Unified Thermal and Power Management in Server Enclosures

Niraj Tolia*, Zhikui Wang, Manish Marwah, Cullen Bash, Parthasarathy Ranganathan, Xiaoyun Zhu

*Now at VMware


2005: $26.5 Billion

"The cost of power and cooling is likely to exceed that of hardware…"

- Luiz Barosso, Google

"In the data center, power and cooling costs more than the IT equipment it supports."

- Christian L. Belady, Microsoft
(former HP)

Zephyr

- Unified Power and Cooling Management

Contributions

- Demonstrate how to do per-blade cooling
- Save both cooling (21-30%) + system power (23-29%)
- Without impacting performance

Blade Servers:

Compact Design
Dense Compute and Storage
Workloads are Virtualized

Blade Servers:

Compact Design
Dense Compute and Storage
Workloads are Virtualized

Virtualization:

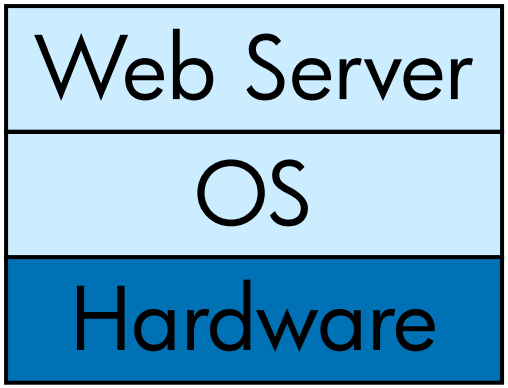

Virtualization:

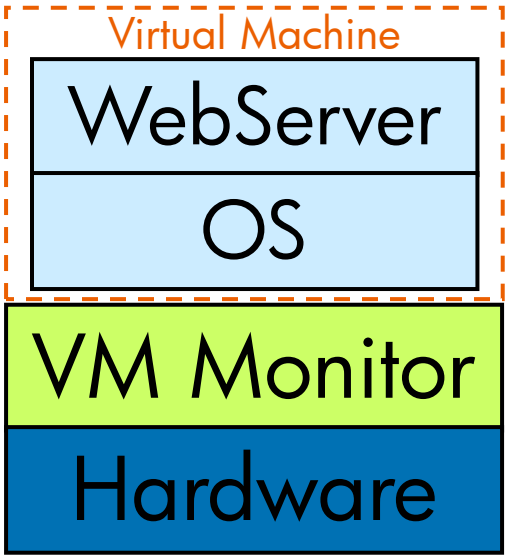

Virtualization:

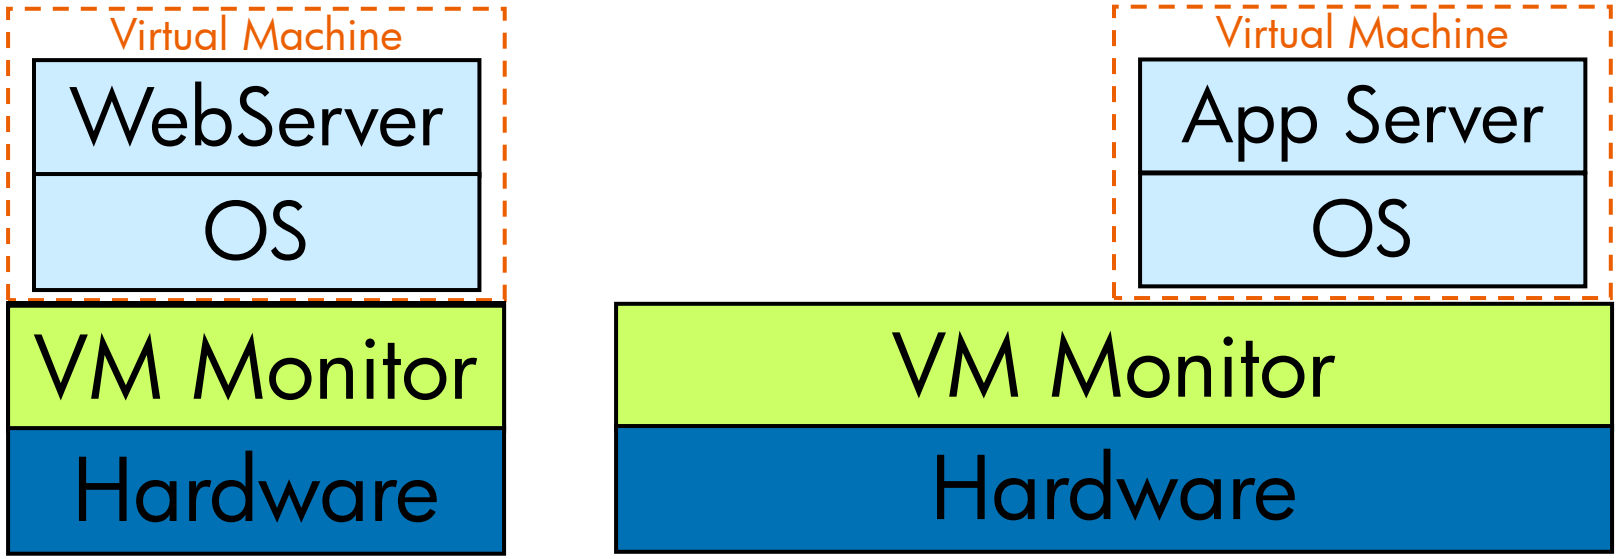

Virtualization:

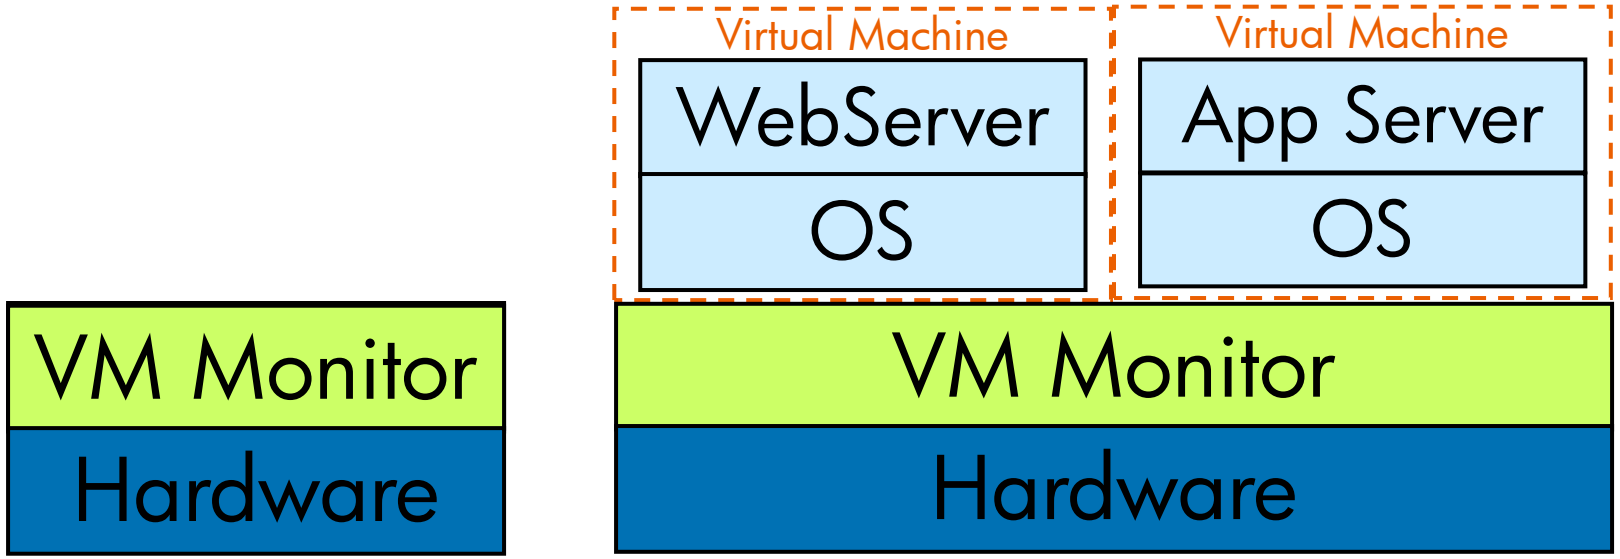

Virtualization:

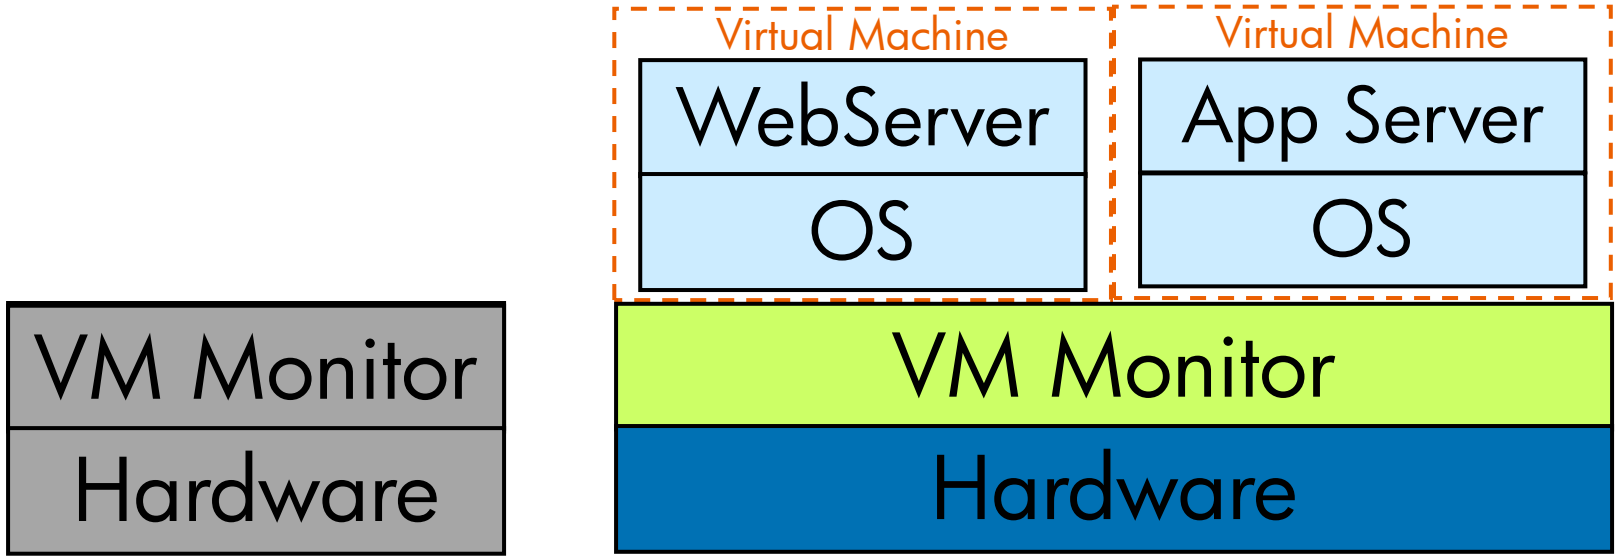

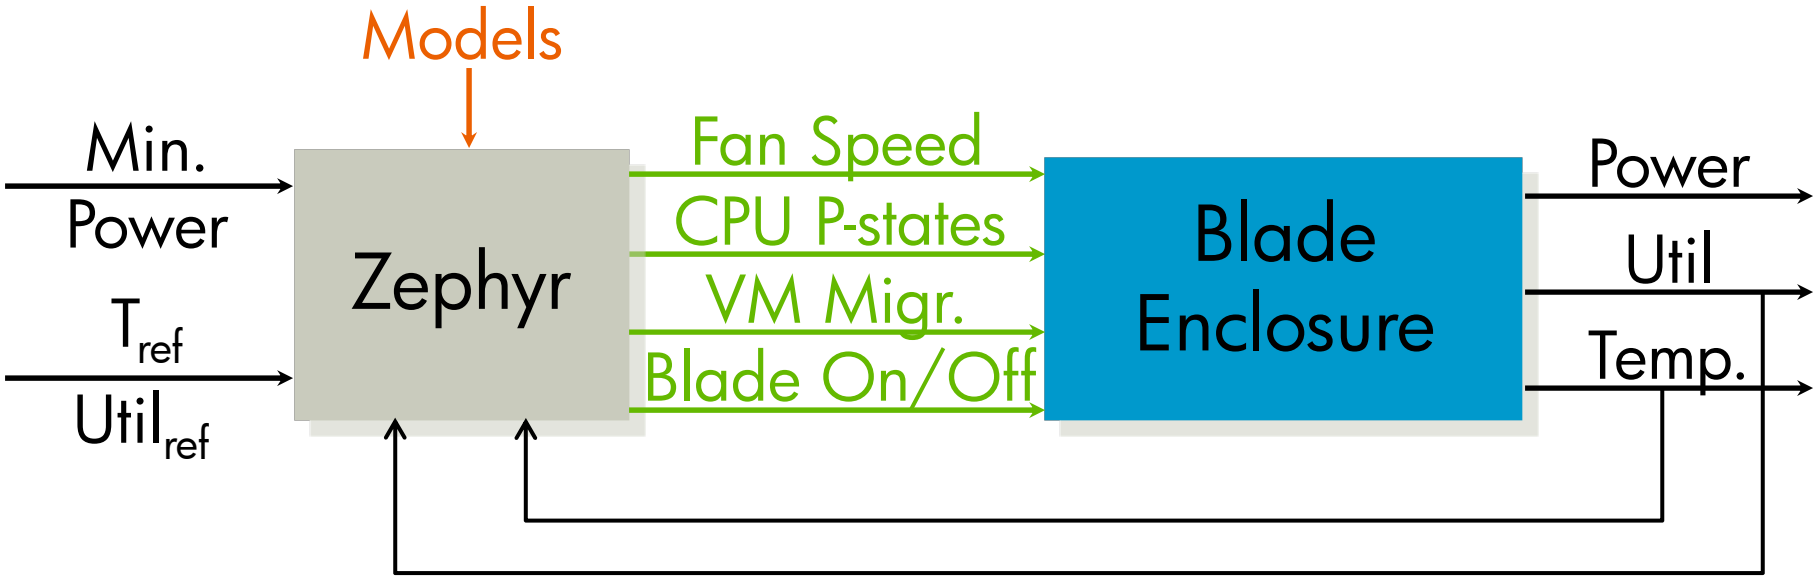

Power Models
Temperature Models

Power Model: Single Blade

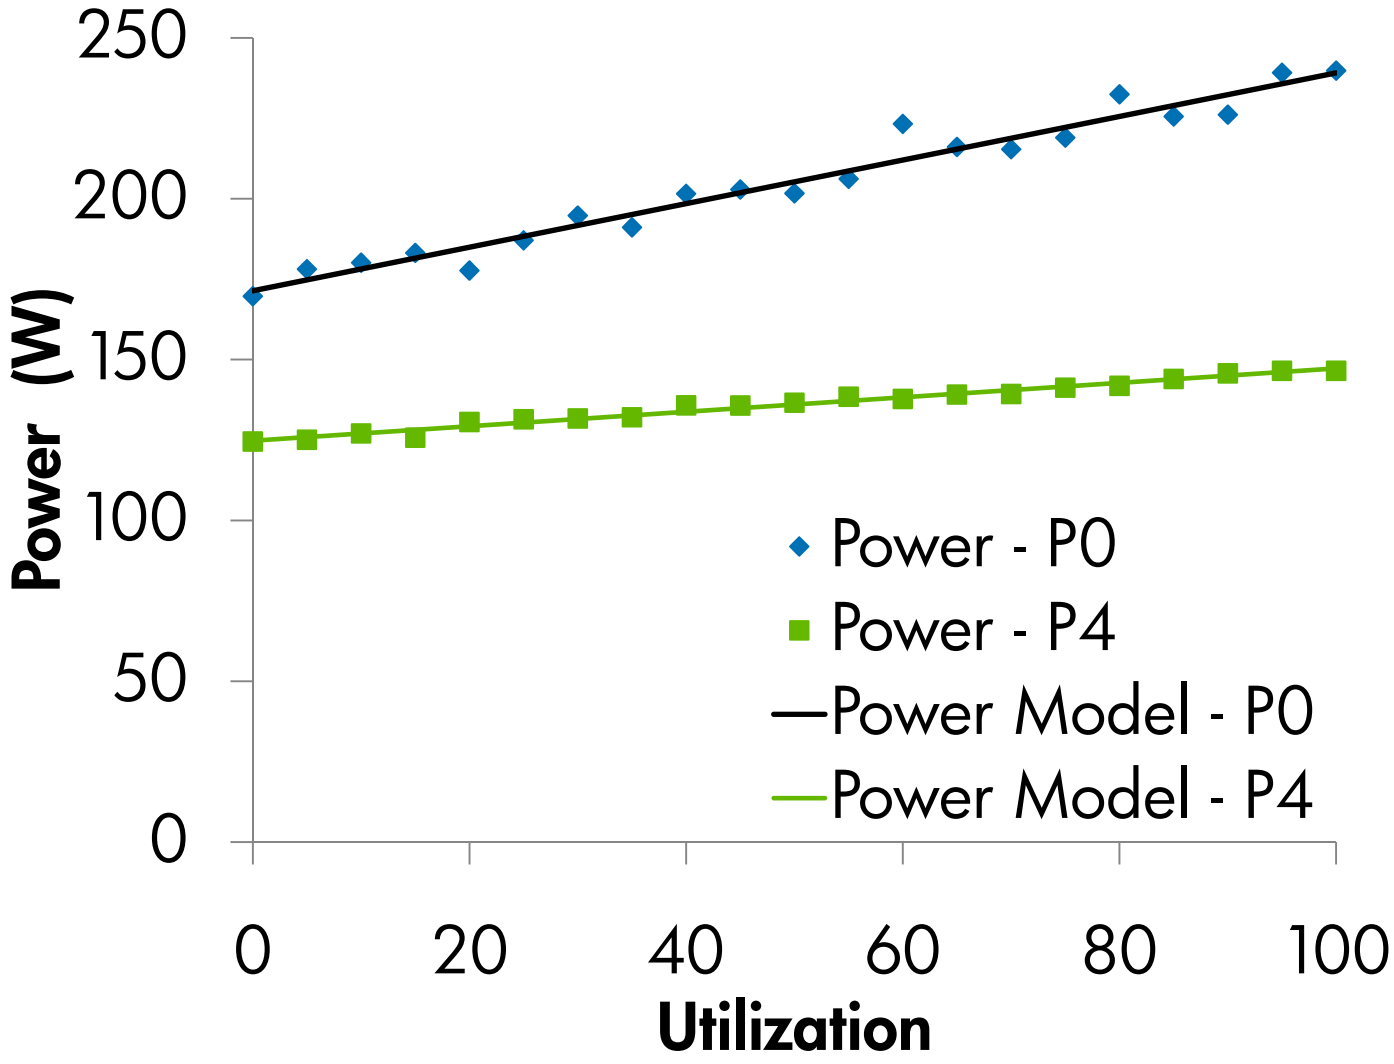

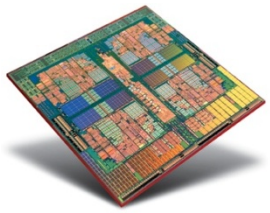

Power Model: Single Fan

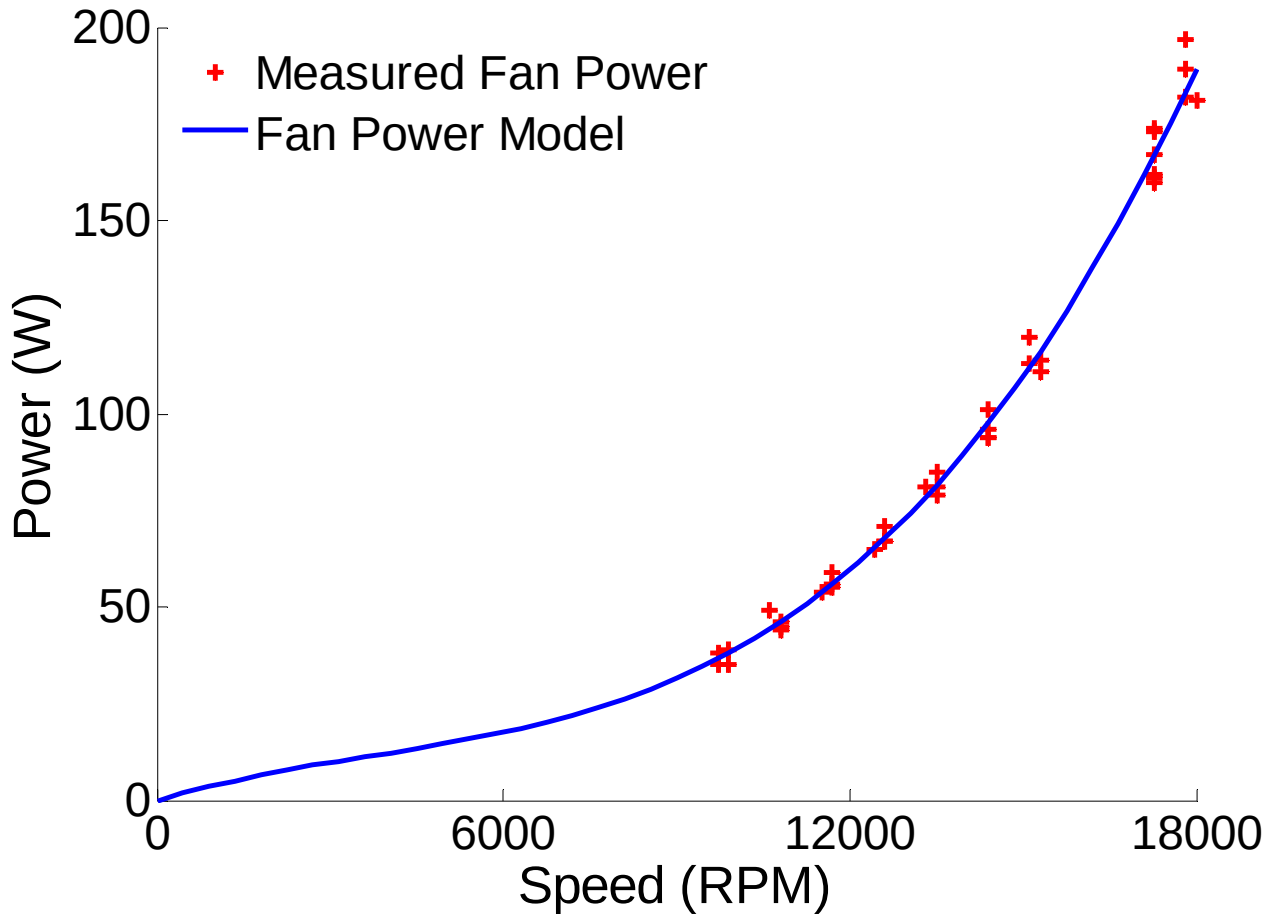

Unfortunately, temperature models are significantly more complex.

Models don't exists
Many things affect temperature
Insufficient sensors in servers
Zonal Variations

Zonal Variation: Fan Influence

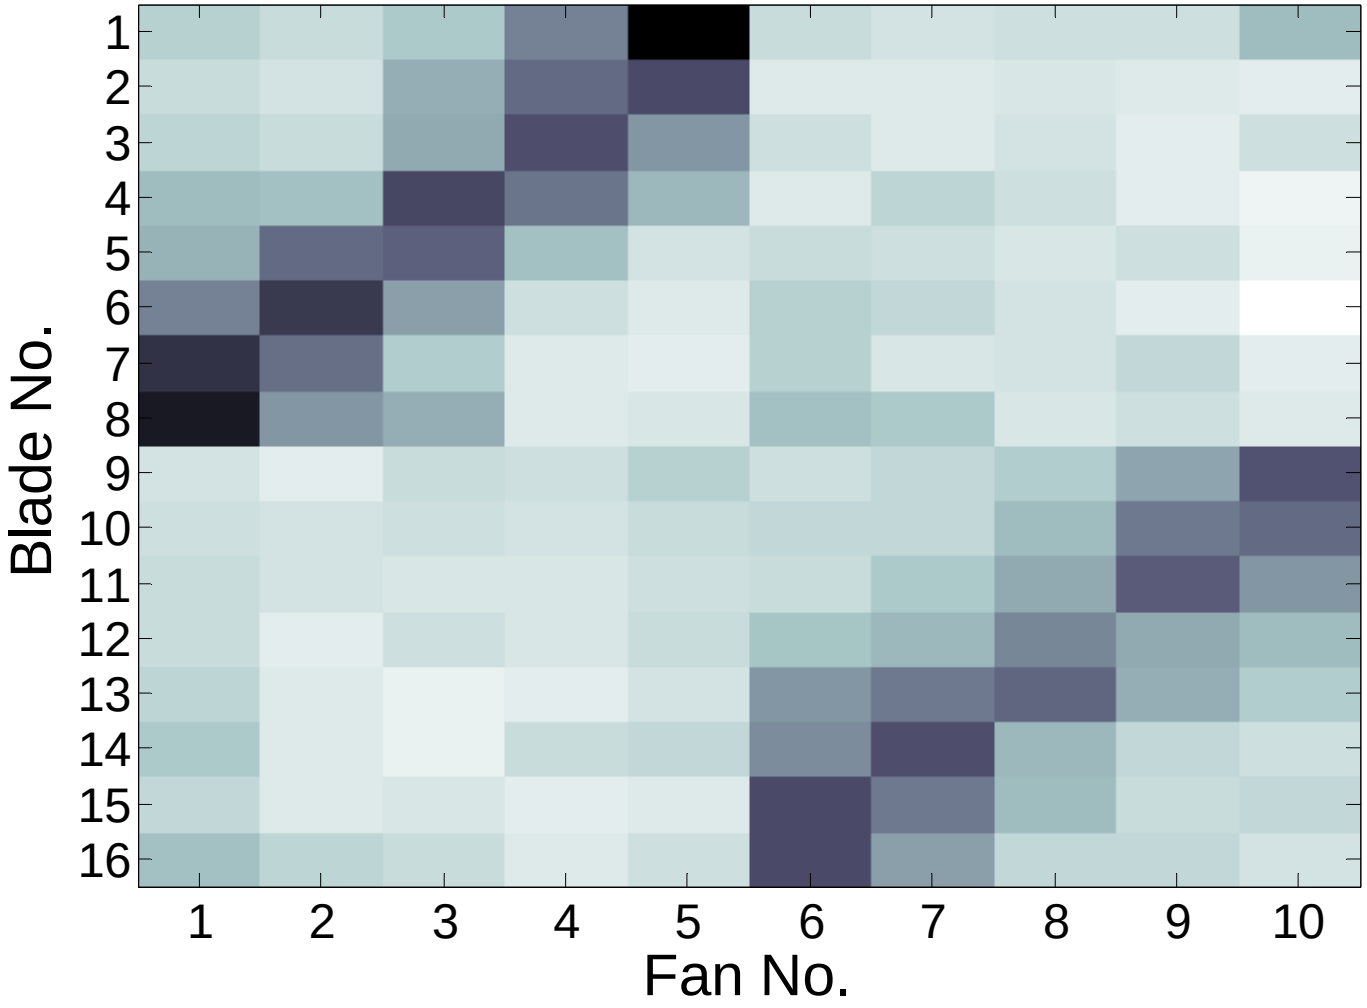

Temperature Models

- Steady State Model
  - Assumes workloads, ambient environment is constant
- Transient Model
  - Finer-grained control in real-world environments

Temperature Models

- Steady State Model
  - Assumes workloads, ambient environment is constant
- Transient Model
  - Finer-grained control in real-world environments

More details in Wang et al.

Data Center Thermal Management IV (Tomorrow)

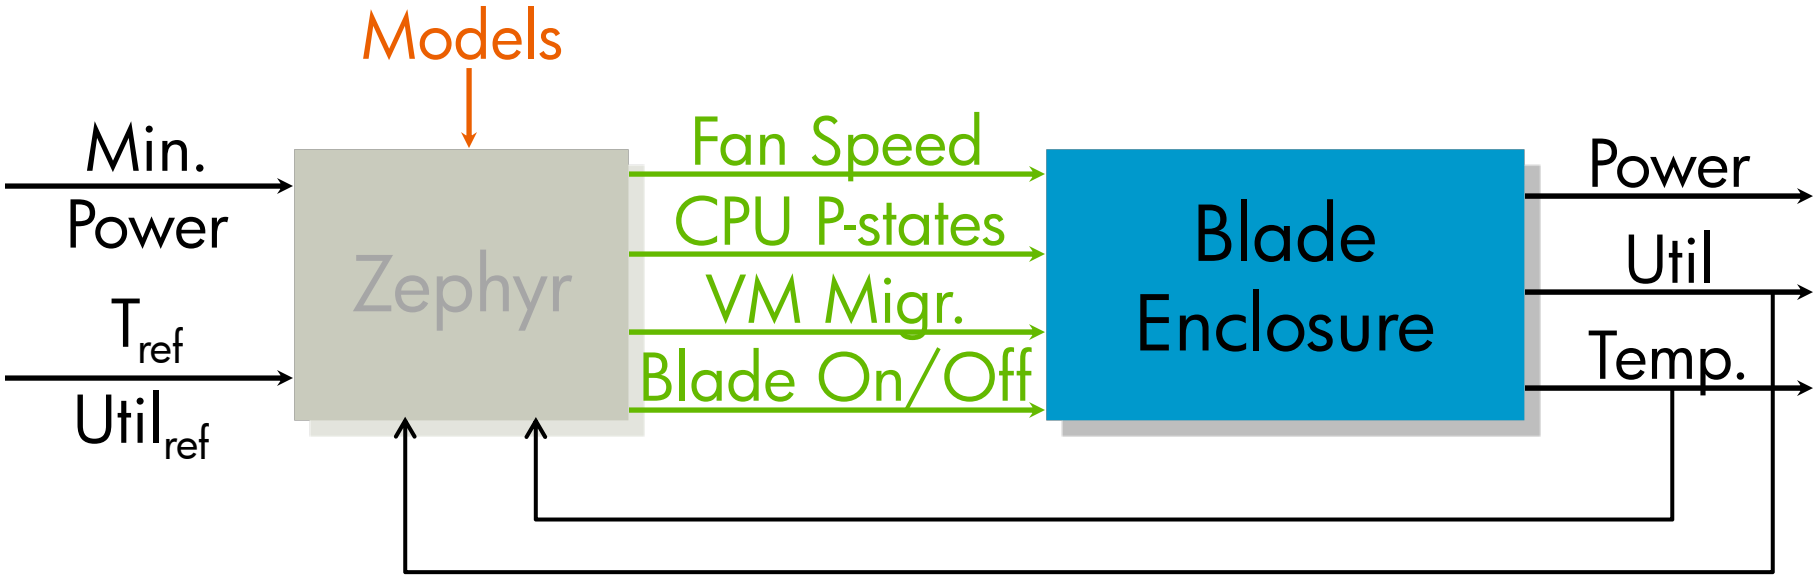
Models
Min.
Power
T_{ref}
$Util_{ref}$
Zephyr
Fan Speed
CPU P-states
VM Migr.
Blade On/Off
Blade
Enclosure
Power
Util
Temp.

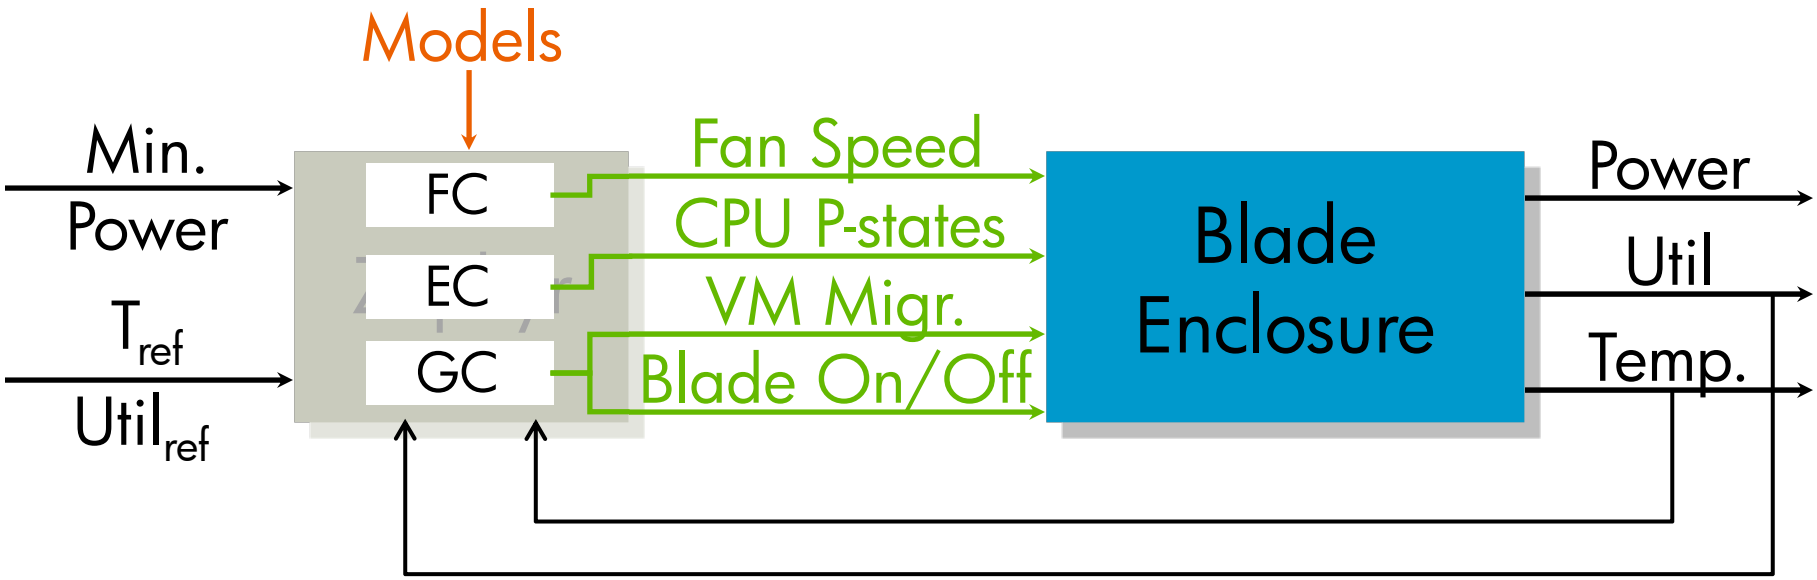
Models
Min.
Power
T_{ref}
$Util_{ref}$
FC
EC
GC
Fan Speed
CPU P-states
VM Migr.
Blade On/Off
Blade Enclosure
Power
Util
Temp.

Group Controller (GC)

- Goal: Minimize Power(Enclosure)

- Changes state every 10 minutes
- Monitors Utilization, Inlet Temperature
- Uses two-step optimization
  - Simulated Annealing for space search
  - Convex Optimization for Fan Cost + Blade Power Model

Experimental Setup

16 Blades (4 cores)
5 P-States
10 Fans, 64 Xen VMs

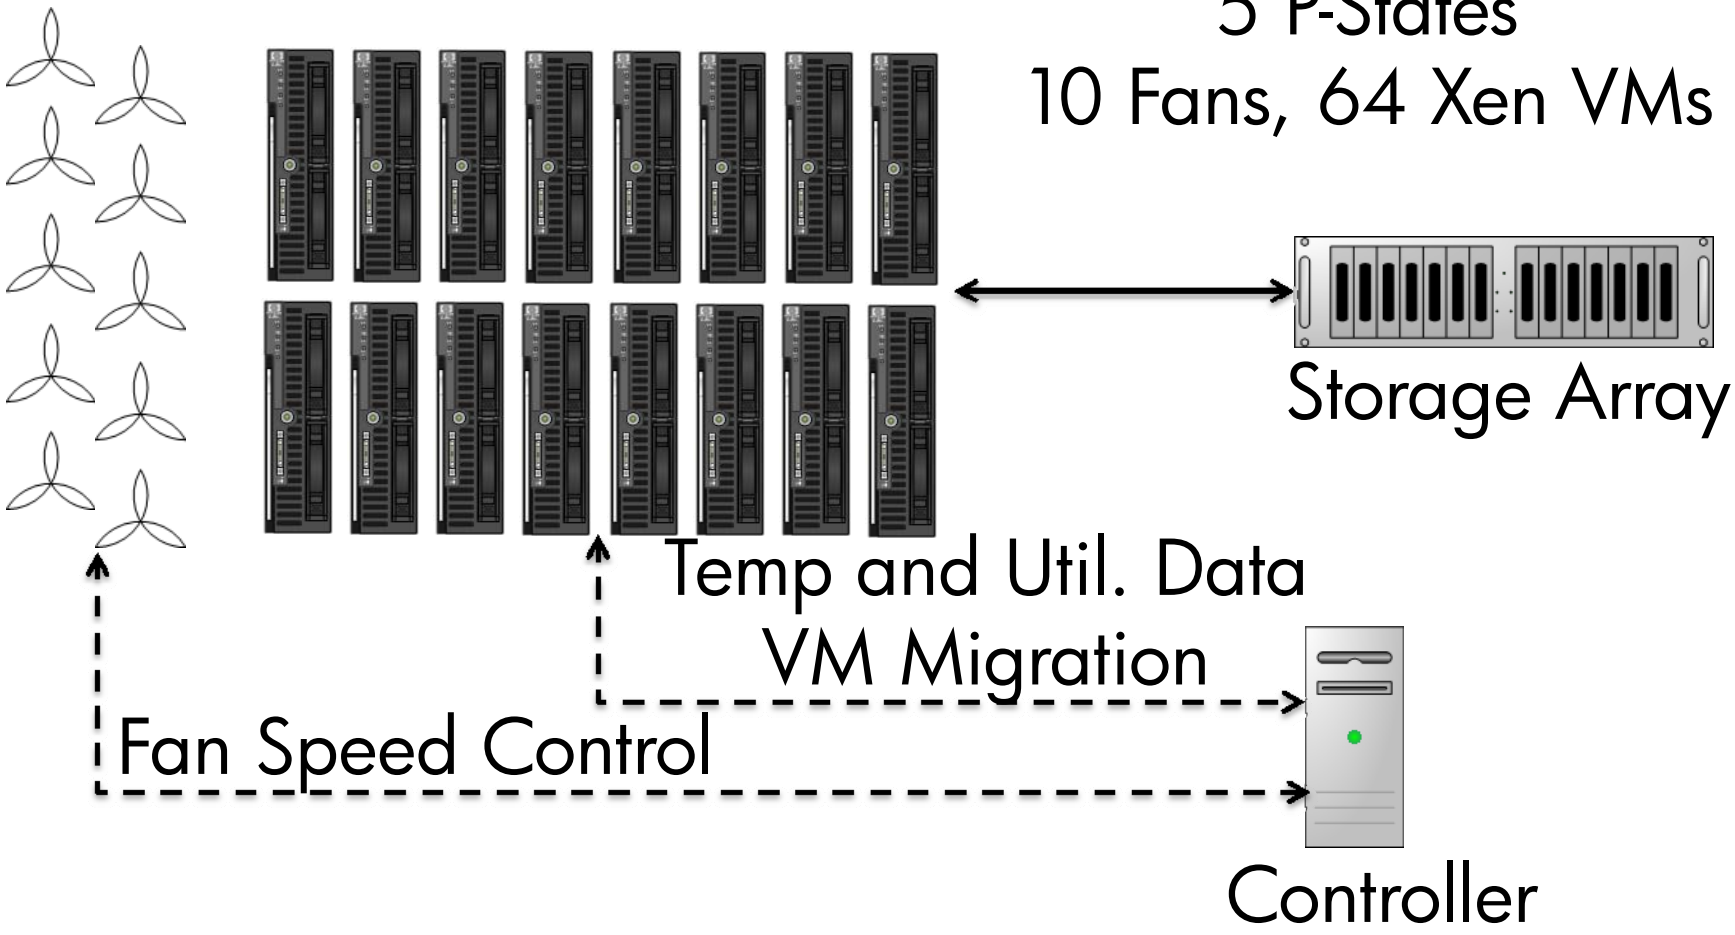

Cooling Power Savings

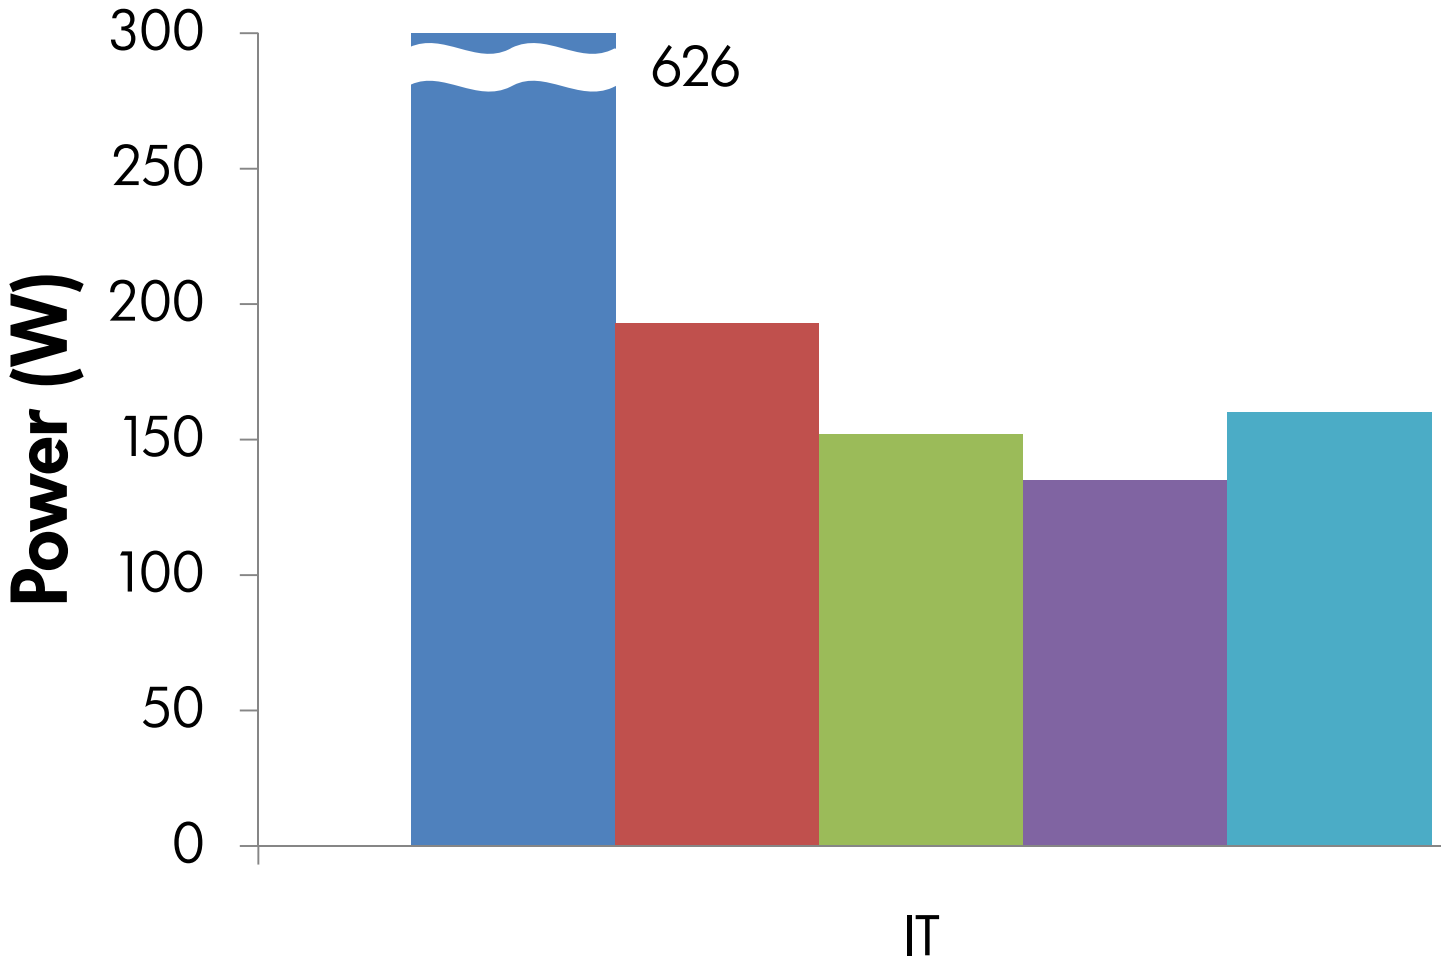

System Power Savings

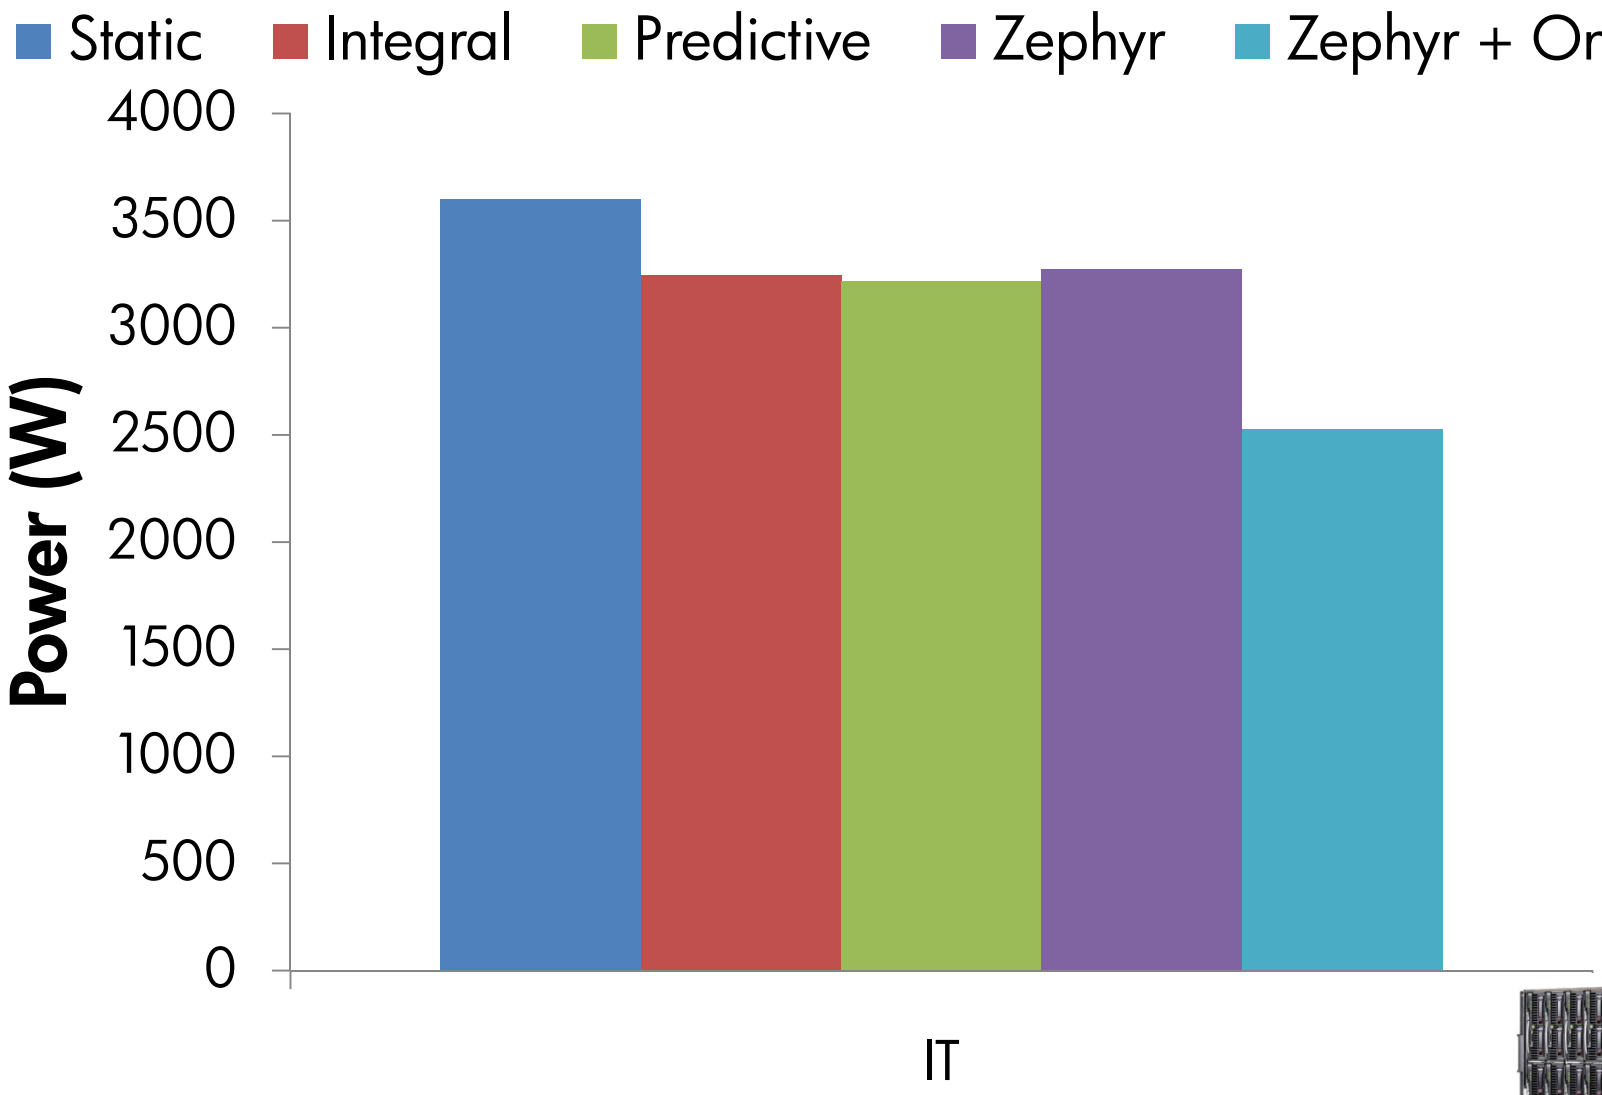

Simulated Servers

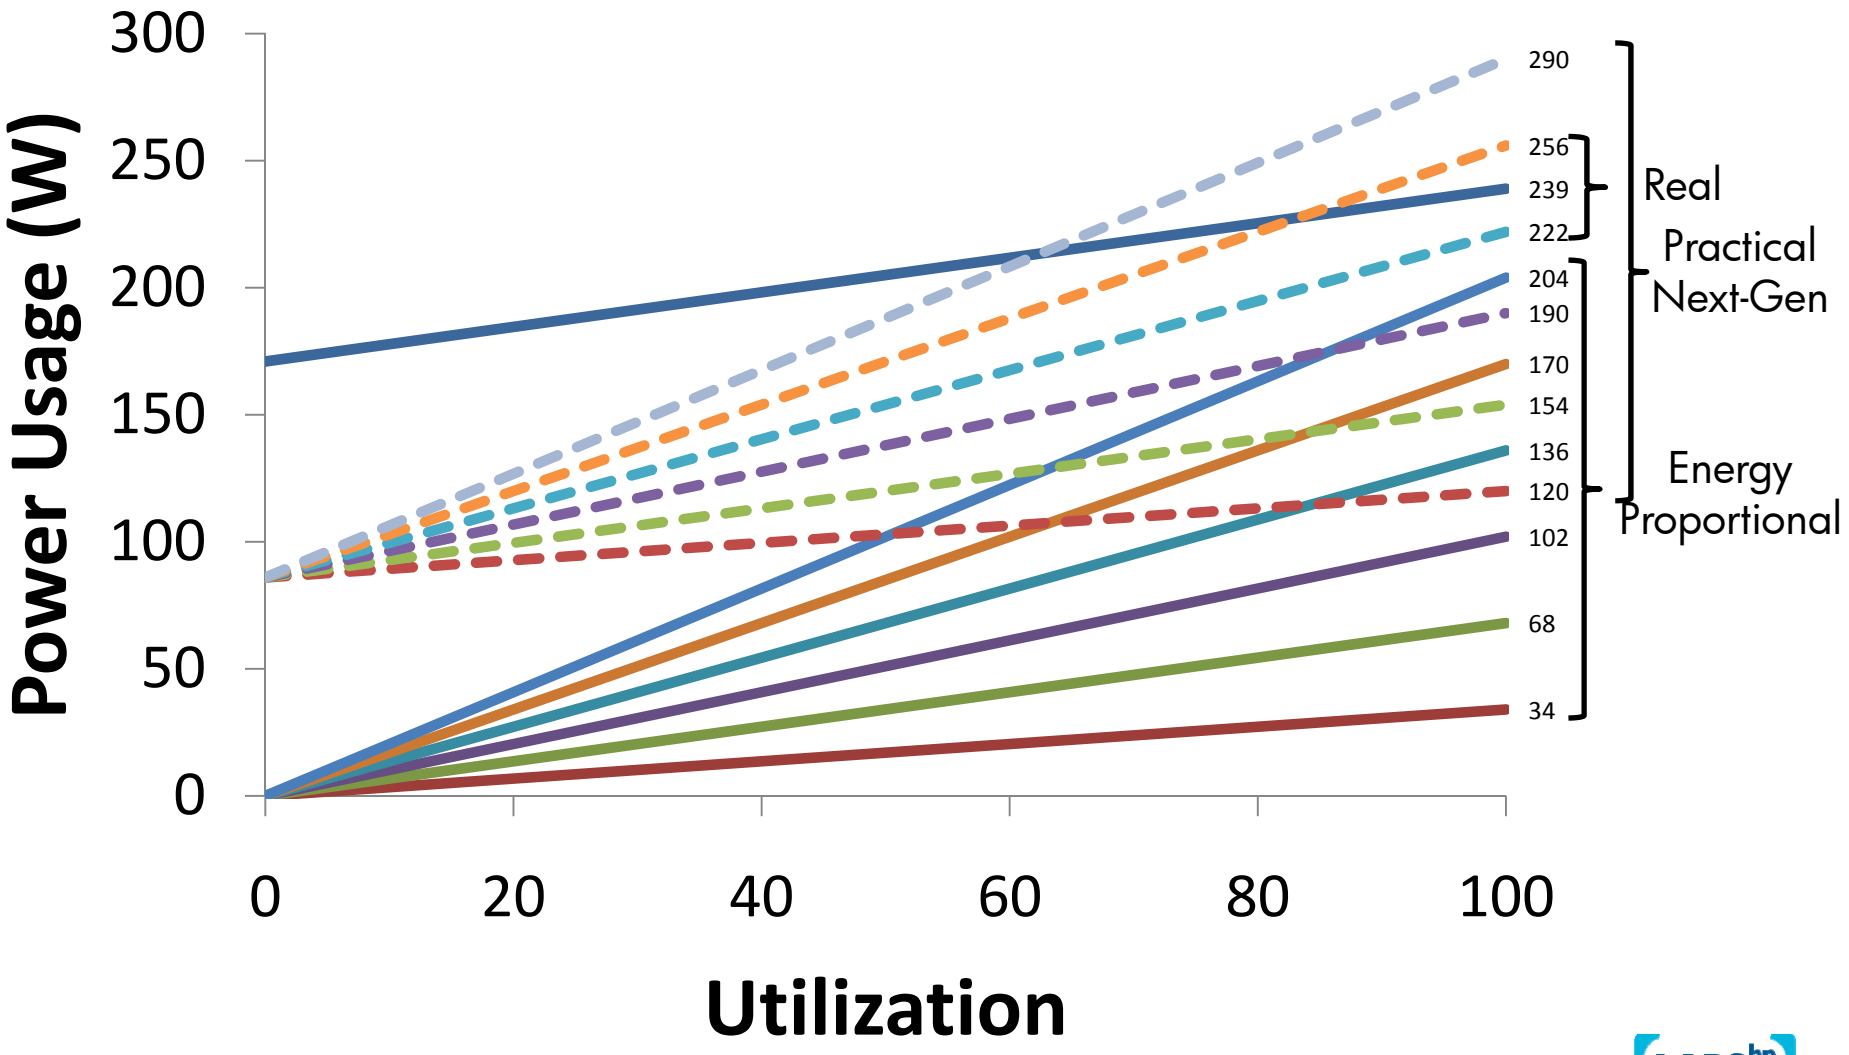

Simulation: T_{amb}+15C &13 servers types

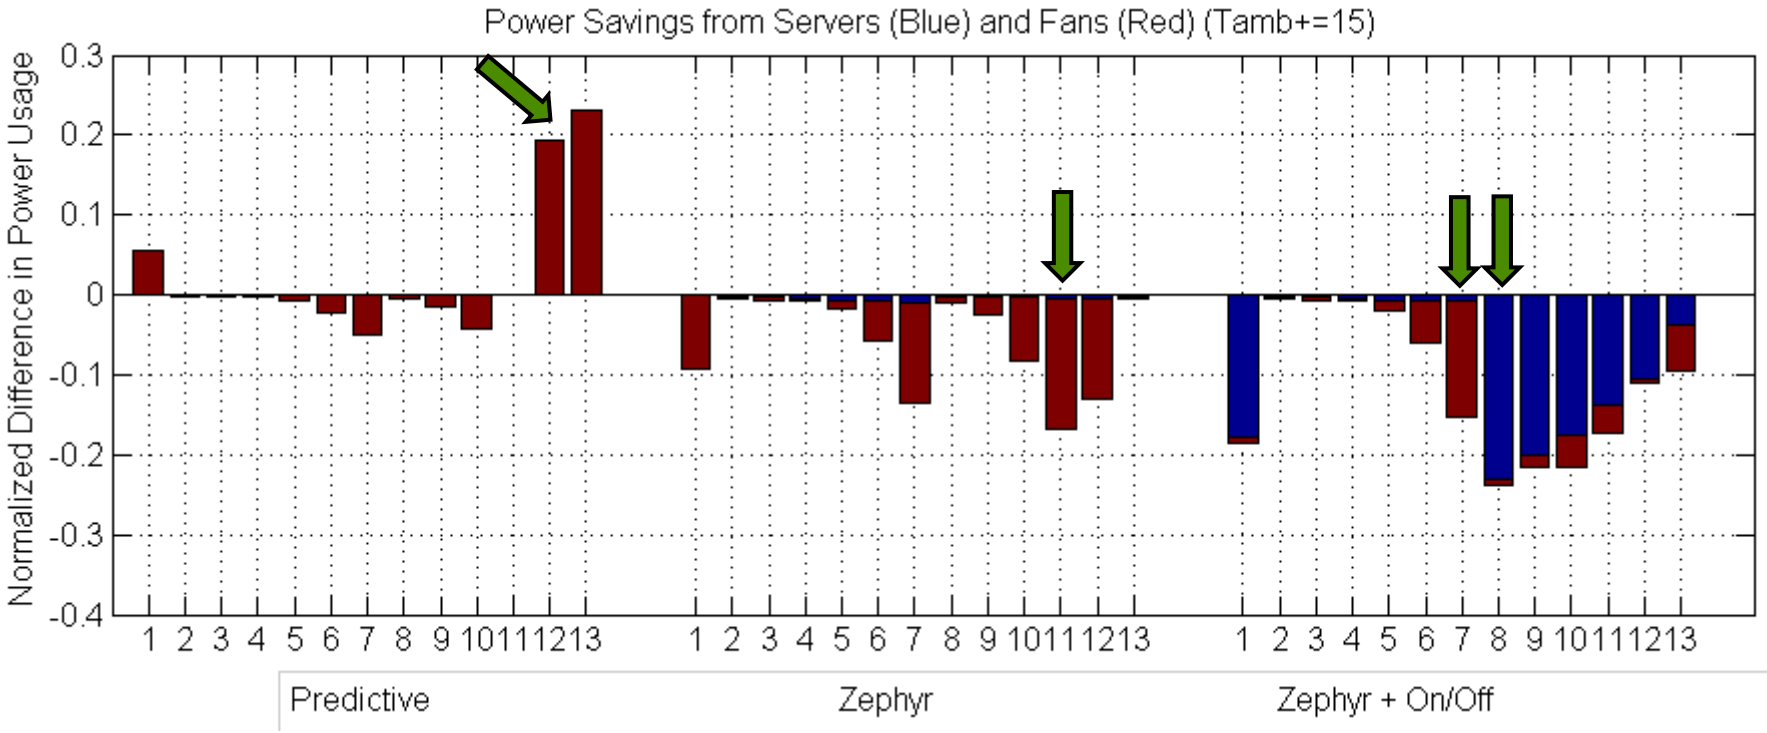

Conclusions

- Zephyr combines:
  - Concepts from Heat Transfer Theory with Systems

- Unified power and cooling management can save:
  - 30% Cooling Power
  - 29% Enclosure Power

Full Length Papers:
www.hpl.hp.com/personal/Niraj_Tolia/

LABS hp