C3-lab: Integrated Innovation for de-carbonizing datacenters

Pat Pannuto[1], George Papen[1], George Porter[1], Barath Raghavan[2], Aaron Schulman[1], and Alex C. Snoeren[1]

[1] *The University of California San Diego, La Jolla, Calif.*
[2] *The University of Southern California, Los Angeles, Calif.*

https://c3lab.net/
Point of contact: George Porter, gmporter@ucsd.edu

November 8, 2021

The Internet Datacenters that power more and more aspects of our life are responsible for a large–and rapidly increasing–share of US carbon emissions. Information and Communication Technology (ICT) accounts for 4% of global carbon emissions today, and that number is expected to double by 2025 [2]. These carbon emissions stem not just from powering the operations of datacenters, but from the supply chains providing the equipment hosted within them. The integrated research vision of the C^3 Lab is to develop a foundation for technological innovations encompassing hardware, software, algorithms, and curricula to enable a realistic achievable path to a zero-carbon cloud infrastructure.

1 Introduction

The recent IPCC report is stark in its conclusions about the impact of human activity on the climate and the need to decarbonize our economy and society. As a source of 4% of the global carbon emissions, ICT–and in particular cloud computing datacenters–must reduce their carbon footprint to help achieve worldwide carbon reduction targets necessary to avoid catastrophic effects of climate change.

In response, the ICT community has seen a slow and steady shift toward greater computing efficiency (e.g. through improvements in datacenter PUE efficiencies), investment in renewable energy (e.g. Power Purchase Agreements and a large build-out in solar, wind, and geothermal power plants), and carbon offsets. Indeed the power grid is rapidly increasing its share of low-carbon energy sources. However, even with this substantial progress, an enormous carbon footprint remains in this rapidly growing sector of the economy necessitating ground-breaking innovations to truly achieve carbonless cloud computing.

The Carbonless Cloud Computing Lab (C^3-Lab) is being established to tackle this problem. C^3-Lab begins with several premises that demand entirely new research advances to achieve a truly zero-carbon cloud: a surfeit of periodically-available renewable power on the grid, a growing physical footprint of computing, and an urgency to achieve total industry-wide decarbonization by the end of the decade. Given these premises, the existing trajectory of decarbonization with the ICT community is insufficient to meet this goal with current approaches producing diminishing returns.

To radically change the existing rate of decarbonization, the lab will study novel architectural responses at every level: from the hardware architecture of individual servers to the architecture of geographically-distributed networked systems and the composition of their services. The Lab will consider *all* elements in the life cycle of computing. Taking a cradle to cradle perspective, our goal is to more effectively amortize the initial carbon footprint of equipment by dramatically extending the productive lifecycle of hardware. Further, the Lab will consider how fundamentally new cloud service offerings (from new tiers to new abstractions to new primitives) can enable and incentivize a shift towards zero-carbon cloud computing.

2 Trends and barriers

Our work is predicated on two trends which will force a fundamental rethinking of how we design, build, and deploy datacenter systems.

Trend 1: The accelerated deployment of renewable energy sources (e.g. solar, wind, geothermal) will result in a dramatic increase in the amount of carbon-free power available. However, as can be seen in Figure 1, this renewable power is intermittent and volatile, varying considerably based on environmental conditions and demands on the local grid. This upends traditional assumptions of static datacenter power levels being highly stable over long timeframes [1].

Trend 2: Today about 25% of the total carbon impact of ICT equipment arises during the manufacturing process (called "Scope 3"), before the device is even put into service. As the energy grid incorporates an increasing share of low-carbon energy sources, the proportion of the *embodied* carbon from manufacturing in the overall carbon footprint will increase [3]. In the limit, a fully decarbonized

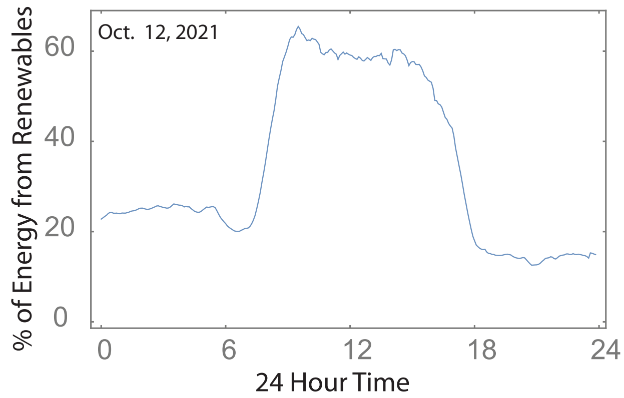

Figure 1: *Percentage of total energy in California for a single day generated by renewables. Between about 9 AM and 3 PM, renewables account for over 60% of the total energy generated.*

energy grid would result in the embodied carbon encompassing 100% of a device's carbon footprint. Accounting for the embodied carbon in terms of required equipment lifecycle to amortize the initial carbon cost is essential for assessing the entire ecological effect of datacenter ICT.

Unfortunately, existing datacenter system design methodology works against these trends, resulting in two primary barriers:

Barrier 1: A limited ability to respond to varying power availability. Datacenters are highly engineered systems that integrate 10s to 100s of thousands of processors, ram and storage, and a large pool of application-specific accelerators such as GPUs, TPUs, and FPGAs. These devices are all interconnected with a high-speed and low-latency datacenter network. The layers of application software, execution runtimes, and virtualization support necessary to implement complex functions relies on a highly-dependent web of interactions between components, typically necessitating an "always on" model of efficiency. Network protocols that deliver data at low latency between components rely on microsecond delivery times, making it hard to scale resources up (or down) in response to transient energy availability. Indeed Google's book on this subject advocates for running systems at high utilization to maximize efficient use of energy [1].

Barrier 2: Short refresh cycles for equipment. Traditional datacenter designs are predicated on energy being expensive and scarce. Powering servers and switches accounts for a significant fraction of a datacenter's overall operational expenses. Indeed, it is common to describe a datacenter using this limit, e.g., “a 50 MW datacenter facility.” This existing design paradigm does not explicitly consider embodied carbon. This exclusion incentivizes operators to replace aging components with newer ones at a rapid pace because newer equipment has historically has been able to do more compute per unit energy. This means that when the impact of embodied carbon is not considered in the overall datacenter design, there is a cross-over point where it is cheaper to purchase new more energy-efficient equipment than to continue to use older, less energy-efficient equipment. The result is that datacenter ICT is replaced relatively frequently without accounting for the cost of the embodied carbon in the equipment.

When the embodied carbon footprint of the ICT equipment is used in the design, the crossover point changes because we are now trading off the *fixed* embodied carbon footprint of the equipment when it is manufactured over its usable operational lifespan. The longer we keep the servers in production, the lower the amortized carbon footprint becomes.

Let η_c be the “operational” carbon per unit time for a given computation. Note that different compute operations and different combinations of energy sources will have different values of η_c.The total ‘operational” carbon expended over a time interval T is then $\eta_c T$. Irrespective of the nature of the computation, the minimum value of η_c is equal to zero and occurs when all of the energy is renewable so that no “operational” carbon is used for computation.

Now consider the effect of the embodied carbon. Let B be the total embodied carbon required to manufacture, deliver and install the equipment. Over the time interval T the total carbon is given by $B + \eta_c T$. The carbon C expended per unit time is then

$$C \;=\; \eta_c + \frac{B}{T}$$

Optimization of this basic equation defines a core attribute of the envisioned research program. As more renewable energy sources are used, η_c decreases and a longer time interval T is required to amortize the effect of the embodied carbon to the level of the “operational” carbon as described by η_c.

We argue that a research program that targets a radical change the existing rate of decarbonization of datacenters —say $2\times$—is a critical component of an overall societal effort to decarbonize. This program will investigate novel architectural responses at every level: from the hardware architecture of individual servers to the architecture of geographically-distributed networked systems and the composition of their services. The success of such a program will reap rewards for decades to come in the form of cloud computing with nearly a zero carbon footprint.

3 Research Vision

Our lab’s research vision aims to enable the dramatic reduction in both the operational and embodied carbon footprint of datacenter systems. As mentioned above, the increased deployment of renewable energy sources will lead to significant variability in the amount of low-carbon power available at any particular location, which we see as both a challenge and opportunity. We will organize our research efforts into three main thrusts.

Thrust 1: Design datacenter systems that can adapt to the variability of renewable energy sources. Here the intermittency of renewable power is a challenge. Just as traditional data processing systems wrestled with whether to move the compute to the data, or the data to the compute, in future datacenter systems the question will be whether to move renewable energy to the datacenter, or whether to move datacenter workloads closer to sources of renewable energy.

There are fundamental and engineering reasons for why it is hard to migrate low-carbon renewable energy over long distances, and so instead we look at the problem of moving compute to low-carbon power. This “movement” could be in time, delaying some or all of a user’s computation until low-carbon renewable energy becomes available, and indeed Google has begun to deploy this approach for their internal workloads [4]. We want to generalize this approach for public cloud workloads.

In addition to time shifting tiers of compute, we could additionally space shift work by migrating jobs from a datacenter fed with high-carbon power to one with low-carbon power. Doing so involves solving a number of research challenges. First, we would need to identify jobs that use a large amount of power. Second, since jobs depend on data and not just executable code, we'd have to identify jobs whose data inputs are small, since wide-area bandwidth is a scarce and expensive resource. Jointly optimizing the selection of jobs based on these criteria would allow a subset of work to migrate in response to differences in grid carbon intensity between datacenters.

User jobs in public clouds rely on a wide variety of infrastructure services, including file and object storage, virtualization platforms such as Docker and Kubernetes, and in-network caches and memory-oriented directories such as Memcached and Redis. Migrating a user's work from one datacenter to another necessitates understanding these dependencies and ensuring either that the relevant infrastructure services are migrated along with the code, or that the services are used infrequently enough that accessing them remotely over the wide-area network is acceptable.

The end goal of Thrust 1 is a datacenter system (both hardware and software) capable of efficiently harnessing the highly variable energy availability of a grid fed primarily by renewable power.

Thrust 2: Reduce the impact of ICT manufacturing on total datacenter carbon footprint. Today, energy purchased from the grid is a significant component of datacenter operational expenses, and thus operators refresh their servers, storage, and networking equipment with the newest equipment to maintain a very high overall *energy efficiency* of equipment. This leads to rapid refresh cycles of ICT equipment (in some cases as short as three years).

While it may be possible to decarbonize the datacenter supply chain to an extent, the technology does not currently exist to fully decarbonize the manufacture and transport of computer equipment. Further, the use of rare earth elements, as well as lithium, graphite, cobalt, and other inputs to the datacenter supply chain result in a high embodied environmental and humanitarian impact.

If you only consider monetary factors, ICT is purchased and deployed for a particular lifespan before it is refreshed with new equipment. That lifespan is calculated as the crossover point where the cost to buy newer, faster equipment is lower than the operational costs of running the older, slower equipment. When we factor in the embodied carbon footprint of the ICT equipment, that changes that crossover point, since now we're trading off the *fixed* embodied carbon footprint of the equipment as it is manufactured over its usable operational lifespan. The longer we keep the servers in production, the lower that amortized carbon footprint becomes.

Figure 1 shows the availability of renewable power in California's CAISO grid on October 12, 2021, and it is representative of the mixture throughout the day of power in this region. If we think of power as fluctuating between a low water mark (when renewables are least available, and much of the grid's power is sourced from fossil fuels) and a high water mark (when low-carbon renewables are producing their maximum delivered output), then we have two regimes. The first regime is closer to the low water mark, where we want very efficient servers to make best use of that power. But as we skew towards the high water mark, it makes sense to bring into production older ICT equipment to take on some aspect of the computing workload. It is true this older equipment won't be as energy efficient as the newer equipment, however it's embodied carbon load is much lower than the newer equipment.

We now have a tradespace between operational and embodied carbon. We envision designing new datacenter systems in which older and newer hardware are *composed* together to build a heterogeneous-aged cluster. Software would be designed to spread compute in a fine-grained way over varying generations of components, taking advantage of the very low embedded carbon of older equipment for tasks that can be delayed to periods of high renewable energy availability, while relying on new, energy-efficient hardware for time-sensitive tasks that take place around the clock.

Building the development tools, runtimes, and virtualization layers necessary to compose older and newer equipment transparently is a major research goal of Thrust 3. In addition, older ICT equipment is likely to fail at a higher rate than newer equipment, necessitating new approaches to fault tolerance and reliability.

Thrust 3: Education. We will develop new course materials aimed at educating students and the public on the environmental effects of mining, manufacturing, operating, and decommissioning the compute equipment underpinning cloud computing and our modern compute ecosystem. We have already piloted new courses at the University of Southern California, the University of California San Diego, and the University of San Diego, and will continue to develop these courses and release the accompanying materials to the public. Further, we observe that optimizing software for resource efficiency needs to be an integral aspect of a modern Computer Science curriculum, and that arming students with the knowledge and skills necessary to make design tradeoffs that focus on a resource-efficient lifecycle is a key aspect of C^3-Lab.

Bibliography

[1] Luiz André Barroso and Urs Hölzle. "The datacenter as a computer: An introduction to the design of warehouse-scale machines". In: *Synthesis lectures on computer architecture* 4.1 (2009), pp. 1–108.

[2] Hugues Ferreboeuf et al. "Lean ICT: Towards digital sobriety". In: *Report for the Shift Project, Paris, March* (2019).

[3] Udit Gupta et al. *Chasing Carbon: The Elusive Environmental Footprint of Computing*. 2020. arXiv: `2011.02839 [cs.AR]`.

[4] Ana Radovanovic. *Our data centers now work harder when the sun shines and wind blows*. `https://blog.google/inside-google/infrastructure/data-centers-work-harder-sun-shines-wind-blows/`. [Online; accessed 7-Oct-2021]. 202.