FAWNdamentally Power-efficient Clusters

Vijay Vasudevan, Jason Franklin, David Andersen,
Amar Phanishayee, Lawrence Tan, Michael Kaminsky*, Iulian Moraru

Carnegie Mellon University, *Intel Research Pittsburgh

Monthly energy statement considered harmful

- Power is a limiting factor in computing
- 3-year TCO soon to be dominated by power cost [EPA 2007]
- Influences location, technology choices

Approaches to saving power

Infrastructure Efficiency	Power generation Power distribution Cooling
Dynamic Power Scaling	Sleeping when idle Rate adaptation VM consolidation
Computational Efficiency	**FAWN**

Approaches to saving power

Infrastructure Efficiency	Power generation Power distribution Cooling
Dynamic Power Scaling	Sleeping when idle Rate adaptation VM consolidation
Computational Efficiency	**FAWN**

Goal of computational efficiency:
Reduce the amount of energy to do useful work

FAWN

Fast Array of Wimpy Nodes

Improve computational efficiency of data-intensive computing using an array of well-balanced low-power systems.

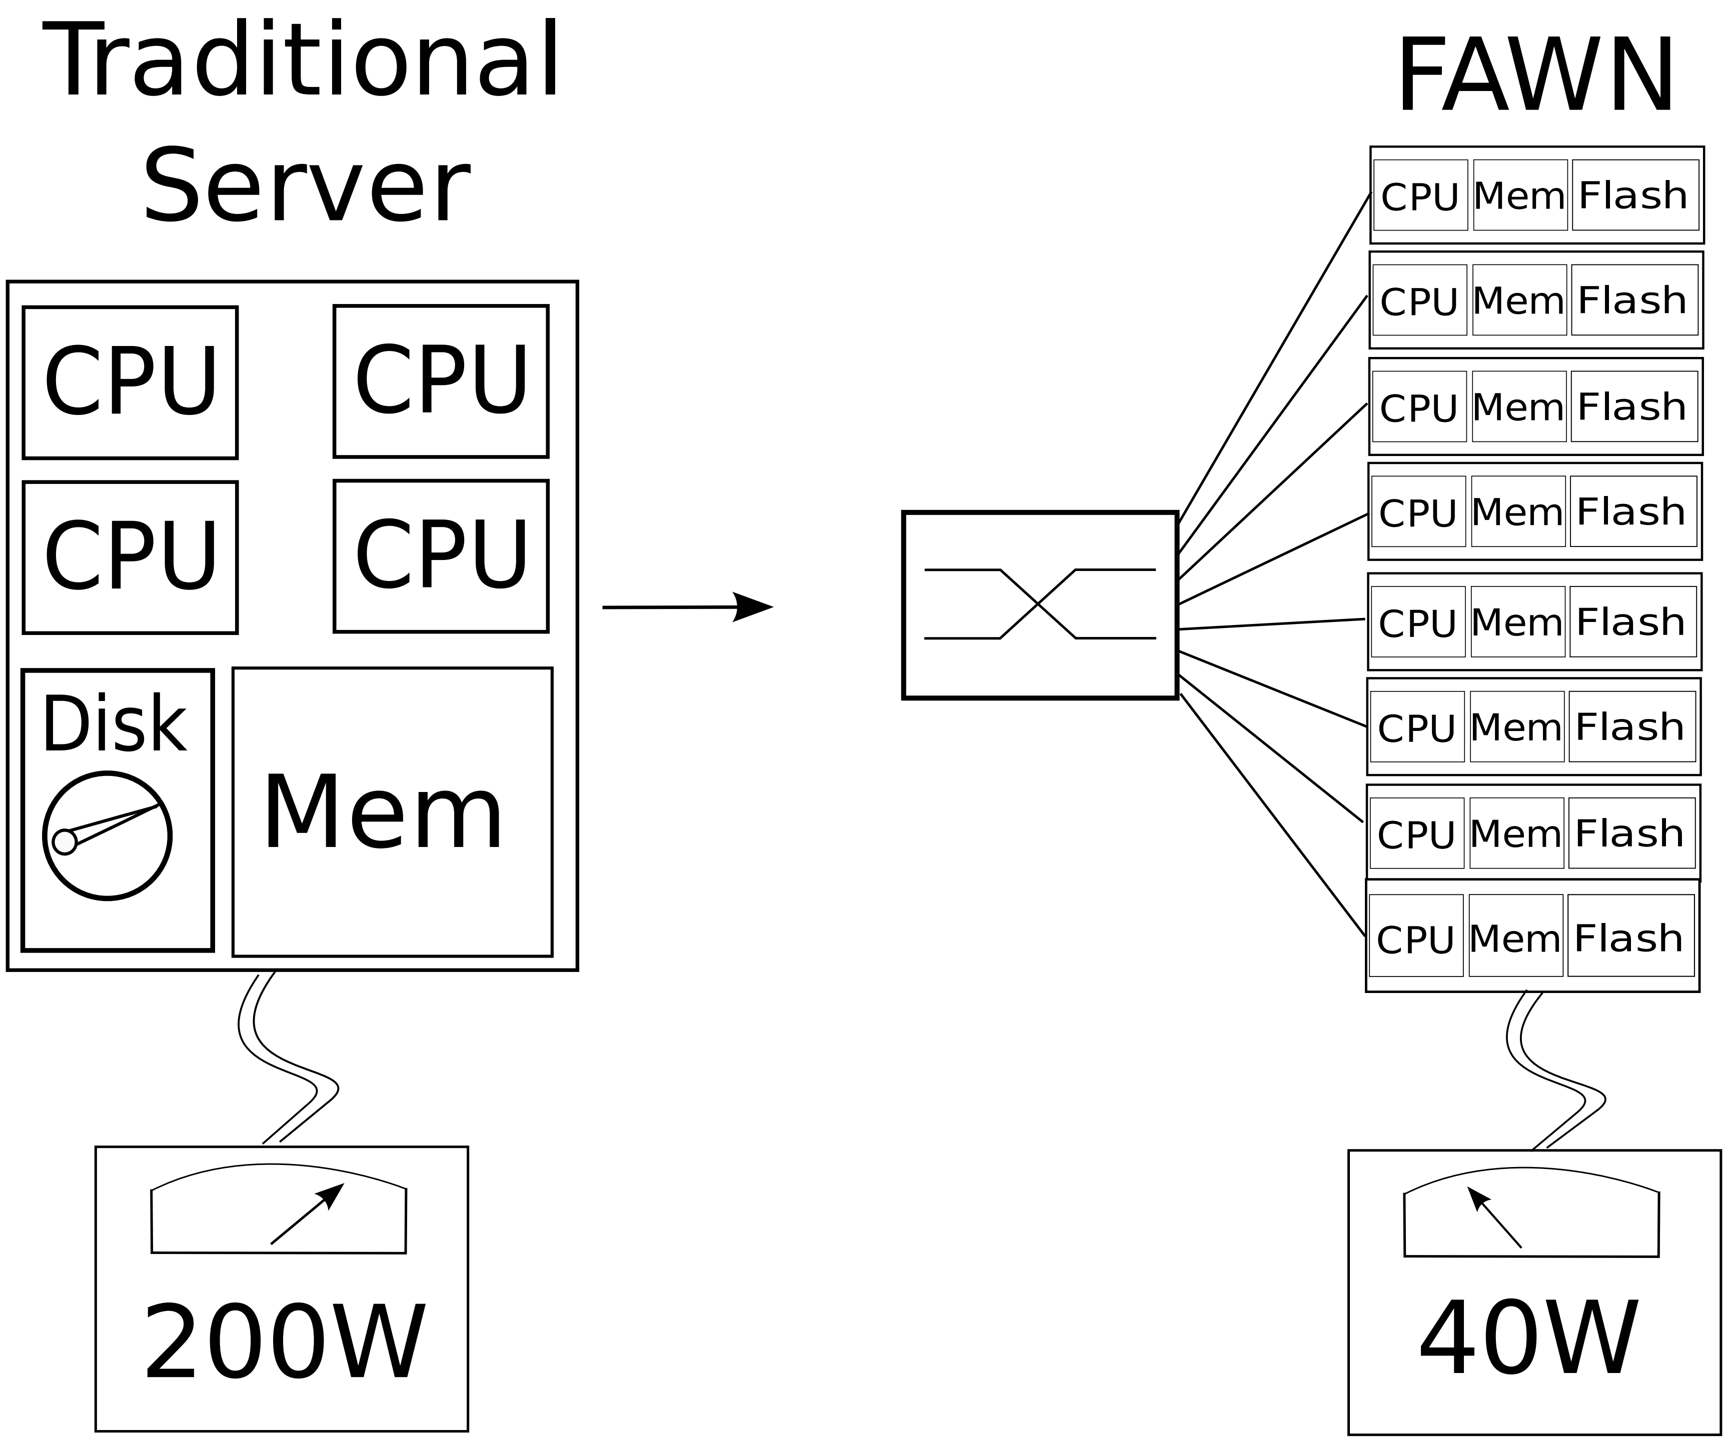

FAWN

Fast Array of Wimpy Nodes

Improve computational efficiency of data-intensive computing using an array of well-balanced low-power systems.

Traditional Server

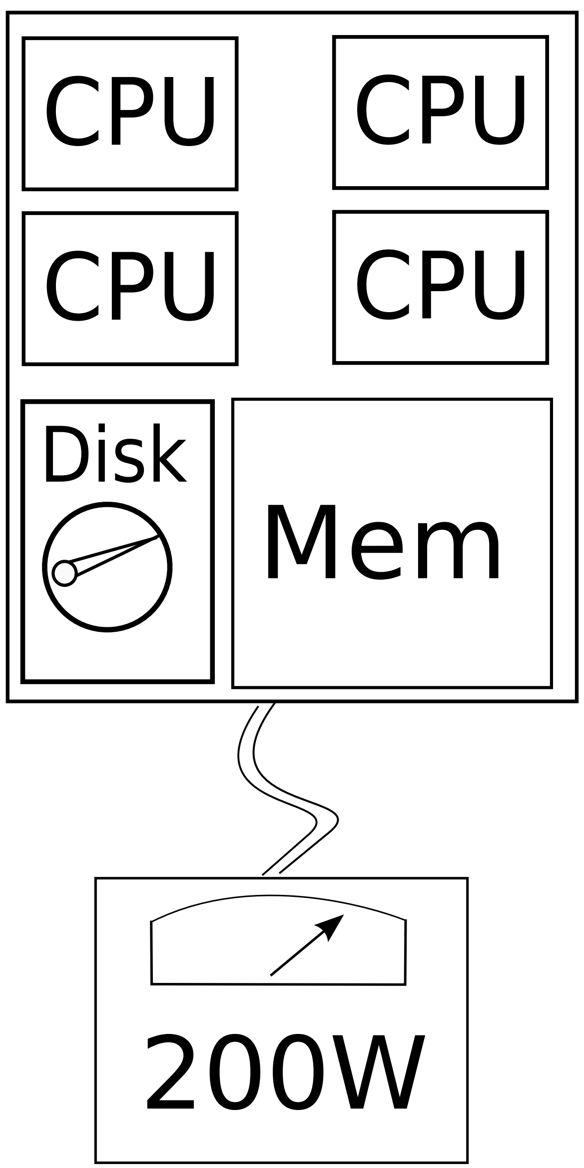

FAWN

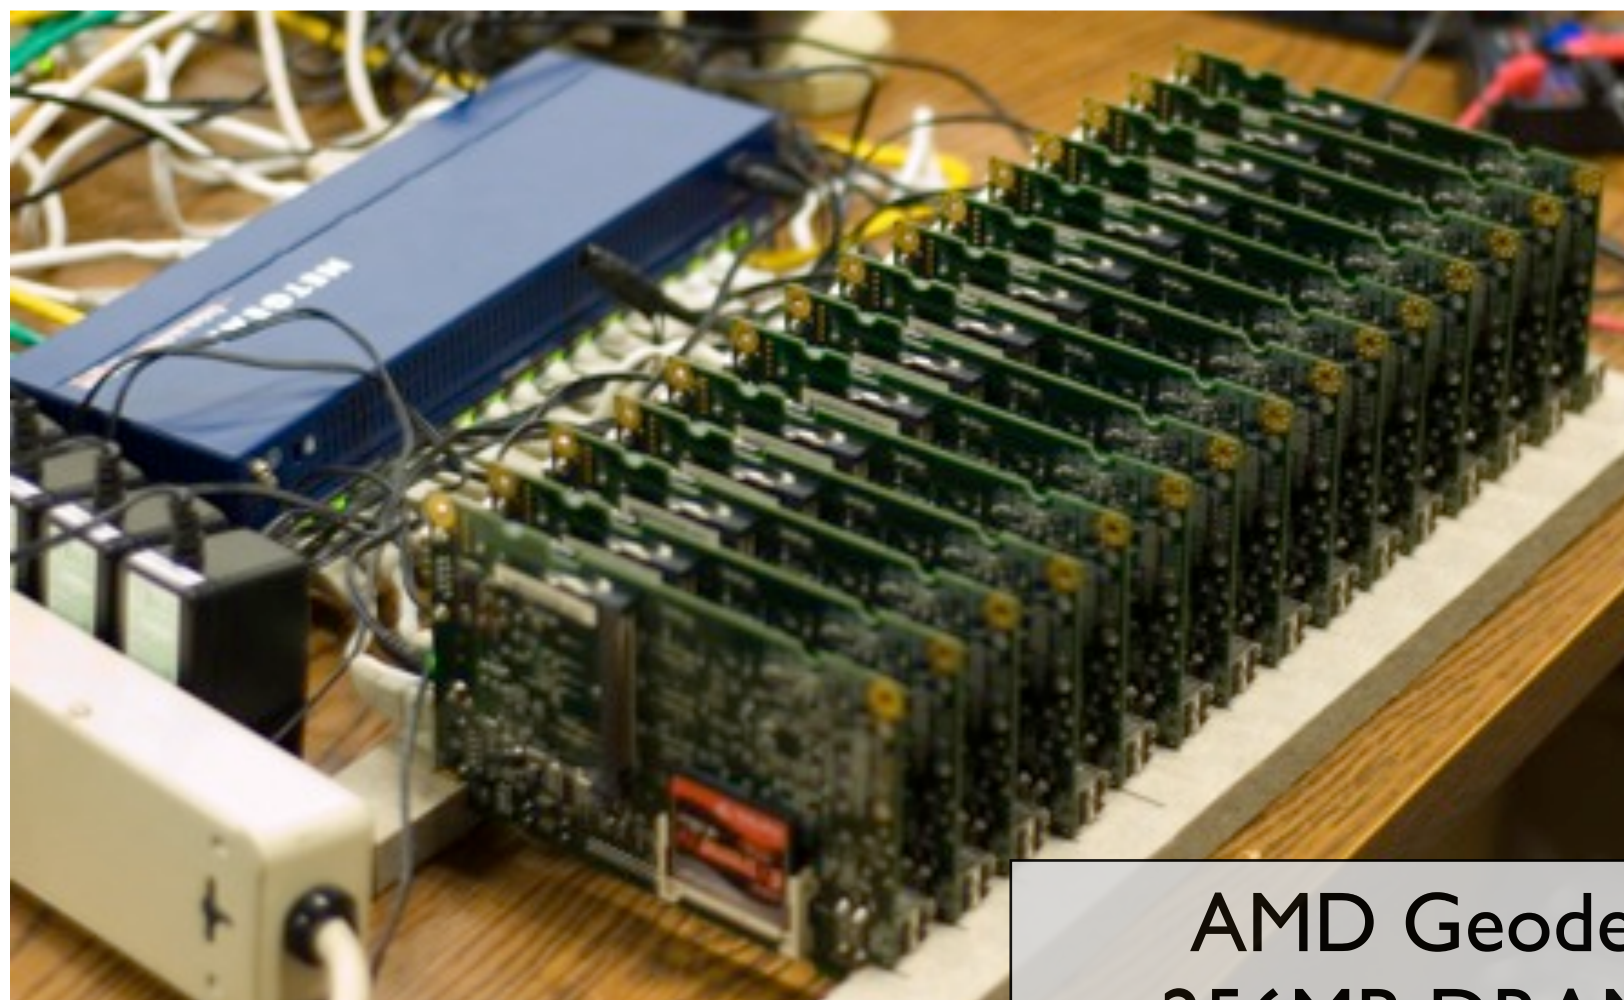

AMD Geode
256MB DRAM
4GB CompactFlash

Target: Data-intensive computing

- Large amounts of data
- Highly-parallelizable
- Fine-grained, independent tasks

Workloads amenable to "scale-out" approach

Outline

- What is FAWN?
- Why FAWN?
- When FAWN?
- Challenges (How FAWN?)

Why FAWN?

1. Fixed costs make dynamic power scaling difficult
2. FAWN balances system to save energy
3. FAWN targets sweet-spot in efficiency
4. FAWN reduces peak power consumption

1. Fixed power costs dominate

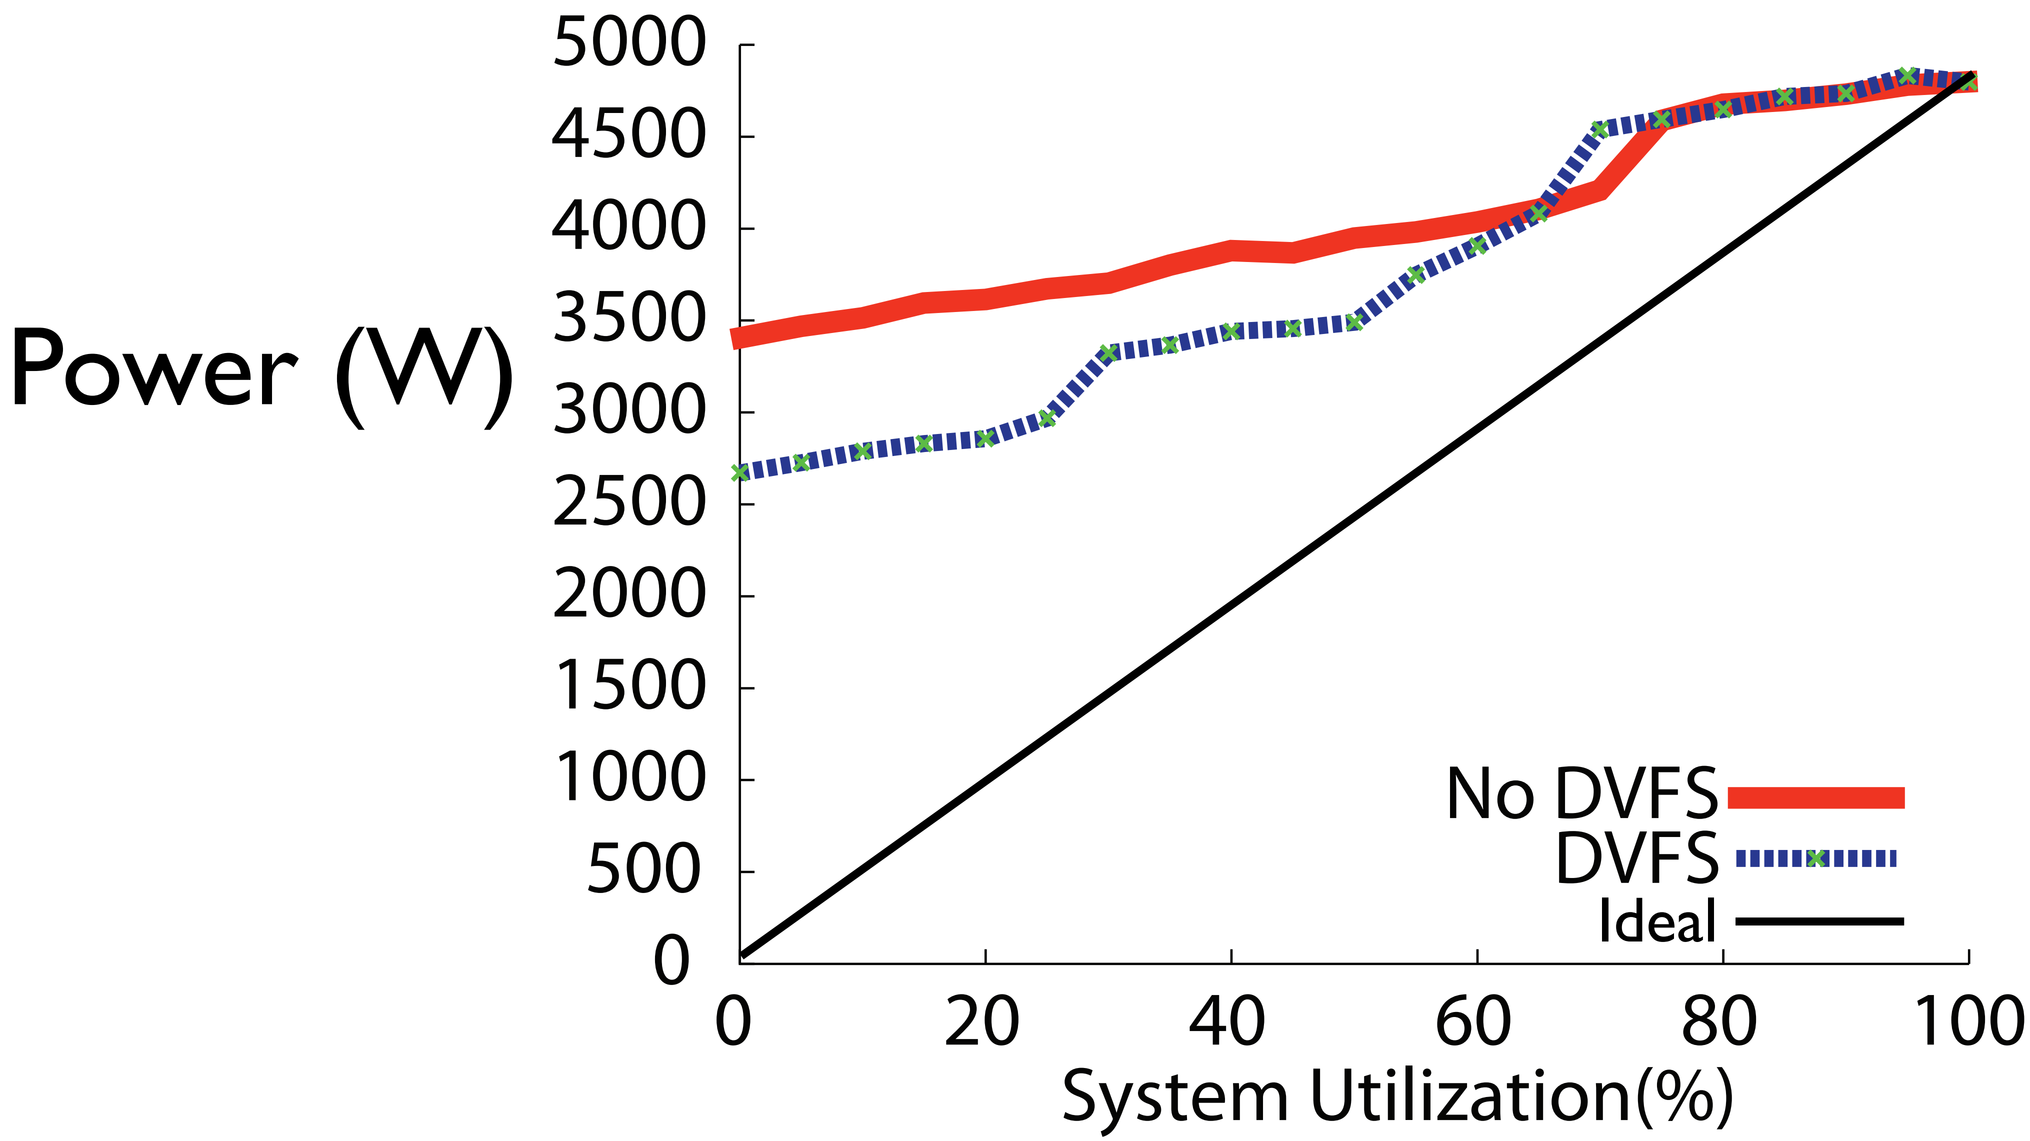

Figure adapted from Tolia et. al
HotPower 08

1. Fixed power costs dominate

70% of peak power
at 0% utilization!

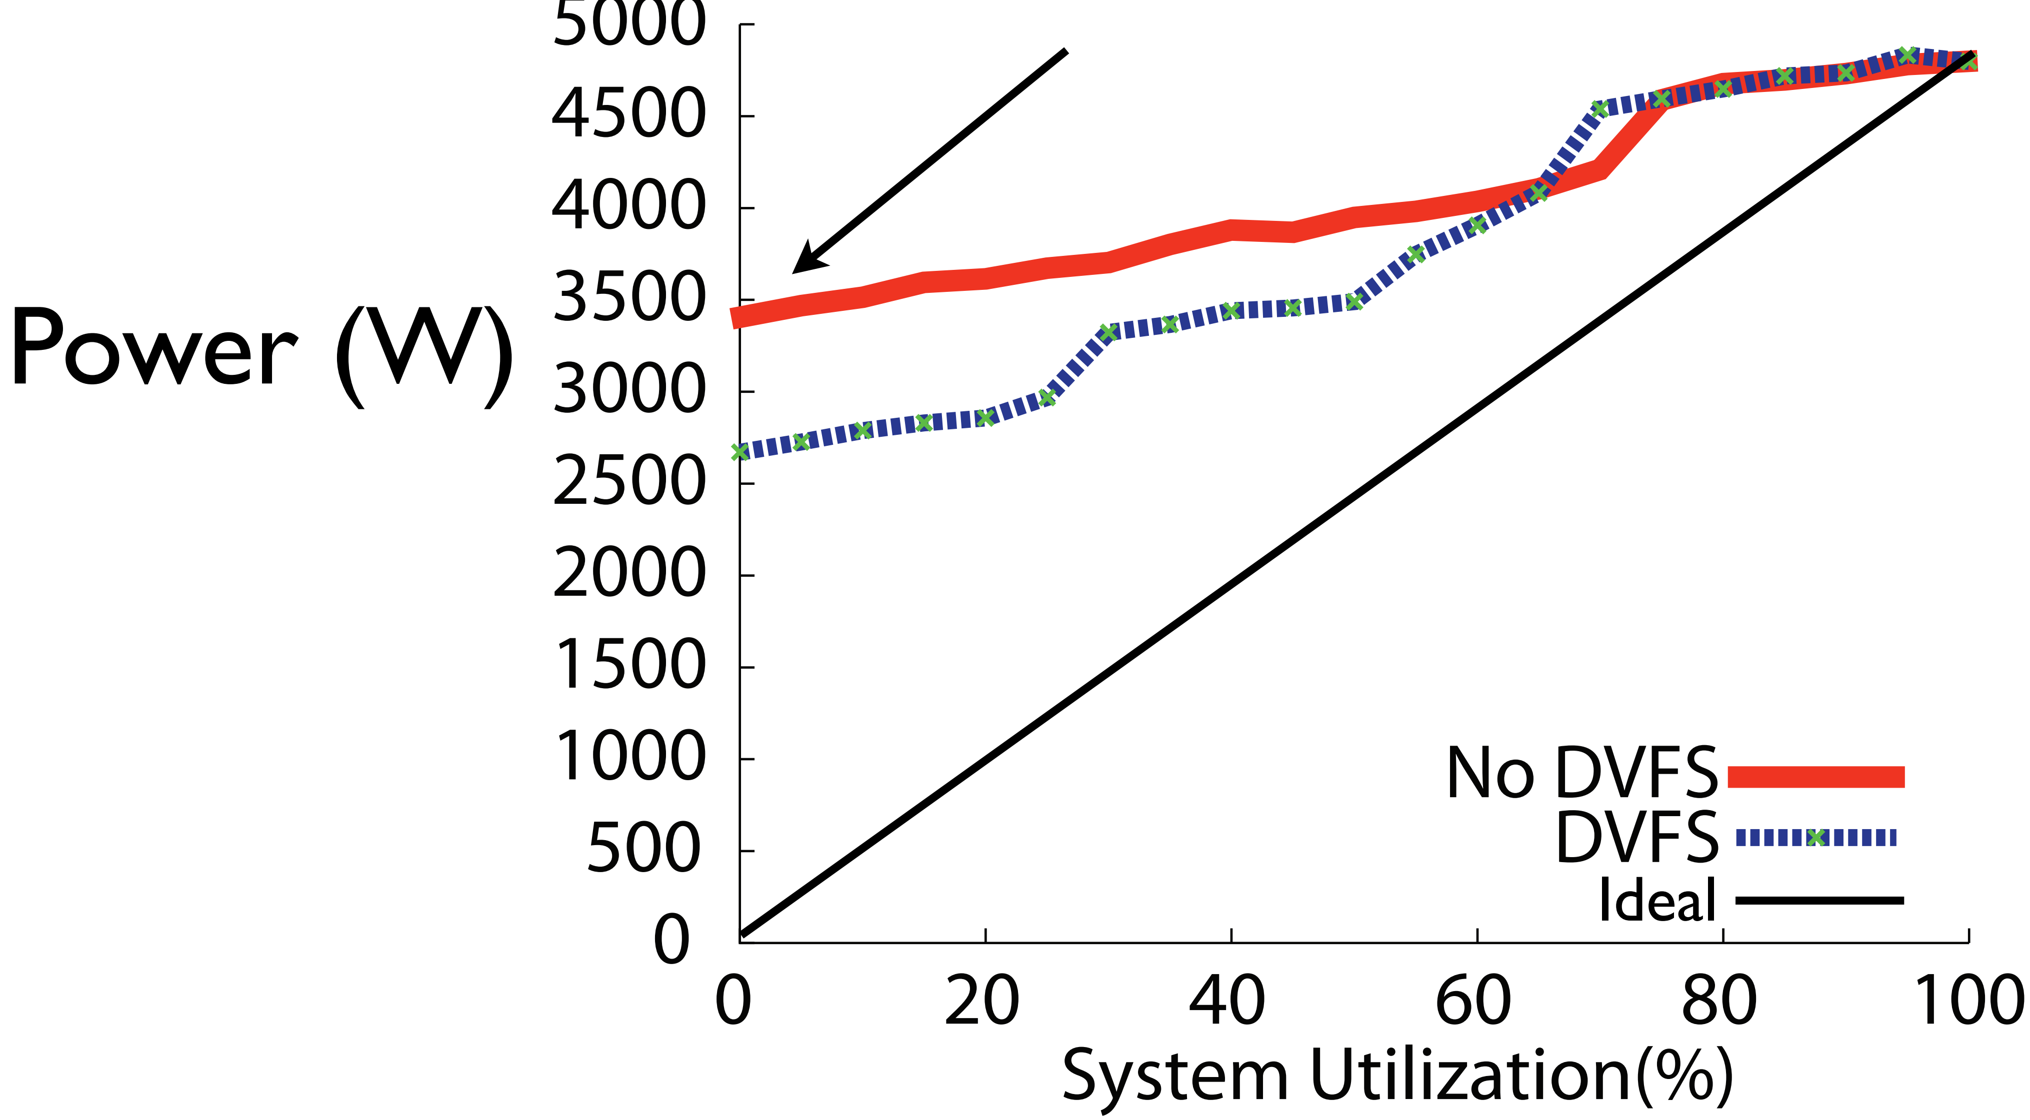

Figure adapted from Tolia et. al
HotPower 08

1. Fixed power costs dominate

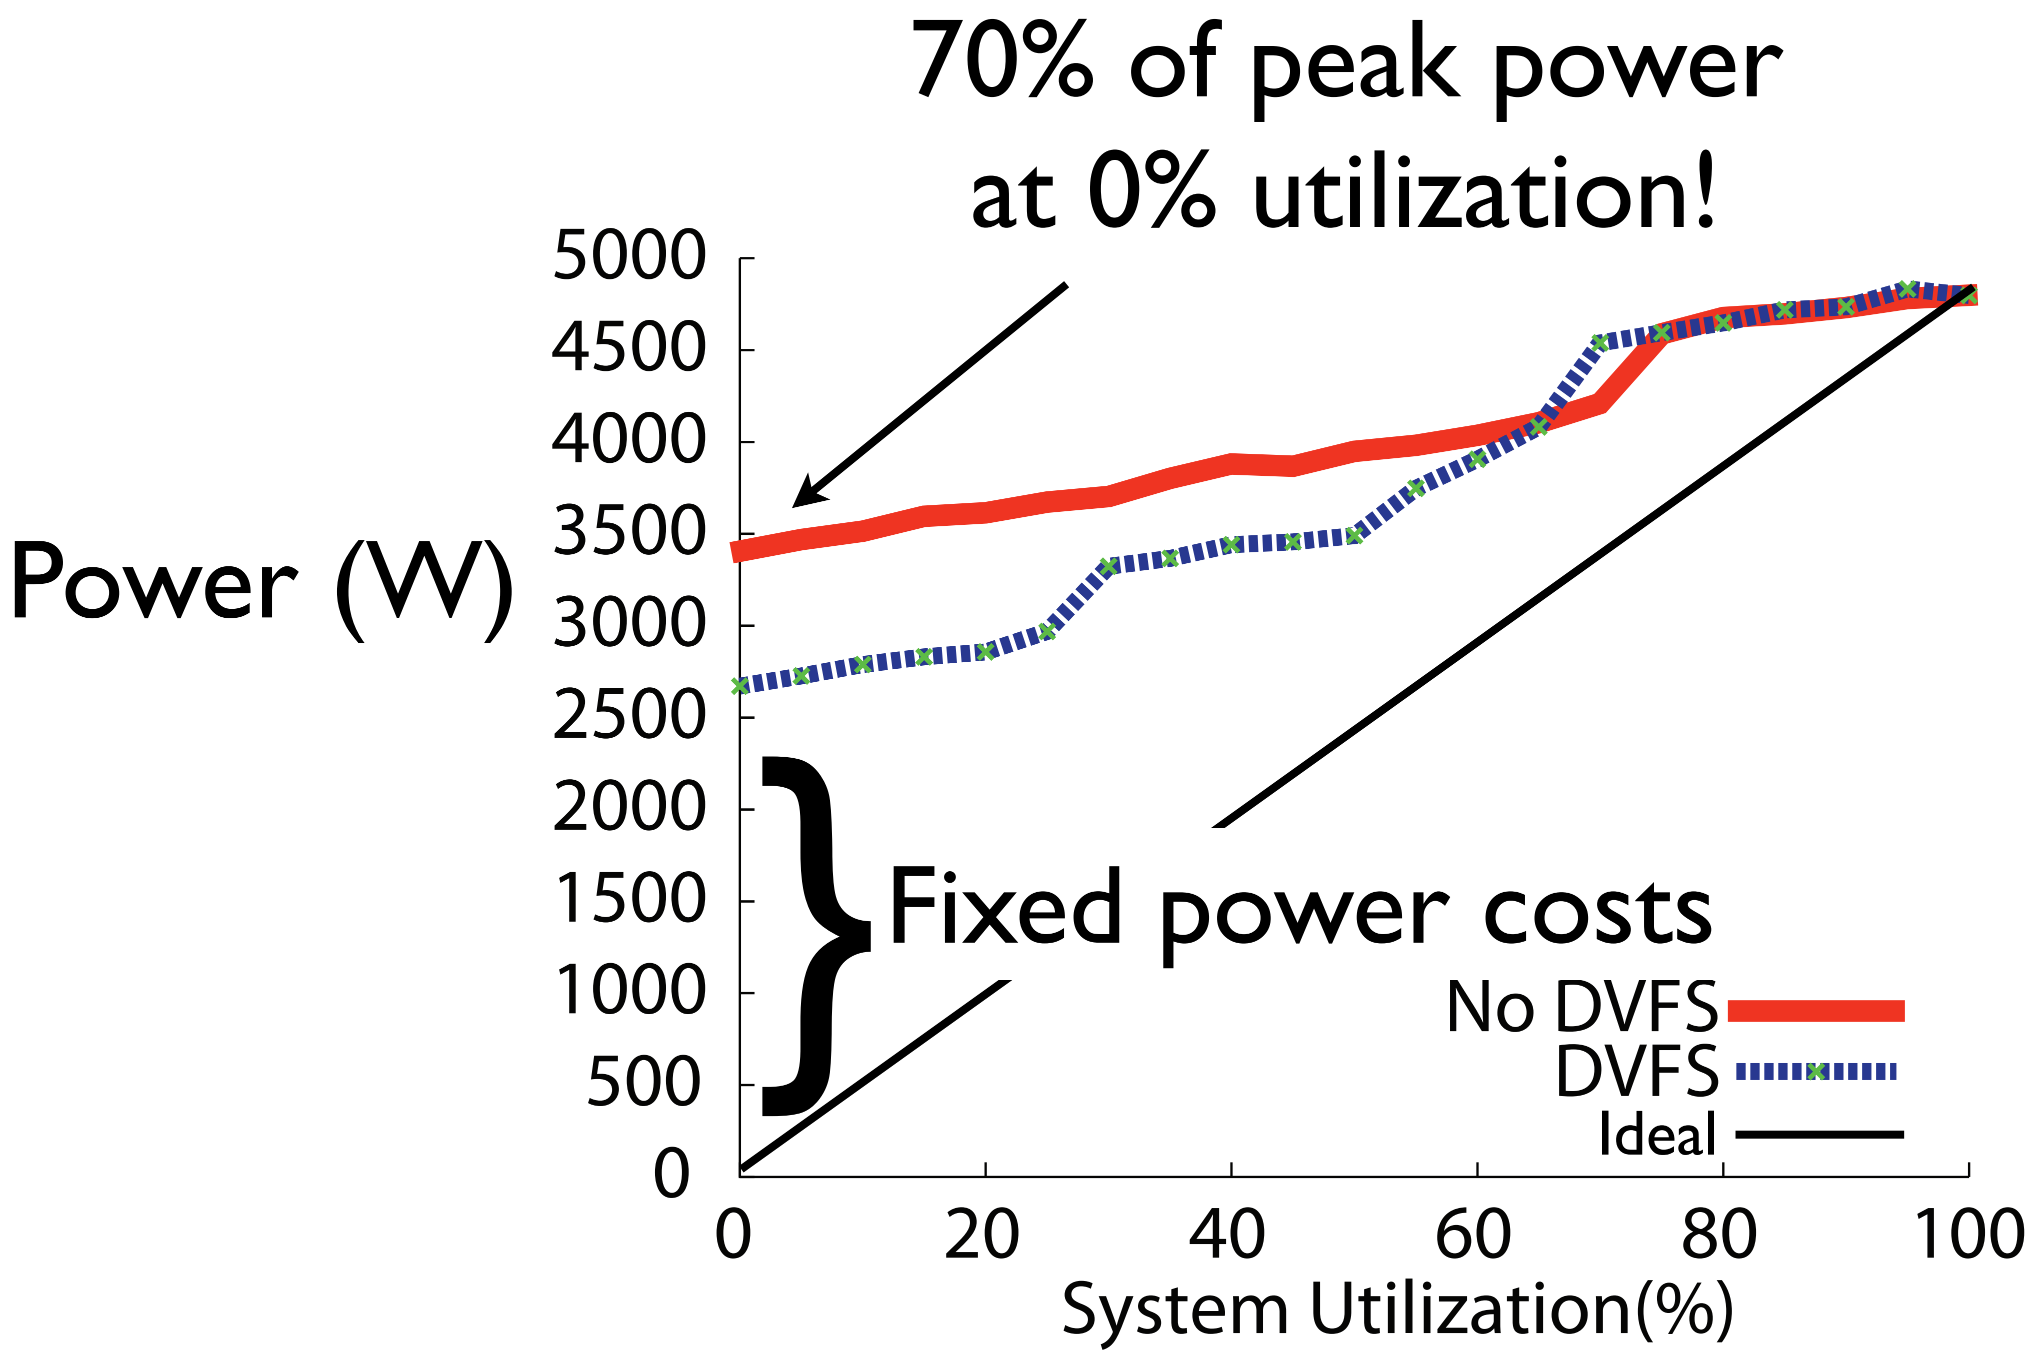

Figure adapted from Tolia et. al
HotPower 08

2. Balancing to save energy

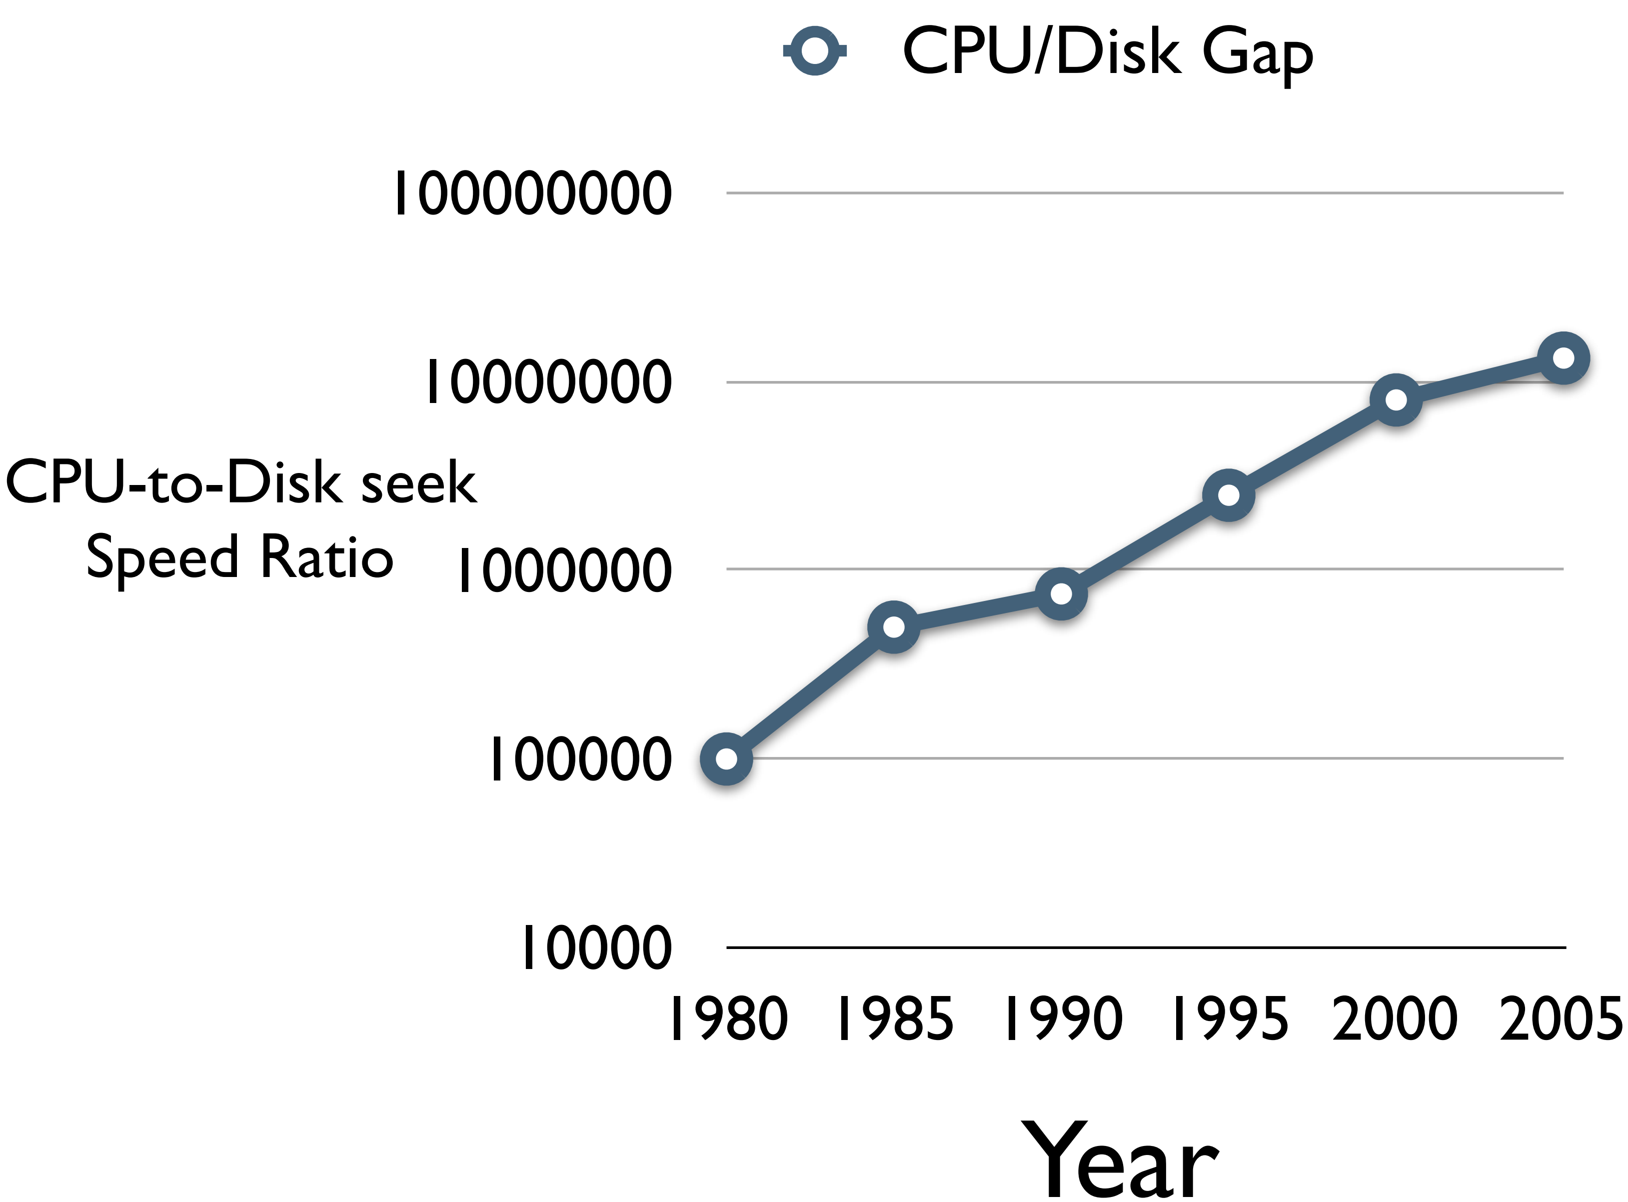

- How do we balance?
  - Big CPUs clocked down?
  - Embedded CPUs?
  - Why not use more disks with big CPUs?

3. Targeting the sweet-spot in efficiency

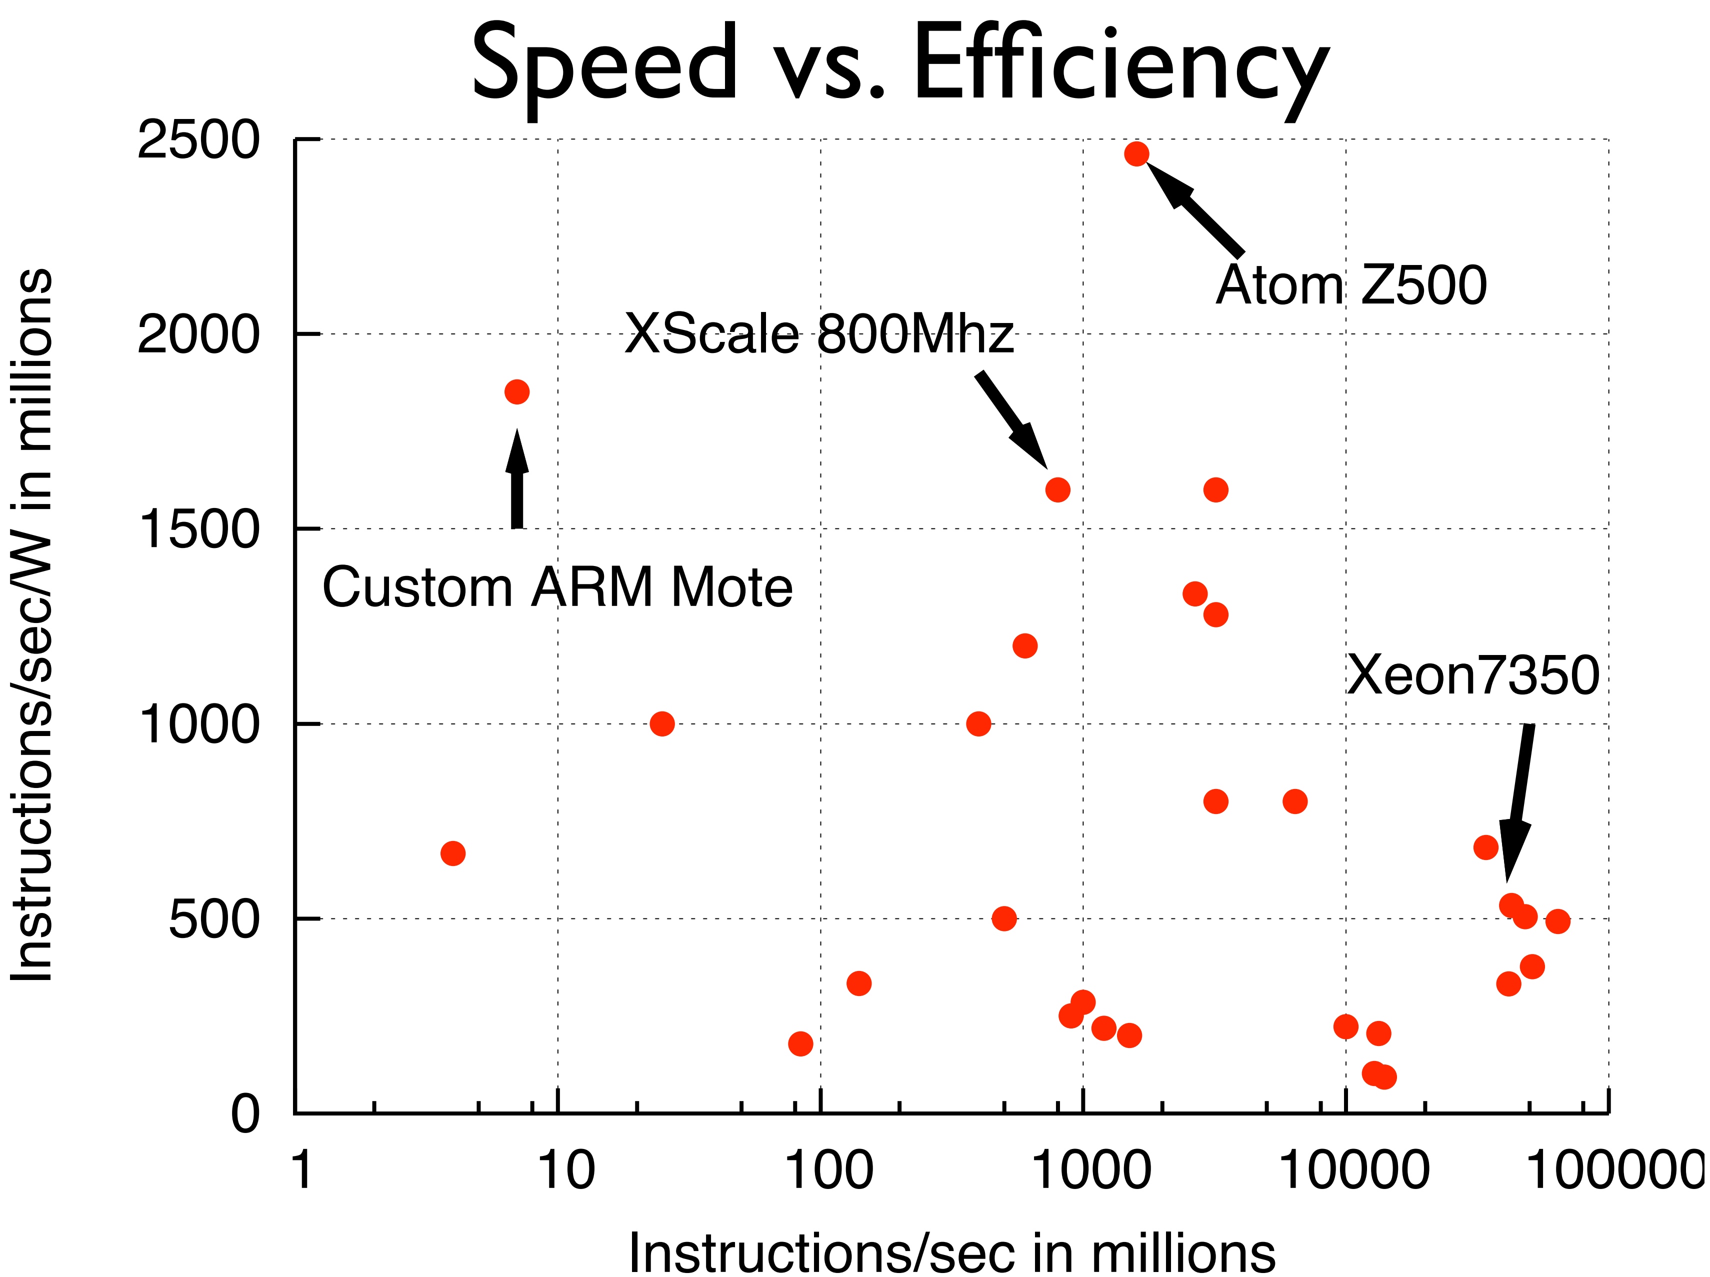

3. Targeting the sweet-spot in efficiency

Fast processors mask memory wall at the cost of efficiency

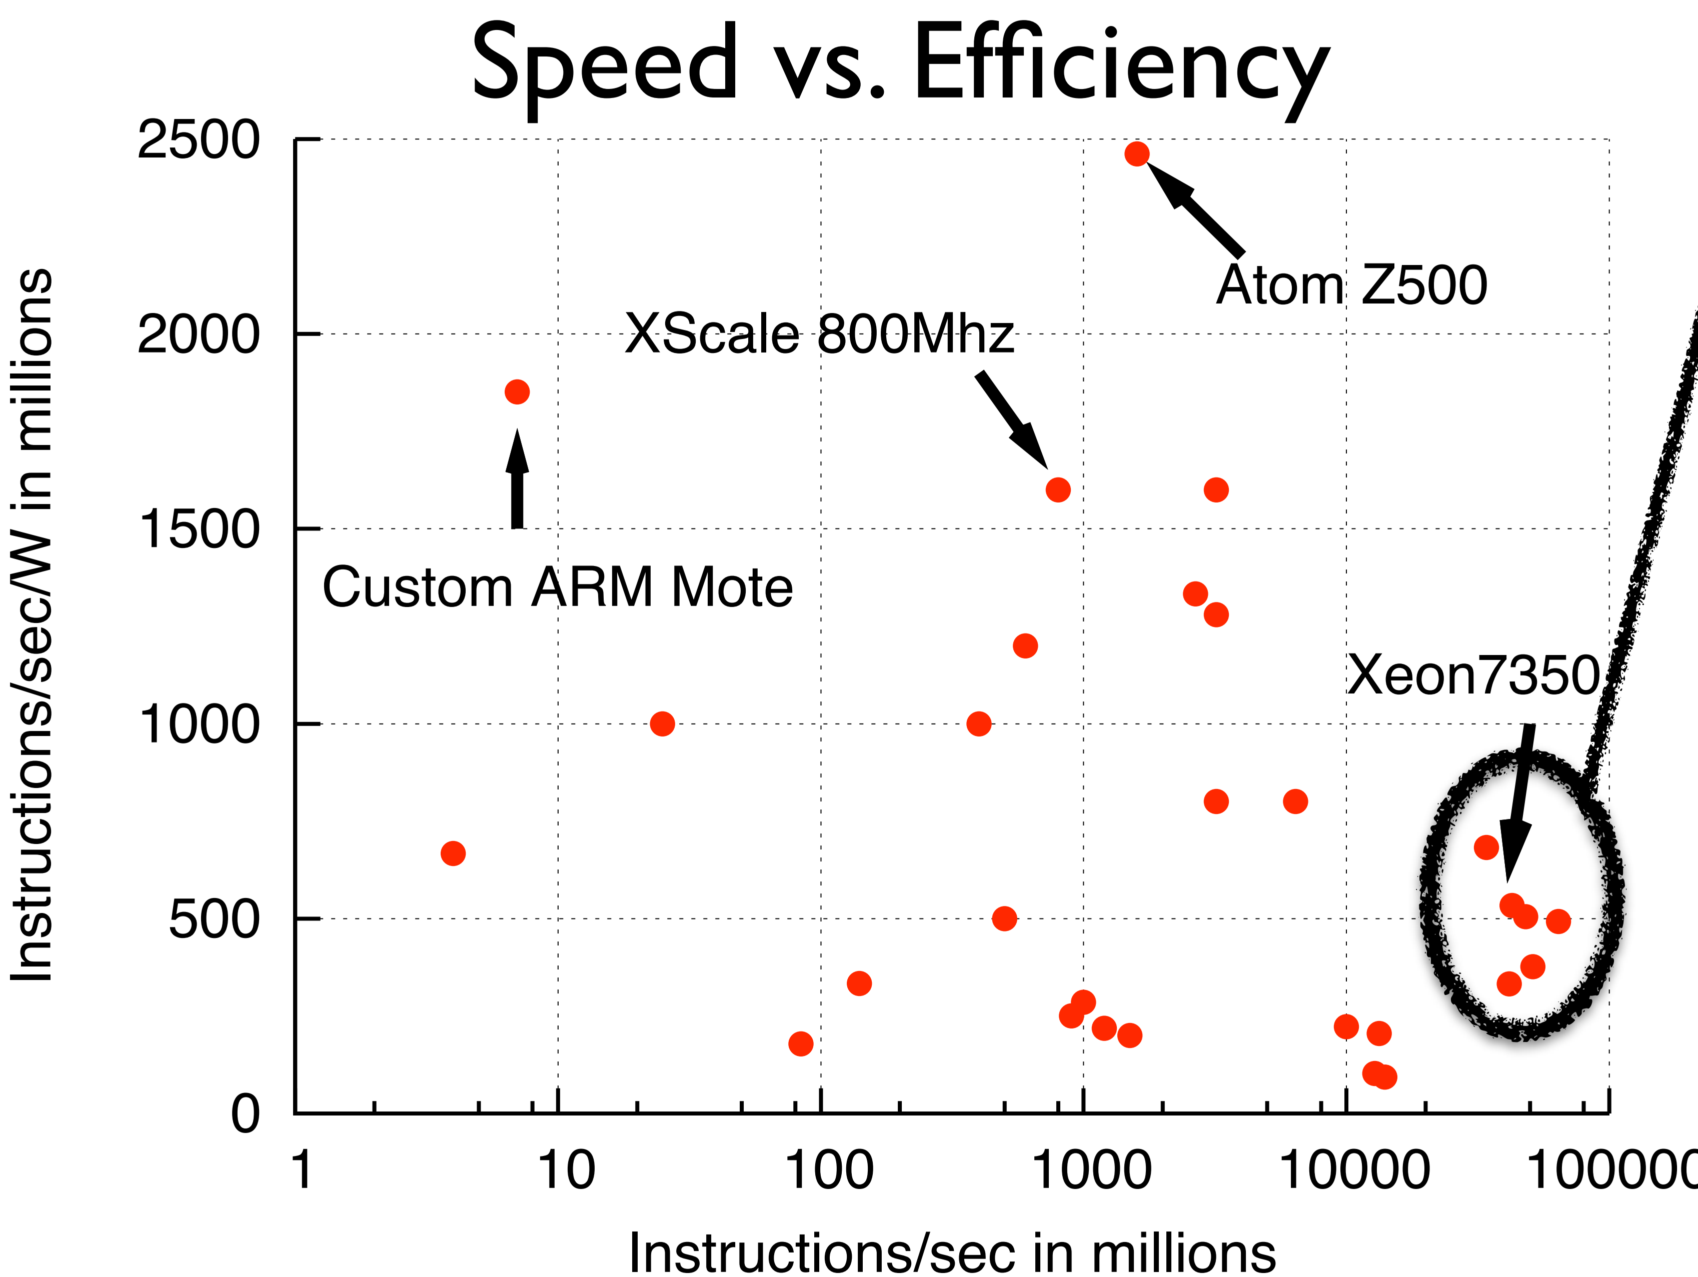

3. Targeting the sweet-spot in efficiency

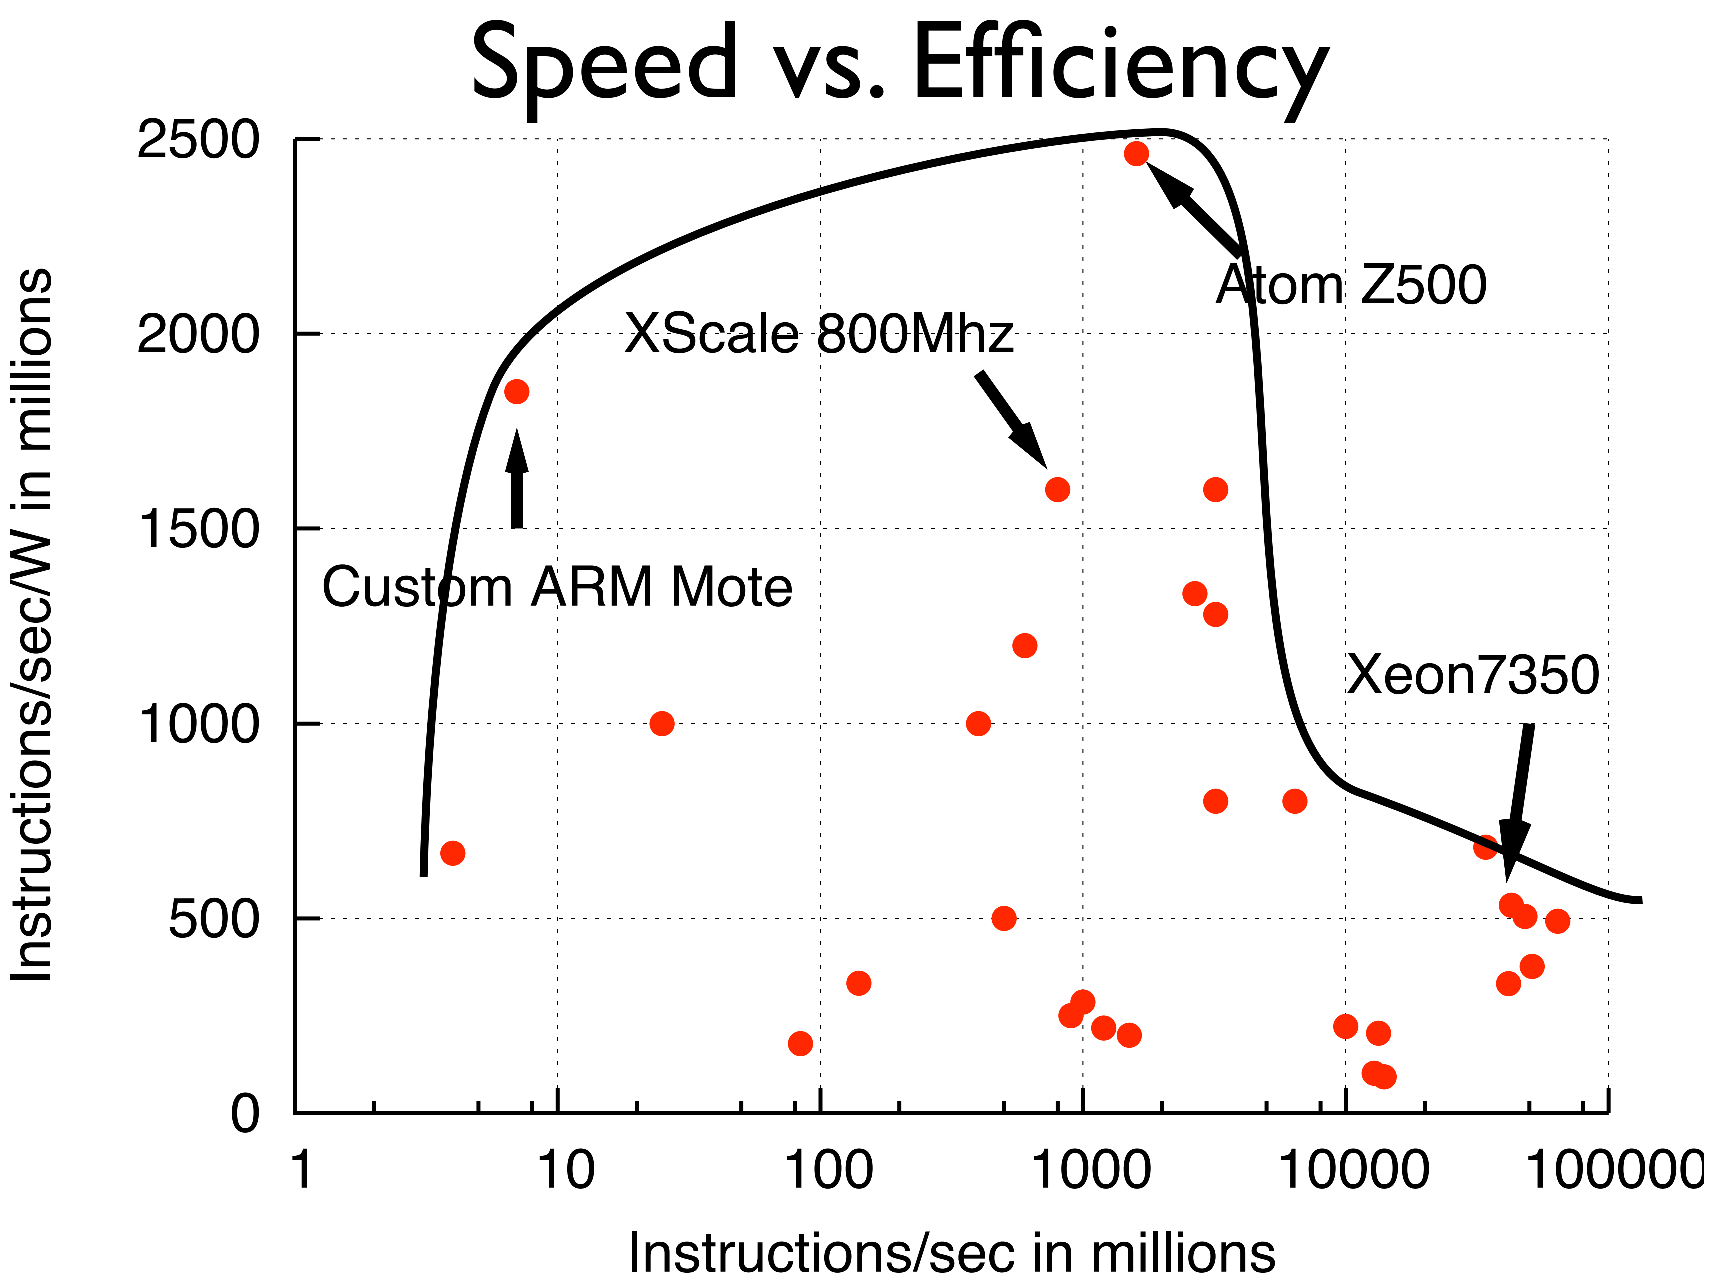

Fast processors
mask memory wall
at the cost of efficiency

3. Targeting the sweet-spot in efficiency

Fast processors
mask memory wall
at the cost of efficiency

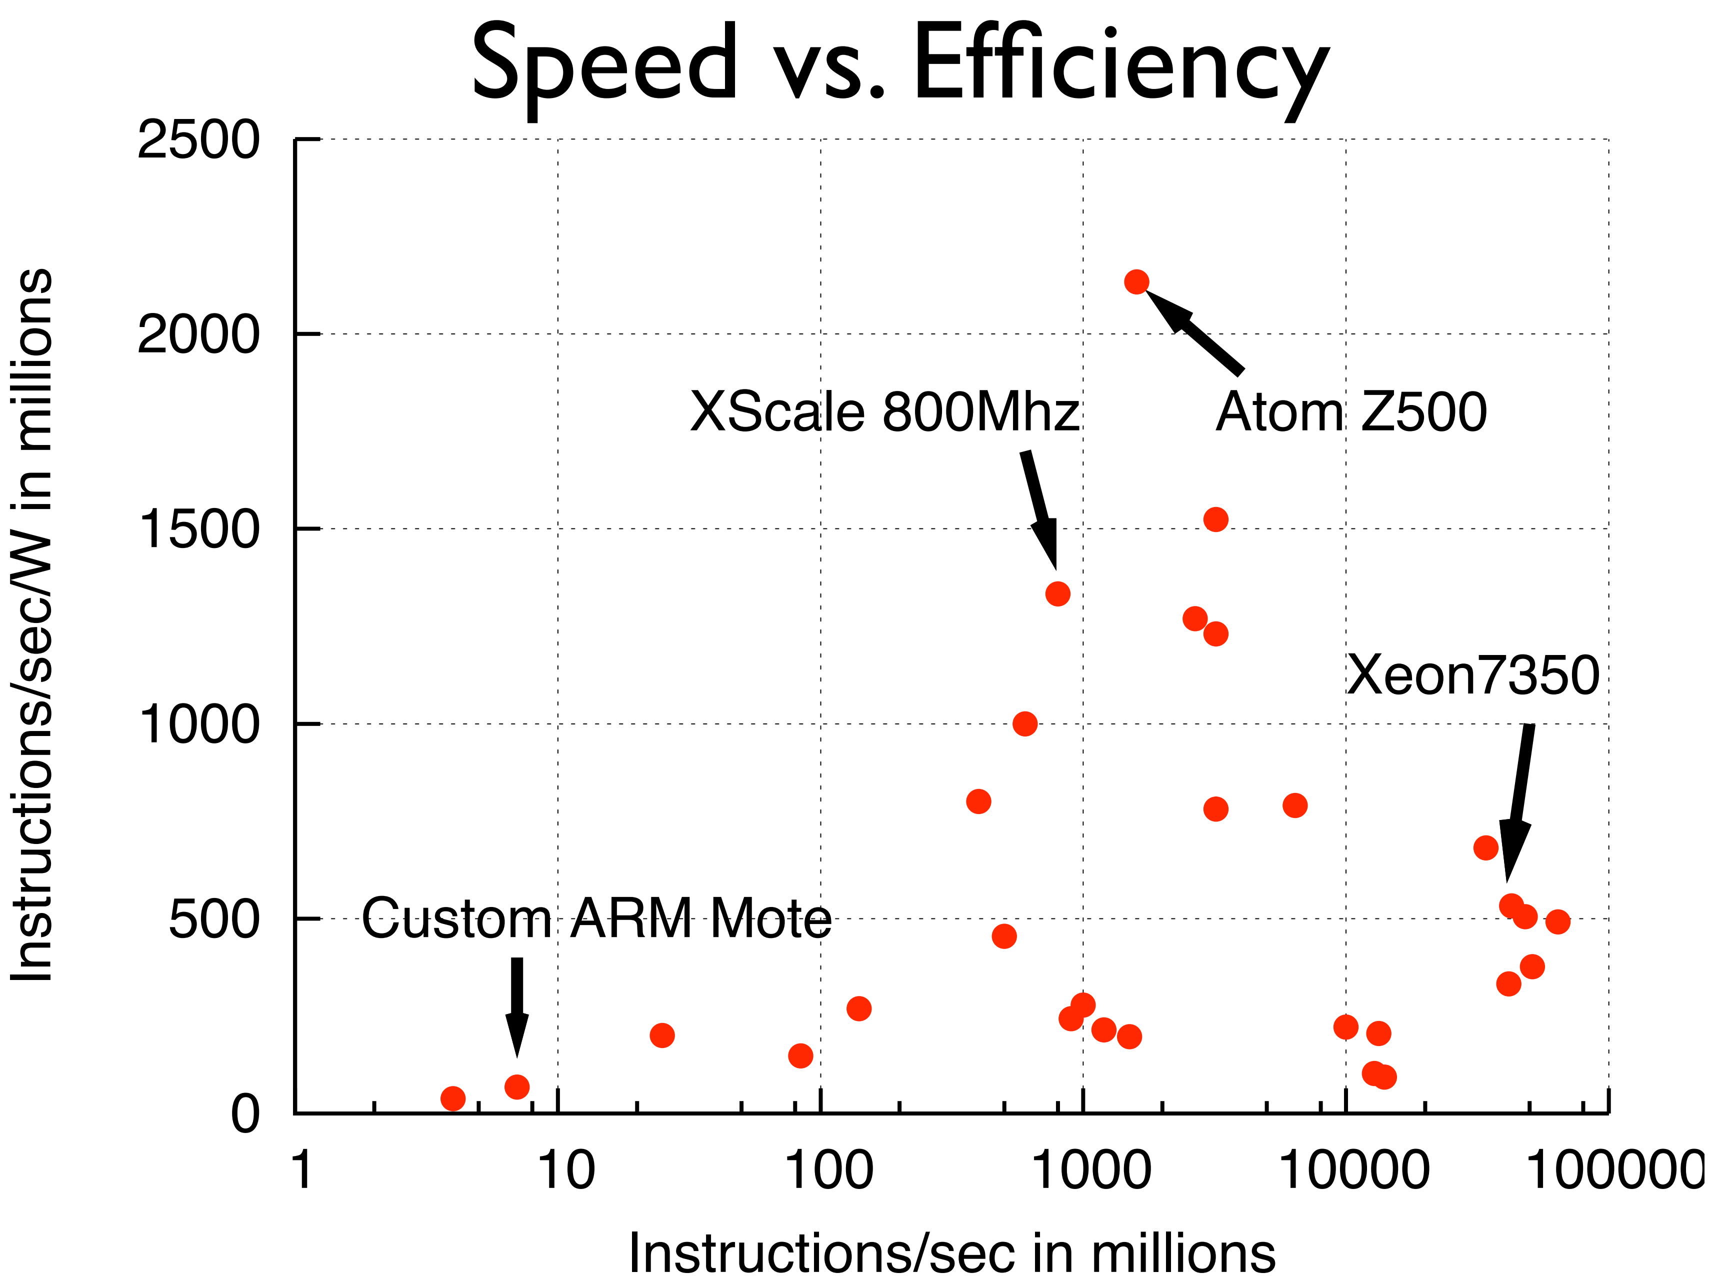

3. Targeting the sweet-spot in efficiency

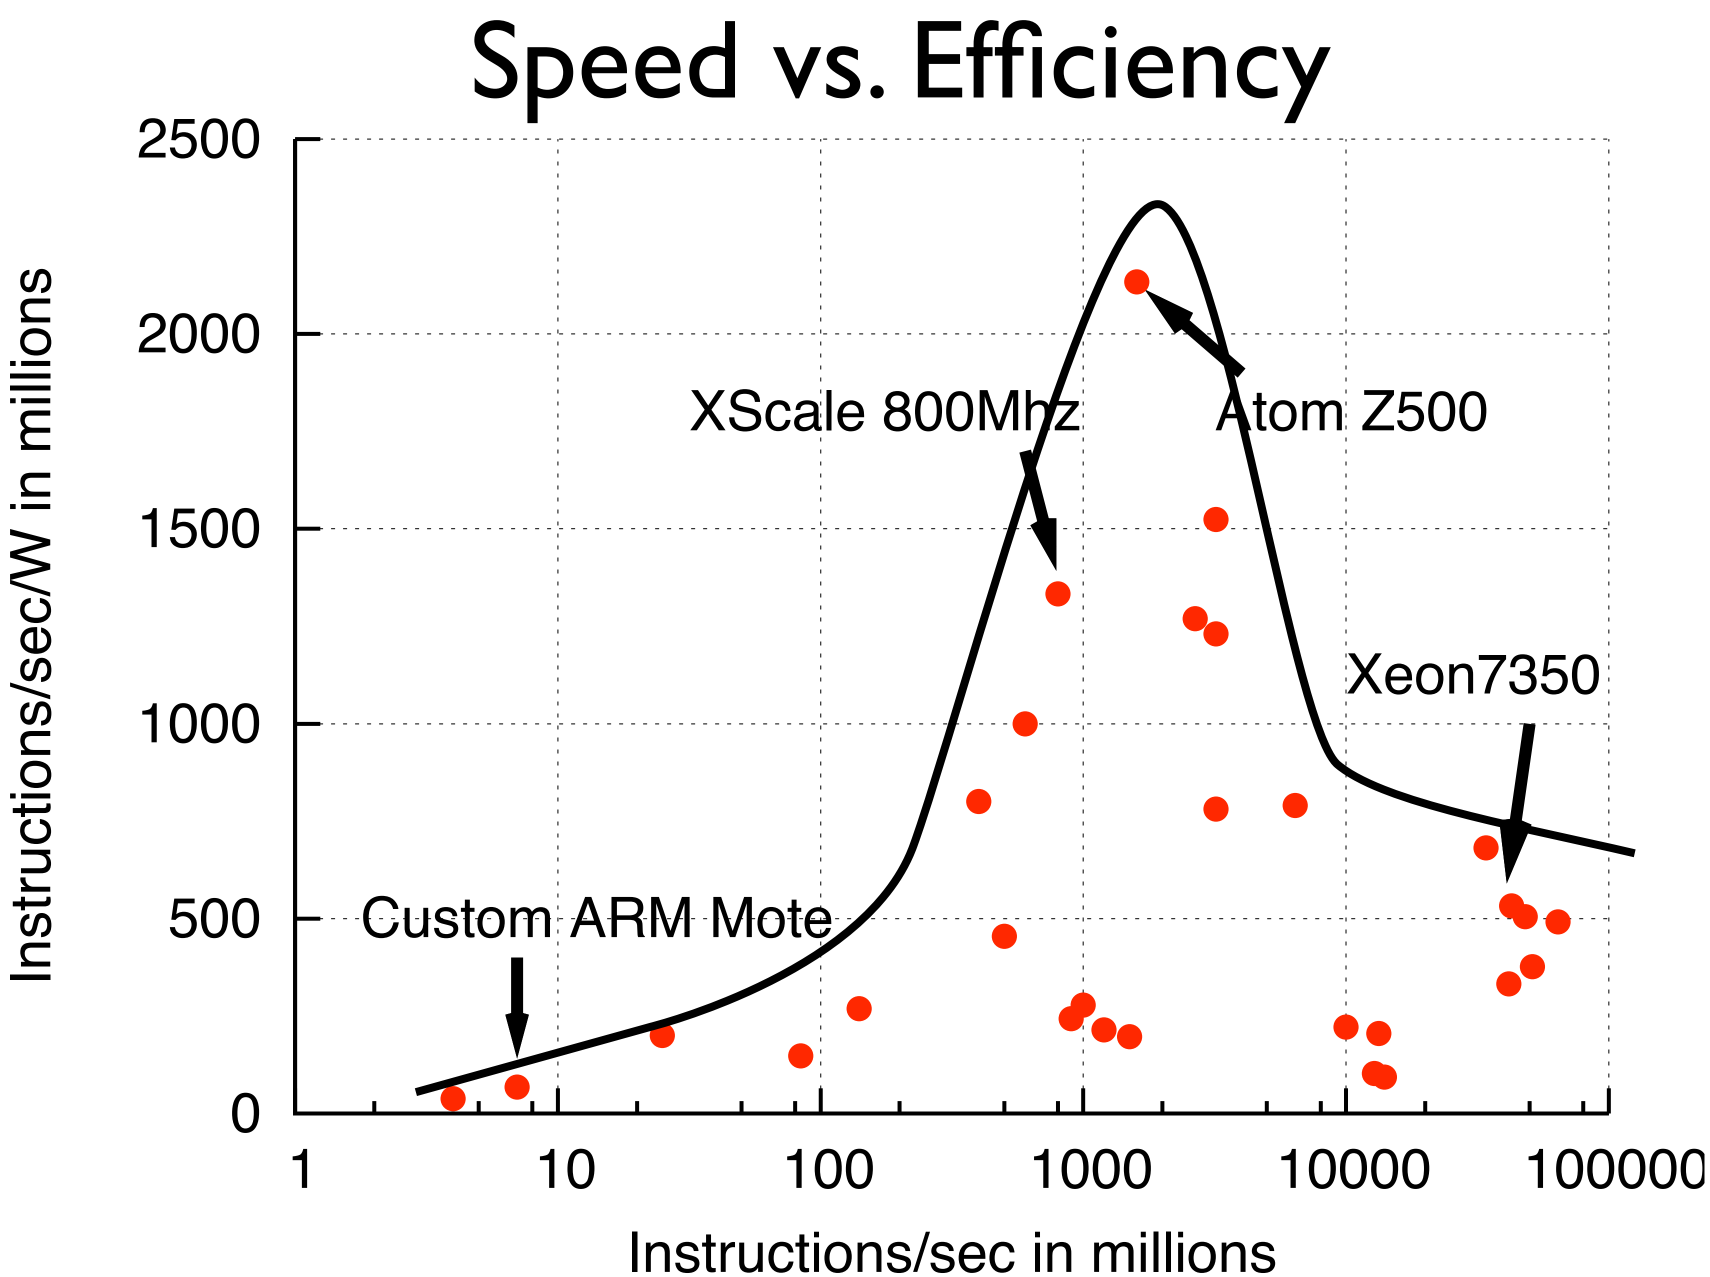

Fast processors
mask memory wall
at the cost of efficiency

3. Targeting the sweet-spot in efficiency

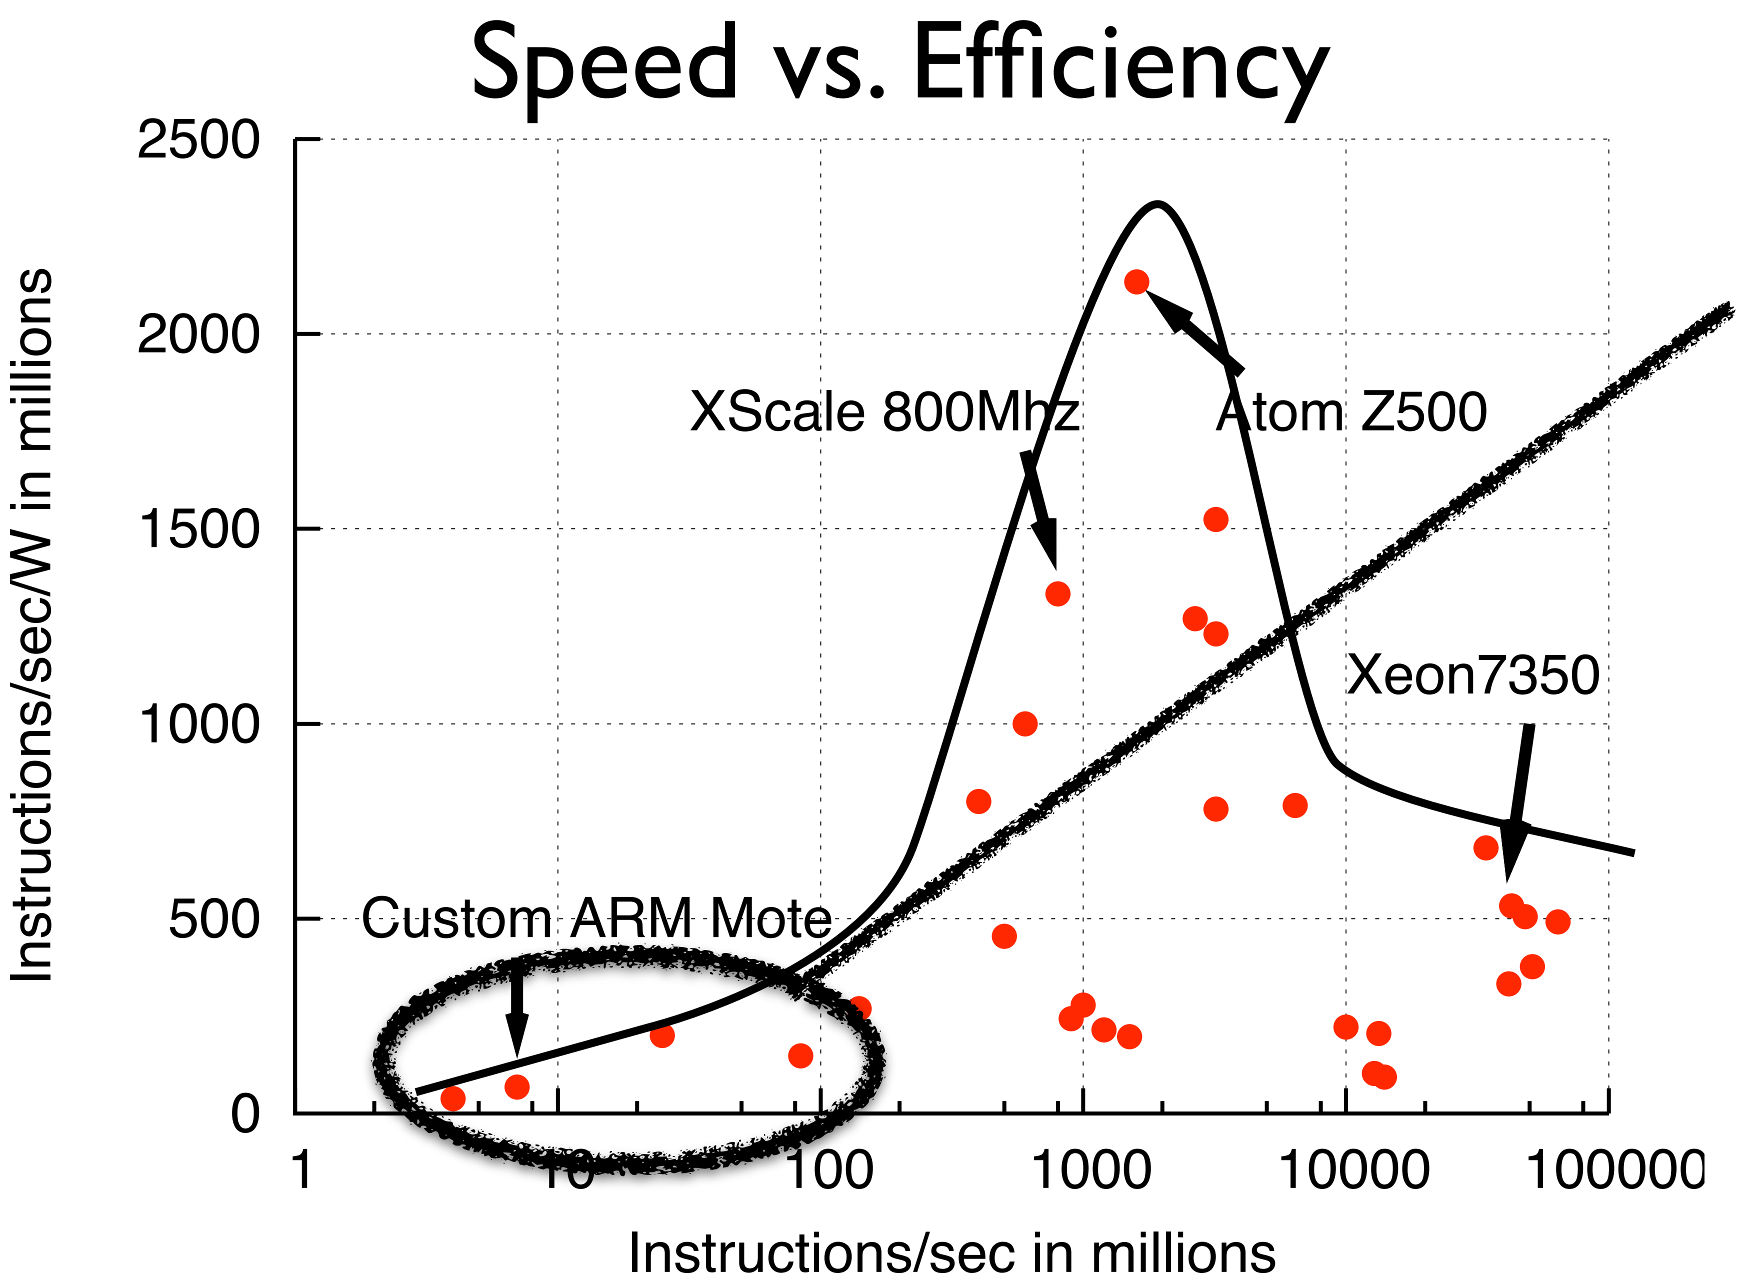

Fast processors mask memory wall at the cost of efficiency

Fixed power costs can dominate efficiency for slow processors

3. Targeting the sweet-spot in efficiency

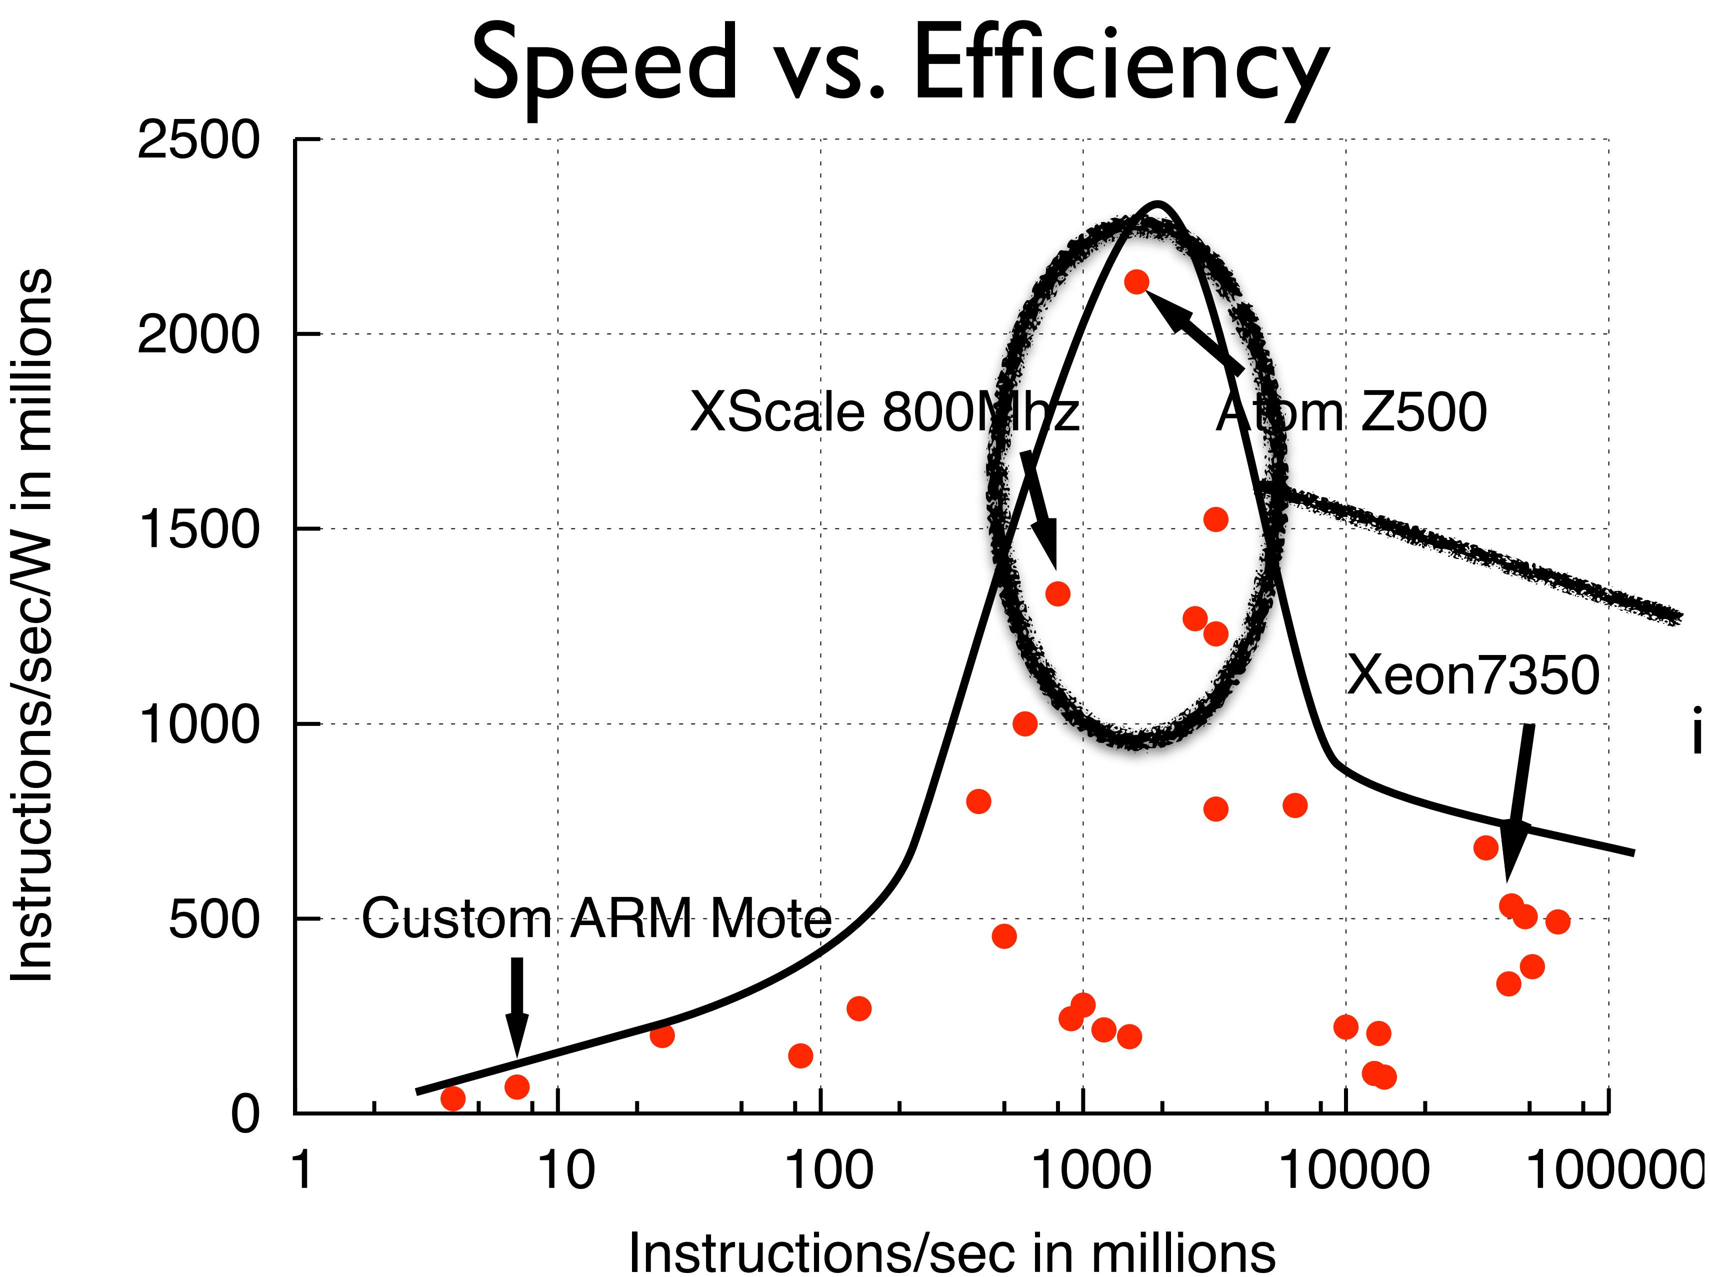

Fast processors mask memory wall at the cost of efficiency

Fixed power costs can dominate efficiency for slow processors

FAWN targets sweet spot in processor efficiency when including fixed costs

4. Reducing peak power consumption

- Provisioning for peak power requires:
  1. worst case cooling requirements
  2. UPS systems upon power failure
  3. power generation and substations

  investment

4. Reducing peak power consumption

- Provisioning for peak power requires:
    1. worst case cooling requirements
    2. UPS systems upon power failure
    3. power generation and substations investment

4. Reducing peak power consumption

- Provisioning for peak power requires:
    1. worst case cooling requirements
    2. UPS systems upon power failure
    3. power generation and substations investment

4. Reducing peak power consumption

- Provisioning for peak power requires:
  1. worst case cooling requirements
  2. UPS systems upon power failure
  3. power generation and substations investment

What is FAWN good for?

- **Random-access workloads (Key-value Lookup)**
- Scan-bound workloads (Hadoop, Data Analytics)
- **CPU-bound workloads (Compression, Encryption)**

Important metrics

Performance	Efficiency	Density	Cost
$\frac{\text{Work}}{\text{time}}$	$\frac{\text{Perf}}{\text{Watt}}$	$\frac{\text{Perf}}{\text{Volume}}$	$\frac{\text{Perf}}{\$}$

Random access workloads

FAWN + CF (4W)

Traditional + HD (87W)

Traditional + SSD (83W)

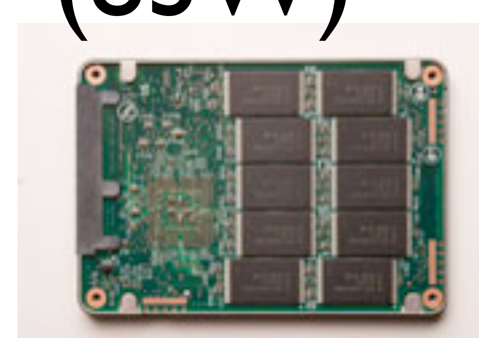

Random access workloads

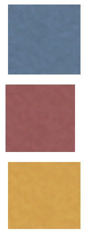

Random access workloads

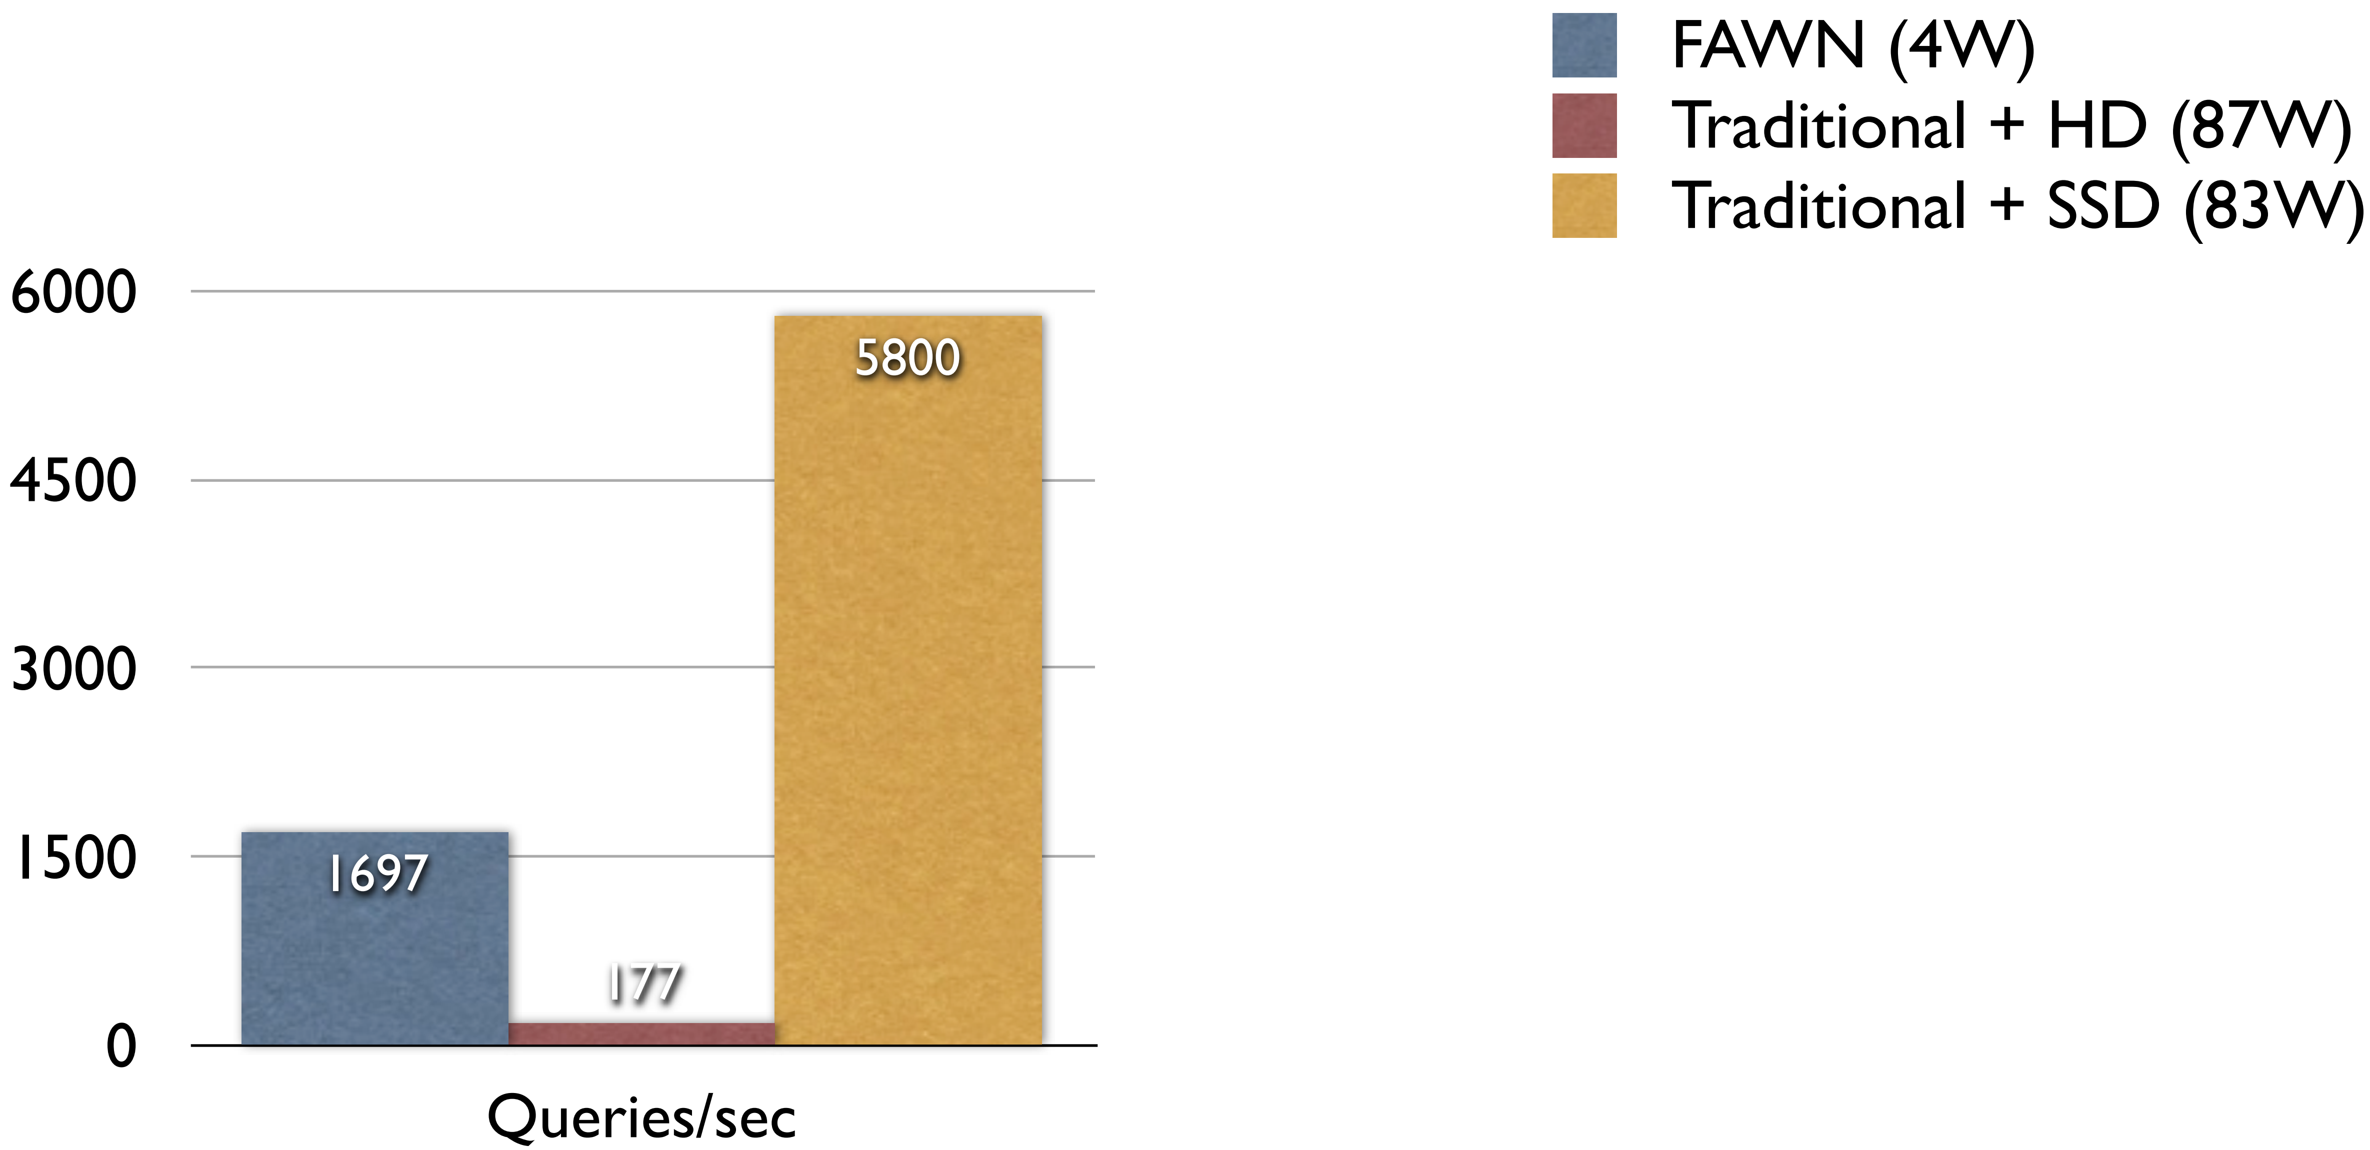

Performance

Random access workloads

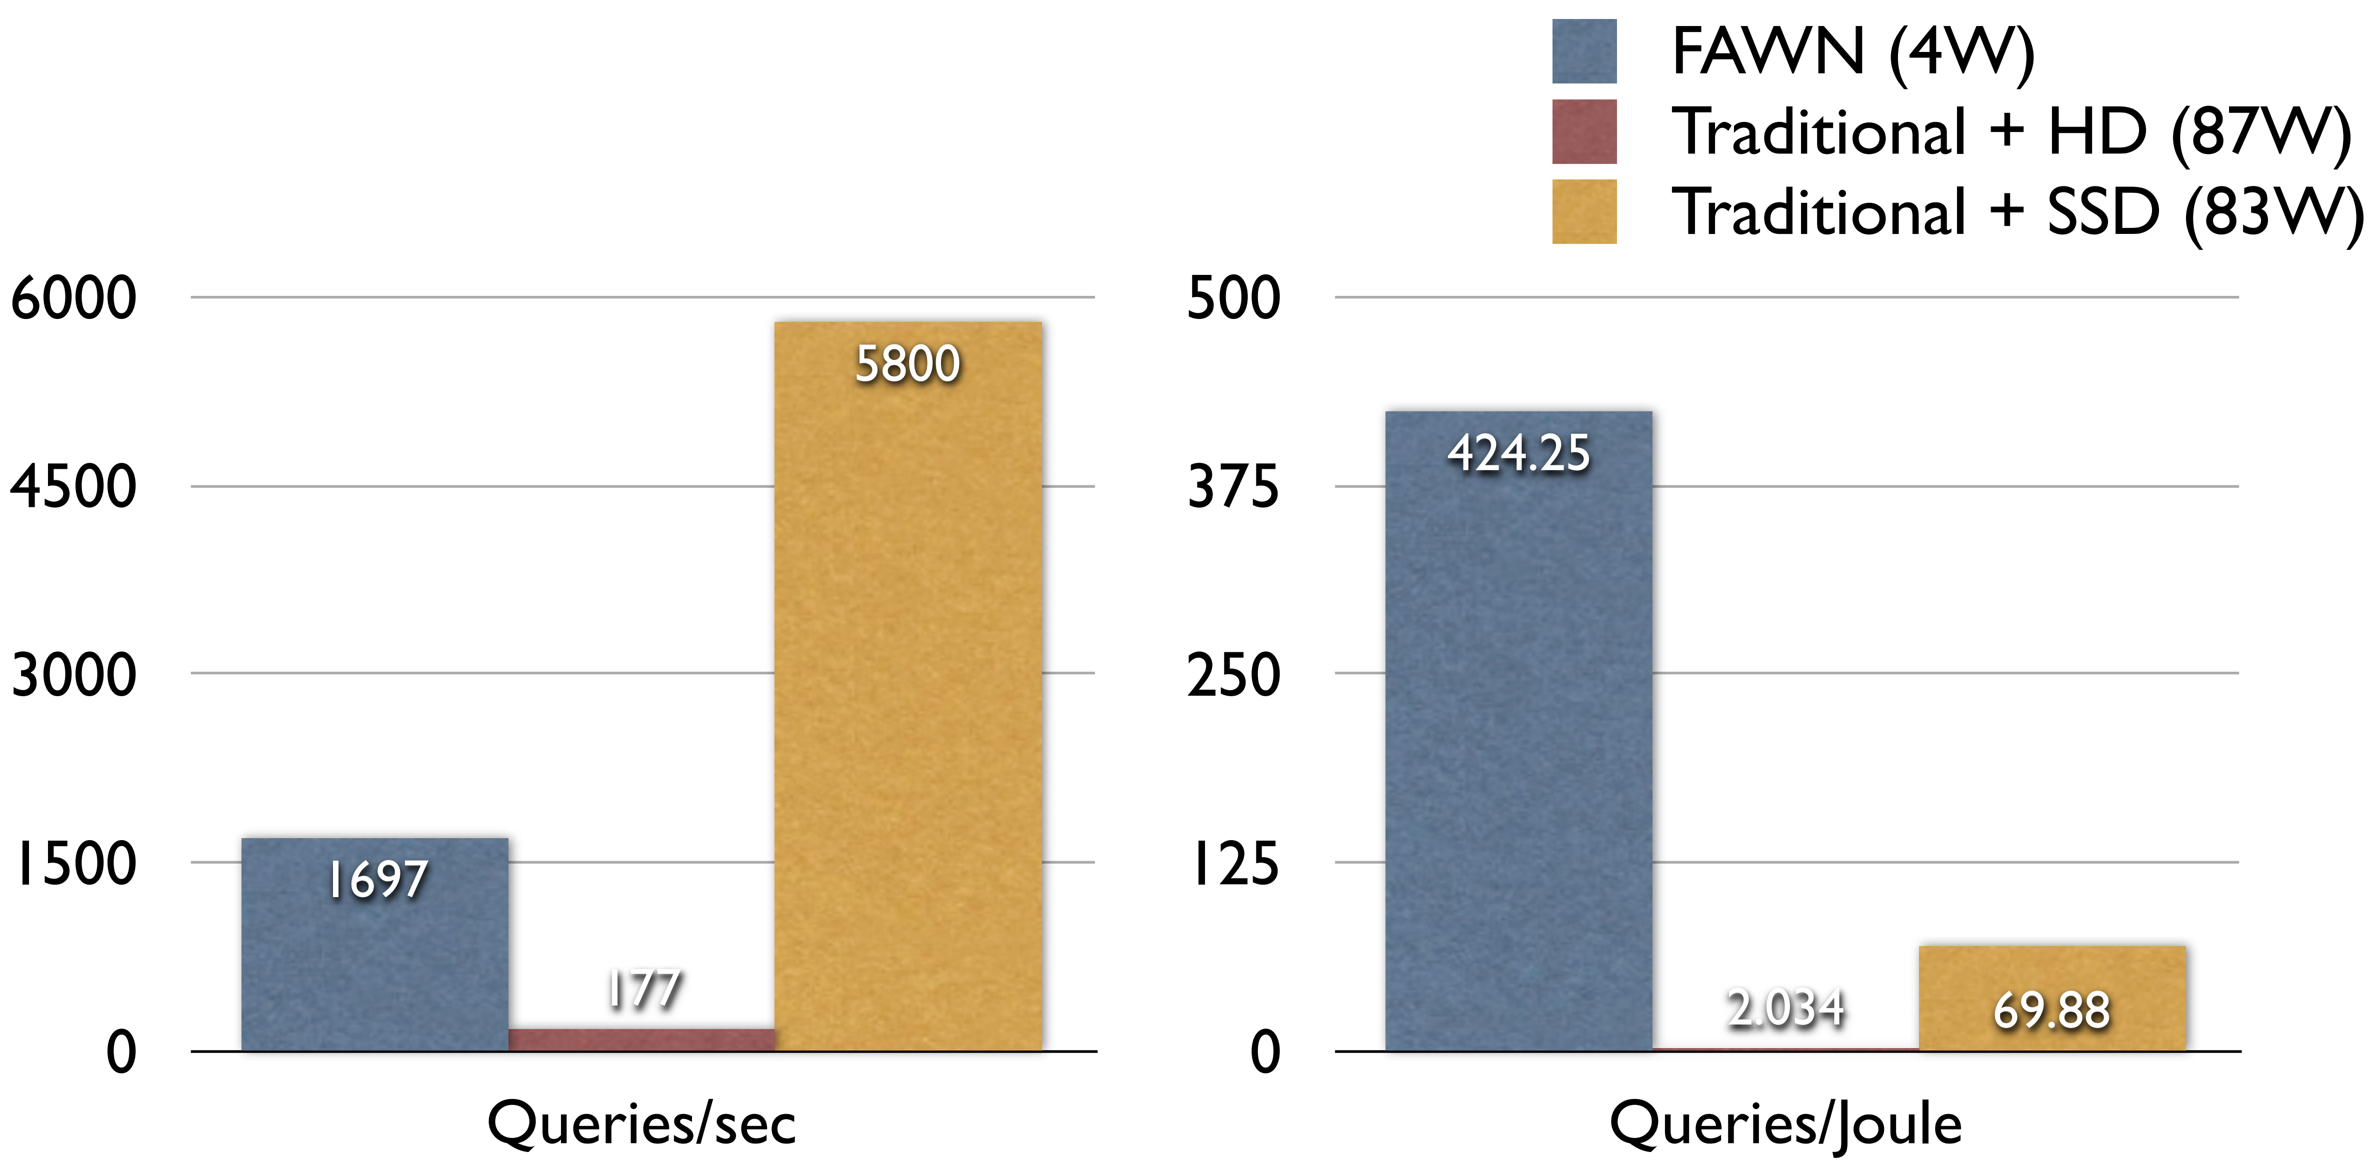

Performance

Efficiency

Random access workloads

FAWN is 6-200x more efficient than traditional systems

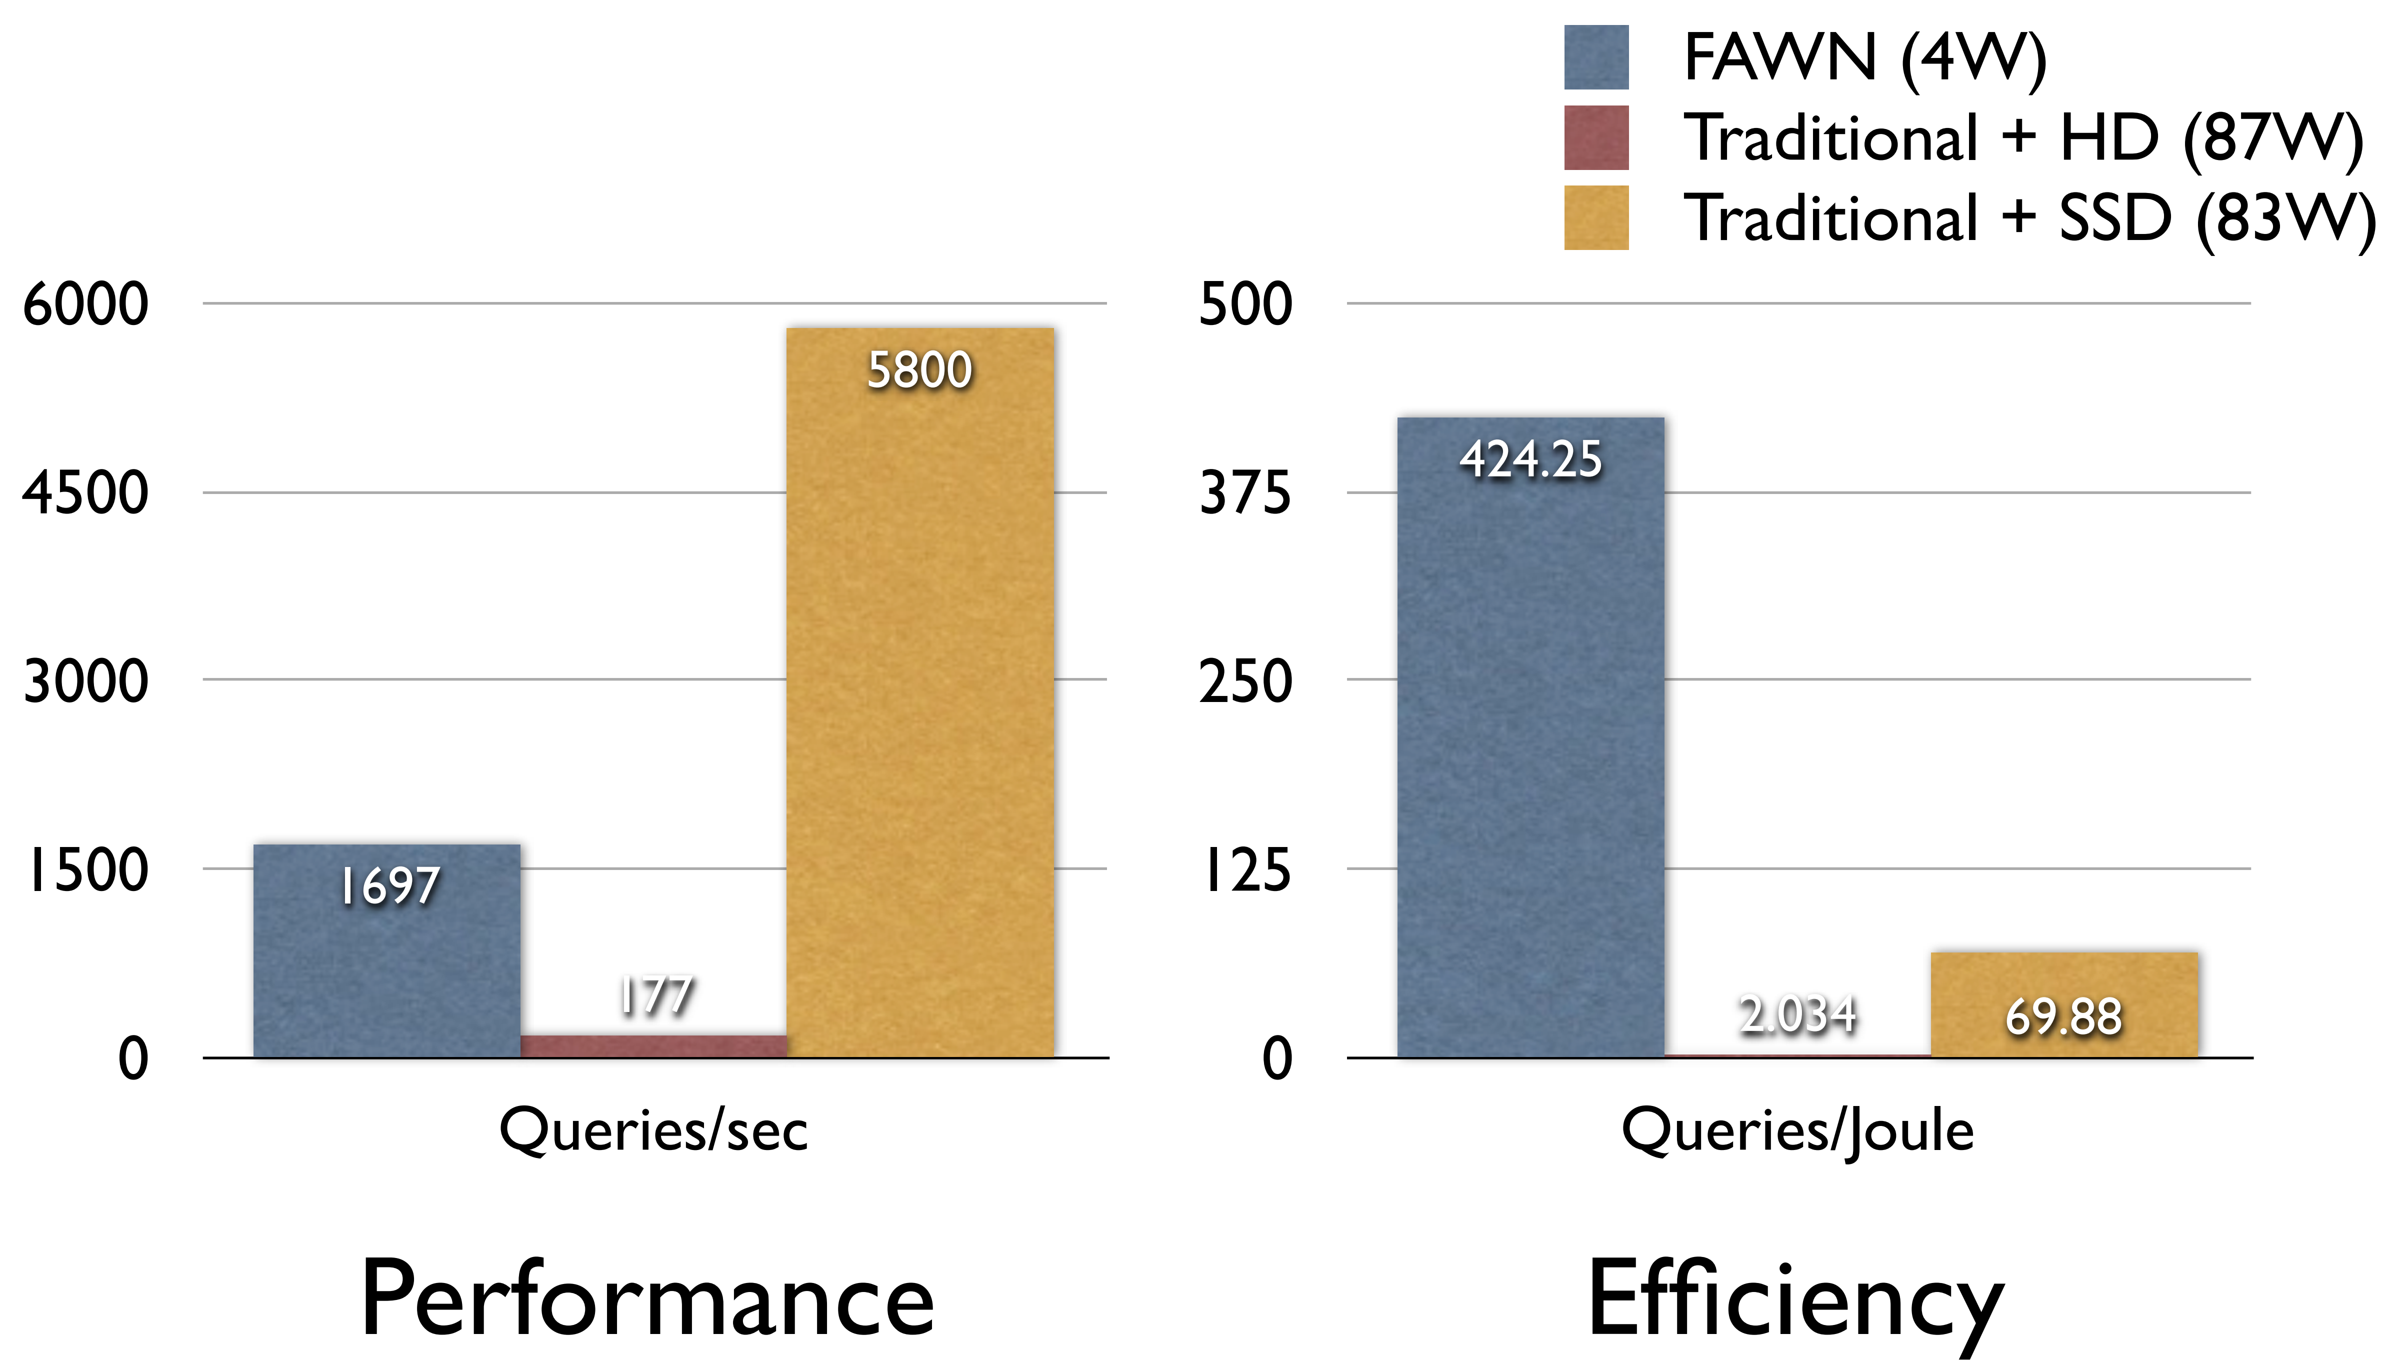

CPU-bound encryption

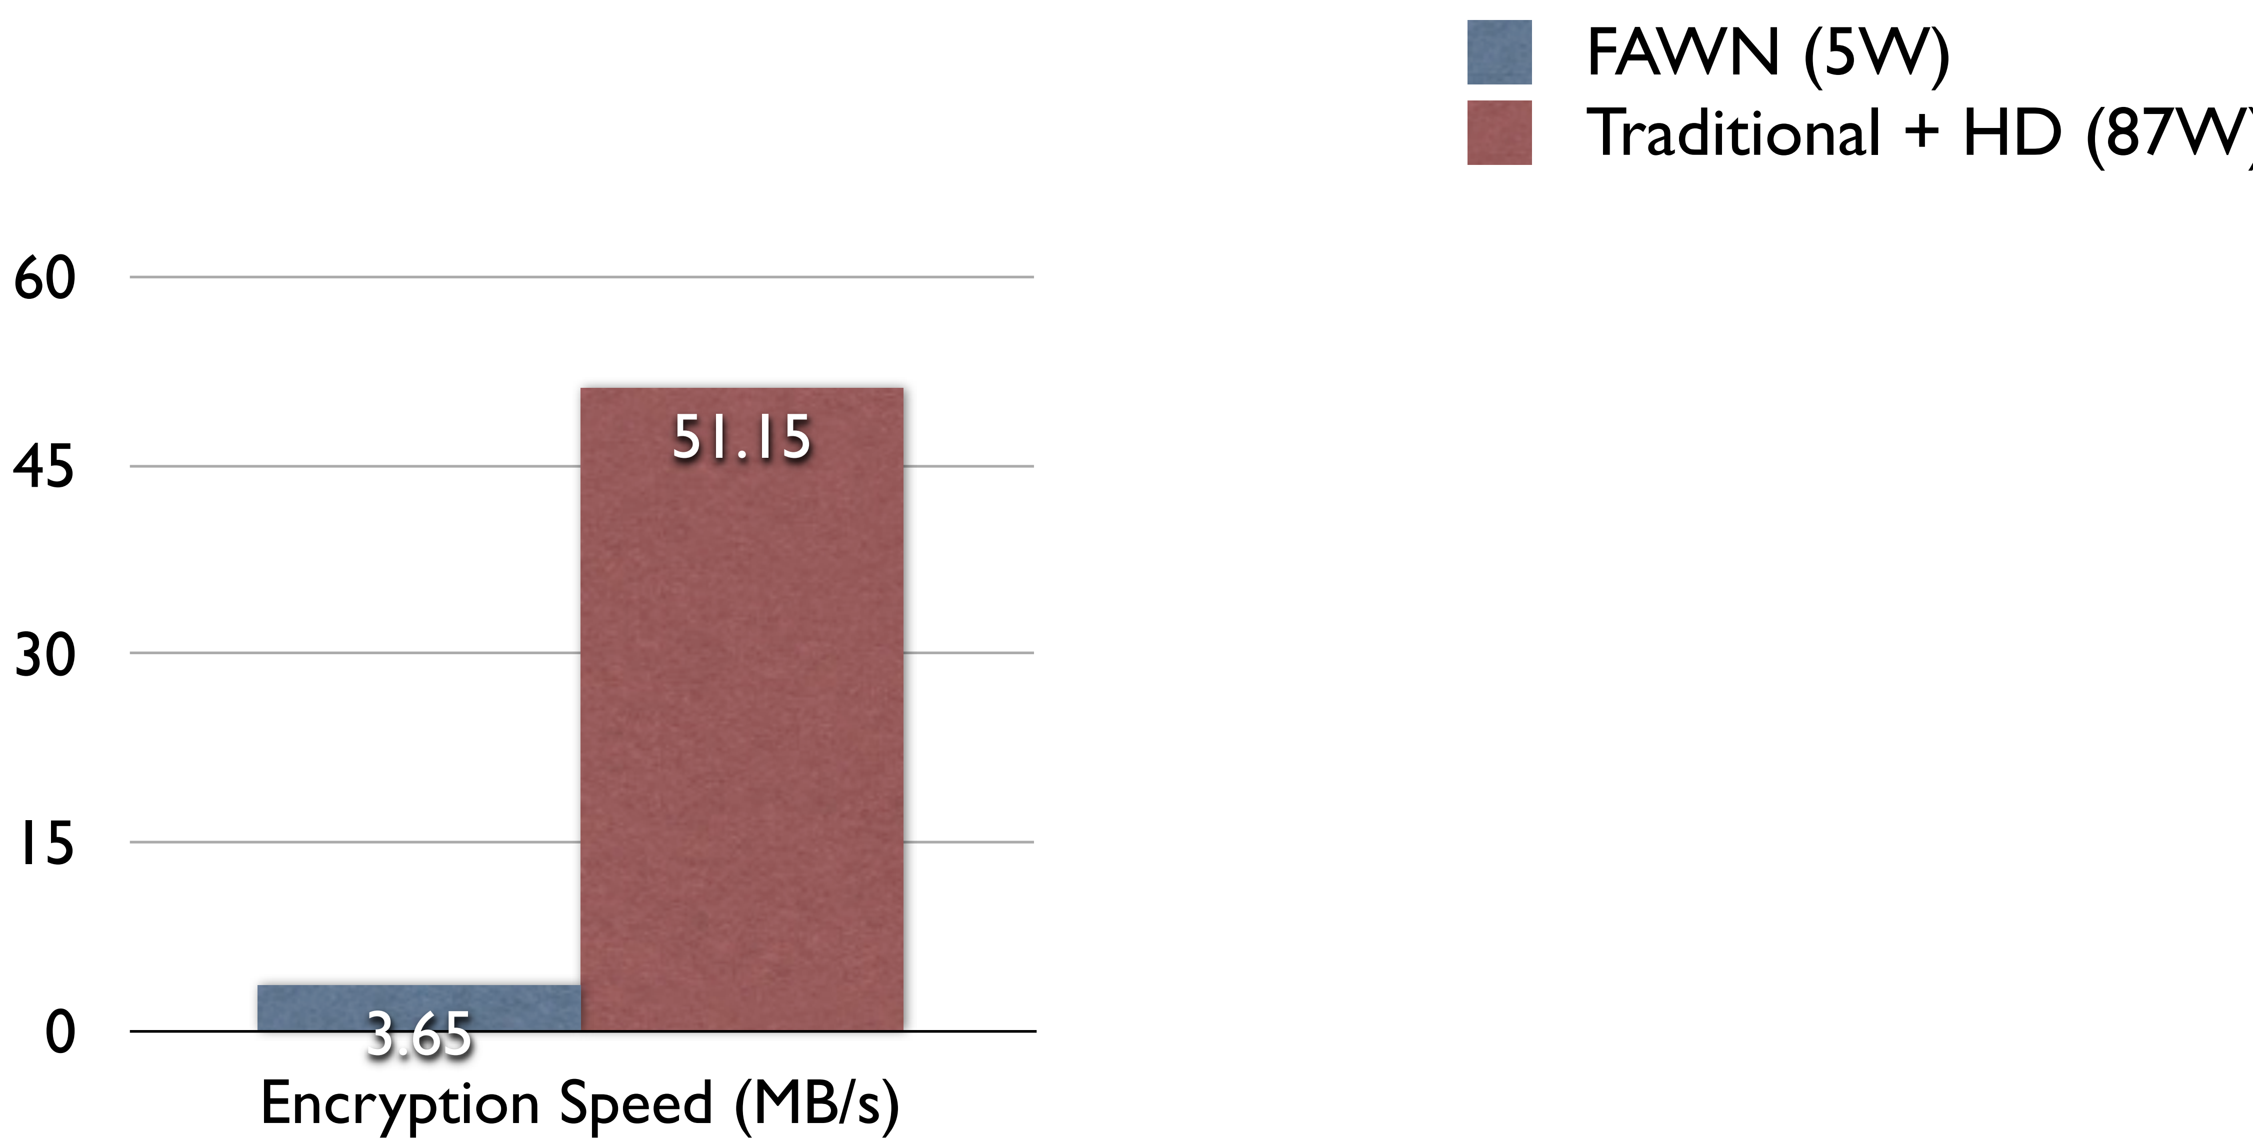

Performance

CPU-bound encryption

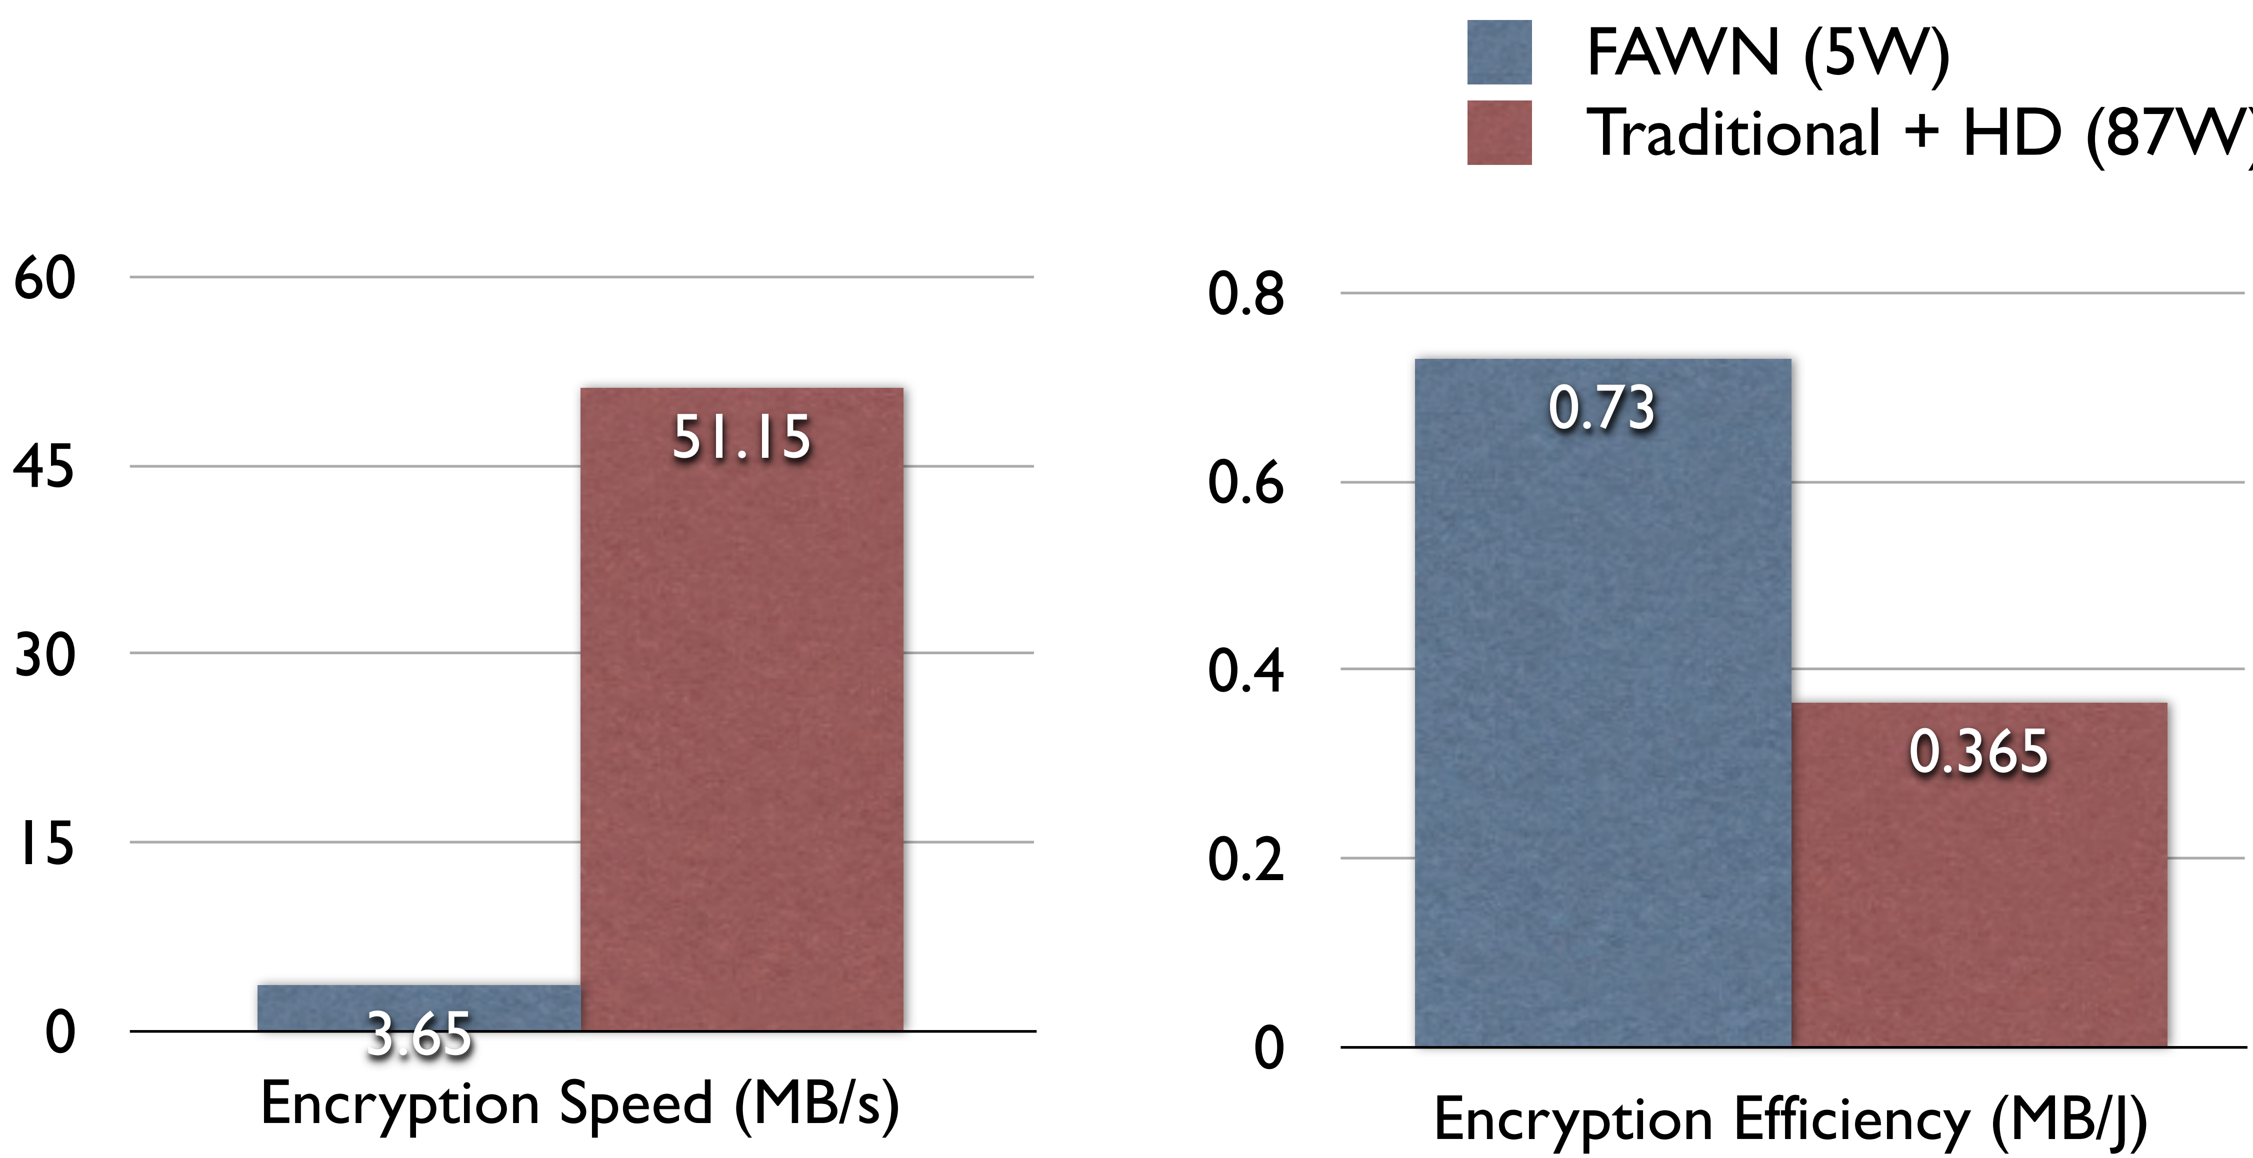

Performance

Efficiency

CPU-bound encryption

FAWN is 2x more efficient for CPU-bound operations!

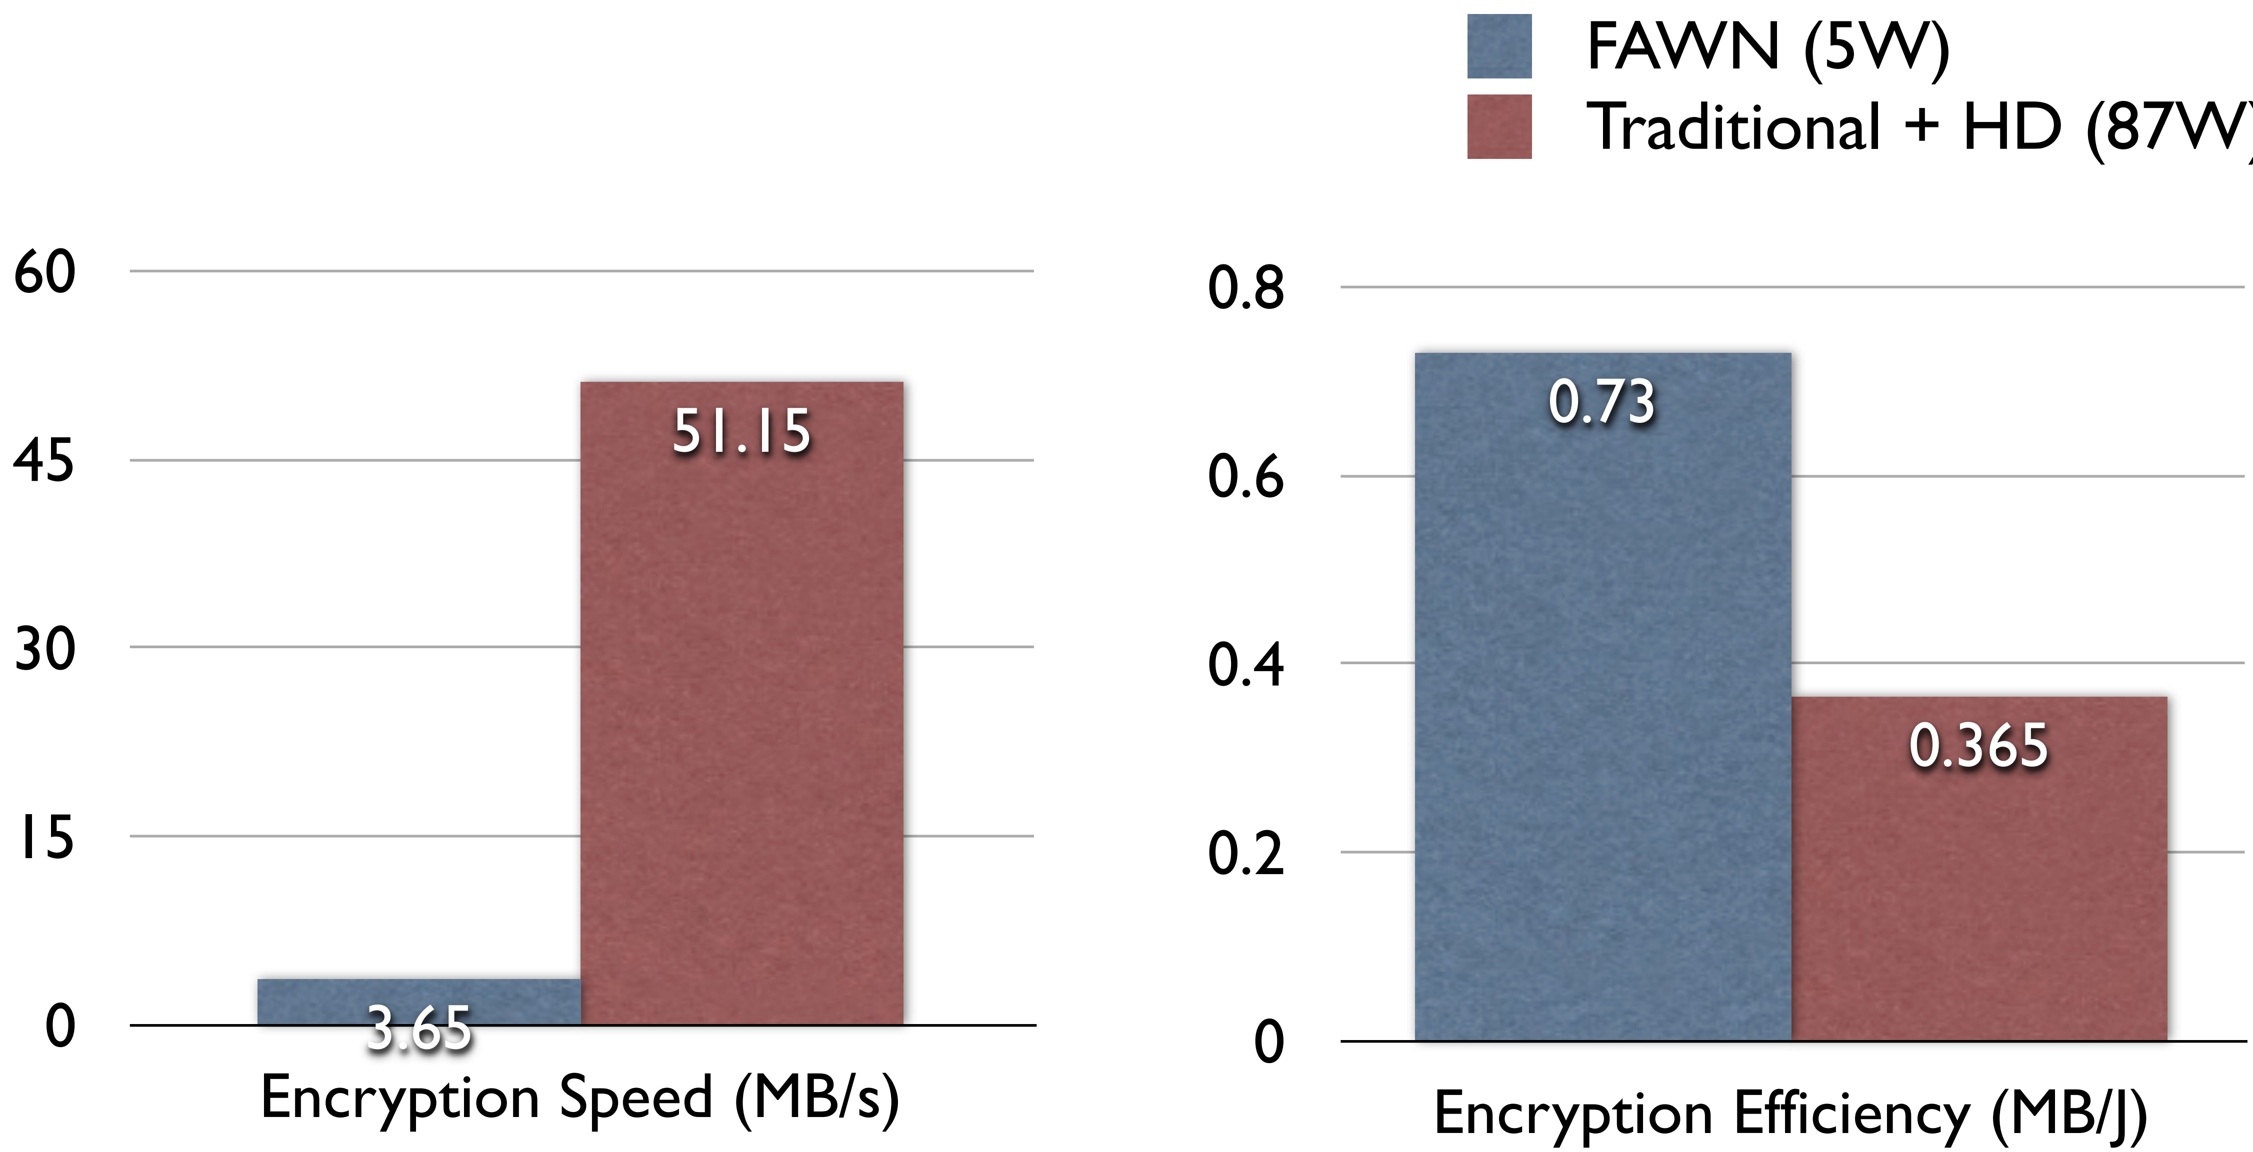

Performance   Efficiency

When to use FAWN for random access workloads?

- Total cost of ownership
  - Capital cost + 3 year power @ $0.10/kWh
- What is the cheapest architecture for serving random access workloads?
  - Traditional + {Disks, SSD, DRAM}?
  - FAWN + {Disks, SSD, DRAM}?

Architecture with lowest TCO

for random access workloads

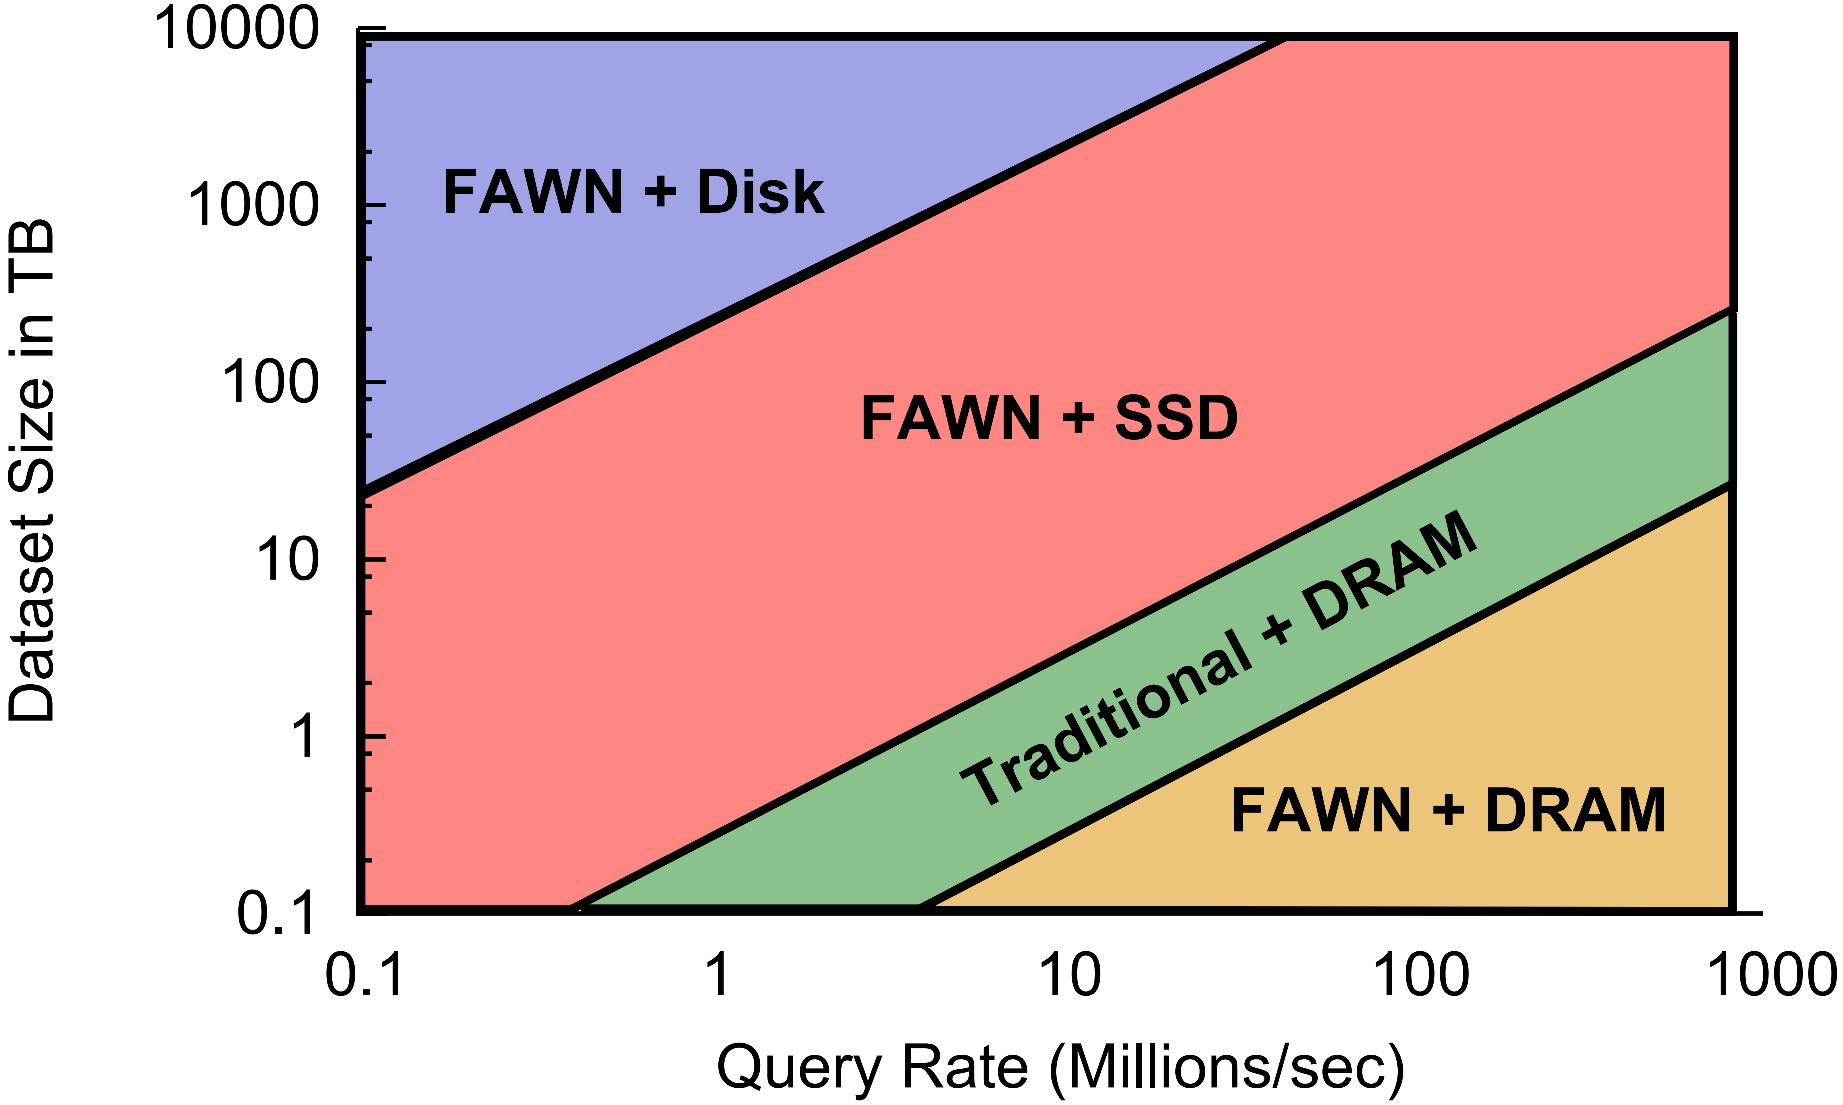

Ratio of query rate to dataset size informs storage technology

Architecture with lowest TCO
for random access workloads

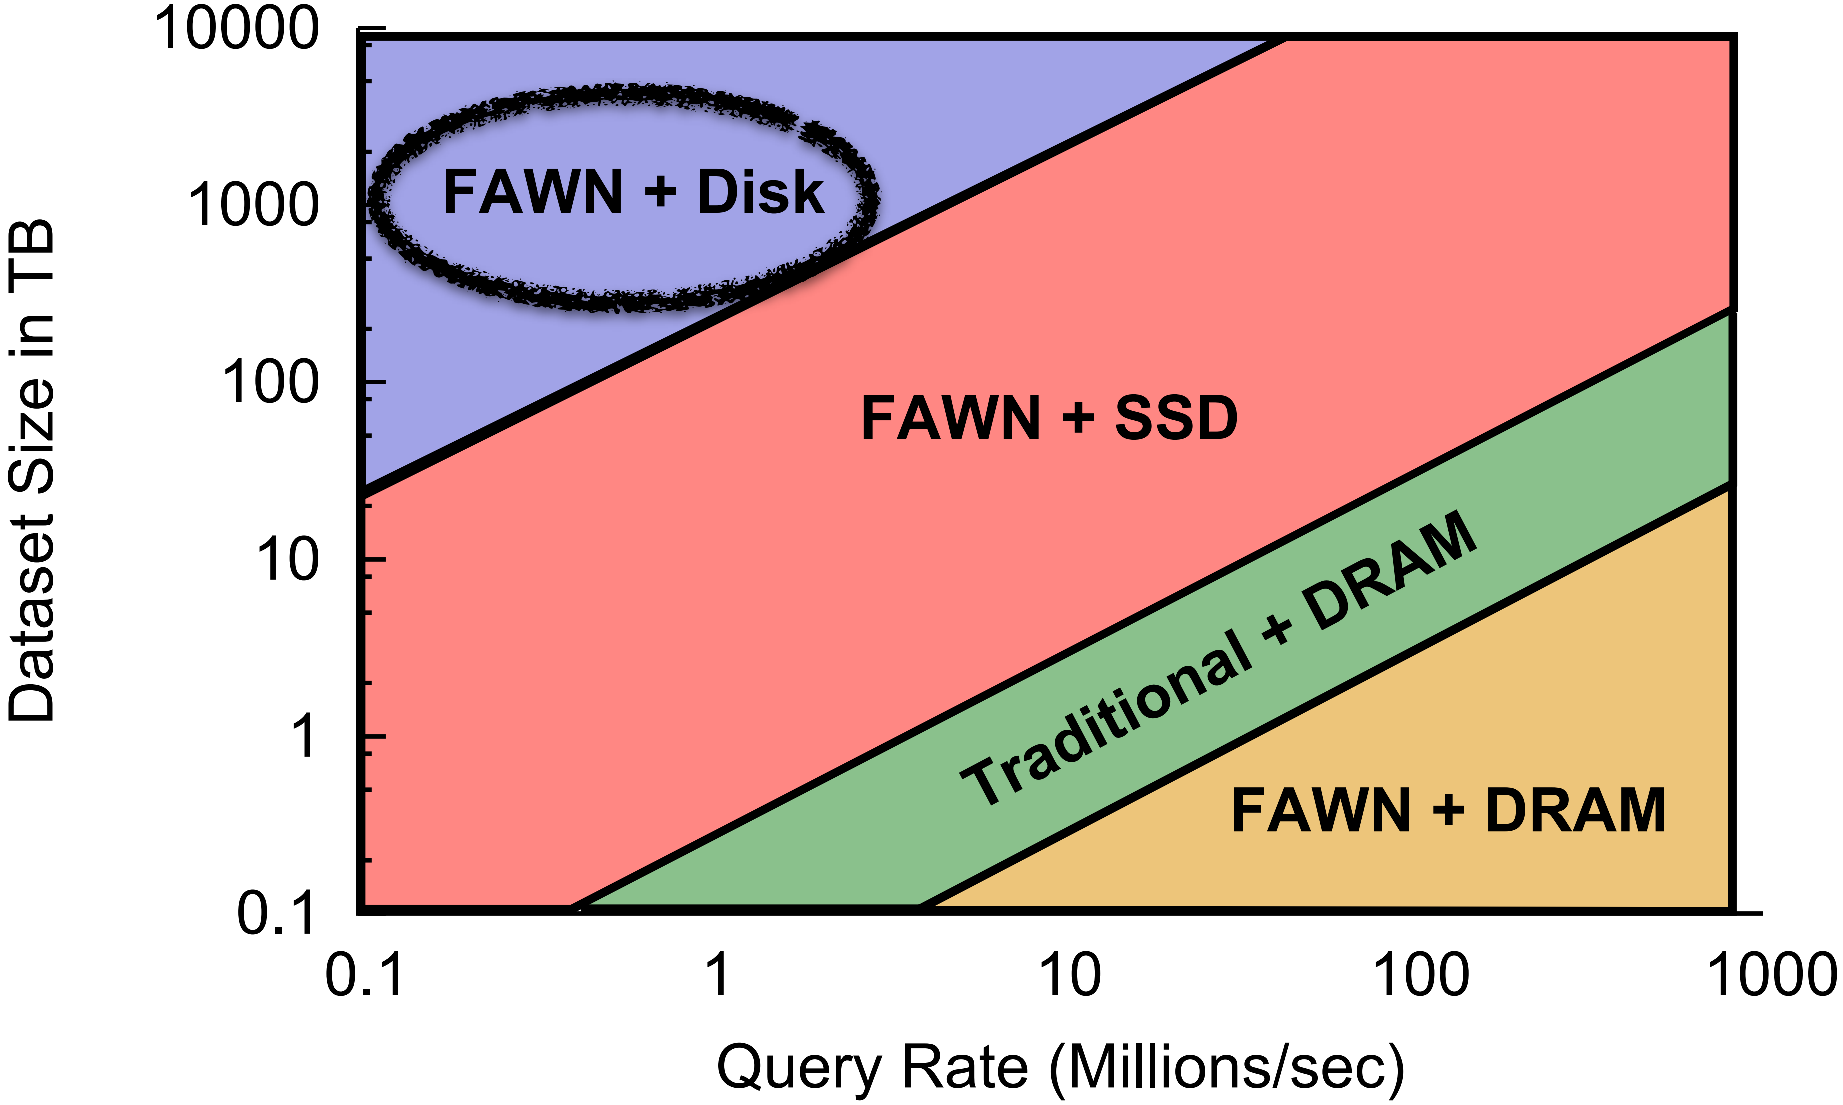

Ratio of query rate to dataset size informs storage technology

Architecture with lowest TCO

for random access workloads

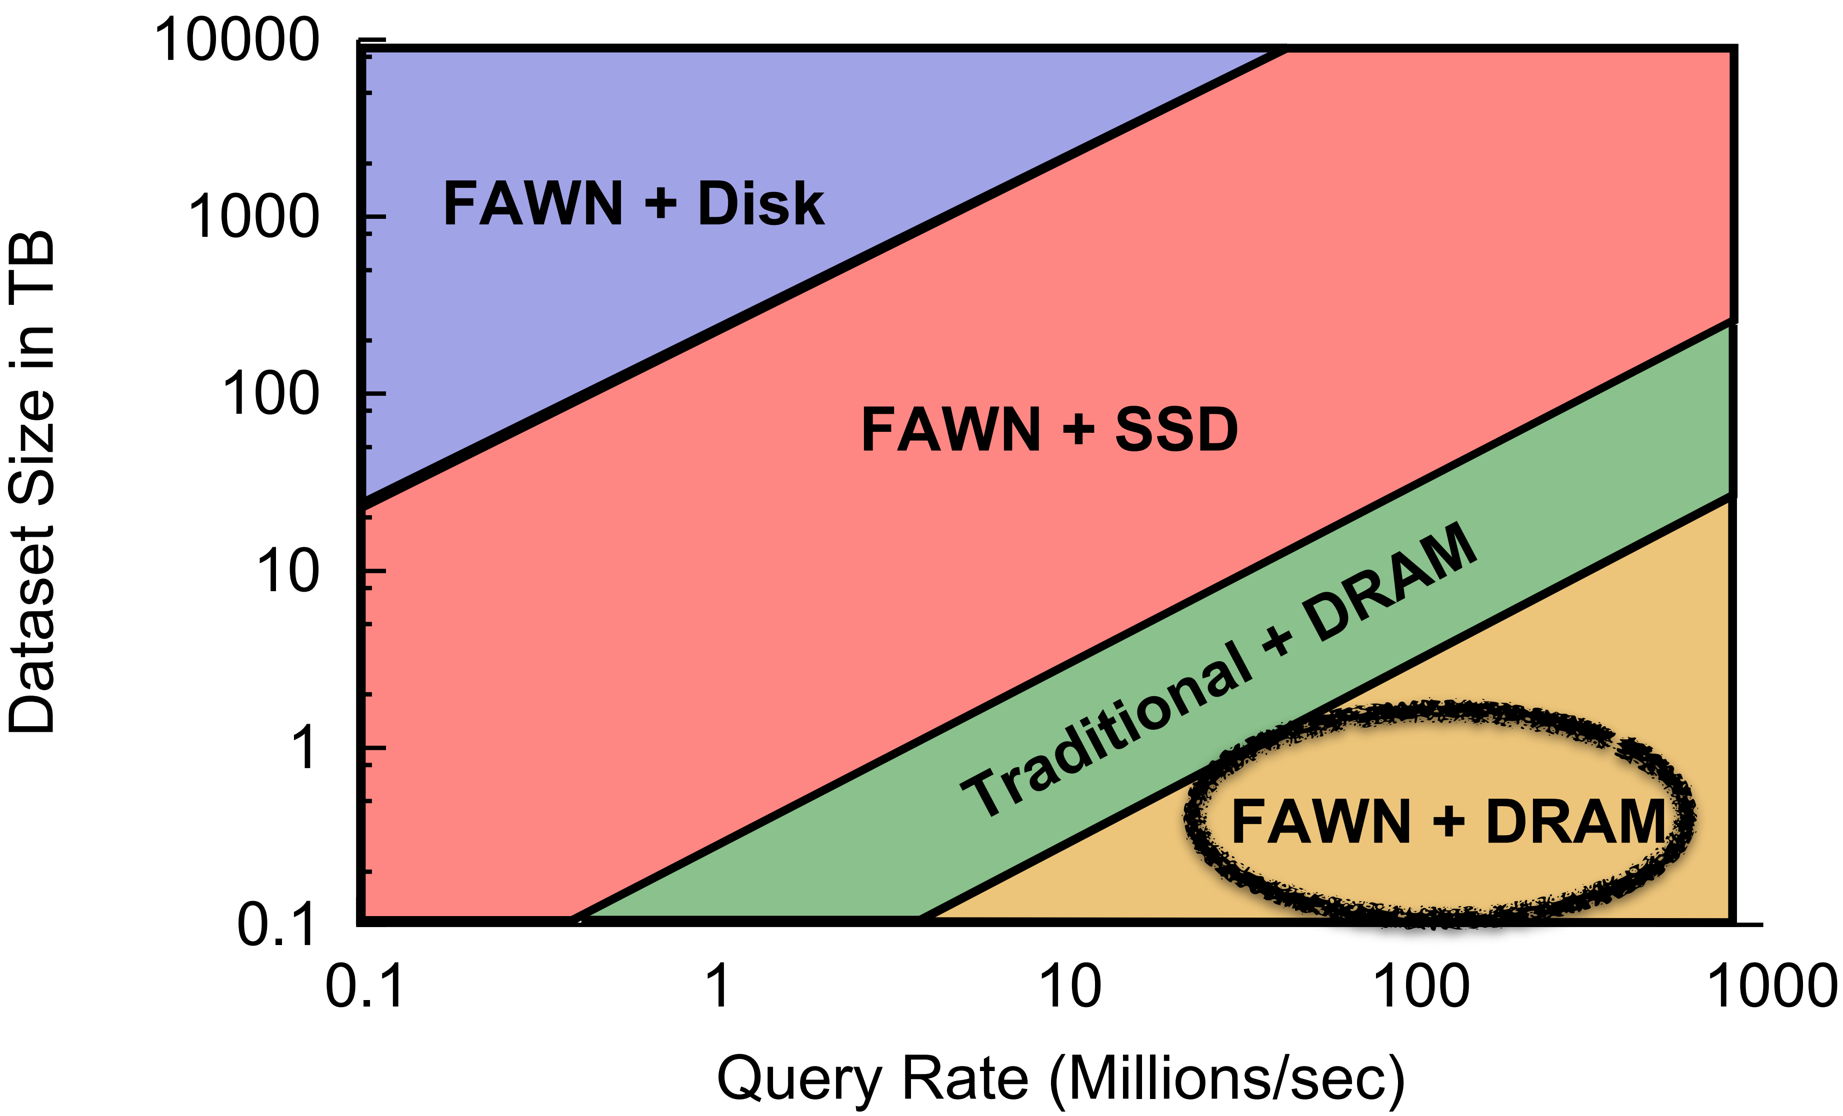

Ratio of query rate to dataset size informs storage technology

Architecture with lowest TCO

for random access workloads

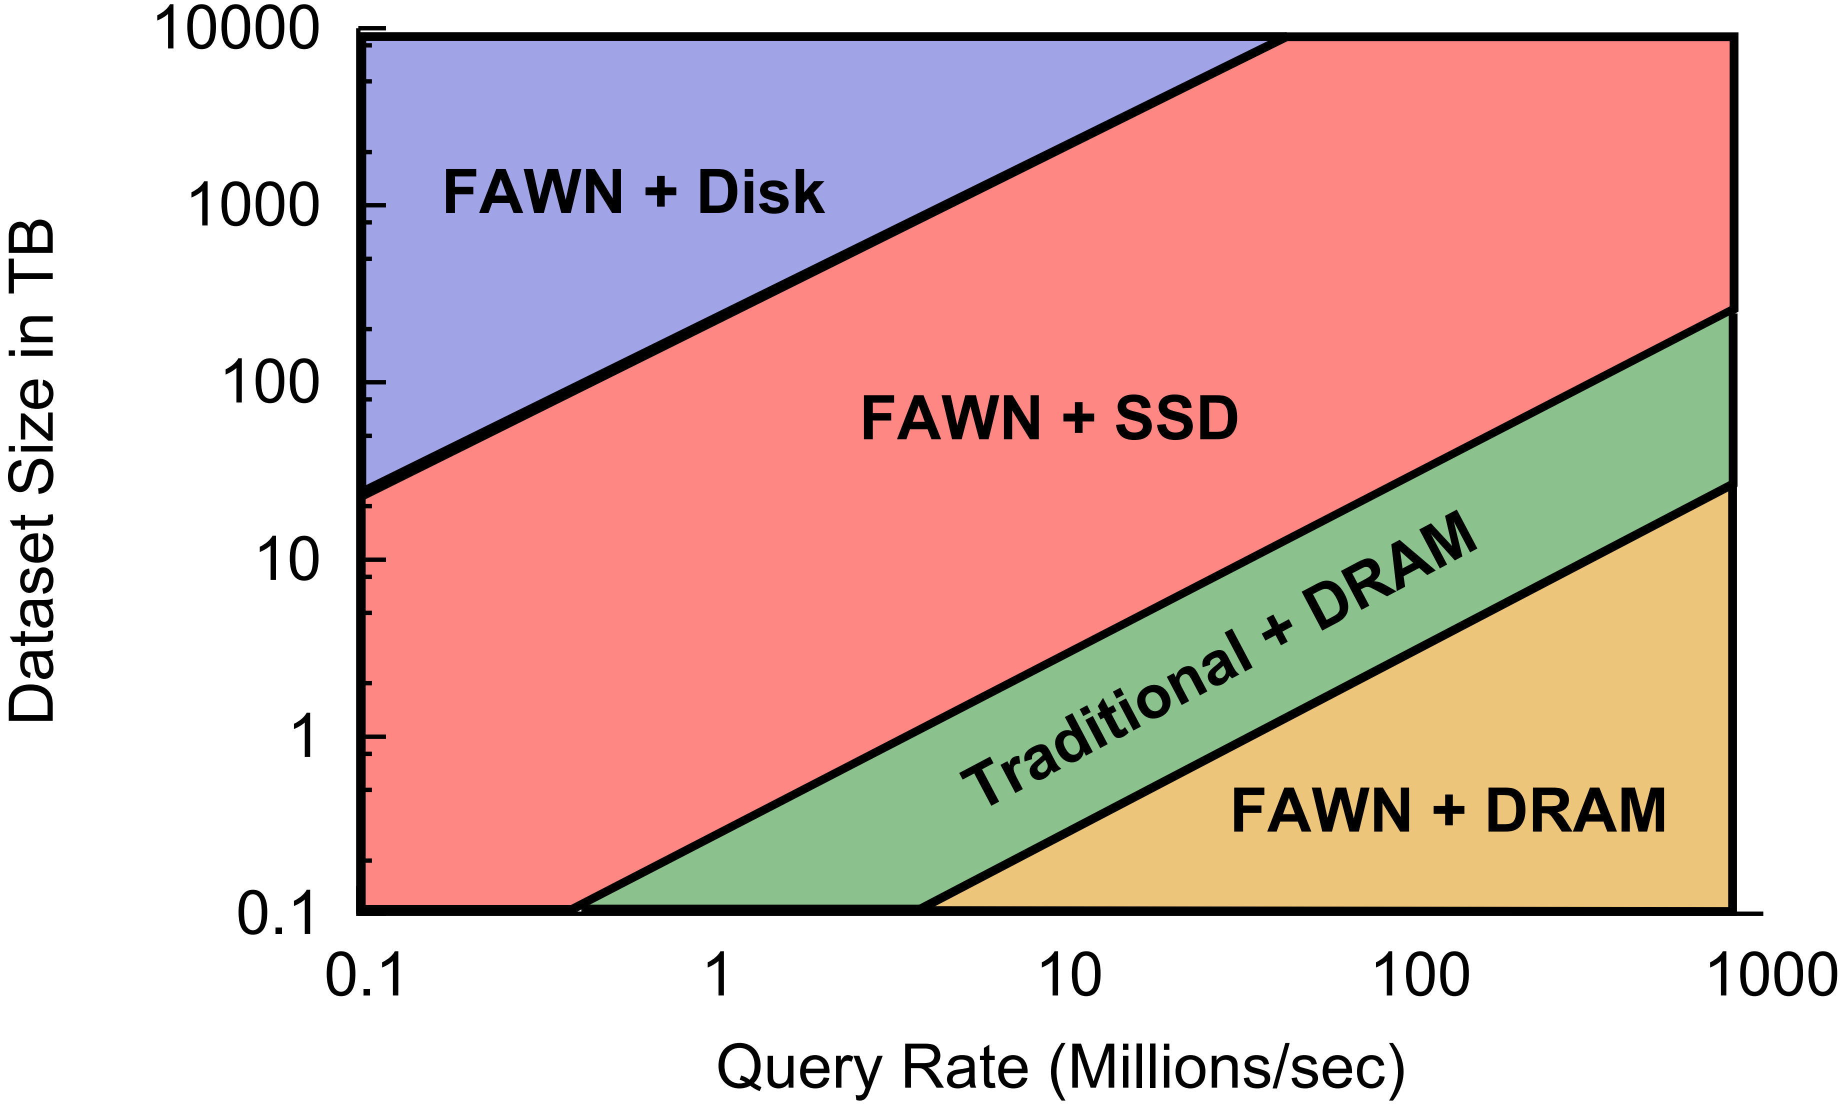

Ratio of query rate to dataset size informs storage technology

FAWN-based systems can provide lower cost per {GB, QueryRate}

Challenges

"Each decimal order of magnitude increase in parallelism requires a major redesign and rewrite of parallel code" - Kathy Yelick

- **Algorithms and Architectures at 10x scale**
  - Dealing with Amdahl's law
- **High performance using low performance nodes**
  - Today's software may not run out of the box
- **Manageability, failures, network design, power cost vs. engineering cost**

Conclusion

- **FAWN** improves the computational efficiency of datacenters
- Informed by fundamental system power trends
- Challenges: programming for 10x scale, running today's software on yesterday's machines...

Hot enough for industry

FEBRUARY 24, 2009

Microsoft trying out netbook processors in datacenters

Though a datacenter would require three times as many netbook processors, the power requirement would still be lower than that of typical server processors

MAY 06, 2009

Intel's Atom chip finding its way into servers

Super Micro and HP are building Atom chips, which were designed for netbooks and low-costs PCs, into server appliances

MAY 15, 2009

Dell puts low-power netbook chip in new server

Via Technologies' Nano processors will power Dell's ultra-light XS11-VX8 servers