work required to have a Python interface and completely Matlab independent core will be assessed in detail. Both a Matlab interface and a Python interface will be provided. This is the ideal scenario as it is removes restrictions for use at some other facilities and scalability to unlimited computing hardware resources, including for example, commercial clouds for peak demands that exceed the resources of DAaaS.

As of March 2020, piloting studies of Matlab/Python interoperability has led to the decision to provide both a Matlab and Python interface, with the core being compiled Matlab (which does not require a Matlab licence) and just a thin Python wrapper around the Matlab objects.

(2) Fitting and resolution function convolution application. This will be a generalization of the prototype translation of TOBYFIT that has already been performed for use in the current HORACE. The prototype has demonstrated the feasibility, but a number of developments are required: (i) a hybrid McStas/ TOBYFIT generalisation is needed to create lookup tables of neutron energy/time/divergence at the sample position for sufficiently accurate modelling of the resolution function now that the Excitations Group instruments all have neutron guides; (ii) it needs to be recast for distributed computing and redesigned for optimal speed; (iii) generalisation of how instrument information is held is required so that resolution function models for other instruments (e.g. IRIS and OSIRIS) can be readily added to the formalism; (iii) more sophisticated algorithms for arbitrary sample mosaic spread, and sampling instrument components need to be implemented. (iv) An API will be developed to interface with models for the scattering function to enable simple user-defined functions or third party codes to be interfaced to the framework and resolution convolution application.

(3) Linear spin wave modelling package. SpinW (https://www.psi.ch/spinw/) will be parallelized and an API (Application Programming Interface) written so that it can be called from the framework or from within the convolution and fitting application. SpinW is written entirely in Matlab. The major numerically intensive computations need to be rewritten in C++ for distributed computing. The user interface needs to remain unchanged – it is the power of the interface that can read Crystallographic Information Files (CIF files), knows of the full space group symmetries and simplifies as much as possible the setting up of magnetic structures and interactions that has made SpinW the *de facto* leading spin wave modelling application. There is the opportunity to collaborate closely with ESS and PSI which has its own resources that will contribute this sub-project.

As of January 2020, SpinW and McPhase parallelisation have been funded as a separate project by the Ada Lovelace Centre at STFC. It will be run a separate project to PACE, in collaboration with the ESS, but an efficient API to PACE is guaranteed, with the management of the SpinW and PACE projects working together.

(4) Phonon cross-section calculation. A distributed computing application will be written to compute phonon cross-sections from the output of CASTEP that contains the force constant matrices. It will be written with an API that will expose it to the rest of the framework in exactly the same way as the one for SpinW. This application will require significant investigation with the collaboration of a CASTEP developer to write an efficient algorithm suitable for the size and non-Cartesian grid of momentum-energy points that are collected on the instruments.

(5) API development as part of (4) and (5). It is important to note that the API to the resolution convolution application and third party modelling codes will be written to make it easy to interface to other applications, for example DMRG (Density Matrix Renormalisation Group) codes, Quantum Monte Carlo (QMC) simulations, codes to compute excitation spectra *ab initio* etc. The reasons the two particular modelling codes SpinW and CASTEP, one for spin waves and one for phonon calculations, have been chosen for the construction phase of PACE are that (i) they cover a large fraction of the needs of the experimental programmes on the instruments, and (ii) the definition of an API that works for two very different models will force consideration of generalities so that the API will be sufficiently flexible that other third party modelling codes can be interfaced to PACE.

(6) GUI workbench. After a several month period of testing on the instruments, feedback from users, and most importantly the new data analysis strategies that will have been developed with experience of items (1) to (4) working together, a GUI will be written that will manage the iterative and exploratory process of gaining a dual understanding of the information content of the data and validity of different models of for example the spin waves in a sample.

(7) Powder data visualisation and manipulation. The experimental programme on MARI is exclusively with polycrystalline materials, and there is a smaller powder programme on LET, MERLIN and MAPS. A development of Mantid currently in progress for the Excitations Group is the implementation of the very popular Matlab-based Mslice application to visualize and to manipulate such data. However, proper analysis of phonon densities of states and powder averaged spin waves for polycrystalline magnetic materials requires the SpinW and CASTEP modelling projects, with proper account of resolution, to be callable from Mantid. PACE will be designed to communicate with Mantid. However, seamless user operations with Mantid will require additional developments and coordination with the Mantid team.

(8) End of construction. The first five deliverables (data analysis framework, fitting and resolution convolution, spin wave calculation, phonon cross-section calculation, API to 3rd party modelling codes) will be available in basic form at the end of the build phase. Iterative development and testing in use by instrument scientists and friendly users will take place from this point to the end of the project.

Note: It is not in the remit of PACE to perform DFT calculations. It is assumed that users will use their own favourite codes to output force constants matricies. PACE will support APIs to selected third party modelling programs, and a generic API for users of PACE to use with the modelling codes of their choice.

3.2 Resources required for the Implementation Phase:

Please see Project Management Plan.

3.3 Timeline (incl. key milestones) for the Implementation Phase:

Please see Project Management Plan.

3.4 Long term operations

Operation and maintenance after construction

Excitations Group instrument scientists will support the use of PACE by users of the instruments as part of their local contact duties. Maintenance of the software, additional functionality in response to the needs of particular experiments, and smaller scale developments will be performed by existing staff effort within the Group: Alex Buts (scientific software expertise, 0.5FTE) with help from scientists with strong computing skills (Duc Le and Toby Perring, 0.2FTE each).

The software will be maintained on a Jenkins server to automate the building and testing of PACE for the supported computing platforms. SCD already has experience of the use of such servers.

To maintain the software on the high performance computing platforms will require some resources within SCD (parallel computing experience for framework maintenance and development, interfacing with new developments in CASTEP, SpinW). At this point it is difficult to estimate this level: up to 0.5FTE is likely to be needed. It is proposed that in the final six months of construction a more accurate assessment is made.

Ongoing support of the DAaaS service and STFC computing hardware on which PACE will run is assumed.

Future major developments

Once the operation phase has been entered, large developments that need effort beyond that described above (examples could include interoperability with McVine, or full integration with Mantid depending on how that project develops etc.) will be proposed as new projects in their own right, with appropriate business cases made. A road map of major long term developments of this kind will be formulated as part of a closing three month review.

Ideal long term scientific support scenario: 'instrument scientist for data analysis'

PACE will make a major difference to the quality, quantity and speed of analysis across the Excitations Group instruments with the support scenario described above.

The most effective use of PACE would be with the addition of a dedicated data scientist (1) to spend time with users after experiments to outline data analysis strategies and create custom functionality for nonroutine experiments, and (2) to provide expertise in the third party modelling codes to support users and the other instrument scientists. Once the construction of PACE is completed and it enters true operation, long term support in the form of an 'instrument scientist for data analysis' will be sought. However, the success of PACE is not dependent on such a dedicated data scientist. It should be noted that other facilities have (SNS) or are planning for (ESS) such staff. The ESS is planning for 0.5FTE per instrument.

3.5 Project management

Please see Project Management Plan