Name: Ali R. Butt
Affiliation: Virginia Tech
Email: `butta@cs.vt.edu`
Topics: 1. Cloud Architectures and Systems
3. Data Portability, Consistency, and Management

Summary of Current Cloud Computing Research Activities

I have been actively involved in cloud computing research for several years. My cloud research has two aspects: (i) designing and developing efficient and cost-effective cloud back-end systems tailored to specific application needs; and (ii) adapting the robust and dynamic cloud model to High-Performance Computing (HPC), especially for supporting the data-intensive components of emerging scientific applications. This research is being supported by the NSF, DOE's Oak Ridge Lab., and through industry collaborations (IBM, NEC, etc.).

For designing efficient cloud back-ends, we focused on MapReduce clusters. MapReduce has emerged as a model of choice for supporting modern data-intensive applications, and is a key enabler for cloud computing. Setting up and operating a large MapReduce cluster entails careful evaluation of various design choices and run-time parameters to achieve high efficiency. However, this design space has not been explored in detail. In this research, we developed an accurate Hadoop simulator, *MRPerf* (awarded Best Paper at MASCOTS'09), to systematically understand the performance of MapReduce setups and to predict expected application performance. Specifically, we employed *MRPerf* to explore and quantify the effect of critical design decisions—such as component inter-connect topologies, data locality, and software and hardware failures—on overall application performance. Thus, *MRPerf* serves as a tool for optimizing existing MapReduce clusters as well as designing new ones.

For adapting the cloud model to HPC, we initially focused on developing a robust and scalable data delivery service for data-intensive scientific applications. A critical research challenge lies in transporting input data to the HPC center from a number of distributed sources, and offloading the result data to geographically distributed, intermittently available end-users, often over under-provisioned connections. Such end-user data services are typically performed using point-to-point transfers that are designed for well-endowed sites and are unable to reconcile the center's resource usage and users' delivery deadlines, unable to adapt to changing dynamics in the end-to-end data path and are not fault-tolerant. To overcome these inefficiencies we developed CATCH, a Cloud-based Adaptive decentralized data Transfer serviCe for HPC (ICS'08, ICS'09, IPDPS'10, IPDPS'11), which leverages a bevy of cloud storage resources to orchestrate a decentralized data transport with fail-over capabilities. Our evaluation shows that CATCH can significantly improve the data transfer times at the HPC center, consequently improving HPC center serviceability.

In addition to the above, I am also involved in collaborative efforts to port critical applications, e.g., epidemics simulations, etc. to the cloud platform. Finally, I am engaged in the teaching and training of cloud technologies, and have designed and taught the first graduate course on Cloud Computing (CS6204) at CS@Virginia Tech in Fall'10.

Future Cloud Computing Research Activities

My current research with the cloud paradigm and interactions with cloud providers have exposed me to a number of experimental system issues, which I plan to explore. In the following, I give a brief overview of some key challenges.

Designing Efficient Cloud Systems Data locality plays a key role in achieving high performance for data-intensive applications. We plan to investigate techniques for improving data locality for given cloud back-end configurations and applications. To this end, one approach is to design task scheduling algorithms that are locality-aware and take into account when and where the needed data is available. An alternative is to explore dynamic data replication that can strategically increase or decrease the number of replicas in the system to create more choices for the scheduler and reduce constraints on where a job can run. However, this requires careful cost-benefit analysis of creating additional replicas, as well as understanding the role of back-end network topologies. We intend to extend and enhance our simulator, *MRPerf*, to enable such investigation and to realize tools that can be used for determining the best replication/scheduler combination suitable for different use-cases.

An additional advantage of systematically studying cloud back-ends is that it *enables realization of efficient private clouds*. With growing concerns about data security and privacy in the cloud, especially for mission-critical and sensitive data, such private clouds are often cited as the key to enabling the cloud model in highly-secure environments.

Integrating Cloud and HPC Modern HPC applications are composed of significant data-intensive components that are cloud-friendly along with traditional compute-intensive kernels. Given the cloud model's demonstrated capabilities in handling data-intensive enterprise applications, it is natural to explore this model in the context of HPC.

We are interested in exploring two aspects of integrating HPC and cloud computing paradigms. First, we want to explore the ability of such an integrated platform to efficiently execute data-intensive components of HPC applications on cloud resources, and conversely, compute-intensive components of cloud applications on HPC resources. *The critical research challenge lies in determining and coordinating the right computation and storage resources to allocate to a job when using both HPC and cloud infrastructure*. For instance, how to support both HPC jobs, e.g, MPI tasks, and cloud applications, e.g., MapReduce tasks, simultaneously on the next generation high-end systems. We plan to develop a unified scheduler to facilitate such cloud and HPC integration. Second, we want to investigate the ability of the platform to use cloud resources for improving performance of HPC tasks. To this end, we aim to extend our current work and develop a cloud-based structured data-management scheme and propose cloud-based solutions for handling data for HPC applications.

Enabling Cloud to use Accelerators Multicore processors with tightly coupled accelerators are becoming common and attractive to high-end systems such as the cloud, as they have the potential to sustain supercomputer-class node performance for dense computations, within a reasonable budget and energy cost. However, the state of knowledge on the use of accelerators in cloud computing is limited; typical cloud back-ends assume homogeneous resources and are not designed for the asymmetry offered by such accelerators. This poses several challenges. First, the effects of alternative workload distributions between general-purpose processors and accelerators are not well-understood. Second, accelerators have limited capabilities for managing external system services, such as I/O, thus requiring support from general-purpose processors in the runtime framework. Finally, in contrast to symmetric clusters, which enjoy a hoard of mature programming models such as MPI and MapReduce, the synthesis of programming models for asymmetric setups, which adapt to the varying capabilities of the components is an open problem. In particular, *hiding the architectural asymmetry from the programming model, and exploiting the vast computational density of accelerators while they communicate with inherently slower components remain major challenges*. We aim to design efficient cloud systems that can transparently utilize accelerators and provide seamless integration of multicore and asymmetric-core components into the cloud. In addition, we aim to extend programming approaches such as MapReduce to efficiently use asymmetric resources as needed by the applications.

Studying the Energy-Cost of Providing Data Reliability in the Cloud Despite the increasing importance of both data reliability and energy efficiency in storage systems, there has been little work investigating the interactions between the two. Cloud data centers have stringent data reliability requirements, and employ many techniques—redundant data copies, backup hardware—to prevent data loss. However, these techniques require more energy either to sustain additional hardware or to perform additional software tasks that keep disks busy longer. Conversely, there are a wide range of storage system designs that provide high performance and energy efficiency but do not consider reliability. Moreover, if reliability constraints are introduced in such energy-aware systems they may result in much higher energy consumption than originally intended. In this research, we plan to explore the energy-cost of providing such reliability, and develop energy-efficient data reliability improving techniques for the cloud.