NSC News
National Supercomputer Centre at Linköpings universitet
2011:2

Chemical weather forecasting

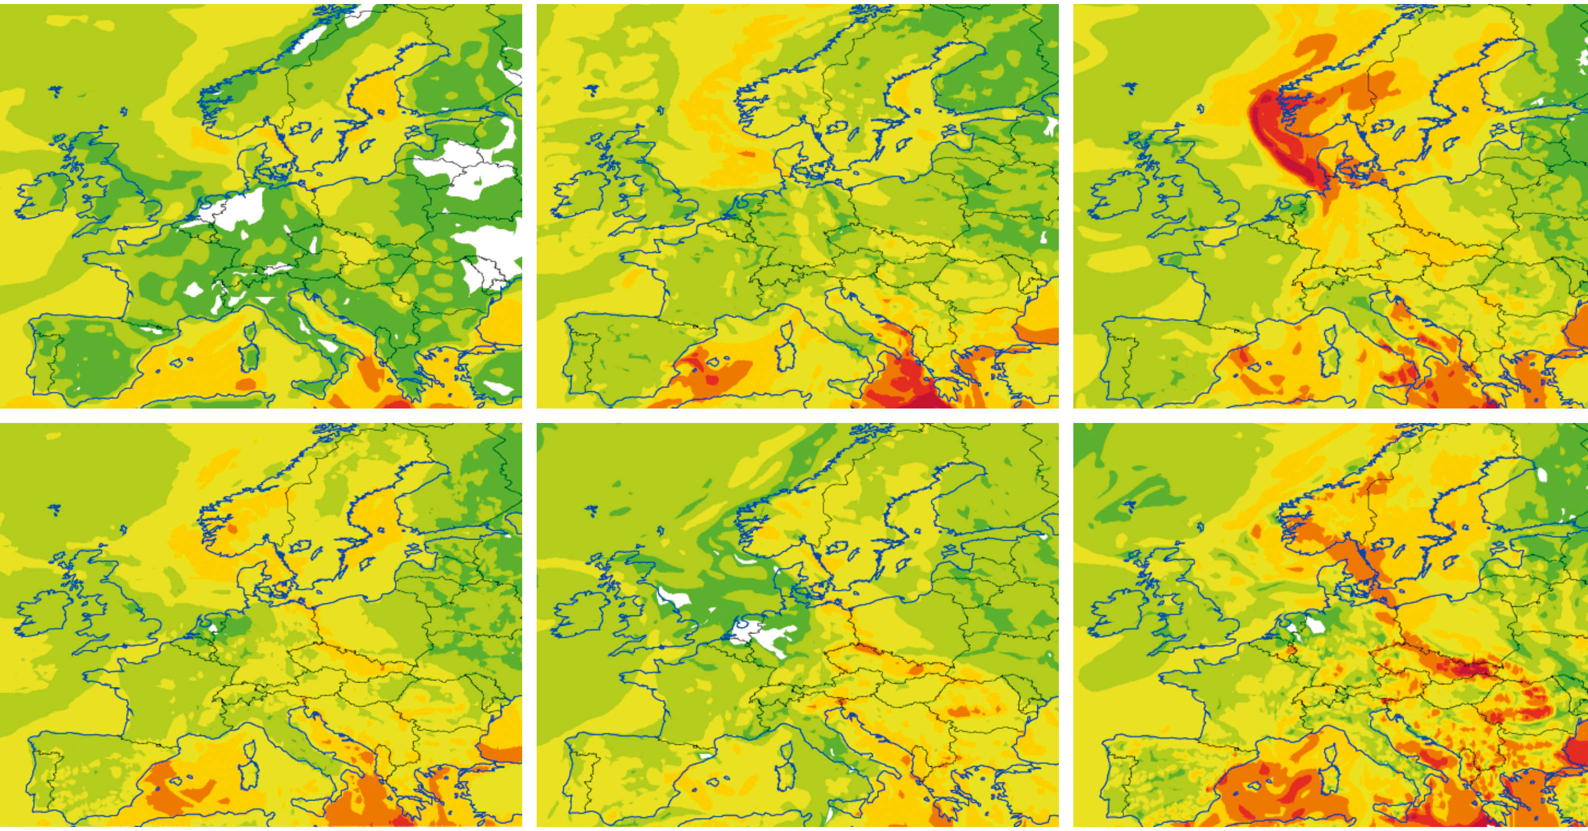

Chemical weather forecasting

How clean is the air that you breathe? Lennart Robertson at SMHI knows how to predict the chemical composition of the atmosphere.

Reed further on page 3

NSC Express

Let NSC's self-service portal NSC Express boost your experience as an user. It facilitates administrative tasks and improves your overview.

Read further on page 6

More HPC needed for Swedish research

Utilisation of high-performance computing is now seen in increasingly more disciplines. From being mainly of interest to physics, chemistry and astronomy, HPC is now present in almost every scientific branch. This is a combined effect of the increased availability of HPC resources and the requirements from researchers to analyse large data sets and use complicated calculations.

As one example, we have recent years seen an enormous increase of HPC in life sciences, as this area is getting more and more data-driven. Currently, the total amount of available data in life sciences doubles every 10 months. Since the development in data storage efficiency only doubles every 18 months, substantial investments are needed in the future to keep up with this increase.

Furthermore, the requirements for computational capacity also increase, since more and more scientific problems are analysed utilising large-scale calculations and simulations. In Sweden, we have seen a close to exponential increase of SNIC allocations, with close to a doubling of allocated core hours every year. Nevertheless, the requests still increase, and for the recent allocation round, all SNIC clusters were on average oversubscribed by a factor of 2. The real needs are much bigger, but researchers try to keep their applications on 'reasonable' levels on par with what is available. The systems at NSC are the most heavily oversubscribed, which we are proud of since it reflects that users have large confidence with the NSC systems. Our capacity system Kappa would need to increase 4-fold in order to satisfy our users.

Thus, in order to meet the increased requirements by the researchers, a substantial increase of SNIC's computational capacity is needed within the next year(s). Extrapolating the current trend from 2007 gives a total need of about 100 million core hours per month at the end of 2012. Currently planned SNIC resources will only provide 60 million core hours at that time, so there is a need for additional 40 million core hours. Furthermore, Sweden has hitherto aimed at having a total capacity corresponding to 10% of the international top system, which has shown to reasonably well correspond to the needs. By autumn 2012, the world's top system is expected to have a peak performance of close to 20 Pflop/s, indicating that the total Swedish capacity then should be close to 2 Pflop/s. Both these ways to predict the future needs give the conclusion that SNIC needs to invest in a new system of about 1 Pflop/s to be up and running before the end of next year.

BENGT PERSSON, NSC DIRECTOR

1 TiB nodes in Kappa

NSC has acquired and installed two nodes with an impressive 1 TiB memory and 64 cores each in one of our main systems, Kappa. These two new nodes will serve as replacement for NSC's previous large shared memory system Mozart. Compared to Mozart each of the nodes has twice the memory and the same number of cores. The nodes were financed through local Linköping University funds and priority will thus be given to users from Linköping University. To get access use NSC Express to apply for a LiU local project with allocation on the fat nodes.

The nodes are HP ProLiant DL980 G7 Servers which contain eight Xeon E7 CPU:s clocked at 2.1333 GHz. Each CPU has 8 cores and the 64 cores in a node give it a peak performance of 546 Gflop/s. The memory architecture is NUMA with an aggregated memory bandwidth of about 130 GiB/s. 6 TiB of local disk is available to users on each node. The nodes are connected to central storage through 10 Gigabit Ethernet and to the other Kappa nodes through an Infiniband interconnect.

PETER MÜNGER

New staff member

I joined NSC in May this year as your application expert in materials science. Originally from Dalarna, I moved to Uppsala for studies and defended my PhD in physics there in 2009, focusing on ab initio simulations of materials for hydrogen storage, batteries, and carbon nanotube catalysis. During and after my PhD, I worked as system administrator, maintaining HPC clusters and local computing services at the department. At NSC, I will work with installing, validating and benchmarking ab initio software such as VASP, Abinit, and Espresso.

PETER LARSSON

Chemical weather forecasting: a dema

The chemical composition of the atmosphere is the aim of a new daily forecasting service. The access to high performance computing plays an important role for the European project where this service is under development.

We have all our own experience of the every day weather. We may though not think that much about the content of the air we breathe or what the wind that may feel refreshing is bringing. The chemical composition of the atmosphere and its variations has, however, a profound role for various health aspects as well as for the climate. Chemical weather forecasting is the aim of the EU FP7 project MACC – Monitoring Atmospheric Composition and Climate – where the every day chemical composition of the atmosphere is in focus (in contrast to more long term climate impact studies). The objectives are to provide a set of services divided into clusters like: global green-house gases, global reactive gases, global aerosols and European air quality forecasts. The services are developed within the European Collaboration GMES (Global Monitoring for Environment and Security) as so called core services, and are to be used by private or public consultants, environmental authorities, or as input to other research activities.

The key elements of the services are:

- Global service lines providing:
  - monitoring of climate, climate forcing and the sources and sinks of key species;
  - monitoring of stratospheric ozone;
  - forecasts of reactive gases and aerosols;
  - boundary conditions for regional models.
- European service lines providing:
  - air quality forecasts from high-resolution regional systems;
  - air quality assessments based on retrospective running of the regional systems using validated observational data;
  - UV radiation assessments and forecasts;
  - solar-energy resource assessments and forecasts.

MACC takes as its input comprehensive sets of satellite data from many tens of instruments supplying information on atmospheric dynamics, thermodynamics and composition, made available by space agencies and institutions with which the agencies collaborate to produce retrieved data products. The satellite data are supplemented by in-situ data from meteorological networks and a limited amount of data from networks providing in-situ measurements of atmospheric composition. Data are processed to provide a range of products related to climate forcing, air quality, stratospheric ozone, UV radiation at the earth's surface and potential for solar power generation (Figure 1).

The project has 45 partners where the European forecast centre, ECMWF, has a key role as coordinator and data provider for all the different parts of the project.

SMHI is participating in the regional cluster that provides chemical weather forecasting and data assimilation for the European scale. This cluster involves some other 5 parties which products are combined to ensemble products that is a composite product of all individual contributions assumed to be of added value then each product separately (Figure 2).

SMHIs part is accomplished by utilising the chemical transport system MATCH including data assimilation of chemical observations. MATCH is the general model tool that we use for various chemical – nuclide atmospheric transport and transformation studies (Robertson et al. 1996, Langner et al., 1998). The model is Eulerian and is for this case spanning a volume covering Europe and up to the tropopause (~10 km) divided into 2.6 million

Fig

Fig

New staff member

Hello, I am a new application expert working 50% at NSC within the MATTER consortium, the other half I'm doing research at IFM, Linköping University. My background is condensed matter physics, more specifically theoretical spectroscopy, and I have studied core-electron binding energies in disordered alloys as well as x-ray absorption spectra in materials. I finished my PhD at Uppsala University in 2005, and after that I have worked as a postdoctoral researcher at Kyoto University, Japan, and at Leoben University, Austria.

WEINE OLOVSSON

and for high performance computing

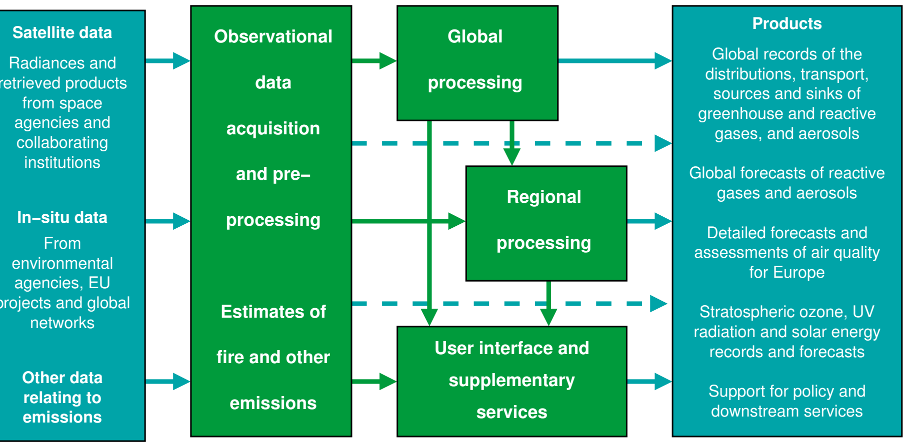

ure 1: Production chain starting with satellite and in-situ chemical measurements and ending in various products.

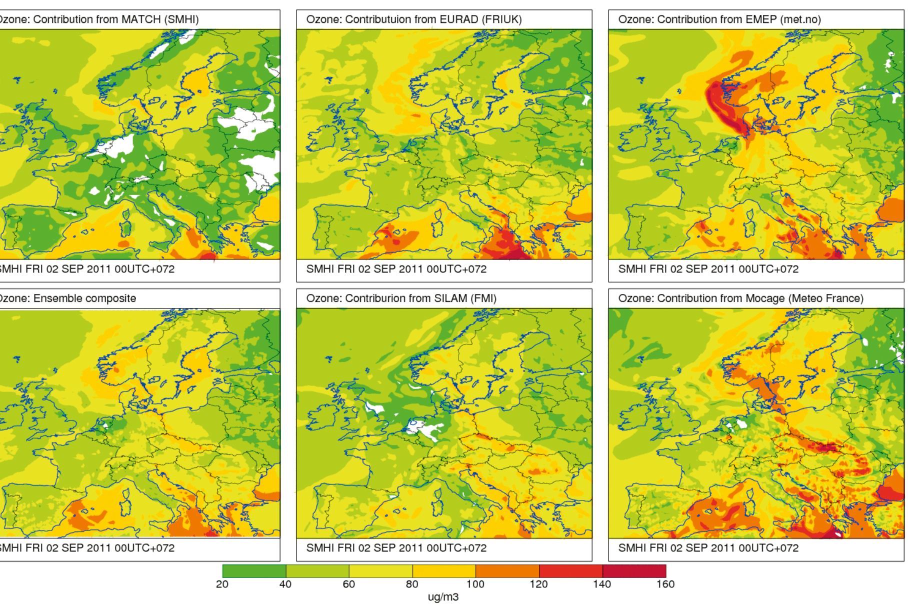

ure 2: Ensemble composite (lower left) from a set of Ozone forecasts from different contributing partners.

grid-cells. Concerning meteorology MATCH is a so called off-line model implying that meteorological information is taken from an external source fed into the model at regular intervals (typically in 3 hour steps). The meteorology used is by nature of MACC from the ECMWF integrated forecast system (IFS). The MATCH tool is also equipped with a number of chemical transformation schemes where the more complex photo-chemistry scheme is used for the MACC products (Simpson et al., 1993). This involves some 100 chemical species like Ozone, NO, NO_2, SO_2, CO, hydro-carbons, various aerosols just to mention a few. The chemical reaction scheme arrives at a set of stiff ODE with a wide range of reaction speeds. This is solved by the Rosenbrook method implemented through a kinetic pre-processor (KPP) that takes an ASCII table of reactions and returns a Fortran code with the complete solver. The atmospheric transport and diffusion is handled by mass-conservative advection and diffusion operators. The transport and chemistry is also constrained by emission inventories of the key species, by the surface characterisation (forests, grass etc.) and boundary conditions, where the latter are prescribed from simulations made on the global scale (the global service in MACC).

MPI course at NSC

A beginners level course in parallel programming using MPI will be hosted by NSC between 26 and 28 October and is open for all SNIC users. The course assumes no prior experience in parallel computing. The concepts behind message passing and distributed memory computing will be introduced and the syntax of the key MPI calls will be explained. The course will include point-to-point communications, non-blocking communication and the collective communications calls. Practical sessions to deepen the understanding of the lectures will be part of the course. Back-to-back with the MPI course, an introductory course on Grid computing and storage targeting new and current HPC users will be held 25 October. An application form and more information regarding these courses is available at URL http://www.nsc.liu.se/mpicourse.

JOHAN RABER

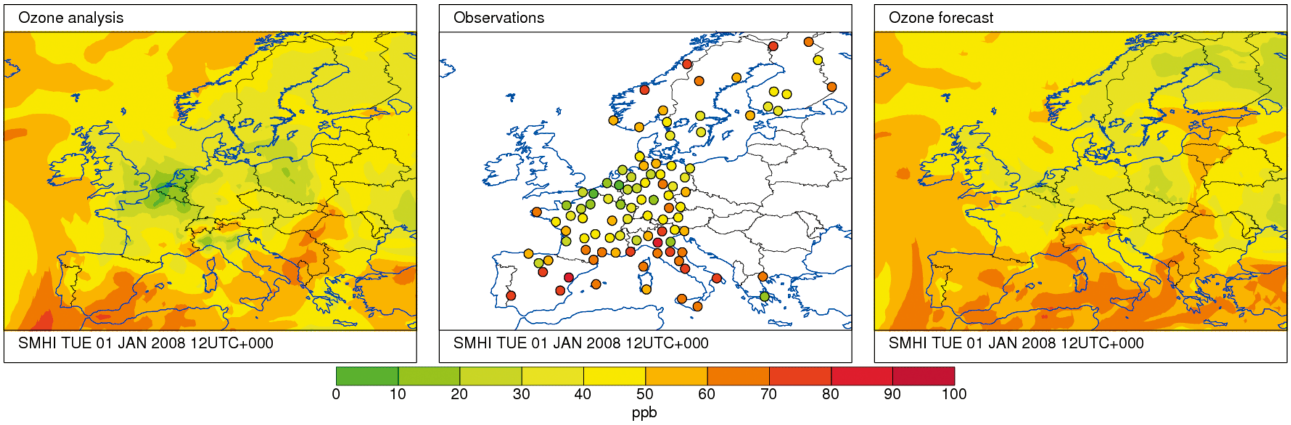

Figure 3: Refinement of an Ozone forecast (right) by inclusion of chemical observations (left and middle).

There is an obvious need for high performance computing where the resources at NSC play an important role to our contribution to the MACC project. In addition to daily forecasting service, hindcasts from 2007 and onwards will be conducted to build up an atmospheric composition archive. A simulation for just a single year demands 4.5 days of computations on 64 cores at gimle just to illustrate the need for the resources that we have available at NSC.

The model calculations are also refined by including observations by using so called data assimilation (Figure 3). This is performed by variational methods minimizing a penalty function representing a balance between model and observation errors (Kahnert 2008).

For the moment only near surface in-situ chemical observations are used but satellite data will soon be incorporated that will add a vertical dimension to the observed chemical state.

All the results from the various clusters in the MACC project are available on a web portal, http://www.gmes-atmosphere.eu, where access is provided by web request, ftp or GIS map-services.

For more information see, http://www.gmes-atmosphere.eu.

References

Kahnert, M. 2008, Variational data analysis of aerosol species in a regional CTM: Background error covariance constraint and aerosol optical observations operators. *Tellus, 60B*, 753–770.

Langner, J., Bergström, R. and Pleijel, K. 1998. European scale modeling of sulphur, oxidized nitrogen and photo-chemical oxidants. Model development and evaluation for the 1994 growing season. SMHI report, RMK No. 82, Swedish Met. And Hydrol. Inst., Norrköping, Sweden.

Robertson, L. , Langner, J. and Engardt, M. 1999. An Eulerian limited-area atmospheric transport model. *J. Appl. Met.* 38, 190–210.

Simpson, D. , Andersson-Sköld, Y. and Jenkin, M.E. 1993. Updating the chemical scheme for the EMEP MSC-W oxidant model: current status. EMEP MSC-W Nore 2/93.

Lennart Robertson is research meteorologist at the research department of SMHI. He has been working with air pollution modelling the last decades with a special interest for data assimilation.

NSC Express

On April 8, NSC launched the NSC Express self-service portal. It makes it easier and faster for you to view your information and to request resources from NSC. At the same time it makes the administration simpler for us at NSC.

As a new NSC user, you register directly in NSC Express. If you are an existing NSC user, you get access by proving that you can log in to a cluster, or by contacting us to get an email with login information.

As an NSC user, you can use NSC Express to view and update your contact information. You can request membership in projects (the principal investigator will get an email with a link used to confirm it) and login accounts on the systems. You can also view core-hour statistics for the projects you are a member of.

As a principal investigator you can add and remove members in your projects. You can even use NSC Express to apply for new small-scale SNAC projects or LiU local projects (the large- and medium-scale SNAC projects are still applied for centrally with SNIC.)

Why not visit www.nsc.liu.se/express to try it yourself? If you have any questions or suggestions for improvements, please send them to support@nsc.liu.se!

KENT ENGSTRÖM AND
TORBEN RASMUSSEN

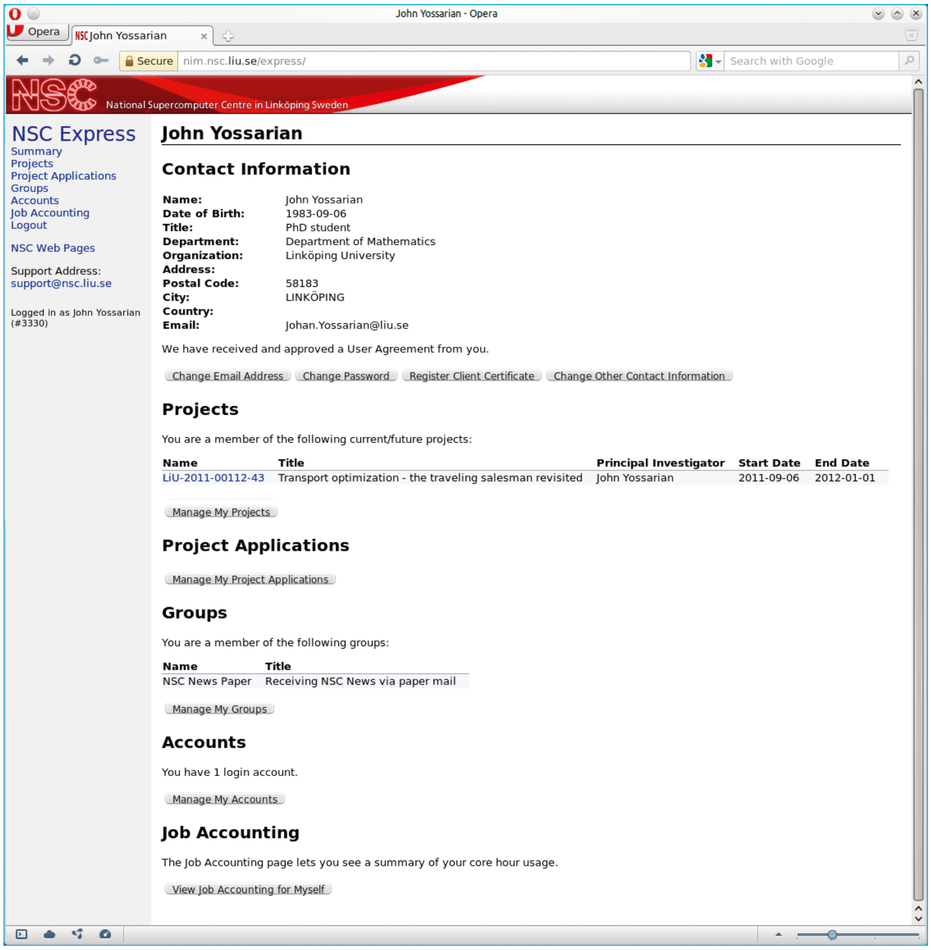

Linköpings universitet

National Supercomputer Centre, Linköpings universitet, 581 83 Linköping, Sweden
Tel: +46 (0) 13-28 26 18, fax: +46 (0) 13-28 25 35, e-mail: info@nsc.liu.se
www.nsc.liu.se

Editor: Peter Münger Use NSC Express, www.nsc.liu.se/express, to subscribe for a paper or electronic version.