Green Server Design: Beyond Operational Energy to Sustainability

Jichuan Chang, Justin Meza, Parthasarathy Ranganathan, Cullen Bash, Amip Shah
Hewlett Packard Labs

"Green" server and datacenter design requires a focus on environmental sustainability. Prior studies have focused on operational energy consumption as a proxy for sustainability, but this metric only captures part of the environmental impact. In this paper, we argue that to understand the total impact, we need to examine the entire lifecycle of the system, beyond operational energy to also include material use and manufacturing. We make two main contributions. We present a methodology that allows such a lifecycle analysis, specifically providing attribution of sustainability bottlenecks to individual system architecture components. Using this methodology, we compare the sustainability tradeoffs between popular energy-efficiency optimizations and discuss sustainability bottlenecks and optimizations for future system designs.

1. Introduction

Environmental sustainability (the manufacturing, operation, and disposal of products to minimize their environmental impact in terms of destruction of natural resources or production of undesired emissions) is fast becoming an important design constraint for Information Technology (IT) systems [2]. The carbon footprint of the IT industry, though only 2% of the world economy, is estimated to be equal to that of the entire aviation industry [27]. Even more importantly, IT is increasingly being used to address the remaining 98% of the carbon emissions of the world economy [27] (e.g., use of video conferencing to avoid travel) and as this trend continues, it will become more important to design "green" IT systems. A recent estimate showed that up to 75% of organizations will soon consider sustainability as one of the criteria in their IT purchases [28]. The UK government is starting a mandatory Kyoto-style cap-and-trade scheme to curb energy consumptions of businesses [4] and the US Congress has similarly been considering various federal cap-and-trade schemes [3].

There has been a large body of prior work on reducing the operational electricity consumption of servers (e.g. [6] [12] [19] [24] [30] [29] [31]). Given that most of the electricity produced in the world comes from carbon-intensive sources, these optimizations can help improve the carbon footprint of servers and datacenters during operation. However, these approaches do not address environmental impact of a system across *all the stages of its lifecycle* such as the extraction of raw materials, manufacturing, transportation, operation, and disposal.

In this paper, we examine the problem of lifecycle-based optimization of future server and datacenter designs. We make two main contributions – (1) a methodology to reason about sustainability from a system architecture perspective and (2) a systematic analysis of the environmental impact of current designs across their entire lifecycle and the tradeoffs with state-of-the-art energy-efficiency techniques.

2. Measuring Sustainability: Using Exergy for Architectural Studies

Numerous schemes exist to quantify the environmental sustainability of systems. Life-cycle assessment (LCA), a field that has been in practice for nearly 50 years [1], involves taking an end-to-end approach to assessing the environmental impact of a system across various stages in its lifecycle.

In this paper, we perform lifecycle assessment using the thermodynamic metric of *exergy* (available energy) consumption to reason about sustainability. A detailed description of exergy is outside the scope of this paper. However, briefly, unlike energy that is neither created nor destroyed (1^{st} law of thermodynamics), exergy is continuously consumed in the performance of useful work by any real entropy-generating process (2^{nd} law of thermodynamics). Several previous studies have discussed how this destruction (or consumption) of exergy is representative of the irreversibility associated with various processes [7] [21] and correspondingly, to a first order, the environmental sustainability [11]. Additionally, models for specific IT systems [18] have shown that optimizations to reduce lifecycle exergy consumption often map fairly well to optimizations based on other types of environmental criteria such as greenhouse gas emissions, pollution, etc. [32].

Unfortunately, previous lifetime exergy characterizations have estimated the total environmental impact in computer systems based on a mapping of the *system mass* or material flows to per-unit estimates of the environmental impact burden [18][34]. Figure 1(a) shows such a breakdown for a typical server (2-socket Xeon-based server with 4 DIMMs and two 72G HDDs, two 1Gb NICs, and 25% utilization) using these methods. Such a model is not very useful for system architects because extending such a breakdown of exergy to systems architecture choices is not clear. Since architectural choices may span multiple stages of the entire system lifecycle, deciding to use one component over another in a system will result in (often non-intuitive) changes to the total system environmental impact due to differences in the manufacturing process, not just for

the chosen component but also for related components that interact at the system or datacenter level. An approach that considers lifecycle exergy consumption from an architectural perspective is required.
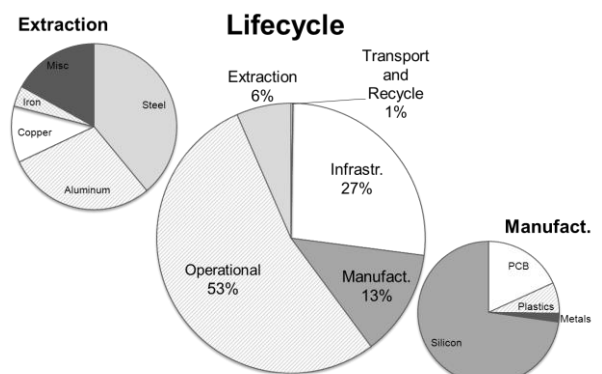
(a) Process-based breakdown of total exergy
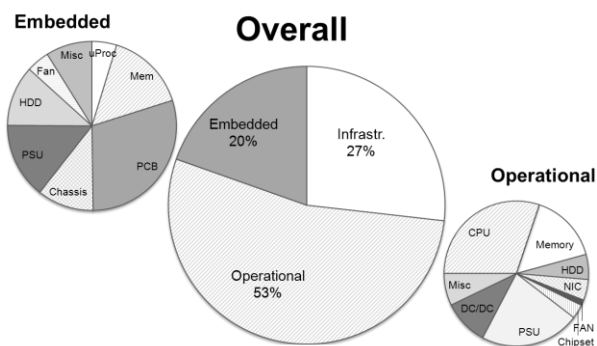
(b) Architecture-based breakdown

Part	Embd. (MJ)	Sources
CPU	158	[9] [13] [26]
Chipset	66	[13] [22] [26]
DRAM	726	[13] [33] [26]
PCB	1400	[13] [18] [34]
Chassis	512	[16] [18] [21]
PSU	683	[13] [18]
HDD	546	[13] [18]
Fan	209	[16] [18]
Misc.	420	[20] [34]

Part	#	TDP (W)	Idle%
Processor	2	95	10%
Memory	4	10	50%
HDD (15K)	2	5	80%
NIC (Gigabit)	2	6	50%
Fan	4	3	0%
Northbridge	1	27.1	0%
Southbridge	1	4.3	0%
PSU	1	33	100%
DC conversion	-	15	100%
Misc.	-	10.6	100%
Total	-	354	-

(c) Sustainability modeling parameters

Figure 1: *(a) illustrates previous process-based approaches to reasoning about sustainability, (b) illustrates our proposed model to reason about sustainability based on system architecture components, (c) summarizes key model parameters.*

Our work attempts to address these issues by adopting an architecture-centric approach to measuring and optimizing the environmental impact of systems. Specifically, we aggregate raw materials at the component level, allowing us to evaluate environmental impact at the granularity of familiar architectural building blocks such as processors,

DIMMs, and hard disk drives, as opposed to their associated raw materials. This enables us to express the environmental impact of complex ensembles of diverse sets of materials succinctly in terms of system architectural choices.

Our approach categorizes exergy[1] into three broad categories – embedded, operational, and infrastructure. *Embedded exergy* is the amount of exergy used to "make" a system component. To a first degree, this is the amount of exergy expended during extraction, manufacturing, transportation, and recycling. For most components, the bulk of the embedded exergy is destroyed during manufacturing as complicated processes use high quality energy to manufacture highly-ordered electronic components, and various chemicals required for making these components themselves require large amounts of energy to manufacture. Our model abstracts out the appropriate exergy destruction values for all of the processes specific to each component, and then aggregates these data to discern the overall exergy consumption related to each architectural component[2].

Operational exergy is the amount of exergy spent during a system's operational lifetime. Although the heat dissipated from the server contains useful work potential, there are currently no practical techniques to harness this waste heat and recover this exergy. In this study, therefore, we assume that operational exergy is equivalent to the electricity consumed during operation. To determine operational exergy, for each component, we use its maximum power rating and model how its power varies with utilization. We determined these values from published sources, internal experiments, and communications with system designers. This model is similar to that used in other recent system studies (e.g. [23]) and provides a high-order estimate of the power consumed across different workloads (varying utilizations). We assume a three-year lifecycle and 99.99% uptime. Figure 1(c) summarizes our model parameters.

In most datacenters, the *cooling and power delivery infrastructure* accounts for a large fraction of the total electricity consumption, and consequently, we account for infrastructure exergy as a separate category. This takes into account the operational energy used by CRAC units, chillers, cooling towers and any other equipment employed in the data center

---

[1] More specifically, it is exergy consumption. In this paper, we loosely use the term exergy to refer to exergy consumption.

[2] We aggregate the embedded data from multiple public sources [13] [34] [33] [20] [9] [22] [16] [26]. Notice these data are derived based on specific supply chain and component models. Modeling embedded exergy in a different context should not directly use these numbers, but rather use the methodology and data sources described here with new, revised assumptions that are appropriate for the system being modeled.
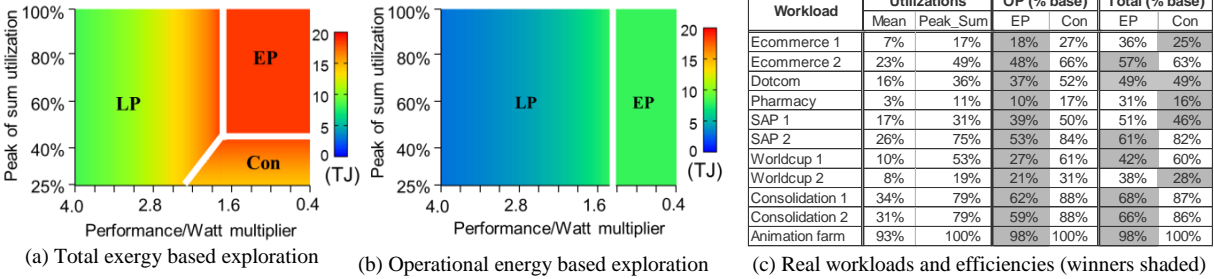
Figure 2: *Illustration of tradeoffs between different energy-efficiency optimizations.*

(a) Total exergy based exploration — **(b) Operational energy based exploration** — **(c) Real workloads and efficiencies (winners shaded)**

Workload	Utilizations		OP (% base)		Total (% base)	
	Mean	Peak_Sum	EP	Con	EP	Con
Ecommerce 1	7%	17%	18%	27%	36%	25%
Ecommerce 2	23%	49%	48%	66%	57%	63%
Dotcom	16%	36%	37%	52%	49%	49%
Pharmacy	3%	11%	10%	17%	31%	16%
SAP 1	17%	31%	39%	50%	51%	46%
SAP 2	26%	75%	53%	84%	61%	82%
Worldcup 1	10%	53%	27%	61%	42%	60%
Worldcup 2	8%	19%	21%	31%	38%	28%
Consolidation 1	34%	79%	62%	88%	68%	87%
Consolidation 2	31%	79%	59%	88%	66%	86%
Animation farm	93%	100%	98%	100%	98%	100%

infrastructure. (Note that on-board fans are considered part of server operational power.) We assume that cooling is provisioned appropriately to handle the maximum power rating, and we use the widely-used power usage effectiveness (PUE) metric [3] [14] to compute infrastructure exergy. The exergy consumption related to building the power and cooling infrastructure in the datacenter is outside the scope of our model; but, when normalized to a datacenter scale and across multiple IT refresh cycles, we expect the allocation of its embedded burden is minimal.

3. Evaluating the state-of-the-art

Exergy breakdown

Figure 1(b) shows the breakdown of total lifecycle exergy using our models. We focus on the same server as in Figure 1(a), and assume a workload utilization of 25%, and a PUE (1.6) based on prior studies [17]. The results show that operational exergy dominates the total exergy of the system (53%), followed by infrastructure exergy (27%), and embedded exergy (20%). Of note is that the embedded exergy contributes a sizable amount to total system exergy. The dominant components of embedded exergy are from silicon-based processes and PCB design. Assuming a datacenter container with 1056 of these servers, the total amount of exergy consumption is 25.4 Tera Joules over a three year timeframe, equivalent to approximately 870 metric tons of coal consumption.

Design space exploration

There has been a large body of prior techniques that address operational energy. However, their impact on total exergy hasn't been studied. Specifically, how do these techniques compare from a sustainability point of view? Are there tradeoffs between operational exergy and embedded exergy that make some of these techniques less effective in improving net sustainability? If these techniques are aggressively applied to future systems, what would the new breakdown of exergy consumption look like? To answer these questions, we studied three broad

categories of optimizations: (i) *Energy proportionality* (EP) [6] in the datacenter space has gained a lot of attention with several optimizations [15] [8] [29] [24] that seek to make the energy consumed by a system be proportional to the activity in the system. (ii) *Consolidation* (Con) is another optimization common in current datacenters. The intuition is that typical utilization on many enterprise services is relatively low and bursty and that across a collection of systems, peaks are often unsynchronized (the peak of the sums of the individual utilizations is lower than the sum of the peak individual utilizations). Multiple virtual machines (or tasks in a task scheduler) on separate servers can be consolidated onto a single server, raising its utilization and reducing the required server count (and total power) [25] [30]. (iii) Recently, there have been several *low-power server solutions* (LP) based on energy-efficient, but lower-power processors [23] [17] [10] [5]. A common idea behind these solutions is to better match the processor architecture to the workload characteristics (primarily around CPU-I/O balance) to leverage significantly better performance/watt.

Figure 2 shows our results from examining these three optimizations for a parameterized design space exploration. The benefits from EP are primarily a function of workload average utilization. Figure 2(a) shows this design space exploration for an average workload utilization held constant at 25%. (We examine other utilization points as well, but omit them for brevity.) For a given average workload utilization, we identify different tradeoffs for the LP designs by using a performance/watt multiplier on the X axis. For some workloads (e.g, enterprise workloads), a lower-power processor may lose more in performance than it saves in power; for these cases the performance/watt multiplier is less than 1 (right side of the figure), indicating the LP solution's performance/watt (or energy efficiency) at peak load is worse than a conventional server. For web workloads, prior studies [23] [17] [10] have found LP to yield better multipliers ranging from 2 to 5 (left side of the axis). As discussed earlier, the effectiveness of consolidation is a function of how many processes can be packed into a single server, which in turn is a function of the peak-of-sum utilization specific to the workload. The Y axis shows

---
[3] PUE = 1+ infrastructure_power / operational_power

this parameter. Lower values indicate that the peaks are completely non-synchronized and consolidation can more readily be leveraged. Different points on the heat map thus represent different workload/system configurations.

For each data point, we individually compute the total exergy for EP, Con, and LP designs providing the same aggregate performance and identify the optimization that achieves the best exergy. (Recall that lower exergy consumption is better.) The heat map's color gradation reflects the absolute value of this best exergy. The division of the heat map into various regimes shows the technique that achieves the best exergy for that region of workload/system configurations. For energy proportionality, we studied a best-case future model where all the hardware shows ideal proportionality (the power consumed in an idle state is zero). For consolidation, we assumed perfect bin-packing that minimizes the number of servers. Figure 2(b) shows a similar picture, but for a case where only operational energy is considered. For EP and Con, we model component power after a conventional server shown in Figure 1(c); for LP, we model an HP BC2500 blade server with maximum component powers similar to [23]. Here we assume a PUE of 1.5 for infrastructure exergy, and adjust the embedded exergy consumption values of components within each system based on a scaling of key physical attributes for each component[4].

Observations

This way of representing the data reveals several interesting high-level trends. First, the figures individually show the different regions when different techniques work best and the cross-over points, as well as the relative magnitude of the benefits. Comparing the two figures allows us to examine the changes to these design tradeoffs when optimizing for just operational energy versus considering total exergy.

Figure 2(a) shows that in general the total exergy of the system is minimized when going towards the bottom left region of the graph—not surprising considering this assumes more power-efficient components and lower resource activity (more consolidation). First comparing EP and Con, we observe that EP outperforms Con when the workloads are not bursty and don't lend themselves to packing (top right part of Figure 2(a)). The break-even point is

roughly corresponding to workloads with peak-of-sum utilizations close to 50%. Below this, Con is a better design alternative. *Interestingly, this conclusion is different than when just focusing on operational energy.* There, given fragmentation in bin-packing, perfect energy proportionality is always better than consolidation. However, when considering total exergy, a reduction in materials associated with fewer servers provides additional reductions in embedded exergy that allow Con to be better than EP[5].

Comparing with LP, we find that after a breakeven point roughly corresponding to 1.6-2.6X improvement in performance/watt, LP designs are always better than both EP and Con. Considering the differences between Figures 2(a) and Figure 2(b), *the inflection point at which LP is better than other alternatives shifts to the left* (requiring even more energy efficiency from lower power processors) *when total exergy is considered.* This is because of the increased embedded exergy from the larger number of lower-power servers required for the same performance. Comparing LP and Con, it is worth noting that there is now a region where consolidation of multiple small processes into one server is better than distributing them into multiple small low-power blades.

Notice that because LP and EP are independent of peak of sum utilization, the break-even point between these solutions is dictated entirely by the performance/watt multiplier. This implies that the optimal choice between these two solutions is dependent on their relative energy efficiencies for the type of workload. The number of machines used in Con depends on the peak-of-sum utilization, but notice that consolidation also raises overall system utilization, increasing operational exergy. This trade-off between fewer machines and higher utilization is shown as the angled line dividing LP and Con.

The table in Figure 2(c) illustrates the tradeoffs between EP and Con with data from various real-world traces. (They correspond to specific real-world points in the bottom right portions of the heat maps.) From an *operational* exergy perspective, EP achieves more savings compared to Con for all the enterprise traces, but, by contrast, from a *total* exergy perspective, in many cases Con outperforms EP.

4. Discussion

The results above illustrate that *focusing on the most efficient system design for operational energy does not always produce the most sustainable solution.* Tradeoffs between operational exergy and embedded exergy need to be considered. The examples in the previous section—requiring larger factors of energy

---

[4] For example, we find that the key physical attribute governing the footprint of a microprocessor is the area of the silicon. Thus, we normalize the impact calculated in Fig. 1(c) by the area to derive an 'impact factor' representing the exergy consumption per unit area. This impact factor can then be scaled as required for processors of different sizes, assuming uniform thickness; fabrication; etc. If other key attributes vary (e.g., a change in the thickness of the package), these can be accordingly parameterized as well. A similar approach can be repeated for each of the different architectural components.

[5] Note that we assume a model where consolidation leads to lower *provisioning* of servers; if consolidation just allowed servers to be turned off, we would not get the embedded savings.

efficiency improvement for low-power servers to be sustainably better, or consolidation being more sustainable than energy proportionality—illustrate this point. The best way to optimize for sustainability is to use power-efficient and material-efficient systems that scale power with resource usage and are utilized fully.

In future systems, as the ratio of embedded exergy to total exergy grows, *new optimizations will be needed that explicitly target embedded exergy*. For example, upcycling (reusing components when they would normally be recycled or discarded) is an effective way to reduce embedded exergy, amortizing the destruction of exergy over a longer period of time. However, this will require new ways of building systems, including designs that allow technology upgrades to be localized only to the components that need to be upgraded, allowing the rest to be upcycled. "Dematerialization" techniques that reduce the material in the solution will also be important. This will require identifying the sweet spot of resources for best performance efficiency. For example, smaller memory configurations could use less silicon and consequently reduce the embedded exergy associated with memory.

Finally, when considering the approaches above, it is important to note that embedded exergy, operational exergy, infrastructure exergy, and performance are not independent variables. For example, dematerialization sometimes reduces infrastructure exergy consumption (e.g., removal of sheet metal in the backplane can enable better designed air flow), but in other cases increases infrastructure exergy (e.g., removal of fans in a server can increase overall cooling energy in the datacenter). Similarly, different optimizations can have different tradeoffs on performance: backplane redesign for dematerialization can impact networking topologies, reductions to cooling infrastructure may lead to performance throttling, and so on. It will therefore be important to *address sustainability holistically across the various components of total lifecycle exergy*.

Overall, as sustainability becomes a more important design consideration for future systems, design methodologies and system optimizations need to correspondingly change to address these emerging challenges. This paper takes the first steps in this direction—around a methodology to reason about sustainability bottlenecks from an architectural viewpoint, and enabling an understanding of tradeoffs and bottlenecks in future designs. We believe, however, that we have only scratched the surface and that these areas offer a rich opportunity for more innovation by the broader community.

References

[1] ISO 14040: Environmental management – Life Cycle Assessment – Principles and framework. ISO, 2006

[2] Revolutionizing Datacenter Energy Efficiency. McKinsey, 2008.

[3] Regional Greenhouse Gas Initiative. http://www.rggi.org. 2008.

[4] UK Government. Carbon Reduction Commitment, July 2009.

[5] D. Andersen, J. Franklin, et al. FAWN: a fast array of wimpy nodes. SOSP 2009.

[6] L. A. Barroso and U. Hölzle. The case for energy-proportional computing. IEEE Computer, 40(12):33–37, 2007.

[7] A. Bejan. Advanced Engineering Thermodynamics (2nd Edition). John Wiley & Sons, 1997.

[8] R. Bianchini and R. Rajamony. Power and energy management for server systems. IEEE Computer, 37(11):68–74, 2004.

[9] S. Boyd, A. Horvath, et al. Life-cycle energy demand and global warming potential of computational logic. Env. Sci. Tech., 2009.

[10] A. Caulfield, L. Grupp, and S. Swanson. Gordon: using flash memory to build fast, power-efficient clusters for data-intensive applications. ASPLOS-XIV, 2009.

[11] I. Dincer and M. Rosen. Exergy: Energy, Environment and Sustainable Development. Elsevier, 2007.

[12] X. Fan, W-D. Weber, and L. Barroso. Power provisioning for a warehouse-sized computer. ISCA 2007.

[13] R. Frischknecht, et al. The ecoinvent database: overview and methodological framework. J. Life Cycle Assessment, 10(1), 2005.

[14] The Green Grid. Green grid metrics: Describing datacenter power efficiency. http://www.thegreengrid.org, 2007.

[15] M. Gupta, S.Singh. Greening of the Internet. SIGCOMM 03.

[16] T. Gutowski, et al. Thermodynamic analysis of resources used in manufacturing processes. Env. Sci. Tech., 43(5), 2009.

[17] J. Hamilton. Cooperative expendable micro-slice servers: Low cost, low power servers for internet-scale services. CIDR 2009.

[18] C. Hannemann, et al. Lifetime exergy consumption as a sustainability metric for enterprise servers. ASME ICES, 2008.

[19] T. Heath, et al. Mercury and Freon: Temperature emulation and management for server systems. ASPLOS-XII 2006.

[20] Y. Huang, C. Weber, and H. Matthews. Carbon footprinting upstream supply chain for electronics manufacturing and computer services. IEEE ISSST 2009.

[21] D. Morris, J. Szargut and F. Steward. Exergy analysis of thermal, chemical and metallurgical processes. Hemisphere, 1988.

[22] N. Krishnan, et al. A hybrid life cycle inventory of nano-scale semiconductor manufacturing. Env. Sci. Tech., 42(8), 2008.

[23] K. Lim, P. Ranganathan, et al. Understanding and designing new server architectures for emerging warehouse-computing environments. ISCA 2008.

[24] D. Meisner, B. Gold, and T. Wenisch. PowerNap: Eliminating server idle power. ASPLOS-XIV, 2009.

[25] R. Nathuji and K. Schwan. VirtualPower: coordinated power management in virtualized enterprise systems. SOSP 2007.

[26] J. Oliver, R. Amirtharajah, et al. Life cycle aware computing: Reusing Silicon Technology. IEEE Computer, 40(12), 2007.

[27] The Climate Group and GeSI, SMART2020: Enabling the low carbon economy in the information age, 2008.

[28] D. Plummer, et al., Gartner's top predictions for IT organizations and users: Going green and self-healing, 2008.

[29] R. Raghavendra, et al. No "power" struggles: Coordinated multi-level power management for the data center. ASPLOS 2008.

[30] K. Rajamani and C. Lefurgy. On evaluating request-distribution schemes for saving energy in server clusters. ISPASS 2003.

[31] P. Ranganathan, P. Leech, D. Irwin, and J. Chase. Ensemble-level power management for dense blade servers. ISCA 2006.

[32] A. J. Shah, C. D. Patel, , and V. P. Carey. Exergy-based metrics for sustainable design. IEEES-4, 2009.

[33] E. Williams. The environmental impacts of semiconductor fabrication. Thin Solid Films, 461(1), 2004.

[34] E. Williams. Energy intensity of computer manufacturing: Hybrid assessment combining process and economic input-output methods. Env. Sci.Tech., 38(22), 2004.