Totally Green: Evaluating and Designing Servers for Lifecycle Environmental Impact

Jichuan Chang Justin Meza Parthasarathy Ranganathan
Amip Shah Rocky Shih Cullen Bash

Hewlett Packard Laboratories, Palo Alto, USA
{jichuan.chang,justin.meza,partha.ranganathan,amip.shah,rocky.shih,cullen.bash}@hp.com

ABSTRACT

The environmental impact of servers and datacenters is an important future challenge. System architects have traditionally focused on operational energy as a proxy for designing green servers, but this ignores important environmental implications from server production (materials, manufacturing, etc.). In contrast, this paper argues for a lifecycle focus on the environmental impact of future server designs, to include both operation and production. We present a new methodology to quantify the total environmental impact of system design decisions. Our approach uses the thermodynamic metric of exergy consumption, adapted and validated for use by system architects. Using this methodology, we evaluate the lifecycle impact of several example system designs with environment-friendly optimizations. Our results show that environmental impact from production can be important (around 20% on current servers and growing) and system design choices can reduce this component (by 30–40%). Our results also highlight several, sometimes unexpected, cross-interactions between the environmental impact of production and operation that further motivate a total lifecycle emphasis for future green server designs.

Categories and Subject Descriptors

C.5.5 [**Computer System Implementation**] Servers

General Terms

Design, Measurement, Performance

General Terms

Environmental sustainability, lifecycle impact, server architecture, datacenter design, disaggregation, dematerialization, green computing, exergy

1. INTRODUCTION

Environmental sustainability—the minimization of environmental impact from destruction of natural resources or production of unwanted emissions—is an important social issue that has expanded to include computing systems as well. For example, the carbon footprint from information and communication technology (ICT) is estimated to be more than that of the aviation industry [1]. This impact is likely to grow as other industries adopt computing-based solutions to reduce their carbon footprint (e.g., the use of video conferencing to avoid travel). Recently, the Environmental Protection Agency announced an EnergyStar rating for servers and is legislating stricter requirements and incentives for sustainable solutions. Several companies (including Cisco, Dell, Google, HP, IBM, and Intel) have announced efforts to reduce the environmental footprint of their product offerings. Such trends around environmental stewardship, regulatory compliance, and business differentiation, make it important to design future servers and datacenters to be "green," that is, with reduced environmental impact.

Prior studies in the systems' community to address this challenge have primarily focused on the operational energy as a proxy for the environmental impact. For example, prior work has addressed the electricity consumption of individual servers and clusters as well as the electricity consumption of the power delivery and cooling infrastructure [2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12]. Such electricity consumption can be significant—a recent study [13] noted that the worldwide datacenter electricity consumption was higher than that of all but four countries in the world! Given that most of this electricity comes from carbon-intensive sources, mitigating this component can significantly address the environmental impact.

However, such focus purely on the operational energy misses additional key components of the environmental impact from other stages in the lifecycle such as the extraction of raw materials, manufacturing, transportation, and recycling. Such sustainability implications from the *production* of the server can be an important contributor to the overall environmental impact and can potentially impact design decisions for sustainability. For example, on small devices such as cell phones, 40–50% of the environmental impact has been attributed to the manufacturing process [14].

It will therefore be important for system architects interested in green designs to study the *lifecycle* environmental impact of their designs. However, a key challenge has been the *absence of good methodologies* to define the net environmental impact of a system design, particularly in a way that can be related back to system design decisions. Additionally, while we can intuitively reason that a server design to address the environmental impact from production would potentially impact other factors such as the operational energy and system performance, *these interactions are not well understood or quantified* for specific designs.

This paper addresses these challenges. Specifically, we make two main contributions. First, we present a methodology that system architects can use to reason about total lifecycle environmental impact of a server design. Our approach is based on the well-studied thermodynamic metric of *exergy* consumption, but adapted to create models usable from an architectural perspective. Applying our approach, we show that production can contribute significantly to the total environmental sustainability of a server (20% on current servers and higher on future systems).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
ASPLOS'12, March 3–7, 2012, London, UK.
Copyright 2010 ACM 1-58113-000-0/00/0010…$10.00.

Second, we use our lifecycle-focused approach to study example designs optimized for environmental sustainability. Our designs include several optimizations such as simpler low-power processors, disaggregation (modular fungible building blocks), dematerialization (reduced-material design), and free cooling (use of ambient outside air). Our results show that system architecture decisions can reduce the environmental impact of production too, by 30–40% for our designs. Also, the environmental impact from production and operation interact in several, sometimes unexpected, ways. For example, reducing the material used in system design reduces the environmental impact from production as intended, but can also improve the operational energy efficiency significantly. Similarly, an energy proportional design, compared to virtual machine consolidation, always achieves better operational energy efficiency, but is not always greener. Overall, our results motivate holistic optimizations across the lifecycle and additional new opportunities to reduce the environmental impact of future designs.

The rest of the paper is organized as follows. Section 2 discusses our methodology to systematically quantify the lifecycle environmental footprint of a server and uses this methodology to quantify the sustainability bottlenecks in existing designs. Section 3 uses our proposed approach to characterize the total environmental impact of sustainability-optimized designs and uses the results to highlight interactions and tradeoffs from a lifecycle-based environmental focus. Section 4 discusses other optimizations and applications for lifecycle analysis and implications for the broader research community. Section 5 discusses related work, and Section 6 concludes the paper.

2. LIFECYCLE SUSTAINABILITY MODELS FOR SYSTEM DESIGNERS

Below, we describe our proposed methodology for system architects to reason about the total environmental impact of their designs. Section 2.1 provides background on *exergy*, the thermodynamic metric on which our model is based. Sections 2.2 and 2.3 discuss how an exergy-based model can be adapted for system designers to measure the environmental impact from production and operation respectively. Section 2.3 applies the model to understand the lifecycle environmental impact of an existing server design and discusses the validity of our model.

2.1 Exergy to measure sustainability: background

What is exergy? The thermodynamic metric, *exergy*, measures the *available energy* in a system in the unit of joules. This is formally defined as the maximum amount of useful work that can be derived relative to a surrounding reference state. For example, consider a piece of coal or a piston containing highly-compressed gas. Both of these have high exergy due to the potential for useful work. In contrast, the combustion products (e.g., ashes) or the same piston with no compression both have near-zero exergy since little useful work can be derived.

The processes of combustion or piston expansion both *conserve energy* (1st law of thermodynamics) but *consume or destroy exergy* (2nd law of thermodynamics). Such destruction of exergy is typically representative of the irreversibility in a process. For example, the combustion products cannot be reversed to form coal by themselves; this irreversibility is what makes a process or system *unsustainable*: if it were somehow physically possible to reverse the state transitions of coal combustion and associated byproducts, the use of fossil fuels like coal may not be a problem. In this context, a sustainable process may be considered as one that transitions between two physical states *reversibly* (zero exergy destruction). Thus the process irreversibilities captured by exergy destruction can also be used to evaluate sustainability.

Why exergy consumption versus other metrics? Environmental impact has been evaluated along a variety of different metrics, including carbon (greenhouse gas) emissions, energy use, pollution, toxicity, impact to human health, etc. [15]. Often, these metrics do not respond uniformly to a particular design change: for example, replacing one component with another might reduce greenhouse gases, but increase toxic emissions. A large body of research in the environmental sciences is devoted to understanding these different impacts, and developing models that reduce a large number of different environmental impacts into a smaller number of variables [16, 17, 18, 19].

Exergy consumption has recently been shown to be one such metric on which such multi-variable optimizations can be based [15, 20], especially for ICT systems. In particular, empirical studies [15] have showed that optimizing for lifetime exergy consumption leads to similar results to those from more complex environmental models that included resource consumption, ecosystem quality, and human health. Other studies [20] have further clarified the boundary conditions under which exergy consumption can be used as a proxy for sustainability (e.g., the quantity of material and energy used by the system, the number of "energy flows" of material transitioning from one state to another, etc.) and showed that these conditions apply for most classes of IT systems, including those that are the focus of this paper.

An additional advantage with exergy consumption is the availability of a large amount of public data accessible to system designers. For example, several studies document generic mass-specific exergy consumption values of different materials as well as the exergy consumption in various common thermal, chemical, metallurgical, manufacturing, and transportation processes specific to building servers [21, 22, 23, 24, 25].

2.2 Adapting exergy consumption for system designers: environmental impact from production

Limitations of prior exergy-based approaches. Prior work has used *lifetime exergy consumption* to study sustainability but these methods do not provide useful insights for computer architects.

As an illustration, Figure 1(a) shows a breakdown of exergy consumption for an HP DL360 G6 server (two-socket Xeon, four DIMMs, two 72GB HDDs, two 1Gb NICs) following the approach discussed by Hannemann et al. [27]. We focus here on *embedded* exergy consumption, the exergy consumed to "make" the server, namely that spent in materials extraction and manufacturing. To obtain embedded exergy consumption, we systematically disassemble the server and recursively analyze each component further to obtain a mapping of material mass on a per-unit basis (e.g., amount of copper in a server). The resulting data is combined with published exergy consumption values for various materials and manufacturing processes that trace the material flow (from initial extraction through end-of-life) [27], to obtain the total breakdown of embedded exergy consumption (Figure 1(a)).

Unfortunately, such an approach is not useful for computer architects. First, to study a different server, or a different architectural choice within a server, the entire process has to be repeated. It is impractical to have such disassembly-based exergy analyses for different architectural choices, particularly for future

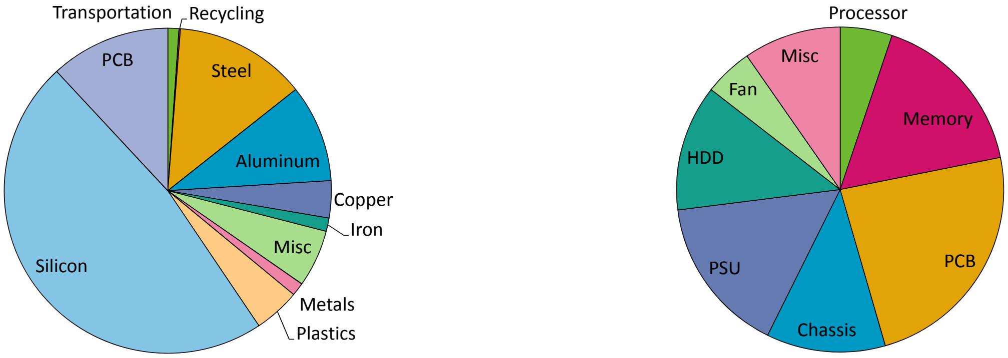

(a) **(b)**

Figure 1. (a) illustrates the breakdown of embedded exergy consumption from existing sustainability models, which do not provide meaningful insights for systems architects. In contrast, our approach, (b), allows for environmental impact to be defined at the granularity of architectural building blocks. (Total embedded exergy consumption for both approaches is approximately 4.6GJ for a single server.)

alternative designs that do not yet exist physically. Second, architectural choices span multiple stages of the lifecycle and multiple categories of environmental impact. Optimizations to one system component can result in non-intuitive changes to the total environmental impact by affecting multiple materials. For example, replacing HDDs with flash memory would reduce aluminum and increase silicon, but the designer would have to go to the point in the supply chain where these materials were extracted and retrace the flow of materials across the lifecycle. Similarly, each material category can include contributions from multiple system parts that are hard to separate out. For example, the copper extraction category is influenced by both the PCB traces and the power supply transformer; the silicon category is affected by the CPU, the chipset, and the memory. It is hard to reason about how different system architectures would affect different materials in a system and, in turn, affect the total embedded exergy consumption. Existing approaches to lifetime exergy consumption can thus be laborious and time-consuming for system designers.

Developing exergy consumption models from an architectural perspective. To address these challenges, we modify the above approach to consider lifecycle exergy consumption from an architectural perspective. Specifically, we aggregate exergy analysis 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 and processes. For each component, in conjunction with environmental engineers, we develop a model that normalizes the *exergy-based impact* in terms of physical characteristics that can be related back to the system architecture across both the embedded and operational phases. As an example, DIMMs may be normalized per MB for a given technology, which in turn can be related to the die size; HDDs can be normalized per unit-storage capacity, which in turn can be related to the platter size; and so on. With this approach, understanding the impact of a design optimization at a systems level (e.g., on the die area) can now lead to a corresponding understanding of the impact on embedded exergy consumption.

To determine such architectural *exergy-based impact factors* (the multipliers used to convert from architectural parameter units into embedded exergy consumption), we examine additional published data from system component spec-sheets. As an example, we plot typical mass versus die size for a number of processors and use this plot to convert the typical per-gram estimates of processors into a per-area estimate. This approach is repeated for key parameters of other components (for example, disks are evaluated per unit area; fans are evaluated per unit radius; etc.).

Figure 2 summarizes the impact factors for our model. With our approach, a system architect now only needs to reason about a few (in this case, nine) high-level architecture-specific parameters – the size and number of processor dies, the size and number of memory dies, the size of the disk, the size of the fan and the dimensions of the printed circuit boards, power supplies, and the chassis[1]. The last column of Figure 2 summarizes the sources from which the impact factors were derived. Note that these numbers are based on public exergy consumption data specific to the supply chain assumptions for the (older) G6 server discussed in this paper. However, the methodology and data sources can be easily extended to derive impact factors for other architectural studies with alternate assumptions.

Component	**Impact factor**	**Unit**	**Sources**
Processor			[23, 21, 25]
Pre-fabrication	0.042	MJ/mm^2	
Fabrication	0.7	MJ/mm^2	
Post-fabrication	0.14	MJ/mm^2	
Memory			[21, 25]
Pre-fabrication	0.042	MJ/mm^2	
Fabrication	0.3	MJ/mm^2	
Post-fabrication	0.03	MJ/mm^2	
PCB	3500.0	MJ/m^2	[21, 27]
Chassis	64.0	MJ/kg	[24, 27]
Power Supply Unit (PSU)	341.5	MJ/kg	[21, 27]
Hard disk drive (HDD)	28.4	MJ/in^2	[21, 27]
Fan	0.87	MJ/mm	[24, 27]
Misc. (LEDs, ASICs, etc.)	4200.0	MJ/kg	[22]

Figure 2. An architecture-centric model for embedded exergy consumption. In contrast to disassembly and materials/manufacturing specific exergy consumption, our model uses a few architecture-specific impact factors to compute embedded exergy consumption.

Figure 1(b) illustrates the breakdown of embedded exergy consumption from our model. The embedded exergy consumption

---

[1] The exergy impact factors for the chassis and power supply are still on a weight-basis in our model. However, the weights for these components are usually documented in spec-sheets. For optimizations around new materials-efficient packaging or chassis redesign—such as those that we study later in this paper—dimensional information can be used to compute the new volume, and correspondingly the new weight.

computed through our model is consistent with that determined through disassembly-based analysis, but with added flexibility for studying system architecture tradeoffs. Comparing Figures 1(a) and 1(b) illustrate the latter point. Unlike the lack of architectural insights in Figure 1(a), Figure 1(b) now shows that approximately three quarters of the embedded exergy consumption can be traced back to the DRAM memory, PCB boards, chassis, and the power supply unit (PSU).

2.3 Adapting exergy consumption for system designers: environmental impact from operation

Server-operational exergy consumption. Exergy consumption can also be used to quantify the environmental impact from the operational phase of the system lifecycle (*server-operational* exergy consumption). When electricity is consumed by a server, it is converted to heat and loses most of the potential for useful work. Consequently, in most cases (unless the waste heat is harnessed), the exergy consumption during operation can be approximated by the total electricity consumption during operation.

To model such server-operational electricity consumption, prior studies have proposed and validated first-order models that estimate the power consumed across different workloads as a function of workload resource utilization (e.g., [28, 29]). We use a similar approach. For each component, we use its maximum power rating and model how its power varies with utilization. Given workload assumptions on resource usage and the time period of operation, this model can then be used to calculate the operational exergy consumption.

Infrastructure-operational exergy consumption. In most datacenters, the cooling and power delivery infrastructure (air conditioners, power delivery units, etc.) account for a large fraction of the total electricity consumption, and consequently, we account for *infrastructure-operational* exergy consumption as a separate category. (Note that on-board fans are considered part of the server-operational power.) This infrastructure exergy consumption can be modeled using the power usage effectiveness (PUE) multiplier [30],

$$\text{PUE} = \frac{\text{operational power} + \text{infrastructure power}}{\text{operational power}}.$$

For many design optimizations, particularly those that just involve changes to the system architecture, a simple pre-determined PUE multiplier can be used over the operational exergy consumption to compute the infrastructure exergy consumption. However, for design optimizations that involve changes to the packaging or cooling architectures, our model uses a thermo-volume based approach developed by Shah et al. [31] to determine the coefficient of performance of the ensemble (COPG) and in turn, the PUE, to compute the infrastructure exergy consumption. This approach is summarized in Appendix A1. Other approaches such as Mercury [10] can also be used to characterize cooling properties for infrastructure-operational exergy consumption.

2.4 Applying the lifecycle exergy consumption model to an existing server

Benefits of proposed approach. The methodology discussed so far enables high-level *lifetime* analysis of the environmental impact, addressing all stages end-to-end, including raw material extraction, manufacturing, and operation. At the same time, the model is *usable by system designers*, with parameters that can be easily related back to architectural design decisions, and used to derive quick answers for high-level design space exploration. Our approach is also *easily extensible* to address new optimizations and future design choices.

Component	**#**	**Peak power (W)**	**% idle/peak power**
Processor	2	95	10%
Northbridge	1	10	50%
Southbridge	1	5	80%
Memory	4	6	50%
HDD (15K RPM)	2	3	0%
NIC (gigabit)	2	27	0%
Fan	4	4	0%
PSU	1	33	100%
DC conversion	-	15	100%
Misc. (LEDs, ASICs, etc.)	-	10	100%
Total	**-**	**354**	**-**

Figure 3. Number of components, peak power, and idle power factors for the system considered in this study. Power numbers are for a single component. Dashes indicate not applicable data.

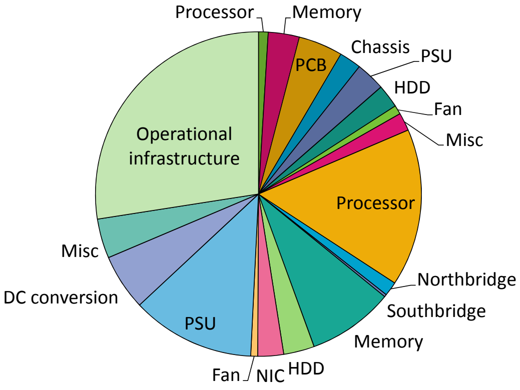

Figure 4. Breakdown of exergy consumption for a traditional server design. Clockwise from the top, the first set of Processor to Misc values are for embedded exergy consumption and the second set of Processor to Misc values are for operational-server exergy consumption.

Sustainability bottlenecks in current server designs. Next, Figure 4 shows the results from applying our methodology to study the total environmental impact of our server, including the operation and production phases. Our choice of parameters for the embedded exergy consumption is detailed in Figure 2, and Figure 3 summarizes the additional model parameters for operational exergy consumption; these are derived from the technical specifications of the server supplemented with additional power characterization experiments. We assume a three year lifecycle and 99.99% uptime, average workload utilization of 25%, and a default PUE of 1.5 (based on [28]).

The overall exergy consumption results show that for a single server, operational exergy consumption dominates the total exergy consumption (54%), followed by infrastructure exergy consumption (27%), and embedded exergy consumption (19%). Notably, the embedded exergy consumption contributes a sizable amount to total system exergy consumption. The components of embedded exergy consumption can be categorized into two broad categories—the silicon electronics in the system (CPU, chipset, memory, ASICs, NICs, etc.) and the physical infrastructure (PCB, chassis, power supplies, fans, etc.). Overall, a datacenter container with a thousand of these servers would consume about 25 terajoules of exergy over a three year period, equivalent to nearly 870 metric tons of coal consumption.

Validation. Section 2.1 discussed the applicability of exergy consumption to measure sustainability. We also performed validation experiments to compare our architecture-focused exergy consumption model with traditional (disassembly-based) exergy consumption models from the environmental community. Comparing the two approaches, we find that the disassembly approach would have suggested a lifetime exergy consumption footprint of around 26.8 GJ per server, while our component-based approach suggests an exergy consumption of approximately 25.4 GJ (a difference of around 5%). In terms of the composition across the lifecycle, both the disassembly-based approach and the architecture-based approach indicate similar breakdown across the server, infrastructure, and embedded phases (within a few percentage points in each category). Thus, we find the two approaches provide quite similar results. However, as discussed earlier, the disassembly-based approach would require on the order of days or even weeks to complete for a new server, while the component-based approach allows for repeated analysis on the order of hours.

3. LIFECYCLE SUSTAINABILITY OPTIMIZATION INSIGHTS: DESIGN STUDIES

This section studies how design tradeoffs change with a lifecycle focus on sustainability. Specifically, how do the various components of the total environmental impact interact with each other? Section 3.1 discusses the sustainability-optimized designs that we consider and Section 3.2 discusses the benefits and insights from evaluating these designs with a lifecycle focus.

3.1 Systems considered

All the designs we consider in this section are container-based solutions. This models a recent trend in the industry where companies like Google, Microsoft, and HP have all built datacenters from shipping containers. A container-based design offers the opportunity to optimize an independent large-sized modular building block in a datacenter, where the only external connections are networking, power, and cooling.

We study three designs—a *baseline* design, similar to containers currently shipping in the market using rack-mounted servers, and two optimized designs, one focusing mainly on *operational energy improvements* using low-power processors in traditional blade servers, and a second *sustainability-optimized* design using more future-looking optimizations around disaggregation, dematerialization, and free-cooling.

Baseline. Figure 5(a) presents our baseline container design, broadly based on the HP POD container design [33]. We assume a standard container size of 12m × 2.33m × 2.35m. The container includes 1056 servers arranged in two aisles of 11 racks each. We assume traditional server designs, specifically using the DL360 server discussed previously in Section 2. The cooling to the container is provided by server-level fans and a container-level direct-expansion heat exchanger unit (DX).

Low-power blade server architectures. The *LP-blade* design is based on proposals in several recent studies [29, 32, 34, 35]. Specifically, these studies argue for low-power server architectures using simpler mobile processors better matched with I/O-bound cloud workloads, and present results showing significant improvements in electricity consumption. We wanted to study the *lifecycle* environmental impact of such operational-energy–optimized designs. Our low-power blade design is based on the HP bc2500 server [36]; the choice of this specific server to derive our parameters was motivated by our access to several of these boards for disassembly for validation. This server uses an Athlon 64X2 3000+ dual-core processor with 2GB DDR2 SODIMM, 80GB SATA drive, and Broadcom 10/100 integrated NIC. We assumed that the server count is scaled to keep the peak throughput performance similar to the baseline (i.e., we benchmarked the performance of both the systems individually and increased the number of *LP-blades* by a factor proportional to the difference in performance). Blades within the enclosure [36] share power and cooling, while outside the enclosure we assume the same container cooling infrastructure.

Dematerialized system architectures. Our final design, *demat,* examines more advanced optimizations targeted at minimizing the environmental impact. Specifically, based on the breakdown of the important sustainability bottlenecks in Section 2, we chose to study three optimization techniques—disaggregation, dematerialization, and free-cooling.

A *disaggregated* system architecture separates system functionality into modular fungible blocks—such as a design that

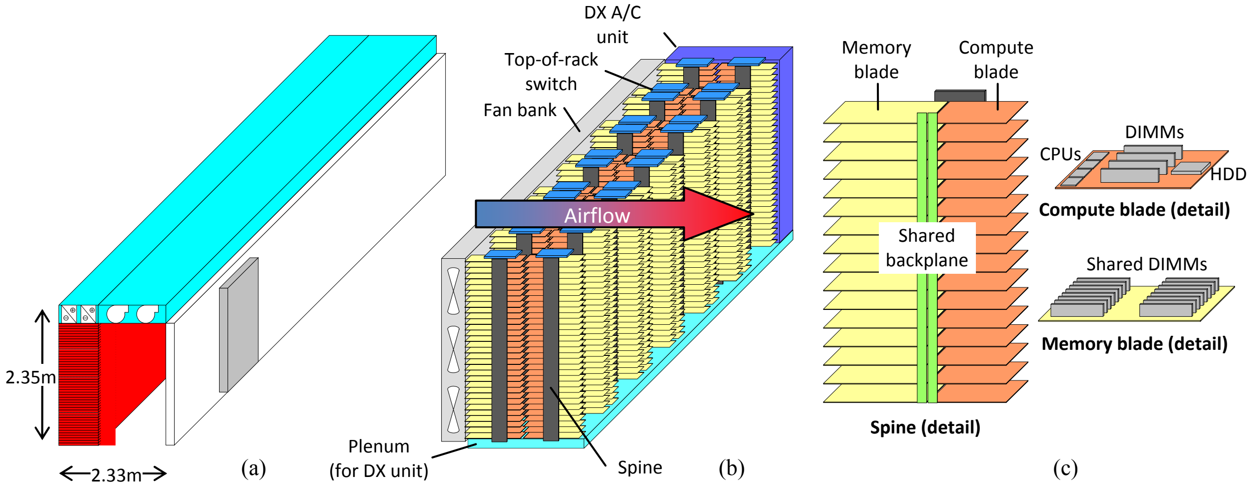

Figure 5: Systems considered. (a) shows the container assumptions for the container design using *baseline* rack-mounted servers; *LP-blade* uses a similar container design except with enclosures of low-power blades instead of rack-mounted servers. (b) and (c) present the design of *demat,* incorporating disaggregation, dematerialization, and free-cooling optimizations.

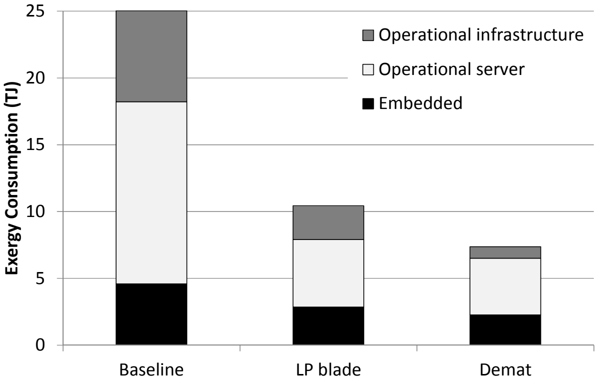

Figure 6. High-level breakdown of total exergy consumption for the three system designs.

uses separate compute blades, storage blades, memory blades, and so on. This improves sustainability by enabling more component-efficient building block designs as well as component reuse across broader ensembles. A *dematerialized* design [1] reduces the amount of material used, particularly around the sheet metal in physical packaging, but also in the size of the printed circuit boards, for example. *Free-cooling* is a technique used in recent datacenters (e.g., those of Facebook, Google, Microsoft) to reuse ambient outside air to cool the datacenter without the infrastructure to chill the air [37].

Figures 5(b–c) present our design incorporating these techniques. The physical packaging and layout of the design consist of multiple "spine" backplanes, with blades attached to the spine. The spines provide communication and power delivery across different elements, and are perforated for air flow. This design reduces sheet-metal, without significantly impacting structural and EMI constraints. We assume two blade types—material-efficient compute blades with low-power processors similar to *LP-blade*, and memory blades that build multi-level memory hierarchies similar to previous proposals [38]. The cooling for the entire container is provided by a single array of shared fans along one side of the container. Blades with similar thermal densities are clustered for improved cooling. The results in this paper assume a location such as San Francisco where typical annual temperatures range between 0 and 30°C.

3.2 Lifecycle sustainability evaluation

3.2.1 Quantification of lifecycle benefits

Our methodology for lifecycle analysis can be used to quantify the total environmental impact of the three designs—*baseline*, *LP-blade*, and *demat*—including both operation and production. Figure 6 summarizes the results. The *baseline* design consumes 25TJ of exergy over its lifecycle. In contrast, *LP-blade* consumes only 10.5TJ, a 58% reduction in the environmental impact (or 2.4× more environmentally sustainable) relative to *baseline*. *Demat* does better, consuming only 7.4TJ, a 71% reduction (and 3.4× better) relative to *baseline*. This translates (using the approach used in [15]) to saving approximately 500–600 metric tons of coal over the system's lifetime.

The overall environmental benefits in the *LP-blade* and *demat* designs come from improvements to the environmental impact from *both operation and production*. Relative to the corresponding numbers for *baseline*, *LP-blade* achieves a 63% reduction in server-operation and infrastructure-operation exergy consumption and a 38% reduction in the embedded exergy consumption for materials and manufacturing. Similarly, relative to the *baseline*, *demat* achieves a 69% reduction in server-operation, 87% in infrastructure-operation, and 51% in embedded exergy consumption.

3.2.2 Cross-category exergy consumption interactions

A further breakdown of the lifecycle environmental impact of the three designs identifies interesting (and in some cases, unexpected) interactions between the different components in the total lifecycle breakdown. Below, we discuss some main results where optimizations to reduce the operation exergy consumption also reduced the embedded exergy consumption and vice versa.

Optimizing operational exergy consumption also benefits embedded exergy consumption. Consider the total exergy consumption benefits of *LP-blade* relative to *baseline*. As expected, the server-operation exergy consumption reduces significantly (63%) due to the use of low-power processors. This is consistent with the 2.3× improvement in performance-per-watt of *LP-blade* over the *baseline*. (There are also some additional, relatively small, improvements in operational exergy consumption due to improved efficiency of shared fans in *LP-blade* relative to *baseline*.) The reductions in the infrastructure-operation exergy consumption are proportional to those in the server-operation exergy consumption (the change from going from rack-mounted servers to blades is relatively minor).

However, in spite of not explicitly targeting the production impact on sustainability, the *LP-blade* design *also reduces embedded exergy consumption by 38%*. To better understand where these benefits come from, Figure 7(a) shows a more detailed breakdown of the embedded exergy consumption across the three designs. Focusing on the *baseline* versus *LP-blade* data for now, we can see that the embedded exergy consumption improvements[2] for LP-blade come from two main factors: (1) the shared infrastructure associated with blade servers compared to rack-mounted servers (e.g., PSU, fan), but also (2) the reduced material associated with lower-power servers, both in reduced silicon overhead for the processor and in reduced material overhead for the printed circuit boards. Processor embedded exergy consumption reduces because the lower power processor choice in *LP-blade* leads to smaller total silicon area (in spite of *LP-blade* needing more of these simpler cores to achieve the same performance as *baseline*). The PCB size reduces due to the smaller form-factor *LP-blades*. These results indicate that prior studies that highlighted the environmental advantages of low-power servers [29, 32, 34, 35] primarily from an operational electricity reduction view point are conservative; a lifecycle methodology can identify additional environmental savings from these designs.

Optimizing embedded exergy consumption also benefits operational exergy consumption. Consider the *demat* design. As expected, the benefits in server-operation exergy consumption are not appreciably different from those with *LP-blade*; the additional improvements can be traced back to the energy benefits from memory disaggregation. The embedded exergy consumption reductions are higher compared to *LP-blade* (51% versus 38% relative to baseline). The breakdown in Figure 7(a) shows that these improvements match our intuition about the benefits from shared infrastructure and reduced sheet metal with

---

[2] Note that there are also (relatively small) degradations in embedded exergy associated with memory and hard disks due to the different parts used in the two real systems we use as reference for this study.

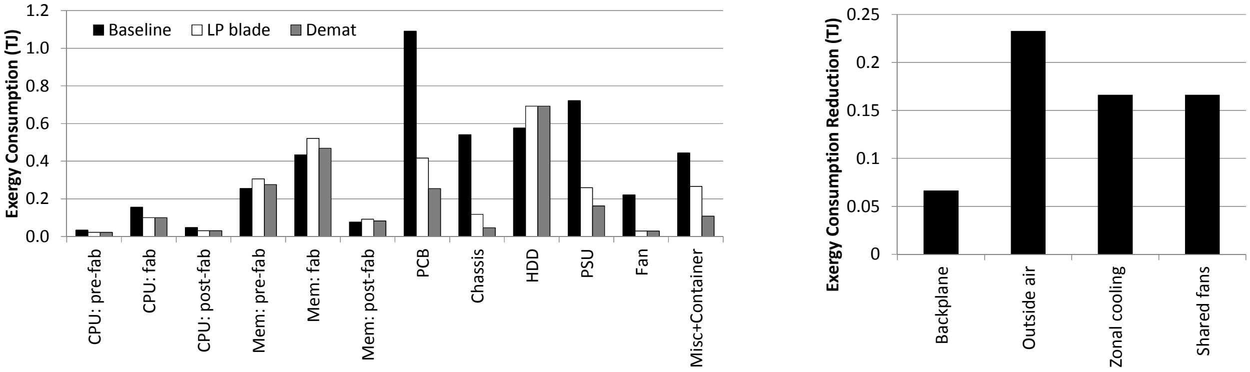

(a) **(b)**

Figure 7. (a) shows a comparison of embedded exergy consumption across the three evaluated server designs and (b) breaks down the reduction in operational infrastructure exergy consumption due to the cooling techniques employed in the *demat* server design.

dematerialization, and resource savings with memory disaggregation.

However, focusing on infrastructure-operation exergy consumption shows some interesting results. Figure 7(b) shows a breakdown of the attribution of various reasons for the benefits in infrastructure exergy consumption. Interestingly, only 36% of the total infrastructure improvements come from free-cooling's elimination of the thermal work associated with DX heat exchangers. The other benefits in infrastructure-operation exergy consumption actually arise as a consequence of optimizations that originally targeted embedded exergy consumption. For example, removing the backplane yields an 11% reduction in infrastructure exergy consumption (from better air flow); the use of shared fans yields an additional 28% improvement in infrastructure efficiency; and, the ability to cluster blades with similar thermal densities with disaggregation contributes to the remaining 25% improvement in cooling efficiency (by placing a low-density memory zone upstream of a high-density compute zone, air pre-heating is reduced).

A lifecycle-focused approach to understanding total environmental impact can identify and optimize for such interactions across different exergy consumption components.

3.2.3 Other results

To validate the physical aspects of the *demat* design and the cooling methodology, we also prototyped a "miniature dematerialized container" using custom PCB boards with heat sources to model different heat loads. Figure 8(a) shows a photograph of this miniature prototype. As an example validation experiment for the cooling models employed in our results, we used the prototype, in the configuration shown in Figure 8(a), to measure the temperature across the four blades in a cross section of the container. Our prototype doesn't include heat sinks and so is not an exact model of the system in Figure 5(b); nevertheless, the results in Figure 8(b) validate the good fit between our model and actual measurements.

We also studied the impact of our design on performance in the context of web workloads. The two main optimizations in our designs that can potentially impact performance are (1) low-power blades, and (2) memory disaggregation. Both these optimizations have been evaluated in prior studies and shown to not impact performance significantly.

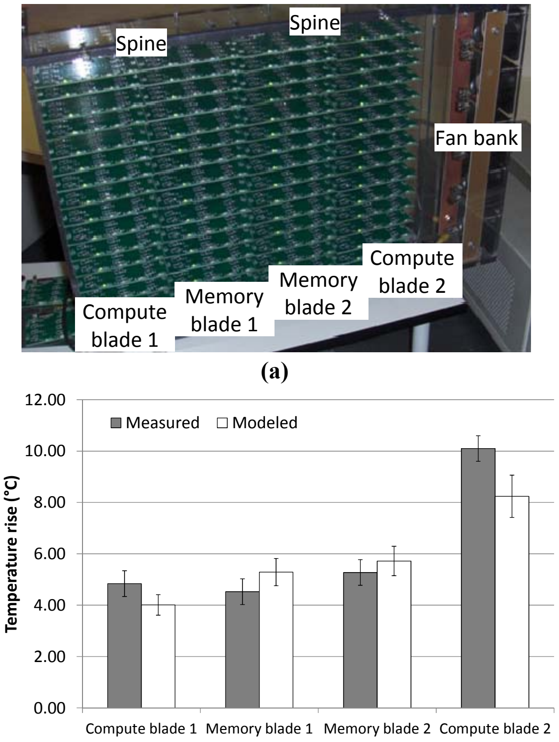

(a)

(b)

Figure 8. (a) shows our prototype miniature *demat* design used to validate physical aspects and cooling methodology. Custom PCB boards and fans can be controlled to generate heat and airflow, respectively. (b) shows the results of a validation experiment, comparing the prototype system (Measured) to our model (Modeled) in a compute-memory-memory-compute configuration.

For the use of low-power blades, we scaled their number to match the same throughput as the *baseline* and our low-power processors are chosen carefully to not violate latency constraints. For memory blades, we duplicated the simulation-based study in a prior paper on this topic [38]. We specifically considered four web workloads—*indexer*, *search*, *image manipulation*, and *MapReduce-based sort,* and a configuration with four DIMMs of local memory per compute blade supplemented by a shared memory blade. Our results are consistent with prior studies, and show that having such a two level memory hierarchy leads to, on average, only a 2% reduction in performance (with a range of 1–4%).

3.2.4 Discussion

There are several design assumptions we could study further. For example, in this paper, we chose a PUE of 1.5 for our *baseline* based on [28]. If we had instead started with an aggressively-optimized infrastructure having, for example, a PUE of 1.2 (most likely already including optimizations like free-cooling to achieve this), our infrastructure benefits would be different. However, even ignoring infrastructure exergy consumption, *demat* still reduces embedded and operational exergy consumption by 2.3× relative to *baseline* and by 1.2× relative to *LP-blade*. The embedded exergy consumption component improves by 25% going from *LP-blade* to *demat*. Similarly, we also studied systems with longer lifetimes (e.g., 5 years), and our results were qualitatively similar. As discussed earlier, rather than a comprehensive evaluation of the entire design space of specific design assumptions, our main goal in this paper was to understand key insights for computer architects when designing with lifecycle sustainability as a key constraint.

4. IMPLICATIONS OF LIFECYCLE ANALYSIS FOR SYSTEM DESIGNERS

So far, we have discussed a methodology that system architects can use to reason about the total environmental impact of their designs, and demonstrated the use of this methodology to evaluate specific designs and identify new tradeoffs. Next, we discuss how a lifecycle-based framework can have broader implications for future green server designs. Section 4.1 builds on the results so far and further makes the case for going beyond operational energy efficiency to target the full lifecycle when optimizing environmental impact. Section 4.2 discusses opportunities for new sustainability optimizations that can be evaluated with our framework.

4.1 The case for holistic lifecycle design

The results in Section 3.2 demonstrated positive interactions between different components of exergy consumption. Below we present another example, where optimizing operational energy alone may not yield the greenest design.

Energy proportionality versus consolidation: We compare two commonly-used categories of optimizations to reduce operational energy: (1) energy proportionality (*EP*) [4], optimizations that match power consumption to resource utilization (e.g., P-states); and (2) consolidation (*Con*), packing multiple workloads or virtual machines (VMs) into a smaller number of systems [40, 6]. Given the large body of work in this space, rather than focus on one particular implementation, we study idealized models that extrapolate these optimizations to the future. Such an approach is also better matched with our intent to explore broad tradeoffs between sustainability and operational energy. Therefore, for this experiment, we start with the *baseline* design from before; *EP* assumes ideal proportionality for the server (i.e., we assumed *baseline* hardware parameters from Figure 3, but modified the last column to zero); *Con* assumes perfect bin packing to minimize the number of servers assuming perfect knowledge and divisibility of workload patterns (the server configuration continues to match that of *baseline*).

Figure 9 illustrates the tradeoffs between *EP* and *Con* with data from 11 real-world enterprise clusters from Raghavendra et al. [7], representing a total of 180 servers. Columns 2 and 3 of the table show the operational energy for the *EP* and *Con* designs relative to the *baseline*. For all the traces, energy proportionality is always better than consolidation. This is because EP has perfect scaling of power with workload utilization. However, with *Con*, given that the bin-packing is optimized for the peak of the total utilization (*peak-of-sum utilization*), the average utilization of individual servers is less than 100% in most cases, even after correction for consolidation. Combined with the non-energy-proportional power profile of the *Con* servers, this leads to reduced cluster energy efficiency for *Con* relative to EP. Columns 4 and 5 show the total exergy consumption of the *EP* and *Con* designs relative to the *baseline.* In contrast to the results on operational energy, from a total exergy consumption perspective, *Con* is best for almost half the workloads. In these cases, the reduced server count with consolidation reduces embedded exergy consumption, an effect that *EP* does not have. These reductions (the reduction factor is equal to the reciprocal of the peak-of-sum utilization) offset *Con's* increased operational exergy consumption relative to *EP*, leading to *Con* ultimately being greener than *EP* in these cases. The specific cases where this effect is seen is a function of the workload behavior—its average and peak-of-sum utilizations. (Note that we assume that consolidation leads to lower provisioning of servers; if consolidation just allowed servers to be turned off, we would not get the savings in embedded exergy consumption.)

Workload	Operational-server (% baseline)		Total (% baseline)	
	EP	Con	EP	Con
E-commerce 1	18%	27%	36%	25%
E-commerce 2	48%	66%	57%	63%
Dotcom	37%	52%	49%	49%
Pharmacy	10%	17%	31%	16%
SAP 1	39%	50%	51%	46%
SAP 2	53%	84%	61%	82%
Worldcup 1	27%	61%	42%	60%
Worldcup 2	21%	31%	38%	28%
Consolidation 1	62%	88%	68%	87%
Consolidation 2	59%	88%	66%	86%
Animation farm	98%	100%	98%	100%

Figure 9. Analysis of the tradeoffs between optimizing for operational energy efficiency and optimizing for sustainability. Shaded cells denote the more energy-efficient or more sustainable system.

Co-design of inter-dependent variables: While this example is relatively intuitive, the results illustrate that *the most efficient system design for operational energy does not always correspond to the most sustainable solution*. The best way to optimize for sustainability is to use power-efficient *and* material-efficient systems that scale power with resources and are utilized fully.

More broadly, embedded exergy consumption, server-operation exergy consumption, infrastructure-operation exergy consumption, and performance are not independent variables. As an illustration, consider dematerialization discussed earlier. This optimization targeted at reducing embedded exergy consumption sometimes reduces infrastructure exergy consumption (when removal of sheet metal in the backplane enables better air flow), but in other cases increases infrastructure exergy consumption (when removal of fans in a server increases overall cooling exergy consumption). This optimization also impacts performance, for example, when backplane redesign for dematerialization impacts networking topologies.

Similarly, design optimizations are not independent, with one optimization enabling others. Going back to the dematerialization example, the departure from traditional notions of enclosures allowed us to increase sharing across a larger number of memory blades, and in turn, the disaggregation enables new optimizations around zones of clustering for optimal thermal density. These

Sustainability opportunity	Why it helps	Possible tradeoffs
CPUs with smaller die areas	Less silicon	Lower operational power, lower performance
Reducing number of CPUs	Less silicon	Lower operational power, lower performance
Upcycling CPUs	Less silicon	Need architectural support, performance implications
Shared CPUs	Shared silicon	Mainly good for consolidation-friendly workloads
Reducing memory	Less silicon/PCB	Lower operational power, lower performance
Shared memory	Shared silicon	Need architectural support, performance implications
Upcycling memory	Less silicon/PCB	Need architectural support, performance implications
Reducing fabrication energy	Less embedded exergy consumption	Need investment in new fabrication technology
Reducing PCB area	Less PCB	Need architectural support, cooling implications
Removing chassis materials	Less steel	Structural issues, need architectural support, cooling impact
Using fewer HDDs	Less copper/other	Lower storage capacity/bandwidth, lower operational power
Upcycling HDDs	Less copper/other	Need architectural support, performance implications
Shared ASICs (e.g., chipset, NIC)	Less silicon	Need architectural support, performance implications
Shared fans	Less materials/power	Potential cooling limitations
Air economizer cooling	Less infrastructure	Limits on heat load, material implications

Figure 10. Some example sustainability opportunities for future system designs and the reasoning from our model as to why they help and the possible tradeoffs involved.

examples argue for more holistic optimizations co-designed across system architecture and physical organization. Indeed, our results in Section 3.2 illustrate that the overall improvements in total exergy consumption are achieved through relatively small improvements in multiple components of the system, through several optimizations that work well with one another.

Consequently, when optimizing for sustainability, a *holistic* approach—co-designed across system architecture, physical organization and packaging, and cooling infrastructure—is needed. A lifecycle model can enable such optimizations.

4.2 Opportunities for future research

4.2.1 Targeting embedded exergy consumption

As embedded exergy consumption becomes a more dominant component of total exergy consumption—as it has for mobile systems or when operational exergy consumption is significantly optimized—*new approaches are needed that specifically target embedded exergy consumption*.

Upcyling: For example, the optimizations in the paper so far that target embedded exergy consumption are primarily along the lines of "reduce and reuse". However, there might be benefits from addressing "recycling" in architecture design as well.

Specifically, our *demat* design can potentially benefit from recycling. Compared to a traditional server where all components are decommissioned during an upgrade cycle, the separation of compute, memory, and storage enabled by disaggregation can support incremental upgrading and "upcycling" for lower exergy consumption. This ability to separate the refresh cycle for specific components like memory and disks from the server refresh cycle can amortize material overhead over longer times, enabling greener designs. For example, with separate memory blades, we could upgrade 75% of the memory with every processor upgrade, but upcycle the remaining 25%. Similarly, disks can be upgraded every other processor generation. Plugging such assumptions into our lifecycle model shows that such upcycling optimizations can indeed be beneficial providing an additional 27% reduction in embedded exergy consumption (or about 15% in total exergy consumption).

However, we did not include these optimizations in our *demat* results earlier for two reasons. First, upcycling has potential implications on performance that are hard to model. Memory speed increases every generation, as does memory density. While deferring the upgrades may offer embedded exergy consumption advantages, performance can potentially degrade due to the reduced speed and capacity. In the absence of specific workloads and deployment scenarios, such degradation is hard to characterize. Additionally, the energy efficiency of components has also traditionally changed every generation. The operational exergy consumption benefits from more energy-efficient memory post-upgrade may be important in the total exergy consumption benefits, offsetting upcycling savings. Reliability variations across generations also need to be factored.

Nevertheless, given the potential benefits, upcyling is worthy of further consideration. New architectural designs that enable upcycling, but without these negative effects, is a promising area of future research.

Silicon-efficient designs: One way to reduce the embedded exergy consumption further is to be more material-efficient. New architectures and fabrication processes that maximize the performance-per-exergy-consumption investment in silicon appear a promising area of research to improve the environmental impact from embedded exergy consumption. However, as discussed before, such optimizations need to be considered in the context of their interactions on other aspects of exergy consumption and net performance. Figure 10 lists other potential systems' optimizations to reduce embedded exergy consumption and their tradeoffs.

4.2.2 Targeting operational exergy consumption

Given the large fraction of operational exergy consumption in the overall environmental impact, *reducing operational energy* will continue to be an important way to make future systems greener. Energy-efficient system designs will continue to be important to reduce server-operational exergy consumption; these designs have a multiplicative effect, also reducing infrastructure exergy consumption and in some cases, even embedded exergy consumption. For infrastructure-operational exergy consumption, our results show that *co-designed physical system design and systems architectures* have significant potential. One area that co-designed thermal and system design can particularly help is around recovering exergy from waste-heat that is dissipated.

Renewable sources of energy (e.g., solar power, fuel cells, wind power) are another promising area for future optimizations. A few studies (e.g., [41, 42]) have already considered renewable energy in the data-center and capping approaches for "brown" energy management, but more research opportunities exist. In particular a lifecycle exergy consumption model can also factor the embedded exergy consumption tradeoffs in developing such new energy sources, helping understand net impact on sustainability. New system architectures to leverage alternate energy sources (a few are illustrated in ongoing work around co-designing battery sources with servers [43] or solar power with multicores [44]) will also be important.

4.2.3 Extending the scope of exergy factors

This study focused on the container as the unit of optimization and drew the boundary of exergy consideration at the container level. That is, our current models focus only on the material and energy *directly* related to production and operation of the *evaluated system.* However, a well-known effect within the environmental sciences is that of recursive and cross-dependent flows. For example, when a system requires some metals to be used, beyond the exergy consumption associated with extracting or manufacturing the metals themselves, an entire metals and mining sector of equipment, roads, employees, etc may have be created to support these manufacturing and extraction processes—additional exergy consumption factors that we currently do not model. Understanding how exergy consumption relates to cost is also an important area of future work, particularly when potential environmental costs—such as carbon taxation—are taken into account.

5. RELATED WORK

To the best of our knowledge, our work is the first to study *systems-focused* modeling and optimization of the *total lifecycle environmental impact* including both production and operation.

There are several studies in the environmental community that have studied lifecycle-based exergy analysis for sustainability [20, 27, 45, 26]. Our work builds on these studies, but develops a model better suited from a system architecture perspective. Various industrial approaches to reason about sustainability have been used by different companies (such as Apple, Dell, Google, and HP), but these are primarily based exclusively on metrics like carbon emissions and often do not address the entire lifecycle. Moreover, most of these approaches are not useful in an architectural context.

A few recent systems' studies have examined lifecycle sustainability. In addition to costs analysis, Anagnostopoulou et al. [12] also study the environment impact of the materials in different cooling configurations for large-scale datacenters. They model each configuration with a detailed parts list, and using the spec-sheets for the various parts, identify the weight of the infrastructure and break it down to the usage of major materials. Their results show that larger infrastructures are more materials-efficient than smaller datacenters. Oliver et al. [25] discuss reusing processors across different devices to amortize the production-related environmental impact of silicon over a longer time frame. They use a simple formula for the energy of semiconductor manufacturing, across the die and the assembly. In contrast to our work, neither of these studies have a unified sustainability model that allows comparison and optimization across the embedded and operational aspects.

A large number of prior systems' studies, including recent papers at specialized venues like SustainIT (`www.usenix.org/events/sustainit10`) and IGCC (`www.green-conf.org`), have focused on operational energy efficiency as a proxy for sustainability. In particular, there has been a large body of prior work on reducing the operational electricity consumption of servers [2, 3, 4, 5, 6, 7, 8, 9, 10, 11]. Other studies have examined renewable energy sources for operational energy. Gupta et al. [46] discuss a conceptual framework to reason about renewable energy sources that addresses workload arrival times, energy proportionality support, and environmental conditions. Le et al. [47, 48] and Qureshi [49] study optimization frameworks and policies to tradeoff "brown" and green energy sources. Govindan et al. examine how stored energy can be managed in datacenters [43]. Li et al. [44, 50] have studied the use of renewable energy in processor and datacenter design. However, none of these studies address the lifecycle environmental impact or the tradeoffs between the environmental impact of production and operation.

6. CONCLUSIONS

System architects are increasingly required to report on, and optimize, the environmental impact of their designs. Current approaches mainly focus only on operational energy, and miss additional important environmental impact caused by the production of the system. Addressing this limitation, our paper develops a framework for, and motivates optimizations around, a lifecycle focus for server architects designing future green servers. Our framework uses the thermodynamic metric of exergy consumption, adapted and validated for easy reasoning of architectural tradeoffs. Our results using the framework show that production-related impact can be important, contributing to 20% of the total environmental impact on current designs, and that system design choices can improve this component by 30-40%. Our results also demonstrate the need to holistically consider cross-interactions in the sustainability impact from production and operation when evaluating specific green techniques—in some cases, the most (operational) energy-efficient design is not the greenest (e.g., energy proportionality versus consolidation), and in others, optimizations targeting small embedded components of total exergy consumption achieve multiplicative benefits by enabling better environmental impact from operation (e.g., net environmental benefits from dematerialization).

However, this work only addresses the surface of the potential benefits from lifecycle optimizations. In addition to continued improvements in operational and infrastructure efficiencies, promising areas of future research include system architectures designed for renewable energy, architectures to better recover wasted heat, and new architectural approaches for better upcycling and improved performance-per-silicon-investment. Our lifecycle-based model can also be extended to broaden the scope of sustainability to look at the broader impact to society including social and economic costs of high performance in future computing systems. Also, our modeling approach of normalizing environmental impact with respect to of architectural building blocks can be applied to other metrics beyond exergy consumption as well. We hope that the contributions of this paper provide a foundation to study these, and other interesting, research questions for future sustainable and green system designs.

7. ACKNOWLEDGEMENTS

Wayne Mack and Tobin Gonzalez assisted in building the prototype discussed in this paper. Ricardo Bianchini, Fred Chong, Kanad Ghose, Anand Sivasubramaniam, Thomas Wenisch, and the anonymous reviewers provided valuable feedback that significantly strengthened this paper.

8. REFERENCES

[1] Chandrakant Patel. Sustainable information technology ecosystem. FAST Keynote Address, 2008.

[2] David Meisner, Christopher M. Sadler, Luiz André Barroso, Wolf-Dietrich Weber, and Thomas F. Wenisch. Power management of online data-intensive services. In *ISCA '11*, 2011.

[3] Susmit Biswas, Mohit Tiwari, Timothy Sherwood, Luke Theogarajan, and Frederic T. Chong. Fighting fire with fire: modeling the datacenter-scale effects of targeted superlattice thermal management. In *ISCA '11*, 2011.

[4] Luiz André Barroso and Urs Hölzle. The case for energy-proportional computing. *IEEE Computer*, 40(12):33–37, 2007.

[5] David Meisner, Brian T. Gold, and Thomas F. Wenisch. Powernap: Eliminating server idle power. In *Proceedings of the 14th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '09)*, Washington, DC, March 2009.

[6] Karthick Rajamani and Charles Lefurgy. On evaluating request-distribution schemes for saving energy in server clusters. In *Proceedings of the IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, pages 111–122, Austin, TX, March 2003.

[7] Ramya Raghavendra, Parthasarathy Ranganathan, Vanish Talwar, Zhikui Wang, and Xiaoyun Zhu. No "power" struggles: Coordinated multi-level power management for the data center. In *Proceedings of the 13th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, pages 48–59, Seattle, WA, March 2008.

[8] Parthasarathy Ranganathan, Phil Leech, David E. Irwin, and Jeffrey S. Chase. Ensemble-level power management for dense blade servers. In *Proceedings of the 33rd International Symposium on Computer Architecture (ISCA 2006)*, pages 66–77, Boston, MA, June 2006.

[9] Xiaobo Fan, Wolf-Dietrich Weber, and Luiz André Barroso. Power provisioning for a warehouse-sized computer. In *Proceedings of the 34th Annual International Symposium on Computer Architecture (ISCA '07)*, pages 13–23, San Diego, CA, June 2007.

[10] Taliver Heath, Ana Paula Centeno, Pradeep George, Luiz Ramos, Yogesh Jaluria, and Ricardo Bianchini. Mercury and freon: Temperature emulation and management for server systems. In *Proceedings of the 12th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, pages 106–116, San Jose, CA, October 2006.

[11] Faraz Ahmad and T. N. Vijaykumar. Joint optimization of idle and cooling power in data centers while maintaining response time. In *ASPLOS*, 2010.

[12] Vlasia Anagnostopoulou, Heba Saadeldeen, and Frederic T. Chong. Quantifying the environmental advantages of large-scale computing. In *First International Green Computing Conference (IGCC)*, 2010.

[13] Greenpeace. "How dirty is your data? A look at the energy choices that power cloud computing." Greenpeace International Report, April 2011.

[14] Mireille Faist Emmenegger, Rolf Frischknecht, Markus Stutz, Michael Guggisberg, Res Witschi, and Tim Otto. Life cycle assessment of the mobile communication system: Towards eco-efficient systems. *International Journal of Life Cycle Assessment*, 11:4, 2006.

[15] Amip J. Shah, Chandrakant D. Patel, and Van P. Carey. Exergy-based metrics for sustainable design. In *Fourth International Exergy, Energy and Environment Symposium (IEEES-4)*, Sharjah, UAE, 2009.

[16] J. Glasson, R. Therivel, and A. Chadwick. *Introduction to Environmental Impact Assessment: Principles and Procedures, Process, Practice and Prospects (3rd Edition)*. Taylor & Francis, 2005.

[17] O. Jolliet, M. Margni, R. Charles, S. Humbert, J. Payet, G. Rebitzer, and R. Rosenbaum. Impact 2002+: A new life cycle impact assessment methodology. *Int. J. Life Cycle Assessment*, 8:324–330, 2003.

[18] M. Goedkoop and R. Spriensma. The eco-indicator 99: A damage oriented method for life cycle assessment. Technical report, PRé Consultants, 2001.

[19] A. Horvath, C. T. Hendrickson, L. B. Lave, F. C. McMichael, and T.-S. Wu. Toxic emissions indices for green design and inventory. *Environment Science & Technology*, 29:86A–90A, 1995.

[20] Amip Shah and Kiara Corrigan. Extending lifecycle exergy analysis beyond resource consumption. In *ASME Energy Sustainability Conference*, Phoenix, AZ, May 2010.

[21] Rolf Frischknecht, Niels Jungbluth, Hans-Jörg Althaus, Gabor Doka, Roberto Dones, Thomas Heck, Stefanie Hellweg, Roland Hischier, Thomas Nemecek, Gerald Rebitzer, and Michael Spielmann. The ecoinvent database: Overview and methodological framework. *International Journal of Life Cycle Assessment*, 10(1):3–9, 2005.

[22] Y. Anny Huang, Chris L. Weber, and H. Scott Matthews. Carbon footprinting upstream supply chain for electronics manufacturing and computer services. In *IEEE International Symposium on Sustainable Systems and Technology (ISSST'09)*, Tempe, AZ, May 2009.

[23] Sarah Boyd, Arpad Horvath, and David Dornfeld. Life-cycle energy demand and global warming potential of computational logic. *Environmental Science and Technology*, 43(19):7303–7309, 2009.

[24] Timothy Gutowski, Matthew S. Branham, Jeffrey B. Dahmus, Alissa J. Jones, Alexandre Thiriez, and Dusan Sekulic. Thermodynamic analysis of resources used in manufacturing processes. *Environmental Science and Technology*, 43:1584–1590, 2009.

[25] John Oliver, Rajeevan Amirtharajah, Venkatesh Akella, Roland Geyer, and Frederic T. Chong. Life cycle aware computing: Reusing silicon technology. pages 58–63, 2007.

[26] David J. Lettieri, Christopher R. Hannemann, Van P. Carey, , and Amip J. Shah. Lifetime exergy consumption as a sustainability metric for information technologies. In *IEEE International Symposium on Sustainable Systems and Technology (ISSST'09)*, Tempe, AZ, May 2009.

[27] Chris R. Hannemann, Van P. Carey, Amip J. Shah, and Chandrakant D. Patel. Lifetime exergy consumption as a sustainability metric for enterprise servers. In *ASME Energy Sustainability Conference*, Jacksonville, FL, 2008.

[28] James Hamilton. Internet-scale service infrastructure efficiency. In *Keynote at the International Symposium on Computer Architecture*, 2009.

[29] Kevin Lim, Parthasarathy Ranganathan, Jichuan Chang, Chandrakant Patel, Trevor Mudge, and Steven Reinhardt. Understanding and designing new server architectures for emerging warehouse-computing environments. In *ISCA*, 2008.

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

[31] Amip Shah, Cullen Bash, and Chandrakant Patel. Optimizing data center cooling infrastructure using exergothermovolumes. In *ITherm*, 2010.

[32] James Hamilton. Cooperative expendable micro-slice servers: Low cost, low power servers for internet-scale services. In *CIDR*, 2009.

[33] HP performance-optimized datacenter (POD). http://h71028.www7.hp.com/enterprise/cache/595887-0-0-0-121.html.

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

[35] David Andersen, Jason Franklin, Michael Kaminsky, Amar Phanishayee, Lawrence Tan, and Vijay Vasudevan. Fawn: a fast array of wimpy nodes. In *SOSP*, 2009.

[36] HP bc2500 product page. http://h20000.www2.hp.com/bizsupport/TechSupport/Document.jsp?objectID=c01125713&lang=en&cc=us&taskId=101&prodSeriesId=3436097&prodTypeId=12454.

[37] Eitan Frachtenberg, Ali Heydari, Harry Li, Amir Michael, Jacob Na, Avery Nisbet, and Pierluigi Sarti. High-efficiency server design. In *SC '11*, 2011.

[38] Kevin Lim, Jichuan Chang, Trevor Mudge, Parthasarathy Ranganathan, Steven K. Reinhardt, and Thomas F. Wenisch. Disaggregated memory for expansion and sharing in blade servers. In *ISCA*, 2009.

[39] Rich Miller. Microsoft's chiller-less data center. http://www.datacenterknowledge.com/archives/2009/09/24/microsofts-chiller-less-data-center/, September 2009.

[40] Ripal Nathuji and Karsten Schwan. Virtualpower: coordinated power management in virtualized enterprise systems. In *Proceedings of the 21st ACM Symposium on Operating Systems Principles (SOSP)*, pages 265–278, Stevenson, WA, October 2007.

[41] Kien T. Le, Ricardo Bianchini, Margaret Martonosi, and Thu D. Nguyen. Cost- and energy-aware load distribution across data centers. In *HotPower*, 2009.

[42] Christopher Stewart and Kai Shen. Some joules are more precious than others: Managing renewable energy in the datacenter. In *HotPower 2009*, 2009.

[43] Sriram Govindan, Anand Sivasubramaniam, Bhuvan Urgaonkar. Benefits and Limitations of Tapping into Stored Energy For Datacenters. In *ISCA*, 2011.

[44] Chao Li, Wangyuan Zhang, Chang-Burm Cho, and Tao Li. Solarcore: Solar energy driven multi-core architecture power management. In *HPCA*, 2011.

[45] Chris R. Hannemann, Van P. Carey, Amip J. Shah, and Chandrakant D. Patel. Lifetime exergy consumption of an enterprise server. In *IEEE International Symposium on Electronics and the Environment (ISEE'08)*, San Francisco, CA, May 2008.

[46] Sandeep K. S. Gupta, Tridib Mukherjee, Georgios Varsamopoulos, and Ayan Banerjee. Research directions in energy-sustainable cyber-physical systems. *Elsevier Comnets Special Issue in Sustainable Computing (SUSCOM)*, 2011.

[47] Kien Le, Ozlem Bilgir, Ricardo Bianchini, Margaret Martonosi, and Thu D. Nguyen. Capping the brown energy consumption of internet services at low cost. In *International Green Computing Conference (IGCC 2010)*, 2010.

[48] Kien Le, Ozlem Bilgir, Ricardo Bianchini, Margaret Martonosi, and Thu D. Nguyen. Managing the cost, energy consumption, and carbon footprint of internet services. In *ACM SIGMETRICS*, 2010.

[49] Asfandyar Qureshi, Rick Weber, Hari Balakrishnan, John Guttag, and Bruce Maggs. Cutting the electric bill for internet-scale systems. In *Proceedings of the ACM SIGCOMM 2009 conference on Data communication*, 2009.

[50] Chao Li, Amer Qouneh, and Tao Li. Characterizing and analyzing renewable energy driven data centers. In *International Conference on Measurement and Modeling of Computer Systems*, 2011.

A1. Thermo-volume infrastructure modeling for new packaging/thermal designs

To reason about cooling improvements, we employ the method developed and validated by Shah et al. [31]. This approach essentially divides the physical infrastructure into "control volumes" of tractable thermofluidic properties, referred to as "thermovolumes." We use CFD models to model air flow for individual blades, determine thermal and flow resistances for each thermovolume in the spine and datacenter, and perform a closed-form optimization to maximize the coefficient of performance (COPG). COPG is defined as the heat removed per unit amount of work input to a given thermovolume; a larger COPG implies lower cooling power spent in removing heat. In situations (such as ours) when power delivery losses are minimal, PUE is approximately equal to (1 + 1 / COPG).

$$\text{Cooling power} = Q_{dc}/COP_G$$

- Q_{dc} is the heat dissipation from the datacenter
- COP_G is the coefficient of performance (efficiency of cooling) and is defined as

$$COP_G = (Q_{dc}/W_{comp})/(1 + A + B + C + D)$$

- W_{comp} is the compressor power for thermal work
- $A = \frac{\text{Total power for cooling infrastructure}}{\text{Compressor power}}$
- $B = \frac{\text{Power for blowers and CRAC units}}{\text{Compressor power}}$
- $C = \frac{\text{Pump power}}{\text{Compressor power}}$
- $D = \frac{\text{Cooling tower power}}{\text{Compressor power}}$

(For our free air cooling design, $B = C = D = 0$.)

$$W_{comp} = f(\text{supply temp from CRAC efficiency curve})$$

$$A = g(\text{flow rate, fan operating curve, system pressure } \Delta)$$

Figure 11. Key thermal model parameters.