Data Assimilation on future computer architectures

Lars Isaksen

ECMWF, Reading, UK
Lars.Isaksen@ecmwf.int

Abstract

The Data Assimilation scalability issues on today's computer architectures are described using the 4-dimensional variational assimilation system (4D-Var) at ECMWF as an example. The scalability issues that affect the ECMWF assimilation system are to a large extend representative for the operational data assimilation systems used at other NWP centres. It is discussed how the future computer architectures are most likely to look. The high performance computers today serve a large variety of users, so numerical weather prediction is now only a small fraction of the user community. The use of 'off-the-shelf' hardware to build high performance computers will continue for the foreseeable future. The challenge is to adapt the data assimilation systems to the computers that are available. It will be addressed if we will be able to use future parallel computers efficiently for data assimilation, and if we can modify our data assimilation methods to utilize future computer architectures better.

1. Introduction

An increasingly important driver for the future developments of our data assimilation systems is the evolution of supercomputer architectures. The main topic of this paper is to discuss data assimilation on future computer architectures. It is now clear that any significant increase of performance of our future systems the next decade will come from increased parallelism, as the future increase in computing performance will not come from increased speed of individual cores, but rather from an increase in the number of cores on individual chips. A recent thorough analysis of the incremental 4D-Var system at ECMWF has indicated that although considerable effort has been (and is) dedicated to the optimisation and parallelisation of its current implementation, there is limited scope for additional tangible gains in scalability to be achieved without a profound revision of the algorithms. Therefore it is natural that the computational aspects of the most viable alternative method, the Ensemble Kalman Filter, also are discussed in this paper.

The last decade has seen very substantial improvement in the quality of analyses and short-range forecasts due to improved forecast modelling and data assimilation techniques and improved observations, especially the space-based component of the global observing system (Simmons and Hollingsworth 2002, Thépaut and Andersson 2010, Andersson and Thépaut 2010). Data assimilation is a central component of any NWP system, and key to continued improvement of forecast skill.

ECMWF has invested significantly in the 4D-Var technique and pioneered its developments and implementation in an operational context, putting ECMWF in a world-leading position and paving the way for many other NWP centres. In addition to requiring considerable initial investment, 4D-Var faces two main challenges. First, as a deterministic data assimilation technique, it does not provide an estimation of analysis and forecast uncertainty, and is limited in its ability to cycle flow-dependent error covariances. Second, the scalability of 4D-Var is limited due to the sequential nature of its algorithm, which makes it unsuitable in its current formulation for future high performance computer (HPC) architectures. In this paper we discuss how to modify the data assimilation system to make better use of the next generations of HPC architectures.

Ensemble Kalman Filtering (EnKF) techniques have recently received a lot of attention in the NWP community, because of their simplicity and their intrinsically scalable nature. They provide fully flow-dependent error covariances. However, these covariances are approximate and noisy, due to the limited affordable size of the ensemble. The competitiveness of the EnKF in a real-size context with a full global observing system remains to be demonstrated.

Hybrid techniques, combining 4D-Var with an ensemble component (ensemble of low resolution 4D-Var assimilations (EDA) or EnKF) seem to be the most attractive way to progress and benefit from the two approaches. At ECMWF we envisage a scenario which will require a profound modification of the 4D-Var algorithm to improve its scalability, accompanied with a radical optimisation of the early delivery analyses from which the operational forecasts are initialised.

The 4D-Var technique provides the flexibility required to extract information from a wide variety of observations, including those that are only indirectly linked with the model/analysis variables. It will remain a vital consideration for the data assimilation system of the future to take full advantage of the information obtained from high resolution conventional data and the multitude of space-based observing system technologies.

The analyses produced through data assimilation have several important uses at NWP centres. These several application areas impose different requirements on the design of the future earth-system data assimilation system that are partly overlapping and partly conflicting: improved analysis quality through the use of higher resolutions and longer assimilation windows, flow-dependent background errors, reliable estimation of analysis and forecast accuracy, increased number of analysed fields, and the degree of coupling to ocean and land-surface analyses.

2. The high performance computer system at ECMWF

The ECMWF High Performance System (HPC) consists at the moment of two identical IBM Power6 systems with 9200 cores each. This system is in the process of being upgraded to two identical IBM Power7 systems with 24400 cores each. Each core consists of two threads on the IBM Power6, and two to four threads on the IBM Power7.

It is interesting to look at the evolution of ECMWF's high performance computer system since the first system, a Cray-1A with one processor, was installed in 1978. Table 1 shows the performance increase and increase in the various system components during the last 30 years. The sustained performance per CPU has only increased by a factor of 24, primarily because the clock speed only has gone up by a factor of 60. The CRAY-1A processor was an advanced highly specialized CPU, produced in small quantities for NWP, fluid dynamics and intelligence applications. Today's supercomputers use commodity processors that are cost-effective and large volume off-the-shelf products with a diverse user base.

Specification	**Cray-1A**	**IBM Power6 System**	**Approx. Ratio**
Year installed	1978	2009	
Architecture	Vector Processor	Cluster of scalar CPUs	
Number of CPUs	1	~9,000	9,000:1
Clock Speed	12.5 nsec (80MHz)	0.21 nsec (4.7 GHz)	60:1
Peak performance per CPU	160 MFLOPS	18.8 GFLOPS	120:1
Sustained performance per CPU	~50 MFLOPS	~1.18 GFLOPS	24:1
Peak performance per system	160 MFLOPS	~320 TFLOPS	2,000,000:1
Sustained performance of operational assimilation system	~50 MFLOPS	~20 TFLOPS	400,000:1
Memory	8 Mbytes	~40 TBytes	5,000,000:1
Disk Space	2.5 GBytes	~1.2 PBytes	500,000:1

Table 1. A comparison of ECMWF's first supercomputer (Cray-1A) with one cluster of the present (2009-12) dual system IBM Power6.
Source *http://www.ecmwf.int/services/computing/overview/supercomputer_history.html*

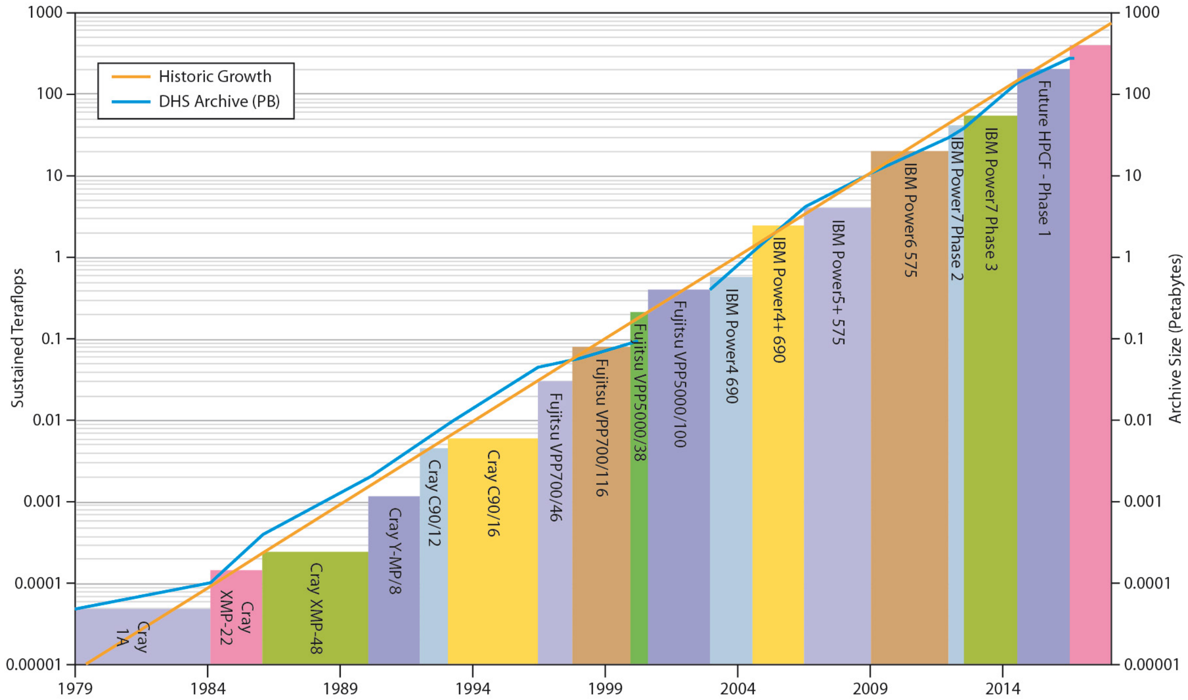

Figure 1: ECMWF's sustained historical supercomputer performance growth and archiving growth from 1979 until today, including the prediction for the next generation system. The vertical axes are logarithmic. The left axis is sustained Teraflops, applicable for the bars. The bars show the name of the supercomputer system. The beige line shows the average exponential growth of sustained computer performance over the period (57% p.a.). The right axis is the unit for the ECMWF DHS (Data Handling System) archive size in PetaBytes, represented by the blue curve. (Source ECMWF: http://www.ecmwf.int/services/computing)

The relative sophistication of processors have been reduced with time, so the primary reason for the overall increased performance of super computers over the years is due to the added parallelism by increasing the number of CPU's (processors or cores) from one to more than 9000.

The amount of memory and disk space has increased by a factor that is very comparable to the overall sustained computational performance increase. This is because the storage systems are using the chip technologies that are responsible for the CPU related improvements. This relationship is evident from Fig. 1 that shows the clear link between ECMWF's archive size and the sustained supercomputer performance.

Fig. 1 shows the almost constant exponential growth of sustained computer power during the last 30 years at ECMWF. It is a manifestation of Moore's law that has ensured that ECMWF has been able to achieve this performance (and archive size) increase for only a small annual increase in real financial cost. A similar historical evolution has been seen at other NWP centres. The same degree of exponential growth has been achieved for the performance of the top 500 super computers (both for the top, the median, and the bottom) over the last 20 years (see http://www.top500.org for further details).

3. The operational data assimilation system at ECMWF

The increase in sustainable supercomputer performance has over the years primarily been used to increase the horizontal and vertical model resolution. But it has also allowed ECMWF to introduce more advanced model parameterizations and use more advanced data assimilation methods. The time spent on running the operational analysis and forecast has been kept more or less the same over the last 30 years. Today, in 2012, the main application is the operational forecast, the 4D-Var assimilation system, and the Ensemble Prediction System (EPS). For this we are using the IFS (Integrated Forecast System, developed by ECMWF and Météo-France. See http://www.ecmwf.int/research/ifsdocs). The operational assimilation is performed twice a day using a 12 hour 4D-Var T1279 (16 km horizontal grid) outer loop, and T159/T255 inner loops (Courtier *et al.* 1994). This is followed by a 10-day T1279L91 forecast. In 2010 an Ensemble of Data Assimilations (EDA) was implemented as part of the operational system to provide flow-dependent estimates of short-range forecasts (Isaksen *et al.* 2010; Bonavita *et al.* 2011; Bonavita 2012 in these proceedings). It is using a 10 member 4D-Var with T399 outer loop and T95 and T159 inner loops. The EPS is using 50 members at T639L62 up to day 10, and T399L62 from day 11-15. Twice weekly the EPS is extended to 32 days, to provide monthly forecasts.

Deborah Salmond from ECMWF has performed extensive evaluation of the scalability of our assimilation system on the present IBM Power6 computer. Some of these results will be presented here. Fig. 2 shows the scalability of the operational T1279 forecast model and the 4D-Var analysis on ECMWF's present IBM Power6 computer. The plots show the sustained performance, based on the wall clock time spent to perform the computations, for a range of threads (two threads per core on IBM Power6) in the range from 2000 to 6000. The operational assimilation system is using 3072 threads. The forecast model (red curve) scales reasonably well compared to the ideal scaling (blue curve). The 4D-Var (black curve) on the other hand shows rather poor scaling characteristics going from 2000 to 6000 threads. The ECMWF assimilation system is using the incremental formulation (Courtier *et al.* 1994) of 4D-Var, where the so-called trajectory step compares observations against a

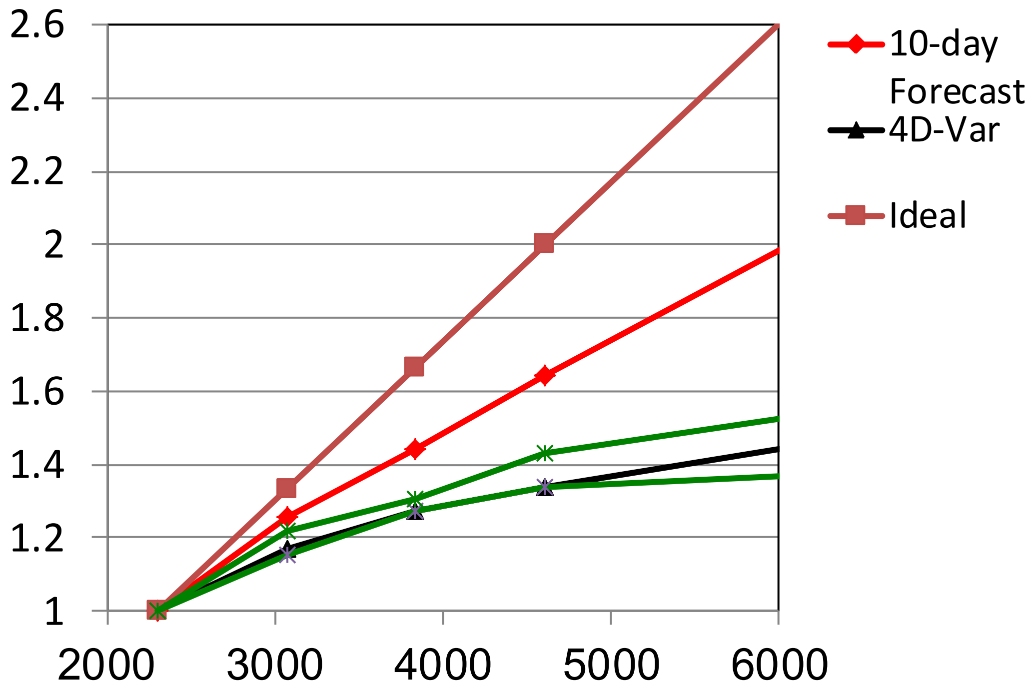

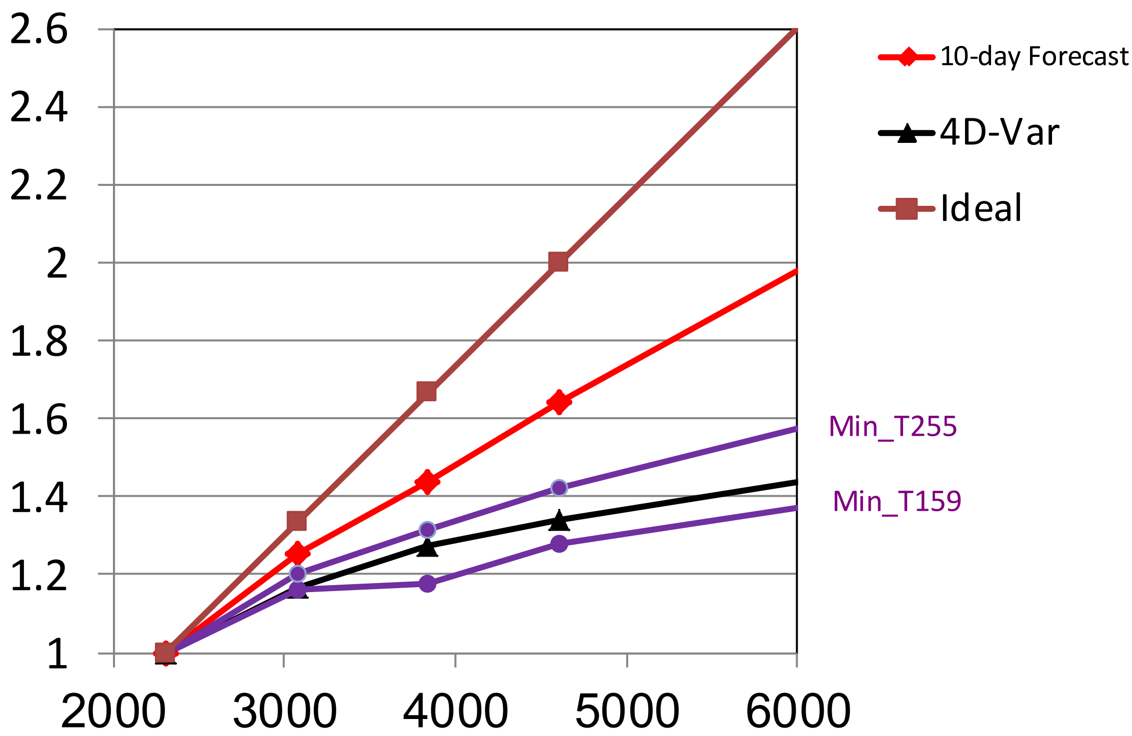

Figure 2: Scalability of the ECMWF forecast model and 4D-Var assimilation system. The upper panel include the trajectory steps (green curves). The bottom panel include the minimization steps. Speed up is shown as function of number of threads. The ideal scaling curve is also shown. Courtesy Deborah Salmond (ECMWF).

T1279 background forecast, followed by a minimization at either T159 or T255 resolution. It is more cost effective to use the incremental formulation where the expensive minimization steps are done at lower resolution. On Fig. 2a the two green curves represents the trajectory steps. The least scalable trajectory step (traj_0) runs a T1279 forecast and compare it against all available observations (of the order 150 million for a 12 hour period), whereas the other trajectory tasks (Traj_1 and Traj_2) only compare model forecasts against the active observations (of the order 12 million for a 12 hour period). The purple curves on fig. 2b present the scalability of the minimisation steps. The T159 minimisation (purple curve labelled Min_159) is less scalable than the T255 (purple curve labelled Min_255). The reason for the difference in scalability for the various steps is primarily due to the difference in the number of grid columns and due to the extra IO and communication for the handling of observations. The T1279 has approximately 2,000,000 grid columns, compared to the 89,000 for T255 and 36,000 for T159. So for a fixed number of cores there is much less work per core for the lower resolution models. The amount of communication is also reduced, but not proportionally as much. Studies of the scalability of the forecast model at T159 resolution (not shown) found that it is even less scalable than the T159 analysis step. This is because the linear adjoint and tangent model is more compute intensive than the non-linear model. These aspects will be discussed in more detail in Section 7.

4. Future computer architectures – hardware issues

In the previous Section it was shown how the major change over the last 30 years of supercomputer design has been the increase in number of cores. We will now discuss how this may change over the next decade. It is evident that the trend towards off-the-shelf hardware will continue, with a mass market for applications that are not exceedingly compute intensive. Most applications will be able to run on one multi-core chip, or are embarrassingly parallel. The focus on power consumption and cooling costs means clock speed and power comsumption per chip will remain more or less the same. But there is no evidence that the expontial growth in computer performance will stop. Moore's law of a doubling in the number of transisters per chip every second year is continuing to hold true. Fig. 3 shows how the increase in chip performance during the last 40 years can be contributed to clock speed (frequency), and to the number of cores per chip. For the period 1970-2003 the improved performance was primarily obtained by increasing the clock frequency. But since 2003 the clock speed has not increased. The increased performance of the chips has been obtained by an increase in the number of cores on each chip. The main reason for this is that power cost increase significantly with higher frequency. In addition to increased power comsumption, it also cause increased cost of cooling the chips. It is much more efficient use of power to move to a multi-core chip concept, rather than increasing the clock speed. The chip design plans for all the major players in the field shows that this trend will continue. It is expected that the number of cores per chip will double each year for the next decade, with the clock speed staying unchanged or slightly reduced. This has significant implications for the design and coding of the future data assimilation systems. Where we now have of the order ten thousand cores, we will very likely see systems for operational NWP with more than one million cores within a decade.

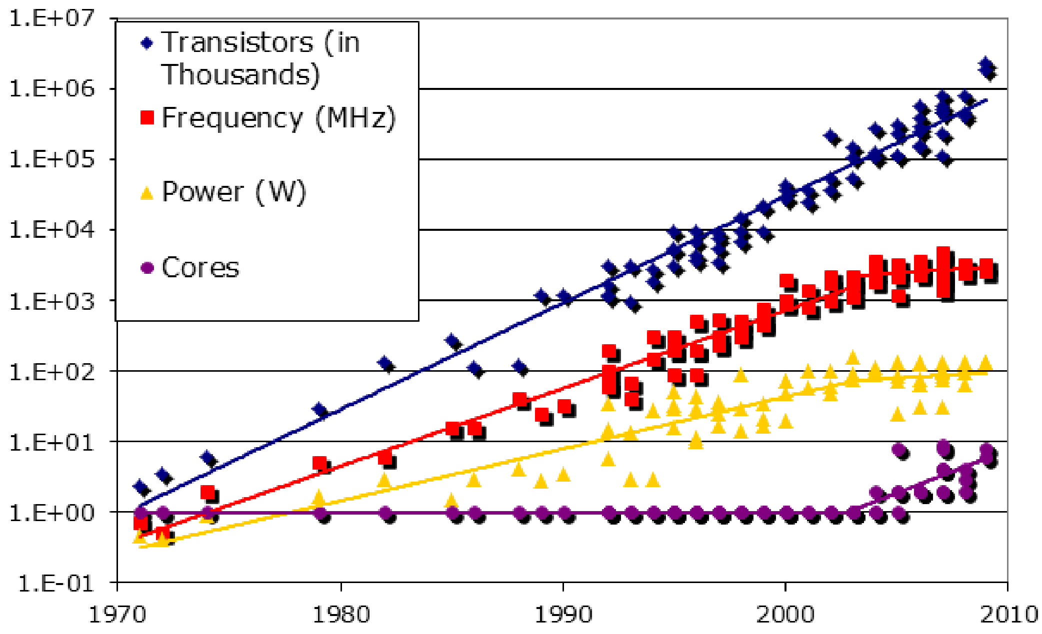

Figure 3: The evolution of computer chips from 1970 until 2009. The curves represent trends of: number of transistors [in thousands] (blue diamonds), clock speed/frequency [mHz] (red squares), power consumption per chip [W] (yellow triangles), and number of cores per chip (purple circles). Data from K. Olukotun, L. Hammond, H. Sutter, B. Smith, C. Batten, and K. Asanoviç. Graphics by Jack Dongarra et al. 2011: Linear Algebra Libraries for High-Performance Computing: Scientific Computing with Multicore and Accelerators. Presented at SC11, Seattle.

We will give a few examples of top of the range super computers. The number one on TOP500 in 2012 is the Japanese K-computer with 10 PFLOPS peak performance. It is a $1.3 billion national funded initiative. The computer consists of 80,000 Fujitsu SPARC64 VIIIfx CPUs. A significant amount of the cost is for the powerful interconnect that ensures this machine is able to run real applications effectively using a large fraction of the system on a single problem.

For petascale computing the interconnect is a very important and potentially expensive component of the future super computers. It is a question which type of topology can scale up to 100,000 cores or more. There are three types of interconnect: cross bar, fat-tree/multi stage, or a mesh/torus. The mesh/torus approach is the most viable solution for systems with more than 10,000 cores. Improvement of the performance, the operability and availability of mesh/torus topology is the biggest challenge.

On the processor design side there is presently debates going on whether future systems will be based on more general purpose CPUs or special purpose GPUs (Graphical Processor Units). The GPUs have been designed by companies like NVIDIA for graphics cards used on workstations and PCs. NVIDIA's Fermi chip is the first to support high performance computing. They have formed partnerships with Cray and IBM on developing GPU based HPC systems. Number 2 and 4 super computer systems on TOP500 is of that type. Other chip manufacturers like AMD/ATI and Intel are also developing and producing GPU based chips for high performance systems.

Fig. 4 shows schematically the comparison of CPUs versus GPUs.

CHIP TYPE	CPU Nahalem	GPU NVIDIA Tesla	GPU NVIDIA Fermi
Cores	4	240	480
Parallelism	Medium Grain	Fine Grain	Fine Grain
Performance Single Precision Double Precision	85 GFlops	933 GFlops 60 GFlops	1040 GFlops 500 GFlops
Power Consumption	90-130W	150W	220W
Transistors	730 million	1.4 bilion	3.0 billion

Figure 4: Comparison of state-of-the-art CPU and GPU chips. Source: M. Govett et al. 2010: Using GPUs to Run Weather Prediction Models, 14th Workshop on Use of High Performance Computing in Meteorology.

The main difference between CPUs and GPUs is that the GPUs have many more cores per chip. This results in increased performance for the same power consumption. The GPU performance is not increased proportionally to the increase in number of cores, because the clock speed in lower and the chip has a reduced and leaner functionality. But GPU based systems would be favoured over CPU based system, if a performance per power unit metric is used. As an example of a next generation system for NWP applications, the GPU camp compare the CPU based **USA Department Of Energy Jaguar system** (2.3 PetaFlops peak performance, 250,000 CPUs, 284 cabinets, 7-10 MW power consumption, price of the order $100 million) versus an **Equivalent GPU System** (2.3 PetaFlops peak performance, 2000 Fermi GPUs, 20 cabinets, 1.0 MW power consumption, price approximately $10 million). The GPU will be more reliable because it consists of only 2000 GPUs versus 250,000 CPUs. The GPU camp argues that large CPU systems (>100,000 cores) are unrealistic for operational weather forecasting due to several issues: power consumption, cooling, reliability, and cost. But in reality it is not so straightforward. GPUs are more difficult to program (as will be addressed in Section 5), and CPUs are gradually becoming more like GPUs, with increased number of cores per chips. The likely scenario is that CPUs and GPUs melt into one technology where the general purpose benefit of CPUs is maintained together with exploitation of the multi core GPU benefit (like increased performance per power unit).

5. Future computer architectures – software issues

GCUs are special purpose build chips that people now are trying to use for more general purpose applications like NWP. The Fortran GPU Compilers still lacks generality: they do not support all Fortran language constructs, so it is required to use lower level languages, like CUDA. It is also necessary to use an extensive set of parallelization directives to guide compiler analysis and optimization. This means it requires additional effort to use GPUs. It has been done at NOAA, where a complete NWP model has been rewritten using CUDA. It is expected to become easier with time to use GPUs, if there is a sufficient user market for GPU based computers, because then computer vendors will have an incentive to develop better and easier usable compilers.

The trend towards more cores per chip will increase the peak performance of chips, but the memory bandwidth and I/O bandwidth will not improve as quickly. Technology improvements may help, but the costs of similarly fast and large volume memory (and I/O) may become unacceptably costly. The trend is decreasing bytes/flop ratios, so locality counts. Packaging technology, like I/O switching costs becomes more of an issue - the relative amount of power needed to move data will increase. Flop metric promises to be an even poorer predictor of sustained performance in the future than it is now. So it will be more difficult to use the next generation computers. Few applications will scale to exascale with their current structure due to things like: lack of sufficient parallelism; load imbalance; lack of data reuse (too few computations per memory access); use of algorithms and data decompositions that require extensive communication. The next generation computer will be limited by bottlenecks in memory, communications, and I/O bandwidths. This not so far future for high performance computing was summarized well by Don Grice (2010, ECMWF HPC Workshop). He concluded that the fundamental programming style is not likely to change much: multi-threaded 'MPI tasks' will be the norm, new languages are emerging to help with extreme scale and a shared memory model at the task level will still exist. The trends will be that the amount of threading will increase; 'more science' will be a way to use CPU cycles. Optimization points will change – computing is 'free', so e.g. more advanced physical parameterizations can be implemented at no extra real time cost.

Today the scaling issues in NWP are primarily due to static and dynamic load imbalance, e.g. convection is more likely in the tropics, leading to imbalance. Jitter in the interconnect also leads to hick-up delays in communication. A problem that increases with the number of processors involved in the communication. MPI communications latency, choice of topology, openMP overheads, I/O and slow sequential shell scripts are all important issues today.

There are many unknown regarding the problems we can expect related to scaling to 100,000 - 1 million cores. It is a question if we can continue to use MPI and OpenMP. Jitter problems will have to be resolved. Overlapping of computations and communications may be the solution to the reduction in communication speed compared to the speed of cores. There is a need for new computer language concepts that allows the user to guide the distribution of data, e.g. Fortran 2008 co-arrays or partitioned global address space (PGAS) languages. Finally there will be an increased need for tools (debuggers, profilers) that work reliably and fast at high core counts, and work with large applications.

6. Assimilation algorithms – their strengths and weaknesses

We will restrict the discussion here to the main assimilation technique, 4D-Var, widely used operationally in the NWP community, and the Ensemble Kalman Filter that has been the focus of much recent active research, plus hybrid techniques that combine the deterministic and ensemble components. Other techniques (e.g particle filters) exist that potentially offer many attractive properties but are not considered practical within this decade. These will therefore not being discussed here.

6.1. 4D-Var

4D-Variational data assimilation approaches (Lewis and Derber, 1985; Le Dimet and Talagrand, 1986) have been pioneered at ECMWF (Courtier and Talagrand 1987; Courtier *et al*. 1994) in the late eighties, and drove the development of the Integrated Forecasting System of the Centre. This assimilation technique was successfully implemented at ECMWF in 1997 (Rabier *et al*. 2000) and has now been used with great success at several operational NWP Centres for a number of years. The principle of 4D-Var can be summarised as follows: given all the information (observations and their associated errors, background and its associated error, atmospheric model, etc.) available over a period of time (assimilation window), 4D-Var looks for the model trajectory which best fits the background and the observations over this assimilation window. This is done through the minimization of a cost function measuring the misfit between the atmospheric state and the observational and background information. A key feature of 4D-Var is the flow dependent influence of observations in space and time controlled by the tangent-linear model dynamics (Thépaut *et al*. 1996). The strengths and weaknesses of 4D-Var have been widely discussed in the literature (e.g. Lorenc 2003) and will not be discussed here.

In the context of this paper it should be mentioned the computational cost of 4D-Var is high, as it requires many integrations of the tangent linear and adjoint models during the minimisation of the cost function. This is fundamentally a sequential process. The incremental technique, using a simplified model in the inner loop (simple physics, low resolution) allows the cost to be reduced. However, the inner modules of the incremental 4D-Var remain the least scalable parts of the scheme, mainly due to the lower resolution applied. The cost and scalability issues of 4D-Var are discussed at length in Section 7.

6.2. Ensemble Kalman Filter

The Ensemble Kalman Filter (EnKF) has been developed from the mid-90s as a Monte Carlo approximation to the Kalman Filter (Evensen 1994; Burgers et al. 1998; Houtekamer and Mitchell 1998). These first implementations, known as *Stochastic EnKF*, were based on the concept of simulating all possible sources of uncertainty in the data assimilation system. This essentially implies running an ensemble of data assimilation systems where random perturbations drawn from their expected error statistics are added to the observations. In this way the estimated analysis covariances asymptotically match the correct values computed by the Kalman filter update equations. Another possibility is to compute a Kalman filter update of the ensemble mean and the analysis covariance and then construct the ensemble members from a 'square root' of the estimated analysis covariance (Anderson 2001; Whitaker and Hamill 2002; Tippet et al. 2003 for a review) This *Ensemble Square Root Filter* (EnSRF), or *Deterministic EnKF*, presents advantages for small size ensembles due to the elimination of sampling error associated with the perturbed observations In order to be computationally efficient, the Deterministic EnKF assimilates observations sequentially, thus implicitly assuming their errors to be statistically independent. This assumption is not needed in the more recent version of the EnKF known as Local Ensemble Transform Kalman Filter (LETKF, Hunt et al. 2007), where the analysis is computed independently for each grid point using all the observations in a predefined local volume. The excellent scalability of the scheme is due to the fact that the analysis is effectively performed in the error subspace spanned by the ensemble perturbations.

The increasing popularity of the EnKF in geophysical NWP applications is based on both practical and theoretical reasons. Compared to a variational algorithm, the EnKF is considerably simpler to code and maintain. This is because the EnKF algorithm is essentially independent of the prognostic model used to cycle the ensemble and does not need perturbation and adjoint versions of the model and observation operators. The EnKF algorithm is also intrinsically parallel, especially in the LETKF flavour, and thus well suited to massively parallel computer architectures.

From a theoretical perspective both the advantages and disadvantages of the EnKF are essentially a consequence of it being a reduced rank approximation of the Kalman Filter. This is discussed in detail in the papers referenced above, and by Bonavita (2012) and Whitaker (2012) in these proceedings.

6.3. Hybrid methods

It has long been recognized that one of the main limitations of strong-constraint 4D-Var is its inability to propagate the error estimates of the state beyond the time span of the analysis window.

It has been shown (Fisher *et al*., 2005) that if the assimilation interval is made sufficiently long (3-5 days), the weak constraint formulation of 4D-Var produces at the end of the assimilation interval an identical analysis to that produced by an extended Kalman filter that has been running indefinitely, i.e. it will implicitly incorporate the error evolution update of the Kalman filter. This is discussed by Fisher and Auvinen (2012) in these proceedings.

At the other end of the spectrum, the EnKF has been demonstrated as a viable low-rank approximation to the extended Kalman Filter. The EnKF explicitly propagates the error covariance estimates. However the sampled B matrix of the EnKF, although fully flow-dependent, suffers from sampling errors and rank-deficiency issues because of the comparatively small size of the ensemble with respect to the local dimension of the error space. On the contrary, 4D-Var allows the initial B matrix to be high rank because error covariances are propagated implicitly through the tangent linear and adjoint of the forecast model. This suggests the idea of a hybrid system where background error estimates from an EnKF (or, more generally, an ensemble of data assimilations, EDA) are used (alone or in combination with a climatological B) for the construction of a regularised flow-dependent B matrix to be used in the variational analysis. This approach has been first proposed by Hamill and Snyder (2000) in the EnKF - 3D-Var framework and successively extended to hybrid Ensemble - 4D-Var configurations (see Buehner *et al*. 2010a and references therein).

Restricting our attention to Ensemble – 4D-Var hybrids, there are two fundamentally different possibilities. One may use the ensemble sample covariance matrix at the start of the 4D-Var assimilation window and let the variational algorithm implicitly evolve the initial covariances through the tangent linear and adjoint model; alternatively one can use the ensemble covariances throughout the assimilation window, without the need of a tangent linear and adjoint forecast models (this approach is known as En-4D-Var, see Buehner *et al.* 2010a).

The hybrid data assimilation system under development at ECMWF aims to use ensemble information to model the background error covariance matrix B at the start of the 4D-Var window. This configuration is being pursued because: a) it is a gradual evolution from the previous system which only made use of a static initial B computed from a climatology of EDA perturbations; b) ECMWF has already developed a tangent-linear and adjoint version of the model, including a simplified physics package; c) it can avoid some of the difficulties involved in the spatio-temporal localization of the ensemble covariances required in En-4D-Var; d) recent results (Buehner *et al*. 2010b) show this method to provide results of comparable or slightly better quality than En-4D-Var.

Another distinctive feature of the ECMWF approach is that the ensemble providing the flow-dependent covariance estimates is not an EnKF but an Ensemble of Data Assimilations (EDA), i.e. a system of N (N=10 at the time of writing) independent, reduced resolution, 4D-Var assimilation cycles which differ by using randomly perturbed observations, sea-surface temperature fields and model physics. If the perturbations are drawn from the true distributions of observation and model error, then the spread of the EDA about the control (unperturbed) analysis will be representative of the analysis error of the 4D-Var system (Isaksen *et al.* 2010). This result is valid for weakly non-linear systems and is based on the assumption that the effective Kalman gain matrix used in the perturbed runs is equal to the one used in the unperturbed, reference system. In other words, the perturbed analysis should be as close as practically and computationally affordable to the reference analysis, in order for their perturbations to simulate the errors of the reference analysis system.

EDA or EnKF perturbations can also be used for the estimation of the B matrix correlation structures. This can be achieved in more than one way. In a simplified experimental configuration Hamill and Snyder, 2000, proposed a hybrid EnKF - 3D-Var scheme where the B matrix is a linear combination of the static B used in the variational scheme and a flow-dependent B sampled from the EnKF. In state of the art, operational implementations of 3-4D-Var the B matrix is never actually constructed, but it is implicitly defined by the operators that perform the transformation from model variables to the control variables (Fisher, 2003). In this framework, it is convenient to do the combination of the two covariance matrices through augmentation of the control vector (alpha control variable, Lorenc, 2003). Barker and Clayton (2012, in these proceedings) discuss this further.

6.4. Review of current/future evolution at other NWP Centres

At present, all operational global deterministic forecasting systems use a variational data assimilation algorithm (3D-Var or 4D-Var), except for one centre (HMC, Russia) which uses an OI algorithm. In addition to the analysis for the deterministic forecast, several centres (including Météo-France and ECMWF) have started to use ensembles of data assimilations to provide analysis uncertainty estimates and to generate all or some of the initial perturbations for their ensemble prediction system. One centre, CMC (Canada), runs a distinct ensemble data assimilation system for initialising the EPS.

Other operational centres share ECMWF's concern regarding scalability of the data assimilation system on future massively parallel computer systems. EnKF is a very scalable method. However, the scalability aspects have to be balanced against the requirements for analyses to be of equal or better quality than those produced by current operational systems. This has not yet been achieved with an

EnKF. It is not yet clear that the significant scientific challenges involved in applying EnKF methods to the current and future observing system (in particular, the use of very high observation densities with a predominance of vertically non-local radiance data) can be overcome.

7. Scalability and computational cost

7.1. Comparing Forecasts and 4D-Var Analyses

We now make a simplified schematic comparison of the scalability of the forecast versus the 4D-Var analysis. The cost of the tangent linear and adjoint model integrations is the dominant cost in 4D-Var, and is likely to remain dominant for the next decade. The tangent linear and adjoint models are parallelised in exactly the same way as the forecast model, and their scaling properties are identical. The difference in scalability between the forecast and 4D-Var is therefore due to the fact that 4D-Var performs a larger number of sequential timesteps at a lower spatial resolution.

The IFS model is parallelised using two-dimensional distributions in both spectral and grid spaces. The current ECMWF operational system run on 3072 threads, as discussed in Section 3, so each processor is responsible for nearly 700 grid columns for the T1279 forecast and for the 4D-Var trajectory calculations. The minimization is run at lower resolution at the same number of processors. For the T255 minimisations this results in around 29 columns per thread, and for the T159 minimisation 12 columns per thread.

The computational work required to process a single grid column for a single timestep is the same for all horizontal resolutions of model. Hence, the minimum number of columns that can be assigned to a processor before communications costs start to dominate is approximately the same for all resolutions of model. If all resolutions are run with as many processors as can be used efficiently, the wall-clock time required to integrate the model for one timestep becomes independent of the resolution of the model. That is, we have perfect weak scaling: the work per processor per timestep is a resolution-independent constant. Given this observation, we can compare different tasks (e.g. the forecast and the analysis) by examining how many sequential timesteps are required to achieve them. It should be noted that the current operational system is not run in this way. The forecast currently uses considerably more wall-clock time than would be required if it were parallelised to the same extent as the analysis.

Let us consider first the deterministic ten-day forecast that is run with a 10-minute timestep. The number of timesteps required for a ten-day forecast is 1440 (10*24*60/10).

The current 12-hour 4D-Var system performs four T1279 outer-loop integrations of length 12 hours, also using a timestep of 10 minutes. The number of timesteps required for the outer loops are 288 (4*12*60/10).

The T159/T255 inner loops use a half-hour timestep. There are three minimisations, which together perform approximately 100 iterations of minimisation. Each iteration performs a 12-hour tangent linear integration followed by a 12-hour adjoint integration. The computational cost is dominated by the cost of the model integrations, so that it is meaningful to characterise the cost of the inner loops in terms of the number of timesteps of integration.

The total number of timesteps required to perform the inner loops is 4800 (100*2*12*60/30).

We see that the current 12-hour 4D-Var system requires three times more timesteps than the ten-day forecast. It is worth noting that the worst-case scenario is that 12-hour 4D-Var requires a wall-clock time equivalent to a few tens of days of high-resolution model integration, assuming that both use similar ratios of gridpoints to processors.

The primary concern for the future is that the timestep length will decrease as model and analysis resolution increases. This will require more timesteps to be performed. In the case of the forecast model, this increase can be absorbed by reducing the number of grid columns per processor. At least an order of magnitude decrease in timestep length (i.e. an order of magnitude increase in resolution) can be accommodated before the forecast model starts to run out of parallelism.

The situation for the 4D-Var minimization is different, since it is not possible to further decrease the number of grid columns assigned to each processor, because the limit of scalability is reached around 20 grid columns per processor on today's computers. If we do nothing to address the problem, we can expect the wall-clock time required to perform a 4D-Var analysis to increase in line with the reduction in timestep length, while the wall-clock time required for the forecast can stay the same. This process reaches a limit when the forecast and 4D-Var use similar ratios of grid columns to processors. At this point, the scaling properties of the model will be similar to those of 4D-Var, and the cost of 4D-Var relative to that of the forecast will not deteriorate further.

Since an order of magnitude increase in the wall-clock time required to run 4D-Var is unacceptable, it will be necessary to drastically reduce the number of sequential timesteps required to perform an analysis. We believe that this is achievable through a combination of algorithmic changes (as will be discussed in section 7.2), and by reducing the number of timesteps in the time-critical part of the analysis (as will be discussed in section 8.3.

7.2. Parallelisation over sub-windows.

Parallelisation by dividing the inner loops into sub-windows (see Fisher *et al*. 2011, Fisher and Auvinen 2012 in these proceedings) has the potential to significantly reduce the sequential cost of 4D-Var. Specifically, in the scalability computation we may replace the length of the analysis window (12 hours) by the length of the sub-window. For example, with a three hour sub-window, the inner loop cost is reduced, from the 4800 timesteps for a 12 hour window, to 1200 (4800/4) timesteps.

Note also that by parallelising over sub-windows, the sequential cost of 4D-Var becomes independent of the total length of the analysis window. There is no impact on wall-clock time of increasing the length of the analysis window, provided the number of iterations required to minimise the cost function stays constant.

It should be stressed that the arguments presented above are simplified. The estimates assume that the numbers of iterations and the total cost per iteration for weak constraint 4D-Var remain similar to those of strong-constraint 4D-Var. This is unlikely to be the case, since exploiting the parallelism of weak constraint 4D-Var may require us to adopt algorithms that require more (or more expensive) iterations.

7.3. Scaling properties and computational cost of an EnKF system

The Ensemble Kalman Filter is often presented as a data assimilation algorithm that does not suffer from the scaling problems associated with 4D-Var. This is true at least from a theoretical stand-point for certain implementations like the LETKF. In the LETKF, each point on the computational grid is analysed independently of each other, leading to a very high number of independent pieces of work. At the resolution of the current ECMWF EDA system (T399) this implies around 2×10^7 points: an order of magnitude more points than the number of cores envisaged during the next decade. Propagation of the covariance matrix is done through running an ensemble of forecasts, which also will be very scalable.

However, from a practical implementation viewpoint there may be scaling issues. To analyze a point in state space the analysis system needs access to all the “local” observations. If the number of cores approaches the number of state locations there will be a high degree of overlap between the “local” observations needed by one task and those needed by tasks analyzing adjacent areas. The distribution of the observations may become a bottle neck. The severity of this issue is crucially dependent on the way the localization is done and to what extent different cores share the same memory. Also, the cost of reading and writing the ensemble states may reduce the scalability significantly.

The total computational cost of an EnKF is decided by the resolution of the ensemble members, the size of the ensemble and the number of observations assimilated. In general running the ensemble of forecasts constitutes the dominant part of the computational cost of the EnKF. It is then clear that the total cost of the EnKF run at a specified resolution is broadly comparable to the cost of running a 4D-Var with the innermost loop at the same spatial resolution and a number of iterations similar to the EnKF ensemble size, the only difference being that the EnKF makes use of the full model while 4D-Var uses its tangent linear and adjoint implementations. It should be noted, however, that while the EnKF forecasts can be run independently and in parallel, the inner loop iterations of 4D-Var are intrinsically sequential.

7.4. Secondary scalability issues for data assimilation

The principal scaling issues for 4D-Var and EnKF have been outlined above. There are also secondary, more technical, scaling issues that affect all data assimilation systems. One such issue concerns computations in observation space. Observations are diverse and inhomogenously distributed in space and time. The cost of computing observation equivalents varies widely between different observations type, e.g. for conventional observations a simple interpolation in space and time is often sufficient whereas for radiances complex and expensive operators are needed. Load-balancing these computations across processors are difficult and involve communications of either observation information or state-space data. Much effort has been spent to optimize this in the current ECMWF 4D-Var system and it does not yet constitute a major obstacle to scaling. However, it will require constant attention.

Another, and perhaps more important, issue is the amount of I/O required. Apart from the first-guess information and observation data, constant datasets (coefficient files for the radiance operator, J_b statistics for 4D-Var, etc.) have to be read before the analysis can proceed. This is in contrast to the forecast model which only needs to read the initial conditions before proceeding. The often large amount of output from a forecast in the form of post-processed fields can normally be written out asynchronously as the forecast proceeds. The scaling properties for I/O are often negative (I/O takes longer when more processors are used due to the increased cost of distributing the data over processors). The importance the question of I/O, as discussed in Section 4, will have for the overall scaling of a future data assimilation system crucially depends on the underlying I/O subsystem and the amount of (mainly observational) data that will be utilized.

Finally it is important to remember that a significant part of our computer will be used for research experiments, for which the computational resources used is as important as scalability. In research mode we need to run a large range of resolutions. This means that the data assimilation algorithms need to be computationally efficient over several orders of magnitude of processors and a large range of resolutions.

8. Possible design of a future assimilation system

Given the considerable technical and scientific uncertainties discussed above, it is clear that flexibility must be at the heart of our plans for the evolution of the ECMWF Data Assimilation system. This should also apply to other NWP centres.

8.1. Flexibility

It is clear that the complexity of future data assimilation systems will increase. At ECMWF the future IFS code will need to accommodate weak constraint 4D-Var with sub-windows running in parallel, ensemble based error estimation (from an ensemble of variational assimilations or an ensemble Kalman filter), and most likely a hybrid 4D-Var/ensemble system. Managing the additional complexity in the code will be another challenge, in addition to the scientific challenges ahead.

Over the last two decades, computer systems and applications in all domains, scientific or otherwise, have become incredibly complex but also increasingly reliable. This has been possible because of improvement in both hardware and software technologies. In particular, object oriented programming has played an important role. This technology has allowed the development of increasingly complex software designs while at the same time helping to ensure efficiency, reliability and better team development work. It is becoming more and more advantageous to adopt similar technology in the IFS in order to successfully manage the increasingly more complex system. This is the goal of the Object Oriented Prediction System (OOPS) project at ECMWF.

The OOPS project is divided in two main tasks. The first task comprises the design of a high level flexible code structure where data assimilation algorithms can be expressed easily. This code is being developed in C++. The second task in the project is to prepare existing parts of the IFS Fortran code to make them callable from the high level code. The vast majority of the code in the new system will come from a more modular version of the existing IFS Fortran code.

The OOPS project will make the data assimilation code more flexible. It will also make it more efficient. In the short term, because the new code structure can handle multiple resolutions within the same executable, 4D-Var can be run in one task. This will save start-up time and I/O that is required to communicate between the existing 5 or 7 jobs that make 4D-Var at the moment (see Fig. 5). OOPS is also required for the implementation of long window 4D-Var with sub-windows running in parallel.

Finally, the productivity of scientists involved in the evolution of complex NWP systems like the IFS should also be taken into account. The learning curve for new scientists contributing to the IFS is very steep and even experienced IFS developers face enormous technical challenges in their everyday activities. The new code structure will allow scientists to focus on their area of expertise by encapsulating the different aspects of the problem and making them more independent. Similar concerns apply to other NWP centres.

8.2. A hybrid system and a longer window

Despite the uncertainties, certain trends seem clear. We will undoubtedly retain and extend the existing ensemble component in any future ECMWF Data Assimilation system. The ensemble is able to estimate analysis uncertainty, which is used both to provide initial conditions for the Ensemble Prediction System and to provide the analysis system with flow-dependent estimates of background error covariance. The ensemble component may well continue to be an ensemble of 4D-Var analyses. However, it is also possible that an Ensemble Kalman Filter may provide sufficiently good error estimates in a more cost effective way.

For the next decade there will continue to be a need for a high-resolution deterministic analysis. We believe that this will continue to be provided by 4D-Var at ECMWF. This analysis will be tightly coupled with the ensemble component, which will provide error estimates from which 4D-Var can generate flow-dependent background error covariances.

There is a strong theoretical argument, backed up by experimental results in simplified systems, in favour of increasing the length of the 4D-Var analysis window. An extension from the current 12-hour window to 24 hours is planned for next year, and we believe further extensions to around 2-3 days will prove beneficial in the future. This is scientifically challenging and will require the development of a significantly more advanced representation of model error covariance. It is likely a long-window EDA/EnKF will be needed to provide flow-dependent model error estimates. Bias estimation for both model and observations is a related important component of the assimilation system. It should be noted that the extension of the 4D-Var analysis window will be of even greater benefit for atmospheric reanalyses. To be computationally viable, this will require algorithmic changes, as discussed in Section 7.2, which enables parallelisation of weak-constraints 4D-Var in the time dimension.

8.3. Optimizing the observation processing

The current operational 4D-Var at ECMWF consists of a backbone of cycling late-cutoff 12-hour 4D-Var analyses, from which simplified “early delivery” 6-hour 4D-Var analyses are spun off.

The backbone late-cutoff suite waits between 5 and 17 hours for the relevant observations to arrive. This ensures that more than 99% of all measurements are available for the analysis.

The early-delivery analyses provide the initial conditions for the operational forecast. These analyses only wait between 1 and 7 hours for the relevant observations to arrive. The early delivery analyses are also much cheaper to run than the late-cutoff analyses, so they are available more than 10 hours earlier, with an approximately 6 hours worse quality. Currently, these early-delivery suites are run twice daily, and use initial conditions from the late cutoff suite. Two additional suites of a similar kind are also run to provide boundary conditions to some of our Member States. These assimilations enable us to benefit from the large fraction of observations that arrive quickly after being measured. With the presently used data cut-off, more than 85% of the available observations are assimilated in these early delivery runs, providing more timely forecasts. As there is continued pressure to deliver forecasts to our Member States earlier, the early delivery analyses need simplification and streamlining. A system is now being designed where the observations will be continuously prepared for the analysis as the observations arrive at ECMWF. This "Continuous Observation Processing Environment" (COPE) is intended to improve the efficiency of the operational suites. Currently, most of the observation related activities are performed in the time critical path. The ongoing increase in the number of observations and increase in model resolution will stretch our ability to deliver the forecasts and associated products on time. Moreover, the observation processing tasks are those parts of the operational suite that are least scalable, making it an even bigger problem on future computer architectures. COPE will enable to move most of the screening tasks out of the time critical window. These tasks are schematically defined in the COPE box in Fig. 5, and in the list of cope activities described in Fig. 6. Fig. 5 also defines the OOPS activities schematically.

An important side issue is that the tight operational schedule leaves very little time to handle inevitable observation related operational failures, leading occasionally to delays in forecast dissemination. Due to the current design of the observation processing, monitoring and diagnostics can only be done after the whole analysis cycle is completed. This makes it difficult to act when observation related problems occur. COPE will enhance the early detection and handling of observation anomalies that could cause failures in the operational suite and dissemination delays.

Using ECMWF as a typical global NWP centre, it is useful to summarize, a possible scenario for the atmospheric data assimilation system at ECMWF by the end of this decade:

- A late cutoff 48 hour high resolution deterministic weak-constraint 4D-Var. This system would use advanced algorithms to split the time windows into three hour sub-windows that are processed in parallel. This would significantly increase the scalability of the overall system. The algorithm would also encompass overlapping windows, with one analysis performed every 12 hours.

- Tightly harnessed to the deterministic 4D-Var, an Ensemble of Data Assimilations (EDA), based on long window weak constraint 4D-Var, to provide both flow-dependent background error and model error estimates.

- Twice or more per day, an early delivery and boundary condition, short-window, high resolution 4D-VAR run. This early delivery run will be heavily streamlined so that the non-scalable work to be performed during the critical path remains minimal.

Variations around this baseline scenario will depend upon technical and scientific considerations, and available HPC resources. Indeed, the benefit of extending the window will have to be balanced against increasing the resolution of the inner loop or increasing the number of outer loops. Also, if EnKF techniques mature enough to fulfil both the required provision of high quality covariances for 4D-Var and initial perturbations for EPS, they will be considered for the ensemble component of the assimilation suite.

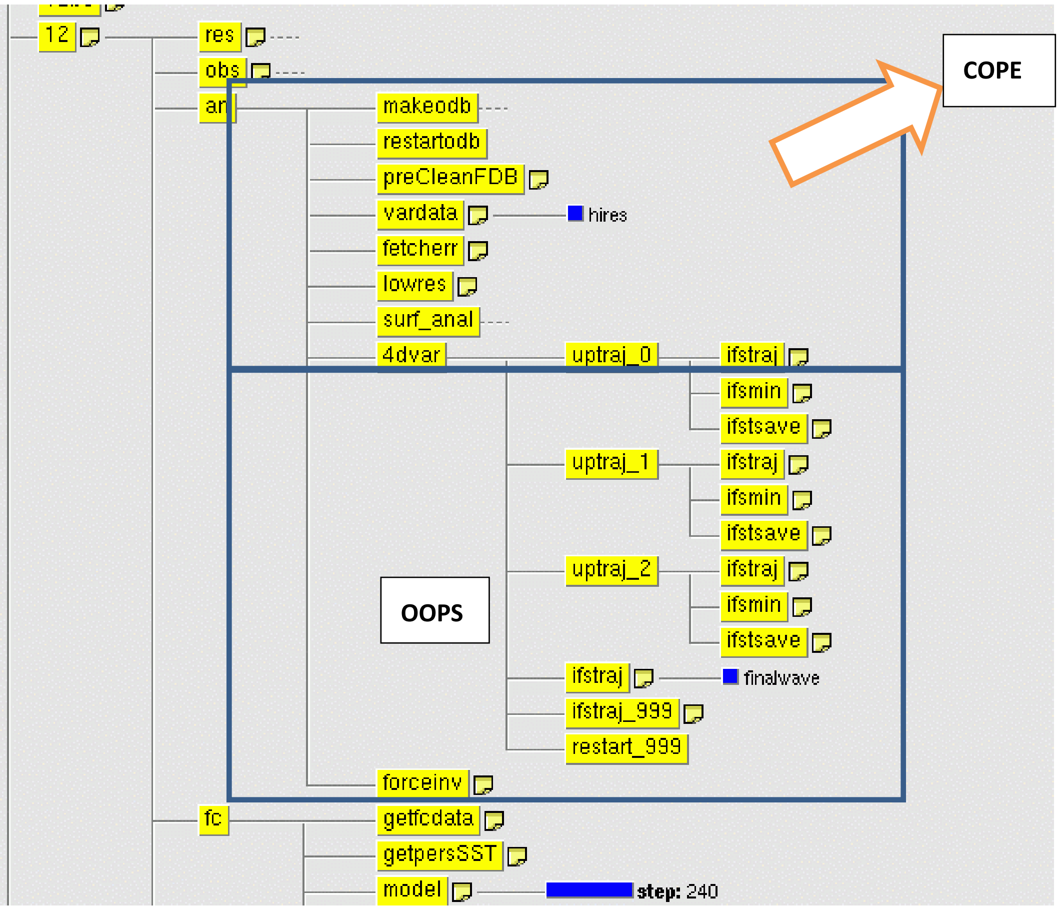

Figure 5: A schematic representation of the data assimilation tasks handled by COPE (Continuous Observation Processing Environment) and OOPS (Object Oriented Programming System) in ECMWF's next generation data assimilation system.

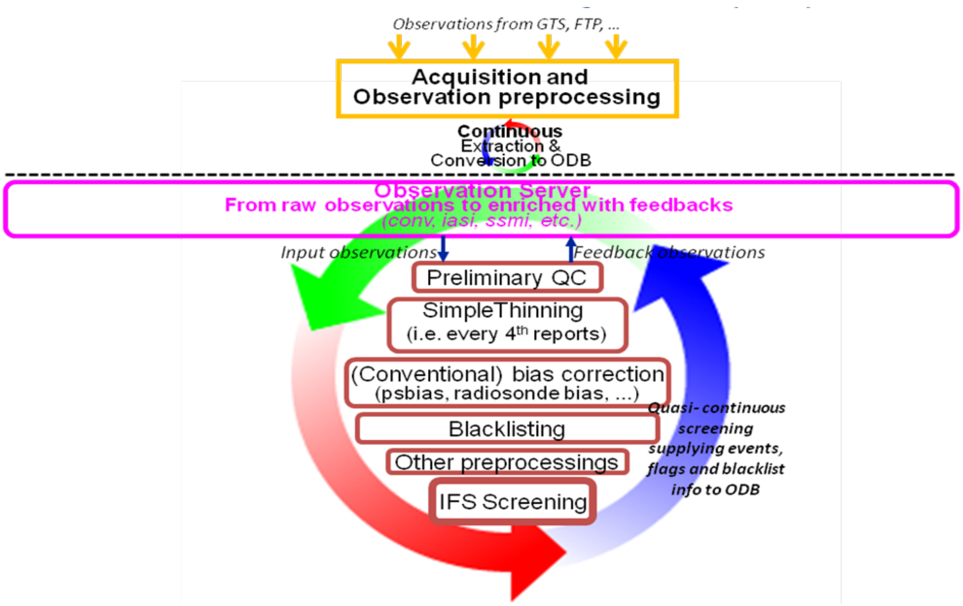

Figure 6: A schematic diagram describing the main tasks handled by the COPE (Continuous Observation Processing Environment) that is under development at ECMWF. Courtesy Anne Fouilloux.

9. Conclusions

The high performance computer resources at ECMWF and other NWP centres have increased exponentially during the last 30 years, with a doubling time of approximately two years. This rate of progress is expected to continue during the next decade. The improved performance will primarily come from more cores on each chip, rather than increased individual core performance. So during the next decade we expect HPC systems with of the one million cores to be common at global NWP centres.

Using the ECMWF data assimilation system as a typical example of a state-of-the-art NWP system, it is shown that there are scalability problems for low resolution forecast models and 4D-Var analysis on today's high performance computer with around 10,000 cores. To be able to use HPCs with an expected increase in the number of cores by a factor of 100 will lead to significant challenges for the data assimilation strategies applied.

ECMWF has invested considerably in 4D-Var over more than two decades and this investment has incontestably paid off, contributing significantly to maintaining ECMWF as the world leader in global NWP. Scientific developments in data assimilation at ECMWF have demonstrated that 4D-Var provides the flexibility required to deal with a large variety of observations, to implement advanced bias correction and quality control, and deal with non-Gaussian error distributions (e.g. cloud and rain assimilation).

We believe that there is still much benefit of 4D-Var to be exploited. In particular, long-window weak-constraint 4D-Var offers considerable opportunities that should be pursued (Fisher *et al.* 2012, Fisher and Auvinen 2012 in these proceedings). These benefits will serve the operational data assimilation suite and the reanalysis requirements.

It is also very clear that ensemble-based data assimilation methods, already used today, in the future will be even more central to the assimilation system and to the ensemble prediction system: for estimation of analysis uncertainty, provision and cycling of flow-dependent error covariances, and provision of initial perturbations. The baseline approach at ECMWF is to further exploit the ensemble of 4D-Var that has already been developed and operationally implemented, while in parallel continuing the development of an Ensemble Kalman filter which could offer in the future an attractive alternative option to the EDA.

The proposed scenario combining a long-window weak-constraint 4D-Var with an EDA is realistic and ambitious. It also presents a number of technical and scientific challenges that will require attention: a profound modification of some algorithms to substantially increase scalability, a reorganisation of the IFS data assimilation code that allows much more flexibility, extensive research in the area of model error specification, and redesign of the time critical part of the operational suite. Addressing these challenges successfully will allow a data assimilation system that can run efficiently on the computers available during the next decade.

Beyond 2020, the HPC architectures could be such that data assimilation algorithms may require even more radical changes to obtain scalability. As discussed in Section 7, the problem will have to be considered from a general point of view, including the scalability of the high resolution forecast model itself.

Acknowledgements

It is acknowledged that parts of the paper are based on an internal ECMWF “Data Assimilation Vision” paper written in September 2011 for the Scientific Advisory Committee by the author together with ECMWF colleagues (Jean-Noël Thépaut, Massimo Bonavita, Michael Fisher, Mats Hamrud and Yannick Trémolet). Further acknowledgements for fruitful discussions and comments on the manuscripts go to Deborah Salmond, George Mozdzynski and Erland Källén. Finally, special thanks to Deborah Salmond and Anne Fouilloux for allowing me to use some of their figures for this paper. All the acknowledged people are working at ECMWF.

References

Anderson, J. L., 2001: An Ensemble Adjustment Kalman Filter for Data Assimilation. *Mon. Wea. Rev.*, **129**, 2884-2903.

Andersson, E. and J.-N. Thépaut, 2010: Assimilation of Operational Data. In *Data Assimilation: Making Sense of Observations*, eds. W.A. Lahoz, B. Khatattov, R. Ménard, Springer, pp 283-299.

Barker, D.M. and A.M. Clayton, 2012: Hybrid variational-ensemble data assimilation. *Proc. ECMWF Seminar on Data assimilation for atmosphere and ocean*, 6-9 September 2011, 31-44. (Available from: http://www.ecmwf.int/publications/)

Bonavita, M.,2012: Ensemble of data assimilations and uncertainty estimation. *Proc. ECMWF Seminar on Data assimilation for atmosphere and ocean*, 6-9 September 2011, 135-160. (Available from: http://www.ecmwf.int/publications/)

Bonavita M., L. Raynaud and L. Isaksen, 2011: Estimating background-error variances with the ECMWF Ensemble of Data Assimilations system: some effects of ensemble size and day-to-day variability. *Quart. J. Roy. Meteor. Soc.*, **137**, 423–434.

Buehner, M., P. L. Houtekamer, C. Charette, H. L. Mitchell and B. He, 2010a: Intercomparison of Variational Data Assimilation and the Ensemble Kalman Filter for Global Deterministic NWP. Part I: Description and Single-Observation Experiments. *Mon. Wea. Rev.*, **138**, 1550-1566

Buehner, M., P. L. Houtekamer, C. Charette, H. L. Mitchell, and B. He, 2010b: Intercomparison of Variational Data Assimilation and the Ensemble Kalman Filter for Global Deterministic NWP. Part II: One-Month Experiments with Real Observations. *Mon. Wea. Rev.*, **138**, 1567–1586.

Burgers, G., P. J. van Leeuwen, and G. Evensen, 1998: On the analysis scheme in the Ensemble Kalman Filter. *Mon. Wea. Rev.*, **126**, 1719-1724.

Courtier, P. and Talagrand, O. 1987. Variational assimilation ofmeteorological observations with the adjoint vorticity equation II: numerical results. *Quart. J. Roy. Meteor. Soc.* **113**, 1329–1347.

Courtier, P., Thépaut, J. N. and Hollingsworth, A. 1994. A strategy for operational implementation of 4DVar using an incremental approach. *Quart. J. Roy. Meteor. Soc.* **120**, 1367–1387.

Evensen, G., 1994: Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *J. Geophys. Res.*, **99** (C5), 10143-10162.

Fisher, M., 2003: Background error covariance modelling. *Proceedings of the ECMWF Seminar on recent developments in data assimilation for atmosphere and ocean*, ECMWF, pages 45–63. (Available from: http://www.ecmwf.int/publications/)

Fisher, M. and H. Auvinen, 2012: Long window 4D-Var. *Proceedings of the ECMWF Seminar on Data assimilation for atmosphere and ocean*, 6-9 September 2011, 189-202. (Available from: http://www.ecmwf.int/publications/)

Fisher, M., Y. Tremolet, H. Auvinen, D. Tan and P. Poli, 2011: Weak-constraint and long-window 4DVAR. *ECMWF Technical Memorandum* 636. (Available from: http://www.ecmwf.int/publications/)

Fisher, M., M. Leutbecher and G. A. Kelly, 2005: On the equivalence between Kalman smoothing and weak-constraint four-dimensional variational data assimilation, *Quart. J. Roy. Meteor. Soc.*, **131**, 3235–3246.

Hamill, T. M. and C. Snyder, 2000: A Hybrid Ensemble Kalman Filter–3D Variational Analysis Scheme. *Mon. Wea. Rev.*, **128**, 2905–2919.

Houtekamer, P. L. and H. L. Mitchell, 1998: Data Assimilation Using an Ensemble Kalman Filter Technique. *Mon. Wea. Rev*., **126**, 796-811.

Hunt, B. R., E. J. Kostelich and I. Szunyogh, 2007: Efficient data assimilation for spatiotemporal chaos: a Local Ensemble Transform Kalman Filter. *Physica* D, **230**, 112-126.

Isaksen, L., M. Bonavita, R. Buizza, M. Fisher, J. Haseler, M. Leutbecher and Laure Raynaud, 2010: Ensemble of data assimilations at ECMWF. *ECMWF Technical Memorandum* 636. (Available from: http://www.ecmwf.int/publications/)

Le Dimet, F. X. and Talagrand, O., 1986: Variational algorithms for analysis and assimilation of meteorological observations: theoretical aspects. *Tellus* **38A**, 97–110.

Lewis, J. M. and Derber, J. C., 1985: The use of the adjoint equation to solve a variational adjustment problem with advective constraints. *Tellus* **37A**, 309–322.

Lorenc, A. C., 2003: The potential of the ensemble Kalman filter for NWP— A Comparison with 4D-Var. *Quart. J. Roy. Meteor. Soc*., **129**, 3183–3203.

Rabier, F., H. Järvinen, E. Klinker, J.-F. Mahfouf, and A. Simmons, 2000: The ECMWF operational implementation of four-dimensional variational assimilation. Part I: Experimental results with simplified physics. *Quart. J. Roy. Meteor. Soc*., **126,** 1143–1170.

Simmons, A. J., and A. Hollingsworth, 2002: Some aspects of the improvement in skill of numerical weather prediction. *Quart. J. Roy. Meteor. Soc*., **128**, 647–677.

Thépaut, J.-N., P. Courtier, G. Belaud and G. Lemaître, 1996: Dynamical structure functions in a four-dimensional variational assimilation. *Quart. J. Roy. Meteor. Soc*., **122**, 535-561.

Thépaut, J.-N. and E. Andersson, 2010: The Global Observing System. In *Data Assimilation: Making Sense of Observations*, eds. W.A. Lahoz, B. Khatattov, R. Ménard, Springer, pp 263-281.

Whitaker, J. S. and T. M. Hamill, 2002: Ensemble Data Assimilation without Perturbed Observations. *Mon. Wea. Rev*., **130**, 1913-1924.

Whitaker, J.S., 2012: Developments in ensemble data assimilation. *Proceedings of the Seminar on Data assimilation for atmosphere and ocean*, 6-9 September 2011, 17-30. (Available from: http://www.ecmwf.int/publications/)