

Towards Systematic Data Center Design

Avrilia Floratou
IBM Almaden Research Center
650 Harry Road
San Jose, California
aflorat@us.ibm.com

Data centers are changing fast. The rapid increase of reliance on cloud computing, particularly in the context of “big data” applications, inevitably leads to the following fact: data centers must constantly update and adapt to the increased user requirements. This fact becomes more evident if we consider today’s cloud environments. The users of these environments have very strict requirements: they expect to have access to specific hardware resources (IOPs, memory capacity, number of CPU cores, network bandwidth), they demand data availability and durability guarantees, defined quantitatively in Service-Level Agreements (SLAs) [1, 2], and they will soon expect to get concrete performance guarantees defined in performance-based SLAs (e.g. [4, 5]). Given these user expectations, the natural question to ask is the following: “*How should we design the data centers of the future?*”

Data center design is a tedious and expensive process. Typically, data center designers determine the hardware setup based on the software that will be deployed, and the expected workload characteristics. For example, in an environment where high availability is important, the data is replicated (e.g., three or five times), and storage is provisioned to accommodate the number of copies. This process leads to an **iterative** approach to data center design: the software configuration is determined first and then the hardware setup (or vice versa). However, such approaches may not always be optimal since they ignore important interactions between these two components [3]. What we propose is an **integrated** approach that explores the hardware/software interdependencies during the data center design process.

We believe that the most cost-effective, extensible and efficient way of performing **integrated** data center design is to use a “wind tunnel” to simulate the entire space of hardware/software co-design. The “wind tunnel” can be used to *explore, experiment, and compare* design choices, and to *reason* about the statistical guarantees that these design choices can provide. There are a lot of interesting research challenges involved in the process of designing, building and using the simulation-based “wind tunnel”. Designing lan-

guages to express queries in the data center design space, making the “wind tunnel” scalable in the number of design variables explored, validating the simulator, managing the simulation data and figuring out how to model the various software and hardware components are necessary steps towards systematic data center design.

We believe that the process of building the simulation-based “wind tunnel” shares similarities with building a specialized data processing system. The similarities are observed end-to-end, and they cover the language, the optimization and the runtime components. From a language perspective, the “wind tunnel” queries bear resemblance to traditional queries issued by database users. For example, a “wind tunnel” user may request the data center configurations that produce a data availability distribution “similar” to a desired one. From a query execution perspective, simulating multiple parts of a data center is likely to be time-consuming. It is possible that a large number of simulation runs may be needed in order to cover the whole configuration space defined in the user’s “wind tunnel” query. We may need to borrow techniques used to scale database queries (dynamic query optimization, parallelization) and apply them in the context of the “wind tunnel” execution. The crucial step of validation of the simulation would benefit from the use of publicly available datasets, and, to this end, we highly encourage researchers and practitioners to invest the time to produce datasets from real-world deployments. Finally, the data management community should work closely with the systems and architecture communities to accurately model the hardware and software components of interest.

We believe that this *systematic approach to data center design* is an exciting new research area with significant implications for future cloud services and potentially fruitful collaboration between diverse research communities.

1. REFERENCES

- [1] Amazon EC2 SLA. <http://aws.amazon.com/ec2-sla/>.
- [2] Windows Azure SLAs. <http://www.windowsazure.com/en-us/support/legal/sla/>.
- [3] A. Floratou, F. Bertsch, J. M. Patel, and G. Laskaris. Towards Building Wind Tunnels for Data Center Design. *PVLDB*, 7(9):781–784, 2014.
- [4] B. Mozafari, C. Curino, and S. Madden. DBSeer: Resource and Performance Prediction for Building a Next Generation Database Cloud. In *CIDR*, 2013.
- [5] V. R. Narasayya, S. Das, M. Syamala, B. Chandramouli, and S. Chaudhuri. SQLVM: Performance Isolation in Multi-Tenant Relational Database-as-a-Service. In *CIDR*, 2013.

This article is published under a Creative Commons Attribution License (<http://creativecommons.org/licenses/by/3.0/>), which permits distribution and reproduction in any medium as well allowing derivative works, provided that you attribute the original work to the author(s) and CIDR 2015.

7th Biennial Conference on Innovative Data Systems Research (CIDR ’15) January 4-7, 2015, Asilomar, California, USA.