CSR:Small:Storage Architecture for the Next Generation of Smart Mobile Platforms

NSF Program 11-555

Activity Category: CSR

Umakishore Ramachandran
Moinuddin K. Qureshi

College of Computing
Georgia Institute of Technology
Atlanta, GA 30332-0280
Phone: (404) 894-5136
FAX: (404) 385-2295
e-mail: rama@cc.gatech.edu
WWW URL: http://www.cc.gatech.edu/~rama

Annual Report via Fastlane June 26, 2013

1 Goals of the Project

Mobile computing is sweeping the world. According to recent statistics, there are up to 4 billion mobile users (more than half of the world's population), with accelerated penetration in regions such as India and Africa. This trend is even more significant in the light of predictions about mobile devices dominating most personal computing landscape in the near future. Recent studies are pointing to the permanent storage in such devices being a weak link. Due to considerations of power and physical footprint, flash has been the storage technology of choice hitherto in mobile platforms. Newer storage technologies such as PCM, and STT-RAM are on the horizon as a competitor to flash. Each of these technologies have their own idiosyncracies both from the point of view of architecting them (from a hardware perspective) and integrating them in the operating system (OS) stack (from a system software perspective). It is time to take a critical look at the OS software stack and architect it to plan for future evolutions of the storage technologies. In this proposal, we focus on the system software and hardware issues, in an integrated fashion, in the construction of high-performance storage architecture for mobile platforms.

The goals of the project are summarized below:

- We will collect traces of applications (both current and futuristic) for mobile platforms and conduct detailed storage-centric measurements of their performance on current and emerging mobile storage technologies.
- We will analyze the source of storage-related performance bottlenecks on mobile platforms at different levels of the software stack (application, system software, and hardware) using a combination of measurement, simulation, and real implementation on existing and emerging mobile platforms.
- We will propose solutions both at the system software and architecture levels, in an integrated manner, for advancing the state-of-the-art in storage for mobile platforms.
- We will conduct detailed quantitative evaluation of our solutions both via simulation and implementation of research prototypes. In addition to validating our research ideas, this effort in of itself may

lead to new insights on how to perform quantitative studies that are closer to real implementation.

2 Activities and Results

2.1 Research and Education

We document below the research accomplishments in the first year of the project (August 15, 2012 to June 1, 2013).

2.1.1 System Studies on Mobile Storage

We have conducted studies on mobile platforms to gauge opportunities for improving the performance. One promising approach is to use the context of the mobile platform to perform "informed" optimizations on storage activities and hence boost application performance. A paper documenting this approach entitled, "*Fjord: Informed Storage Management for Smartphones*", appeared in **MSST 2013** [1].

Abstract of Fjord paper in MSST 2013: Smartphone applications are becoming more sophisticated and require high storage performance. Unfortunately, the OS storage software stack is not well engineered to support flash-based storage used in smartphones. On top of that, storage software stack is configured to be too conservative due to the fear of sudden power failures. We believe that this conservatism with respect to data reliability is misplaced considering that many of the popular apps (e.g., Web browsing, Facebook, Gmail) that run on todayŠs smartphones are cloud-backed, and the local storage on the smartphone is often used as a cache for cloud data. In this paper, we propose Informed Storage Management framework, named *Fjord*, for mobile platforms. The key insight is to use *system-wide dynamic context* information to improve the storage performance on mobile platforms. We implement a set of mechanisms (write buffering, logging, and fine-grained reliability control), and through judicious use of these mechanisms based on system context, we show how we can achieve significant improvement in storage performance. As proof of concept, we implement Fjord in two Android smartphones and experimentally validate the performance advantage of informed storage management with multiple smartphone applications.

2.1.2 Storage Access Traces on Mobile Platforms

Continuing the promising results of Fjord, we are undertaking a detailed study of several applications to understand their storage activities from the point of view of determining opportunities for relaxing reliability to gain performance. We are in the process of instrumenting the Android/Linux system stack on Nexus-7 platforms to intercept the storage system calls so that we can carry out such a comprehensive study. We expect to have the results of this study by the end of Summer 2013 or early Fall 2013.

2.1.3 Impact of Storage Class Memories on Memory Virtualization

Virtualization technology is a powerful software construct used in a variety of domains, ranging from cloud computing infrastructures and data centers to modern day Smartphones. One of the exciting opportunities, with the evolution of new storage class memories with access latencies closer to DRAM than the disk and fine-grain granularity of access similar to DRAM, is a re-evaluation of the entire memory management structure of operating systems. We are investigating the impact of storage class memories on virtualization technology in general and in particular for mobile platforms in the context of this proposal.

2.1.4 Multi-tiered Storage

Storage architecture of high-end servers are interesting to study to understand the role played by Flash memory in the storage hierarchy. It will also give us insight into architecting the storage of Smartphones. Video is the dominant workload on the Internet. Smartphones have become the *de facto* client devices for viewing video. One of the research questions we have been studying is the exploration the architecture of multi-tier storage that serve the video to the clients. A paper entitled, "*FlashStream: A multi-tiered storage architecture for HTTP video streaming incorporating MLC Flash Memory*", is to appear in **ACM Multimedia 2013** [2].

Abstract of FlashStream to Appear in ACM Multimedia 2013: We present FlashStream, a multi-tiered storage architecture for HTTP video streaming that incorporates inexpensive MLC Flash memory as a cache between RAM and hard disk to deliver high throughput at reduced deployment and operational costs compared to a disk-based system. The key architectural elements of FlashStream include object-level caching, optimal write granularity to overcome the small random write problem of flash, and a QoS-sensitive caching strategy that monitors the storage activity of the disk and the flash to ensure that video streaming performance is not hampered by the caching activity. We experimentally compare FlashStream with three different systems for HTTP video streaming: hard disks only, ZFS (from Sun Solaris), and Flashcache (from facebook). We show that FlashStream performs two times better than the best case performance of the three systems. Moreover, FlashStream reduces the total cost of operation (TCO) for a 3 year period by 33.3% for the same quality of service compared to an HDDs only system.

2.1.5 Mobile Embedded OS - A new Undergraduate Course

In Spring 2013, Professor Ramachandran offered a new junior undergraduate level course entitled, “Mobile Embedded OS" [5].

The course employs a new pedagogical style which we believe is ideal for project-oriented classes. The basic idea is to create a “sandbox" for students to play in. The class meets only *once a week for 3 hours* continuously. There is very little in-class lecturing by the professor. Instead the contact hours are structured around group interactions: (a) for students to work on the projects, debug, test, share experiences, seek peer help, etc.; (b) for the instructor and/or the TAs to present important conceptual or implementation details pertaining to the project being worked on by the students or for a future project; (c) for brainstorming to come up with novel “big project" ideas related to the course; (d) for students (by prior assignment by the instructor) to explain some specific subsystem pertaining to the course to the entire class. Put differently, the class is structured the way an industry team would work on a large product development project. The stress is on peer learning from fellow students. The principle is *collaborative* as opposed to *competitive* learning. Students are assessed based on completing their projects and demonstrating that they completely understand what they have built to the satisfaction of the instructor. There is no partial credit, it is all or nothing. But the students can seek help from their peers and the TAs to successfully complete the project.

Mobile Embedded OS is an exemplar for a course that uses this pedagogical style. The course is a deep immersion into kernel hacking for embedded platforms. The students are assumed to have the conceptual knowledge of systems (architecture and OS) from prior course-work. The course uses an embedded platform (such as Nexus-7) and the Android stack for the experimental platform. The course is structured around a set of projects, each of which involves some significant changes to the Linux Kernel which is at the base of the Android stack. Flavor of projects include developing device drivers for onboard sensors in the Nexus-7 platform, wireless protocol stack, modifying and/or extracting information from kernel subsystems (e.g., memory manager, scheduler) and exporting it to the application. Using Android as the basis for this course is a deliberate choice since it is an open commercially supported software platform. Further, it is a complete software stack from kernel level mechanisms to higher level abstractions for robust application development.

The course serves as a recruiting ground for identifying undergraduates who are passionate about kernel hacking. Indeed it is from this course that we have recruited Daniel Whatley and Andrew Wilder to work in our lab this summer.

3 Training and Development

3.1 Graduate Students

We currently have five students associated with the project though only 2 of them is supported by the funding. *Wonhee Cho* and *Yeonju Jeong* are the students funded from this grant. In addition to them, three more students contribute to the project. Here is a complete list and their work assignments:

1. Wonhee Cho and Yeonju Jeong are the lead students on the project, and are working on software support for increasing the efficiency and lifetime of emerging storage class memories, which is the technology of choice especially for Smartphones.
2. Dushmanta is working on virtualization technologies for personal mobile platforms that has implications for the storage architecture of Smartphones.
3. Moonkyung Ryu is working on P2P video streaming algorithms and multi-tiered storage architecture for CDN servers that help characterize the workloads on Smartphones and their implications on the system software stack for Smartphones.
4. Lateef Yusuf is working on constructing transient social networks using Smartphones. Given that

social networks drive the use of Smarphones, his research plays a vital role in informing the kinds of applications and their storage access patterns on Smartphones, which are essential for the proposed research.

We are also connecting these students with appropriate internship opportunities as noted below.

3.2 Undergraduates

Currently, we have two undergraduates, Andrew Wilder Daniel Whatley working (for course credit) under the guidance of my PhD students (Wonhee and Yeonju) to instrument the Linux kernel under Android for monitoring the storage access calls made by the applications. Another undergraduate Nick Olive (supported by NSF REU) is working under the guidance of my PhD student (Moonkyung Ryu) to study multi-tier storage.

3.3 Internships

We have successfully placed many of our PhD students on internships in Summer 2013 that are relevant to the proposed research.

Wonhee Cho. He is currently working as an intern for *American Megatrends Inc.*. His work entails dual OS implementation on Samsung tablets to support both Android and Windows-8.

Yeonju Jeong. She is also currently working as an intern for *American Megatrends Inc.*. Her work entails porting gaming applications to run on top of dual OS (Android and Windows-8) on Samsung tablets.

Dave Lillethun. He is working as an intern for CISCO, where he is researching issues related to supporting latency sensitive applications (such as video streaming) on mobile platforms.

Kirak Hong. He is working as an intern for Google, where he is researching programming paradigms for live interactive applications targeting mobile platforms as end devices.

4 Outreach Activities

4.1 Industrial Collaboration

VMware is very interested in aspects of this research and has given an unrestricted gift jointly to Professors Ramachandran and Qureshi in support of their research entitled, “Rethinking Operating System Structure with Heterogeneous Main Memory." This collaboration will enhance the research output from this project beyond the NSF funding.

Dr. Qureshi spent several years at IBM Watson Research before joining the faculty at Georgia Tech. He has strong connections with them on PCM technology. These connections will help further the quality of the research results we will be able to produce from this project.

4.2 International Collaboration

Dr. Ramachandran has ongoing collaborative project with Professor Kurt Rothermel of Institute for Parallel and Distributed Systems (IPVS, `www.ipvs.uni-stuttgart.de`) entitled, “Complex Event Processing in the large.” This project is funded by the state of Baden Württemberg, Germany for four years from 2010-2013. Dr. Ramachandran plays a consultative role in the project, which enables him to spend time at IPVS every year. He spent two weeks (June 1-16, 2013) as a guest of the institute.

Through the German Academic Exchange Service (DAAD), Professors Rothermel and Ramachandran have secured funding for three graduate students (every year) from the University of Stuttgart to visit Georgia Tech for a year-long stay involving course and project work, which can be used towards their MasterŠs degree from the University of Stuttgart. The first batch of students will arrive in Fall 2013. Involvement of such students in this project will enhance the intellectual output beyond what is possible with the available budget that we have asked for.

4.3 CERCS

Georgia Tech has an NSF I/UCRC grant administered by the the Center for Experimental Research in Computer Systems (CERCS). We have an advisory board meeting for CERCS that comprises researchers from leading industries. We have a board meeting once every 6 months as part of this I/UCRC award. One of the fixed items on the agenda during these meetings is a poster and demo session by the students. This serves as a great opportunity for students on this project to meet and discuss their research with leading researchers from industries.

4.4 Open Source Software

Through the efforts of one of the students who graduated recently from our group (Dr. Hyojun Kim, currently with IBM Almaden Research), we have been successful in open-sourcing our *FlashFire* [3, 4] software, which to date has a user community of over 100,000.

4.5 Harman Innovation Council

PI Ramachandran is a member of the **innovation council** of Harman International (`www.harman.com`).

5 Plans for the Second Year

We plan to continue research as proposed. The second year plans includes:

Measurements and Analysis of Storage Access Traces of Apps on Android. The results from this task will serve as a guide in identifying opportunities for sacrificing storage reliability to gain performance without sacrificing correctness of apps running on Smartphones. This study will also help to prune and select the design space exploration for both software and hardware solutions in the context of emerging storage class memories.

System Software Investigations and Architectural Innovations. We will continue working on these two tasks in an integrated manner so that the solutions reinforce each other. The evaluation methodology (next task) will be heavily used in quantifying the efficacy of the solutions explored.

Performance Evaluation Studies. This task will be ongoing through the lifetime of the project. The methodologies developed through this task will be used as a general framework in the entire project.

6 Publications and Products

6.1 Publications

See the *highlighted* references at the end of this document for publications that appeared or in preparation as a consequence of this grant.

7 Contributions

The activities we are currently undertaking are starting to bear fruits and there are papers in the publication pipeline. These are listed in the references (highlighted) at the end of this report.

7.1 Human Resources

There are five graduate students and three undergraduates associated with this grant.

8 Special Requirements

None.

References

[1] Hyojun Kim and Umakishore Ramachandran. Fjord: Informed storage management for smartphones. In *MSST 2013: 29th IEEE Symposium on Massive Storage Systems and Technologies*, May 2013.

[2] Moonkyung Ryu and Umakishore Ramachandran. FlashStream: A multi-tiered storage architecture for HTTP video streaming incorporating MLC flash memory. In *ACM Multimedia 2013: The 21st ACM International Conference on Multimedia*, October 2013. To Appear.

[3] Hyojun Kim. FlashFire SSD Accelerator. `http://flashfire.org/xe/`.

[4] Hyojun Kim and Umakishore Ramachandran. Flashfire: Overcoming the performance bottleneck of flash storage technology. Technical Report GT-CS-10-20, `http://smartech.gatech.edu/`

handle/1853/36335, School of Computer Science, College of Computing, Georgia Institute of Technology, 2010.

[5] CS 4803 - Mobile Embedded OS, Spring 2013. http://www.cc.gatech.edu/~rama/CS4803-External/.