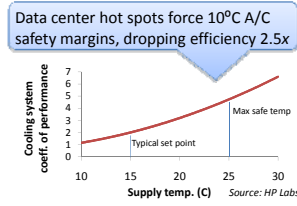
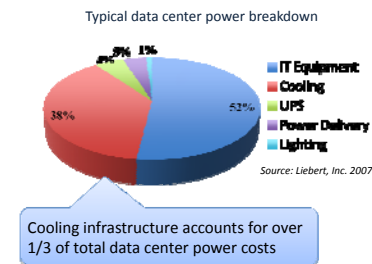
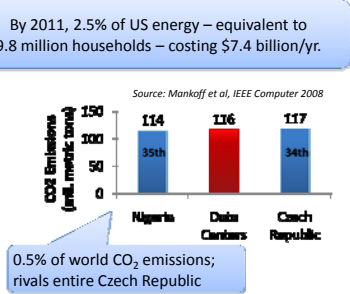
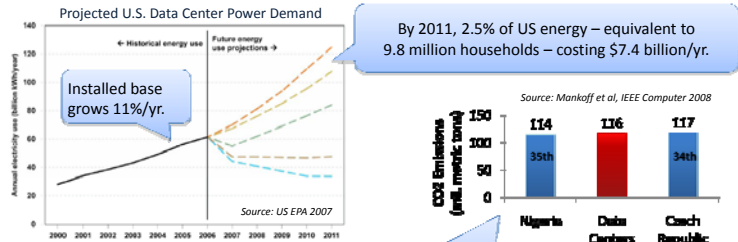


# Maelstrom: Real Time Data Center Temperature Modeling

Steven Pelley, Yifan YangGong, David Meisner, Thomas F. Wenisch (Advanced Computer Architecture Lab, University of Michigan), and James W. VanGilder (APC by Schneider Electric)

## 1. Data Center Cooling Efficiency: A Critical Challenge

A massive fraction of data center power is consumed by cooling infrastructure. System architects can realize enormous energy efficiency gains through thermal-aware mechanisms and cooling-conscious system design. Both offline modeling and real-time control rely on *fast, accurate, and usable* models of data center thermodynamics.



## 2. Our Objectives

- Build a fast, accurate data center thermal modeling kernel based on computational fluid dynamics (CFD)
- Scale CFD solver performance to enable real-time temperature tracking for online thermal management

## 3. Why Computational Fluid Dynamics?

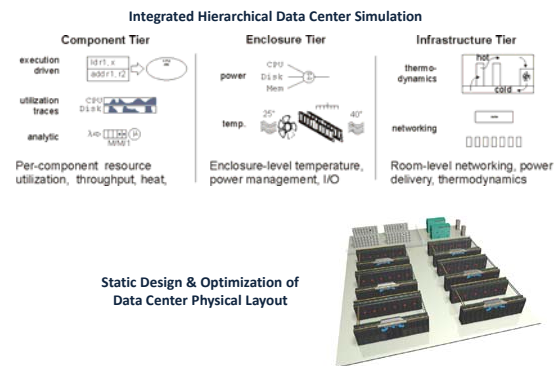
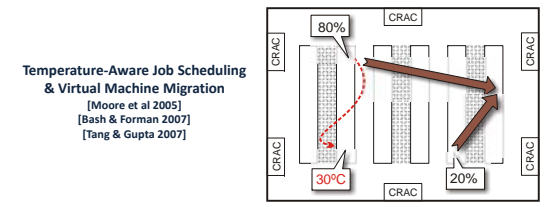
The thermodynamic environment in data centers is chaotic, with many high speed jets and large temperature gradients. These result in markedly non-linear flows. Although linear estimators are fast and simple, they fail to predict temperatures accurately in complex environments. Furthermore, existing linear approximations require either empirical results, which must be updated whenever topology changes, or extensive training data (e.g., derived from CFD simulations). CFD can provide physically-accurate temperature estimates in arbitrarily complex data centers. Maelstrom builds on the open-source *Dolfin* CFD package from TU Delft.

## 4. Related Research

[Moore et al 2006]	Machine-learning based modeling of temperature coupling among servers. Scalability limited by rapid growth in training data requirements.
[Tang et al 2006]	Abstract heat flow modeling of data centers. Drastically lower computational requirements than CFD, but relies on fixed topology and stable airflow patterns.
[Choi et al 2007]	CFD modeling within system and rack enclosures. We apply similar approaches to model entire data center environments.
[Ramos & Bianchini 2008]	Energy management based on predictive thermal modeling. Relies on a linear model similar to flow networks.
[VanGilder & Zhang 2008]	Coarse-grid methods to accelerate data center CFD.

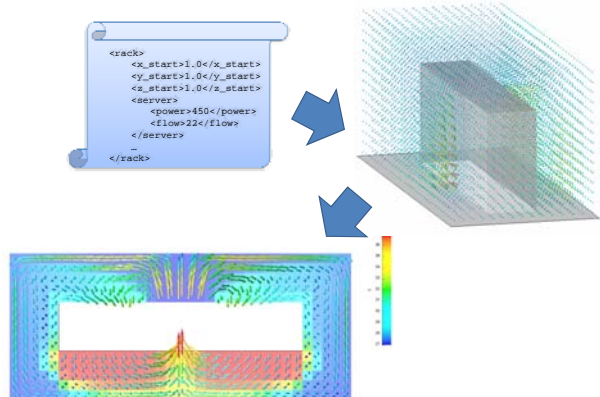
## 5. Applications

Maelstrom can act as the computational core for offline modeling tools, to enable research on next generation data center hardware, software, and physical infrastructure, and real-time online models, to complement sensor data with predictions for "what-if" scenarios.



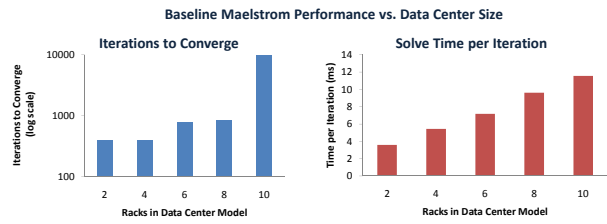
## 6. Improving Usability

Conventional CFD tools require substantial mechanical engineering expertise to operate. To make CFD-based data center modeling accessible to the systems architecture community, Maelstrom accepts an XML specification that describes a data center at the granularity of racks, clusters and servers. Maelstrom then automatically builds an appropriate mesh for its CFD solver.



## 7. Improving Speed

Data center CFD modeling can be substantially accelerated through domain-specific CFD solver customization. For example, omitting turbulence modeling can substantially reduce computational requirements. Though turbulence modeling is essential in general CFD applications, it tends not to affect server inlet temperatures, typically the only temperatures of interest. Furthermore, CFD is highly parallelizable, and many conventional optimization techniques can provide substantial benefits, but are ignored in existing solvers.



- |  |   |
|--|---|
| <b>Techniques to reduce iterations</b> <ul style="list-style-type: none"> <li>• coarse grid</li> <li>• multigrid</li> <li>• block-structured mesh</li> <li>• initial value prediction</li> </ul> | <b>Techniques to accelerate solve time</b> <ul style="list-style-type: none"> <li>• improved linear solvers</li> <li>• domain decomposition</li> <li>• block-structured mesh</li> <li>• solving pressure and temperature in parallel</li> </ul> |
|--|---|

## 8. Maintaining Accuracy

Existing real-time data center thermal models assume linear relationships among critical flow variables. Though computationally expedient and accurate for constrained, low-velocity flows, such models omit key physical phenomena, such as the non-linear pressure-airflow relationship, buoyancy, and momentum effects, which are critical in data centers with multiple, interacting high-speed jets. In contrast, Maelstrom seeks to accelerate conventional CFD, omitting only physical phenomena shown to have negligible impact in data centers.

Accuracy impact of omitting turbulence from CFD solver (FLOVENT) For the Michigan Academic Computer Center (~250 racks)

Average server inlet temperature prediction error	0.25°C
Worst-case server inlet temperature prediction error	2.63°C
Fraction of inlets with error > 1.5°C	4.7%