Big Data and HPC Convergence: The Cutting Edge and Outlook

Sardar Usman[1(✉)], Rashid Mehmood[2], and Iyad Katib[1]

[1] Department of Computer Science, FCIT, King Abdulaziz University, Jeddah 21589, Saudi Arabia
usmansardar@hotmail.com, iakatib@kau.edu.sa

[2] High Performance Computing Center, King Abdulaziz University, Jeddah 21589, Saudi Arabia
RMehmood@kau.edu.sa

Abstract. The data growth over the last couple of decades increases on a massive scale. As the volume of the data increases so are the challenges associated with big data. The issues related to avalanche of data being produced are immense and cover variety of challenges that needs a careful consideration. The use of (High Performance Data Analytics) HPDA is increasing at brisk speed in many industries resulted in expansion of HPC market in these new territories. HPC and Big data are different systems, not only at the technical level, but also have different ecosystems. The world of workload is diverse enough and performance sensitivity is high enough that, we cannot have globally optimal and locally high sub-optimal solutions to all the issues related to convergence of big data and HPC. As we are heading towards exascale systems, the necessary integration of big data and HPC is a current hot topic of research but still at very infant stages. Both systems have different architecture and their integration brings many challenges. The main aim of this paper is to identify the driving forces, challenges, current and future trends associated with the integration of HPC and big data. We also propose architecture of big data and HPC convergence using design patterns.

Keywords: HPC · Big data · Hadoop · HPDA · Design patterns
IoT · Smart cities · Cognitive computing

1 Introduction

Over the years, HPC has contributed a lot in scientific discoveries, improved engineering designs, enhanced manufacturing, fraud detection, health care, and national security, thus played crucial role towards quality of human life. The world has seen exponential data growth due to social media, mobility, E-commerce and other factors. Major chunk of data has been generated in the last few years alone and is even growing at more rapid rate [1]. To deal with ever growing volume of data, researchers have been involved in developing algorithms to accelerate the extraction of key information from massive data. Big data is a buzzword, which catches lots of attention in the recent years. It means massive amount of structured, semi structured and unstructured data

R. Mehmood et al. (Eds.): SCITA 2017, LNICST 224, pp. 11–26, 2018.
https://doi.org/10.1007/978-3-319-94180-6_4

collected from different resources and is not possible to store and process this data by traditional databases and software techniques.

Historically only the largest companies, government research organizations and academic computing centers have had an access to the computing power necessary to get to valuable conclusions in a reasonable amount of time. All that is rapidly changing with vast improvement in the price, performance, availability and density of compute power beyond the human imagination.

The categorization of data vs. computing affected by solution urgency i.e. real time solution, and also depends on what we trying to achieve. As the volume of data is growing bigger, it brings more challenges to process that data in real time. As projected, in 2018 over 4.3 Exabyte of data will be created on daily basis [2]. Over the years HPC community have not been deprived of huge volume of data i.e. climate modeling, design and manufacturing, financial services etc. that resulted in high fidelity models and interdisciplinary analysis to explore data for deeper insights. The use of High Performance Data Analytics HPDA is increasing at brisk speed in many industries resulted in expansion of HPC market in these new territories.

Powerful analytics is a key to extract a value from data by confronting budget and marketing challenges and plays huge roles in making plans, predicting business trends and understanding customer demands. Choosing a right solution depends on the size of data, urgency of results, prediction about the needs of more processing power as the size of data increases, fault tolerance for applications in case of hardware failure, data rate and scalability etc. A real time application with high response time especially when dealing with huge volume of data, is still a challenging task and is one of the driving forces towards the convergence of big data and HPC.

Both HPC and Big data are different system not only at the technical level but also have the different ecosystem. Both have different programming model, resource manager, file system and hardware. HPC are mainly developed for computational intensive applications but recently data intensive applications are also among the major workload in HPC environment. Due to recent advancements of data intensive applications, number of software frameworks has been developed for distributed systems, cluster resource management, parallel programming models and machine learning frameworks. High performance computing have very well established standard programming model e.g. Open MP/MPI. Big data analytics have been grown up in different perspective and have different population of developers that uses java and other high level languages with primary focus on simplicity of use, so that problem domain can be solved without a detailed knowledge of HPC. These difference in the infrastructure, resource manager, file system and hardware makes the system integration a challenging task.

As the data is getting bigger and bigger in volume so is the need of high computing. HPC community has been dealing with massive amount of data and big data analytics for years. The solutions evolved over the years to deal with large volume of data, should be useful for big data analytics [3]. The main aim of this paper is to identify motivation and driving forces towards the integration of HPC and big data. Also highlighting the current trends, challenges, benefits and future aspects of unified integrated system. We also present architecture for the convergence of HPC and Big data using design patterns.

The rest of the paper is organized as follows. The next section examines the difference between HPC and Hadoop framework with respect to hardware, resource management, fault tolerance and programming model. Literature survey is presented in Sect. 3 and Convergence challenges are discussed in Sect. 4 followed by the future directions in Sect. 5. The architecture using design pattern for the convergence of HPC and big data is presented in Sect. 6 and paper is concluded in the final section.

2 HPC and Big Data Frameworks and Their Differences

Different solutions emerged over the years to deal with big data issues and are successfully implemented. But never the less, all these solutions do not satisfy the ever-growing needs of big data. The issues related to big data are immense and cover variety of challenges that needs a careful consideration, for example data representation, data reduction/compression, data confidentiality, energy management, high dimensionality, scalability, real and distributed computation, non-structured processing, analytical mechanism and computational complexity etc. The exponential outburst of data and rapidly increasing demands for real time analytical solutions urges the need for the convergence of high-end commercial analytics and HPC. Business intelligence/analytical solutions today lack the support for predictive analytics, lack of data granularity, lack of software flexibility to manipulate data, lack of intuitive user interface, relevant information is not aggregated in a required manner and slow system performance [4].

HPC community have been dealing with complex data and compute intensive applications, and solutions have been evolved over the years. As the volume of data is increasing at brisk speed so are the associated challenges i.e. data analysis, minimizing data movement, data storage, data locality and efficient searching. As we are heading towards exascale era, the increase in system concurrency introduced a massive challenge for system software to manage applications to perform at extreme level of parallelism. Large-scale applications use most widely deployed message-passing programming model MPI along with traditional sequential languages, but with the introduction of architectural changes (many core chip) and high demand in parallelism make this programming model less productive for exascale systems. Billion-fold parallelism is required to exploit the performance of extreme scale machines and locality is critical in terms of energy consumption. As the complexity and scale of software requirements is on a rise, simple execution model is a critical requirement, which ultimately reduce the application programming complexity required to achieve the goals of achieving extreme scale parallelism. A current trend in HPC market includes use of advanced interconnects and RDMA protocols (Infinity Band, 10–40 Gigabits Ethernet/iWARP, RDMA over converged Enhanced Ethernet), enhanced redesign of HPC middleware (MPI, PGAS), SSDs, NVRAM and Burst buffer etc. Scalable parallelism, synchronization, minimizing communication, task scheduling, memory wall, heterogeneous architecture, fault tolerance, software sustainability, memory latencies, simple execution environment and dynamic memory access for data intensive application are some of the core areas that requires considerable time and efforts to address Exascale challenges [5]. The difference between Hadoop and HPC framework is highlighted in the following section.

2.1 Hardware

Most of the modern HPC and Hadoop clusters are commodity hardware. In HPC environment, Compute nodes are separated from data nodes. There are two types of data storage, temporal file system on local nodes and persistent global shared parallel file system on data nodes. The existing HPC clusters have limited amount of storage on each compute node. LUSTRE is most widely used parallel file system in HPC and almost 60% of the top 500 supercomputers use LUSTRE as their persistent storage. Data needs to be transferred from data nodes to the local file system on each compute node for processing. Data sharing is easy with distinct data and compute nodes but spatial locality of data is an issue [6, 7].

Hadoop cluster uses local disk space as a primary storage. The same node serves as a data node and compute node. The computational task is scheduled on same machine where data is resided resulting in enhanced data locality. Hadoop is write-once and read-many framework. I/O thorough put of Hadoop is much higher, due to co-locating of data and compute node on the same machine [7].

2.2 Resource Management

Another major difference between Hadoop and HPC cluster is resource management. Hadoop's Name node has Job tracker daemon. Job tracker supervised all map-reduce tasks and communicates with the task trackers on the data node. Compared to Hadoop's integrated job scheduler, HPC scheduling is done with the help of specialized tools like Grid engine, Load leveler etc., [8] with controlled resources (memory, time) provided to the user.

2.3 Fault Tolerance

HPC resource scheduler use checkpoint mechanism for fault tolerance. In case of node failure, it reschedule job from the last stored checkpoint. It needs to restart the whole process if the checkpoint mechanism is not used. On the other hand, Hadoop uses job tracker for fault tolerance. As data and computation are co-located on same machine, job tracker can detect a node failure on run time by re-assigning a task on a node where duplicate copy of data is resided [8, 9].

2.4 Programming Model

Hadoop uses map-reduce programming model, which makes life easier for the programmers as they just need to define map step and reduce step, when compared to the programming efforts needed for HPC applications. In HPC environment, programmer needs to take fine-grained responsibilities of managing communication, I/O, debugging, synchronization and checkpoint mechanism. All these tasks needs considerable amount of efforts and time for effective and efficient implementation. Hadoop does provide a low level interface to write and run map-reduce applications written in any language, although Hadoop is written in Java. Following Table 1 summarizes the difference between HPC and Hadoop framework [7].

Table 1. HPC vs. Hadoop eco system

	Big Data	HPC
Programming model	Java applications, SparQL	Fortran, C, C++
High level programming	Pig, Hive, Drill	Domain specific language
Parallel run time	Map-reduce	MPI, Open MP, OpenCL
Data management	HBase, MySQL	iRODS
Scheduling (Resource management)	YARN	SLRUM (Simple LINUX utility for resource management)
File system	HDFS, SPARK (Local storage)	LUSTRE (Remote storage)
Storage	Local shared nothing architecture	Remote shared parallel storage
Hardware for storage	HDDS	SSD
Interconnect	Switch ethernet	Switch Fiber
Infrastructure	Cloud	Supercomputer

Both Hadoop and Spark are big data frameworks and do perform the same tasks, are not mutually exclusive and able to work together. Spark is mostly used on the top of Hadoop and advance analytics of spark are used on data stored in Hadoop's distributed file system (HDFS). Spark has the ability to run as Hadoop's module through YARN and as a standalone solution [10] and can be seen as an alternative to map-reduce rather than a replacement to Hadoop framework. Spark is much faster compared to Hadoop because it handles in memory operations by copying data from distributed file systems in to faster logical RAM. Map-reduce writes all data back to distributed storage system after each iteration to ensure full recovery whereas Spark arranges data in resilient distributed datasets that are capable of full recovery in case of failure. Spark capability of handling advance data analytics in real time stream processing and machine learning is a much more advance that gives Spark edge over Hadoop. The choice of selecting either of the data processing tool depends on the needs of an organizations e.g. Dealing with big structured data can be done efficiently with map-reduce and there is no need to installed a separate layer of Spark over Hadoop [11]. Spark on demand allows users to use Apache Spark for in situ data analysis of big data on HPC resources [12]. With this setup, there is no longer to move petabytes of data for advance data analytics.

3 Research Related to HPC and Big Data Convergence

The integration of HPC and Big data started at different levels of their Eco systems and these integrated solutions are still at very infant stages. The convergence of both these technologies is the hottest topic for the researcher over the last few years. In [6] Krishnan et al. proposed a myHadoop framework using standard batch scheduling system for configuring Hadoop on-demand on traditional HPC resources. The overhead in this setup includes site-specific configuration, keeping input data into HDFS and then staging results back to persistent storage. HDFS is heavily criticized for its I/O bottleneck. Availability of limited storage is big challenge to integrate Hadoop with HPC clusters. Islam et al. [13] proposed a hybrid design (Triple-H) to reduce I/O bottleneck in HDFS and efficient resource utilization for different analytics system performance and cluster efficiency with overall low system cost.

Data intensive applications have been intensively used in HPC infrastructure with multicore systems using Map-reduce programming model [14]. With increase in parallelism, the overall throughput increases resulted in high-energy efficiency as the task is completed in shorter span of time. When Hadoop runs on HPC cluster with multiple cores and each node is capable of running many map/reduce tasks using these cores. This ultimately decreases the data movement cost and increase throughput but due to high disk and network accesses of Map-reduce tasks, the energy consumption and through put cannot be predicted. High degree of parallelism may or may not affect energy efficiency and high performance.

Tiwari et al. [15] studied the Hadoop's energy efficiency on HPC cluster. Their study shows that energy efficiency of map-reduce job on HPC cluster changes with increase in parallelism and network bandwidth. They determine the degree of parallelism on a node for improving the energy efficiency and also benefits of increasing the network bandwidth on energy efficiency by selecting configuration parameters on different types of workloads i.e. CPU intensive and moderate I/O intensive, CPU and I/O intensive workloads, also energy and performance characteristics of a disk and network I/O intensive jobs. When the number of map slots reached beyond 40, number of killed map tasks almost doubled. Thus increasing the parallelism to certain extent has positive impact on energy efficiency.

Scientific data sets are stored in back end storage servers in HPC environment and these data sets can be analyzed by YARN map-reduce program on compute nodes. As both compute and storage servers are separated in HPC environment, the cost of moving these large data sets is very high. The High-end computation machine and analysis clusters are connected with high-speed parallel file system. To overcome the shortcomings of offline data analysis, "in situ" data analysis can be performed on output data before it is written to parallel file system. The use high-end computation node for data analysis results in slowing down simulation job by the interference of the analysis task and inefficient use of computation resources for data analysis tasks. Spark on demand allows users to use Apache Spark for in situ data analysis of big data on HPC resources [12]. With this setup, there is no longer to move petabytes of data for advance data analytics.

According to Woodie [16], the use of InfiniBand for large clusters is most cost effective then standard Ethernet. The performance of HPC oriented Map-reduce solutions (Mellanox UDA, RDMA-Hadoop, DataMPI etc.) depends on the degree of change in Hadoop framework as more deep modification means an optimal adaption to HPC systems. Hadoop with IPoIB (IP over InfiniBand) and Mellanox UDA requires minimal or no changes in Hadoop implementation and only requires minor changes in Hadoop configuration. RDMA-Hadoop and HMOR are the HPC oriented solutions to take advantage of high speed interconnects by modifying some of the subsystems of Hadoop. DataMPI is a framework that developed from the scratch, which exploits the overlapping of map, shuffle and merge phases of map-reduce framework and increases data locality during the reduce phase. DataMPI provides the best performance and an average energy efficiency [17]. The use of InfiniBand improved the network bandwidth, as InfiniBand being widely used in HPC environment. Communication support in Hadoop relies on TCP/IP protocol through Java sockets [17]. So it is difficult to use high performance interconnects in an optimal way so different HPC oriented map-reduce solutions came that addresses the problem of leveraging high performance interconnects RDMA –Hadoop, DataMPI etc. Wang et al. [18] compared the performance of 10 GigaBit Ethernet and InfiniBand on Hadoop. With small intermediate data sizes the use of high speed interconnect, increased the performance by efficiently accelerating jobs but doesn't shows the same performance with large intermediate data size. The use of InfiniBand on Hadoop provides better scalability and removes the disk bottleneck issues. As the Hadoop cluster is getting bigger, organizations feel the need of specialized gear like solid-state drives (SSDs) and the use of InfiniBand instead of standard Ethernet. The use of infiniBand with RDMA (remote direct memory access) allows 40 Gigabits/s raw capacity out of Quad Data Rate (QDR) infiniBand port which is four times as much bandwidth as 10 GigaBit Ethernet port can deliver [16].

The use of infiniBand allows maximum scalability and performance while overcoming the bottlenecks in the I/O. Islam et al. [19] proposes an alternative parallel replication scheme compared to pipelined fashioned replication scheme by analyzing the challenges and compared its performance with existing pipelined replication in HDFS over Ethernet, IPoIB, 10 GigE and RDMA and showed performance enhancement with parallel model for large data sizes and high performance interconnects.

4 Challenges of Convergence

The world of workload is diverse enough and performance sensitivity is high enough that, we cannot have globally optimal and local high sub-optimal solution to all the issues related to convergence of HPC and big data. HPC and Hadoop (big data) architectures are different and have different eco system. The cross fertilization of HPC and Big data is the hottest topic for the researchers over the last few years. Most of the research related to the convergence of HPC and big data started at distinct levels of eco system but do not address the problem of moving data especially in HPC environment. The integration of data intensive applications in HPC environment will bring many challenges. In Exasacle environment cost of moving big data will be more then cost of floating point operations. There is a need for high energy efficient and cost effective interconnects for high bandwidth data exchange among thousands of processors. We also need a data locality aware mechanism especially when dealing with big data in HPC shared memory architecture. The cost of moving big data for processing also brings another challenge of high power consumption. With massively parallel architecture with hundreds of thousands of processing nodes, the cost of moving data will be very high. According to Moore et al. [20], and energy efficiency of 20 pJ (Pico Joules) per floating point operation is required for exascale system where as current state of art multicore CPUs have 1700 pJ and GPUs have 225 pJ per floating point operation.

Minimizing the data movements means the innovation in memory technologies with enhanced capacity and bandwidth. To deal with 3Vs (volume, velocity, veracity) of big data, efficient data management techniques need to be investigated included data mining and data co-ordination [13] as most of the HPC platforms are compute centric, as opposed to the demands of big data (continuous processing, efficient movement of data between storage devices and network connections etc.). To deal with massive parallel architecture and heterogeneous nature of big data, innovation needed at the programming model to deal with the next generation of parallel systems. Thus reducing the burden of parallelism and data locality for application developer as MPI leave it to the programmer to handle issues related to parallelism. Hadoop being widely used as a big data framework, achieve fault tolerance by the replication of data on multiple nodes and job tracker assign job to other node in case of node failure. Fault tolerance in HPC is by means of checkpoint mechanism, which is heavily criticized and not suitable for exascale environment. In exascale systems hardware failure will be a rule not an exception. The MTBF (mean time between failures) window in current Peta-scale system is in days and for exascale systems it will be in minutes or may be few seconds. So there is need for a comprehensive resilience at the different levels of exascale eco system. Exascale systems will be constrained by power consumption, memory per core, data movement cost and fault tolerance. The integration between HPC and big data must address the issues of scalability, fault resilience, energy efficiency, scientific productivity programmability and performance [21].

Resilience, power consumption and performance are inter-related to each other. High degree of resilience or fault tolerance is achieved but on the expense of high power consumption. As we are heading towards exascale era, convergence of both HPC and big data will make energy efficiency a core issue to handle. Severs and data-centers are facing the same problem of power consumption including companies like Google, Amazon and Facebook etc. According to an estimate the actual cost of exascale system will be less then cost of power consumption for maintaining and running exascale system for one year [22].

The energy efficiency techniques in big data can be broadly categorized as software/hardware based energy efficient techniques, energy efficient algorithms and architectures. As set of commodity hardware is used in both HPC and Big data platforms for processing of data. The integrated hardware solution for data intensive applications and computational intensive applications wouldn't work for exascale systems as hardware solution helps to achieve fault tolerance but on the expense of high energy consumption. The current Peta scale high performance computing with checkpoint mechanism to achieve fault tolerance and energy efficiency does not suit well for the integrated solution of HPC (Exascale) and big data. Soft, hard and silent errors in exascale environment will be rule not an exception. Thus collaborative efforts are needed at system level or application level resilience to deal with fault tolerance and energy efficiency for the integrated solution.

As we have seen that both HPC and Hadoop (big data) architectures are different and have different eco system. Both have different programming model, resource manager, file system and hardware. These difference in the infrastructure, resource manager, file system and hardware makes the system integration a challenging task. As the data is getting bigger and bigger in volume so is the need of high computing. One of the biggest challenges, that both big data and HPC community facing is energy efficiency. Exascale Parallel computing system will have thousands of nodes with hundreds of cores each and is projected to have billions of threads of execution. The frame of Main Time between Failures MTBF in super computers is in days and weeks. But for Exascale computing with million times more components, the perception of MTBF is in hours or minutes or may be in seconds. Each layer of Exascale Eco system must be able to cope with the errors [23].

Real time data analysis is also a driving force behind the urgency of the need for the necessary convergence of the analytics, big data, and HPC when dealing with computation, storage and analysis of massive, complex data sets in high scalable environment. Scalability issues addressed by the HPC community by capitalizing the advancements in network technologies (low latency network), efficient and large memory should also address the scalability issues of the data analytics [24].

5 Driving Forces and Future Aspects

High performance data analytics HPDA includes tasks involving massive amount of structured, semi-structured and unstructured data volumes and highly complex algorithms that ultimately demands the needs of HPC resources. Companies now have the computing power they need to actually analyze and act upon their data. This translates into numerous benefits for the company, environment and society over all. In the energy sector companies are now able to more accurately drill for oil. Automobiles and airlines are much safer due to rapid modeling of operational data design optimization and aerodynamics analysis, allowing them to deliver more cost effective products that operate safer and are more fuel-efficient. In the financial sector banks and card issuers can do fraud detection in real time. Stock investors can quickly track trends in the market to better serve their investing customers. Retailers and advertisers can now review historic purchasing data to better deliver the right products and advertisement to their customers and whether researchers can study thousands of years of weather data in hours or days instead of weeks or months, improving the quality of predictions and safety of people worldwide. HPC industry has been dealing with data intensive simulations and high performance analytics solutions also evolved over the years urges the commercial organizations to adopt HPC technology for competitive advantage to deal with time critical and highly variable complex problems. The chasm between data and compute power is becoming smaller all the time. The global HPDA market is growing rapidly and according to forecast HPDA global market size was US 25.2 billion and with the growth of nearly 18%, it is projected to be around US 82 billion in 2022 [25] (Fig. 1).

Fault tolerance, high power consumption, data centric processing, limitations of I/O and memory performance are few of the driving forces that are reshaping the HPC platforms to achieve Exascale computing [26]. Data intensive simulations, complex and time critical data analytics requires high performance data analytics solutions for example Intelligence community, data driven science/engineering, machine learning, deep learning and knowledge discovery etc. These competitive forces have pushed relatively new commercial companies (Small and Medium scale Enterprises SMEs) into HPC competency space. Fraud/anomaly detection, affinity marketing, business intelligence and precision medicine are some of the perusable new commercial HPC

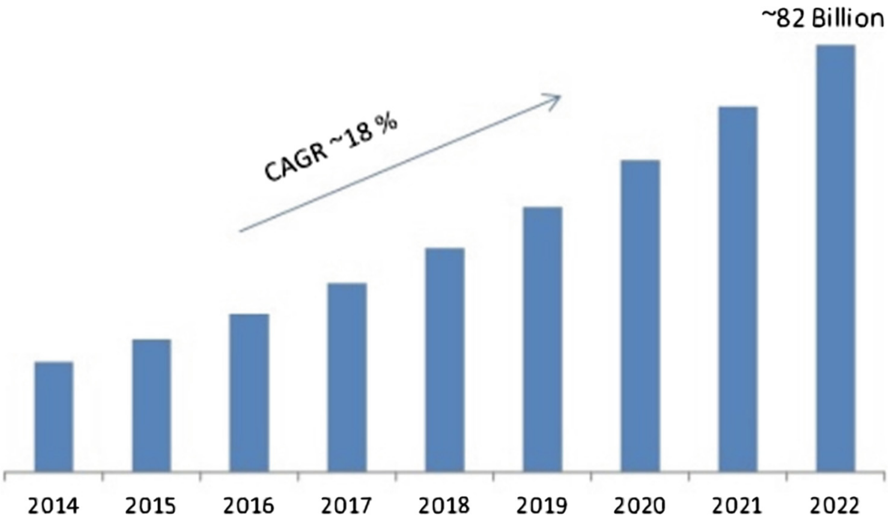

Fig. 1. HPDA market forecast [25]

market segments that require high performance data analytics. The use of HPDA will increase with time in future demanding convergence of HPC and big data. HPDA is becoming an integral part of future business investments plans of enterprises, to enhance customer experience, anomaly detection marketing, business intelligence, security breaches etc. and discovery of new revenue opportunities.

5.1 The Internet of Things IoT and Smart Cities

IoT links physical devices (computers, sensors, electronics,) equipped with sensors to the Internet and network connectivity enabling them to communicate. The common IoT platform brings heterogeneous information together and facilitates communication by providing common language. According to Gartner [27] IoT units installed base will reach 20.8 billion by 2020 resulted in massive amount of data which will further highlight the security, customer privacy, storage management and data centric networks challenges. Smart city demands better and more inventive services to run whole city smoothly and improve people's life through the innovative use of data.

Smart cities and IoT are some of the emerging HPDA application areas. HPC has been involved in managing power grids and transport for the upstream design of vehicles and urban traffic management in smart cities for quite some time and its use over time will increase in the markets of cognitive computing/AI, driverless vehicles and healthcare organizations. Baz [28] investigated the connection between IoT and HPC by highlighting some of the challenges in smart world applications (smart building management, smart logistics and smart manufacturing) and possible opportunities with HPC enable solutions. China's HPC-IoT plan 2030 is based on the use of HPC in IoT network wellness management and security [29].

5.2 Cognitive Technology

Cognitive systems are capable of understanding complex language constructs, correlate the association and help to rationalize information and discover insights. The key in cognitive systems is learning, adaptability and how the system is evolving, helps in decision-making process, discovery of new ventures, improved production and operation systems, optimizing resources, proactive identification of faults ahead of failure etc. The motive of cognitive computing is to handle complex problems without no or little human intervention. According to IBM estimate 80% of data is unstructured and is of no use for the machines and not fully exploited. The cognitive computing can be seen as a potential candidate for the exploration of unstructured data to get more useful information insights and efficient decision-making. The rapid growth of data from multidisciplinary domains requires powerful analytics but lacks human expertise to tackle the diverse and complicated problems. The cognitive computing allows people with less experience to interact with machine thanks to the advancement in natural language processing and Artificial intelligence technologies e.g. Google DeepMind and Qualcomm's Zeroth Platform. The advancement in cognitive technology with the integration of AI and machine learning for big data tools and platforms will increase the quality of information, dealing with the complex data analytics with lesser human intervention but requires rapid data access (low latency), faster time to insights, hardware acceleration for complex analytics [2]. Extracting information from vast amount of data requires innovation in compute and storage technologies, which should provide cost effective storage, improved performance in a desired time frame. The infrastructure required cognitive storage with learning ability for computers to store only relevant and important data. The computing requires efficient processing which demands high memory bandwidth and extreme scale parallelism for efficient resource utilization within energy efficiency constraints. Open power foundation [2] is an initiative towards partnering technology solutions with diverse companies coming together to provide technology solutions to a variety of problems. With data centric computing, time to solution will be dramatically reduced. Cognitive computing though still at its infancy stages but in future will be a key technology for the success of modern businesses, to get insights of the vast amount of unstructured data by leveraging computing technology to work better with the way humans want to work and smoothing the natural relationship between human and the computer.

6 Design Patterns

The need for HPDA demands innovative ways, to accelerate data and predictive analysis to target above-mentioned complex challenges by revolutionary and evolutionary changes in programming models, computer architecture and runtime systems to accommodate potential interoperability and scaling convergence of HPC and Big data eco systems [2]. There is growing need for the efficient exploration of novel techniques to allow HPC and Big data applications to exploit billion-fold parallelism (Exascale systems), improved data locality, unified storage systems, synchronization and ultimately the single system architecture to overcomes the cost and complexity of moving data which also improves the total cost of ownership and brings in flexibility to manage workflows and maximize system utilization. Design patterns and skeletons are the potential candidates to address above-mentioned challenges to design scalable, robust software development and applicable proved solutions in both HPC and big data community.

The parallel programing problem has been an active area of research for decades focusing primarily on programming models and their supporting environments. As we move towards Exascale (millions of components, billions of cores) programming parallel processors and handling billion-way parallelism is one of the major challenge that research community is facing. Software architecture and design plays a vital role in designing robust and scalable software. Common set of design elements (derived from domain expert's solutions), are captured in a design pattern of that particular domain to assist the software designer to engineer robust and scalable parallel software. These patterns define the building blocks of all software engineering and are fundamental to architect parallel software. The design problem at different level of software development is addressed by developing layered hierarchy of patterns by arranging patterns at different levels. These design patterns have been developed to assist software engineers to architect and implement parallel software efficiently. Our Pattern Language OPL is one of prominent source of cataloguing and categorizing the parallel patterns [30]. A design pattern provides a clean mechanism to cater common design problems using generic guidelines.

Big Data design patterns provide the concrete representation of analysis and technology centric patterns of most common occurring problems in BigData environment [31]. These design patterns provides the building blocks for the efficient design of big data architecture. The standardization and integration of design patterns can be seen as the potential candidates for the efficient and effective convergence of HPC and big data. Figure 2 shows the logical architecture of different layers and design patterns (HPC & BigData) can then be applied at distinct levels to address the issues related to big data and HPC convergence. One of the challenges associated with data visualization and interactive management is huge volume, variety and velocity of data and is often hard to evaluate and reapply the design solution. The visualization and management layer involves applying patterns for distributed and parallel visualization, interactive data exploration, rendering data visualization, real time monitoring for live analysis and recommendations.

The analytics/processing layer includes patterns for analytics and depending on the problem domain includes in-situ, in-transit, real time or batch processing. Advanced analytics requires predictions, advance algorithms, simulations and real time decisions that require high performance computing for processing and managing massive volume of data [32].

There is a trade-off between Performance, resilience and power consumption. Trade-off patterns needs to identify and accommodate these trade-offs in best possible way by indulging the best practices from both HPC and Big data communities. The processing pattern includes analytics patterns for unstructured and structured data, algorithms for conversion of unstructured to structured data, large-scale batch and graph based processing patterns and also parallel design patterns. The access/storage layer includes design patterns for the effective and efficient retrieval and storage mechanism for parallel

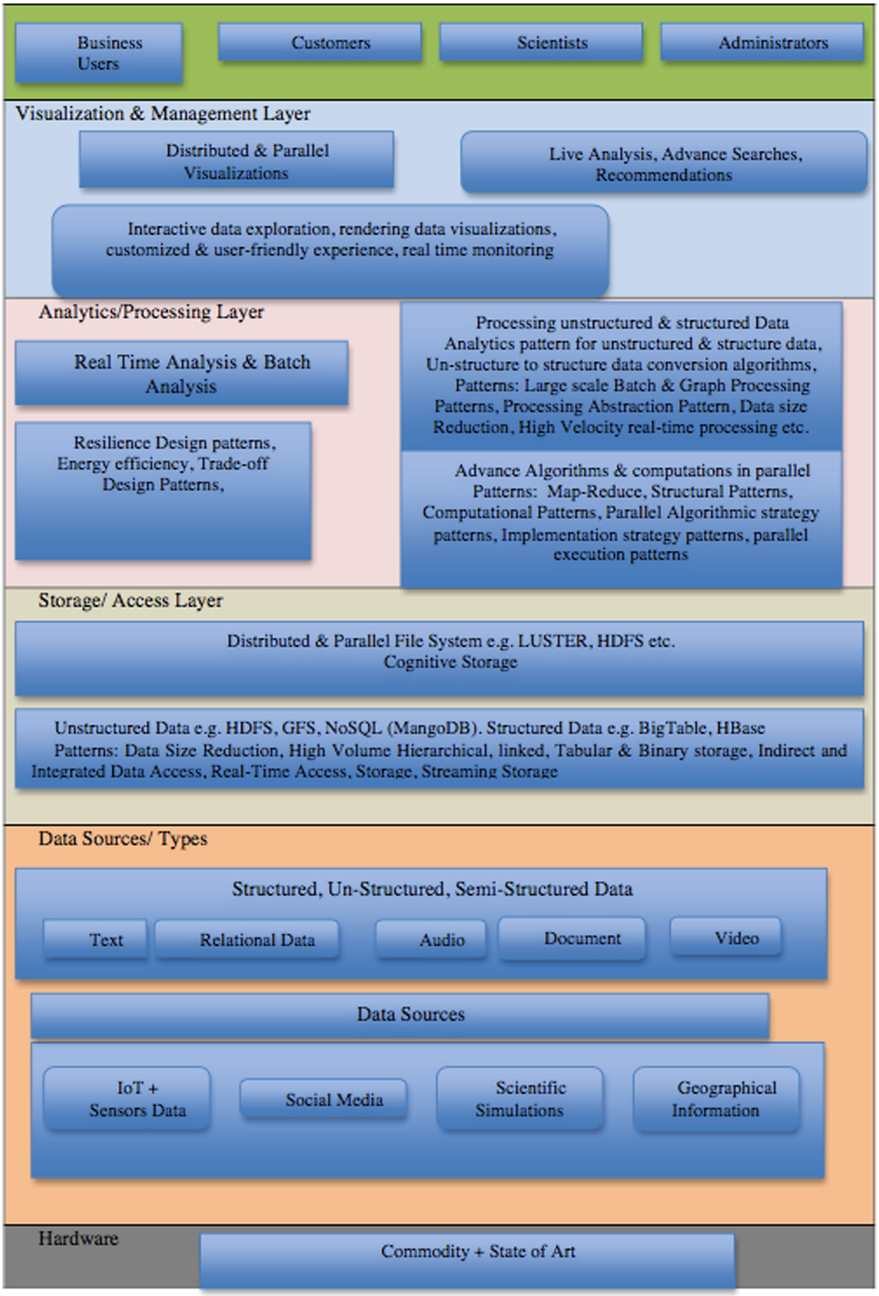

Fig. 2. Logical layered architecture of design patterns

and distributed file systems. This includes data size reduction for high volume hierarchical, linked, tabular and binary cognitive storage for real time in-direct and integrated access. The cognitive storage with learning ability to automate the process of data purging by keeping only relevant and important data for cost effective storage and improved performance.

HPC software development community lack the expertise of software engineering principles as these patterns define the building blocks of software engineering and are fundamental to architect parallel software. There is a need to invest the research efforts towards exploration of innovative approaches to make use of design patterns and skeletons to overcome scalability, elasticity, adaptability, robustness, storage, parallelization and other processing challenges of the unified HPC and big data environment.

7 Conclusion

The increased processing power, emergence of big data resources and real time analytical solutions are the prime drivers that pushing the realm of big data. As both HPC and big data systems are different and have different architecture. The challenges associated with inevitable integration of HPC and big data are immense and solutions are starting to emerge at distinct levels of eco system. As we are heading towards convergence of both, we will have to deal with modality, complexity and vast amount of data. Currently we have distinct and perhaps overlapping set of design choices at various levels of infrastructure. A single system architecture but with enough configurability in it that you can actually serve different design points between compute intensive and design intensive. The single system architecture overcomes the cost and complexity of moving data. It also improves the total cost of ownership and brings in flexibility to manage workflows and maximize system utilization. Realizing these benefits requires coordinated design efforts around key elements of the system i.e. compute (multicore, FPGA), interconnect (next generation fabric), memory (Non Volatile memory, storage burst buffer, Luster file system). This coordinated effort may result in useable, effective and scalable software infrastructure.

The connected and ubiquitous synergy between HPC and Big data is expected to deliver the results which cannot be achieved by either alone. There is a need for the leading enterprises to use HPC technology to explore efficiently huge volume of heterogeneous data to surpass static searches into dynamic pattern discovery for the competitive advantage. The integration of computing power in HPC and demands for a quick and real time analytics for big data with cognitive technology (computer vision techniques, Machine learning, natural language processing) are considered as reshaping the future technology for accelerating analytics and deriving meaningful insights for efficient decision-making.

Acknowledgments. The authors acknowledge with thanks, the technical and financial support from the Deanship of Scientific Research (DSR) at the King Abdul-Aziz University (KAU), Jeddah, Saudi Arabia, under the grant number G-661-611-38. The work carried out in this paper is supported by the HPC Center at the King Abdul-Aziz University.

References

1. Singh, K., Kaur, R.: Hadoop: addressing challenges of big data. In: 2014 IEEE International Advance Computing Conference (IACC), pp. 686–689. IEEE (2014)
2. Charl, S.: IBM - HPC and HPDA for the Cognitive Journey with OpenPOWER. https://www-03.ibm.com/systems/power/solutions/bigdata-analytics/smartpaper/high-value-insights.html
3. Keable, C.: The convergence of High Performance Computing and Big Data – Ascent. https://ascent.atos.net/convergence-high-performance-computing-big-data/
4. Joseph, E., Sorensen, B.: IDC Update on How Big Data Is Redefining High Performance Computing. https://www.tacc.utexas.edu/documents/1084364/1136739/IDC+HPDA+Briefing+slides+10.21.2014_2.pdf
5. Geist, A., Lucas, R.: Whitepaper on the Major Computer Science Challenges at Exascale (2009)
6. Krishnan, S., Tatineni, M., Baru, C.: myHadoop-Hadoop-on-Demand on Traditional HPC Resources (2011)
7. Xuan, P., Denton, J., Ge, R., Srimani, P.K., Luo, F.: Big data analytics on traditional HPC infrastructure using two-level storage (2015)
8. Is Hadoop the New HPC. http://www.admin-magazine.com/HPC/Articles/Is-Hadoop-the-New-HPC
9. Katal, A., Wazid, M., Goudar, R.H.: Big data: issues, challenges, tools and good practices. In: 2013 Sixth International Conference on Contemporary Computing (IC3), pp. 404–409. IEEE (2013)
10. Hess, K.: Hadoop vs. Spark: The New Age of Big Data. http://www.datamation.com/data-center/hadoop-vs.-spark-the-new-age-of-big-data.html
11. Muhammad, J.: Is Apache Spark going to replace Hadoop? http://aptuz.com/blog/is-apache-spark-going-to-replace-hadoop/
12. OLCF Staff Writer: OLCF Group to Offer Spark On-Demand Data Analysis. https://www.olcf.ornl.gov/2016/03/29/olcf-group-to-offer-spark-on-demand-data-analysis/
13. Islam, N.S., Lu, X., Wasi-ur-Rahman, M., Shankar, D., Panda, D.K.: Triple-H: a hybrid approach to accelerate HDFS on HPC clusters with heterogeneous storage architecture. In: 2015 15th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing, pp. 101–110. IEEE (2015)
14. Ranger, C., Raghuraman, R., Penmetsa, A., Bradski, G., Kozyrakis, C.: Evaluating MapReduce for multi-core and multiprocessor systems. In: 2007 IEEE 13th International Symposium on High Performance Computer Architecture, pp. 13–24. IEEE (2007)
15. Tiwari, N., Sarkar, S., Bellur, U., Indrawan, M.: An empirical study of Hadoop's energy efficiency on a HPC cluster. Procedia Comput. Sci. **29**, 62–72 (2014)
16. Woodie, A.: Does InfiniBand Have a Future on Hadoop? http://www.datanami.com/2015/08/04/does-infiniband-have-a-future-on-hadoop/
17. Veiga, J., Exp, R.R., Taboada, G.L., Touri, J.: Analysis and Evaluation of Big Data Computing Solutions in an HPC Environment (2015)
18. Wang, Y., et al.: Assessing the performance impact of high-speed interconnects on MapReduce. In: Rabl, T., Poess, M., Baru, C., Jacobsen, H.-A. (eds.) WBDB-2012. LNCS, vol. 8163, pp. 148–163. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-642-53974-9_13
19. Islam, N.S., Lu, X., Wasi-ur-Rahman, M., Panda, D.K.: Can parallel replication benefit Hadoop distributed file system for high performance interconnects? In: 2013 IEEE 21st Annual Symposium on High-Performance Interconnects, pp. 75–78. IEEE (2013)

20. Moore, J., Chase, J., Ranganathan, P., Sharma, R.: Making scheduling cool: temperature-aware workload placement in data centers (2005)
21. Reed, D.A., Dongarra, J.: Exascale computing and big data. Commun. ACM **58**, 56–68 (2015)
22. Rajovic, N., Puzovic, N., Vilanova, L., Villavieja, C., Ramirez, A.: The low-power architecture approach towards exascale computing. In: Proceedings of the Second Workshop on Scalable Algorithms for Large-Scale Systems - ScalA 2011, p. 1. ACM Press, New York (2011)
23. Cappello, F.: Fault tolerance in petascale/exascale systems: current knowledge, challenges and research opportunities. Int. J. High Perform. Comput. Appl. **23**, 212–226 (2009)
24. Gutierrez, D.: The Convergence of Big Data and HPC – insideBIGDATA. https://insidebigdata.com/2016/10/25/the-convergence-of-big-data-and-hpc/
25. High Performance Data Analytics (HPDA) Market-Forecast 2022. https://www.marketresearchfuture.com/reports/high-performance-data-analytics-hpda-market
26. Willard, C.G., Snell, A., Segervall, L., Feldman, M.: Top Six Predictions for HPC in 2015 (2015)
27. Egham: Gartner Says 8.4 Billion Connected “Things”; Will Be in Use in 2017, Up 31 Percent From 2016. http://www.gartner.com/newsroom/id/3598917
28. El Baz, D.: IoT and the need for high performance computing. In: 2014 International Conference on Identification, Information and Knowledge in the Internet of Things, pp. 1–6. IEEE (2014)
29. Conway, S.: High Performance Data Analysis (HPDA): HPC - Big Data Convergence - insideHPC (2017)
30. Keutzer, K., Tim, M.: Our Pattern Language_Our Pattern Language (2016). Keutzer—EECS UC Berkeley, Tim—Intel. file:///Users/abdulmanan/Desktop/Our Pattern Language_Our Pattern Language.htm
31. Bodkin, R., Bodkin, R.: Big Data Patterns, pp. 1–23 (2017)
32. Mysore, D., Khupat, S., Jain, S.: Big data architecture and patterns, Part 1: Introduction to big data classification and architecture. https://www.ibm.com/developerworks/library/bd-archpatterns1/index.html