Conceptualizing A Computing Platform for Science Beyond 2020: To Cloudify HPC, or HPCify Clouds?

Geoffrey C. Fox
School of Informatics and Computing
Indiana University, Bloomington
IN 47408, USA

Shantenu Jha
RADICAL, Rutgers University
Piscataway
NJ 08854, USA

ABSTRACT

A primary challenge of the cyberinfrastructure research community is the need to define the Platforms for Science beyond 2020. We analyze major current trends and propose that in order to deliver the Platform for Science in 2020 the dominant research challenge is to manage the convergence of capabilities of traditional HPC systems with richness of Apache Big Data systems. In this vision paper, we purport to examine the relationship between infrastructure for data-intensive computing and that for High Performance Computing and examine possible "convergence" of capabilities.

1. COMPUTING PLATFORMS: STATUS QUO

Two important current major trends in computing systems are: (i) A growth in high performance computing (HPC) with an international exascale initiative, and (ii) Software systems to support data-intensive applications with an accompanying cloud infrastructure of well publicized dramatic and increasing size and sophistication.

Traditional exascale simulations involve applications of differential equation based models that need very fine space and time steps and this leads to numerical formulations that need the memory and compute power of an exascale machine to solve individual problems (capability computing). Big data problems do not clearly need a full exascale system to address a single job. Typically any platform will be running lots of jobs that are sometimes pleasingly parallel/MapReduce (Cloud) and sometimes small to medium size HPC jobs which in aggregate are exascale (HPC Cloud) (capacity computing).

Understanding the trends has a practical consideration, as engineering the merger for HPC and data-intensive problems will allow: (i) More efficient sharing of large scale resources running simulations and data analytics; (ii) The need for higher performance Big Data algorithms; (iii) Richer software environment for research community building on many "big data" tools, and (iv) Facilitate a sustainability model for HPC, as it does not have resources to build and maintain a full software stack.

We use term cloud broadly without choosing particular implementations: public/private clouds, OpenStack/Docker virtualization. In fact many public clouds now offer features characteristic of HPC including GPU's, high performance networks and FPGA accelerators. This is clear for deep learning in the cloud where the value of GPU's Is well understood.

Probably public clouds either offering IaaS or those running today's Internet are the lowest cost (but not necessarily lowest price) solution and in aggregate are far more powerful than the systems used in science research. Of course all systems require a significant ecosystem with many people developing, testing and running software.

2. UNDERSTANDING APPLICATIONS

We have examined extensively the landscape of applications across the HPC and data-intensive spectrum. For example in Ref. [1, 2] on examining applications with common characteristics, we introduced the concept of Ogres and 64 Convergence Diamonds (features). Ogres provides a classification and structure including, (i) classic MPI-based simulations, (ii) pleasingly parallel and workflow systems, and (iii) data-intensive applications epitomized by deep learning.

We highlight the primary differences and similarities between data-intensive problems and traditional high-performance applications.

Some machine learning like topic modeling (LDA), clustering, deep learning, dimension reduction, graph algorithms involve Map-Collective or Map-Point to Point iterative structure and benefit from HPC. However, in general, deep learning doesn't exhibit massive parallelism due to stochastic gradient descent using small mini-batches of training data, but deep learning does use small accelerator enhanced HPC clusters. If this were to change, this would have important implications for Deep learning on HPC platforms.

Further differences between data-intensive and simulation applications worth a brief mention are: (i) Classic Non-iterative MapReduce is major paradigm in data-intensive sciences, but it is not a common simulation paradigm except where "reduce" summarizes pleasingly parallel execution as in some Monte Carlo simulations; (ii) There are many data-intensive applications that have a master-worker or embarrassingly parallel flavor; examples are all the event- based applications from distributed devices or the separate interactions on a science instrument – accelerator, telescope, light source; (iii) There are important similarities between Grid Computing (from HPC field) and some big data applications and technologies; (iv) Data intensive applications often have large collective communication, whereas classic simulation has a lot of smallish point-to-point messages which motivates the MapCollective model, and (v) Simulations tend to need high precision and very accurate results (partly because of differential operators), however, data-intensive problems often don't need high accuracy as seen in trend to low precision (16 or 32 bit) deep learning networks, as there are no derivatives and the data has inevitable errors.

There are similarities between graph based data intensive applications and particle simulations with a particular cutoff force. Both are MapPoint-to- Point problem architecture many data-intensive problems involve full matrix algorithms and thus are easy to parallelize similar to "long range force" (as in gravitational simulations) as all points are linked. This suggests an opportunity for fast multipole ideas in data-intensive applications that was for example highlighted in the NRC report [3].

3. COMPUTING PLATFORMS: TRENDS

There are at least three trends that we see represented in any Future Platform for science research:

- The increasing power and complexity of modern HPC systems as exemplified by those involved in drive to build exascale class machines.
- The increasing use and sophistication of commercial and open cloud infrastructure (that can be used as IaaS, PaaS, SaaS, FaaS etc.)
- The increasing functionality and use of Big Data software systems in conjunction with HPC.

In addition, there is growing interest in streaming data and event based computing models such Amazon Lambda, IBM OpenWhisk and Function-as-a-Service (serverless computing). These general trends are likely to continue independent of specific technology trends. The specifics of the technology will manifest across the following:

Microscopic Architecture: The three primary microscopic architecture are: (i) Continuation of X86 systems, (ii) Many core systems (e.g., KNL) and (iii) non-traditional architectures (e.g., GPU, FPGA) etc.

Macroscopic Architecture: The three primary macroscopic architectures are: (i) Data Center Model, (ii) Traditional supercomputers and, (iii) Clusters (with virtualization) such as those represented by NSF Comet.

4. CONVERGENCE OF CAPABILITIES

Armed with an understanding of the spectrum of applications and platform trends, we will now examine the convergence of high-performance computing and data-intensive platforms[4, 5, 6].

In order to meet the requirements of future science applications, there must be a greater/richer set of analysis-as-a-service than currently available. Future analysis and associated middleware must utilize traditional performance capabilities, yet expose fundamentally new capabilities. Thus a prudent, if not only approach, is to (re-)design the software stack for analysis.

Future platforms be usable by applications from both ends of the spectrum: traditional HPC applications that need Big Data ("All Exascale Applications are data-intensive problems"), as well as data-intensive applications that will increasingly need HPC (e.g., Deep Learning with HPC capabilities). Given the current separation of characteristics of HPC and data-intensive applications this requires a convergence of capabilities.

We deduce that any Future Platform must satisfy the following constraints:

- It must allow easy integration of public and private clouds and allow HPC and cloud approaches to run well and run together
- It must allow the powerful features of modern clouds such as ABDS, XaaS to be useable on HPC hardware
- It must support distributed data sources and repositories
- It should support modern workflow and portals including Python based front ends; an area where simulations and Big Data have similar requirements.

We call such platforms HPC Cloud platforms. Independent of whether we consider "cloudification" of HPC or "HPCfication of Clouds".

Independent of the directionality, the future platform will be a software- defined system that works across different types of macroscopic and microscopic architectures as well as for different applications systems. This requires the selective integration of the Apache Big-Data Stack (ABDS) capabilities appropriately implemented for supercomputing platforms. We have examined High Performance Computing Enhanced Big Data Stack (HPC-ABDS) where we examined the addition of high performance runtime and components to Apache systems. We have highlighted the importance of the Big Data systems associated with Apache Foundation, such as Hbase, Hadoop, Spark, Storm etc., which we term the Apache Big Data Stack (ABDS), even though important components such as MongoDB and Tensorflow are not Apache projects. We note that most of these technologies are in principle usable on both HPC and Cloud IaaS systems, though in practice many challenges remain. Independent of the hardware infrastructure, there are even stronger forces driving the adoption of ABDS technologies. They offer usability, functionality and sustainability that is not available in the HPC ecosystem. A realization of the HPC-ABDS concept is provided by the SPIDAL project [7, 8] and discussed in publications [9, 2].

5. REFERENCES

[1] S. Jha, J. Qiu, A. Luckow, P. Mantha, and G. C.Fox, "A Tale of Two Data-Intensive Paradigms: Applications, Abstractions, and Architectures," in *Big Data (BigData Congress), 2014 IEEE International Congress on*, 2014, pp. 645 – 652. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6906840

[2] G. C. Fox, S. Jha, J. Qiu, and A. Luckow, "Towards an understanding of facets and exemplars of big data applications," in *Proceedings of the 20 Years of Beowulf Workshop on Honor of Thomas Sterling's 65th Birthday*, ser. doi: 10.1145/2737909.2737912. New York, NY, USA: ACM, 2015, pp. 7–16. [Online]. Available: http://doi.acm.org/10.1145/2737909.2737912

[3] Committee on the Analysis of Massive Data and Committee on Applied and Theoretical Statistics and Board on Mathematical Sciences and Their Applications and Division on Engineering and Physical Sciences, *Frontiers in Massive Data Analysis.* The National Academies Press, 2013.

[4] G. Fox, J. Qiu, and S. Jha, "High performance high functionality big data software stack," http://www.exascale.org/bdec/sites/www.exascale.org/bdec/files/whitepapers/fox.pdf, 2014.

[5] G. C. Fox, S. Jha, J. Qiu, and A. Luckow, "Ogres: A Systematic Approach to Big Data Benchmarks," http://www.exascale.org/bdec/sites/www.exascale.orgbdec/files/whitepapers/OgreFacets.pdf, Barcelona, January 29-30 2015.

[6] Geoffrey Fox, Judy Qiu, Shantenu Jha, Supun KamburugamuveSaliya Ekanayake, "Big Data, Simulations and HPC Convergence" Big Data and Extreme-Scale Computing (BDEC) Frankfurt June 15 to June 17, 2016 http://dx.doi.org/10.13140/RG.2.1.3112.2800.

[7] J. Qiu, S. Jha, A. Luckow, and G. C.Fox, "Towards HPC-ABDS: An Initial High-Performance Big Data Stack," March 18-21 2014. [Online]. Available: http://grids.ucs.indiana.edu/ptliupages/publications/nist-hpc-abds.pdf

[8] G. C. Fox, S. Jha, J. Qiu, and A. Luckow, "Towards an Understanding of Facets and Exemplars of Big Data Applications," October 14 2014. [Online]. Available: http://grids.ucs.indiana.edu/ptliupages/publications/OgrePaperv9.pdf

[9] G. C. Fox, J. Qiu, S. Kamburugamuve, S. Jha, and A. Luckow, "Hpc-abds high performance computing enhanced apache big data stack," in *2015 15th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing*, May 2015, pp. 1057–1066.