

HW Databases, Where art Thou?

Sangjin Han Aurojit Panda Sylvia Ratnasamy
UC Berkeley
{sangjin, apanda, sylvia}@cs.berkeley.edu

ABSTRACT

Information about computer hardware is invaluable not only for enterprises and consumers, but also for the research community [2]. A lot of hardware components and systems, such as microprocessors, memory, disks, network links, servers, and consumer devices, appear and disappear. Archives for the technical specifications of such hardware allow us to understand the historical trajectory of technological advancement, identify urgent problems we face today, and predict the future by extrapolating from current trends, for example:

- Are the number of cores really doubling every x years?
- Are mobile devices evolving faster than workstations?
- How to make trade-offs between DRAM bandwidth, capacity, and latency?
- Are the assumptions made by a paper still applicable at present, and with a larger dataset?

One can easily find information about particular hardware on the Internet. Unfortunately, old hardware is often missing, the information is fragmented across multiple sources (e.g., Wikipedia, online markets, datasheets, and vendor websites), and the unstructured data hinders systematic analysis. Usually data collected by individual researchers is not public or kept up-to-date. Prior work such as CPU DB [1] has attempted to lay a foundation, however much remains unresolved.

We propose building a collaborative store for such information in an open-source approach. The data should be machine readable (through an API) for both analysis and decision making for automated tools (e.g., optimized code generation for target systems). The human readability is meant to aid with maintaining the database, and allows people to perform point queries with relative ease.

BODY

We need an open HW specification database that is comprehensive, extensible, maintainable, and accessible by both humans and machines.

REFERENCES

- [1] A. Danowitz, K. Kelley, J. Mao, J. P. Stevenson, and M. Horowitz. Cpu db: recording microprocessor history. *Commun. ACM*, 55(4), Apr. 2012.
- [2] J. Gray. What next?: A dozen information-technology research goals. *Journal of the ACM (JACM)*, 50(1), 2003.

Volume 2 of Tiny Transactions on Computer Science

This content is released under the Creative Commons Attribution-NonCommercial ShareAlike License. Permission to make digital or hard copies of all or part of this work is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page.
CC BY-NC-SA 3.0: <http://creativecommons.org/licenses/by-nc-sa/3.0/>.