INVITED TALK: SELF-AWARENESS AND ADAPTIVE TECHNOLOGIES: ARE THEY THE FUTURE OF OPERATING SYSTEMS?

Lamia Youseff

Google, Inc.
email: lyouseff@google.com

Talk summary

Trends in multicore architectures point to an ever-increasing number of cores available on a single chip. Moore's Law predicts an exponential increase in integrated circuit density; in the past, this increase in circuit density has translated into higher single-stream performance, but recently single-stream performance has plateaued and industry has turned to adding cores to increase processor performance. In a few years, multicores have gone from esoteric to commonplace.

Given exponential scaling, it will not be long before chips with hundreds of cores are standard, with thousands of cores following close behind. This new architecture trend is providing an exciting opportunity for exploring different research directions in scaling operating systems. At the same time, a similar, independent trend can be seen in the growth of cloud computing. Rather than consolidating a large number of cores on a single chip, cloud computing consolidates multicore machines within a data center. There is much commonality between constructing OSs for clouds and multicores, such as large-scale resource management, heterogeneity, and possible lack of widespread shared memory [1]. These similarities allow operating systems to be designed for both multicore and cloud computers.

The primary question facing OS designers over the next ten years will be: *What is the correct design of OS services that will scale up to hundreds or thousands of cores?* We argue that the structure of monolithic OSs is fundamentally limited in how they can address this problem. In contrast, our work explored a new factored structure for the OS, which we dubbed *fos* for "factored operating system" [2]. The structure of fos brings scalability concerns to the forefront by decomposing an OS into services, and then parallelizing within each service. To facilitate the conscious consideration of scalability, fos system services are moved into userspace and connected via messaging. In fos, a set of cooperating system servers which implement a single system service is called a *fleet*. However, a fundamental research challenge in this design is to identify the characteristics of such a fleet. Given the unprecedented variability in demand of the system resources, the OS fleets have to deploy elastic techniques to adapt to this variability at runtime.

We argue that the OS services have to deploy elastic techniques to adapt to this variability at runtime. In this talk, we advocate for elastic OS services, illustrate their feasibility and effectiveness in meeting the variable demands through our prototype system, dubbed *elastic fos* or "*e-fos*", which provides elastic technologies for OS services in the fos operating system [3]. We furthermore showcase a prototype elastic file system service in *e-fos* and illustrate its effectivness in meeting variable demands.

About the speaker

Dr. Youseff holds a Ph.D. degree in Computer Science from the University of California, Santa Barbara and has received her post-doctoral training at MIT working at CSAIL with Professor Anant Agarwal on research in cloud computing, operating systems and next generation exascale computing paradigm. Dr. Youseff has received several awards, including the international ACM-UPE'02, the AUC presidential cup'03, XHPC best paper award'06 and IPDPS best poster award'08 at the TCPP forum. In addition, she has served as a program committee member, program co-chair and organizer to several top conferences and workshops in computing systems, including SOCC, DAC and USENIX Middleware. She also has tens of technical publications and book chapters with more than 1000 non-self citations. She is currently a research software engineer in cloud computing at Google Seattle, WA office. Dr. Youseff worked on fos and e-fos with the fos Carbon group between 2009 and 2011 when she was affiliated with CSAIL, MIT.

1. REFERENCES

[1] K. Modzelewski, J. Miller, A. Belay, N. Beckmann, C. Gruenwald III, D. Wentzlaff, L. Youseff, and A. Agarwal, "A Unified Operating System for Clouds and Manycore: fos." [Online]. Available: http://mit.dspace.org/handle/1721.1/49844

[2] D. Wentzlaff, C. Gruenwald, III, N. Beckmann, K. Modzelewski, A. Belay, L. Youseff, J. Miller, and A. Agarwal, "An operating system for multicore and clouds: mechanisms and implementation," in *Proceedings of the 1st ACM symposium on Cloud computing*, ser. SoCC '10. ACM, 2010, pp. 3–14.

[3] L. Youseff, N. Beckmann, H. Kasture, C. Gruenwald, D. Wentzlaff, and A. Agarwal, "The case for elastic operating system services in fos," in *Proc. of the 49th Annual Design Automation Conference*, ser. DAC '12. ACM, 2012, pp. 265–270.