OpenInfra

Co-simulation Framework for the Infrastructure Nexus

Jiaheng Lu, Yunming Xiao, Shmeelok Chakraborty, Silvery Fu, Yoon Sung Ji, Ang Chen, Mosharaf Chowdhury, Nalini Rao, Sylvia Ratnasamy, Xinyu Wang

HotInfra' 24

What is the Infrastructure

Infrastructure:

- Cloud Infrastructure: Datacenter
- Industrial Infrastructure: Power Plant, Water System

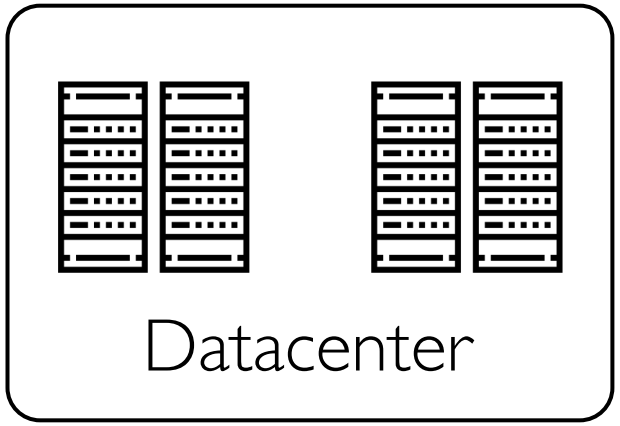

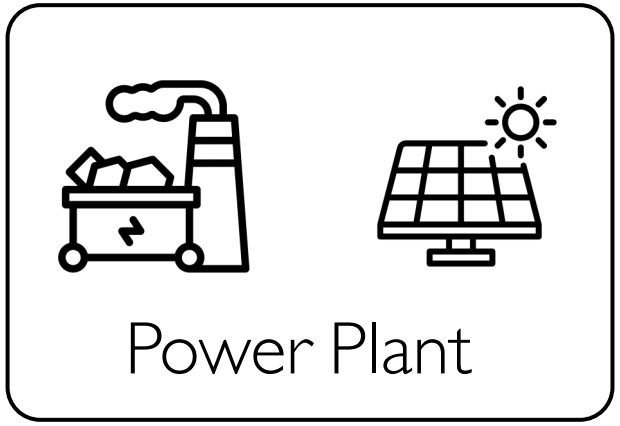

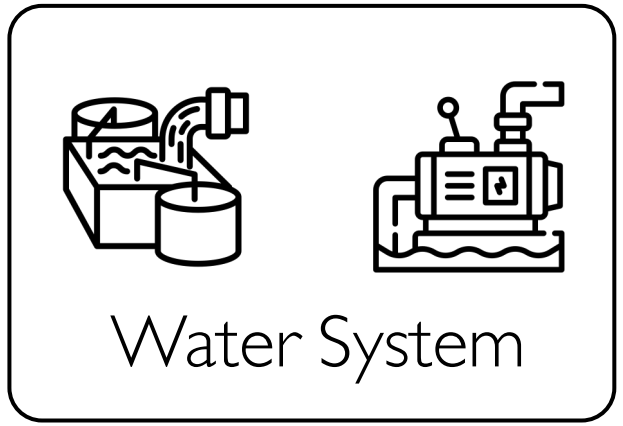

What is the Infrastructure Nexus

Datacenter/X Nexus: Datacenter/Energy, Datacenter/Water:

- Datacenter requires energy to power its facilities
- Datacenter requires water to cool IT equipment

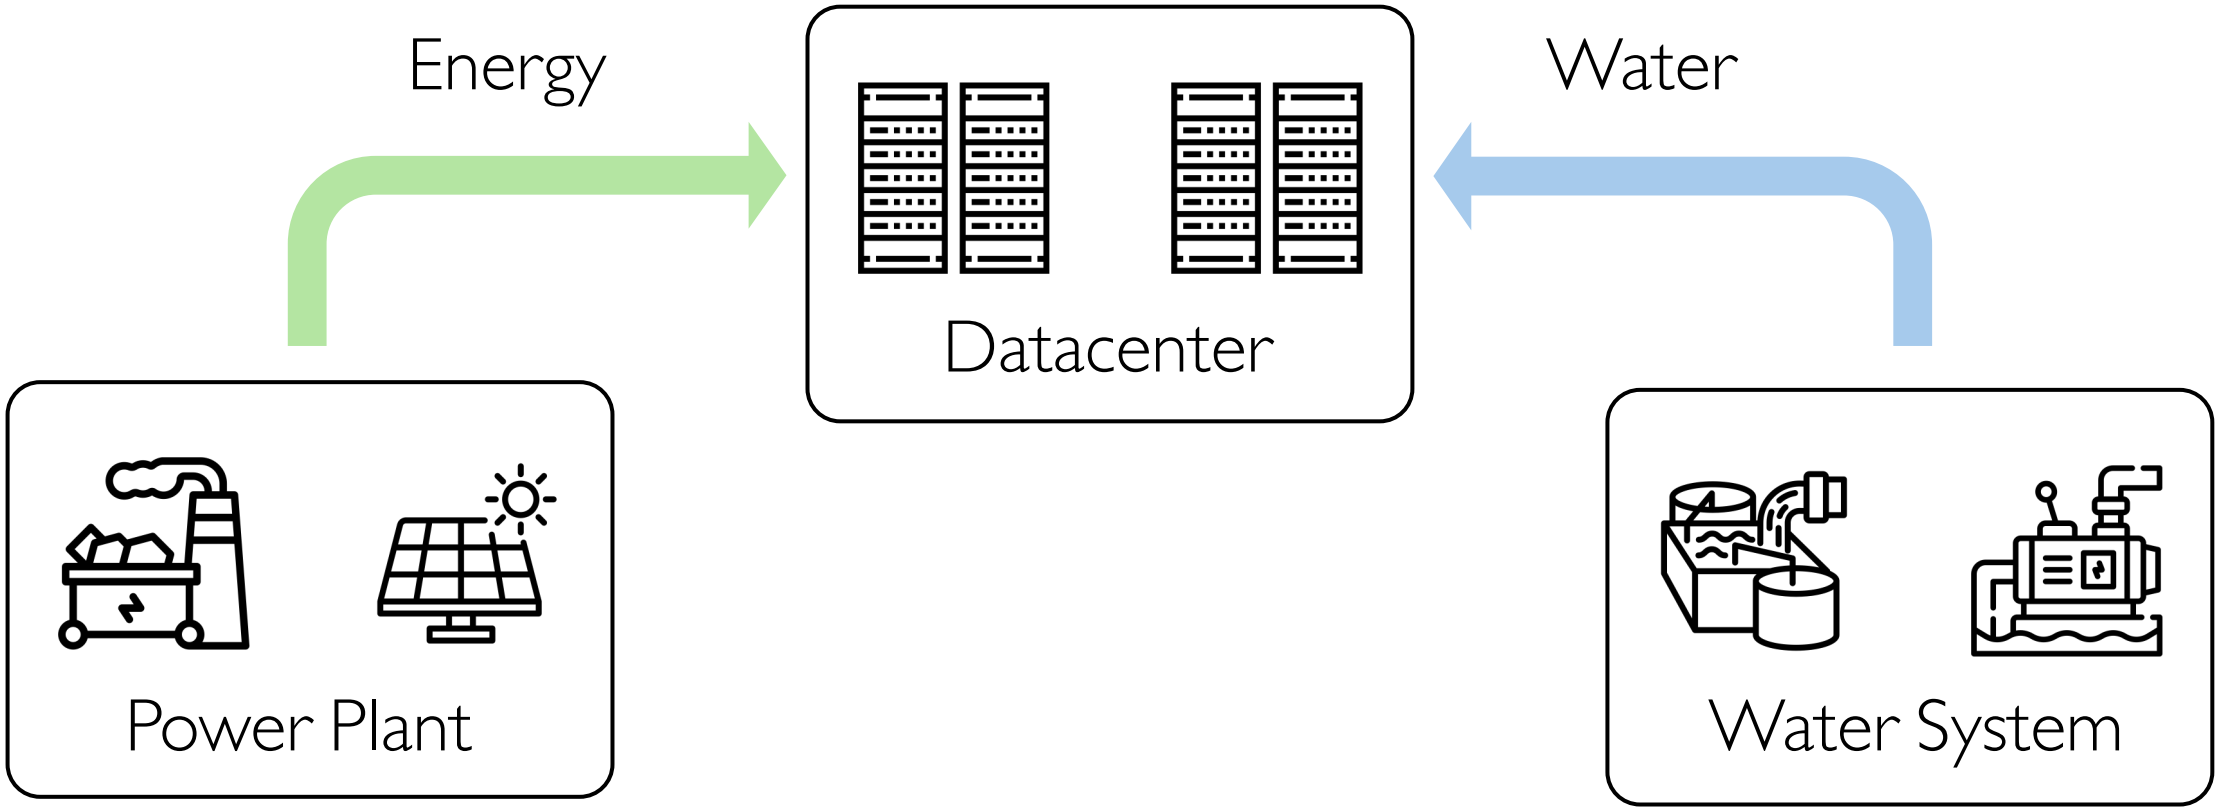

What is the Infrastructure Nexus

Power/Water Nexus:

- Power Plant requires water to cool its equipment
- Water System requires energy to pump and filer water

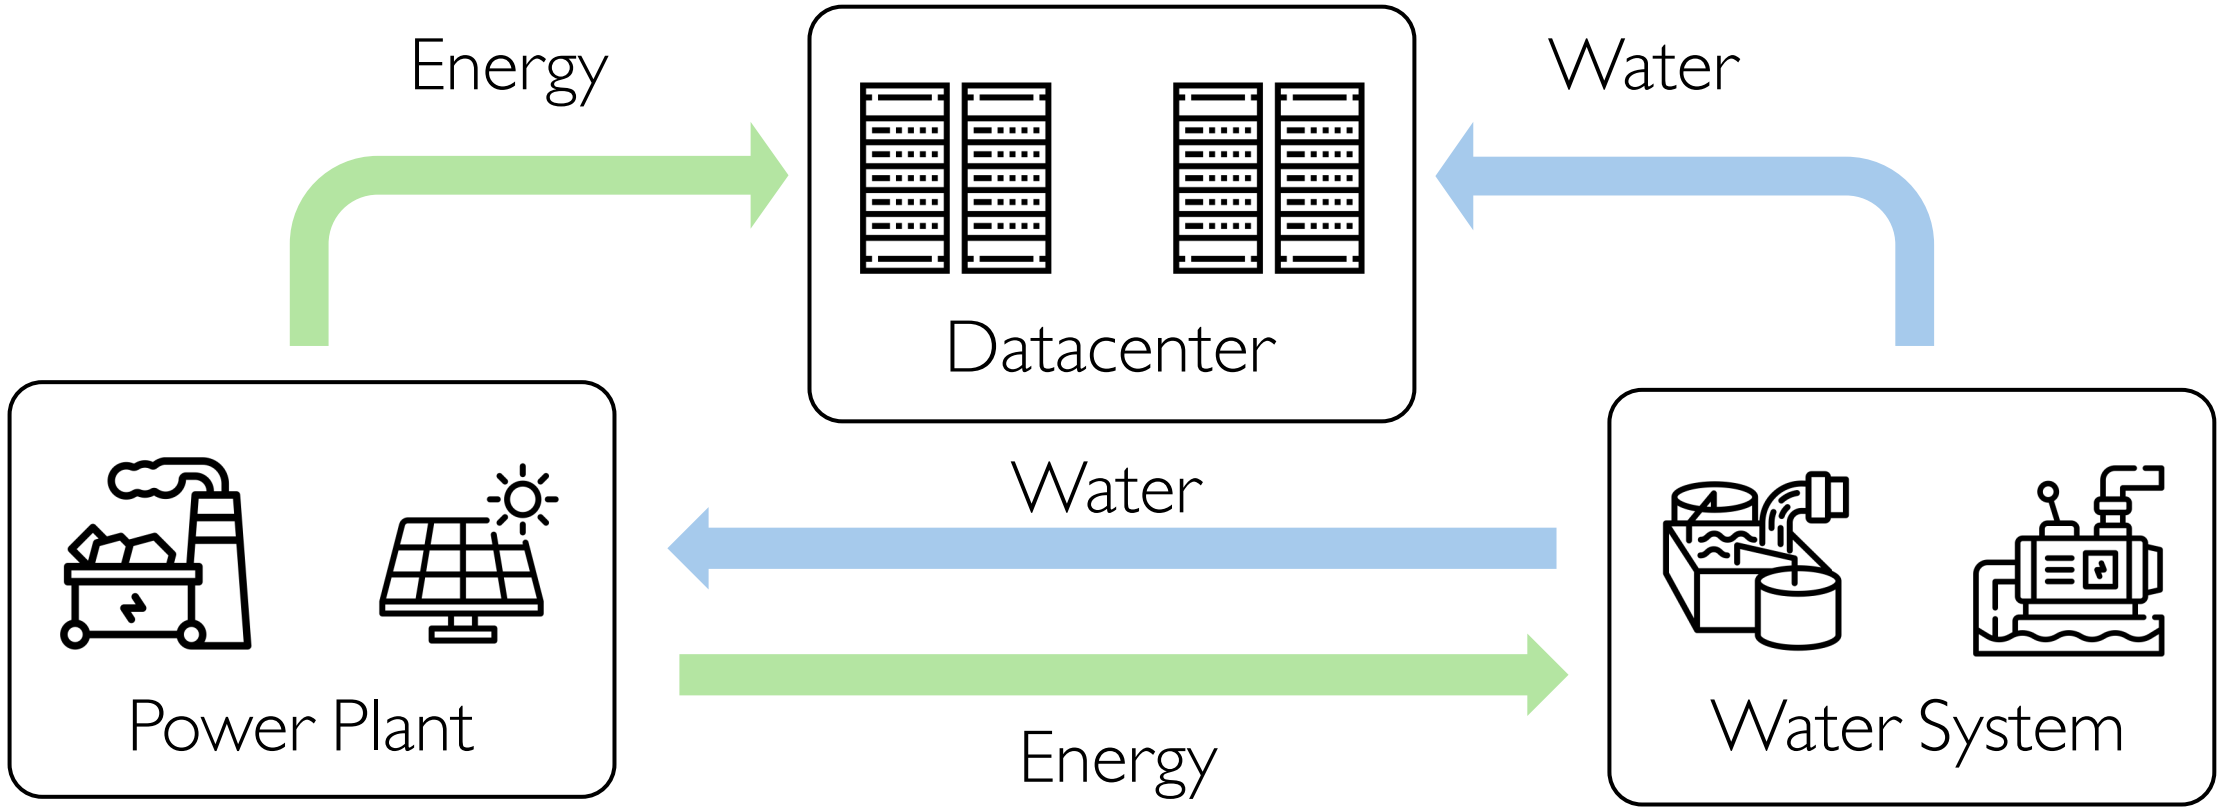

Why is the Infrastructure Nexus important?

- Datacenter water usage increases quickly

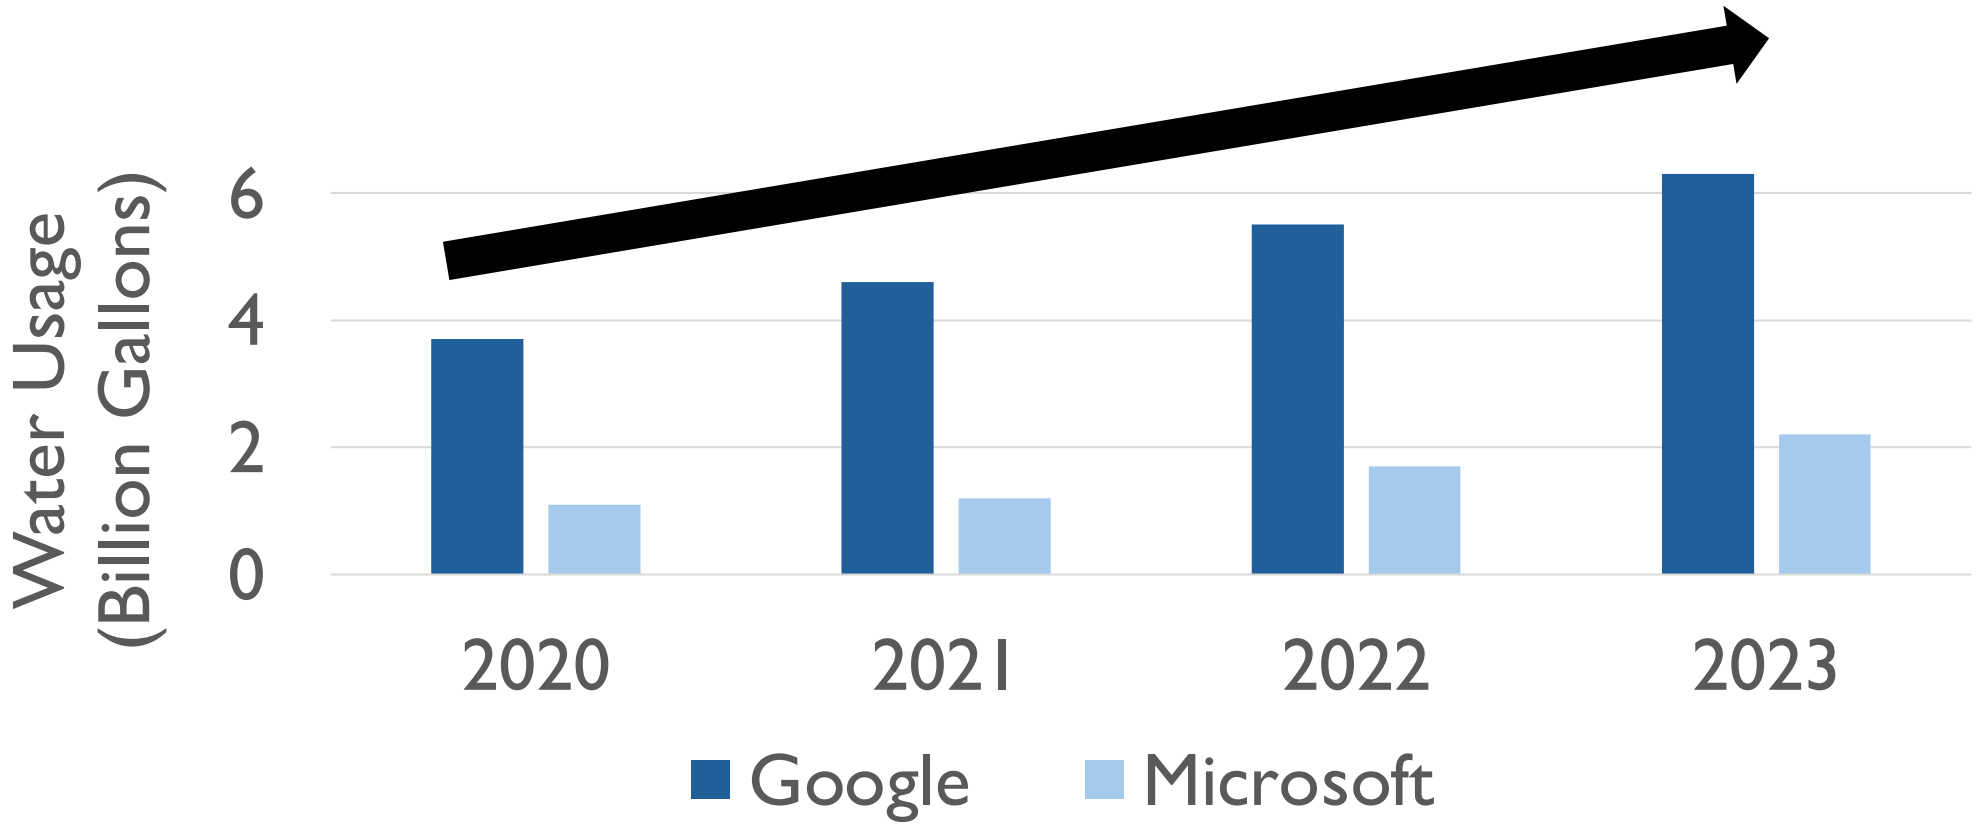

Source:
https://www.ft.com/content/1d468bd2-6712-4cdd-ac71-21e0ace2d048?segmentId=2c1df321-36a4-1206-2c08-112c059dd69d

Why is the Infrastructure Nexus important?

- Datacenter power usage keeps increasing

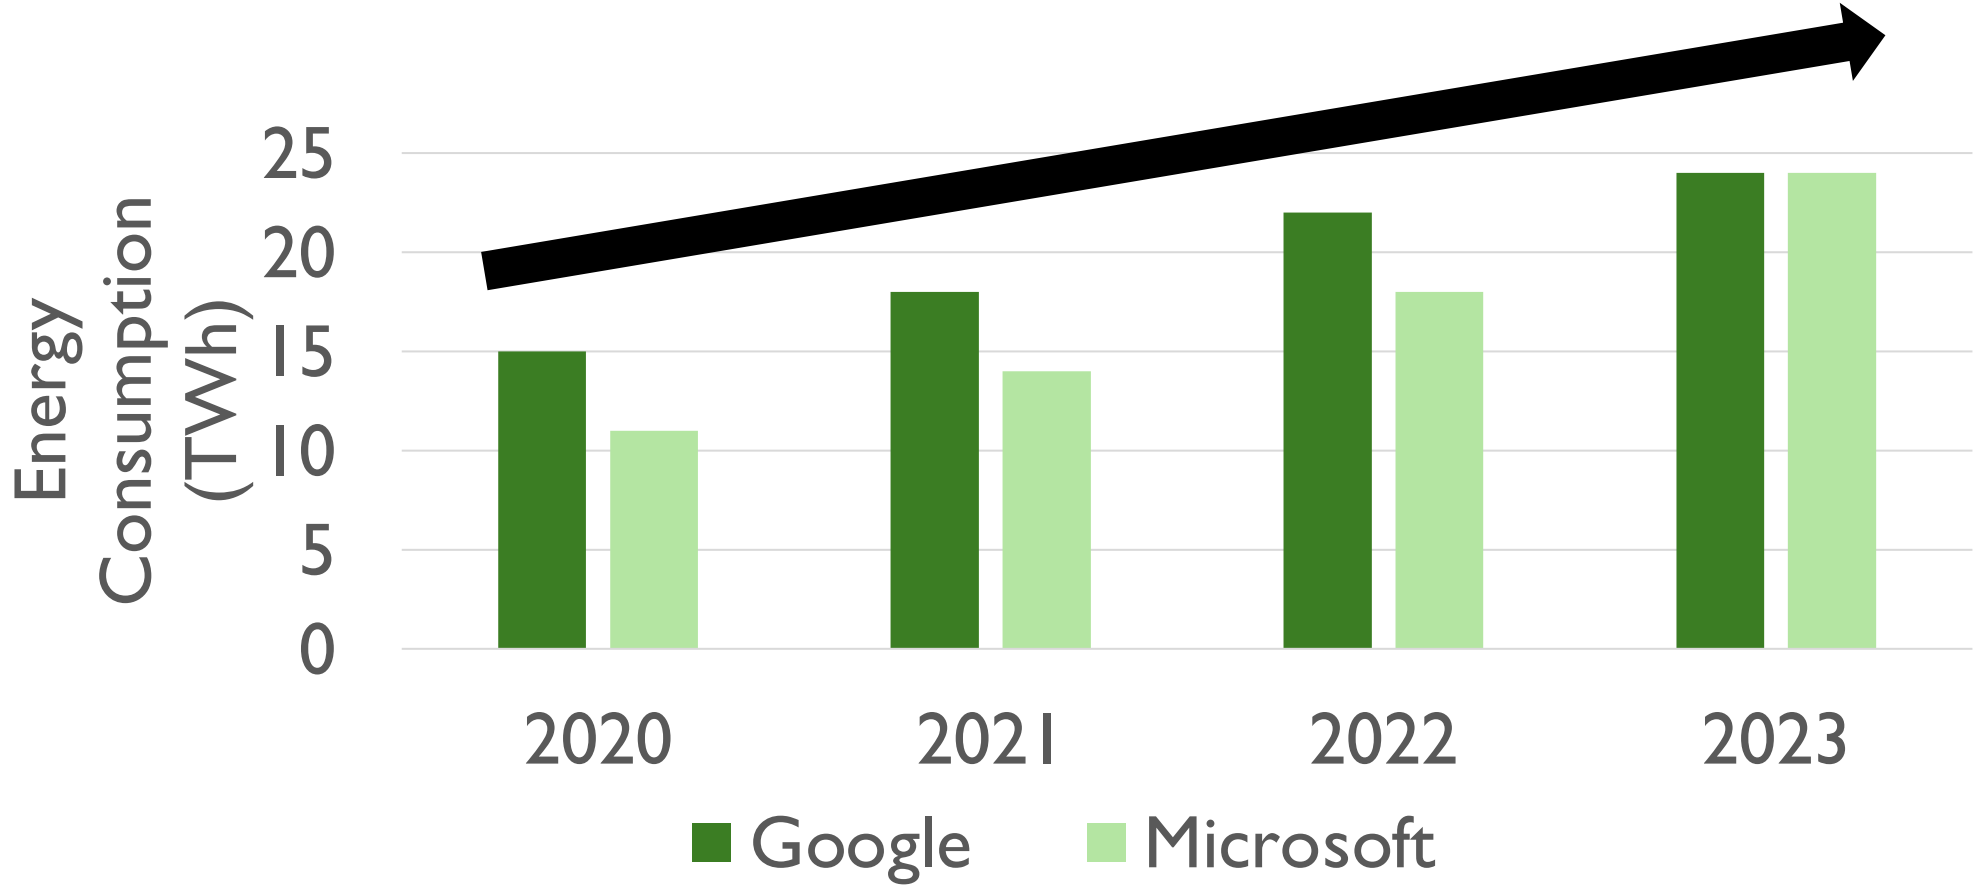

Source:
https://www.goldmansachs.com/insights/articles/AI-poised-to-drive-160-increase-in-power-demand
https://www.visualcapitalist.com/microsofts-electricity-use-has-doubled-between-2020-2023/
https://www.statista.com/statistics/788540/energy-consumption-of-google/

Opportunities at the Infrastructure Nexus

Co-optimization to balance power and water usage trade-offs

- Cooling datacenters at night is more water-efficient because of low temperature
- Powering datacenters at night is less carbon-efficient because of less renewable energy

Why we need the co-simulation framework

Problem:

- Direct evaluation on multiple physical infrastructures is hard!
- Requires expertise across multiple domains.

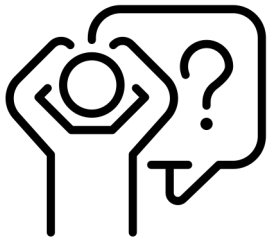

What is co-simulation:

- A complex system composed of individual simulators.

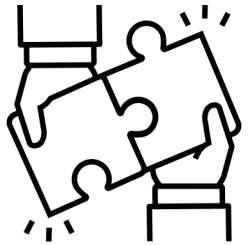

Why co-simulation:

- Easy integration of high-fidelity simulators from multiple domains.

OpenInfra goals

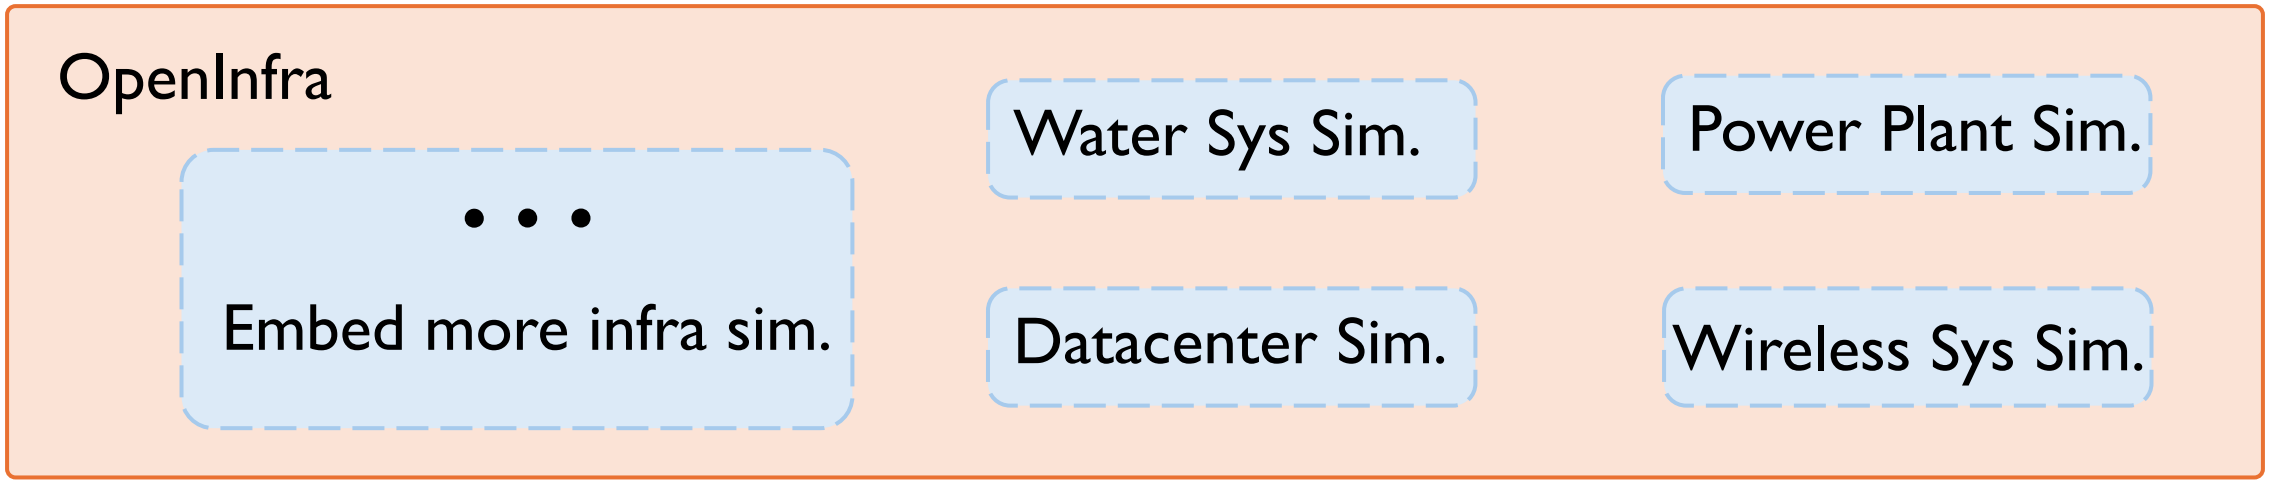

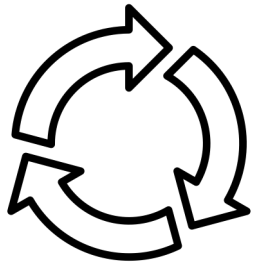

Reuse Simulators

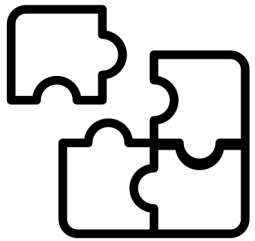

Seamless Integration

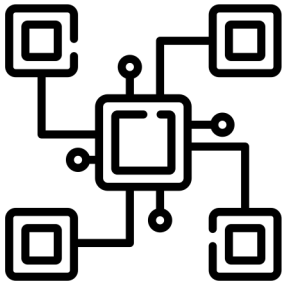

Simulator Interactions

Co-simulation for Infrastructure Nexus is hard

Heterogenous types of simulators required

Each types has many simulators

OMNet++

Datacenter simulator

Power Plant simulator

5G simulator

System level challenges

- Language:

Requiring an intuitive language to define nexus experiments

- Execution:

Performance limited by slowest simulators and resource constraints.

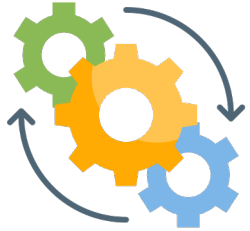

- Scalability:

High communication overhead in distributed simulators.

System level challenges - Cont'd

- Algorithm

Easy integration with existing libraries for various simulation goals.

- Synchronization

Synchronize and translate events across simulators.

- Data management

Handle diverse data formats across simulators.

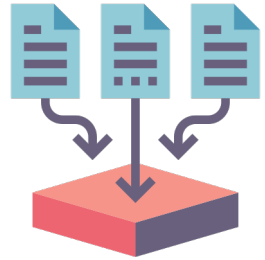

OpenInfra Architecture

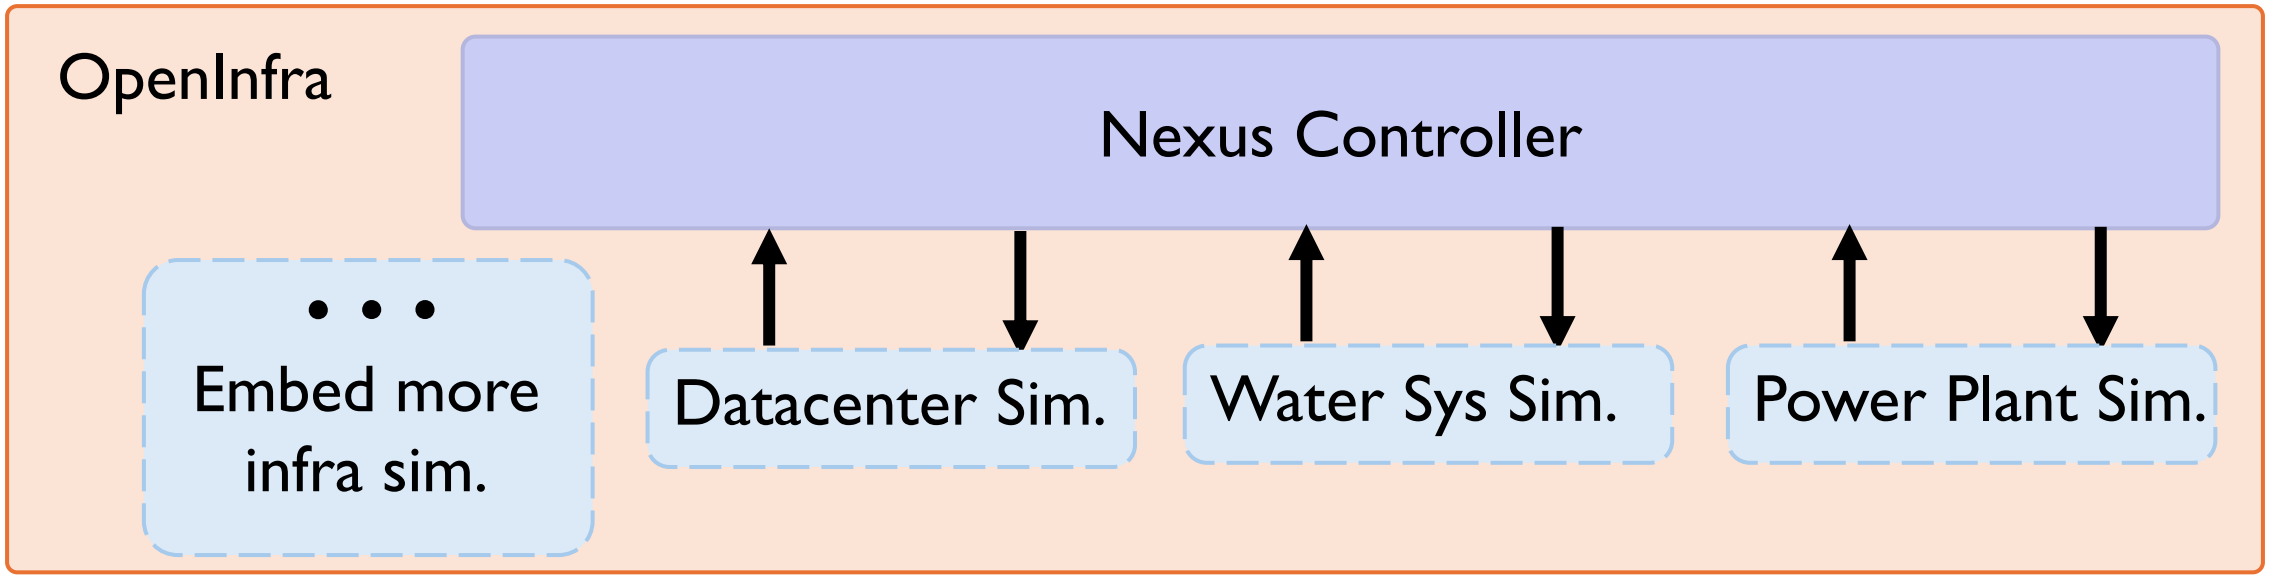

- State(): expose internal state variables from a simulator to OpenInfra
- Actuate(): trigger specific actions within a simulator
- **Capabilities**: Co-optimization and specific event creation (e.g., infrastructure failure)

OpenInfra Architecture - Cont'd

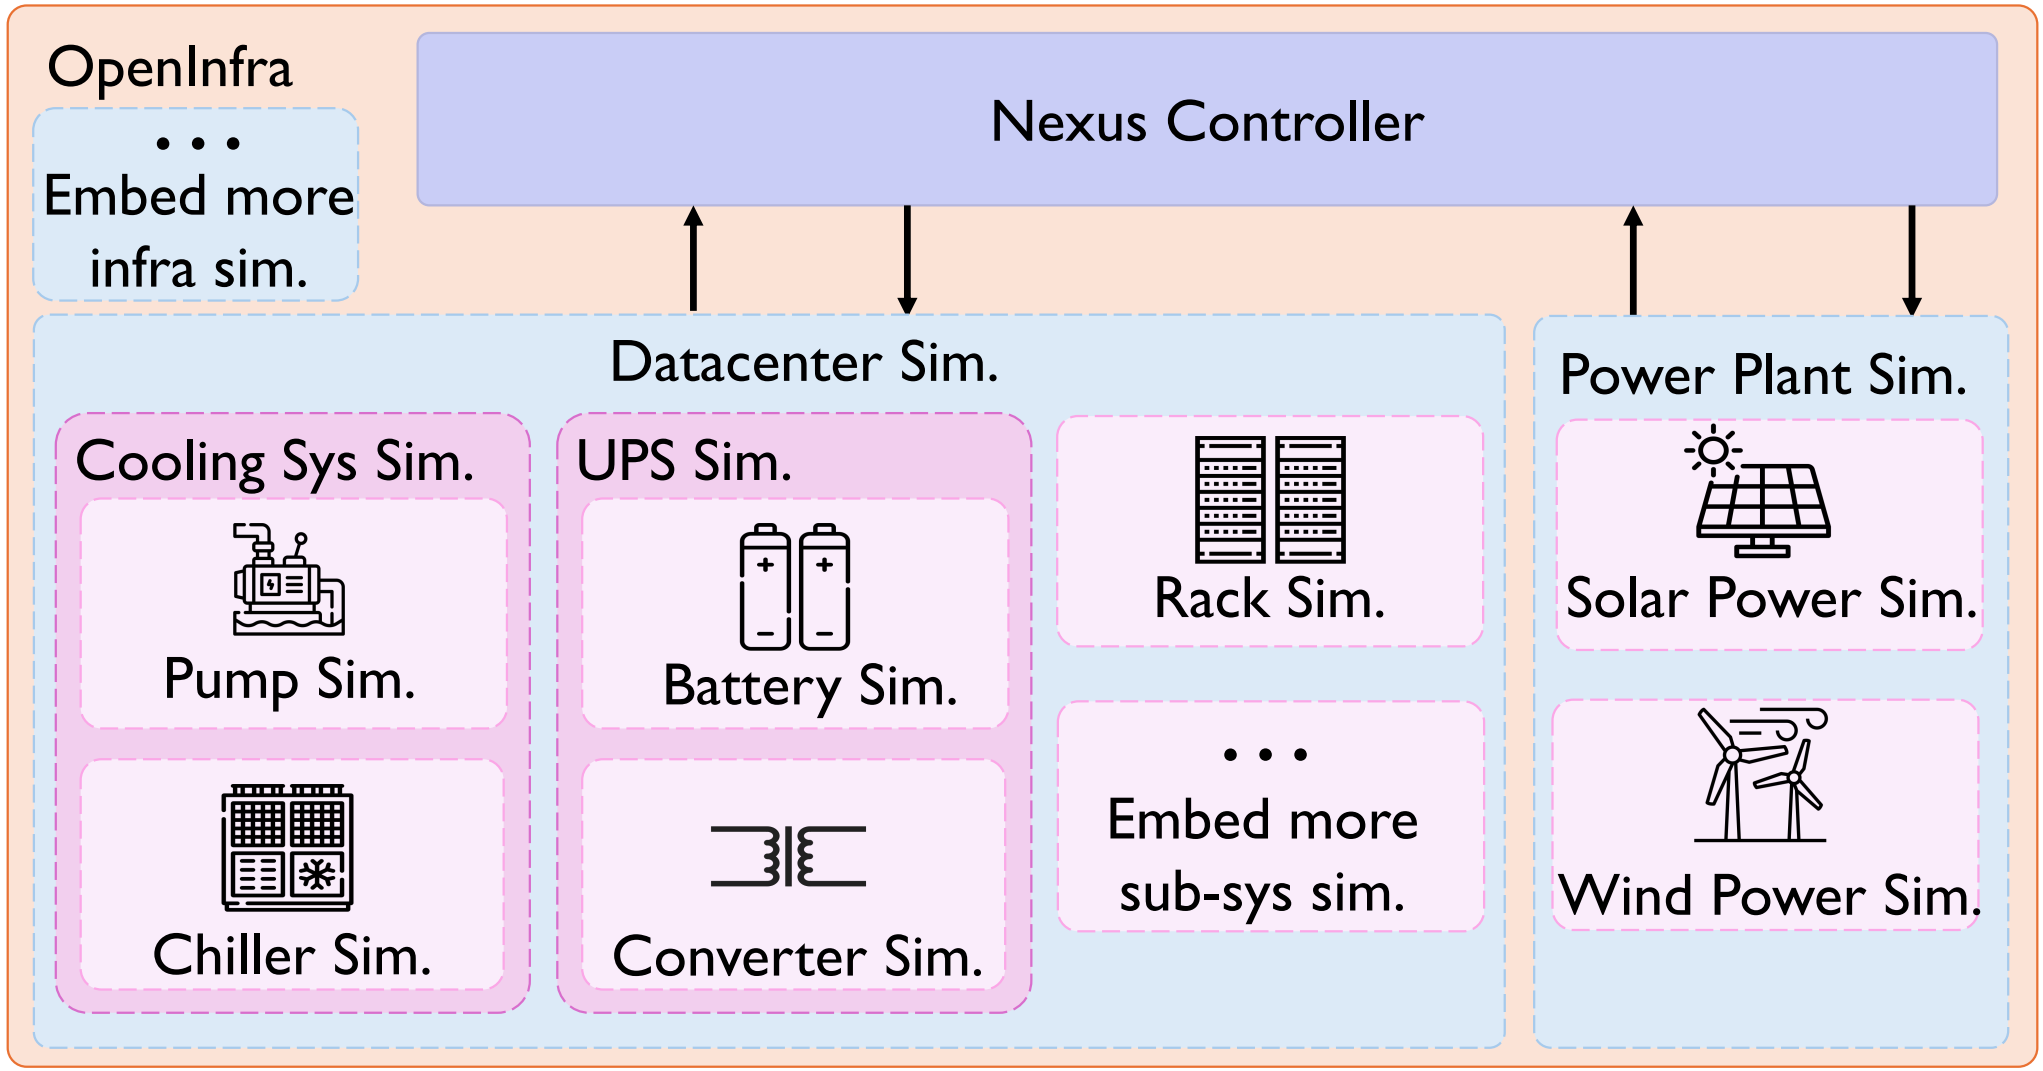

Infrastructure-as-Code management

- Multi-level IaC to fit multiple hierarchies of simulation

```
1  "datacenter": {
2    "granularity": "rack_level",
3    "rack_number": 2,
4    "server_number_per_rack": 52,
5    "UPS_redundancy": "n+1",
6    "UPS_battery_attribute": "...",
7    "connected_rack_number_per_PDU": "1",
8    "PDU_redundancy": "2n"
9  }
```

Initial Stage IaC file

```
1  // Other components ...
2  "Racks": [
3      {
4        "rack_id": 1,
5        "priority": 1,
6        "number_of_servers": 52
7      },
8      {
9        "rack_id": 2,
10       "priority": 2,
11       "number_of_servers": 10 // original 52
12     }
13 ]
```

Final Stage IaC file

What we have supported

- **Integrate 17 simulators**
  - 8 types of power plant simulators
  - 3 types of Uninterruptible Power Supply (UPS) simulators
  - Datacenter rack (load/power usage) simulators
  - Pump/Chiller simulators
  - Datacenter/Power plant water usage simulators
- Working prototype available on GitHub: https://github.com/JhengLu/OpenInfra
  - Written in Python with 4,000 lines of code
  - Active expansions and improvements are in progress
  - All contributions are welcome!

Simulation results

Setting:

- Data center with 7,392 servers and 5 UPS units, running for 100 hours
- Two renewable power plants: wind and solar
- Power plant follows the EIA trace[1]
- Server load follows the Google trace[2]

Speed:

- Runs in 5 seconds on MacBook Pro M2 Max

Source:
[1]: https://www.eia.gov/
[2] M. Tirmazi, A. Barker, N. Deng, M. E. Haque, Z. G. Qin, S. Hand, M. Harchol-Balter, and J. Wilkes. Borg: the next generation. In Proceedings of the fifteenth European conference on computer systems, pages 1–14, 2020

Simulation results

0-10h

- consumption > renewable supply, battery SoC drops

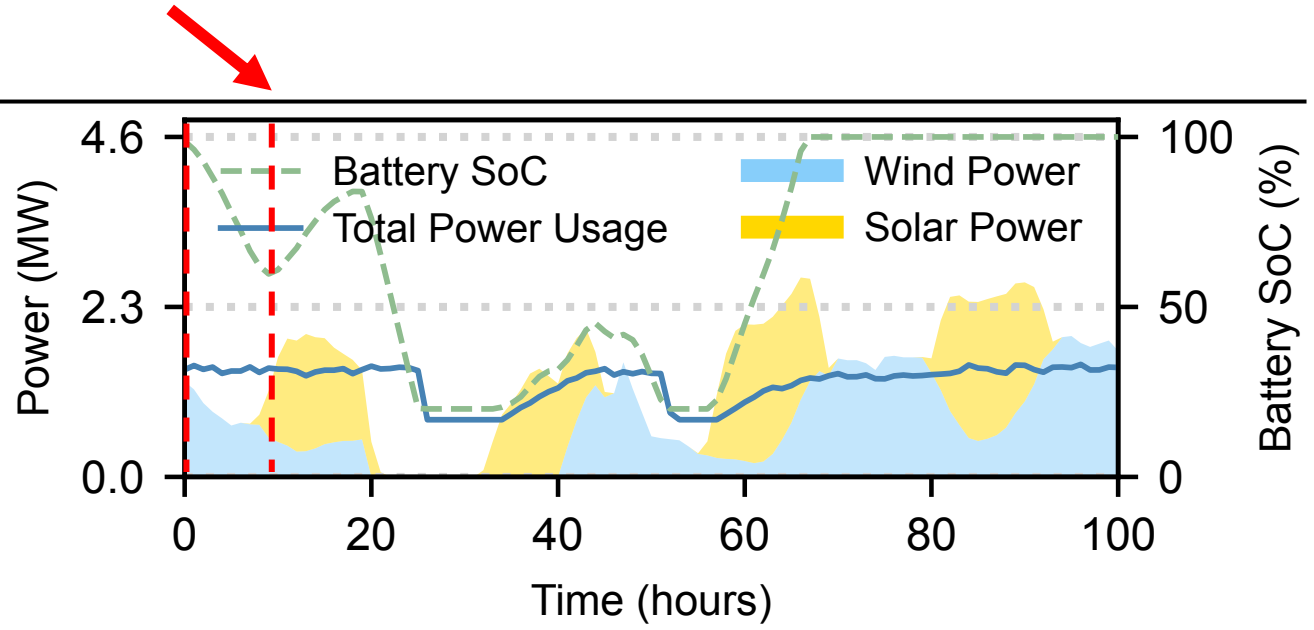

(a) Power Usage and Generation

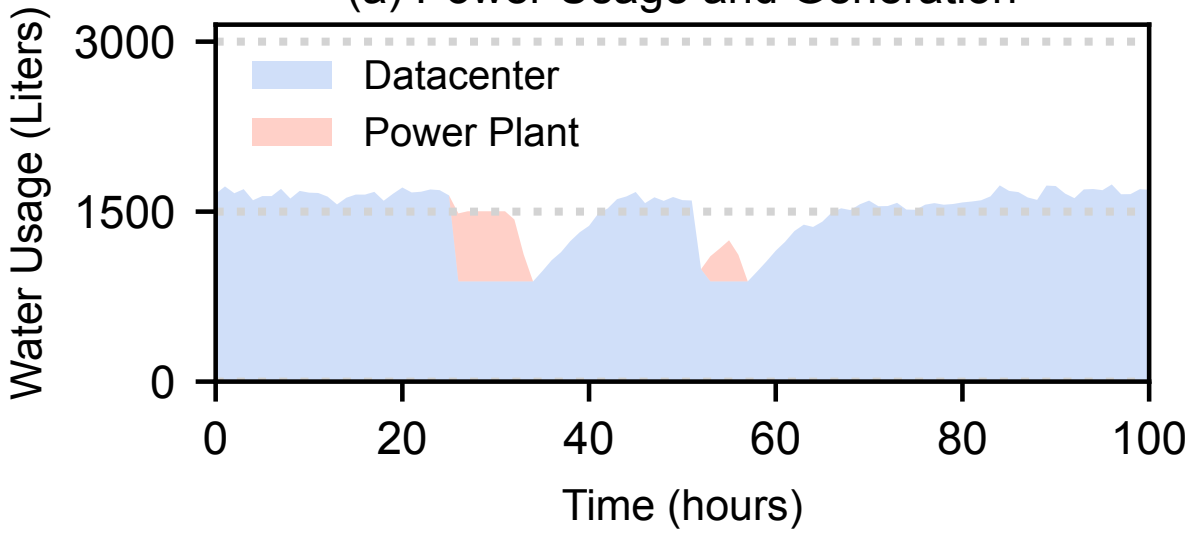

(b) Water Usage (Datacenter and Power Plant)

Simulation results

10-20h

- Consumption < renewable supply, battery recharge

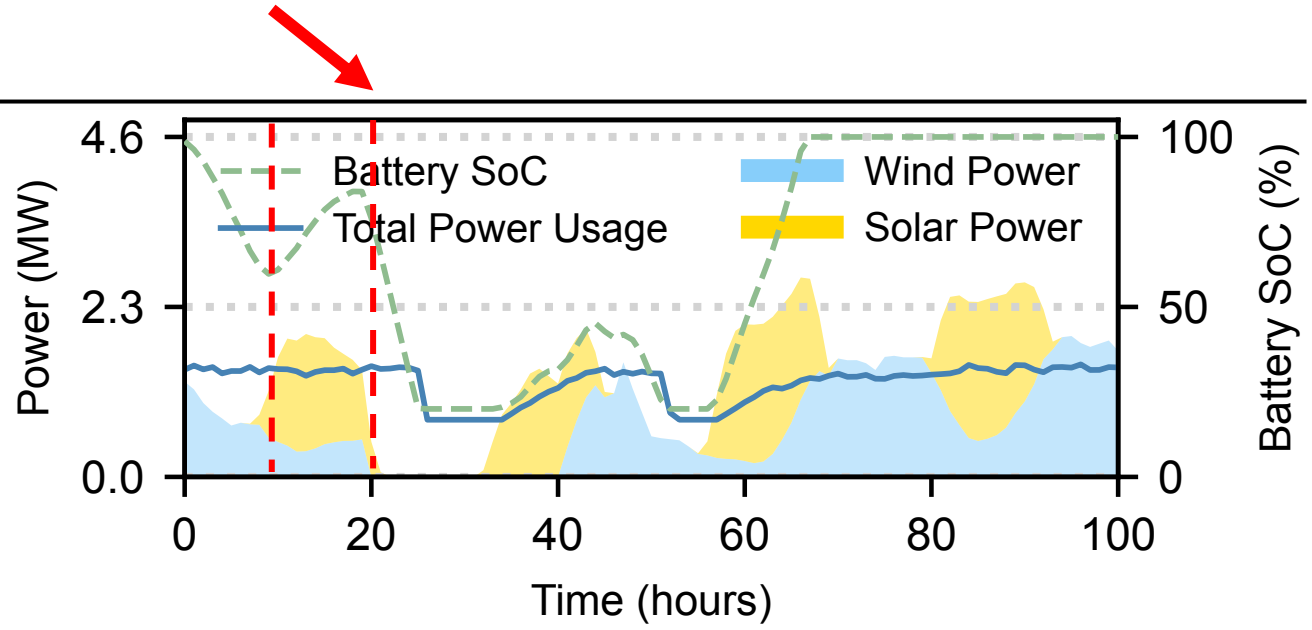

(a) Power Usage and Generation

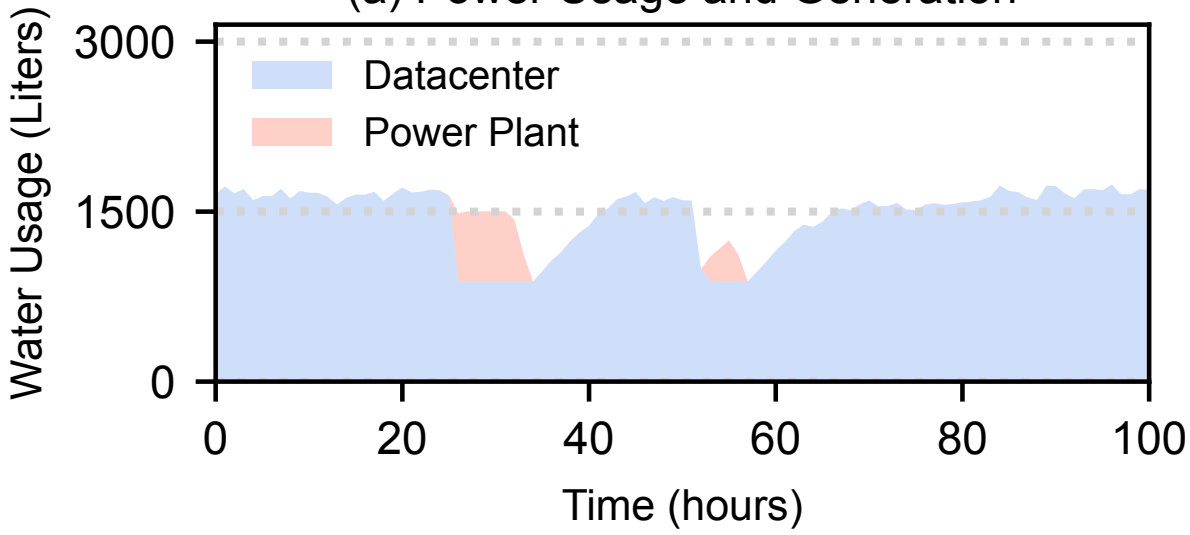

(b) Water Usage (Datacenter and Power Plant)

Simulation results

20-32h

- The wind power plant failure is intentionally created, and the solar power naturally disappeared because it is nighttime

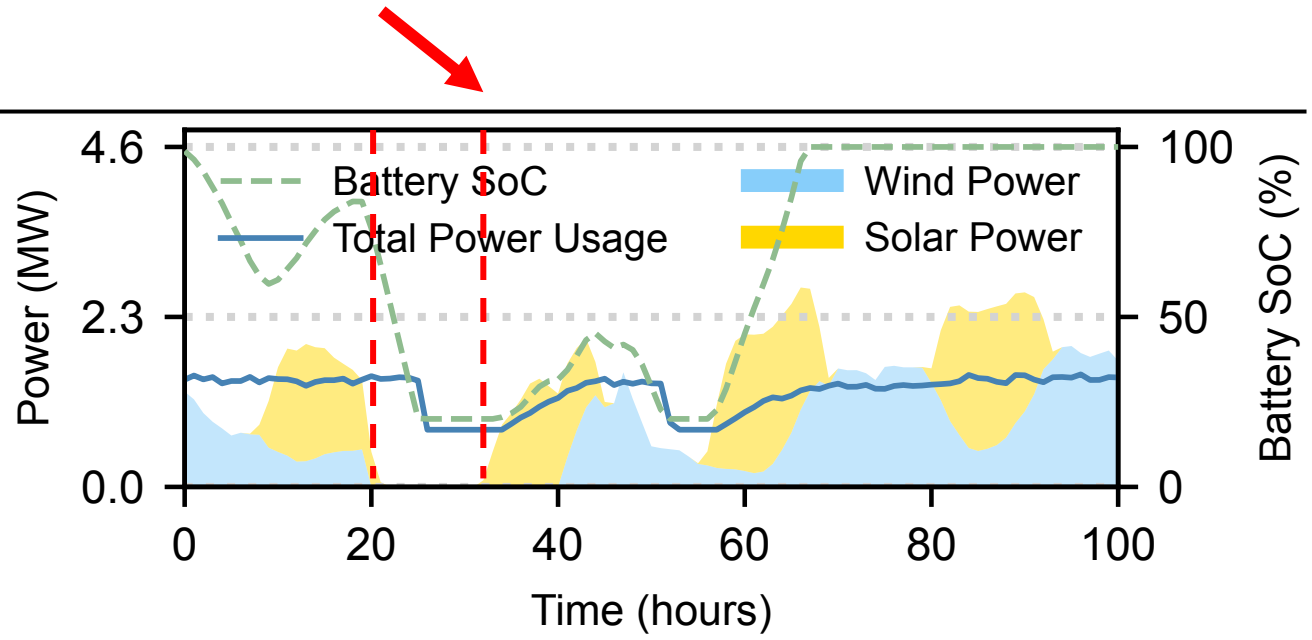

(a) Power Usage and Generation

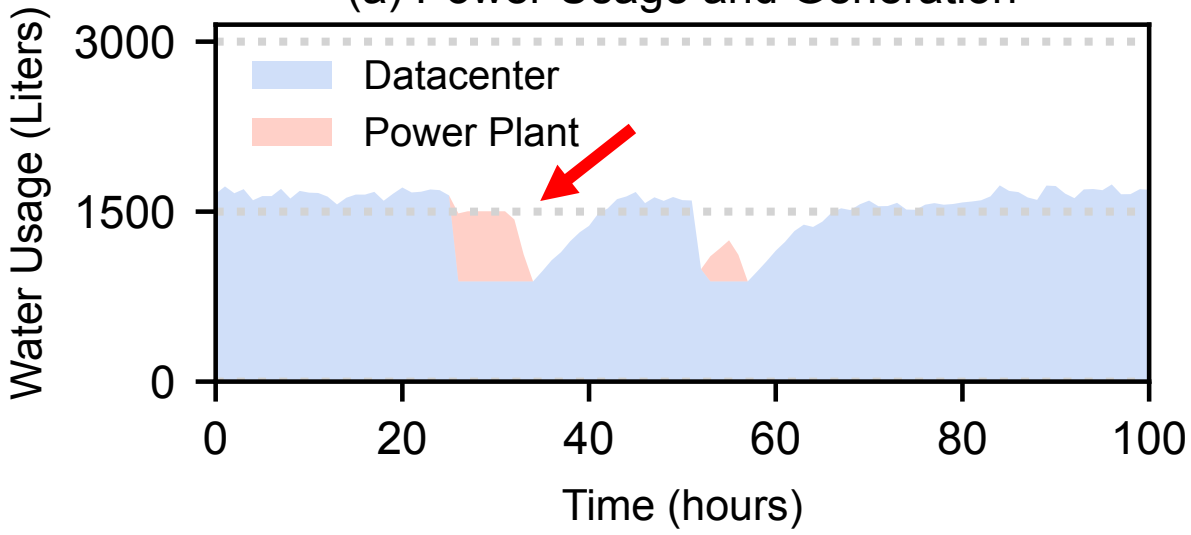

(b) Water Usage (Datacenter and Power Plant)

Simulation results

32-40h

- consumption < renewable supply, solar power recovers, battery gets charged

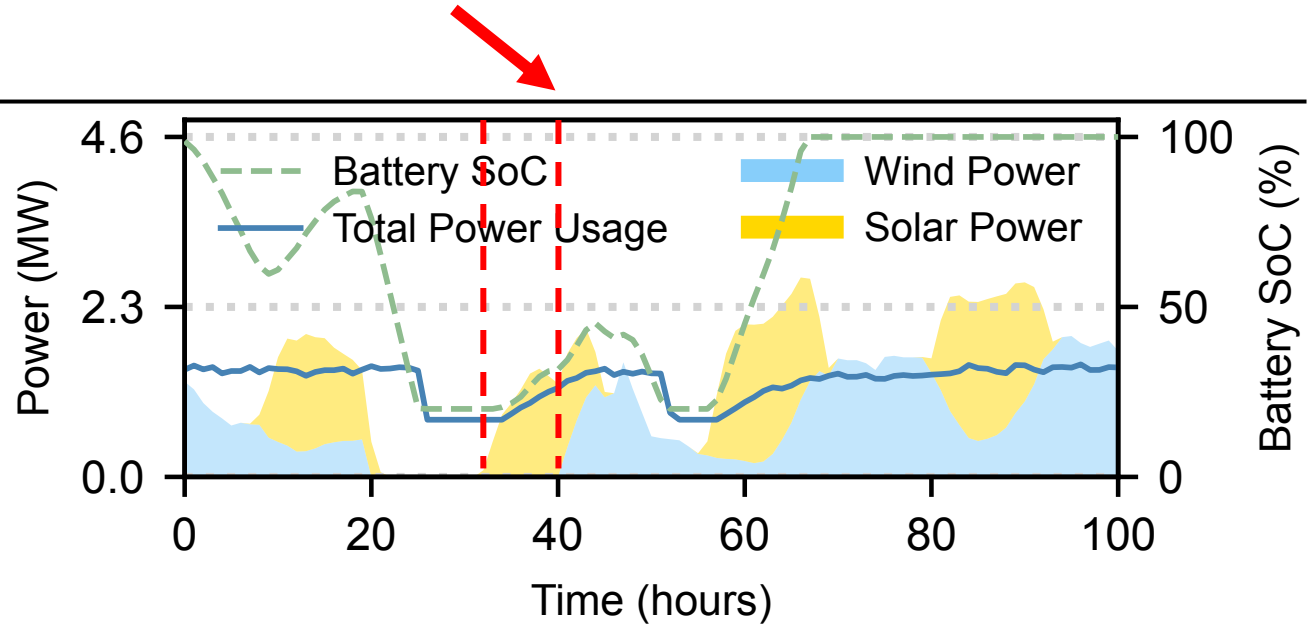

(a) Power Usage and Generation

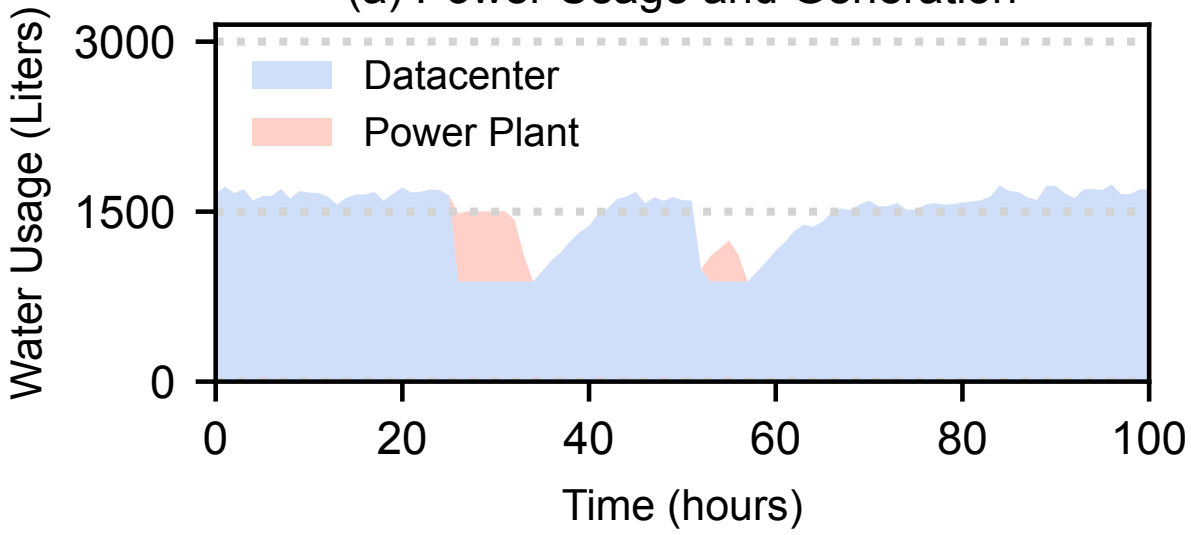

(b) Water Usage (Datacenter and Power Plant)

Simulation results

40h onwards:

- Wind power failure ends

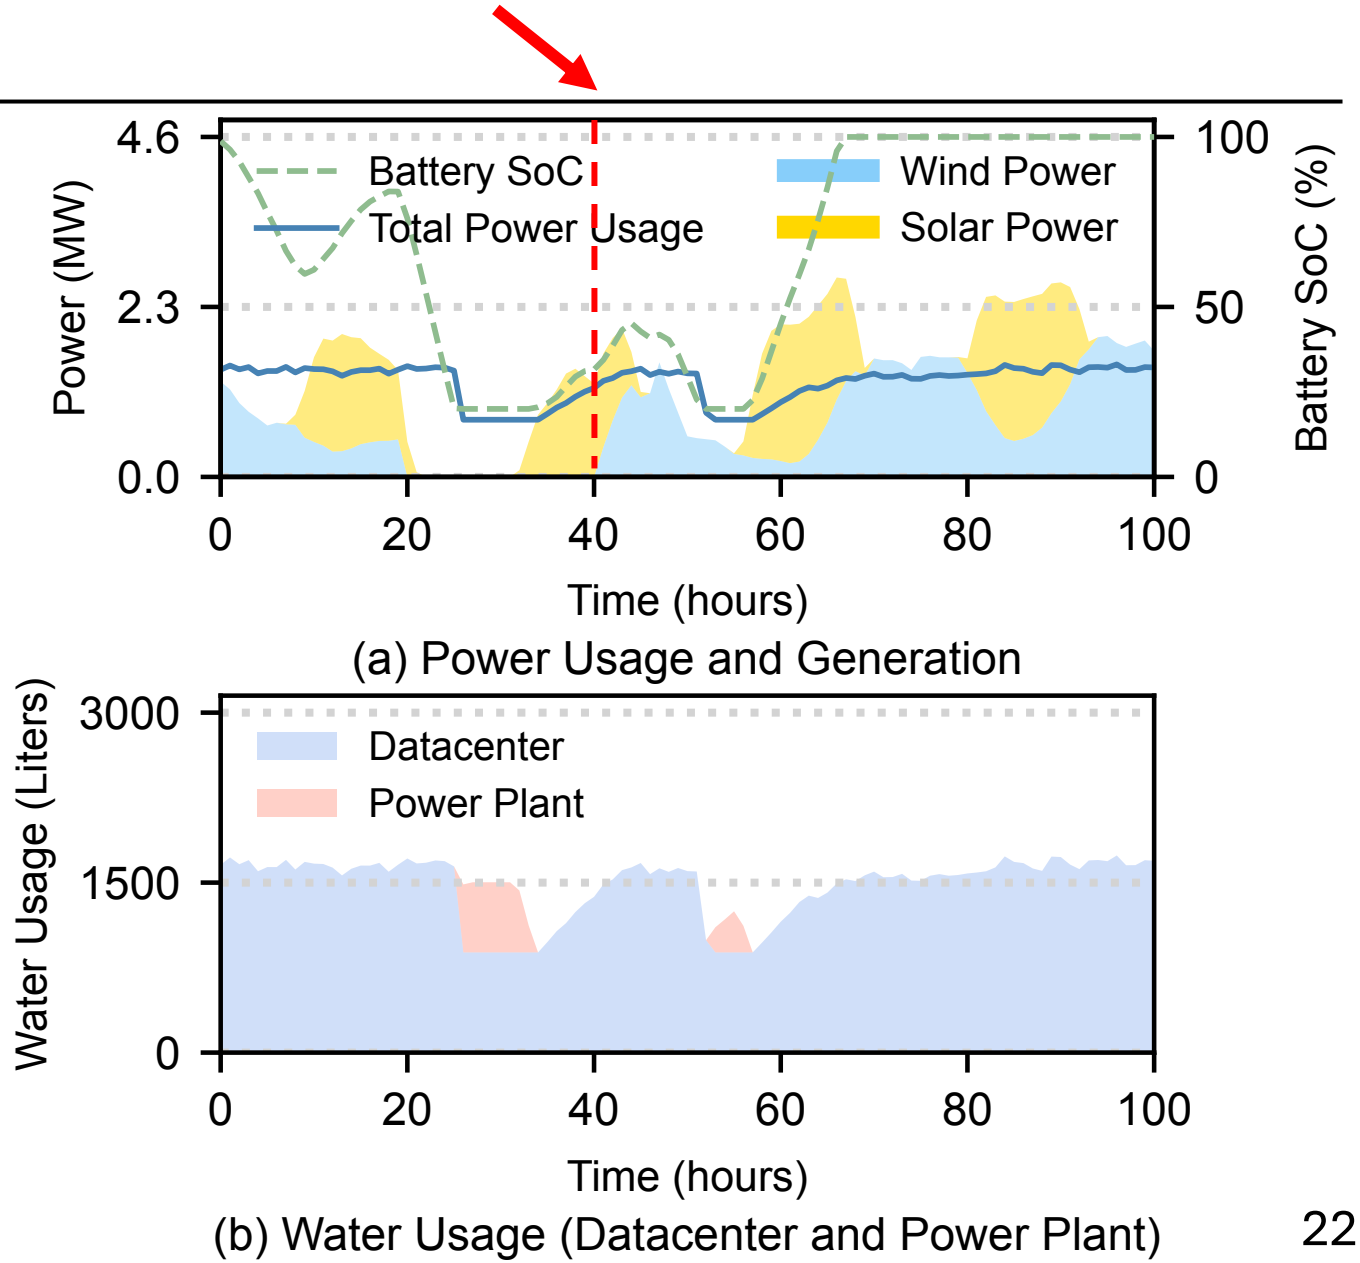

Simulation results

At 50h:

- Power supply drops

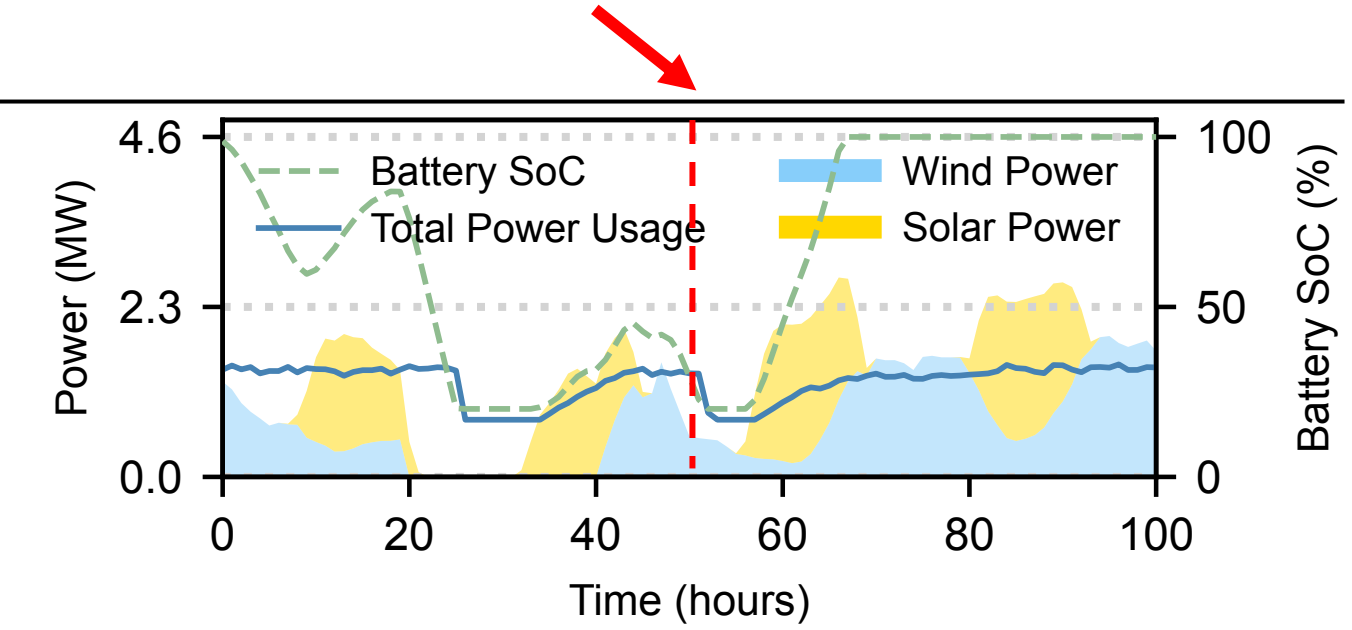

(a) Power Usage and Generation

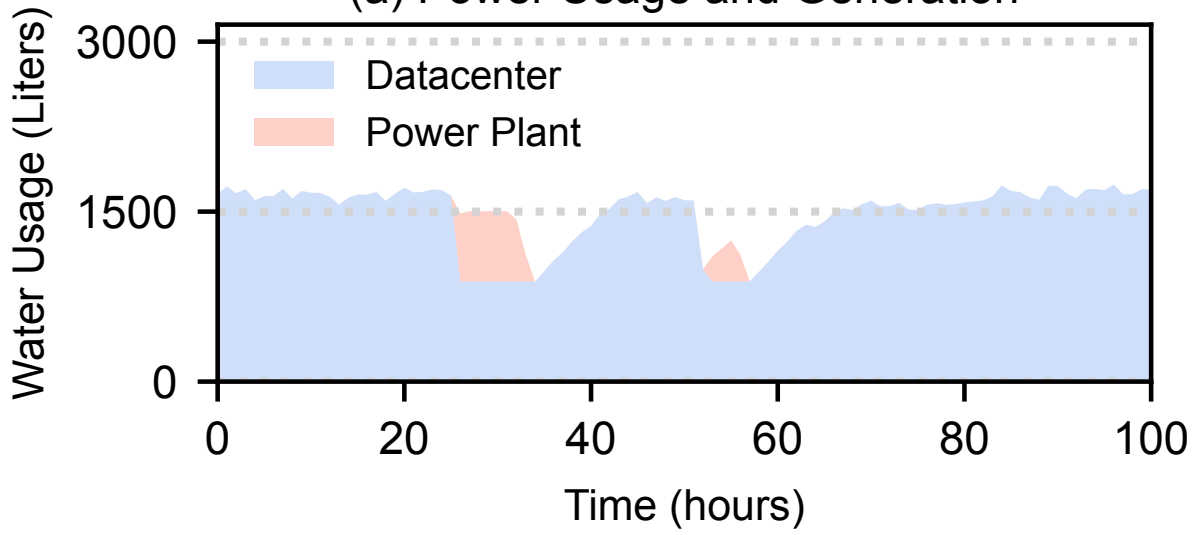

(b) Water Usage (Datacenter and Power Plant)

Simulation results

After 55h:

- Power supply stabilizes

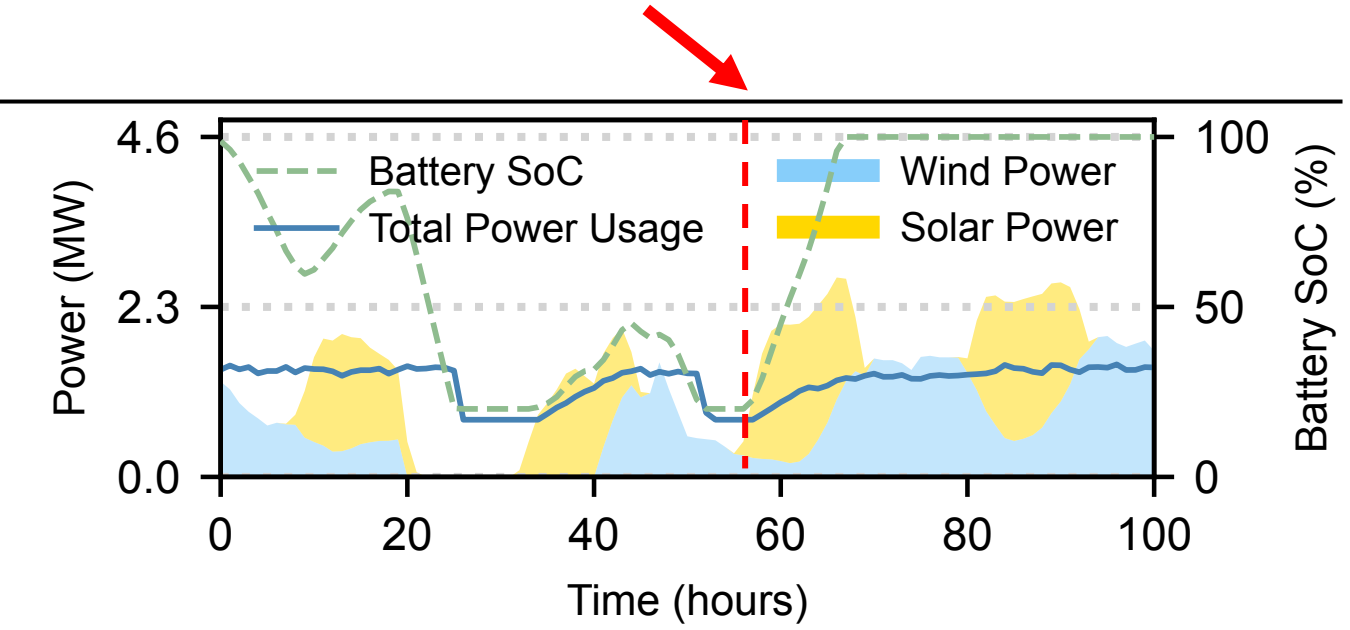

(a) Power Usage and Generation

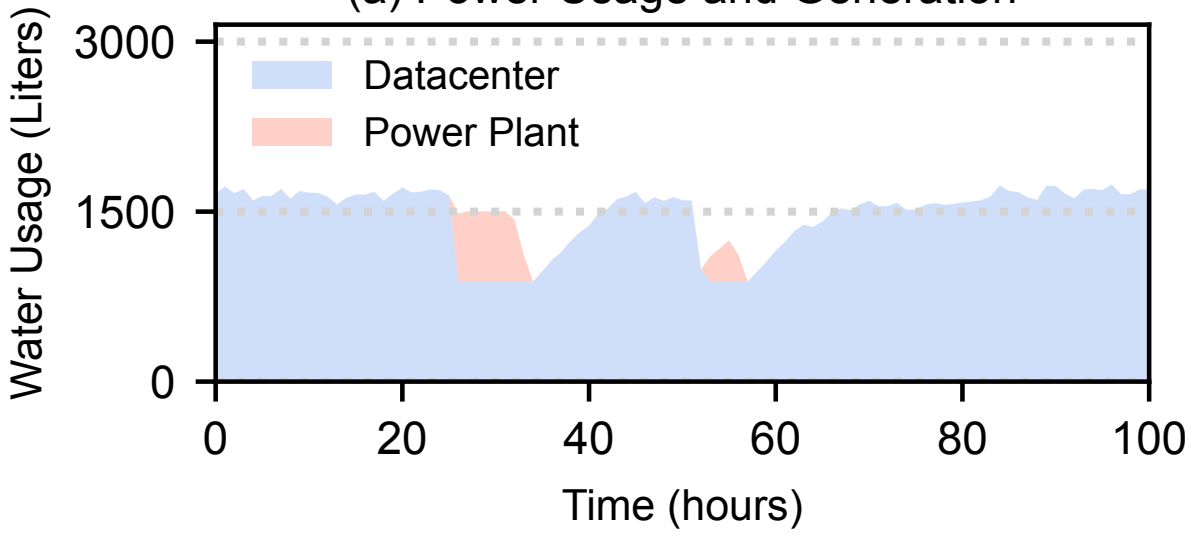

(b) Water Usage (Datacenter and Power Plant)

Conclusion

- What is the Infrastructure/Infrastructure Nexus
- Goals, Challenges, and Architecture of OpenInfra
- Future work:
  - Developing a programming model for the co-simulation framework
  - Exploring containerized simulator management using Kubernetes
  - Enabling zero-modification integration for new simulators
  - Expanding simulation use cases
- I am actively searching for PhD positions for Fall 2025!

Thank you! Feel free to reach out for a chat!

Jiaheng Lu jiahenglu.com