Data-centric Systems

Vijay Nagarajan and Boris Grot

January 2017

1 Summary

System design has, until now, been compute centric. The focus has always been on speeding up computation within each layer of the system stack. However, the advent of Big Data has given rise to new memory-intensive applications, such as graph processing, analytics and search, where data movement (rather than computation) is consuming the bulk of the time and energy.

Meanwhile, the breakdown of traditional technology scaling (i.e., Moore's law and Dennard Scaling) has coincided with the emergence of several disruptive memory and communication technologies, including: fast non-volatile memory (NVM), that has blurred the the difference between storage and memory; 3D stacking, that has enabled massive memory bandwidth and has also provided a vehicle to bring computation physically close to the memory elements in a form of *near-memory processing*; and chip-integrated photonic interfaces, that has enabled ultra-low-latency and high-bandwidth communication.

The combination of these technological considerations, combined with the rapid growth in data volume and the demand to process and serve this data at online speeds, is driving a need to design, from the ground up, a new class of highly distributed *data-centric systems*. Critical questions that need to be answered include the following.

- What should the new memory hierarchy look like?
- What form should file system and networking abstractions take in the age of super-fast storage and interconnect?
- How should data-centric systems be programmed?
- How to ensure consistency and availability in a cost-effective and high-performance manner?
- The collapse of familiar abstractions would significantly increase system design and verification complexity. This motivates an approach in which verification should very much be a part of the design process and not a post-hoc concern. How to manage design complexity and verify these future systems?

Clearly, answers to the above questions require a vertical approach spanning programming languages and verification down to systems and architec-

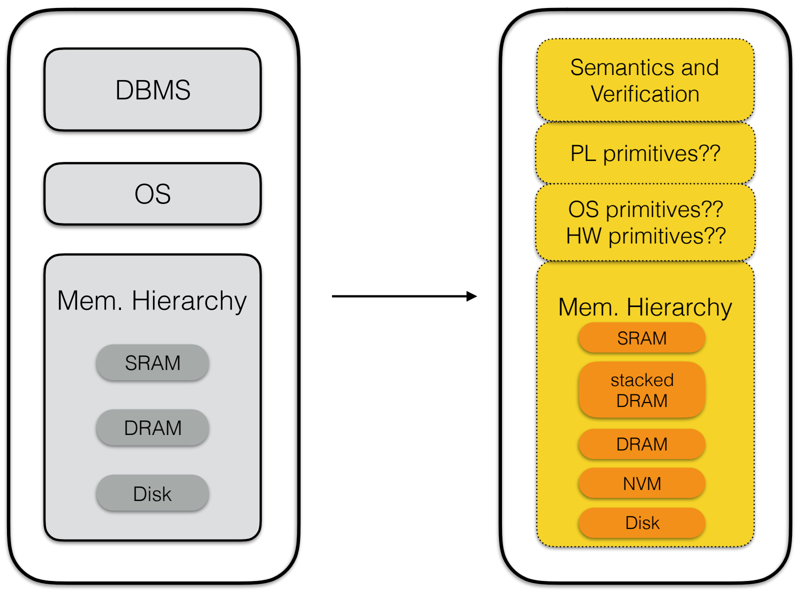

Figure 1: Transitioning from a compute-centric approach to a data-centric approach.

ture. Figure 1 shows how today's compute-centric system stack should evolve into tomorrow's vertically-integrated data-centric one.

2 Researchers

This research programme spans researchers working in different layers of the stack including Computer Architecture (Cintra, Grot, Nagarajan, Topham), Systems (Bhatotia, Sarkar), Programming Languages (Cheney, Cole, Dubach, O'Boyle, Wadler), Databases (Buneman, Fan, Libkin, Viglas), Semantics and Verification (Fleuriot, Jackson, Stark), across LFCS, ICSA and CISA.

3 Activities

We plan to hold monthly reading group meetings involving all of the above researchers and their students. We also plan to hold semi-annual 1-day research retreats in a topic of common interest to all of the researchers. For example, data replication and consistency is of common interest to researchers across the stack – with each community having its own terminology and approach. This would be a great opportunity to speak a common vocabulary and exchange ideas. The purpose of research retreat is to ensure that the researchers actually talk to each other without distractions; a single day (within reach from Edinburgh) will ensure costs are kept minimal. We also plan to invite at least one external speaker to each of these retreats. It is worth noting that this area has already started to receive interest in the form of external visits and sabbaticals from internationally leading researchers: Prof. Dan Sorin (Duke University) and Dr. Paul Gratz (Texas A and M) both working on memory systems, are visiting Edinburgh starting July 2017 for one year sabbaticals.

4 Outcome

This research programme will force researchers from databases, systems, computer architecture, programming languages and semantics to interact, paving the way for future researchers from these communities to understand each other and avoid replicated work. More ambitiously, it could allow for the development of a cross-stack, verified, performant data system (e.g. Key-value store) which requires the pooling of expertise in each of the above areas, something that industry has succeeded in doing but not academia. With significant disruptions in both the application level and technology, a rethink of the system stack is inevitable and Informatics is in a great position to be involved in this high impact endeavour. Finally, this research programme is aligned with two of EPSRC's new cross-ICT priorities: *Data-Enabled Decision Making* and *Cross-Disciplinarity and Co-Creating.*