# HPC Meets Real-Time Data: Interactive Supercomputing for Urgent Decision Making BoF summary

This BoF was held at SC18 on Thursday, November 15th, 2018 from 12:15 - 13:15 and the general idea was that big data and digitalization have created lots of new opportunities, that have the potential to move HPC well beyond traditional computational workloads. New technologies, such as social and sensor networks, generate very large volumes of data and a major challenge is how to best process this to turn it into valuable collateral. This issue is especially challenging when there are specific time constraints imposed on the data and the capability provided by HPC is potentially very useful here.

However, the potential of HPC goes much further than this, because it is now possible to combine simulation with this high velocity data and live analytics to aid in real-time decision making. Whilst there is considerable existing work by the community in using HPC to process data *offline*, the idea of processing data in (or near) real-time and then using this in combination with simulations to aid in urgent decision making, such as how to tackle a forest fire, is far less mature. It is our opinion that interactive supercomputing that fuses HPC with real-time data for urgent decision making is likely to be of great benefit when tackling solving societal, economic and environmental problems in the future. To some small extent the community is already starting to consider this and an example of that is the plan by the EuroHPC project to reserve a number of cycles of their exascale demonstrator to urgent decision making.

But this is not simple and there are numerous challenges all the way up the stack that must be addressed. These range from hard-core technical challenges, such as how to support the dynamic nature of resource usage, all the way to human computer interaction, such as how to best present the data in a way in which decision makers can best leverage this to make correct decisions each and every time.

In this BoF we focussed on three general areas that tackled challenges at different levels, these were

- Interactive supercomputing and applications: Top down, the use of interactive supercomputing by different applications and to understand what other applications there are out there already interested in this space.
- Interactive data exploration challenges: Experiences around the interactive exploration of data and challenges arising from this. Also exploring members of the HPC community who would like to do this but are not currently.
- Technical challenges arising from fusing HPC with real-time data: Driven bottom up, how are we going to actually fuse HPC and real-time data so people can take advantage of this in their codes? What challenges arise from the assumptions made by current generation HPC technologies and how might we address this?

After an initial presentation, the BoF split into three sub-groups, each of which was holding discussions focussing on one of these areas. At the end of the session we summarised the findings of each sub-group and this document provides further details and explorations of the themes discussed.

## Group 1: Interactive supercomputing and applications

The leaders of this group had three general aims

1. To formulate the definition and requirements for interactive supercomputing.
2. Define possible use cases of such a technology.
3. Outline any technical issues foreseen by the community at this early stage.

It was felt that formulating the definition of interactive supercomputing in this context is really important as terminology can easily become overloaded and the meaning obscured. We believe that this is especially likely in novel, high impact areas because people naturally want to tie their own research activities to that theme, which can then distort the key purpose of that specific technique. The group decided that interactive supercomputing should, at least, satisfy three criteria: (1) *be faster than physical time*, (2) *comprise a feedback loop so that simulation could be adjusted on-the-fly*, (3) *be accessible and easy to use by the community*.

These three criteria deserve more exploration. Being faster than physical time (we avoid real-time as this is an overloaded term) means that HPC is being leveraged in contrast to simply running on a cluster. It is this feature, the fact that HPC provides significant benefit, that is the crucial thing here. But because the fusion of HPC and real-time data is to be used for urgent decision making, there needs to be some sort of bound on the performance as results are required ASAP. The feedback loop is to enable urgent decision makers to interactively explore aspects of the simulation whilst it is running, for instance, experimenting with different scenarios to understand the full impact of their decisions. This feedback loop is not just necessarily feeding back into an individual simulation, but also on an ensemble basis where the user can easily start and stop simulations as needed. The last requirement, to be accessible, is important as urgent decision makers are likely to be domain experts rather than computing experts. A number of technologies already exist, such as ParaView/Catalyst for in-situ data analytics and visualisation, and accessible supercomputing via web interfaces, such as http://openondemand.org/

The group then discussed possible use cases, a number of which were identified, including forest fires, disease spreading, medical simulations, space weather, traffic, finance, nuclear reactor operation, earthquake response, geophysics and hydrocarbon (oil and gas) exploration. From the discussions held we were surprised at how varied these different applications are, and this might raise a challenge when it comes to standardisation. For instance, how easy is it to standardise interfaces to sensor data and interactivity with such a diverse set of application areas, each with their own specific input data formats and interactivity patterns? However, in many ways, this is a nice problem to have. The fact that the group believed such a wide ranging set of areas could take advantage of interactive supercomputing for urgent decision making, is really quite positive because it illustrates the impact that this could have in many areas.

In terms of the foreseen technical difficulties, the adaption of resources was seen as a significant challenge because HPC machines are simply not set up to support a dynamic allocation of resources. How best to support elastic compute is a major question as it is unacceptable for this sort of application to be sitting in a batch queue for an unbounded amount of time, when users need an answer ASAP to make urgent decisions. Instead of adapting resources, the group discussed how one might adapt the accuracy of the simulation depending on the available resources in such cases where any decision, regardless of specific accuracy, is better than no decision at all. The group also discussed virtualisation and the fact that this sort of interactive supercomputing might imply a significant number of dependencies when processing the sensor data, running the simulations and supporting the interactivity technology. Therefore virtualising this, maybe via a containers based approach, could be useful as it would allow for easier portability between machines and might help in urgent situations where other HPC resources need to be brought online in order to meet the demand placed upon it by users.

Group 2: Interactive data exploration challenges

The purpose of this group was to discuss the challenges arising from interactive data exploration from the technology level all the way up to issues around human-computer interaction. We felt this is a hugely important area because, whilst the amount of data and processing power is growing very significantly, human cognition is constant. Therefore, whilst we have access to unprecedented amounts of data and computing power to process it, if the data is not comprehensive and users can not interact with it, then it is of limited value. This challenge is especially significant in the field of urgent decision making, where time pressure is likely to play a role and the user's need to be able to quickly understand the information in order to make the correct decision every time.

The background of the group comprised individuals from computational chemistry backgrounds, an individual from NASA who is interested in fiber placement with robots and tomography data edge computing, people who focus on molecular dynamics and protein folding, some interested in metadata management and an individual with expertise in stream processing.

The first topic of conversation was that around the reduction of data and feature extraction, machine learning was the biggest stand out technology mentioned here. NASA is already doing this very successfully based on machine learning techniques for structure identification. For protein folding, proteins are represented as graphs and then machine learning techniques can be applied to these (in fact, AI has already beaten humans in protein exploration [https://www.theguardian.com/science/2018/dec/02/google-deepminds-ai-program-alphafold-predicts-3d-shapes-of-proteins](https://www.theguardian.com/science/2018/dec/02/google-deepminds-ai-program-alphafold-predicts-3d-shapes-of-proteins)). There was also a discussion around stream processing and how machine learning can be used to extract features here too.

Data exploration for decision making was then discussed and it was highlighted that aircraft manufacturers, such as Boeing, and NASA consider this a very important feature, but there are also other application areas such as MD codes. The ability to exploit provenance data, and understand how a specific situation has arisen given its history of progression, was also highlighted as important. Being able to fuse the data successfully in with the visualisation is also important for NASA, it isn't just the visualisation but also being able to drill down and understand the numbers behind this in more detail too.

The ability for models to adapt was also considered important. The group highlighted two general drivers to this, firstly adaption in response to user input during visual analytics. For instance potentially identifying a feature or situation and wanting the simulation to *zoom in* on this, whether it being starting new ensembles or modifying an existing simulation to focus in on a part of the domain. The other general driver is that driven by data assimilation where external sensor data automatically kicks off the execution of new simulations (potentially ensembles) in order to process this.

On the topic of ensemble simulation, some of the group members make frequent use of this already. For instance NASA performs uncertainty bounding and in MD codes the use of trajectories and close trajectories is beneficial. The last topic the group discussed was the importance of visualisation, generally the group agreed that this is a very powerful approach to interacting with data and simulations. There are numerous example of this, such as visual inspection and the group members believed that ray tracing could assist in MD codes.

<u>Technical challenges arising from fusing HPC with real-time data</u>

The group started by an observation that edge computing had not been mentioned in the introduction and a number of participants considered this surprising. There is lots that can be done on the edge to make life easier, and examples were discussed such as converting data into a consistent format, reducing data and compressing it. This then makes life easier at your central HPC machine, as some of the work is already done and potentially the amount of data that must be handled is reduced. However nothing comes for free and data reduction at the edge is a really big challenge in itself!

A related question is whether you give the same priority to data at the edge, or not, and the group discussed an example of doing partial analysis on the edge data and using this – which can be better than having no data (and hence no answer) at all. Things could then be improved iteratively as you get more data and give more confidence to your decisions.

The group then moved onto a question around what the possible machine architecture for this looks like and one member of the group has already developed a distributed architecture that aims to do exactly this. They described their architecture via a diagram which contains more limited processing at the edge and then refinement to performed to combine data together. Eventually this data reaches the large resource in the middle, but still the data volume this must handle is much less than if all the raw data were simply dumped over to it.

The group asked what volume (size) and velocity of data we are talking about here, and an observation was made that it is really important to make this explicit as *big data* can really be applied to many different data sets. In response to this the group talked about different data sizes and it is clear in different application spaces the volume and velocity of data varies greatly, but in many cases the data can still be challenging (it might be unreliable or require lots of processing to be workable.) One of the group members highlighted that sensors that generate large volumes of data are very frequent now and for the entry price of around $50, one can buy an array of sensors and quickly generate significant volumes of data from it.

A discussion was then held about real time, and a key question is what we actually mean by real time in this context. A participant noted that true real time often involves lots of money! It was highlighted that HPC machines are simply not designed for the guaranteed execution of a code within a specific time frame and there are numerous challenges to supporting this, not just technical but also from a policy perspective. Technically speaking, the ability to kill or suspend/resume specific jobs is important and this could go someway towards elastic compute (e.g. Met Office do this to some extent with their hot backup Cray for weather forecasting). Many existing HPC machines expect some sort of locally accessible disk but even with state of the art parallel filesystems these can suffer from contention and be slow, which is an issue if you need an answer ASAP. It might also be possible to accept running simulations on smaller numbers of cores if this is *all you can get* rather than wait a longer time for resource to become available and maybe then scale up later if the application allows it (this could require elastic applications, which provide the ability to dynamically scale the amount of parallelism whilst they are running.)

An in depth discussion was then held again around the idea of starting with less accurate data and refining the solution based on this as more accurate data arrives. Of course this approach relies upon a simulation code that can handle inaccurate data (either a single model where data is updated on the fly as it is running, or a new ensemble execution for each new item of more accurate data.) There is also a danger that simulations running with less accurate data can result in poor decisions being made and sometimes it might be better to make no decision rather than the wrong one. One of the participants is in the field of earthquake response and they said what that community tends to do is to pre-record many different scenarios (around 16,000) and try to match up the appropriate scenario based on the current data and results from the solver. In this way you can reduce the chances of poor decisions being made.

The group also highlighted that the software infrastructure is key to this, especially considering there might be high levels of heterogeneity (FPGAs out on the edge was given as an example) and bringing this all together will place a burden on the software. From this the group discussed data formats and it was suggested that data formats and standards are good, and there are plenty already out there including HDF5 and NetCDF. So instead of reinventing a new standard, it is probably best to rely on existing technologies and standards which the sensors and/or computation codes might already support.

A challenge raised by the group was the security and privacy of data, for instance certain areas (the group highlighted self driving cars) already collect significant amounts of data but this is considered sensitive and individuals would be wary of sharing this, hence at times you need to bring the computation to the data. In these areas it might be possible to encrypt data and communicate it, or even encrypt it and process the data encrypted (there is related work here but there are latency penalties which might be unacceptable).

Conclusions and next steps

As far as we are aware, this BoF was the first time that the fusion of HPC and real-time data for urgent decision making has been addressed at an HPC conference. Whilst some previous work has focused on this topic ([1] and [2]), prior work is rather sparse, and we believe that there are significant opportunities for further development. This BoF has revealed there are a number of different groups and individuals interested in this topic and, driven by their own specific activities, are currently tackling some of the accompanying challenges. A central theme throughout the sessions was that we, as a community, should be working together to address those challenges in a more consolidated manner which not only involves the sharing of technologies but also agreements over aspects such as standardisation.

As the next step, we will conduct a survey analysing various applications that can benefit from the fusing of data and HPC for urgent decision making. We believe that this initiative would be of significant interest to many of these present at the BoF and beyond.

*[1] Leong, Siew Hoon, Anton Frank, and Dieter Kranzlmüller. "Leveraging e-infrastructures for urgent computing." Procedia Computer Science 18 (2013): 2177-2186.*

*[2] Beckman, Pete, et al. "SPRUCE: A system for supporting urgent high-performance computing." Grid-Based Problem Solving Environments. Springer, Boston, MA, 2007. 295-311.*