Does Low-Power Design Imply Energy Efficiency for Data Centers?

David Meisner
meisner@umich.edu
Advanced Computer Architecture Lab
University of Michigan

Thomas F. Wenisch
twenisch@umich.edu
Advanced Computer Architecture Lab
University of Michigan

Abstract*—Data center efficiency has quickly become a first-class design goal. In response, many studies have emerged from the academic community and industry using low-power design to help improve the energy efficiency of server hardware. Generally, these proposals hold the assumption that low-power design is inherently better for energy efficiency; this preconception stems mostly from great success in the mobile space with building low-power, energy-efficient systems. We observe that unlike mobile devices, constraining a data center server to a low power budget is arbitrary and higher power design choices can be more energy efficient. We analyze the energy efficiency design space of past commercial server designs and find that high-power servers are generally more energy efficient than low-power ones. Furthermore, we evaluate building low- or high-power server clusters and find that the small increase in the cost of cooling high-powered servers is easily outweighed by their greater efficiency.

keywords - servers, data centers

I. Introduction

Energy efficiency has quickly become a first-class design goal for data centers [5], motivating a concerted effort by government, academia, and industry to investigate potentially transformative strategies to improve efficiency. A growing chorus of studies in academia [3], [16], [17], [21] and industry [11] espouse *low-power design* to improve the energy efficiency of server hardware. The motivation for low-power design is two-fold. First, years of engineering effort have been expended optimizing mobile/embedded-class designs for energy-constrained deployments. Second, these designs omit complex features (e.g., massive caches and wide out-of-order issue) that conventional wisdom suggests provide diminishing performance returns, providing an inherent efficiency advantage.

We observe two major caveats to adopting low-power components in the data center. First, scaling software to run on weaker systems presents a significant challenge as it implies distributing the work of one high-power server across several low-power servers. The greater demand for parallelism makes scale-out more difficult; systems with finite parallelism run into Amdahl bottlenecks. Furthermore, though comparable throughput may be achievable, latency may suffer greatly as noted by several authors [18], [14], [16].

We discuss a second caveat that recent literature has not addressed: limiting the system power envelope might *decrease* the energy efficiency of the hardware itself. This paper argues that constraining a system's power budget (e.g., by selecting low-power components) eliminates points in its design space that may be more efficient. In essence, we call into question the notion that low power correlates with energy efficiency in the data center context. Rather, we argue that designers often pay an efficiency price in exchange for a lower peak power. While mobile platforms have inherent peak power budgets (e.g., driven by form-factor), for servers there is no inherent advantage to using a low-power design. In this study, we examine the question of whether low power correlates with energy efficiency for data center servers from two aspects. First, we investigate the recent

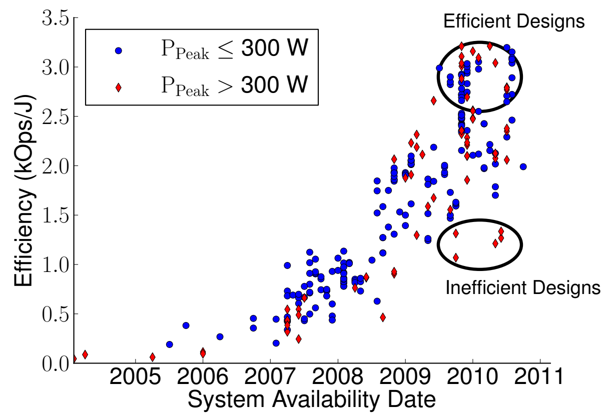

Fig. 1: Historical trend in server efficiency. For both high-power (> 300 W) and low-power (≤ 300 W) designs, efficiency is generally increasing, however there have been highly efficient and inefficient designs recently in both classes. Data from [1].

historical record by leveraging the rich set of power/performance data available in the audited SPECpower2008 results [1]. Using the historical data, we investigate what factors of server design, if any, have been correlated to energy efficiency. Our first central observation is summarized in Figure 1, which shows efficiency results against system availability date. We classify a design as low-power if its peak power is less than 300W and high-power if it is greater. Although it is clear that there is a general trend toward greater efficiency in both classes, even in the latest submissions, efficient and inefficient designs appear in both. Whereas peak power is not particularly correlated with efficiency, we find instead that (1) energy efficiency at peak utilization, (2) wide dynamic range (equivalently, low idle power), (3) peak performance and, to a lesser degree, (4) the number of cores per socket correlate strongly.

A second strong motivation for low peak power in the data center is the higher cost of cooling high-power systems. It is well-known that cooling costs grow rapidly with power density. Hence, intuitively, one might expect power-dense servers to exact an unpalatable increase in cooling energy. We investigate the question: to make up for its disadvantage in cooling requirement, how much more energy efficient must a high-power server be to break even on total energy with a low-power server? Using the SPECpower data, we build on a previously published model of data center thermodynamics [24] to contrast cooling costs of high-power and low-power 1U servers. Our second central result is that, to build an iso-power and iso-throughput data center, high-power-density servers must be only a *few percent* more

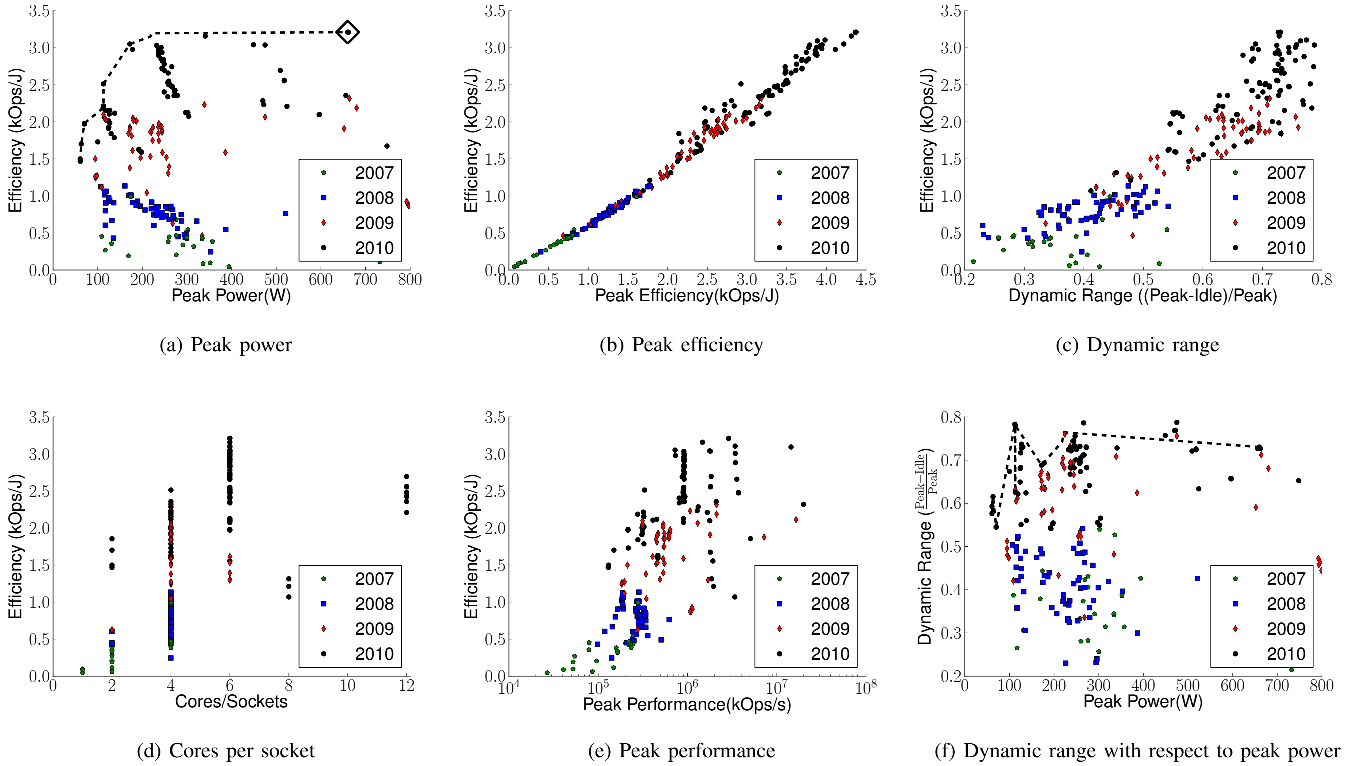

(a) Peak power (b) Peak efficiency (c) Dynamic range

(d) Cores per socket (e) Peak performance (f) Dynamic range with respect to peak power

Fig. 2: Server Efficiency Design Space. Each point represents a unique server design submitted to SPECPower2008. Subfigure a) demonstrates that the peak power of a server has little to no correlation to average efficiency. Note that the most efficient server design (highlighted with a diamond) is also one of the highest-peak-power designs. Subfigure b) demonstrates that server peak-efficiency strongly correlates with average efficiency. Similarly, subfigure c) shows that the dynamic range (i.e., energy-proportionality) of a server also greatly affects average efficiency, but correlates less strongly than peak-efficiency. In subfigure d), one can see that, in general, adding more cores to a system increases efficiency, but the range of efficiencies for a given core integration is large and systems with less cores can easily be more efficient than those with more. The peak performance of a system also correlates well with efficiency as seen in subfigure e), but less strongly than peak efficiency or dynamic range. Subfigure f) illustrates the relative lack of correlation between peak power and dynamic range. We annotate this graph with a line connecting all the pareto optimal designs from subfigure a). Data collated from [1].

energy-efficient than the low-power-density alternatives. Therefore, the cooling-cost objection to high-power servers is a red herring.

II. Related Work

Many studies have investigated the energy efficiency of processors through design space exploration [4], [20]. These studies provide insight into how to design the most efficient processor given a performance goal; however, these studies do not consider the total power cost of a server platform (e.g., memory, disks, etc.) nor the implications of deploying these designs in a data center.

Other studies have evaluated the efficiency of a specific workload on differing server platforms [16], [27], [21], [17], [3]. These studies conclude that low-power designs may sometimes lead to latency degradation [16] and that the highest-performing server may be the most energy efficient [27]. These studies provided the intuition to motivate our own work; however, they each look at only a handful of server designs and do not consider total data center power.

Several proposals present schemes for reducing data center infrastructure costs. Power capping schemes [9] oversubscribe data center infrastructure and PowerRouting [25] corrects dynamic power distribution imbalances to amortize capital costs. Cooling power can be minimized through dynamic management [4], [23]. Such proposals are orthogonal to our investigation of high- vs. low-power servers.

III. Server-level Efficiency

To better understand the factors affecting server-level efficiency, we analyze the recent history of the server design space based on results submitted to SPECPower [1]. This data set includes over 200 server configurations that have been finely tuned by their vendor to maximize both performance and minimize power consumption. Although this data set evaluates only one benchmark, there is no other data set which provides insight into hundreds of *real* carefully-optimized server designs.

Low-power and high-power. We classify servers as *low-power* or *high-power* based on their relative *maximum instantaneous* power consumption (peak power). Hence we say a system is low-power if its instantaneous power draw is small compared to other design alternatives.

Energy efficiency. We define the *energy-efficiency* of a system as the work done per unit energy. This metric is a function of the *average* power draw of a system. Over a long period, this can be expressed as:

$$\text{Energy Efficiency} = \frac{\text{Throughput}}{P_{\text{Avg.}}} \tag{1}$$

SPECPower provides an overall efficiency metric in terms of operations per second per watt, or (equivalently) operations per joule. This average efficiency is determined by running the server across

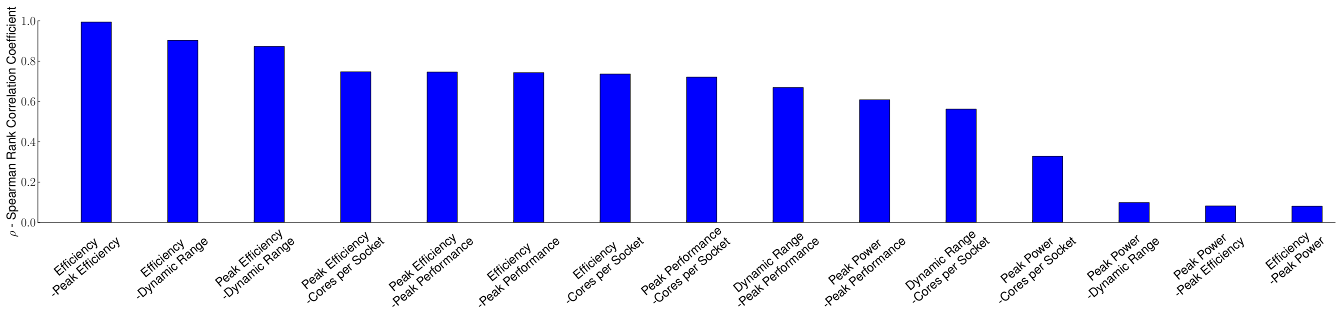

Fig. 3: Correlation between design features. This figure provides the spearman rank correlation coefficient of each pair of design features we examine. This metric measures monotonic covariance between two variables, but does not rely on linearity. The pairs are ranked from most- to least-correlated. This figure quantifies the trends in Figure 2. Peak efficiency, dynamic range and cores per system kWatt have the strongest correlation with efficiency. On the other hand, peak and average efficiency correlate little with peak power.

its entire range of utilization; determining efficiency across the entire utilization spectrum is critical because servers do not spend most of their time at peak utilization [5]. We correlate efficiency against a number of metrics to determine which factors affect efficiency and which are irrelevant.

A. *Analysis*

We choose a feature set from SPECPower and use scatter plots to present graphically how the features correlate in Figure 2. We have selected these particular features to (1) demonstrate our main thesis, that peak power is uncorrelated to efficiency; and (2) to reveal other interesting factors that do seem to correlate well to efficiency. Each design is classified by year based on the disclosed system availability date; note that for simplicity the year "2007" also includes a few pre-2007 designs.

Peak power. Figure 2(a) illustrates our main result: low-power servers do not necessarily provide greater efficiency than their high-power counterparts. Many of the low-power designs have far worse efficiency than the high-power designs. In fact, the most efficient server is one of the highest power designs (highlighted with a diamond).

Peak efficiency. Perhaps unsurprisingly, a server's energy-efficiency at peak (efficiency at 100% utilization) correlates more strongly against its *average* efficiency than any other factor we examined. Figure 2(b) shows that servers with a higher peak efficiency have greater average efficiency. The average efficiency is always less than the peak efficiency; modern servers are only maximally efficient at 100%. Intuitively, the efficiency of these systems has been increasing over time.

Dynamic range. Next, we show that the dynamic range (peak power minus idle power normalized by peak power) of a server correlates well against average efficiency. A system's dynamic range is a good proxy for its energy-proportionality [5]. Figure 2(c) illustrates that servers with a larger dynamic range are generally more efficient, but that the correlation is not as strong as for peak efficiency. Some designs with a smaller dynamic range have the same efficiency as those with larger dynamic range; moreover some designs with the same dynamic range have significantly different efficiencies. Interestingly, the dynamic range of servers seems to be increasing over time.

Multicore integration. A recent trend is to use multicore designs to maximize throughput per watt. Figure 2(d) shows the impact of mutlicore integration (i.e., number of cores per socket) on server efficiency. The overall trend is that more cores per socket increases efficiency; however, the range of efficiency at each of these design points is large. Many four core designs are more efficient than 6 or 8 cores alternatives.

Peak performance. We observe that servers that have higher peak performance tend to be more efficient as seen in Figure 2(e). This observation has also been noted with servers running database work-loads [27]. It is possible to have systems with orders of magnitude difference in performance and the same efficiency, the general trend suggests that building a faster server yields a more efficient system. Over time, the peak performance of these systems has been increasing as has the efficiency.

Dynamic range with respect to peak power. Finally, we would like to know if the peak power of a server has an effect on its dynamic range. For example, do systems with more powerful processors increase the dynamic range because these components tend to be more energy proportional than DRAM and disk? Subfigure f) shows that this hypothesis is not supported by the data. In fact, we see that peak power has little correlation with dynamic range. We annotate a line connecting all the pareto optimal designs chosen from Figure 2(a); designs that are pareto-optimal with respect to peak power and efficiency tend also to have high dynamic range.

B. *Feature Correlation*

Figure 3 shows the spearman rank correlation coefficient between all pairs of the features in the design space we consider. The spearman rank coefficient measures monotonic covariance between two variables, but does not rely on linearity. That is, a ρ = 1.0 indicates that two features move in the same direction monotonically; whereas ρ = 0 indicates no correlation between the features. We rank these pairs from highest to lowest correlation.

We highlight several interesting data points. Foremost, the results indicate that peak power and (peak) efficiency are the *least correlated* among the factors we investigate, confirming our main thesis. Contrary to our expectations, peak power and dynamic range are also relatively uncorrelated—we had expected servers with higher absolute peak power to have generally larger dynamic range, as there is greater scope for power variability. Several of the correlations confirm that the industry trend towards multicore is improving energy efficiency, at least for throughput-oriented workloads like SPECpower. Finally, the historical record indicates that energy efficiency at peak is the single most important factor in determining average efficiency.

IV. Data Center-level Efficiency

We turn, now, to the second half of our study, where we investigate the implications of high-power servers on data-center-level energy

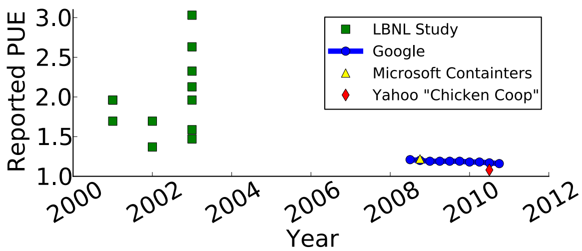

Fig. 4: Industry has achieved drastic improvements in PUE. This figure shows the trend in PUE as reported by various institutions over time [19], [6], [10], [13]. Most of the gains in efficiency can be attributed to better cooling containment systems (e.g, simple curtain systems, containerized data centers, hot-aisle containment). Note that there is a noticeable gap in published data between 2002 and 2008. This gap is likely due to innovations in PUE remaining confidential because they gave a significant competitive advantage.

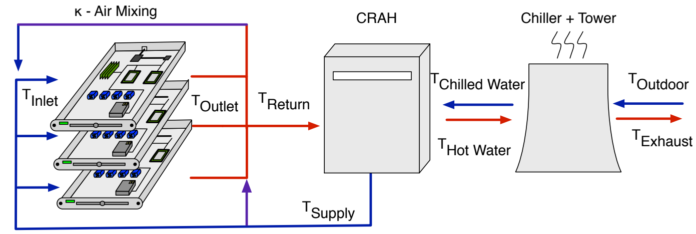

Fig. 5: Data Center Cooling Model. Racks of servers are cooled by CRAHs. We assume a fixed, maximum server outlet temperature set by reliability standards across all servers. CRAHs cool air to T_{Supply} which is optimized to minimize cooling costs between fans and CRAHs. The constant κ determines the amount of mixing between cold and hot air. This value is directly influenced by the PUE of the data center (higher PUE implies higher κ).

efficiency. One metric for measuring data center efficiency is the *Power Usage Effectiveness* (PUE), defined as:

$$\text{PUE} = \frac{\text{Power Delivered to Data Center}}{\text{Power Delivered to Servers}} \tag{2}$$

While the merit of this metric has been debated [12], it is widely used and reported in industry and serves as a useful approximation of the efficiency of data center power infrastructure and cooling.

A. Historical Trends in Cooling

Initially, when data center efficiency first emerged as an important design consideration, historical data suggested industry-average PUE was around 2.0. This result implied that as much energy is spent in power delivery and cooling infrastructure as was consumed by the servers themselves. Recently however, industry leaders have drastically improved the PUE of their latest installations. Figure 4 presents the PUEs disclosed by various institutions over time [19]. Within the last year, reports of PUE below 1.1 have been commonplace. Such a drastic shift in infrastructure efficiency changes ones' approach to data center design and reduces the importance of minimizing heat generation. For example, even very recent studies, such as [2] assume PUEs of 2.0; this assumption grossly overstates the energy overhead of cooling in modern data center design. As of the last three months of 2010, of the ten data center PUEs disclosed by Google, all but one is under 1.2 [10]. This drastic shift also illustrates the difficulty of predicting industry trends, the US EPA's 2007 study of data center efficiency predicted that, by this year (2011), PUE would reach only 1.4 with wide-scale deployment of liquid cooling. Instead, improved containment (separating hot and cold air), air-side economization (using outside air) and high-temperature operation (raising servers operating temperature) have driven gains.

B. Power and Cooling Model

We now construct a power and cooling model to evaluate the impact of using high- or low-power design to build a server cluster. Our model builds upon the work in [24]; for simplicity we assume homogeneous servers. We explain each of the components of the model needed to provision a cluster and determine its power consumption.

Figure 5 illustrates the heat flow in a typical data center. Servers are cooled with cold air at a temperature of T_{Inlet}, which is drawn through the chassis by fans and expelled at a temperature of T_{Outlet}. Hot air returns through a *Computer Room Air Handler* (CRAH) at a

Variable	Value	Description
f	Variable	CRAH Volume Flow Rate
E	0.5	Heat Transfer Efficiency
κ	[0.5, 0.95]	Containment Index
$C_{P_{\text{air}}}$	1.012 $\frac{\text{J}}{\text{g}\cdot\text{K}}$	Heat Capacity of air
f_{Max}	6900 CFM	Maximum flow rate of CRAH
$\dot{m}_{\text{ServerMax}}$	70 CFM	Maximum fan flow rate through a server
P_{FanMax}	9 W	Maximum server fan power [8]
$\dot{m}_{\text{FanMax}}$	86 CFM	Maximum fan flow rate [8]
P_{CRAHPeak}	3 kW	Maximum power draw of CRAH
P_{CRAHIdle}	0.1 kW	Minimum power draw of CRAH
T_{Outlet}	90°F	Server outlet temperature
T_{Outside}	85°F	Outside air temperature

TABLE I: Cooling Parameters. This table provides the value and a brief description the fixed variables used in our study.

temperature of T_{Return}, where heat is removed and cooled to the cold air temperature T_{Supply}. CRAHs exchange the removed heat with a chilled water loop with a cold water temperature of $T_{\text{Chilled}_{\text{Water}}}$ and hot water temperature of $T_{\text{Hot}_{\text{Water}}}$ (The hot water temperature does not appear in our model because it is fully determined by other variables, hence, for brevity, we refer to $T_{\text{Chilled}_{\text{Water}}}$ simply as T_{Water}). Finally, the cold water loop temperature is maintained by a chiller and cooling tower, which expels data center heat into the outside air.

Historically, a major source of cooling waste is in the mixing of warm and cold air. This phenomenon is captured by the parameter κ, which is the *containment index* of the data center [24]. A containment index of 1.0 represents perfect isolation of hot and cold air streams; any value between 0 and 1.0 represents various degrees of mixing.

$$\kappa = \frac{\Delta T_{\text{Server}}}{T_{\text{Outlet}} - T_{\text{Supply}}} \tag{3}$$

Servers. We use a simple, yet accurate, server power model [26]:

$$P_{\text{Server}} = (P_{\text{Max}} - P_{\text{Idle}}) \cdot U + P_{\text{Idle}} \tag{4}$$

Where U is the average server utilization. This model has been validated to work well for coarse-grain power estimates [26]. More sophisticated models can be used to estimate peak power [22] or to account for inter-node variation [7].

The thermodynamics of cooling a server are captured by the

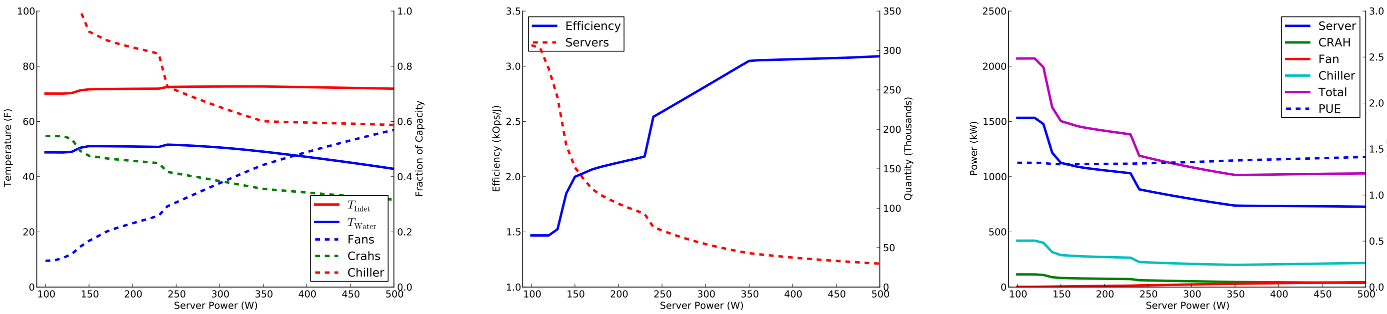

(a) Fan Speed and Temperature (b) Server Quantity and Efficiency (c) Power Consumption and PUE

Fig. 6: Impact of server power on cooling. These three figures illustrate key variables in our data center cooling optimization model as a function of server power for iso-throughput data center designs. Subfigure (a) illustrates server inlet and chilled water loop temperature (left axis) and server and CRAH fan speeds as a fraction of peak (right axis). Subfigure (b) shows the efficiency curve derived from SPECpower results that we input to our model (left axis) and the number of servers (right axis). Subfigure (c) shows the total power dissipation in the servers and each part of the cooling system (left axis) as well as the data center PUE (right axis).

following equation:

$$\dot{Q}_{\text{Server}} = \kappa \cdot \dot{m}_{\text{Server}} \cdot C_{P_{\text{Air}}} \cdot (T_{\text{Outlet}} - T_{\text{Inlet}}) \tag{5}$$

Where $\dot{Q}_{\text{Server}}$ is the rate of heat removal (in watts) and must equal the power draw of the server to maintain stable temperatures, κ is the containment index, $\dot{m}_{\text{Server}}$ is the mass flow rate of air through the server (in kg/s), $C_{P_{\text{Air}}}$ is the specific heat capacity of air (in $\frac{\text{J}}{\text{gK}}$ and T_{Outlet} and T_{Inlet} are the server outlet and inlet temperature, respectively (in Kelvin).

Server reliability depends on maintaining temperatures of each component within the server below some critical threshold. Hence, we fix a temperature constraint at the server outlet to allow our model to solve for the optimal (with respect to cooling cost) combination of inlet temperature and mass flow rate to cool the server. Fixing the temperature at the outlet rather than the inlet ensures all system components operate at a temperature at or below this temperature. Maximum acceptable server operating temperatures are currently hotly debated; we fix T_{Outlet} at 90°F, which is aggressive, but in line with current industry practice. We will shortly describe how we determine T_{Inlet}, which depends on the operation of other cooling components. However, given a T_{Inlet}, we can determine the necessary flow rate, $\dot{m}_{\text{Server}}$, through the server by:

$$\dot{m}_{\text{Server}} = \frac{P_{\text{Server}}}{\kappa \cdot C_{P_{\text{Air}}} \cdot (T_{\text{Outlet}} - T_{\text{Inlet}})} \tag{6}$$

Note that we have equated P_{Server} with $\dot{Q}_{\text{Server}}$ to maintain temperature stability; all server heat generated must be removed by air cooling.

Fans. The required flow rate through a server is achieved by forced air provided by fans. Fan power varies cubically with respect to the flow rate of air pushed (or sucked) by the device. Since flow rate varies linearly with respect to fan speed, the power is:

$$P_{\text{Fan}} = P_{\text{Fan}_{\text{Max}}} \cdot \left(\frac{\dot{m}_{\text{Fan}}}{\dot{m}_{\text{Fan}_{\text{Max}}}} \right)^3 \tag{7}$$

In our case $\dot{m}_{\text{Fan}} = \dot{m}_{\text{Server}}$.

CRAHs. The amount of heat a CRAH removes from the air is governed by the equation:

$$\dot{Q}_{\text{CRAH}} = E\,\kappa \dot{m}_{\text{CRAH}} C_{P_{\text{Air}}} f^{0.7} (\kappa T_{\text{Outlet}} + (1-\kappa) T_{\text{Inlet}} - T_{\text{Water}}) \tag{8}$$

Here, E is the transfer efficiency, κ the containment index, $\dot{m}_{\text{CRAH}}$ the mass flow rate of the CRAH. The value f represents the fractional fan speed of the CRAH (from 0 to 1.0) and determines the power required to run the CRAH itself:

$$P_{\text{CRAH}} = P_{\text{CRAH}_{\text{Idle}}} + P_{\text{CRAH}_{\text{Dyn}}} f^3 \tag{9}$$

Which can be determined by solving Equation 8 for f.

Chiller. We model a chiller capable of removing 8MW of heat, with a peak power consumption of 3,200KW at maximum load. The power model is based on empirical regression curves detailed in [15]. Chiller power only relies on three variables: Q_{Chiller}, T_{water}, and T_{Outside}

Cluster sizing and cooling provisioning. We assume one CRAH is responsible for cooling a unique set of servers; accordingly, a given CRAH removes heat equal to aggregate power dissipation of this set of servers:

$$Q_{\text{CRAH}} = \sum_{i=1}^{N_{\text{Server}}} P_{\text{Server}} = N_{\text{Server}} \cdot E[P_{\text{Server}}] \tag{10}$$

The total power of all CRAHs in the system is:

$$P_{\text{CRAH}_{\text{Total}}} = \sum P_{\text{CRAH}} = N_{\text{CRAH}} \cdot P_{\text{CRAH}} \tag{11}$$

We assume that load is spread across all sets of servers evenly.

C. Understanding the Model

Figure 6 provides the output of our model for various values of server peak power for a cluster with a fixed throughput. We leverage the peak power, dynamic range and efficiency characteristics from Section 3 for the server model. At each server design point, we determine the optimal setting for fan speeds and temperature. Figure 6 (a) provides the temperature of water and air as well as the CRAH and server fan speeds for each design. Increased server power causes an increase in server fan speed rather than altering temperature. In Figure 6 (b) we provide the relationship between peak power and efficiency, and the number of servers needed to achieve the fixed throughput. Finally in Figure 6 (c) we show how each component of the cluster's power varies with server design and how PUE is affected.

D. Cooling Analysis

To understand how the overheads in cooling affect the choice of high- or low-power design, we perform a case study contrasting the two design philosophies. Because low-power servers lower the power density of a data center, their cooling solutions are generally less

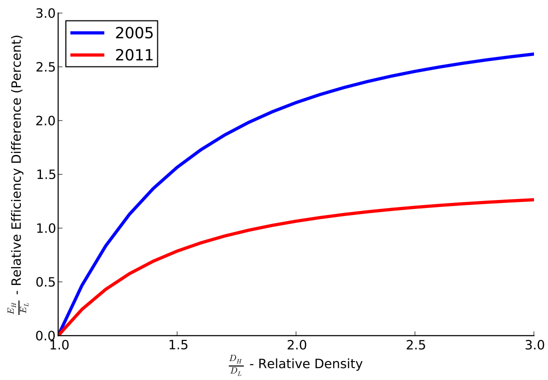

Fig. 7: Relative server efficiency increase necessary to offset the difference in cooling for less dense servers. Clusters are Iso-power and Iso-throughput.

costly. The goal of this study is to understand how much more efficient high-power servers need to be than their low-power counterparts such that the server power savings outweighs the increase in cooling cost. The constants used in this study are summarized in Table I. We evaluate a data center operating at 30% utilization, which is typical of modern facilities [5].

The tradeoff in moving to low-power servers is density. Although low-power servers can tolerate a higher inlet temperature (or lower mass flow rate), more servers must be installed to reach a target throughput. To quantify this tradeoff, we ask the question – for a given reduction in density, how much more efficient must a high-power server design be such that the total cluster power consumption is equal to a low-power design?

Figure 7 illustrates this tradeoff between density and efficiency. The metric $\frac{D_H}{D_L}$ provides the ratio of high-power server density D_H to low-power D_L; at 2.0 this means low-power servers will occupy twice as much space (equivalently, each 1U server draws half as much power). Varying this ratio impacts $\frac{E_H}{E_L}$, the ratio between the high- and low-power efficiency (E_H and E_L respectively). We provide two scenarios to evaluate the break-even point in cooling. First, we evaluate an older data center with a low κ of .9 and conservative T_{Outlet} of 80°F. The tradeoff for this configuration ("2006") demonstrates that reducing density by a factor of three requires just under a 2% gain in server efficiency. Second, we contrast with a modern data center with a higher κ representing hot aisle containment, an aggressive and conservative T_{Outlet} of 95°F and free cooling from air-side economization. In this configuration ("2011"), high-power servers effectively do not need an efficiency advantage over low-power servers to break even in cooling, because the gap in cooling efficiency is negligible. It's important to note that in both scenarios the operating points of the cooling system are at their optimal settings. Therefore, these results should be viewed as using a scheme like [23]; many data centers from the early 2000s may have far worse efficiency due to static temperature and fan control.

V. Conclusion

We have provided a case that low-power design does not always imply an increase in energy efficiency. This study is intended to provide motivation for future research to delve into the details of system design rather than rely on the notion that low-power systems are always more efficient.

Our study provides two key insights challenging the assumption that low power should be favored for energy efficiency. The first is that, in fact, data from real servers shows that lower power servers actually may be far *less* efficient than their high-power counterparts. Peak efficiency, dynamic range and high core-count tend to correlate far more strongly with server efficiency. Second, we demonstrate that high-power servers should not be eschewed in favor of low-power servers due to cooling overheads. Great leaps in infrastructure efficiency have been made in the recent past and the small difference in cooling overheads can easily be overcome by the greater efficiency of the higher-power server. These insights suggest that low-power design may not be key objective for server designers and that the full design space of possible power envelopes should be explored.

This work was supported by grants from Google, Intel, HP Labs, and NSF grants CNS-0834403, CCF-0811320, and CCF-0815457.

References

[1] SPEC power_ssj2008, 2011.
[2] F. Ahmad and T. Vijaykumar. Joint optimization of idle and cooling power in data centers while maintaining response time. *Architectural Support for Programming Languages and Operating Systems*, 2010.
[3] D. G. Andersen, J. Franklin, M. Kaminsky, A. Phanishayee, L. Tan, and V. Vasudevan. FAWN: a fast array of wimpy nodes. *Symposium on Operating Systems Principles*, 2009.
[4] O. Azizi, A. Mahesri, B. C. Lee, S. J. Patel, and M. Horowitz. Energy-Performance Tradeoffs in Processor Architecture and Circuit Design: A Marginal Cost Analysis. *International Symposium on Computer Architecture*, 2010.
[5] L. A. Barroso and U. Hölzle. The Case for Energy-Proportional Computing. *IEEE Computer*, (December):33–37, 2007.
[6] Data Center Knowledge. Microsoft: PUE of 1.22 for Data Center Containers. 2008.
[7] J. D. Davis, S. Rivoire, M. Goldszmidt, and E. K. Ardestani. Accounting for Variability in Large-Scale Cluster Power Models. *Exascale Evaluation and Research Techniques*, 2011.
[8] Ebmpapst. 3200J Series, 2008.
[9] X. Fan, W.-D. Weber, and L. A. Barroso. Power provisioning for a warehouse-sized computer. *International Symposium on Computer Architecture*, 2007.
[10] Google. Data center efficiency measurements, 2011.
[11] J. Hamilton. Cooperative Expendable Micro-Slice Servers (CEMS): Low Cost, Low Power Servers for Internet-Scale Services. *Conference on Innovative Data Systems Reasearch*, 2009.
[12] J. Hamilton. PUE is Still Broken and I still use it, 2010.
[13] J. Hamilton. Yahoo! Compute Coop Design. 2011.
[14] U. Hölzle. Brawny cores still beat wimpy cores , most of the time. *IEEE Micro*, 30(4), 2010.
[15] M. Hydeman and K. L. Gillespie. Tools and Techniques to Calibrate Electric Chiller Component Models. *ASHRAE Transactions*, 2002.
[16] V. Janapa Reddi, B. C. Lee, T. Chilimbi, and K. Vaid. Web search using mobile cores. *International Symposium on Computer Architecture*, 2010.
[17] L. Keys, S. Rivoire, and J. D. Davis. The Search for Energy-Efficient Building Blocks for the Data Center. *Workshop on Energy-Efficient Design*, 2010.
[18] W. Lang. Wimpy Node Clusters: What About Non-Wimpy Workloads? *Workshop on Data Management on New Hardware*, 2010.
[19] Lawrence Berkeley National Laboratory. Facilities: High-Performance Buildings for High-Tech Industries, 2011.
[20] Y. Li, B. Leez, D. Brooks, Z. Huyy, and K. Skadron. CMP Design Space Exploration Subject to Physical Constraints. *High-Performance Computer Architecture*, 2006.
[21] K. Lim, P. Ranganathan, J. Chang, C. Patel, T. Mudge, and S. Reinhardt. Understanding and Designing New Server Architectures for Emerging Warehouse-Computing Environments. *International Symposium on Computer Architecture*, 2008.
[22] D. Meisner and T. F. Wenisch. Peak Power Modeling for Data Center Servers with Switched-Mode Power Supplies. *International Symposium on Low Power Electronics and Design*, pages 319–324, 2010.
[23] J. Moore, J. Chase, P. Ranganathan, and R. Sharma. Making scheduling cool: Temperature-aware workload placement in data centers. In *USENIX Annual Technical Conference*. USENIX Association, 2005.
[24] S. Pelley, D. Meisner, T. F. Wenisch, and J. W. VanGilder. Understanding and abstracting total data center power. *Workshop on Energy-Efficient Design*, 2009.
[25] S. Pelley, D. Meisner, P. Zandevakili, T. F. Wenisch, and J. Underwood. Power Routing: Dynamic Power Provisioning in the Data Center. *Architectural Support for Programming Languages and Operating Systems*, 2010.
[26] S. Rivoire, P. Ranganathan, and C. Kozyrakis. A comparison of high-level full-system power models. *HotPower*, 2008.
[27] D. Tsirogiannis, S. Harizopoulos, and M. A. Shah. Analyzing the energy efficiency of a database server. *SIGMOD*, 2010.