Research Statement

Bryan Mills

Department of Computer Science, University of Pittsburgh

`bmills@cs.pitt.edu`

Computing and information systems evolved from single machines, to clusters of modular systems, to large-scale, networked systems of heterogeneous systems. This evolution was enabled by unprecedented progress in almost every aspect of computing, which paved the way to order of magnitude increase in CPU capabilities, storage capacity and communication speeds. This incredible growth, however, has come at the cost of increasingly sophisticated, complex systems whose behavior is increasingly difficult to specify, predict and manage. The complexity that is inherent in large-scale systems stems from a variety of causes. Large-scale systems typically encompass highly heterogeneous computing, storage and communications components, with varying degrees of capabilities and behavior consistency. They are often designed to support applications that cannot be completely specified. They are called upon to support various computational models ranging from massively parallel high performance computing, to utility computing over deeply virtualized cloud and data-centers infrastructure, to data-centric Internet services operating across multiple heterogeneous systems and management infrastructure, to tightly controlled, time-sensitive systems.

The complexity of these systems is expected to exceed todays systems complexity by every measure, including the number of computing nodes, the nature and quality-of-service requirements of the supported applications and services, the amount of data to be processed, stored, accessed and communicated to transform information into knowledge and the number of connections and interdependencies among the hardware and software components of the system. Furthermore, future computing and information systems must deliver a high-level of availability, dependability and performance. These systems, however, are likely to exhibit a level of interconnected complexity that makes it extremely difficult, if not impossible, to predict the effects of behavioral changes of individual components on the overall performance of the system. Consequently, system-level behavior emerges from parallel non-linear interaction of multiple system components that is difficult to define and specify at the time the system is designed and deployed.

There is a high risk that we would find ourselves reliant on information technology systems that are fragile and cannot deliver the services and performance for which they were engineered and call for new frameworks, methods and tools of large-scale, complex system design and development. Previous research has focused on developing models to study complex information technology infrastructures and gain better understanding of how their structure impact their emergent behaviors. In my research, I take a different prospective that assumes that unpredictable changes to system performance will occur and actively seek to develop uniform adaptive frameworks, methodologies and tools to achieve *system resilience* in future, large-scale computing and information systems. In this context, *system resilience* characterizes the ability of the system to mitigate the impact of and dynamically adapt to changing conditions and perform gracefully in face of disturbance in order to provide appropriate QoS support for diverse types of applications and services in heterogeneous computing environments.

Preliminary Research. My past research was concerned with the tradeoff between data availability and privacy of the data being stored. This work became my masters thesis in which I developed a new technique for scattering, concealing and recovering data, named SCAR [5, 4]. A main contribution of this work was a novel way of securing data in an un-trusted distributed storage system that does not rely on central authorities, while ensuring high availability and security. At the time, the idea of securely scattering data across multiple un-trusted data locations and maintaining availability and reliability was unique in itself.

The main tenet of SCAR is based on the observation that it is harder to break encrypted information if the attacker has no way of obtaining the encrypted data. SCAR achieves this using a simple, yet powerful concept of concealment through random distribution. The goal is to break data into pieces and randomly distribute the data throughout the network so that only authorized users can locate the pieces and recover the original data. Hence, our approach was to scatter data such that its locations are concealed, and ensure that only authorized users can recover the data being stored. In order to achieve this goal, SCAR combines hash chaining and erasure coding to ensure data privacy and data availability. This work provided an implemented in a distributed hash table (DHT). Furthermore, we provided an analytical model that describes the security and availability of data being stored using SCAR. These models are then used to study the inherent tradeoffs between data security and availability and show how SCAR attempts to balance these competing goals.

Current Research. Althought the race to build the worlds first exascale-class system has been underway for the last 10 years, two of the biggest challenges are power and resilience, each a direct result of the massive amount of parallelism necessary to achieve this goal. Delivering exascale performance could require a system with over a million sockets, each supporting many cores [1]. This would result in a system with many-millions of components including increases in memory modules, communication networks and storage devices. With this explosive growth in component count will come a sharp decrease in the overall system reliability and an increase in system power requirements. In my Ph.D. thesis, I show that it is possible to develop a solution that addresses both aforementioned challenges.

The research began with the study of power and energy consumption of existing coordinated and uncoordinated checkpointing and exploring possible enhancements. Through this work I developed a new power-aware replication technique called *shadow replication*. Using analytical models, simulations and experimentations we demonstrate that shadow replication provides system resilience more efficiently than checkpointing and traditional replication.

The basic idea of *shadow replication* is to associate with each process a suite of "shadow processes", whose size depends on the "criticality" and performance requirements of the underlying application. A shadow process is an exact replica of the main process. In order to overcome failure, the shadow is scheduled to execute concurrently with the main process, but at a different computing node. Furthermore, in order to minimize power, shadow processes initially execute at decreasingly lower processor speeds. The successful completion of the main process results in the immediate termination of all shadow processes. If the main process fails, however, the primary shadow process immediately takes over the role of the main process and resumes computation, possibly at an increased speed, in order to complete the task. Moreover, one among the remaining shadow processes is then promoted to be the primary shadow process. The main challenge in realizing the potential of the shadow replication stems from the need to compute the speed of execution of the main process and the speed of execution of its associated shadows, both before and after a failure occurs, so that the target response time is met, while minimizing energy consumption.

Since the failure of an individual component is much lower than the aggregate system failure, it is very likely that most of the time the main processes complete their execution successfully. Successful completion of a main process automatically results in immediatly halting its associated shadow processes, providing a significant savings in energy consumption. Furthermore, the number of shadow processes to be instantiated in order to achieve the desired level of fault-tolerance must be determined based on the likelihood that more than one process fails within the execution time interval of the main task. The completion of a main or its shadow results in the successful execution of the underlying task.

The contributions of my thesis fall into two main categories. The first is a better understanding of both the power and energy requirements of existing resilience methods. I provide a detailed power analysis of existing checkpointing methods and explore ways of conserving both power and energy during both the writing and restoring phase of checkpoints [3]. This understanding enabled us to make our second major contribution which is the development of a new power-aware fault tolerant method called *shadow replication* [6, 7]. I developed an analytical model to describe the energy and power consumption of this new technique and demonstrate the potential energy savings of this technique in an exascale system. I further developed a simulation framework again showing the potential energy savings achievable. I then implemented *shadow replication* within the Message Passing Interface (MPI) and provided experimental data to show the viability of this technique.

Future Research. In the broadest sense my future research will continue to better understand how to harness unreliable systems to provide reliable and predictable quality of service (QoS) to applications. This work spans several different computing areas, including high performance computing, cloud computing, big data analytics and mobile computing. The growing complexity of future computing and information systems threatens to undermine the very benefits they seek to provide. My primary focus will continue to be the development of system software and middleware to enable useful work to be efficiently completed in such environments.

Resilience and Power Consumption in High Performance Computing. I will continue to investigate how to enable extreme scale systems to reach the exascale computing goals given the predicted increase in system faults. In the next 5 to 7 years we will witness the development of high performance computing systems which will be plagued with both hard and soft failures inherent to such large parallel systems [1, 8, 9]. This research will focus upon solving the problem of system resilience for tightly coupled applications that characterize the workloads for the U. S. Department of Energy. I've collaborated in the past with researchsers at U. S. Department of Energy and will continue this collaboration for the foreseeable future. If power was not a concern, we could easily solve the resilience challenge by building more redundancy into the system, either at the hardware or software level [2]. I believe the challenge of resilience and power are intrinsically related to one another and need to be explored in unison. As a result, my work in system resilience will also be concerned with power and energy management. I suspect this work will also led to many other explorations involving system software and power management.

Soft Real-time Tightly-coupled Applications in Cloud Computing. One of the hallmarks of the cloud computing environment is its support of decoupled applications, which allows massive parallelism. This is ideal for tasks requiring timely but not necessarily precise responses such as indexing content and providing shopping suggestions. The inherent laxity in precision allows massively parallel applications to continue to respond in the face of failures and other system events. As these solutions begin to be applied to areas traditionally dominated by centralized database servers, such as medical analysis and banking, the results will need to become more accurate and precise while still demanding fast responses. This will result in tightly coupled real-time applications being executed in the inherently unreliable cloud environments. My work in high performance computing will enable me to explore these issues and offer valuable contributions to cloud computing research.

Shadow Computing. To enable efficient application execution in these faulty and unreliable environments I believe a new computational model will be required, which I call *Shadow Computing*. Instead of traditional replication or re-execution this new computational model will provide a goal-based adaptive execution in order to provide *system resilience* targeted at meeting the requirements of complex applications executing within complex systems. Adaptive execution is the ability of the system to dynamically harness all available resources to achieve the required level of QoS for a given application. Additionally, the application will have the ability to change its QoS requirements causing adaptive execution to adjust accordingly. The challenge is to maintain or exceed the applications QoS while minimizing the system resources in spite of systems-level changes, such as failures or the availability of additional system resources. In order to achieve adaptive execution, the shadow computing model would associate a set of shadows to the main execution, which are dynamically instantiated and adjusted in order to address the current state of the system and maintain the application's QoS requirements. Shadows are not necessarily processes but objects which can evolve into processes if necessary to achieve the applications requirements. My previous work in *shadow replication* [6] is a specific instance of this broader vision of a new computational model I call *Shadow Computing*.

Reliable Private Cloud-based Storage. Recent work in the area of Oblivious RAM (ORAM) [10] has caused me to revisit my previous on SCAR provided secure, private, distributed storage using DHTs. This more recent work attempts to solve the issue of obfuscating the data access patterns. My vision is to use modern storage systems such as GFS or Amazon's S3 with SCAR to provide truly secure data storage in the cloud. The work with ORAM will provide a new mechanism to solve the data tracking problem.

Highly Distributed Mobile Calculations. As processing power continues to increase on mobile devices there will be an opportunity to perform calculations in a cluster of mobile nodes. This marks an extreme version of moving the calculation to the data, in which the nodes collecting the data would collaborate with one another, in physical proximity, to perform complex calculations that any single node would be unable to complete. Such an environment would be plagued with reliability concerns and will require new programming models and systems software to provide the *system resilience* to deliver the required QoS to these applications.

References

[1] S. Ahern, A. Shoshani, K.-L. Ma, A. Choudhary, T. Critchlow, S. Klasky, V. Pascucci, J. Ahrens, E. W. Bethel, H. Childs, J. Huang, K. Joy, Q. Koziol, G. Lofstead, J. S. Meredith, K. Moreland, G. Ostrouchov, M. Papka, V. Vishwanath, M. Wolf, N. Wright, and K. Wu. Scientific discovery at the exascale, a report from the doe ascr 2011 workshop on exascale data management, analysis, and visualization, 2011.

[2] K. Ferreira, J. Stearley, J. H. Laros, III, R. Oldfield, K. Pedretti, R. Brightwell, R. Riesen, P. G. Bridges, and D. Arnold. Evaluating the viability of process replication reliability for exascale systems. In *Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '11, pages 44:1–44:12, New York, NY, USA, 2011. ACM.

[3] B. Mills, R. E. Grant, K. B. Ferreira, and R. Riesen. Evaluating energy saving for checkpoint/restart. In *First International Workshop on Energy Efficient Supercomputing (E2SC) in conjunction with SC13: The International Conference for High Performance Computing, Networking, Storage and Analysis*, Nov 2013.

[4] B. Mills and T. Znati. Increasing DHT data security by scattering data (invited paper). In *Computer Communications and Networks, 2008. ICCCN '08. Proceedings of 17th International Conference on*, pages 1–5, 2008.

[5] B. Mills and T. Znati. SCAR - scattering, concealing and recovering data within a DHT. In *Proceedings of the 41st Annual Simulation Symposium (Anss-41 2008)*, ANSS-41 '08, pages 35–42, Washington, DC, USA, 2008. IEEE Computer Society.

[6] B. Mills, T. Znati, and R. Melhem. Shadow computing: An energy-aware fault tolerant computing model. In *Proceedings of the International Conference on Computing, Networking and Communications (ICNC)*, 2014.

[7] B. Mills, T. Znati, R. Melhem, R. E. Grant, and K. B. Ferreira. Energy consumption of resilience mechanisms in large scale systems. In *Parallel, Distributed and Network-Based Processing (PDP), 22st Euromicro International Conference*, Feb 2014.

[8] S. R. Sachs. Tools for exascale computing: Challenges and strategies, 2011.

[9] B. Schroeder and G. A. Gibson. A large-scale study of failures in high-performance computing systems. In *Proceedings of the International Conference on Dependable Systems and Networks*, DSN '06, pages 249–258, Washington, DC, USA, 2006. IEEE Computer Society.

[10] E. Stefanov, M. van Dijk, E. Shi, C. Fletcher, L. Ren, X. Yu, and S. Devadas. Path oram: An extremely simple oblivious ram protocol. In *Proceedings of the 2013 ACM SIGSAC Conference on Computer and Communications Security*, CCS '13, pages 299–310, New York, NY, USA, 2013. ACM.