A Vision for Computational Decarbonization of Societal Infrastructure

David Irwin, Prashant Shenoy, Mohammad Hajiesmaili, Walid A. Hanafy, Jimi Oke, and Ramesh Sitaraman, *University of Massachusetts Amherst, Amherst, MA, 01003, USA*

Yuvraj Agarwal, Geoffrey J. Gordon, and Zico Kolter, *Carnegie Mellon University, Pittsburg, PA, 15213, USA*

Deepak Rajagopal and Mani Srivastava, *University of California Los Angeles, Los Angeles, CA, 90095, USA*

Vivienne Sze and Priya Donti, *Massachusetts Institute of Technology, Cambridge, MA, 02139, USA*

Andrew Chien, John Birge, and Ali Hortacsu, *University of Chicago, Chicago, IL, 60637, USA*

Line Roald, *University of Wisconsin, Madison, WI, 53706, USA*

Modern society is at a critical inflection point with rapidly accelerating demand for energy due to growth in domestic manufacturing, datacenters, artificial intelligence (AI), electric vehicles, and electric heat pumps. Sustaining this growth while also reducing society's carbon emissions will necessitate a shift beyond our long-standing focus on improving energy efficiency to optimizing carbon efficiency. This paper lays out a vision for a new field of computational decarbonization, which focuses on optimizing and reducing the lifecycle carbon emissions of complex computing and societal infrastructure systems. We identify an important class of decarbonization problems that arise from interdependencies across multiple infrastructure domains, including computing, transportation, the built environment, and the electric power grid. As we discuss, solving these problems will require developing novel computational techniques, algorithms, systems, and AI methods that sense, optimize, and reduce the operational, embodied, and lifecycle greenhouse gas emissions of societal infrastructure over long temporal and spatial scales.

There is now a broad consensus that society must undertake a rapid *energy transition* to reduce and, ultimately, eliminate its carbon emissions. However, despite promising advances in low-carbon energy technologies, such as renewable energy and energy storage, scientific experts and industry leaders generally acknowledge that many of the technologies necessary to achieve ambitious decarbonization goals do not yet exist.[15] To reduce carbon emissions, researchers have long focused on reducing society's energy consumption by optimizing energy efficiency in various domains, including computing, the built environment, and transportation. Unfortunately, while improving energy efficiency increases productivity and economic output—by enabling more to be done with less energy at lower cost—it alone is not sufficient to decarbonize society for multiple reasons. In particular, 1) energy-efficiency improvements are subject to rebound effects that generally serve to increase energy usage (referred to as Jevon's Paradox[3]) and 2) there are fundamental bounds to any task's energy efficiency that, in many cases, we are approaching. For example, if historical trends continue, we will reach Landauer's limit—the physical bound of computing's energy efficiency—by the 2040s.[12]

Accelerating society's decarbonization will require increasing our emphasis on directly optimizing for

Digital Object Identifier 10.1109/MIC.2025.3575016
Date of publication 30 May 2025; date of current version 30 June 2025.

carbon efficiency. While energy efficiency captures the work done per unit of energy consumed, carbon efficiency captures the work done per amount of carbon [and other greenhouse gases (GHGs)] emitted. Thus, a task's carbon efficiency is a function of not only its energy efficiency but also its *energy's carbon intensity*, i.e., the emissions from generating each unit of energy. As a result, energy-efficient systems can be carbon *inefficient*—if their energy derives from carbon-intensive, or "brown," sources, i.e., by burning fossil fuels—and carbon-efficient systems can be energy *inefficient*—if their energy derives from zero-carbon renewables. Notably, optimizing carbon efficiency does not suffer from the aforementioned drawbacks: 1) it is not subject to Jevon's Paradox since carbon is not an economic resource but an energy byproduct, and 2) it is possible to be infinitely carbon efficient by using zero-carbon energy. However, since zero- or low-carbon energy largely comes from intermittent renewable sources, such as wind, solar, and hydro, its availability, and thus energy's carbon intensity, varies substantially, e.g., more than an order of magnitude, over time and space (see Figure 1). As a result, optimizing carbon efficiency requires societal systems to be fundamentally rethought to exploit temporal and spatial variations by doing more work when and where low-carbon energy is available.

Despite its increasing importance, optimizing carbon efficiency has seen much less research attention, and many fewer advances, than optimizing energy efficiency for both social and technical reasons. Socially, since energy incurs a monetary cost, there has always been a strong economic incentive to reduce cost by increasing energy efficiency. By contrast, in the absence of a carbon tax, there is no direct economic incentive to optimize carbon efficiency. However, government subsidies,[2] proposed mandatory reporting,[7] and continuing declines in renewable energy costs, primarily for solar and wind, are starting to introduce some

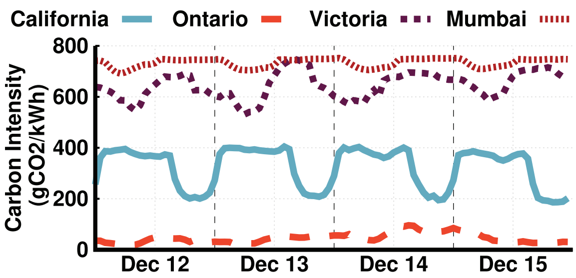

FIGURE 1. Grid carbon emissions vary 6–43× temporally and spatially.

indirect incentives. Technically, there are multiple barriers to optimizing carbon efficiency including the following:

1) a lack of *visibility* into the carbon emissions of the energy we use, i.e., "operational" emissions, and the products we consume, i.e., "embodied" emissions
2) a lack of *flexibility* in responding to, and optimizing for, changes in operational and embodied carbon emissions, e.g., by adapting how we use energy and the products we buy
3) a lack of *programmability* in exposing software-defined interfaces to automatically monitor and control systems to optimize carbon efficiency by leveraging energy flexibility.

The absence of social incentives combined with this lack of *visibility*, *flexibility*, and *programmability* has largely prevented sophisticated carbon-efficiency optimizations.

The field of *Computational Decarbonization* (or CoDec) focuses on overcoming these barriers by optimizing carbon efficiency to reduce the lifecycle carbon emissions—the sum of operational and embodied carbon amortized—of computing and societal infrastructure using computational and data-driven techniques. Since CoDec research targets the foundations of society's infrastructure—for computing, electricity, the built environment, and transportation—which all represent different forms of inter-connected *cyberphysical systems*, we view it as a special type of *sense–optimize–reduce* problem that operates both within and across these domains, as well as over multiple temporal and geographical scales. We envision new research addressing the aforementioned technical challenges by developing novel 1) *sensing* approaches to provide *visibility* into systems' operational and embodied carbon over their lifetime, 2) *optimization* methods grounded in theory and artificial intelligence (AI) to exploit new dimensions of energy *flexibility*, which are emerging in modern infrastructure, for optimizing carbon efficiency, and 3) software-defined interfaces and systems for *programmatically* deploying these optimizations to *reduce* carbon emissions.

OVERVIEW

There are multiple trends that motivate a new focus on computational decarbonization across society.

- *Carbon–energy–performance (CEP) gap*: Optimizing for carbon efficiency introduces fundamental tradeoffs with multiple metrics, including performance, energy efficiency, cost, fairness, and many others. In particular, we hypothesize that it may be *impossible* to simultaneously optimize carbon efficiency, energy efficiency, and performance in societal infrastructure domains. Thus, better understanding and quantifying these complex tradeoffs across domains is critically important in developing carbon-efficiency optimizations, which must navigate them in reducing carbon emissions.

- *Inter-dependency gap*: While independent efforts have begun to decarbonize different sectors, such as cloud computing, the built environment, and transportation, they ignore complex interdependencies that exist between sectors, as well as across time and space. These sectors are connected both physically—via the electric grid—and economically—through supply chains. As one simple example of this interdependency, while shifting to remote work reduces transportation emissions, it likely increases residential building emissions. Siloed approaches that ignore such couplings are likely to achieve limited success. In contrast, new research needs to rigorously address, and exploit, such interdependencies through cross-system and cross-domain approaches, mediated by computation.
- *Scale gap*: The scale at which society must address decarbonization problems is fundamentally different than conventional energy-efficiency problems. Specifically, traditional energy optimizations have broadly focused on relatively small spatiotemporal scales. As shown in Figure 2, efforts within computing, for instance, have focused on subsecond optimizations within servers, e.g., chip- and OS-level power management,[16,19] to hours-level methods within datacenters, e.g., optimizations based on diurnal workloads.[10] In contrast, many decarbonization problems manifest at much longer temporal scales, ranging from minutes to months or years, and much larger spatial scales, ranging from individual systems or communities to entire regions. Decarbonization problems at these intermediate "mesoscales"[a] abound, and include optimizing the annual emissions of a hyperscale cloud datacenter, increasing equipment lifetimes to reduce embodied emissions, and decarbonizing a regional electric vehicle charging network. Addressing these intermediate challenges is critical to bridging the "scale gap" that exists between low-level energy optimizations and global-scale decarbonization goals.[11]

[a]Inspired by mesoscale meteorology, representing weather phenomena over a few kilometers to 2000 km. Meso means intermediate scale.

Importantly, computational decarbonization research needs to address problems at mesoscales through a novel class of computational techniques, algorithms, systems, and AI methods designed to sense, optimize, and reduce the lifecycle (i.e., operational and embodied) GHG emissions of societal infrastructure over the intermediate time scales of minutes-to-years and spatial scales of communities-to-countries. This research will also need to recognize the unique role computing will play in decarbonizing society both as a "means" to automate and coordinate carbon-efficiency optimizations across time, space, and sectors, and as a "medium" that consumes increasingly significant amounts of energy but also has substantial temporal, spatial, and performance flexibility.

The potential benefits of CoDec research can be quite significant. EPA studies show that society's infrastructure generates 89% of its GHG emissions[1]—with 25% from electricity grids, 27% from transportation, 13% from buildings and 24% from industry. By accelerating the decarbonization of these sectors through its unique multidomain approach, CoDec can help bend the emissions curve downwards. For a single sector,

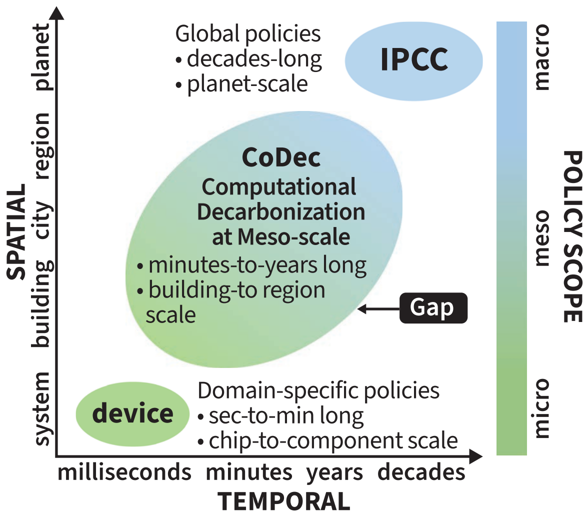

FIGURE 2. Computational decarbonization fills the gap between global-scale, decades-long policies, and domain-specific, short-term decisions.

such as computing, this can eliminate 730 megatons of annual GHG emissions in perpetuity, with each 1% additional emission reductions across domains yielding 30 billion tons of savings.[1]

CODEC VISION

We envision a future decarbonized world that will generate energy from a mix of mostly low-carbon, but volatile, energy sources. This transition is poised to alter the foundation of modern society by transforming energy from a resource that is largely centralized, costly, and dirty, but highly stable, to one that is largely distributed, cheap, and clean, but highly intermittent and variable. Thus, in this future world, clean energy will be plentiful at times, but precious at others, making coordination and orchestration of energy's supply and demand critical for maintaining quality of life, sustaining the global economy, and ensuring energy availability during times of scarcity. Importantly, all of society will need to make a fundamental shift to adapt to this future by exploiting multiple dimensions of energy-flexibility—the degree to which a system can modulate its energy usage across time, space, and uses—to consume low-carbon energy when and where it is available.

The future world we envision will also need to employ human-centered approaches where individuals, and the organizations to which they belong, have the ability to make decisions ("green choices") about when and how they consume carbon as part of their normal activities, such that societal systems respect and implement these decisions in fair, accountable, transparent, and equitable manner.

Some early efforts have already begun to address various facets of decarbonization, which we briefly discuss through the lens of visibility, flexibility, and programmability. Electrical energy's carbon intensity—the emissions per unit of electrical energy generated (in $g \cdot CO_2/kWh$)—depends on the mix of generation sources that supply energy to each regional grid. Since many grid operators now publish data on their generation sources in real time, it is now possible to estimate the grid energy's carbon intensity. For example, third-party services, such as WattTime and ElectricityMap, provide real-time estimates of grid energy's carbon-intensity via cloud application programming interfaces. As depicted in Figure 1, grid energy's carbon intensity can vary by $6\times$ over a day (in California) and by $43\times$ across regions (between Ontario and Mumbai). Cloud providers, including Google Cloud Platform (GCP) and Azure have, in turn, begun to use such data to estimate and expose their cloud resources' carbon intensity to end users.[9] However, the current data's temporal and spatial resolution remains coarse and backward-looking (e.g., average carbon intensity for large regions over the prior month), and is thus insufficient for larger forward-looking mesoscale approaches. In addition to directly optimizing for changes in grid energy's carbon efficiency, many companies also purchase carbon offsets to reduce their net carbon footprint on an annualized basis.[14] Unfortunately, carbon offsets are only effective as a transitional mechanism since we must ultimately reduce absolute global carbon emissions to near zero, where there is little carbon left to offset. To reach zero carbon, we must eventually move beyond carbon offsets and focus on *changing operations* to always run on low-carbon energy, e.g., from solar, wind, hydro, nuclear, geothermal, and so on. Many types of carbon offsets may actually delay these operational changes by providing a means for reducing carbon emissions for less than it would cost to make such changes.[5]

Reducing lifecycle carbon emissions by improving carbon efficiency requires systems to fundamentally alter their short- and long-term operation by exploiting flexibility. For example, since cloud workloads have considerable temporal and spatial flexibility, recent approaches have quantified emissions reductions from shifting work across time and cloud regions.[17,18] In addition, the grid has long operated demand response programs[13] that turn off flexible loads, often manually, to reduce infrequent seasonal peak usage, which disproportionately impacts cost and carbon emissions. Apple's iOS also recently incorporated a green charging option for all iPhones that shifts charging times based on changes in the grid's carbon intensity.[4] However, these efforts, while promising, are isolated, uncoordinated, and fail to consider and exploit interdependencies and couplings within and across domains. Further, these efforts largely focus on optimizing operational, and not embodied, emissions, which are also important. In contrast, CoDec research should focus on optimizing lifecycle carbon, i.e., a system's emissions amortized over its entire lifetime, which includes both embodied carbon from systems' manufacturing, transportation, installation, and decommissioning, and operational carbon resulting from their day-to-day operation. While prior efforts have addressed these two components through largely independent approaches,[17,20] optimizing both introduces dependencies (or couplings) that necessitate a joint approach. For example, while increasing a server or car's lifetime amortizes its embodied carbon over a longer duration, it also increases its operational carbon—as aging equipment may be less efficient than newer models.

Further, systems in different domains have varied lifetimes and vastly different ratios of operational and embodied carbon. For example, smartphones tend to have short 3–4 year lifetimes with *embodied* carbon representing 80% of their lifecycle costs, while buildings have lifetimes of 50 years or more with *operational* carbon constituting 80% of lifecycle carbon (see Figure 3). While some industry efforts for optimizing embodied carbon seek to recycle products to reduce e-waste or resell them to extend lifetimes,[6] little work has considered jointly optimizing operational and embodied carbon holistically, while also being cognizant of domain-specific differences. Finally, with the emergence of industrial Internet of Things technologies, systems, such as buildings, vehicles, and factories increasingly expose programmatic control, providing the basic building blocks for lifecycle carbon management. However, these programmatic interfaces only allow control of individual systems, but do not enable managing and orchestrating larger collections of systems for decarbonization at scale.

CEP Impossibility Conjecture

Optimizing lifecycle carbon for societal infrastructure cannot be done in a vacuum without considering these systems' other objectives—such as their energy efficiency, cost efficiency, and performance (for users). Complex systems that must satisfy multiple independent goals are generally forced to make nontrivial tradeoffs, as the goals often conflict. As one example, in large distributed systems, the consistency-availability-partition Tolerance (CAP) theorem states that distributed systems can only achieve two of the three desirable goals of consistency, availability, and network partition tolerance at any time.[8] The CAP theorem, originally posed as a conjecture but later proven,[8] is a seminal result that now underpins the design of nearly every large distributed system.

Societal infrastructure is subject to similar conflicts between the goals of carbon efficiency, energy efficiency, and performance. Specifically, in scenarios where energy is not entirely carbon free and energy's carbon intensity varies over time and space, it is not possible for systems to simultaneously maximize carbon efficiency, energy efficiency, and performance. Analogous to the CAP theorem, we postulate this fundamental tradeoff as an impossibility result called the *CEP conjecture*. As shown in Figure 4, the CEP conjecture is a fundamental principle that states that systems across a broad range of domains can achieve at most two out of the three properties of carbon efficiency, energy efficiency, and performance. Note that energy efficiency implicitly correlates with cost efficiency since using energy incurs a monetary cost.

Although a conjecture, preliminary work reveals examples in domains from computing to transportation where this property arises. For example, an intelligent electric vehicle (EV) charging network might slow charging when grid carbon intensity is high, reducing emissions and increasing energy efficiency, but doing so increases charging time, reducing user-perceived performance. Likewise, prior work on zero-carbon clouds trades operational carbon for increased embodied emissions.[20] The CEP conjecture has profound implications on decarbonization optimizations since it suggests tradeoffs across conflicting goals are required: understanding how to navigate these tradeoffs can serve as a unifying theme across a broad range of CoDec research.

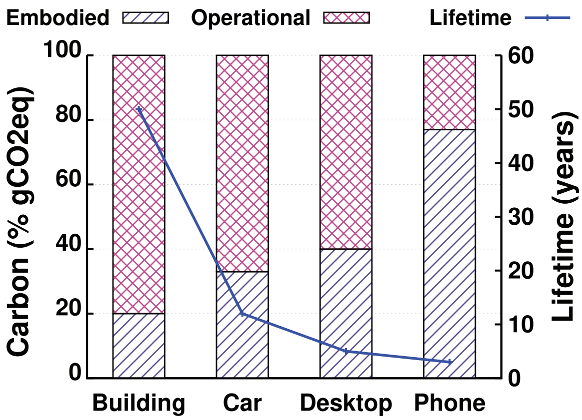

FIGURE 3. Lifetimes and embodied carbon across domains.

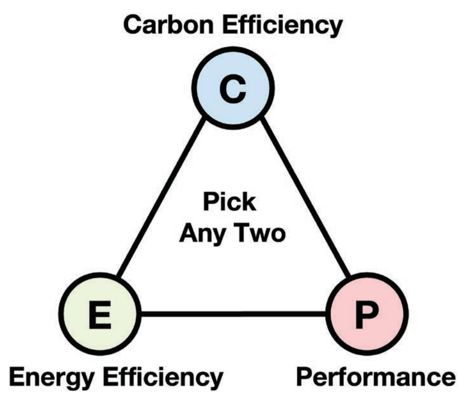

FIGURE 4. The CEP impossibility conjecture: "Pick any two of carbon efficiency, energy efficiency, and performance."

Multiscale Temporal and Spatial Dependence

Computational decarbonization, when viewed as a sense–optimize–reduce control problem, exhibits multiscale temporal and spatial dependence. Specifically, we can decompose the decarbonization of any system into a set of subproblems, each with its own sense–optimize–reduce loop (see Figure 5). Within a system, these decarbonization loops arise hierarchically at multiple temporal and spatial scales. For example, a hyperscale cloud datacenter may operate a carbon-aware cloud scheduler that optimizes emissions at the scale of minutes to hours, a resource manager that optimizes annual emissions, and a lifetime manager that manages emissions over several years.

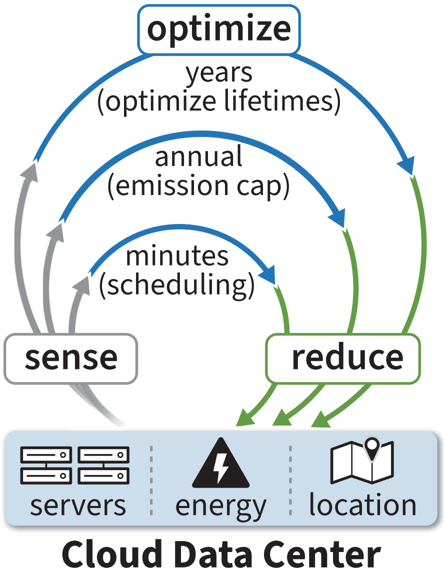

FIGURE 5. Sense-optimize-reduce loops at multitime-scales.

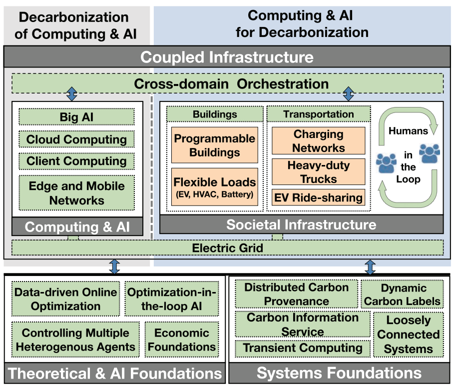

FIGURE 6. Research directions for our CoDec project.

Similarly, a city-wide EV charging network or a distributed Content Distribution Network (CDN) may coordinate across different spatial locations to optimize their network emissions. These concurrent spatiotemporal control loops are linked through their contribution to lifecycle carbon and hence must also coordinate with one another—a form of joint optimization—to achieve a common decarbonization goal. Such multiscale dependent spatiotemporal sense–optimize–reduce decarbonization loops are a distinguishing characteristic of CoDec problems.

Further, since decarbonization decisions in one domain can affect emissions in another, these sense–optimize–reduce loops also have interdependencies across domains. For instance, decarbonizing transportation by shifting to clean electric cars and trucks increases grid demand and potentially emissions from "dirty" peaker plants. Moreover, Jevons' Paradox[3] can introduce rebound effects where increasing the availability of cheap low-carbon energy may encourage society to consume more of it, causing demand to outstrip supply. To address these problems, future research in CoDec will need to address and exploit cross-domain interdependencies through new decarbonization approaches that are coupled with and mediated by computation.

Figure 6 provides a layered view of the different areas of research in our ongoing research project in computational decarbonization. The lowest layer focuses on both 1) foundational algorithms and policies for optimizing CEP tradeoffs grounded in theory and AI and 2) foundational systems mechanisms capable of securely implementing these policies. Foundational theory and AI approaches for optimizing carbon efficiency at mesoscales include learning-driven online optimization, optimization-in-the-loop learning, and multiagent learning that also considers economic incentives. Foundational systems mechanisms include general software platforms and carbon services for improving distributed infrastructure systems' visibility, flexibility, and programmability to monitor and respond to changes in both their operational and embodied carbon emissions. The next layer then focuses on enabling the previous automated policies for multiple domains of societal infrastructure to optimize lifecycle carbon emissions. This includes computing infrastructure, e.g., AI applications, large-scale cloud platforms, edge networks, and client devices, as well as other large-scale infrastructure, including the built environment, electric transportation networks, and human-in-the-loop systems. Finally, the highest layer focuses on cross-domain optimizations that coordinate between different sectors, including computing, buildings, transportation, the grid, and so on.

CONCLUSION

This article lays out a vision for a new research area, called computational decarbonization, that focuses on applying computational techniques to jointly decarbonize, not just computing, but multiple sectors of society, e.g., computing, buildings, transportation, the electric grid, supply chains, and so on. Computational decarbonization differs from prior efforts in multiple ways, particularly its focus on jointly optimizing lifecycle carbon emissions across multiple sectors and time scales.

ACKNOWLEDGMENTS

This work was supported by a National Science Foundation (NSF) Expeditions award CNS-2325956.

REFERENCES

1. "Sources of greenhouse gas emissions," US Environmental Protection Agency, Washington, DC, USA, 2021. [Online]. Available: https://www.epa.gov/ghgemissions/sources-greenhouse-gas-emissions
2. "How the inflation reduction act will drive global climate action," Center for American Progress, Washington, DC, USA, Aug. 2022. [Online]. Available: https://www.americanprogress.org/article/how-the-inflation-reduction-act-will-drive-global-climate-action
3. B. Alcott, "Jevons' paradox," *Ecol. Econ.*, vol. 54, no. 1, pp. 9–21, 2005.
4. "Use clean energy charging on your iPhone." Apple Support. Accessed: Mar. 2025. [Online]. Available: https://support.apple.com/en-us/HT213323
5. N. Bashir et al., "Enabling sustainable clouds: The case for virtualizing the energy system," in *Proc. ACM Symp. Cloud Comput. (SoCC)*, New York, NY, USA: ACM, Oct. 2021, pp. 350–358.
6. V. Forti, C. P. Balde, R. Kuehr, and G. Bel, "The global E-waste monitor 2020: Quantities, flows and the circular economy potential," United Nations, New York, NY, USA, 2020. Accessed: Mar. 2025. [Online]. Available: https://ewastemonitor.info/wp-content/uploads/2020/11/GEM_2020_def_july1_low.pdf
7. G. Gensler, "Statement on proposed mandatory climate risk disclosures," US Government Securities Exchange Commissions, Washington, DC, USA, Mar. 2022. [Online]. Available: https://www.sec.gov/news/statement/gensler-climate-disclosure-20220321
8. S. Gilbert and N. Lynch, "Brewer's conjecture and the feasibility of consistent, available, partition-tolerant web services," *ACM SIGACT News*, vol. 33, no. 2, pp. 51–59, 2002, doi: 10.1145/564585.564601.
9. "Carbon footprint." Google Cloud. Accessed: Mar. 2025. [Online]. Available: https://cloud.google.com/carbon-footprint
10. S. Govindan, A. Sivasubramaniam, and B. Urgaonkar, "Benefits and limitations of tapping into stored energy for datacenters," in *Proc. 38th Annu. Int. Symp. Comput. Architect. (ISCA)*, New York, NY, USA: ACM, Jun. 2011, pp. 341–352.
11. IPCC, "Climate change 2022: Impacts, adaptation and vulnerability – Summary for Policymakers," Cambridge Univ. Press, Cambridge, UK, 2022. [Online]. Available: https://www.ipcc.ch/report/ar6/wg2/downloads/report/IPCC_AR6_WGII_SummaryForPolicymakers.pdf
12. R. Landauer, "Irreversibility and heat generation in the computing process," *IBM J. Res. Dev.*, vol. 5, no. 3, pp. 183–191, Jul. 1961, doi: 10.1147/rd.53.0183.
13. J. Medina, N. Muller, and I. Roytelman, "Demand response and distribution grid operations: Opportunities and challenges," *IEEE Trans. Smart Grid*, vol. 1, no. 2, pp. 193–198, Sep. 2010, doi: 10.1109/TSG.2010.2050156.
14. "Climate." Meta. Accessed: Mar. 2025. [Online]. Available: https://sustainability.fb.com/climate/
15. S. Pichai, "Our third decade of climate action: Realizing a carbon-free future," *Google Blog*, Sep. 14, 2020. [Online]. Available: https://blog.google/outreach-initiatives/sustainability/our-third-decade-climate-action-realizing-carbon-free-future
16. K. Shen, A. Shriraman, S. Dwarkadas, X. Zhang, and Z. Chen, "Power containers: An OS facility for fine-grained power and energy management on multicore servers," *ACM SIGPLAN Notices*, vol. 48, no. 4, pp. 65–76.
17. A. Souza et al., "Ecovisor: A virtual energy system for carbon-efficient applications," in *Proc. 28th ACM Int. Conf. Archit. Support Program. Lang. Operating Syst. (ASPLOS)*, New York, NY, USA: ACM, Mar. 2023, pp. 252–265.
18. E. Strubell, A. Ganesh, and A. McCallum, "Energy and policy considerations for deep learning in NLP," in *Proc. 57th Annu. Meeting Assoc. Comput. Linguistics*, Association for Computational Linguistics, Jul. 2019, pp. 3645–3650.
19. M. Weiser, B. Welch, A. Demers, and S. Shenker, "Scheduling for reduced CPU energy," in *Proc. 1st USENIX Conf. Operating Syst. Des. Implementation (OSDI)*, USENIX Association, Berkeley, CA, USA, Nov. 1994, pp. 13–23.
20. F. Yang and A. A. Chien, "Large-scale and extreme-scale computing with stranded green power: Opportunities and costs," *IEEE Trans. Parallel Distrib. Syst.*, vol. 29, no. 5, pp. 1103–1116, May 2018, doi: 10.1109/TPDS.2017.2782677.

DAVID IRWIN is a professor of electrical and computer engineering at the University of Massachusetts Amherst, Amherst, MA, 01267, USA. His research interests include computer systems, energy-efficiency, and sustainability. Irwin has a Ph.D. in computer science from Duke University. Contact him at deirwin@umass.edu.

PRASHANT SHENOY is a distinguished professor of computer science at the University of Massachusetts Amherst, Amherst, MA, 01267, USA. His research interests include distributed systems, networking, and cloud computing. Shenoy earned his Ph.D. from Computer Science, University of Texas at Austin. Contact him at shenoy@cs.umass.edu.

MOHAMMAD HAJIESMAILI is an associate professor of computer science at the University of Massachusetts Amherst, Amherst, MA, 01267, USA. His research interests include optimization, machine learning, and algorithms. Hajiesmaili earned his Ph.D. degree in computer engineering from the University of Tehran. Contact him at mhajiesmaili@umass.edu.

WALID A. HANAFY is a postdoctoral research associate in computer science at the University of Massachusetts Amherst, Amherst, MA, 01267, USA. His research interests include distributed systems, cloud computing, and carbon-efficiency. Hanafy earned his Ph.D. degree in computer science from the University of Massachusetts Amherst. Contact him at whanafy@cs.umass.edu.

JIMI OKE is an assistant professor of civil and environmental engineering at the University of Massachusetts Amherst, Amherst, MA, 01267, USA. His research interests include transportation engineering, network science, and machine learning. Oke earned his Ph.D. degree in civil engineering from Johns Hopkins University. Contact him at jboke@umass.edu.

RAMESH SITARAMAN is a distinguished professor of computer science at the University of Massachusetts Amherst, Amherst, MA, 01267, USA. His research interests include content delivery networks, edge computing, and networking. Sitaraman earned his Ph.D. degree in computer science from Princeton University. Contact him at ramesh@cs.umass.edu.

YUVRAJ AGARWAL is an associate professor in the school of computer science at Carnegie Mellon University, Pittsburgh, PA, 15213, USA. His research interests include computer systems, networking, and embedded systems. Agarwal earned his Ph.D. degree in computer science from the University of California San Diego. Contact him at yuvraj@cs.cmu.edu.

GEOFFREY J. GORDON is a professor in the school of computer science at Carnegie Mellon University, Pittsburgh, PA, 15213, USA. His research interests include robust, safe, and secure machine learning. Gordon earned his Ph.D. degree in computer science from Carnegie Mellon University. Contact him at ggordon@cs.cmu.edu.

ZICO KOLTER is a professor in the school of computer science at Carnegie Mellon University, Pittsburgh, PA, 15213, USA. His research interests include multiagent planning, reinforcement learning, and decision-theoretic planning. Kolter earned his Ph.D. degree in computer science from Stanford University. Contact him at zkolter@cs.cmu.edu.

DEEPAK RAJAGOPAL is an associate professor in the Institute of the Environment and Sustainability at the University of California Los Angeles, Los Angeles, CA, 90095, USA. His research interests include life cycle assessment, industrial ecology, and energy/agricultural economics and policy; Rajagopal earned his Ph.D. degree in energy and resources from the University of California Berkeley. Contact him at rdeepak@ioes.ucla.edu.

MANI SRIVASTAVA is a professor of electrical and computer engineering at the University of California Los Angeles, Los Angeles, CA, 90095, USA. His interests include embedded systems, wireless networks, and cyberphysical systems. Srivastava earned his Ph.D. degree in electrical engineering and computer science from the University of California Berkeley. Contact him at mbs@ucla.edu.

VIVIENNE SZE is a professor of electrical engineering and computer science at the Massachusetts Institute of Technology, Cambridge, MA, 02139, USA. Her research interests include VLSI, low-power design, and machine learning. Sze earned her Ph.D. degree in electrical engineering from the Massachusetts Institute of Technology. Contact her at sze@mit.edu.

PRIYA DONTI is an assistant professor and the Silverman (1968) Family Career Development Professor of Electrical Engineering and Computer science at the Massachusetts Institute of Technology, Cambridge, MA, 02139, USA. Her research interests

include machine learning for forecasting, optimization, and control in high-renewable grids. Donti earned her Ph.D. degree in computer science and public policy from Carnegie Mellon University. Contact her at donti@mit.edu.

ANDREW CHIEN is the William Eckhardt Distinguished Service Professor of Computer Science at the University of Chicago, Chicago, IL, 60637, USA. His research interests include cloud software, sustainable computing, and computer architecture. Chien earned his Sc.D. degree in computer science from the Massachusetts Institute of Technology. Contact him at aachien@uchicago.edu.

JOHN BIRGE is the Hobart W. Williams Distinguished Service Professor of Operations Management at the University of Chicago, Chicago, IL, 60637, USA. His research interests include operations management and mathematical modeling of systems under uncertainty. Birge earned his Ph.D. degree in operations research from Stanford University. Contact him at john.birge@chicagobooth.edu.

ALI HORTACSU is the William M. Ogden Distinguished Service Professor of Economics at the University of Chicago, Chicago, IL, 60637, USA. His research interests include economics and market efficiency. Hortacsu earned his Ph.D. degree in economics from Stanford University. Contact him at hortacsu@uchicago.edu.

LINE ROALD is an associate professor of electrical and computer engineering at the University of Wisconsin–Madison, Madison, WI, 53706, USA. Her research interests include energy systems, renewable energy, and optimization. Roald earned her Ph.D. degree in electrical engineering from ETH Zurich. Contact her at roald@wisc.edu.
