Interactive Weather Simulation and Visualization on a Display Wall with Many-Core Compute Nodes

Bård Fjukstad, Tor-Magne Stien Hagen, Daniel Stødle, Phuong Hoai Ha, John Markus Bjørndalen, and Otto Anshus

Dept. of Computer Science, University of Tromsø, N-9037 Tromsø, Norway
{baard.fjukstad,phuong.hoai.ha,john.markus.bjorndalen,otto.anshus}@uit.no,
{tormsh,daniels}@cs.uit.no

Abstract. Numerical Weather Prediction models (NWP) used for operational weather forecasting are typically run at predetermined times at a predetermined resolution and a fixed geographical region. The period between each run is a function of waiting for observational data and the availability of compute resources. The resolution is a function of the geographical region, the available processing power and operational forecasting time constraints. The geographical region is defined by being a region with known need or interest for forecasts. These characteristics make it hard to interactively produce and visualize on-demand high-resolution forecasts for a small and arbitrarily located region. This paper documents a system achieving this, using a high-resolution tiled 22 mega pixel display wall, a 16 node PC cluster and a HP BL 460c blade server with two quad core processors. We document the performance characteristics experimentally. The results show that using 10 km resolution background data, the system produces a 6 hour forecast for a 117 x 123 km small region with 3 km resolution, in 3 minutes. Visualizing the forecast takes between 3 - 75 seconds. An informal survey among operational forecasters indicate that the majority is willing to wait up to 3 minutes for higher resolution forecasts. This paper identifies and documents some of the bottlenecks and computational challenges created by combining interactivity and traditional batch oriented computing. The main bottlenecks in the system are identified as the execution time of the NWP and the preparation of data for visualization.

Keywords: Interactive Numerical Weather Model, WRF, Tiled Display Walls, Live Data Sets, On-Demand Computation.

1 Introduction

Numerical Weather Prediction models for use in weather forecasting centers are often computed for a fixed static region at a fixed resolution. One example is the very high-resolution turbulence forecasting system called SIMRA [5], in daily operational use by the Norwegian Meteorological Institute [10]. The SIMRA system uses the wind field from a coarser model to estimate the detailed current

K. Jónasson (Ed.): PARA 2010, Part I, LNCS 7133, pp. 142–151, 2012.

turbulence levels around specific locations. Available compute resources limit the number of locations this model can make available to the weather forecaster. This reduces the number of airports where the forecaster can assess the current level of risk for severe turbulence. Therefore, only places of interest with a previously known high level of risk have pre-computed models available. At any given day, this may or may not be the actual trouble spots. Results from NWPs are interactively visualized on a typical PC display. Standard screen resolution is in the order of 1200x1024 pixels.

For any given area and selected parameter the visualization software often has to render the available data using a subset of the original data points. One example is viewing wind fields. These are often visualized using small arrows at each data point, which is not possible to do for a large area on a small screen without either reducing the readability of the plot or displaying only a subset of the available data plots. Using very large high-resolution displays gives the user the option of both viewing large areas, and at the same time all available data points. This has previously been shown to be advantageous using standard visualization software [6].

This paper presents WallWeather, an interactive system and approach for visualizing state-of-the-art numerical meteorological models using a wall-sized high-resolution tiled display [11]. The idea is that the user does not know a priori where high-resolution forecasts would be most useful, and that the user based on available coarser models, select the area and desired resolution. Initially the resolution is a function of the available background meteorological data. The user can select a region of interest by zooming in on that region and have NWP done on-demand for the selected area at the desired resolution. The many-core compute clusters will provide the on-demand weather forecasts for the selected areas.

The ability to select smaller regions of interest with high-resolution forecasts, combined with a display wall supported by on-demand computing, enables a close to interactive experience for the user, at resolutions orders of magnitudes larger than regular desktop displays.

WallWeather is a platform for further experimenting with various ways to divide the total workload and also to investigate the many bottlenecks such complex combined systems present. WallWeather is also a system that both generates and visualizes datasets on demand, as opposed to existing batch-oriented systems where datasets are created at pre-defined times.

This work is based on an idealized use case shown in Table 1.

2 The Numerical Weather Prediction model

In this paper the WRF NWP model [2] is used. WRF is currently a very popular research model for high-resolution weather forecasting systems. WRF is available in numerous settings and is extensively used in many meteorological research and operational centers [3].

A simpler downscaling of the wind field for each time-step, like the SIMRA system, may reduce the workload but does not provide the forecaster with the fully integrated set of parameters available from the NWP models.

To simplify the prototype, the resolution of the WRF model is limited to a to a fixed set of discrete resolutions. This is a necessity given the available topographical, meteorological and environmental data. NWP models are usually downscaled by a factor of 3-5, so that when using 50 km background data, 9.9, 3.3 and 1.1 km resolution models would be the natural levels for stepwise increase of the resolution for the NWP model. Even higher resolution models are possible with access to high-resolution background data. To ensure numerical stability of the model with the steep topography in the area of interest, the time step of the model must be reduced more than recommended.

In the prototype, a static set of background meteorological data from a date with locally severe weather in the area of interest where chosen. An independent start analysis using actual observations is not used in the system. For the small areas in which the NWP model is run, normally only a few actual observations would be available, and a long time-period is needed to include the necessary observation error statistics for the analysis. The prototype still incurs most of the workload that an operational system would require.

Figure 1 shows a possible scenario with several trouble spots. For Areas A, B and F the requested resolutions are large enough for running WRF directly using the background meteorological data. Areas D and E are requested with a higher resolution and require an intermediate step, area C, to be computed. Once area C is computed, all higher resolution areas that fall within C require no extra intermediate computations. The effect of these scenarios on the perceived latency for the user is shown in Figure 2.

3 Experimental Platform

3.1 The Display Wall

The display wall [12] consists of 28 projectors driven by 28 computers arranged in a 7x4 grid yielding a total resolution of 7168 x 3072 pixels. When using WallGlobe the user perceives the display wall as one single coherent display.

Table 1. Idealized use case

1	The forecaster browses a coarse resolution model for possible trouble spots.
2	The forecaster zooms in to view details and triggers a new NWP run.
3	The forecaster views the results from a high-resolution model for the specific area.
4	The forecaster pans the view to include nearby trouble spots, or zooms out and focuses on a new area.

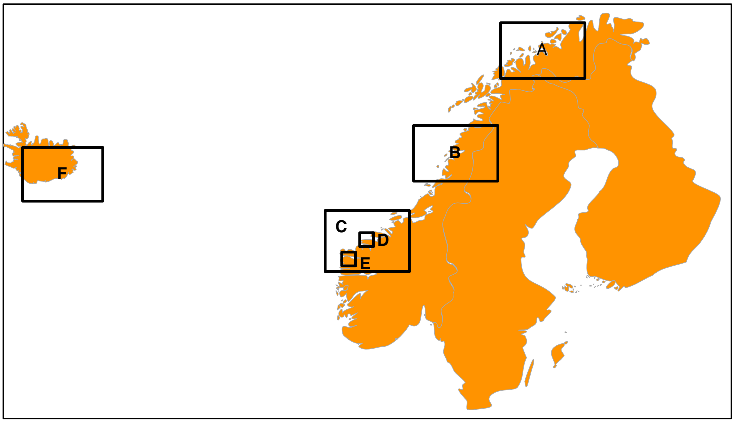

Fig. 1. Cases A, B, D, E and F are the trouble spots in this situation. A background model is assumed available in the whole area of interest.

Zoom and pan is implemented using a touch-free interface [11]. Figure 3 shows two persons interacting with the display wall and the WallGlobe visualization system.

3.2 WallScope

The work presented in this paper is implemented as part of the WallScope [7] visualization and computation platform. WallScope uses a Live Data Set (LDS) architecture. Visualization clients run on each computer in the display wall cluster, all synchronized by a separate state server. Each client requests data from the LDS, which initiates local or remote computations to satisfy the request. LDS may also return a cached copy if the computation has been performed earlier. The LDS architecture is shown in Figure 4. The architecture separates visualization from data management, and data management from the data producer. For interactive visualization of weather forecasts the WallScope system is extended with an on-demand simulation and visualization backend using the WRF NWP model.

The visualization system used in this paper is WallGlobe, a system for visualizing the Earth by combining data from different compute resources in the WallScope system. WallGlobe requests images of size 512x512 pixels from the LDS, which are used as part of the final rendering. Each tile in the display wall requests the images it needs to complete the visualization. Until the maximum zoom level is reached, each tile will at all times use images with a resolution as high as or higher than the resolution of the local resolution. The tile resolution is 1024x768 pixels. The worst-case scenario is that a tile has just reached a new zoom level and that the available images are offset as illustrated in Figure 5. In this case, 25 images have to be requested and retrieved from the LDS.

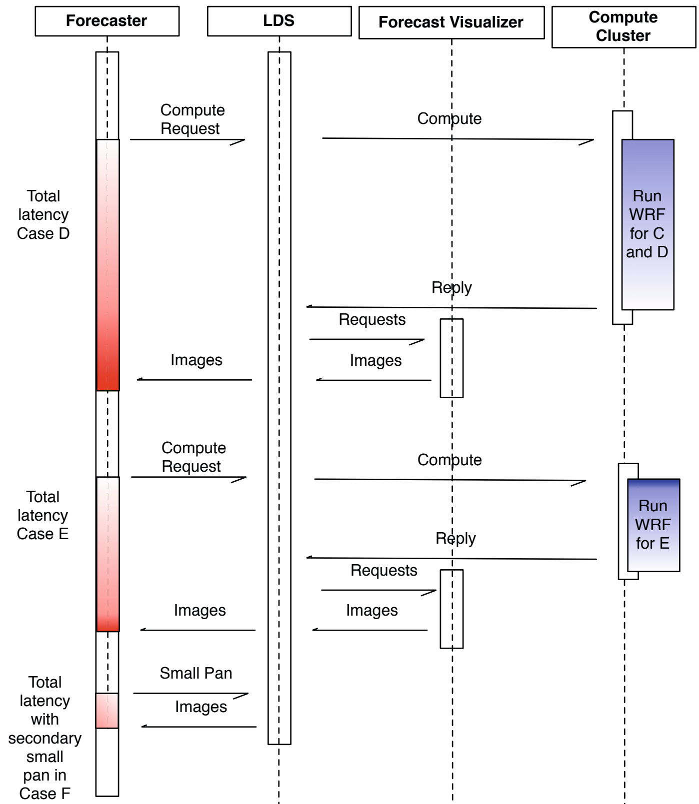

Fig. 2. Three different cases are shown. Case D with no intermediate level available, Case E, with the C area available, and Case F where the model is first run on a slightly larger area than requested so that minor Pans does not trigger a full generation of a new area.

3.3 Compute Clusters

The prototype utilizes two clusters. One is a local 32 node 3.2 GHz Pentium 4 cluster, "Rocks"; the other is a 704 node 1408 CPU 5632 core "Stallo" [4] high-performance cluster. The Rocks cluster is a dedicated cluster and jobs submitted are immediately executed. Stallo uses a standard batch job queuing system and is therefore not very well suited for interactive use. An express queue with limitations on the number of cores available for each job can be used for a near real-time interactive use.

3.4 Network

Every node in the display wall are interconnected using gigabit Ethernet. The display wall is connected to the compute clusters over a gigabit Ethernet link.

Fig. 3. Using the display wall and the WallGlobe visualization system

4 Experiments

4.1 Methodology

To evaluate the WallWeather system, two experiments were conducted.

In experiment one, a small informal survey of the operational forecasters at the Norwegian Meteorological Institute in Tromsø, Norway, was conducted, to establish a limit on how long a forecaster would be willing to wait for higher resolution forecasts for a selected area. 14 out of 18 possible participants responded to the questionnaire.

In experiment two, the actual total latency of the system was measured, using both compute clusters. These experiments showed the effect of running the data producing services on a multi node, multi-core platform. The WRF model is expected to scale well and perform well on these platforms [8].

The ECMWF ERA-Interim data used in this study have been obtained from the ECMWF data server [1]. A specific date with severe weather in the area were used for this study. The data has a spatial resolution of around 50 km. The model was run for a 6 hour forecast for a small 39x41 grid, 28 vertical levels with 9.9 km resolution using a time step of 30 sec. The timestep were shortened due to the very steep topography in the model area and to keep the model numerically stable.

The perceived latency after triggering a NWP model run, depends on the availability of the background data the model needs at the requested resolution. Figure 2 illustrate this. As explained in chapter 2, a run of the WRF model may require several steps with increasing resolution before the requested resolution is computed, figure 2 illustrates this. In the top part of Figure 2 area C has to be computed first, and then area D. If the next requested area falls inside the

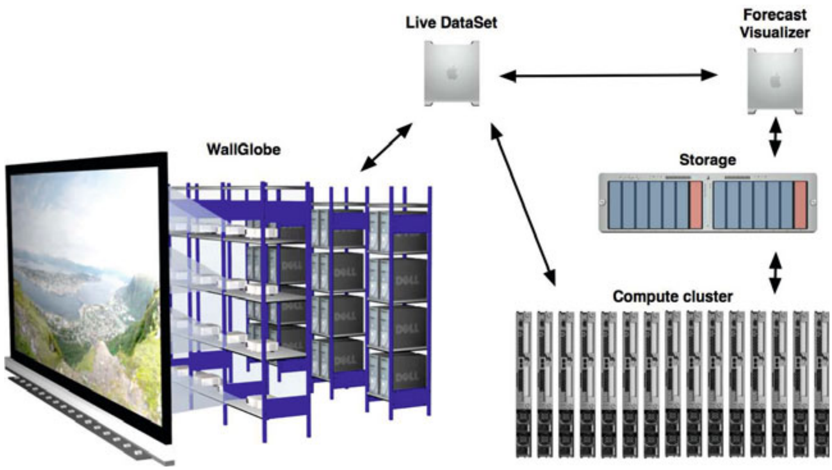

Fig. 4. Architecture with the communication paths indicated

already computed area of C, then the request can be answered by running the model only for the new area E. If the request is only for a small pan within the areas already computed and visualized, then the request will be satisfied from the LDS cache, as shown in the lowest part of Figure 2.

4.2 Results

Table 2 shows the results of experiment one. Almost 60% of the forecasters were willing to wait more than one minute for higher resolution forecasts. Less than 30% would wait more than 5 minutes.

Table 3 show the results of experiment two. For the actual computation, the times are in separate columns. Transferring the resulting data files and retrieving one parameter from the forecast visualizer is identical for both, and are therefore merged into one column.

Table 2. How long a forecaster is willing to wait for higher resolution forecasts

Time	Count
5-14 sec	2
15-44 sec	0
45 sec - 1 min	3
1-2 min	2
3 - 5 min	3
5 - 10 min	2
More than 10 min	2
Total N	14

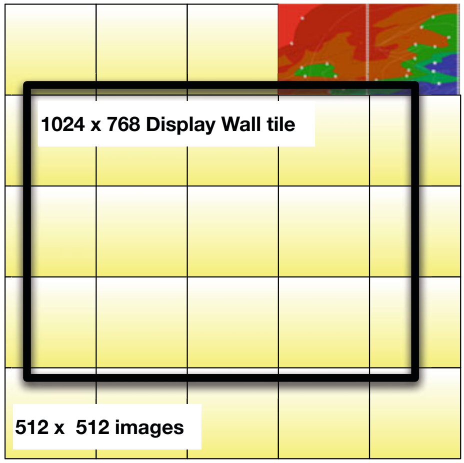

Fig. 5. Illustration of the parts needed for one tile of the display wall. Each image requested from the LDS is 512x512 pixels. Each display wall tile has a resolution of 1024x768 pixels. The WallGlobe will always use images with higher or equal resolution to the tile's resolution.

Table 3. Average run-times Case E using the Stallo cluster using 8 cores on 1 node and the Rocks cluster using 1 core on 16 nodes. Models domain is 39 x 41, 28 vertical levels, 9.9 km resolution, 6 hour forecast with 30 sec time steps.

Task	**Time on "Stallo"**	**Time on "Rocks"**
Running pre-processing on cluster front-end	13 sec	13 sec
Running the WRF model	56 sec	174 sec
Transferring result file to visualization host	0.4 sec	
Retrieving one parameter for visualization	3 sec	

5 Discussion

Table 3 indicates that the largest bottleneck is the execution of the WRF forecast model. When the numerical forecast model is completed, the next bottleneck is the generation of visualization data from the model output. The time listed for visualization in Table 3 is for one single image of size 512x512 pixels.

The system was not intended as a system for delivering high-resolution numerical forecasts each day or at a specific schedule for large areas. For such use the traditional batch oriented systems would be better. The system was created to provide additional high-resolution forecasts for smaller user-selected areas, based on existing coarser resolution NWP model data available to the forecaster.

The WallWeather system provides a practically interactive system even if the latency times for the user are longer than some operational use will tolerate. The system has the ability to display high-resolution visualizations from user defined areas using on demand numerical weather prediction models. This enables possible new insight into relevant meteorological problems, as well as better and more accurate forecasts.

One major bottleneck is the use of one single node for forecast visualization. When each image used by the LDS would come from a single visualizing node, all images needed for covering one single tile on the display wall would take up to 75 sec to retrieve. Since the LDS uses caches, most images that are shared with other tiles on the display wall would be retrieved much faster.

One observation is that using fixed grid sizes with variable spatial resolution in the NWP model, the workload on the computational components varies only with the spatial resolution and time steps needed in the model.

Based on experiment one the latency of the system falls within the acceptable waiting time for the forecasters.

6 Related Work

The triggered WRF forecasts part of the LEAD project [13], presents a similar use case to the WallWeather system. Higher resolution WRF model runs were generated automatically using positions of known severe weather systems from the NOAA NWS news feed. By changing the workflow brokering on a powerful computation cluster to increase the scheduling priority of the model run, timely forecasts were provided. The project identified several problems regarding reliability problems on the compute cluster and the effect on the lack of provided forecasts. No end-user latencies were reported.

7 Conclusions

This paper has presented a prototype of an interactive numerical weather model system, used for on-demand high-resolution visualization on a high-resolution display wall. New numerical weather prediction models are relatively easy to set up with a large range in resolutions, limited mostly by available environmental data, and available computing resources. The experiments conducted on the WallWeather system demonstrates that interactive running of NWPs on high-resolution display walls is coming closer to a practical solution for operational weather forecasting.

8 Future Work

Using GPUs in WRF may improve the runtime significantly [9]. Utilizing GPUs may also improve the visualization performance.

There are various obvious ways to speed up the forecast visualization part of the system. Implementing a distributed system using a compute cluster with single forecast visualization node on each compute node is one possible solution. Depending on the number of nodes, this may reduce the forecast visualization delay to 3 seconds.

Acknowledgements. The authors would like to thank the technical staff at the Computer Science department, University of Tromsø, and the HPC group at the Computer Center, University of Tromsø.

This work has been supported by the Norwegian Research Council, projects No. 159936/V30, SHARE - A Distributed Shared Virtual Desktop for Simple, Scalable and Robust Resource Sharing across Computers, Storage and Display Devices, and No. 155550/420 - Display Wall with Compute Cluster.

References

1. ECMWF, European Center for Medium Range Weather Forecasting, http://dataportal.ecmwf.int/data/d/interim_daily/
2. The WRF model, http://www.wrf-model.org
3. WRF Real Time Forecasting, http://wrf-model.org/plots/wrfrealtime.php
4. The Stallo cluster (2010), http://docs.notur.no/uit/stallo_documentation/user_guide/key-numbers-about-stallo
5. Eidsvik, K., Holstad, A., Lie, I., Utnes, T.: A prediction system for local wind variations in mountainous terrain. Boundary-Layer Meteorology 112(3), 557–586 (2004), http://dx.doi.org/10.1023/B:BOUN.0000030561.25252.9e
6. Fjukstad, B., Anshus, O., Bjørndalen, J.: High resolution numerical models on a Display Wall. In: The 7th Annual Meeting of the European Meteorological Society (EMS) and the 8th European Conference on Applications of Meteorology (2007)
7. Hagen, T.M.S., Stødle, D., Anshus, O.: On-demand high-performance visualization of spatial data on high-resolution tiled display walls. In: Proceedings of the International Conference on Information Visualization Theory and Applications, pp. 112–119 (2010)
8. Michalakes, J.: (2003), http://www.cisl.ucar.edu/dir/CAS2K3/CAS2K3
9. Michalakes, J., Vachharajani, M.: Gpu acceleration of numerical weather prediction. In: IEEE International Symposium on Parallel and Distributed Processing, IPDPS 2008, pp. 1–7 (April 2008)
10. Midtbø, K., Bremnes, J.B., Homleid, M., Ødegaard, V.: Verification of wind forecasts for the airports hammerfest, honningsvåg, sandnessjøen, ørsta, værnes, sandane and narvik (2008), http://met.no/Forskning/Publikasjoner/metno_report/2008/filestore/report02_20081.pdf
11. Stødle, D., Troyanskaya, O., Li, K., Anshus, O.J.: Tech-note: Device-Free Interaction Spaces. In: 3DUI 2009: Proceedings of the IEEE Symposium on 3D User Interfaces, pp. 39–42 (March 2009)
12. Wallace, G., Anshus, O.J., Bi, P., Chen, H., Chen, Y., Clark, D., Cook, P., Finkelstein, A., Funkhouser, T., Gupta, A., Hibbs, M., Li, K., Liu, Z., Samanta, R., Sukthankar, R., Troyanskaya, O.: Tools and applications for large-scale display walls. IEEE Comput. Graph. Appl. 25(4), 24–33 (2005)
13. Wilhelmson, R., Alameda, J., Rossi, A., Hampton, S., Jewett, B., Weber, D., Thomas, K., Wang, Y., Droegemeier, K.: Lead: Automatic triggering of high resolution forecasts in response to severe weather indications from the noaa storm prediction center. Urbana 51, 61801