Proceedings of the ASME 2010 International Mechanical Engineering Congress & Exposition
IMECE2010
November 12-18, Vancouver, British Columbia, Canada

IMECE2010-39340

LIFECYCLE-BASED DATA CENTER DESIGN

Justin Meza, Rocky Shih, Amip Shah, Parthasarathy Ranganathan, Jichuan Chang, Cullen Bash
Hewlett-Packard Co. Laboratories
1501 Page Mill Road
Palo Alto, CA 94304

ABSTRACT

Environmental sustainability is an increasingly important design constraint for next-generation servers and datacenters. Unlike prior studies that focus on operational energy use, we study the environmental impact of current designs across the entire lifecycle, including embedded impact factors related to material use and manufacturing. Based on the insights provided by this study, we propose a solution co-designed across system architecture and physical packaging, including (1) material-efficient physical organization, (2) environmentally-efficient cooling infrastructures, and (3) effective design of system architectures to reuse components – all working together to improve sustainability. We provide a detailed evaluation of our proposed solution in terms of sustainability, thermal manageability, and computational performance. Our results show that the proposed approach is effective in addressing the (often non-intuitive) tradeoffs between performance and different components of sustainability.

INTRODUCTION

The manufacturing, operation, and disposal of products to minimize their environmental impact in terms of destruction of natural resources or production of undesired emissions are fast becoming an important design constraint for Information Technology (IT) systems [1]. The carbon footprint of the IT industry, though only 2% of the total economy, is estimated to be equal to that of the entire aviation industry [2]. For example, in 2006, datacenters in the US alone were estimated to consume more than 60 billion kWh of electricity, translating to a correspondingly large carbon footprint [3]. Furthermore, regulatory and competitive pressures will also increase the push to design more sustainable systems. A recent estimate is that up to 75% of organizations will soon consider sustainability as one of the criteria in their IT purchases [4].

As shown in Fig. 1a, a typical data center consists of a row of computing racks which draw in cold air from a supply plenum. Each of these racks contain a large number of individual computer servers along with auxiliary storage and networking equipment. The electricity used by the servers, storage and networking gear is dissipated in the form of heat, which gets picked up by the cold air and is returned to the Computer Room Air Conditioning (CRAC) unit for further refrigeration. Traditionally, such data centers have been deployed in large warehouses containing thousands of servers. More recently, however, the ability to rapidly assemble data centers for scaling out an enterprise compute infrastructure has given rise to the notion of ‘containerized’ data centers. Essentially, these types of data centers consist of a shipping container (as shown in Fig. 1b). Such a container may hold two rows of IT equipment from Fig. 1a, so that cold air is delivered to the intake of each row and exhausted through a single hot aisle. Multiple such containers may be placed side-by-side to meet the computational demands of large numbers of users.

There has been a large body of prior work on reducing the operational electricity consumption of servers and data centers. Given that most of the electricity produced in the world comes from carbon-intensive sources, these optimizations can help reduce the carbon footprint of servers and datacenters. However, these approaches do not address environmental aspects related to the extraction of raw materials, manufacturing, transportation, operation, and disposal.

In this paper, we examine the problem of lifecycle-based optimization of future servers and datacenters. Based on a systematic analysis of the environmental impact of current designs across the entire lifecycle, we propose a new datacenter solution that incorporates novel approaches to address material and infrastructure impact on sustainability. This solution is co-designed with the computer system architecture to address performance-related constraints. To the best of our knowledge, this is the first such work that examines lifecycle-based designs for datacenters in such holistic fashion.

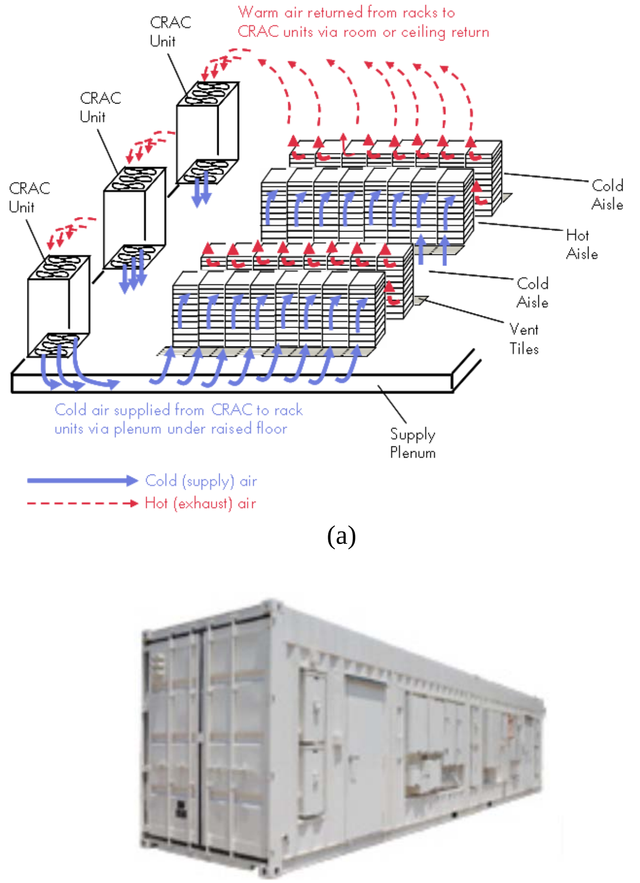

(a)

(b)

Figure 1. (a) Example of a typical raised floor data center; (b) a containerized data center.

ASSESSMENT OF EXISTING SYSTEMS

Numerous schemes exist to quantify the environmental impact of computer systems. Life-cycle assessment (LCA), a field that has been in practice for nearly 50 years [5,6], involves taking an end-to-end approach to assessing environmental impact across various lifecycle stages, such as the extraction of raw materials, manufacturing, transportation, operation, and disposal. In this paper, we utilize lifecycle assessment using the thermodynamic metric of *exergy* (available energy) consumption to reason about sustainability. Several previous studies have elaborated on how destruction of exergy is representative of the irreversibility associated with various processes [7,8] and correspondingly, to a first order, the environmental sustainability [9]. In addition, prior work has successfully developed lifecycle exergy consumption models for select instantiations of IT systems [10], and shown that optimizations based on lifecycle exergy consumption often map fairly well to optimizations based on other types of environmental criteria such as greenhouse gas emissions, pollution, etc [11], particularly for systems where the magnitude of material and energy use are the primary drivers of environmental impact (as opposed to the type of materials or energy use).

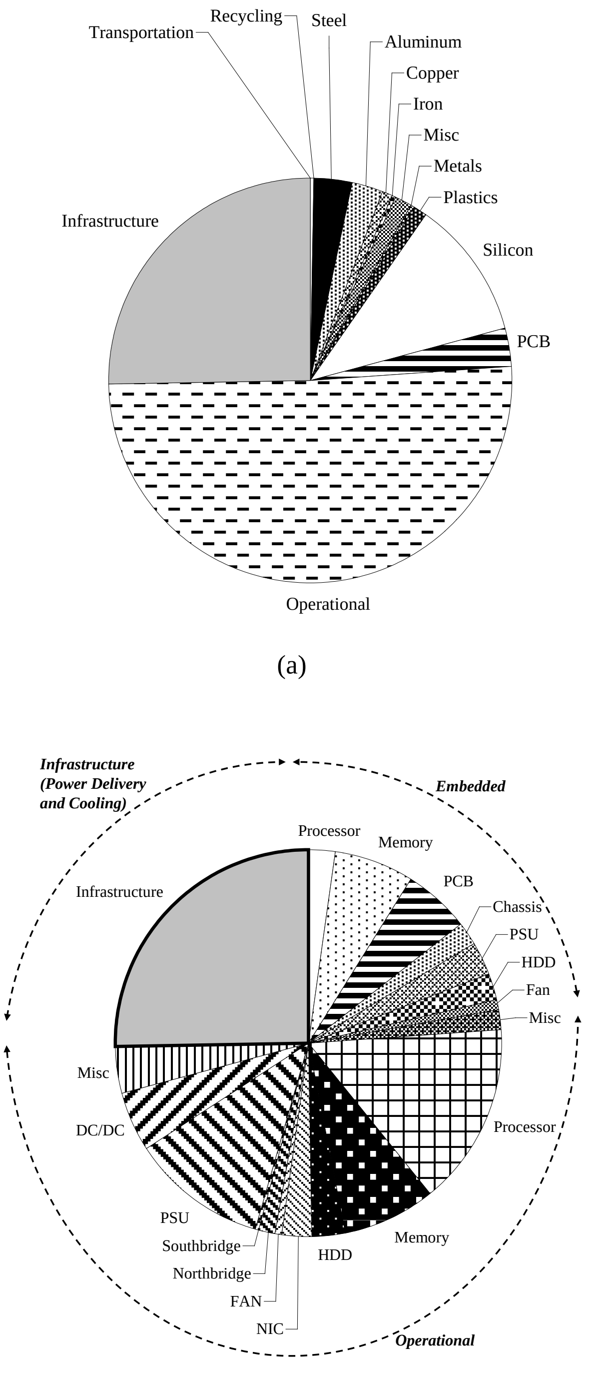

(a)

(b)

Figure 2. Breakdown of an existing computer server, (a) by material; and (b) by component.

The primary advantage of such an exergy-based analysis, relative to comparable approaches based on the first law, is that the former provides for a framework which quantitatively enables inclusion of issues related to both material and energy use. In traditional approaches, as an example, one may be required to trade off the mass of materials (in kg) relative to the electricity used during runtime (in kWh) to reach an optimal design point. By situating such optimizations on the common platform of exergy consumption, however, the savings resulting from using fewer materials is quantified in the same units as the savings resulting from energy efficiency savings. Thus, it becomes possible to evaluate such trade-offs quantitatively.

Unfortunately, previous lifetime exergy characterizations have estimated the total environmental impact of electronics in computer systems based on a mapping of the system mass or material flows to per-unit estimates of the environmental impact burden [10,12]. Figure 2(a) shows such a breakdown for a typical server (2-socket Xeon-based server with 4GB DRAM and two 72GB HDDs, two 1-Gb/s NICs, assuming 25% average utilization and a PUE of 1.6 [13], where PUE= [operational power + infrastructure power] / operational power) following the approach of Hannemann et al. [10]. Such breakdowns are typically achieved through system disassembly. Each stage of the lifecycle is broken down into its key components, and each component is then measured in terms of its constituents. For example, in Fig. 2a, the cradle-to-gate stages (which include extraction of raw materials, manufacturing, and transportation) is broken up into underlying materials such as silicon, plastics, metals, etc. It should be noted that for convenience, the manufacturing processes associated with these materials are included in the same category, viz. the ‘silicon’ legend includes environmental impacts related to both the extraction of silicon as well as the processing of silicon into its final electronically active state. Similarly, the runtime phase of the system is broken down into the exergy consumed during operation of the system as well as the exergy consumed by the supporting infrastructure which must operate to support the system as well. Assuming a datacenter container or “pod” with 1056 of these servers, the total exergy consumption is 26.8 TJ over a three year timeframe.

While such a decomposition of the lifetime exergy consumption of a system is useful, extending such a breakdown of exergy consumption to system architecture choices is not clear. For example, if one wishes to reduce the amount of silicon used in the system, it is not apparent whether a better choice would be to reduce the amount of silicon in the processor or in the memory (both of which contain similar materials). Similarly, it is not obvious how a choice around the amount of silicon in memory or processor for a particular system might relate to the runtime exergy consumption of the system, as these components are interlinked. In other words, since architectural choices may span multiple stages of the entire system lifecycle, deciding to use one component over another in a system based on the type of breakdown shown in Fig. 2a may result in (often non-intuitive) changes to the total system environmental impact due to their 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 instead considers lifecycle exergy consumption from an architectural perspective is required. Such an architectural view of the system lifetime exergy consumption is shown in Fig. 2b for the same baseline server as before.

For convenience, we categorize exergy consumption into three broad categories – embedded, operational, and infrastructure. *Embedded exergy consumption* 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 consumption is due to complicated manufacturing processes for highly-ordered electronic components, and the various chemicals required for making these components themselves also require large amounts of energy to manufacture. While a true process-based analysis of each component would require estimating the exergy destroyed during each of the processes specific to that component, we have instead abstracted impact factors based on available data [12-20].

Operational exergy consumption is the amount of exergy consumed by the system during its operational lifetime. In this study, we assume that this is the same as the *energy* consumed during operation. (Although the heat dissipated from the server contains useful work potential, there are currently no effective techniques to use this waste heat and recover this exergy.) To determine the operational energy used by each component, we determine its maximum power rating and model how its power varies with resource usage. 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 [21] and allows us to have a high-order model for the power consumed across different workloads (varying utilizations).

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 consumption* 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 infrastructure. (Note that server fans are considered part of the server operational power.) We assume that cooling is provisioned appropriate to the maximum power rating and employ the widely-used power usage effectiveness (PUE) metric [22] to compute infrastructure exergy consumption. The exergy consumption related to building the power and cooling infrastructure in the datacenter is outside the scope of this study; moreover, when amortized over the several thousand servers that are supported by the facility, the incremental contribution of the exergy consumption embedded within the infrastructure is generally small.

The benefit of the approach underlying Fig. 2b, which is primarily geared towards portraying an appropriate disaggregation of cumulative exergy consumption data, is that a system designer can easily understand how different architectural choices relate to each other. For example, while the overall breakdown of both Fig. 2a and Fig. 2b is identical – the operational exergy consumption dominated the total exergy consumption of the system (51%), followed by infrastructure exergy consumption (25%), and embedded exergy consumption (24%) – Fig. 2b provides additional insight around how different architectural choices in terms of memory, processor design, etc. may affect other stages of the lifecycle. The next section explores conceptual redesigns of the system based on the insight gained from the above analysis.

LIFECYCLE-BASED SYSTEM REDESIGN

We consider system redesign in three stages: embedded design; operational design; and infrastructural design. It is important to note that these are not independent variables; therefore, it becomes necessary to embrace a holistic design approach. We propose evaluating different design trade-offs in terms of *performance per total exergy consumption* to ensure that localized changes in one aspect of the architecture do not negatively impact changes elsewhere in the system. We employ a structured approach to this evaluation by following the method of exergothermovolumes developed by Shah et. al. [22]. First, however, we discuss several design approaches within each of the three categories discussed above.

Embedded Design

The embedded exergy consumption associated with system architecture choices can be addressed either by simplifying the manufacturing processes (outside the scope of this paper; but also hard to redesign given huge existing investments in current processes) or by reducing components. It follows from Fig. 2a and Fig. 2b that the most effective way to reduce embedded exergy consumption within components is to reduce the amount of materials used. For example, a smaller memory configuration would use less silicon and consequently reduce the embedded exergy consumption associated with memory at all the process stages. We call this approach at reducing materials *"dematerialization."* From an architectural standpoint, existing systems tend to be fully loaded with all the necessary functionality. For example, each system may have its own

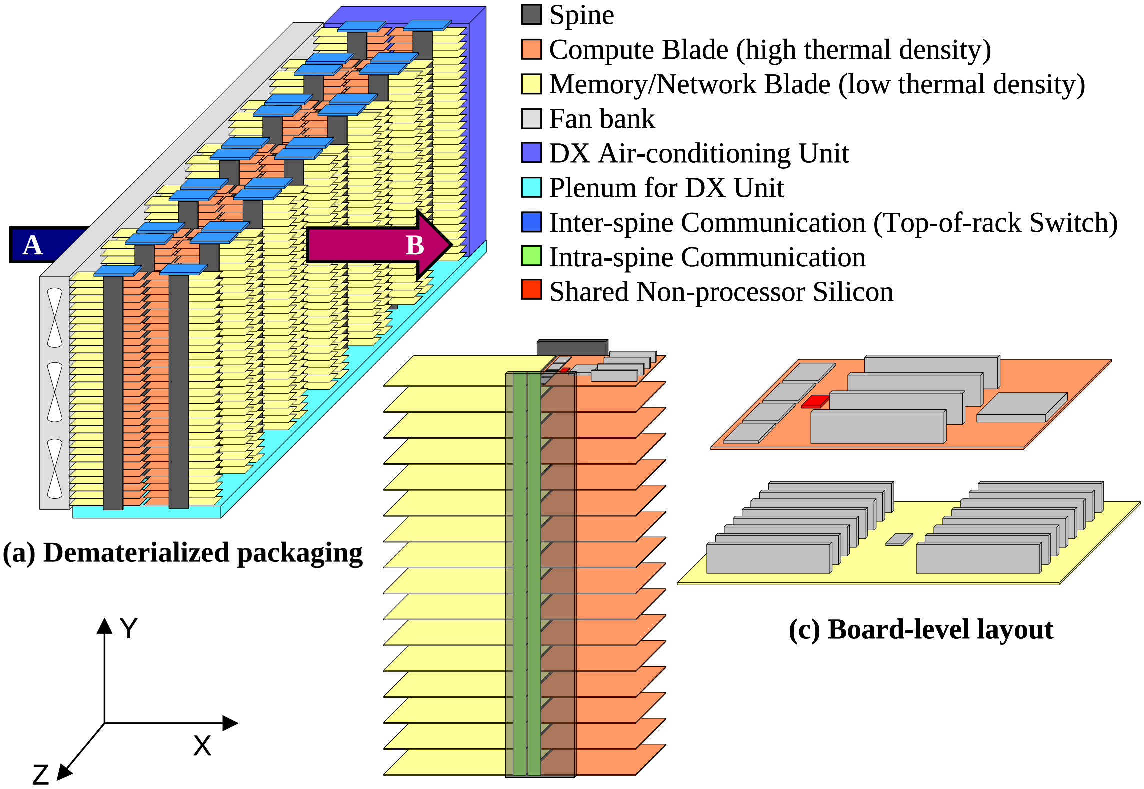

(a) Dematerialized packaging

(c) Board-level layout

(b) Spine as the communication and power delivery backbone

Figure 3. Conceptual sketch of proposed data center design.

dedicated memory, storage, fans, etc. Most of these systems have been provisioned for peak capacity, but in practice, these systems rarely see full load. As a result, in a data center where thousands of such systems have been deployed, there is significant wastage of material when most of the components are only being utilized at partial load. Instead, by only providing the materials which are necessary to meet average load and pooling resources from multiple systems together in order to meet rarely occurring peak load requirements, a dematerialized system design may be achieved.

Reusing components when they would normally be recycled or discarded (*"upcycling"*) is another effective way of reducing embedded exergy consumption, as this would essentially amortize the destruction of exergy over multiple generations of systems (i.e., we gain an offset credit for avoiding the need to remanufacture components in a next-generation system). Designing system architectures that enable updates to the system to be localized only to functionality that needs to be updated provides more opportunities for upcycling. For example, in current system architectures, the entire motherboard (containing all the electronic components) needs to be disposed at end-of-life. Often, however, additional functionality in the next-generation system is only desired for a specific component – such as a higher frequency (faster) processor. Then, by designing *disaggregated* systems where the memory and the processor have been separated on to different physical media, we could enable upcycling of the memory in cases where only the processor needs to be upgraded.

Given the above two principles, in order to reduce the embedded exergy consumption of the physical infrastructure, we consider the "container" as the packaging volume and seek to keep the amount of physical materials to a minimum. We propose a new "boxless" design (Fig. 3a) that just consists of multiple "spines" with individual printed circuit boards (effectively, small form-factor blade servers) attached to the spine. The spines provide communication and power delivery across different elements, and are perforated for air flow. Such a design significantly eliminates sheet-metal in the datacenter, reducing the embedded exergy consumption associated with sheet metal, PCBs, and box-level fans. They also allow improved airflow and cooling (discussed later). The design also seeks to aggressively share resources or dynamically use shared resource pools to minimize the amount of material required (Fig. 3b). Disaggregation allows for shared resources (e.g., shared memory pools, shared storage) and consequently lower provisioning of material. Our physical organization also allows for shared power delivery and cooling. In addition, we also optimize intra-blade system architecture to share resources that do not have too much contention – for example, Fig. 3c shows a compute blade optimized to share non-processor silicon such as the IO hub, management processor, NIC, etc. In addition, by using separate compute blades, storage blades, network blades, etc., the proposed design allows for specific components to be individually replaced. We also leverage recent proposals for separate "memory" blades [24] to enable memory disaggregation. Given that memory is a dominant contributor to the embedded exergy consumption in existing systems, the ability to separate the memory refresh cycle from the server refresh cycle enables new upcycling opportunities.

Operational Design

Numerous efforts already exist to facilitate reduction of energy use within servers [25-38]. We leverage these existing techniques. In particular, we take advantage of the ability to *consolidate* workloads aggressively. The intuition is that typical utilization on many enterprise services is relatively low, and that across a collection of systems, peaks are often unsynchronized. This allows multiple machines (or tasks in a task scheduler) to be consolidated on to a single server to reduce the number of required servers (and consequently their power usage). In addition, there have recently been several proposals and solutions around *energy-efficient servers* based on lower-power processors. 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 at the processor level. We utilize these approaches previously proposed in the literature in the current design. For example, the compute components shown in Fig. 2c are similar to the low-power components proposed in the recent literature referenced above.

Cooling Infrastructure

To reduce infrastructure exergy consumption, we utilize airside economization to minimize the requirement for cooling infrastructure for the data center. This approach uses ambient outside air to cool the datacenter without the need to expend work on chilling outside air. Figure 3a shows how air from the external environment can be brought in (point A) and exhausted (point B). Note that the design includes a direct expansion (DX) air-conditioning unit, which can refrigerate the supply-side air in case of deployment in high-temperature ambient environments. While the DX unit is not necessarily required, we find that inclusion of such air-conditioning capability provides an additional control knob. As discussed later, this helps optimize the balance between flow work and thermodynamic work for overall sustainability improvements. In addition, to minimize the inlet air requirements for each system blade described in Fig. 3, we include an additional layout optimization. Specifically, the blades attached to the spines in each aisle follow a specific sequence – memory blade, compute blade, compute blade, memory blade (light and dark-colored blades in Fig. 3). The rationale behind such an arrangement is to take advantage of differing component temperature requirements. In traditional data center design, the desired temperature to be maintained is an inlet temperature to the systems, and the cooling infrastructure must support this specification. However, because we have disaggregated the system components in the present design, we are able to specify higher inlet temperatures than might otherwise be suggested. Specifically, at typical system flow rates, a memory blade can be cooled with less than 2 °C rise across the blade. In comparison, in existing architectures, each system or blade contains a high-power compute node and lower power memory which causes 15-25 °C rise in temperature for each system. Thus, by placing a low-density zone upstream of a high density compute zone, we are able to minimize preheating of the air.

Thus, based on the lifecycle assessment performed previously, we obtain a new data center design that takes into account the potential for both architectural and thermal enhancements. In the next section, we evaluate the performance of this design relative to traditional architectures.

EVALUATION

To evaluate the proposed datacenter design, we look at the metrics of total exergy consumption and performance-per-total-exergy-consumption. For performance evaluation, we use the standard practice of simulating the computational behavior of the system against representative benchmarks. Before looking at these overall metrics, however, we first explore the differences between the proposed and existing design in terms of embedded, operational, and infrastructural considerations.

Infrastructural evaluation

Figure 4a shows the temperature profile heatmap for a compute blade server dissipating 245 W with 0.0283 m^3/s (60 CFM) of air moving through the server. We note that the maximum case temperature is about 46 °C, well below the specified allowable (design) case temperature of 75 °C (which in itself is conservative). Figure 4b shows the temperature rise from point A to point B in Fig. 3 for a system that has been optimized for minimum fan work without exceeding the allowable case temperature. Figure 4c shows COP_G for different outside air temperatures. COP_G, the 'grand' coefficient of performance, is essentially a measure of the heat removed from the infrastructure as a function of the *total* infrastructure cooling power [39]. COP_G is similar to the previously discussed PUE metric, but eliminates consideration of power delivery losses within the infrastructure. (Note that higher COP_G is better, but lower PUE is desirable.) As the outside air temperature increases, the difference between the inlet temperature and the maximum allowable case temperature is reduced. To compensate for this reduced thermal margin, the fans are required to deliver a higher flow rate of air to maintain system temperatures within operating limits. This causes the infrastructure efficiency to decrease as outside air temperature increases. Because the fan power scales non-linearly with the air speed, after some threshold outside air temperature, it is more efficient for the infrastructure to rely upon air-conditioning (DX units) rather than further increasing fan speeds; this is where the slope of the COP_G curve flattens out.

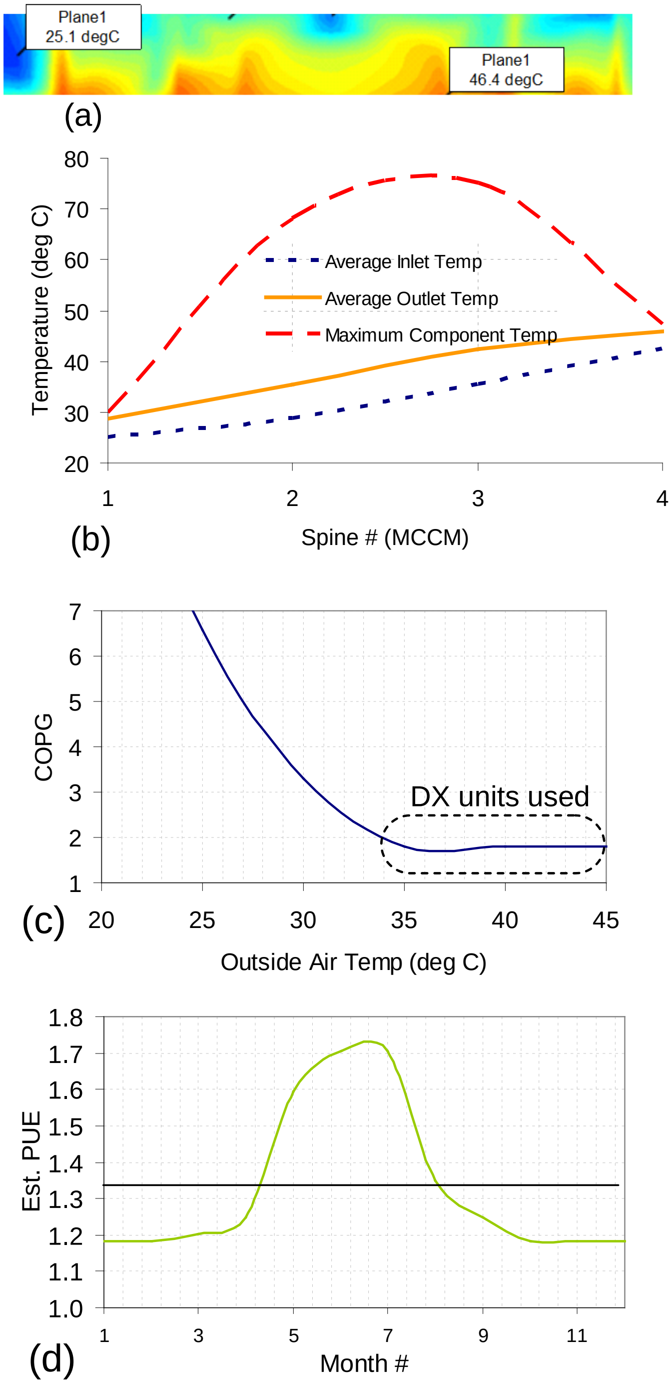

Figure 4. Evaluation of thermal efficiency of proposed design. (a) Temperature heatmap of a compute blade at 25 °C inlet temperature; (b) Temperature profile across an array of spines containing compute, memory and storage blades; (c) COP curve for datacenter containing such spines; (d) Simulation of PUE as a function of time for datacenter (flat line represents annual average for environment considered).

Having obtained the optimal COP_G curve, we then simulate the data center efficiency in terms of PUE over an entire year for a location where the outside air temperature ranges from wintertime lows of around 0 ^oC to summertime highs of 45 ^oC, which encompasses a wide range of locations around the globe. The result is shown in Figure 4d, where the annually averaged PUE is about 1.3. Note that if a more favorable location had been chosen – such as one where the outside air temperature mostly ranged between 0 ^oC and 30 ^oC (e.g., San Francisco; many places in the UK or Western Europe; etc.) – then the annually-averaged PUE would drop to below 1.2.

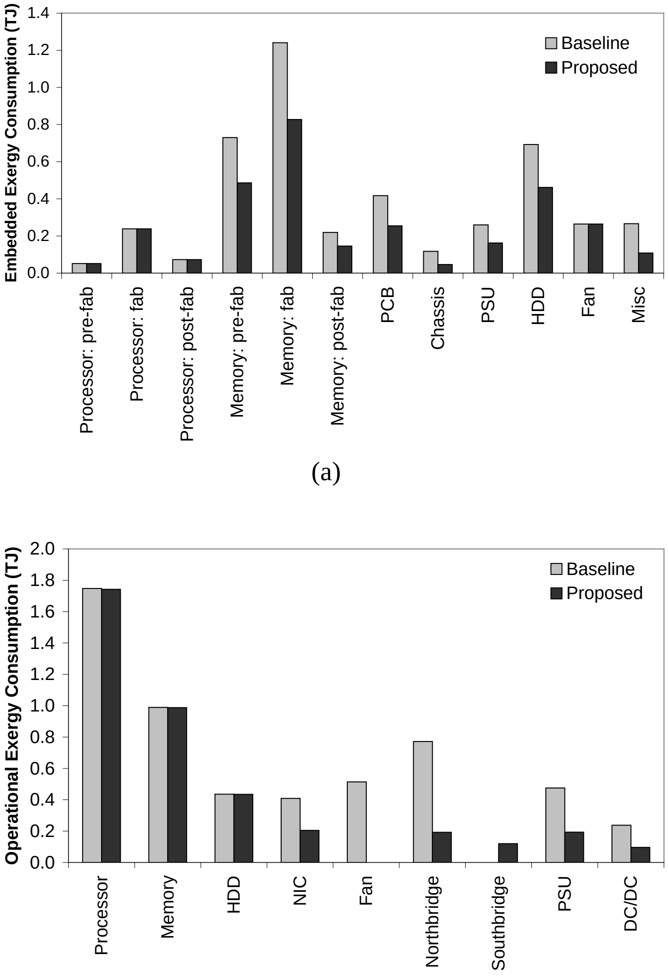

(a)

(b)

Figure 5. (a) Embedded Exergy Consumed and (b) Operational Exergy Consumed in Baseline (existing) and Proposed Design.

Table 1. Performance for proposed system along different benchmark workloads.

Workload	Description	Performance
Indexer	Nutch indexing wikipedia pages	98%
Search	Nutch search engine serving in-memory inde	99%
Image	Photo montage and manipulations	99%
MR-sort	MapReduce local sort phase	96%
H-mean	Harmonic mean of the above workloads	98%

To summarize, we find that the benefits in cooling come from various aspects of the design.

- The dematerialization of the servers allows us to eliminate the backplane. Compared to state-of-the-art systems, this enables us to reduce the pressure drop by about 4%, yielding a corresponding savings in fan power.
- We are also able to pool together the fluid delivery mechanism for multiple spines to a single bank of fans. By eliminating inefficiencies related to fluid handling across multiple systems in the infrastructure, we are able to reduce the electricity required in the cooling infrastructure by about 10%.
- The use of outside air eliminates the chiller and related auxiliary equipment (pumps, cooling tower etc.). This reduces the thermal work required by about 14%.
- The ability of zone-based thermal density optimizations to leverage differences in cooling high-density components (such as CPU) from low-density components (such as memory or storage) provides about 16% savings in cooling.

Embedded evaluation

Relative to the existing baseline data center, the proposed design is found to have an embedded exergy consumption that is approximately 31% lower. As shown in Fig. 5a, these benefits come from several different components. Reductions in memory and disk embedded exergy consumption stem from upcycling benefits of disaggregation. Disaggregation also leads to smaller board sizes and corresponding reduction in the materials and energy used to manufacture the PCB. The new package design leads to lower materials in the chassis and the miscellaneous category.

Operational evaluation

Interestingly, though we did not explicitly focus on reducing operational exergy consumption, Fig. 5b shows about 28% improvement in runtime energy consumption for the systems. In particular, the IO hub and NIC components see lower energy due to reduction in the number of components via dematerializing of systems. Elimination of system-level fans via aggregated delivery of cooling also provides some savings to the systems.

Impact on Performance

We next consider the impact of the proposed design in terms of system computational performance. There are two architectural changes that affect performance. The first change pertains to shared board-level resources. For the workloads we consider, this does not make a significant impact on performance. The second change pertains to remote memory usage and is based on the individual workload's memory working set size and access patterns. This *can* affect performance, and merits further investigation. We study this latter effect using a methodology similar to prior studies of this topic [24]. We specifically considered four web workloads – *indexer, search, image manipulation*, and *mapreduce-based sort*. Each of these workloads are representations of how the system behaves computationally with different tasks that a user might perform. For example, the *search* workload exercises the system in a manner that's similar to what is encountered when a user may be searching the web. Table 1 presents the results for the above four workloads, corresponding to a configuration with four DIMMS of local memory per compute blade supplemented by a shared memory blade. Our results 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% degradation). For most users, this degradation in performance will likely be tolerable.

Overall Evaluation

Figure 6 shows a comparison of the exergy consumption for different components of the life-cycle shown in Fig. 2(b). Overall, we find in Fig. 6(a) that the proposed design has a lifetime exergy consumption that is about 41% lower. More important than the absolute number, however, the evaluation of Fig. 6(b) points to some interesting observations. First, the overall improvements in total exergy consumption are achieved through relatively small improvements in multiple different components of the system: no single component improves by more than 16%. Similarly, it is hard to attribute a large fraction of the improvements to any single design approach. For example, the dematerialization of the sheet metal and the departure from traditional notions of enclosures allows increased sharing across a larger number of memory blades, and in turn, the disaggregation enables new optimizations around clustering of zones for optimal thermal density. Ultimately, these infrastructural savings far outweigh the benefit derived from savings in embedded exergy consumption in the sheet metal. That is, the biggest contributor to improved sustainability within enterprise computing systems seems to be derived from holistic optimizations co-designed across system architecture and physical organization. Often times, focusing on embedded exergy can enable new optimizations that address operational and cooling energy as well.

Overall, for the configuration presented in Fig. 3, we find that the net *performance per total exergy consumption* improves relative to the baseline by a factor of 1.6X.

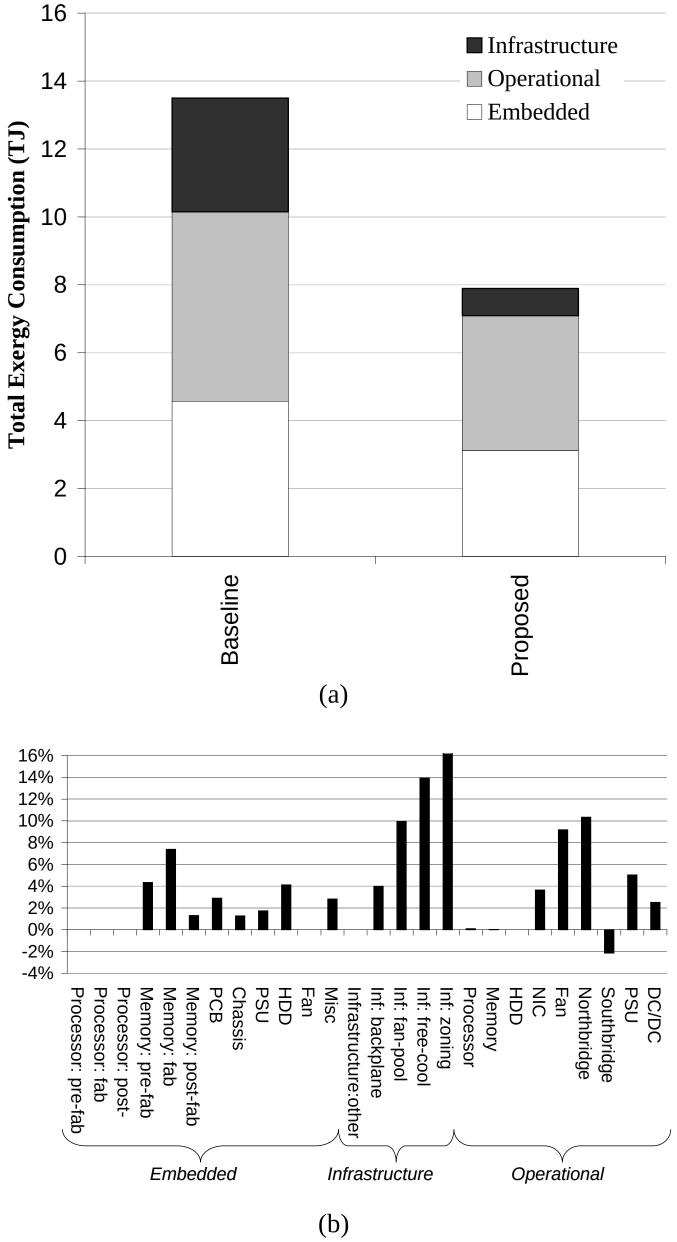

(a)

(b)

Figure 6. Evaluation of exergy consumption savings. (a) Comparison of exergy consumption between proposed and baseline systems. (b) Component-level contributions to the total savings in exergy consumption between proposed and baseline systems. The vertical axis represents fractional savings normalized to the total savings across the proposed design. For example, a savings of 16% in the above figure implies that 16% of the total savings comes from the specified component.

CONCLUSIONS

In this paper, we address the design of datacenters from a lifecycle perspective. We use the metric of lifecycle exergy consumption for system design studies to allow environmental impact evaluation at the level of architectural building blocks like processors, DIMMs, etc. Such a second-law approach allows for fusing information regarding material use and energy use into a single overall metric, and thus provides a generalized approach that can be applied to optimize the environmental footprint of any type of computing infrastructure. Using this analysis, we systematically examine bottlenecks and propose a new containerized datacenter design targeted at reducing exergy consumption across the lifecycle, particularly in terms of embedded impact factors related to material use and manufacturing as well as their impact on burdened infrastructure costs.

As part of ongoing and future work, we intend to examine other aspects of the design space. Silicon-based components constitute a dominant fraction of embedded impact, and we are looking at design alternatives that achieve similar performance at lower transistor count. In this paper, we assumed a constraint of minimal performance loss. However, designs that trade off performance for sustainability might achieve net improvements in overall performance-per-exergy and this direction is worth pursuing. A few studies (e.g., [40, 41]) have considered renewable energy in the datacenter and capping approaches for “brown” energy management. These approaches are complementary to our work and can be used in conjunction with the proposed solution. Exploring alternate power delivery and cooling infrastructure for the data center (e.g., fuel cells or wind power, hybrid liquid/air-cooling) offer another interesting set of design points for holistic lifecycle optimizations.

The key design constraints for future computing systems are evolving from a pure focus on traditional metrics like performance to also include emerging metrics like power and environmental sustainability. Design approaches and the scope of architectural design need to correspondingly change to address these emerging challenges. The results in this paper demonstrate how such a cross-domain approach across system architecture, physical packaging, and infrastructure can lead to new designs with lower impact on the environment and improved energy efficiency at equivalent performance. We believe more such work is needed.

REFERENCES

[1] *Revolutionizing Data Center Energy Efficiency: Key Analyses*. Uptime Institute and McKinsey, 2008.

[2] Paul Otellini. Building a continuum of computing. Intel Developer Forum Keynote, 2009.

[3] *EPA Report on Server and Data Center Energy Efficiency*. US Environmental Protection Agency, August 2007.

[4] “Gartner Highlights Key Predictions for IT Organizations and Users in 2008 and Beyond,” January 31, 2008, available http://gartner.com/it/page.jsp?id=593207

[5] *ISO 14040: Environmental management – Life Cycle Assessment – Principles and framework*. International Organization for Standardization, 2006.

[6] Henrikke Bauman and Anne-Marie Tillman. *The Hitch Hiker’s Guide to LCA*. Studentlitteratur AB, 2004.

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

[8] David R. Morris Jan Szargut and Frank R. Steward. *Exergy analysis of thermal, chemical and metallurgical processes*. Hemisphere, 1988.

[9] Ibrahim Dincer and Marc Rosen. *Exergy: Energy, Environment and Sustainable Development*. Elsevier, 2007.

[10] Chris R. Hannemann, Van P. Careyand 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.

[11] 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.

[12] Eric Williams. Energy intensity of computer manufacturing: Hybrid assessment combining process and economic input-output methods. *Environmental Science and Technology*, 38(22):6166–6174, 2008.

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

[14] Rolf Frischknecht et al. The ecoinvent database: Overview and methodological framework. *International Journal of Life Cycle Assessment*, 10(1):3–9, 2005.

[15] Eric Williams. The environmental impacts of semiconductor fabrication. *Thin Solid Films*, 461(1):2–6, 2004.

[16] 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.

[17] 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.

[18] Nikhil Krishnan, Sarah Boyd, Ajay Somanni, Sebastian Raoux, Daniel Clark, and David Dornfeld. A hybrid life cycle inventory of nano-scale semiconductor manufacturing. *Environmental Science and Technology*, 42(8):3069–3075, 2008.

[19] 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.

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

[21] Kevin Lim, Parthasarathy Ranganathan, Jichuan Chang, Chandrakant Patel, Trevor Mudge, and Steve Reinhardt. Understanding and designing new server architectures for

emerging warehouse-computing environments. In *International Symposium on Computer Architecture (ISCA)*, 2008.

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

[23] Shah, A., Bash, C.E., Patel, C.D., “Optimizing Data Center Cooling Infrastructures Using Exergothermovolumes”, ITherm, Las Vegas, NV June 2010 (pending publication)

[24] 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 ’09: Proceedings of the 36th annual international symposium on Computer architecture*, pages 267–278, New York, NY, USA, 2009. ACM.

[25] Ganesh Ananthanarayanan and Randy H. Katz. Greening the switch. In *HotPower 2008*, 2008.

[26] M. Gupta and S.Singh. The greening of the internet. In *SIGCOMM*, 2003.

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

[28] Yiyu Chen, Amitayu Das, Wubi Qin, Anand Sivasubramaniam, Qian Wang, and Natarajan Gautam. Managing server energy and operational costs in hosting centers. In *Proceedings of the International Conference on Measurements and Modeling of Computer Systems (SIGMETRICS)*, pages 303–314, Banff, Canada, June 2005.

[29] E. N. Elnozahy, Michael Kistler, and Ramakrishnan Rajamony. Energy-efficient server clusters. In *Proceedings of the 2nd International Workshop on Power-Aware Computer Systems (PACS)*, pages 179–196, Cambridge, MA, February 2002.

[30] Xiaobo Fan, Wolf-Dietrich Weber, and Luiz Andre 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.

[31] Taliver Heath, Bruno Diniz, Enrique V. Carrera, Wagner Meira Jr., and Ricardo Bianchini. Energy conservation in heterogeneous server clusters. In *Proceedings of the ACM Symposium on Principles and Practice of Parallel Programming (PPOPP)*, pages 186–195, Chicago, IL, June 2005.

[32] 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.

[33] Gunjan Khanna, Kiek Beaty, Gautam Kar, and Andrzej Kochut. Application performance management in virtualized server environments. In *Proceedings of the 2006 IEEE/IFIP Network Operations and Management Symposium (NOMS)*, Vancouver, Canada, April 2006.

[34] 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.

[35] Eduardo Pinheiro, Ricardo Bianchini, Enrique V. Carrera, and Taliver Heath. Dynamic cluster reconfiguration for power and performance. *Compilers and Operating Systems for Low Power*, pages 75–93, 2003.

[36] 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.

[37] 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.

[38] 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.

[39] C. Patel, et al., “Energy Flow in the Information Technology Stack: Introducing the Coefficient of Performance of the Ensemble,” Proc. ASME IMECE, November 2006.

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

[41] K. Le et al. Cost- and energy-aware load distribution across data centers. In *HotPower*, 2009.