

Move Bits not Watts: Algorithms for Sustainable Data Centers

By Adam Wierman, *Professor of Computer Science*

I was a latecomer to computer science. When I started as an undergraduate at Carnegie Mellon University, I was a civil engineering major. Then I went to psychology, statistics, math, and finally I found what I was really passionate about: computer science. There was an amazingly engaging professor, Steven Rudich, who got me excited about the fundamental challenges in the area, and the perspective of “focusing on the fundamentals” has defined my work since.

These days, my primary focus has been on finding places where computer science can help to address the problem of burgeoning energy consumption. Computer science can play a major role in this area, but it’s also a major part of the problem.

Research in this area puts me at the boundary of several fields, such as economics, electrical engineering, and control, because one ends up using tools from many different areas to gain a clearer understanding of these concerns.

I’m particularly interested in the effects and potential of cloud computing, where so much of our lives seem to have migrated. The data centers that make the cloud work the way it does are huge energy users: there are about 2,000 medium to large data centers in the United States, and those 2,000 buildings make up 2 to 3% of the country’s

energy usage. That’s still a small percentage today, but energy usage in data centers is growing at about 10% a year, while the energy usage of the United States as a whole is growing at about 1% a year. The problem is that their servers are basically always on. There are 10,000 or 100,000 servers sitting there idling at 10% capacity most of the time.

My students and I started out just thinking about how we could make data centers more efficient. Can we use renewable energy in powering them? Can we make them more efficient at using renewable energy?

A key observation that guides our research is that not all the work that the cloud is doing is email, search, and other such things where an immediate answer is required. A lot of it is “delay-tolerant,” like scientific computing, where if it’s going to take a week to do a simulation, what matters to the user is that it’s done in about a week, not that it’s done in a week and a minute versus a week and ten minutes. That affords the flexibility to run the computations when it’s sunny out or when it’s windy out, or when the grid sends a signal that there’s a huge demand because it’s a hot day.

But this approach requires a lot of dynamic control of the workload, of the servers, and of the cooling for them. That’s a big challenge, and it’s terrifying to the data center operators because they care about reliability. If Gmail or Netflix or Flickr goes down for 10 minutes, that’s a disaster. One has to be very careful not to sacrifice reliability, which makes getting flexibility out of the services a really challenging problem.

What we’re doing is designing the algorithms that determine when to

move work from server to server in a data center, when to turn a server off or on, what power state the server should be running in. It’s actually a very dangerous decision—the costs involved in turning something off are nearly the same as the cost of leaving it running for an hour or two. One way to think about this is as a rent-or-buy problem. If there’s not enough workload to keep it busy, the server should be turned off to take advantage of the savings in cost and energy. But if the server is uncertain whether it might be needed, great care must be exercised in turning it off. So, keeping it on is like renting and turning it off is like buying. Since the cost of “buying” is so high, it becomes essential to make the choice carefully. And, in reality, it’s not just a binary decision of renting or buying. One has very fine-grained control of which servers to turn on, where to keep data depending on which servers are kept on, which cooling systems are kept on—and everything’s correlated over time, so things have to be in certain orders for certain jobs to work. It becomes a very complicated decision about which things to turn off and when.



Our goal is to design algorithms that we can take to HP and Apple and Google and say: Using this algorithm will manage your capacity in a sophisticated way so that when you have a solar farm next to your data center, you can take advantage of that to save money on the grid, and be net zero, or close to it, by using renewable energy as much as possible.

In fact, we’ve been working with HP for three years, and they’re now pretty convinced. They have our algorithms implemented as part of their “net-zero data center architecture,” and so the ideas have made the initial transition from academia to industry.

There’s no way we could have made that kind of progress if my Caltech students hadn’t connected with HP. That was pivotal, and it’s been great to have the Caltech community really help in making those connections. Our work is quite mathematical, and so it took a lot of effort from the students to convince HP

that the algorithms we had developed actually made sense for their system. Having students go onsite was what made the transition out of academia possible, because there’s nothing we can do here that will convince them that for their data center, something like this can work. They need to see that even though it’s their architecture, their design, and the things they do are specialized, the models still apply.

Going forward, there’s still a lot to do. We think that data centers can actually be a key to helping integrate renewable energy more efficiently into the grid itself. The problem with renewable energy is that it fluctuates. In a grid, you have to match demand—which you basically don’t have any control over—with supply at every given instant, and that’s really hard if you can’t predict the availability of wind energy and solar energy. But data centers, if they’re sophisticated in the way we’ve been

talking about, can give you some control over demand, because you can say to a data center, we need an adjustment of demand of a megawatt to help balance our energy sources. According to HP and some smaller companies that we’re working with, they can very easily give 10 to 25% flexibility in their energy usage at any given point during the day, which means that in a 20-megawatt data center, you basically have two to five megawatts of storage. That’s like having a two-megawatt battery that grid operators could just plug and play and control, if the market is set up so that it makes sense for the data center to provide this flexibility to the system. And this is where economics comes in: How can we design markets to extract this flexibility? So, at this point, we’re working on both the control schemes for the data centers and market design for demand response to try to understand how they can work together.

One of the harshest realities of going from the data-center world to the electricity-market world is the difference between talking to engineers about how to design a system and trying to have a policy impact on how markets are regulated. There is just a complete difference in how changes are realized. In the data center, a test bed can show that things are working, but there’s no parallel to this in the policy arena.

But the outcome, as I tell the students when they start, is that if we can make it possible for data centers to provide such services for the grid, that would basically save a few power plants. That’s a very different form of impact than a computer scientist typically has. It’s not just that people will use your system, but you can have an impact on a crucial challenge facing society. ■ ■ ■

Adam C. Wierman is Professor of Computer Science.

Visit rsrg.cms.caltech.edu.