Algorithmic Foundations of Grid-Integrated AI Infrastructure

December 2025 @NSF Compute-Energy Nexus
Paper presented at ACM SIGCOMM 2025

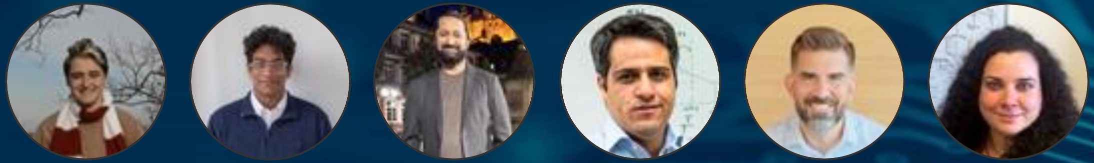

Adam Lechowicz, Rohan Shenoy, Noman Bashir, Mohammad Hajiesmaili, Adam Wierman, Christina Delimitrou

University of Massachusetts Amherst
Massachusetts Institute of Technology
Caltech
Berkeley University of California

US power demand is skyrocketing

- Unprecedented (absolute) demand growth
  - 150–250 TWh per year
- Data centers, running AI workloads, are expected to add ~450 TWh/yr by 2040, 10–12% of total growth.

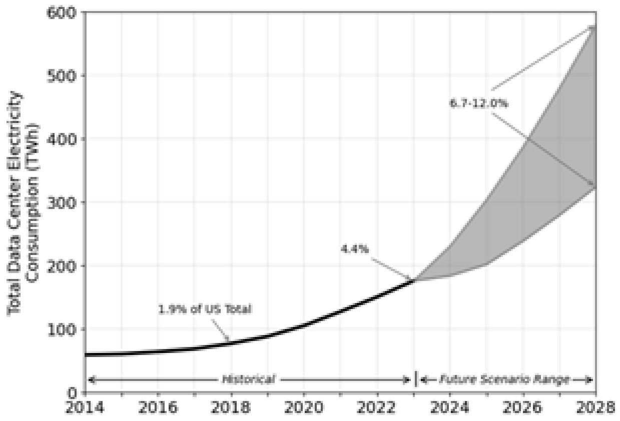

US power demand is expected to increase up to 3.5 percent annually through 2040.

US power demand, by sector, current trajectory scenario,[1] terawatt-hours

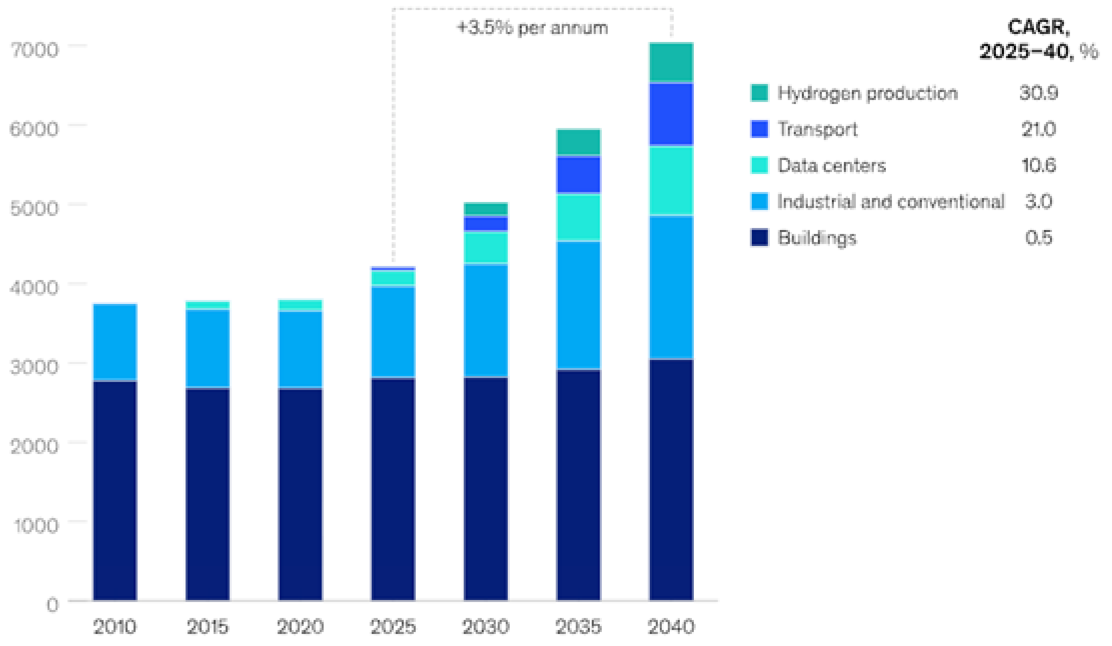

[1]Volumetric power demand load is total downstream sales, comparable to data that operators are required to provide on the US Energy Information Administration's Form EIA-861. It represents the gross of distributed generation, behind-the-meter systems, and transmission and distribution losses.
Source: McKinsey Power Model

Shehabi, A., Smith, S.J., Hubbard, A., Newkirk, A., Lei, N., Siddik, M.A.B., Holecek, B., Koomey, J., Masanet, E., Sartor, D. 2024. ***2024 United States Data Center Energy Usage Report****. Lawrence Berkeley National Laboratory, Berkeley, California. LBNL-2001637*

AI is forcing a redesign of the compute backbone

and calls for the integration of compute and energy infrastructure

- AI workloads are pushing compute, cooling, and power systems beyond traditional limits
- Power access is becoming a primary bottleneck for scaling AI compute
- Integration with energy infrastructure (renewables, microgrids, storage) is essential

Why the AI era is forcing a redesign of the entire compute backbone

Amin Vahdat, Google

August 3, 2025

- To meet AI's massive scaling, we must scale energy supply, delivery, and adaptiveness
- e.g., turning down compute when power is constrained and using real-time telemetry
- otherwise, efficiency, cost, and feasibility will suffer

concerns about the environmental impacts of this massive growth …

https://venturebeat.com/ai/why-the-ai-era-is-forcing-a-redesign-of-the-entire-compute-backbone

CoDec: a new field of **computational decarbonization**

to unify, transform, and accelerate societal decarbonization across domains

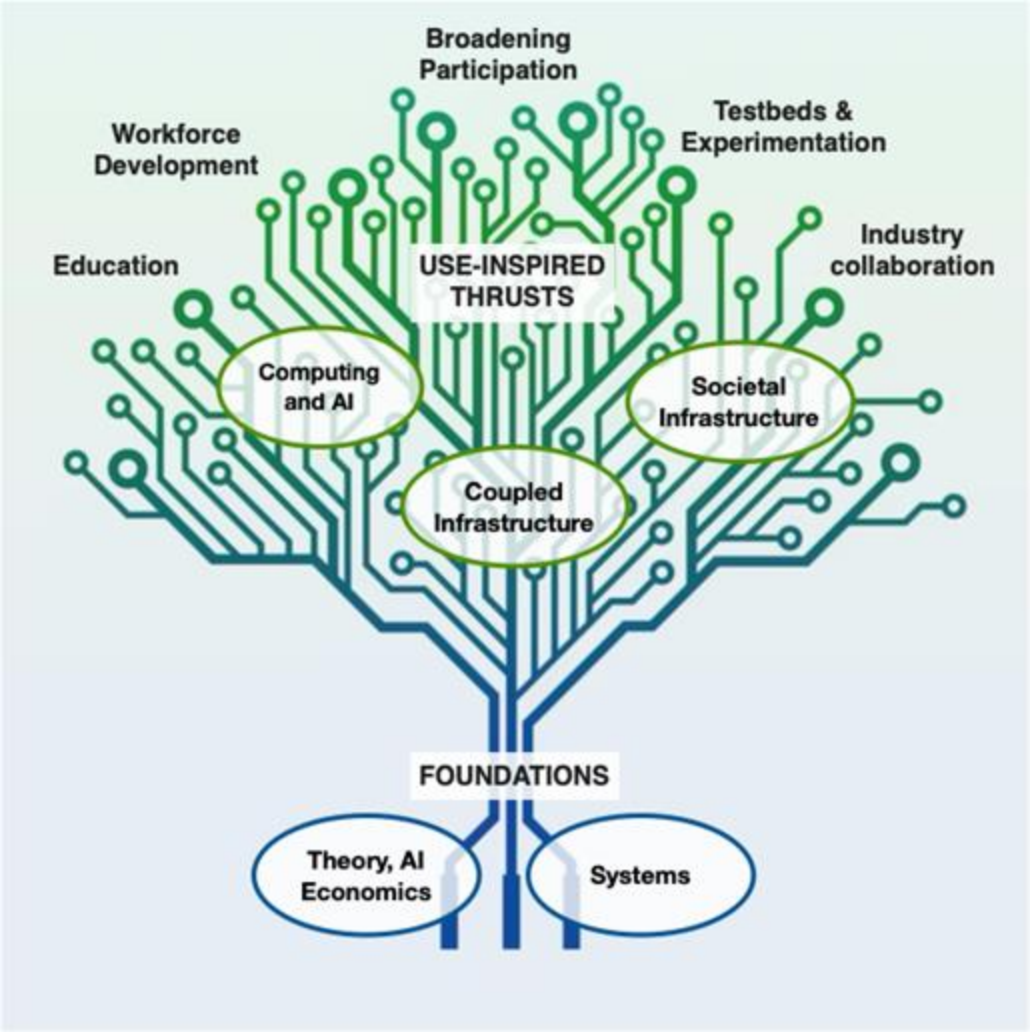

Theory and AI Foundations

- Learning & optimization, Incentive design

Systems Foundations

- Carbon services, coordination abstractions

Decarbonize Computing and AI

- AI, data center servers, clients, edge networks

Decarbonize Societal Infrastructure

- Built environment, transportation, human factors

Coupled Infrastructures

- Grid, computing couplings, buildings + transport

Computational decarbonization of societal infrastructures

EXPEDITIONS: CoDec

SOLAR Lab

Sustainability, Optimization, Learning, and Algorithms Research

The CoDec Expedition will **develop the new field of Computational Decarbonization,** focusing on optimizing and reducing the lifecycle carbon emissions of complex computing and societal infrastructure systems.

Sensing to provide visibility into systems' operational and embodied carbon over their lifetime

Exploit new dimensions of **energy flexibility** in modern infrastructure for optimizing carbon-efficiency

Software-defined interfaces and systems for programmatically **deploying these optimizations** to reduce carbon emissions

Intuitive interfaces that **account for concerns of privacy**, fairness, equity, incentives, and user preference

Cost of **flexibility**!

The war of the efficiencies: Understanding the tension between carbon and energy optimization

ACM SIGENERGY Energy Informatics Review 2024

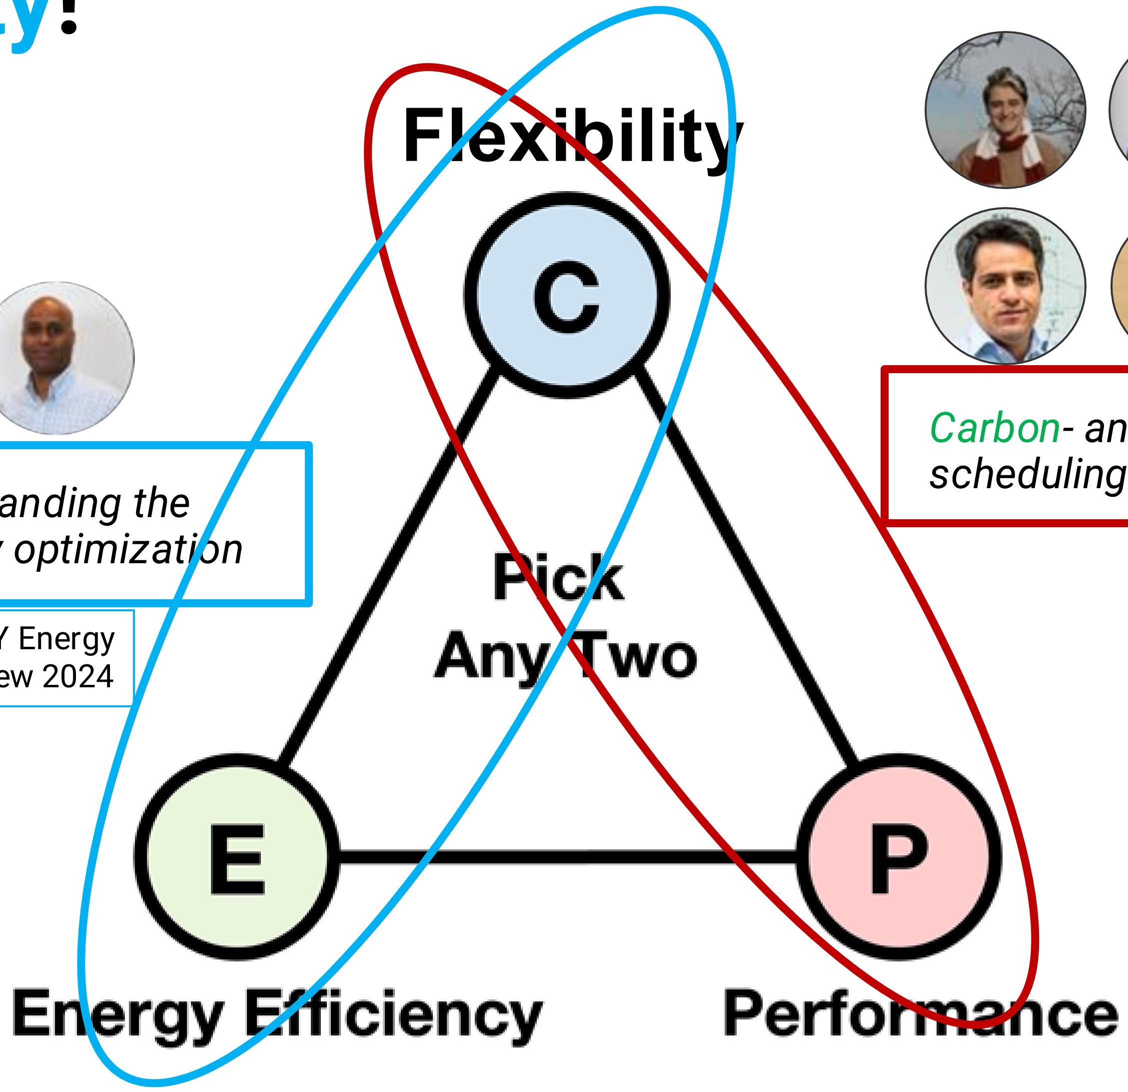

Carbon- and precedence-aware scheduling for data processing clusters

SIGCOMM 2025

- W. Hanafy, et al. "**The war of the efficiencies: Understanding the tension between carbon and energy optimization**." in ACM SIGENERGY Energy Informatics Review 2024
- R. Bostandoost, et al. "**Data-driven Algorithm Selection for Carbon-Aware Scheduling**." in *ACM HotCarbon* 2024
- A. Lechowicz, et al. "**Carbon- and Precedence-Aware Scheduling for Data Processing Clusters**." in *ACM SIGCOMM* 2025 – short paper – full paper under review

Signal- and precedence-aware scheduling for data processing clusters

33% of the lifetime carbon footprint of a foundation model can be attributed to *data processing workloads*.

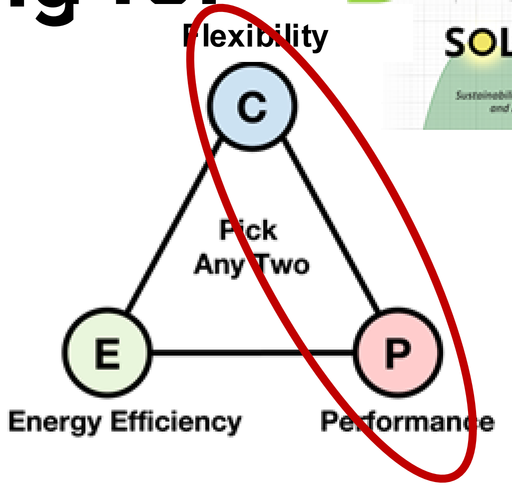

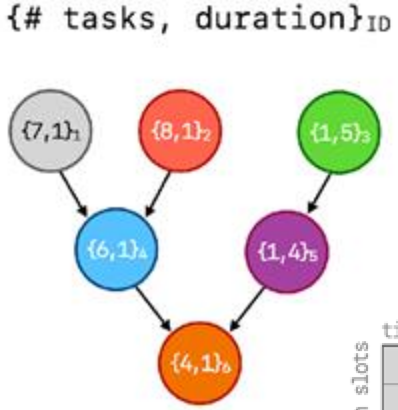

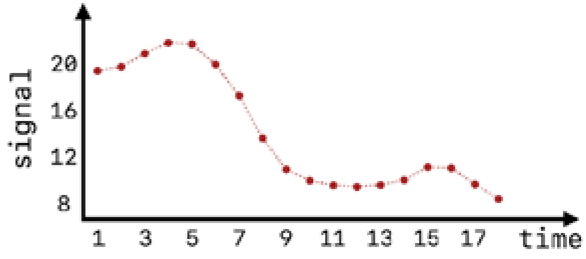

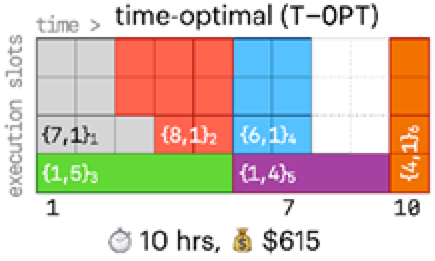

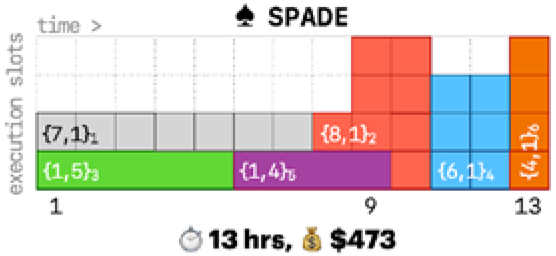

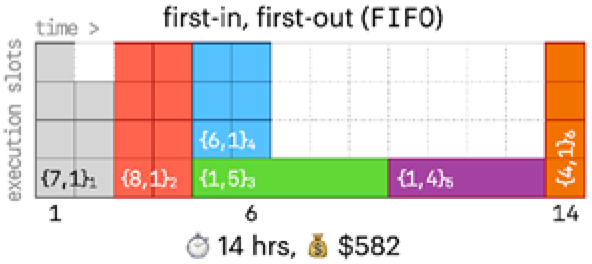

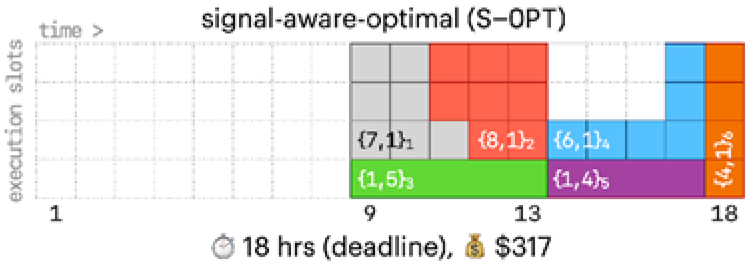

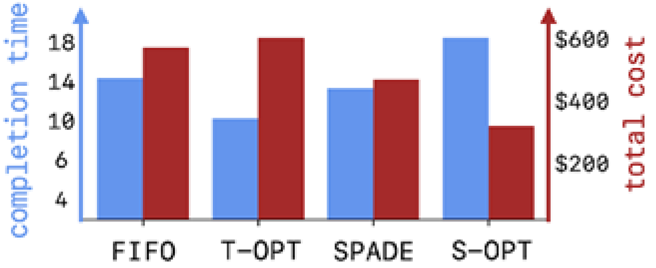

- A. Lechowicz, et al. "**Carbon- and Precedence-Aware Scheduling for Data Processing Clusters**." in *ACM SIGCOMM* 2025

Our main system: SPADE

Scheduling and Provisioning for Adaptive DAG Execution

- Key idea for joint scheduling and provisioning: introduce a metric of *relative importance* to relate the DAG structure against the time-varying signal

Our main system: **SPADE**

- Key idea for joint scheduling and provisioning: introduce a metric of *relative importance* to relate the DAG structure against the time-varying signal

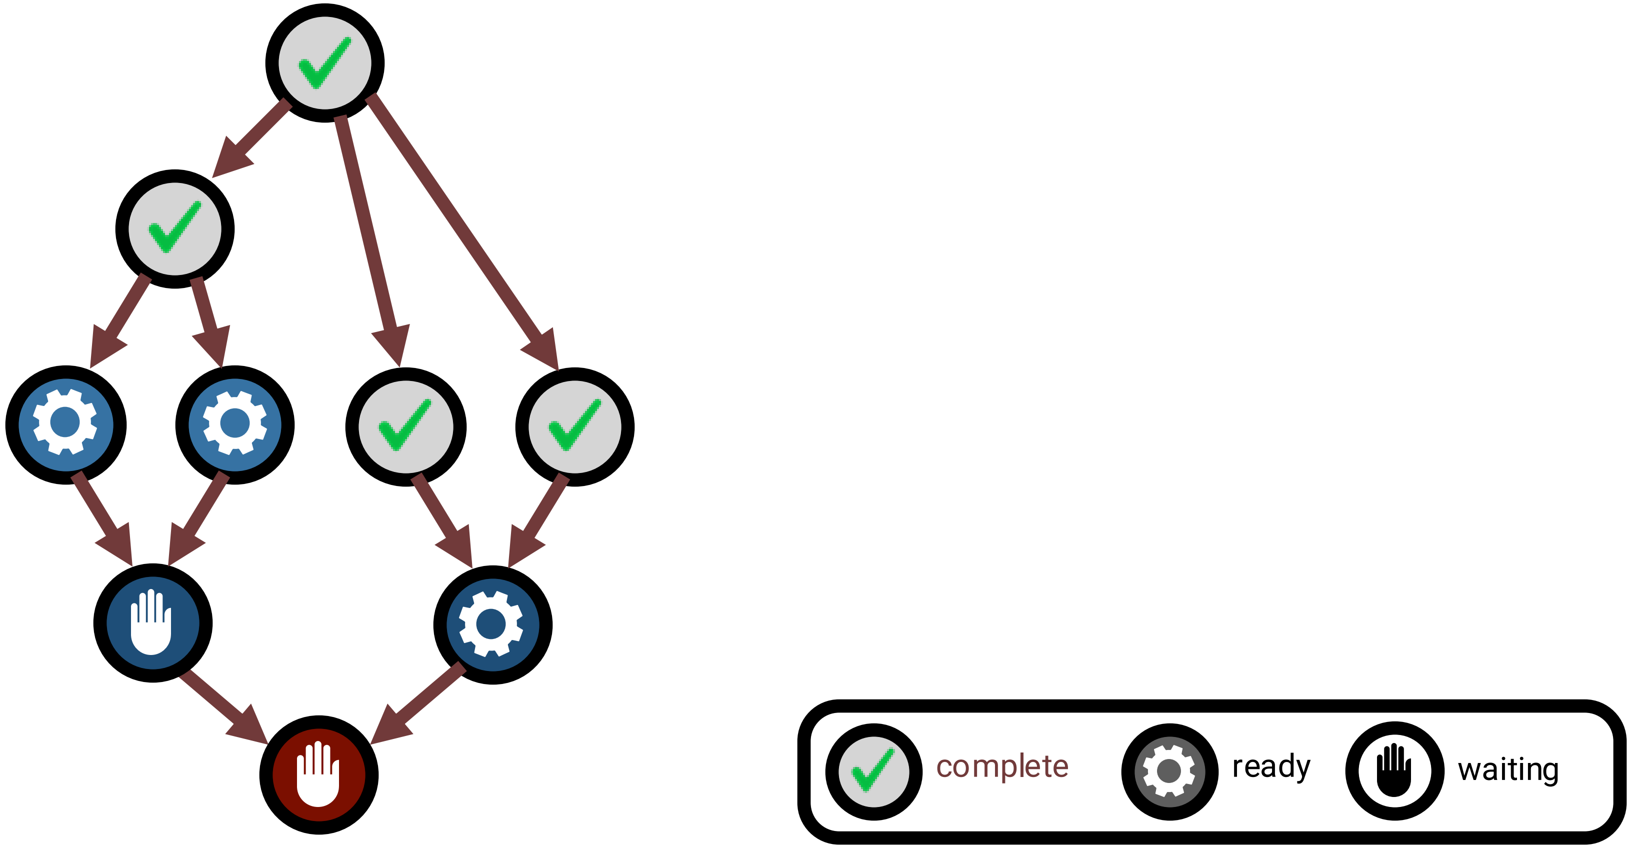

Our main system: SPADE

- Key idea for joint scheduling and provisioning: introduce a metric of *relative importance* to relate the DAG structure against the time-varying signal

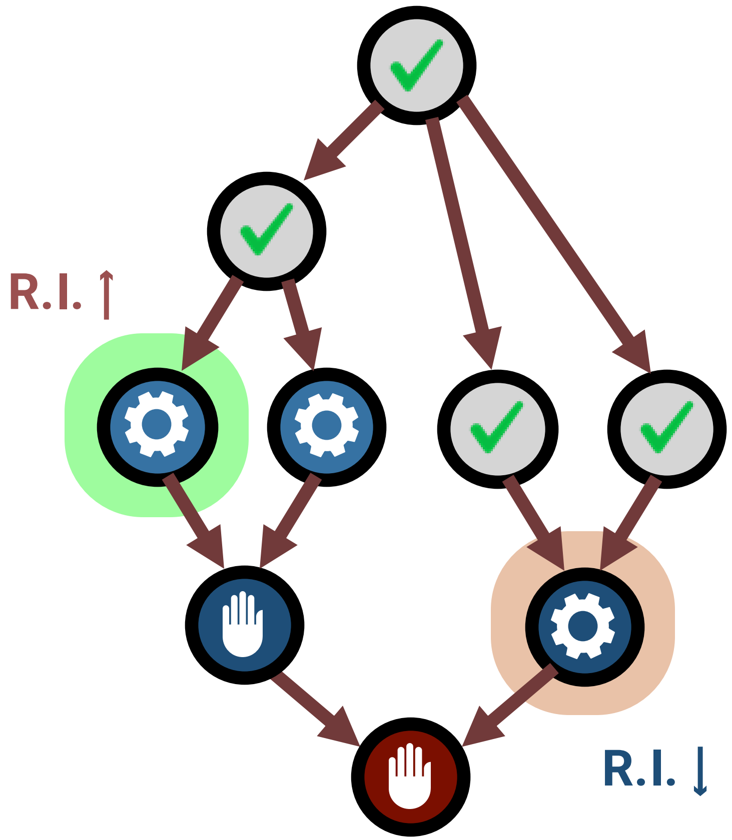

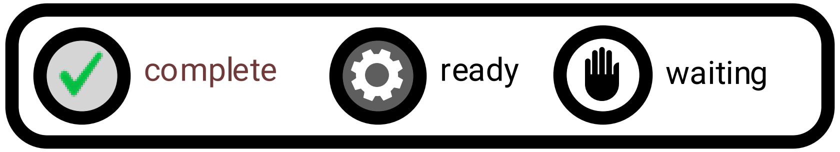

Our main system: SPADE

- Key idea for joint scheduling and provisioning: introduce a metric of *relative importance* to relate the DAG structure against the time-varying signal

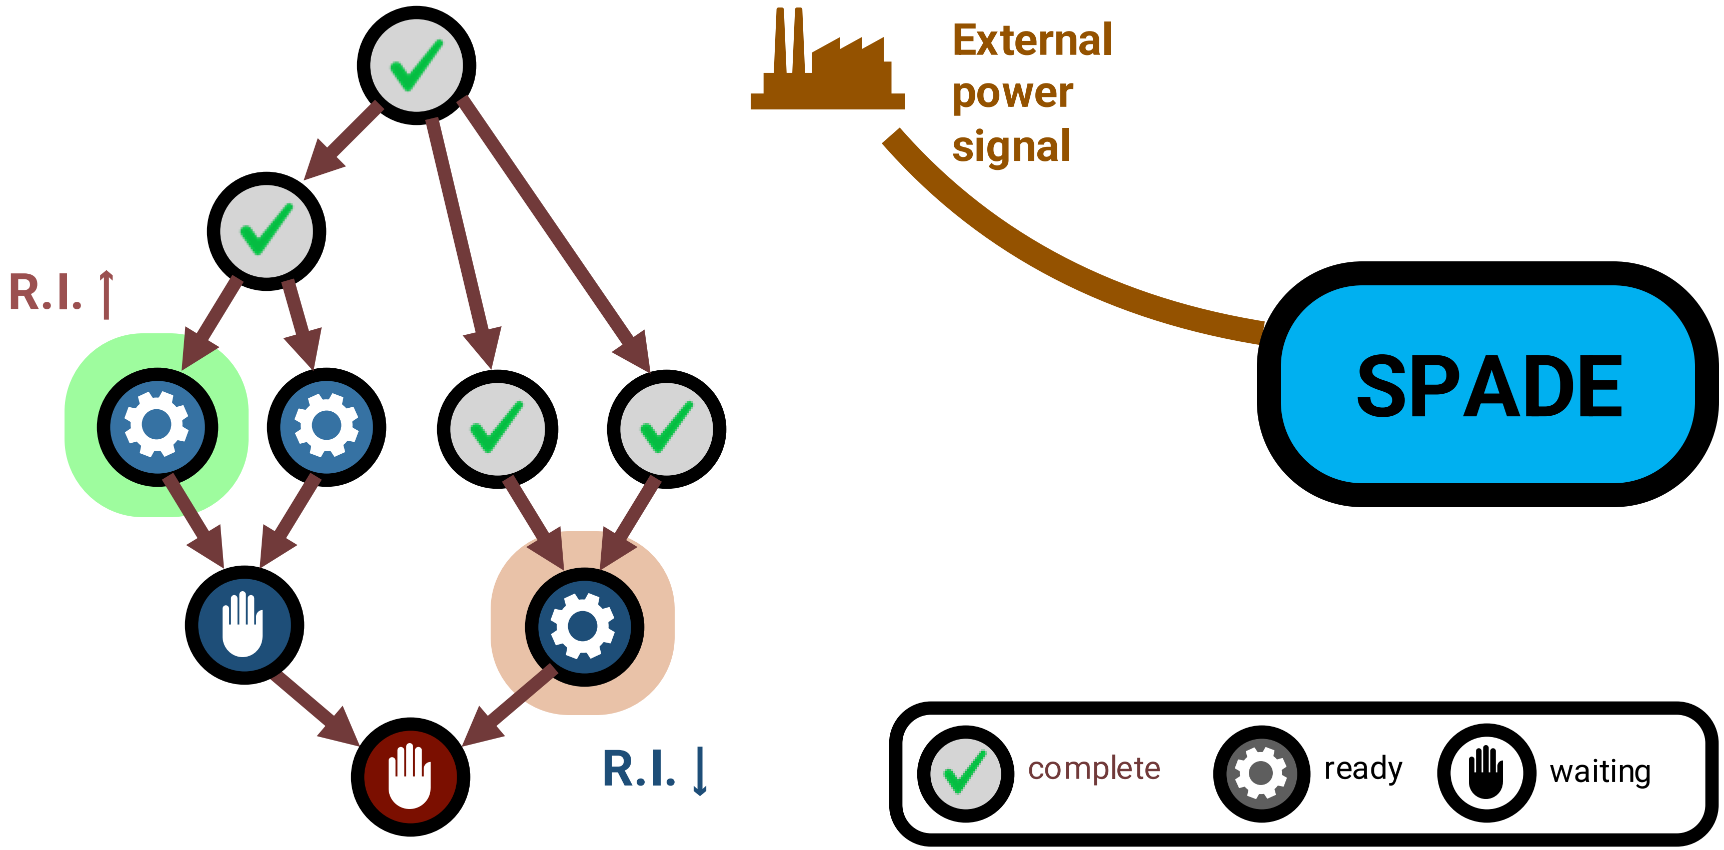

Our main system: SPADE

- Key idea for joint scheduling and provisioning: introduce a metric of *relative importance* to relate the DAG structure against the time-varying signal

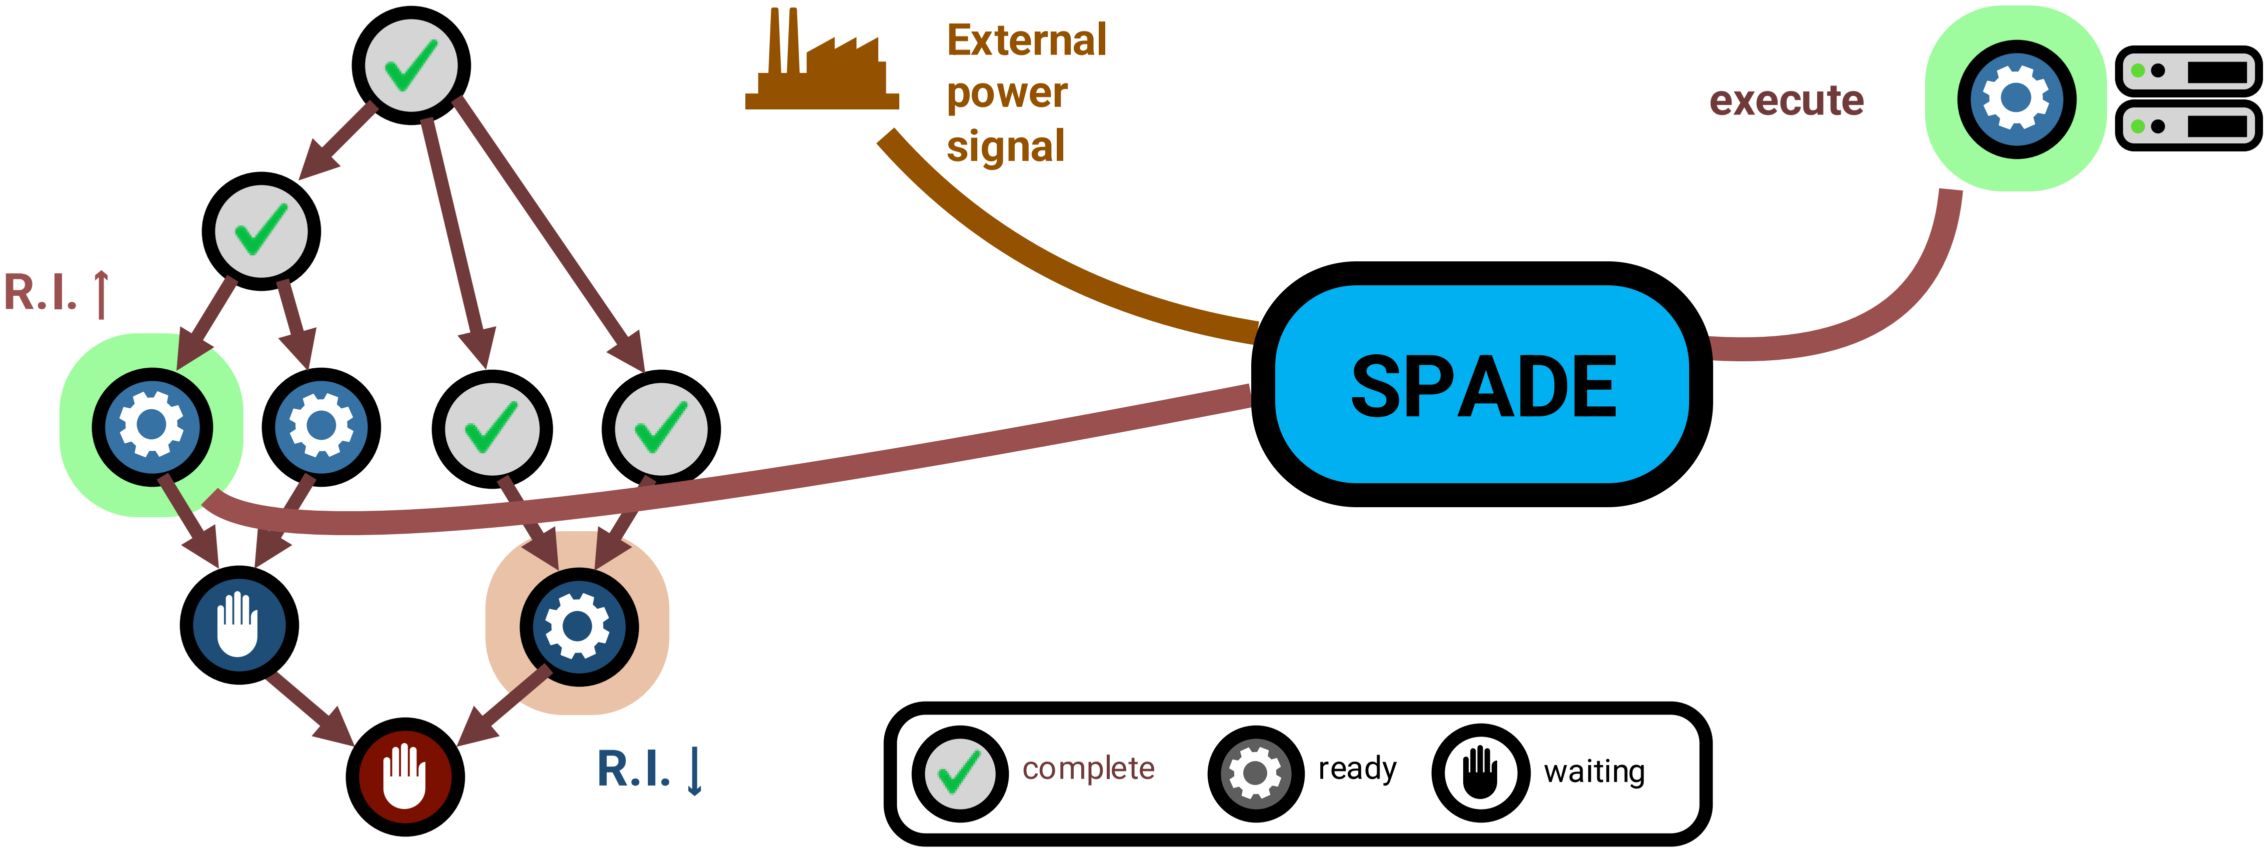

Our main system: **SPADE**

- Key idea for joint scheduling and provisioning: introduce a metric of *relative importance* to relate the DAG structure against the time-varying signal

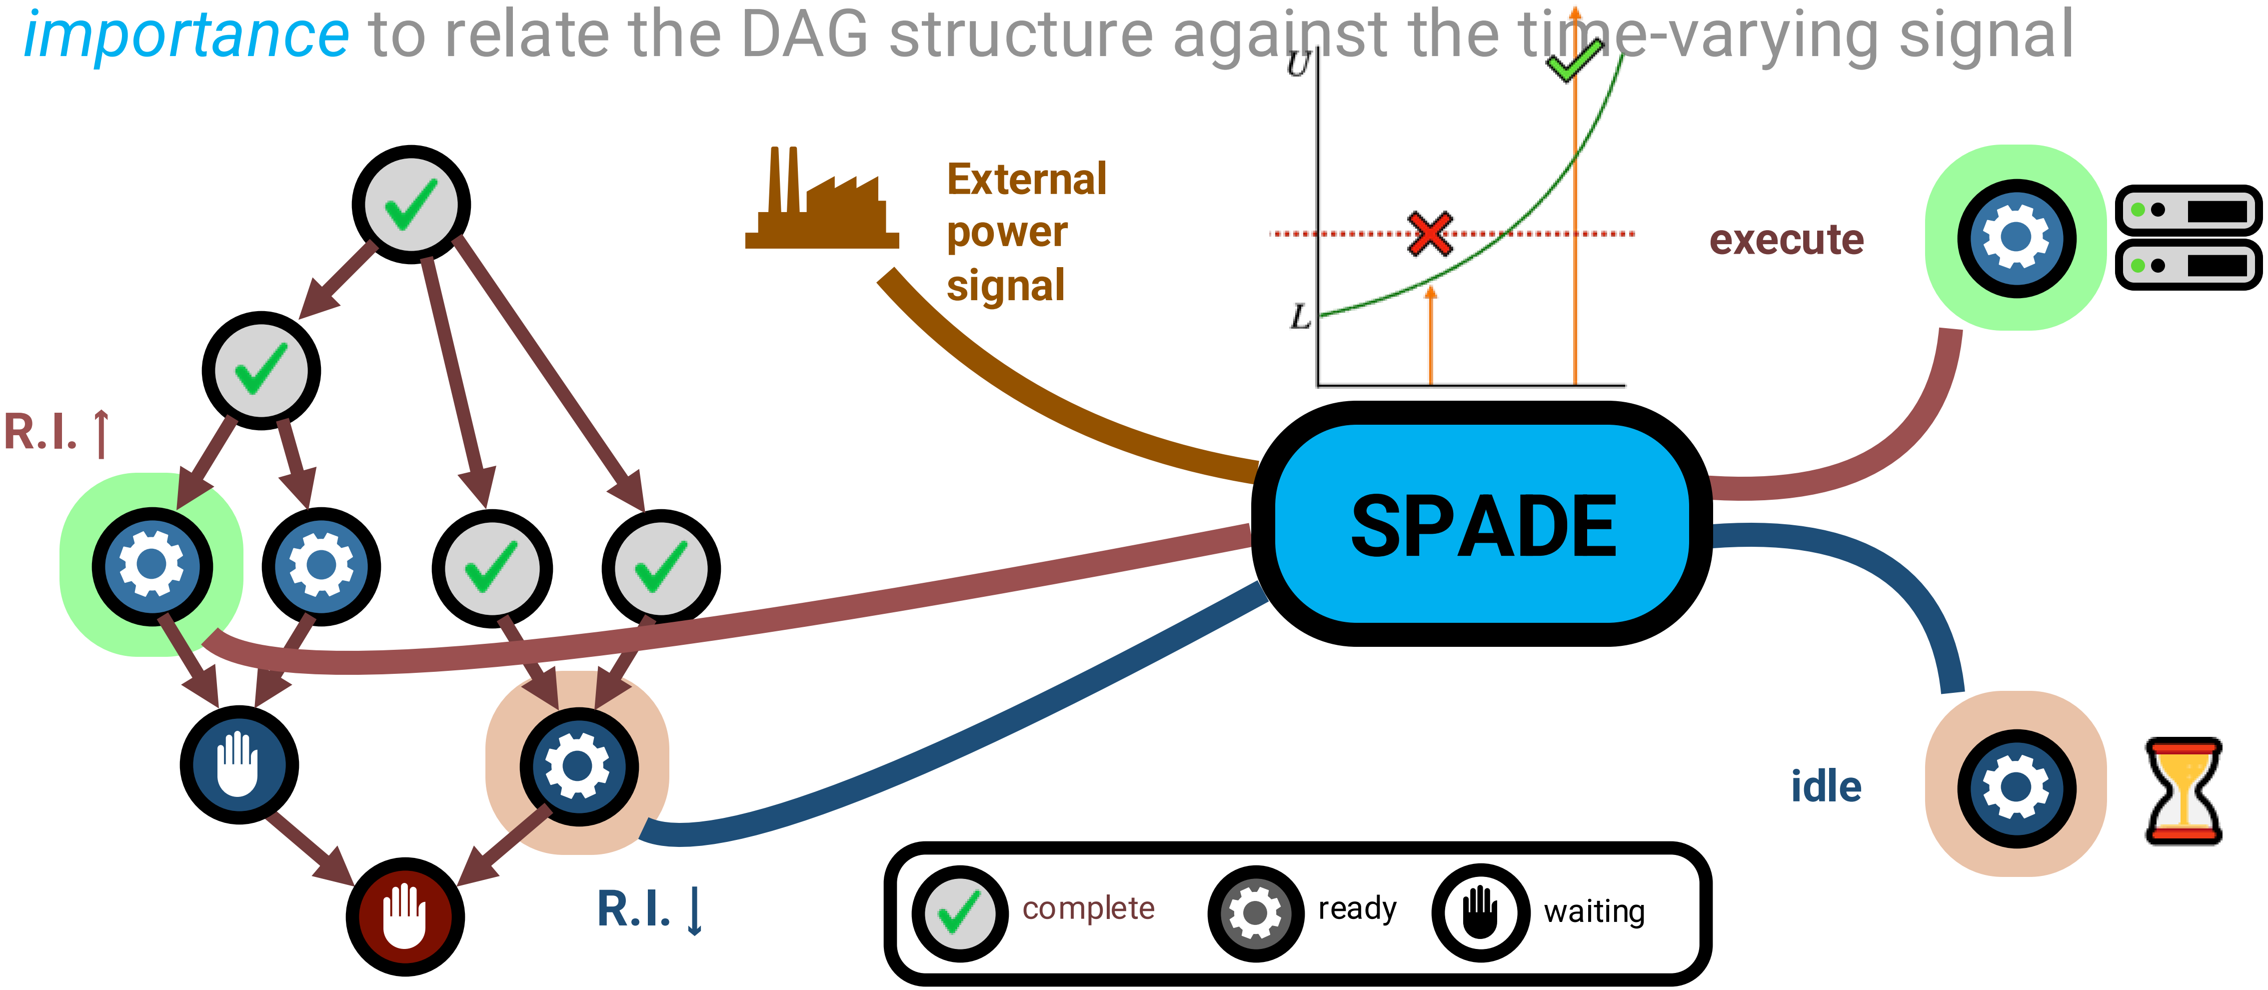

Theoretical results

- **Stretch Factor**: for a signal-aware schedule, the worst-case (multiplicative) increase in makespan compared to a signal-agnostic schedule.
- [Informal Theorem]: Given a time-varying signal s and K resources, the stretch factor of SPADE is $\leq 1 + \frac{D(\gamma,s)K}{2\text{-}1/K}$ for a *signal-awareness parameter* $\gamma \in [0,1]$
  - $D(\gamma, s)$ describes the *fraction of tasks shifted because of a high signal*
  - As γ increases, the signal cost of SPADE decreases, at the expense of a larger stretch factor in the worst case.

Performance evaluation

- Implemented SPADE in simulation env. and 100-node k8s testbed
- SPADE reduces carbon footprint by 32.9% with <5% increase in overall makespan; more noticeable impact on tail JCT
- Strictly better trade-off between performance and carbon compared to "provisioning-only" approach (SAP)

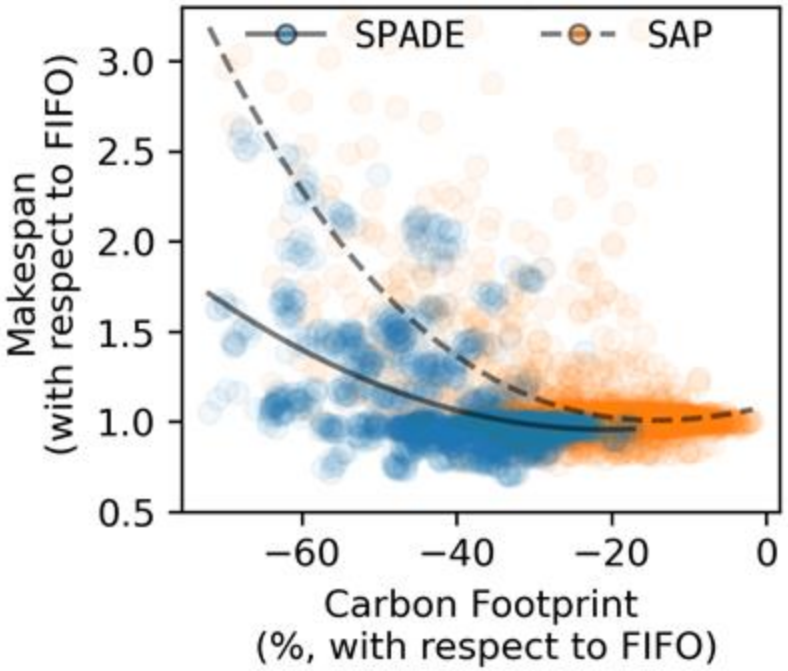

Flexible scheduling and resource allocation

- **The Online pause and resume problem**
  - Carbon-aware workload scheduling (Sigmetrics 2024)
  - Carbon-aware electric vehicle charging management (Sigmetrics 2024)
  - Carbon-aware spatiotemporal workload management (Sigmetrics 2025)

Photo from Google

- A. Lechowicz, et al. "**The online pause and resume problem: optimal algorithms and an application to carbon-aware load shifting**." In *ACM SIGMETRICS* 2024.
- A. Lechowicz, et al. "**Online conversion with switching costs: Robust and learning-augmented algorithms**." In *ACM SIGMETRICS* 2024.
- A. Lechowicz, et al. "**Learning-augmented competitive algorithms for spatiotemporal online allocation with deadline constraints**." In *ACM SIGMETRICS* 2025.

Flexible scheduling and resource allocation

Theoretical Foundations (bridging two lines of research)

Convex body chasing

(local and short-term switching costs)

Online packing/covering

(enforcing long-term capacity/budget or demand constraints)

Foundational CoDec problems feature both requirements in a unified problem

- A. Lechowicz, et al. "**Chasing convex functions with long-term constraints: optimal consistency-robustness tradeoffs**." In ICML 2024

Conclusion

Yes, we are dealing with order-of-magnitude larger problems!

Yes, the problems are at the intersection of multiple infrastructures!

But, the problems calls for fundamentally different ideas and techniques to solve!

This talk: joint scheduling and provisioning for joint minimization of first-order metrics (job completion time), and second-order metrics (the cost of the external signal)









Expeditions: Computational Decarbonization of Societal Infrastructures at Mesoscales (CoDec)

University of Massachusetts Amherst

Carnegie Mellon University

Massachusetts Institute of Technology

UCLA

The University of Chicago

Wisconsin
University of Wisconsin–Madison