# Allinea DDT helps Weather and Climate Codes scaling up for the Cray XC40

## While the predictability – or seemingly lack thereof – of weather patterns is a headline-grabbing topic, great strides are being made to increase the fidelity and range of forecasts on both a local level and global scale.

While the predictability – or seemingly lack thereof – of weather patterns is a headline-grabbing topic, great strides are being made to increase the fidelity and range of forecasts on both a local level and global scale.

Leaders in numerical weather prediction - at the UK's Met Office and the European Centre for Medium-range Weather Forecasts (ECMWF) – are accelerating their software development projects to take advantage of the capability and scale of their new Cray XC40 systems.


Tied together within the same community, the fields of weather forecasting and climate modelling may differ in both focus and approach, but the fundamental goal of developing more accurate predictions through the use of higher fidelity modelling and simulation tools remains the same. And critical to this goal is high-performance computing (HPC).

Dr Paul Selwood, manager of HPC optimisation at the Met Office, one of the world's leading weather forecasting and climate research organisations, emphasised the impact of HPC: "Over the past 30 years, we have been able to improve the accuracy of our forecasts by roughly one day per decade. This means that every 10 years, a forecast will have the same level of reliability at three days out that it previously did at two days out," he said. "HPC provides the means to actually do the complex calculations that our science has enabled, and so every forecast and prediction is reliant on those resources. It's absolutely central to what we do."

HPC in weather forecasting is not without its challenges. According to Mike Hawkins, head of the High Performance Computing and Storage Section at ECMWF, the main computing issue in this field has been the scalability of the codes. In order to increase the accuracy of forecasts, higher resolutions and more accurate simulations of physical processes are required. This needs more computing power and the only practical way to achieve this is through the use of more processors, as Hawkins explained: "There hasn't been a dramatic increase in the speed of processors and so we have to make use of an increasing number of them."

"This presents us with a real challenge because we then have to put a forecast system that's been developed over a long time and that we know produces very good results, into a new architecture. Not only that, but we have to ensure it will work for the next decade at least." Hawkins added that with each new architecture the Centre has to look at the scalability and portability of the codes, and how they can be maintained for long periods of time. This represents a major and long-term development programme for the centre.

## Honing in on the hardware

In the latter half of 2014, The Met Office announced that it had awarded a contract to Cray for the provision of multiple Cray XC supercomputers and Cray Sonexion storage systems. These systems will replace the current pair of IBM Power7 supercomputers that are currently coming to the end of their operational lives at the centre. The £97 million ($128 million) contract will consist of three phases over several years and will expand Cray's presence in the global weather and climate community.


Consistency and reliability are more than buzzwords within weather forecasting, where unique demands are placed on the HPC systems. At The Met Office, based in the UK, two separate systems operate in unison in order to provide operational resilience as data from more than 10 million weather observations per day are used alongside an advanced atmospheric model to create 3,000 tailored forecasts on a daily basis.

"Most high-performance computers exist in academia where the set-up essentially consists of one large machine being used to run the most expansive leading-edge science possible. And if that one machine goes offline for a week, the impact is often minimal," commented Dr Paul Selwood. "At The Met Office, however, we need to get multiple forecasts out on time every single day and if any one of those forecasts is late, it's instantly of no use." Selwood continued by saying that forecasting remains an unusual field within HPC in that the machine's primary function is as an operational system rather than a research system.

The European Centre for Medium-range Weather Forecasts (ECMWF) has also recently deployed two new Cray systems. Following a competitive tendering exercise, ECMWF selected two Cray XC30 systems which arrived on site at the end of 2013. The systems came into production last year and have been running the Centre's operational forecast since September 2014. "Changing from one supercomputing system to another is an incredibly large job - getting to the point of turning it on requires months of work,' ECMWF's Mike Hawkins explained.

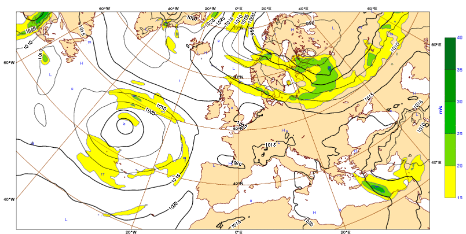

"We have to produce an operational weather forecast, to a tight timescale, several times a day so before switching systems we need to ensure that the forecasts run reliably on the new machine offer the same standard of results as those run on the previous system. This is quite a time-consuming process that contributes to fairly long periods of parallel running between the old and new systems," he added.

The two Cray XC30 systems at ECMWF offer about 3,500 nodes of parallel processing based on two 12-core Intel processors, with a theoretical peak performance of about 1.8 Petaflops per machine. The Centre has a benchmark for measuring the performance of the systems based on a stripped down version of its code that is easier for both the Centre and vendors to run and test. Based on that measure, the two systems have a combined sustained performance of 200 Teraflops – a boost of three times the overall processing power compared to the previous machines deployed at the Centre.

## Ensuring code fidelity

Hardware is only part of the puzzle: the reliability of the code can be of paramount importance, operational runs at the centre must produce results to very tight timeframes and ensuring that all bugs are removed from the system is critical. With this in mind, both The Met Office and ECMWF chose to deploy a powerful debugger, Allinea DDT.

"The best recommendation we received for Allinea DDT from the developers using it is that it does what good debugger is supposed to do, it does it fast and it does it in a very intuitive way," said Mike Hawkins at ECMWF. "We needed a solution that could not only handle our very complicated code, but that could be used in a wide range of scales. These scales range from people who are developing pieces of code on their desktop machines – or perhaps even on one or two nodes on the supercomputer – right up to a full configuration, which takes up to 100 nodes of our machine. Having the same tool work across that entire spectrum is of definite benefit to us."

Likewise, the demands at The Met Office meant that any chosen debugger needed to be able to handle complex code development and porting work. "Allinea DDT was one of the very few parallel debuggers that could fit in well with the types of workflows we have," said Dr Paul Selwood. "The right debugger can save everyone quite a bit of time, and that's what we're really trying to do with Allinea DDT – improve the time to solution for a development or investigation of a scientific problem." He added that within the new system, Allinea DDT has found a number

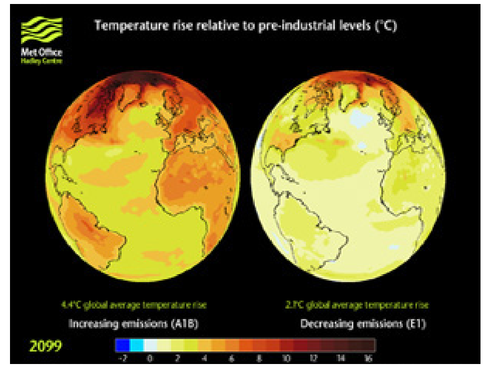


of problems relatively quickly, which in turn has enabled those bugs to be fixed equally as fast.

Allinea DDT ensures that even the most complex multi-threaded or multi-process software problems can be solved quickly and easily. Users see process and thread anomalies instantly with the parallel stack display – the scalable view that narrows down problems for any level of concurrency: from a handful of processes and threads to hundreds of thousands.

Allinea DDT also lowers the risk of major changes during code development – by working with software version control systems to highlight code changes, or to automatically log values of variables across all processes at each changed section of code. This enables users to track down exactly how and why a particular change introduced problems to the code. This is particularly useful for The Met Office which is taking the rather unusual step of utterly rewriting its codes with new algorithms. "It's a different approach to programming compared to what we have traditionally used," Selwood explained. "It's a long-term project and it's going to take a while before we really see the benefit, but given the direction HPC architectures are heading in with ever increasing demands on parallelism and scalability, it's something we need to do. Our current codes are 30 years old and showing their age."

As HPC architectures and the code bases in numerical weather prediction evolve, the forecast is clear that Allinea DDT will continue to be the world's most scalable and sought-after debugger.