Success Story

National Service for Computational Chemistry Software

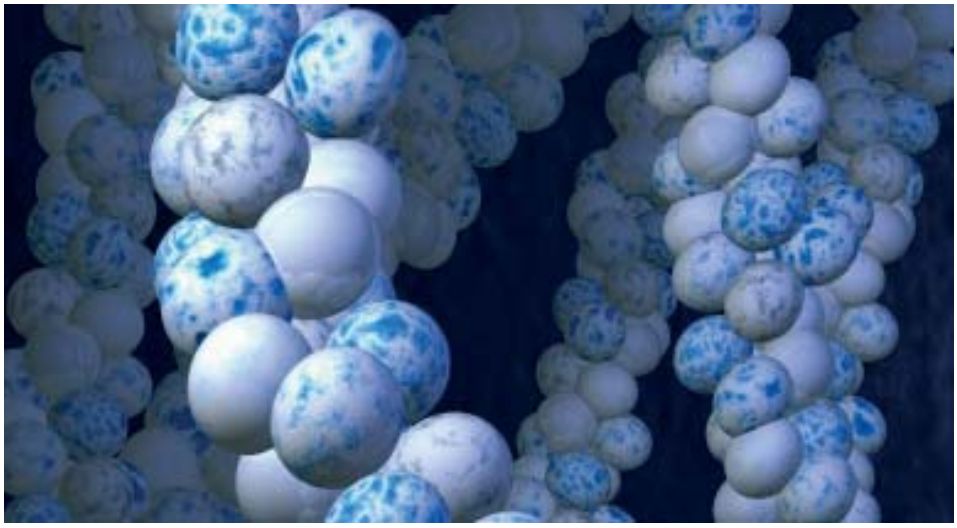

Delivering a Complete Service for Computational Chemists

"To support our users, we are currently providing access to more than 20 different software packages. In computational terms, these have very different demands, so we really have to handle an extremely diverse mix of application requirements."

– Sarah Wilsey, National Service for Computational Chemistry Software

The National Service for Computational Chemistry Software (NSCCS) provides access to software packages, computing resources, specialist scientific consultation and software training for UK academics working across all fields of chemistry. The service is funded by the Engineering and Physical Sciences Research Council (EPSRC), and any researcher eligible to apply for an EPSRC grant can apply to use its resources, which are provided free of charge to the user.

The NSCCS started life in 1997 as the UK Computational Chemistry Facility (UKCCF), whose users were predominantly quantum chemists accessing the facility for its hardware resources. In 2001 it was reclassified as a national service, and today, while around two-thirds of NSCCS users still run quantum chemistry codes, the service is expanding its remit to include biological and materials chemists.

A wide range of projects are carried out on the service's machines. The physical chemistry projects tend to involve studying systems with only a few atoms at very high levels of theory to obtain extremely accurate energetic and spectroscopic properties; the inorganic and organic chemistry projects normally deal with the structure and reactivity of medium-sized systems containing tens to hundreds of atoms; the biological chemistry projects typically require molecular dynamics simulations on tens to hundreds of thousands of atoms to study the structural changes that occur in biological systems on a ps-ns timescale; while periodic methods such as plane-wave density functional theory are used by the materials chemists to calculate the structure and properties of infinite crystalline systems.

"To support our users, we are currently providing access to more than 20 different software packages," explains Sarah Wilsey, who manages the NSCCS. "For historical reasons, around half of these are quantum chemistry packages such as Gaussian®, ADF® and MOLPRO®, but we also provide access to biological chemistry packages

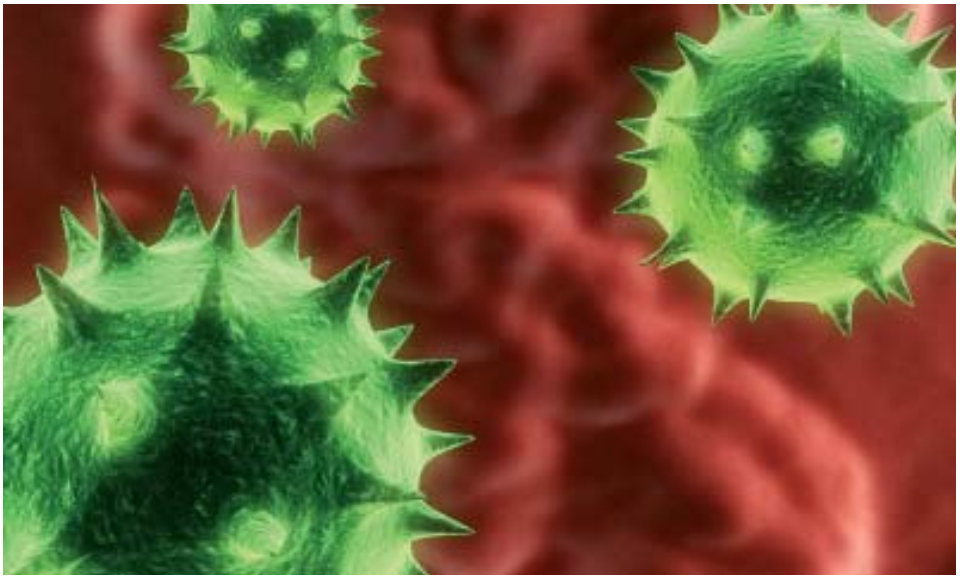

"SGI® Altix® gave us the best value for money, the most shared memory, and SGI were also able to provide the support we needed to help us compile, test and optimize our packages."

– Sarah Wilsey, National Service for Computational Chemistry Software

such as AMBER®, GROMACS® and NAMD®; and materials chemistry packages such as CRYSTAL®, GULP® and CASTEP®.

"In computational terms, these packages have very different demands. The quantum chemistry calculations often need large amounts of memory in addition to CPU time. When run in parallel, these codes normally only scale well on up to 16 processors, but the scaling is significantly enhanced if the memory is shared across all the processors. Molecular dynamics simulations, on the other hand, do not need much memory, but are very CPU-intensive. These codes are often specifically designed to run efficiently in parallel over 100's of processors in order to get a reasonable turnaround time on jobs."

One Machine or Two?

When the UKCCF was established, its computing resource was based on a DEC Alpha® machine, to which an Opteron® cluster was added in 2004. Although the DEC Alpha provided a stable platform for the service, it was reaching the end of its usable life; and the cluster had limited shared memory capabilities.

"When we were considering how best to upgrade our computing resources, we felt it was very important to have a large shared memory machine, but we also wanted to be able to support a lot of CPU cycles," continues Sarah Wilsey. "When we went out to tender for the project, from all the available options, the SGI® Altix® seemed by far the best way to achieve this. It gave us the best value for money, the most shared memory, and SGI were also able to provide the support we needed to help us compile, test and optimize our packages, which is invaluable in ensuring we're getting the best performance out of our software.

"We considered buying two machines – one machine to give us the shared memory we need, and an additional standard Linux® cluster to give us mainly CPU cycles. However, there are significant advantages to having just one machine in terms of support, storage and ease of use."

The NSCCS is now based on a 224 core SGI Altix 4700 server with 896GB memory and 15TB disk storage. The service is managed from Imperial College, with hardware support provided by the Rutherford Appleton Laboratory. At any one time the NSCCS has around 100 active users, who between them are involved in running 40-50 projects. Typical examples of these users are Carmen Domene, a researcher at Oxford University who has been using the service for the last two years; and Tim Wright, a Reader in Physical Chemistry at the University of Nottingham, who has been a user ever since the service was established.

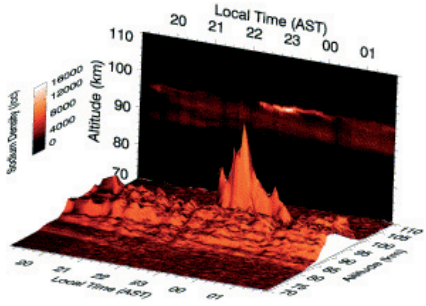

A sudden sodium layer observed during the Coqui II campaign at the Arecibo Observatory in Puerto Rico, February 20 1998.

[S. C. Collins et al. Journal of Atmospheric and Solar-Terrestrial Physics 64 (2002) 845-860].

K channel in a lipid bilayer

Meeting the Needs of Problem-Oriented Researchers

The types of project that Carmen carries out involve using classical and quantum simulations of biological systems to study gating and permeation in ion channels – changes in which are key to diseases such as diabetes, Parkinson's and Alzheimer's. "The transport of ions and water across membrane cells and organelles is a prerequisite for many of life's processes. Understanding how membrane proteins function so that drugs or antibiotics, ions or water molecules can travel in and out of cells is therefore vital," she explains.

"My research is aimed at understanding the dynamics of these proteins – for example, how they open and close, i.e. the gating process which is essential for their biological function. Gating is not yet fully understood, but is the reason why these molecules can play major roles in a large and diverse number of physiological processes, and why there is also immense interest in ion channels as targets for pharmaceutical drugs.

"Without the resources provided by the NSCCS, I simply wouldn't be able to afford to run the calculations I need. And, because the service's new computer is so much more powerful, I'll be able to study the integration between different levels of complex biomolecular systems, as well as running longer simulations and tackling problems where bigger systems are involved."

Tim Wright's work, meanwhile, involves producing highly accurate potential energy surfaces and highly accurate models of the interactions between small molecules. One example is the study of pre-reactive complexes (i.e. the first interactions between weakly-bound molecules before they react together), where he has been looking at how a nitric oxide molecule interacts with small molecules such as carbon monoxide and methane. A very high level of theory is needed to understand the one-to-one interactions involved and to gain insights into their behavior.

Another area that he and his team have been involved in has been modelling the chemistry of the upper atmosphere, in collaboration with Professor John Plane from the University of Leeds. They have

Software packages provided by the NSCCS include the following:

QUANTUM CHEMISTRY	BIOLOGICAL CHEMISTRY	MATERIALS CHEMISTRY
ACES2	AMBER®	CASTEP®
ADF®	AUTODOCK	CRYSTAL®
Dalton	GROMACS®	DL_POLY
GAMESS-US	NAMD	GULP®
GAMESS-UK		SIESTA
Gaussian®		
Jaguar™		
MOLPRO®		
MOLCAS		
NWChem		
ORCA		
Q-Chem		
TURBOMOLE		

The NSCCS also supports a range of visualization packages. A complete listing of the packages provided by the service can be found at: www.nsccs.ac.uk/software_full.php

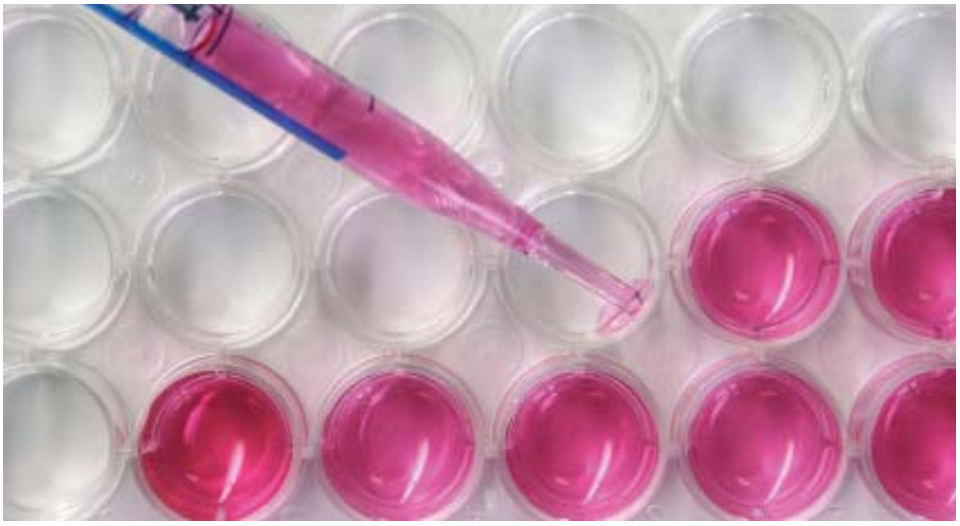

"Without the resources provided by the NSCCS, I simply wouldn't be able to afford to run the calculations I need."

– Carmen Domene, Oxford University

been looking at the chemistry of sodium and potassium, which can suddenly form clouds. The team set up a theoretical model of the ionic chemistry of charged sodium and potassium clusters, plugged the results into an atmospheric model, and compared their predictions with observations taken with a ground-based telescope in Puerto Rico. They found a very good correspondence between the two – demonstrating the power of the model for understanding systems that cannot easily be probed by any other method.

"Projects such as studying the interaction between weakly-bound species really are pushing the most sophisticated quantum chemistry methods to the limits," says Tim. "It would be impossible for every researcher in the UK who uses the NSCCS to afford a similar size of computer for their own department, so the fact that we can access it remotely, and know that it's well supported with all the most up-to-date software packages really is a benefit to our research."

"The crucial thing about the NSCCS is that it's targeted at problem-oriented researchers, and provides a complete service for 'non-specialist' users," adds Professor Mike Robb from Imperial College. "There are many people who develop software themselves and do all their calculations on their own hardware. Most of our users are not like that. They are people who are working in the lab, doing experiments, and we are giving them access to the computational chemistry support they need for those experiments."

"This is one of the reasons why our training service has become increasingly important," concludes Sarah Wilsey. "About two-thirds of our users are doing experiments, and close to a third are primarily experimentalists who started using our service because of the support and training we are able to provide. This can be one-to-one training in the use of a specific software package; more general software workshops for up to 40 people that we run periodically throughout the year; or our specialist scientific consultation service, through which people can ask us for advice on what sort of calculations to run.

"As a measure of the value of the NSCCS to the UK computational chemistry community, around 50 scientific papers are published each year by the service's users – many of which might otherwise never see the light of day – so the service really is making a significant contribution to computational chemistry research."

More information on the NSCCS is available at www.nsccs.ac.uk

Corporate Office
1140 E. Arques Avenue
Sunnyvale, CA 94085
(650) 960-1980
www.sgi.com

North America +1 800.800.7441
Latin America +55 11.5185.2860
Europe +44 118.912.7500
Japan +81 3.5488.1811
Asia Pacific +1 650.933.3000

3989 [01/07]

J15235