

6

Warehouse-Scale Computers to Exploit Request-Level and Data-Level Parallelism

The datacenter is the computer.

Luiz André Barroso,
Google (2007)

A hundred years ago, companies stopped generating their own power with steam engines and dynamos and plugged into the newly built electric grid. The cheap power pumped out by electric utilities didn't just change how businesses operate. It set off a chain reaction of economic and social transformations that brought the modern world into existence. Today, a similar revolution is under way. Hooked up to the Internet's global computing grid, massive information-processing plants have begun pumping data and software code into our homes and businesses. This time, it's computing that's turning into a utility.

Nicholas Carr
*The Big Switch: Rewiring the World, from
Edison to Google (2008)*

6.1 Introduction

Anyone can build a fast CPU. The trick is to build a fast system.

Seymour Cray

Considered the father of the supercomputer

The warehouse-scale computer (WSC)¹ is the foundation of Internet services many people use every day: search, social networking, online maps, video sharing, online shopping, email services, and so on. The tremendous popularity of such Internet services necessitated the creation of WSCs that could keep up with the rapid demands of the public. Although WSCs may appear to be just large datacenters, their architecture and operation are quite different, as we shall see. Today's WSCs act as one giant machine and cost on the order of \$150M for the building, the electrical and cooling infrastructure, the servers, and the networking equipment that connects and houses 50,000 to 100,000 servers. Moreover, the rapid growth of cloud computing (see Section 6.5) makes WSCs available to anyone with a credit card.

Computer architecture extends naturally to designing WSCs. For example, Luiz Barroso of Google (quoted earlier) did his dissertation research in computer architecture. He believes an architect's skills of designing for scale, designing for dependability, and a knack for debugging hardware are very helpful in the creation and operation of WSCs.

At this extreme scale, which requires innovation in power distribution, cooling, monitoring, and operations, the WSC is the modern descendant of the supercomputer—making Seymour Cray the godfather of today's WSC architects. His extreme computers handled computations that could be done nowhere else, but were so expensive that only a few companies could afford them. This time the target is providing information technology for the world instead of high-performance computing (HPC) for scientists and engineers; hence, WSCs arguably play a more important role for society today than Cray's supercomputers did in the past.

Unquestionably, WSCs have many orders of magnitude more users than high-performance computing, and they represent a much larger share of the IT market. Whether measured by number of users or revenue, Google is at least 250 times larger than Cray Research ever was.

¹ This chapter is based on material from the book *The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale Machines*, by Luiz André Barroso and Urs Hölzle of Google [2009]; the blog Perspectives at mvdirona.com and the talks "Cloud-Computing Economies of Scale" and "Data Center Networks Are in My Way," by James Hamilton of Amazon Web Services [2009, 2010]; and the technical report *Above the Clouds: A Berkeley View of Cloud Computing*, by Michael Armbrust et al. [2009].

WSC architects share many goals and requirements with server architects:

- *Cost-performance*—Work done per dollar is critical in part because of the scale. Reducing the capital cost of a WSC by 10% could save \$15M.
- *Energy efficiency*—Power distribution costs are functionally related to power consumption; you need sufficient power distribution before you can consume power. Mechanical system costs are functionally related to power: You need to get out the heat that you put in. Hence, peak power and consumed power drive both the cost of power distribution and the cost of cooling systems. Moreover, energy efficiency is an important part of environmental stewardship. Hence, work done per joule is critical for both WSCs and servers because of the high cost of building the power and mechanical infrastructure for a warehouse of computers and for the monthly utility bills to power servers.
- *Dependability via redundancy*—The long-running nature of Internet services means that the hardware and software in a WSC must collectively provide at least 99.99% of availability; that is, it must be down less than 1 hour per year. Redundancy is the key to dependability for both WSCs and servers. While server architects often utilize more hardware offered at higher costs to reach high availability, WSC architects rely instead on multiple cost-effective servers connected by a low-cost network and redundancy managed by software. Furthermore, if the goal is to go much beyond “four nines” of availability, you need multiple WSCs to mask events that can take out whole WSCs. Multiple WSCs also reduce latency for services that are widely deployed.
- *Network I/O*—Server architects must provide a good network interface to the external world, and WSC architects must also. Networking is needed to keep data consistent between multiple WSCs as well as to interface to the public.
- *Both interactive and batch processing workloads*—While you expect highly interactive workloads for services like search and social networking with millions of users, WSCs, like servers, also run massively parallel batch programs to calculate metadata useful to such services. For example, MapReduce jobs are run to convert the pages returned from crawling the Web into search indices (see Section 6.2).

Not surprisingly, there are also characteristics *not* shared with server architecture:

- *Ample parallelism*—A concern for a server architect is whether the applications in the targeted marketplace have enough parallelism to justify the amount of parallel hardware and whether the cost is too high for sufficient communication hardware to exploit this parallelism. A WSC architect has no such concern. First, batch applications benefit from the large number of independent datasets that require independent processing, such as billions of Web pages from a Web crawl. This processing is *data-level parallelism* applied to data in storage instead of data in memory, which we saw in Chapter 4. Second, interactive Internet service applications, also known as *software as a service* (SaaS), can benefit from millions of independent users of interactive Internet

services. Reads and writes are rarely dependent in SaaS, so SaaS rarely needs to synchronize. For example, search uses a read-only index and email is normally reading- and writing-independent information. We call this type of easy parallelism *request-level parallelism*, as many independent efforts can proceed in parallel naturally with little need for communication or synchronization; for example, journal-based updating can reduce throughput demands. Given the success of SaaS and WSCs, more traditional applications such as relational databases have been weakened to rely on request-level parallelism. Even read-/write-dependent features are sometimes dropped to offer storage that can scale to the size of modern WSCs.

- *Operational costs count*—Server architects usually design their systems for peak performance within a cost budget and worry about power only to make sure they don't exceed the cooling capacity of their enclosure. They usually ignore operational costs of a server, assuming that they pale in comparison to purchase costs. WSCs have longer lifetimes—the building and electrical and cooling infrastructure are often amortized over 10 or more years—so the operational costs add up: Energy, power distribution, and cooling represent more than 30% of the costs of a WSC in 10 years.
- *Scale and the opportunities/problems associated with scale*—Often extreme computers are extremely expensive because they require custom hardware, and yet the cost of customization cannot be effectively amortized since few extreme computers are made. However, when you purchase 50,000 servers and the infrastructure that goes with it to construct a single WSC, you *do* get volume discounts. WSCs are so massive internally that you get economy of scale even if there are not many WSCs. As we shall see in Sections 6.5 and 6.10, these economies of scale led to cloud computing, as the lower per-unit costs of a WSC meant that companies could rent them at a profit below what it costs outsiders to do it themselves. The flip side of 50,000 servers is failures. Figure 6.1 shows outages and anomalies for 2400 servers. Even if a server had a mean time to failure (MTTF) of an amazing 25 years (200,000 hours), the WSC architect would need to design for 5 server failures a day. Figure 6.1 lists the annualized disk failure rate as 2% to 10%. If there were 4 disks per server and their annual failure rate was 4%, with 50,000 servers the WSC architect should expect to see one disk fail per *hour*.

Example Calculate the availability of a service running on the 2400 servers in Figure 6.1. Unlike a service in a real WSC, in this example the service cannot tolerate hardware or software failures. Assume that the time to reboot software is 5 minutes and the time to repair hardware is 1 hour.

Answer We can estimate service availability by calculating the time of outages due to failures of each component. We'll conservatively take the lowest number in each category in Figure 6.1 and split the 1000 outages evenly between four components. We ignore slow disks—the fifth component of the 1000 outages—since

Approx. number events in 1st year	Cause	Consequence
1 or 2	Power utility failures	Lose power to whole WSC; doesn't bring down WSC if UPS and generators work (generators work about 99% of time).
4	Cluster upgrades	Planned outage to upgrade infrastructure, many times for evolving networking needs such as recabling, to switch firmware upgrades, and so on. There are about 9 planned cluster outages for every unplanned outage.
1000s	Hard-drive failures	2% to 10% annual disk failure rate [Pineiro 2007]
	Slow disks	Still operate, but run 10x to 20x more slowly
	Bad memories	One uncorrectable DRAM error per year [Schroeder et al. 2009]
	Misconfigured machines	Configuration led to ~30% of service disruptions [Barroso and Hölzle 2009]
	Flaky machines	1% of servers reboot more than once a week [Barroso and Hölzle 2009]
5000	Individual server crashes	Machine reboot, usually takes about 5 minutes

Figure 6.1 List of outages and anomalies with the approximate frequencies of occurrences in the first year of a new cluster of 2400 servers. We label what Google calls a cluster an *array*; see Figure 6.5. (Based on Barroso [2010].)

they hurt performance but not availability, and power utility failures, since the uninterruptible power supply (UPS) system hides 99% of them.

$$\begin{aligned} \text{Hours Outage}_{\text{service}} &= (4 + 250 + 250 + 250) \times 1 \text{ hour} + (250 + 5000) \times 5 \text{ minutes} \\ &= 754 + 438 = 1192 \text{ hours} \end{aligned}$$

Since there are 365×24 or 8760 hours in a year, availability is:

$$\text{Availability}_{\text{system}} = \frac{(8760 - 1192)}{8760} = \frac{7568}{8760} = 86\%$$

That is, without software redundancy to mask the many outages, a service on those 2400 servers would be down on average one day a week, or *zero nines* of availability!

As Section 6.10 explains, the forerunners of WSCs are *computer clusters*. Clusters are collections of independent computers that are connected together using standard local area networks (LANs) and off-the-shelf switches. For workloads that did not require intensive communication, clusters offered much more cost-effective computing than shared memory multiprocessors. (Shared memory multiprocessors were the forerunners of the multicore computers discussed in Chapter 5.) Clusters became popular in the late 1990s for scientific computing and then later for Internet services. One view of WSCs is that they are just the logical evolution from clusters of hundreds of servers to tens of thousands of servers today.

A natural question is whether WSCs are similar to modern clusters for high-performance computing. Although some have similar scale and cost—there are HPC designs with a million processors that cost hundreds of millions of dollars—they generally have much faster processors and much faster networks between the nodes than are found in WSCs because the HPC applications are more interdependent and communicate more frequently (see Section 6.3). HPC designs also tend to use custom hardware—especially in the network—so they often don't get the cost benefits from using commodity chips. For example, the IBM Power 7 microprocessor alone can cost more and use more power than an entire server node in a Google WSC. The programming environment also emphasizes thread-level parallelism or data-level parallelism (see Chapters 4 and 5), typically emphasizing latency to complete a single task as opposed to bandwidth to complete many independent tasks via request-level parallelism. The HPC clusters also tend to have long-running jobs that keep the servers fully utilized, even for weeks at a time, while the utilization of servers in WSCs ranges between 10% and 50% (see Figure 6.3 on page 440) and varies every day.

How do WSCs compare to conventional datacenters? The operators of a conventional datacenter generally collect machines and third-party software from many parts of an organization and run them centrally for others. Their main focus tends to be consolidation of the many services onto fewer machines, which are isolated from each other to protect sensitive information. Hence, virtual machines are increasingly important in datacenters. Unlike WSCs, conventional datacenters tend to have a great deal of hardware and software heterogeneity to serve their varied customers inside an organization. WSC programmers customize third-party software or build their own, and WSCs have much more homogeneous hardware; the WSC goal is to make the hardware/software in the warehouse act like a single computer that typically runs a variety of applications. Often the largest cost in a conventional datacenter is the people to maintain it, whereas, as we shall see in Section 6.4, in a well-designed WSC the server hardware is the greatest cost, and people costs shift from the topmost to nearly irrelevant. Conventional datacenters also don't have the scale of a WSC, so they don't get the economic benefits of scale mentioned above. Hence, while you might consider a WSC as an extreme datacenter, in that computers are housed separately in a space with special electrical and cooling infrastructure, typical datacenters share little with the challenges and opportunities of a WSC, either architecturally or operationally.

Since few architects understand the software that runs in a WSC, we start with the workload and programming model of a WSC.

6.2 Programming Models and Workloads for Warehouse-Scale Computers

If a problem has no solution, it may not be a problem, but a fact—not to be solved, but to be coped with over time.

Shimon Peres

In addition to the public-facing Internet services such as search, video sharing, and social networking that make them famous, WSCs also run batch applications, such as converting videos into new formats or creating search indexes from Web crawls.

Today, the most popular framework for batch processing in a WSC is MapReduce [Dean and Ghemawat 2008] and its open-source twin Hadoop. Figure 6.2 shows the increasing popularity of MapReduce at Google over time. (Facebook runs Hadoop on 2000 batch-processing servers of the 60,000 servers it is estimated to have in 2011.) Inspired by the Lisp functions of the same name, Map first applies a programmer-supplied function to each logical input record. Map runs on thousands of computers to produce an intermediate result of key-value pairs. Reduce collects the output of those distributed tasks and collapses them using another programmer-defined function. With appropriate software support, both are highly parallel yet easy to understand and to use. Within 30 minutes, a novice programmer can run a MapReduce task on thousands of computers.

For example, one MapReduce program calculates the number of occurrences of every English word in a large collection of documents. Below is a simplified version of that program, which shows just the inner loop and assumes just one occurrence of all English words found in a document [Dean and Ghemawat 2008]:

```
map(String key, String value):
    // key: document name
    // value: document contents
    for each word w in value:
        EmitIntermediate(w, "1"); // Produce list of all words

reduce(String key, Iterator values):
    // key: a word
    // values: a list of counts
    int result = 0;
    for each v in values:
        result += ParseInt(v); // get integer from key-value pair
    Emit(AsString(result));
```

	Aug-04	Mar-06	Sep-07	Sep-09
Number of MapReduce jobs	29,000	171,000	2,217,000	3,467,000
Average completion time (seconds)	634	874	395	475
Server years used	217	2002	11,081	25,562
Input data read (terabytes)	3288	52,254	403,152	544,130
Intermediate data (terabytes)	758	6743	34,774	90,120
Output data written (terabytes)	193	2970	14,018	57,520
Average number of servers per job	157	268	394	488

Figure 6.2 Annual MapReduce usage at Google over time. Over five years the number of MapReduce jobs increased by a factor of 100 and the average number of servers per job increased by a factor of 3. In the last two years the increases were factors of 1.6 and 1.2, respectively [Dean 2009]. Figure 6.16 on page 459 estimates that running the 2009 workload on Amazon's cloud computing service EC2 would cost \$133M.

The function `EmitIntermediate` used in the Map function emits each word in the document and the value one. Then the Reduce function sums all the values per word for each document using `ParseInt()` to get the number of occurrences per word in all documents. The MapReduce runtime environment schedules map tasks and reduce task to the nodes of a WSC. (The complete version of the program is found in Dean and Ghemawat [2004].)

MapReduce can be thought of as a generalization of the single-instruction, multiple-data (SIMD) operation (Chapter 4)—except that you pass a function to be applied to the data—that is followed by a function that is used in a reduction of the output from the Map task. Because reductions are commonplace even in SIMD programs, SIMD hardware often offers special operations for them. For example, Intel’s recent AVX SIMD instructions include “horizontal” instructions that add pairs of operands that are adjacent in registers.

To accommodate variability in performance from thousands of computers, the MapReduce scheduler assigns new tasks based on how quickly nodes complete prior tasks. Obviously, a single slow task can hold up completion of a large MapReduce job. In a WSC, the solution to slow tasks is to provide software mechanisms to cope with such variability that is inherent at this scale. This approach is in sharp contrast to the solution for a server in a conventional data-center, where traditionally slow tasks mean hardware is broken and needs to be replaced or that server software needs tuning and rewriting. Performance heterogeneity is the norm for 50,000 servers in a WSC. For example, toward the end of a MapReduce program, the system will start backup executions on other nodes of the tasks that haven’t completed yet and take the result from whichever finishes first. In return for increasing resource usage a few percent, Dean and Ghemawat [2008] found that some large tasks complete 30% faster.

Another example of how WSCs differ is the use of data replication to overcome failures. Given the amount of equipment in a WSC, it’s not surprising that failures are commonplace, as the prior example attests. To deliver on 99.99% availability, systems software must cope with this reality in a WSC. To reduce operational costs, all WSCs use automated monitoring software so that one operator can be responsible for more than 1000 servers.

Programming frameworks such as MapReduce for batch processing and externally facing SaaS such as search rely upon internal software services for their success. For example, MapReduce relies on the Google File System (GFS) (Ghemawat, Gobioff, and Leung [2003]) to supply files to any computer, so that MapReduce tasks can be scheduled anywhere.

In addition to GFS, examples of such scalable storage systems include Amazon’s key value storage system Dynamo [DeCandia et al. 2007] and the Google record storage system Bigtable [Chang 2006]. Note that such systems often build upon each other. For example, Bigtable stores its logs and data on GFS, much as a relational database may use the file system provided by the kernel operating system.

These internal services often make different decisions than similar software running on single servers. As an example, rather than assuming storage is reliable, such as by using RAID storage servers, these systems often make complete

replicas of the data. Replicas can help with read performance as well as with availability; with proper placement, replicas can overcome many other system failures, like those in Figure 6.1. Some systems use erasure encoding rather than full replicas, but the constant is cross-server redundancy rather than within-a-server or within-a-storage array redundancy. Hence, failure of the entire server or storage device doesn't negatively affect availability of the data.

Another example of the different approach is that WSC storage software often uses relaxed consistency rather than following all the ACID (atomicity, consistency, isolation, and durability) requirements of conventional database systems. The insight is that it's important for multiple replicas of data to agree *eventually*, but for most applications they need not be in agreement at all times. For example, eventual consistency is fine for video sharing. Eventual consistency makes storage systems much easier to scale, which is an absolute requirement for WSCs.

The workload demands of these public interactive services all vary considerably; even a popular global service such as Google search varies by a factor of two depending on the time of day. When you factor in weekends, holidays, and popular times of year for some applications—such as photograph sharing services after Halloween or online shopping before Christmas—you can see considerably greater variation in server utilization for Internet services. Figure 6.3 shows average utilization of 5000 Google servers over a 6-month period. Note that less than 0.5% of servers averaged 100% utilization, and most servers operated between 10% and 50% utilization. Stated alternatively, just 10% of all servers were utilized more than 50%. Hence, it's much more important for servers in a WSC to perform well while doing little than to just to perform efficiently at their peak, as they rarely operate at their peak.

In summary, WSC hardware and software must cope with variability in load based on user demand and in performance and dependability due to the vagaries of hardware at this scale.

Example As a result of measurements like those in Figure 6.3, the SPECpower benchmark measures power and performance from 0% load to 100% in 10% increments (see Chapter 1). The overall single metric that summarizes this benchmark is the sum of all the performance measures (server-side Java operations per second) divided by the sum of all power measurements in watts. Thus, each level is equally likely. How would the numbers summary metric change if the levels were weighted by the utilization frequencies in Figure 6.3?

Answer Figure 6.4 shows the original weightings and the new weighting that match Figure 6.3. These weightings reduce the performance summary by 30% from 3210 `ssj_ops/watt` to 2454.

Given the scale, software must handle failures, which means there is little reason to buy “gold-plated” hardware that reduces the frequency of failures. The primary impact would be to increase cost. Barroso and Hölzle [2009] found a factor of 20 difference in price-performance between a high-end




Figure 6.3 Average CPU utilization of more than 5000 servers during a 6-month period at Google. Servers are rarely completely idle or fully utilized, instead operating most of the time at between 10% and 50% of their maximum utilization. (From Figure 1 in Barroso and Hölzle [2007].) The column the third from the right in Figure 6.4 calculates percentages plus or minus 5% to come up with the weightings; thus, 1.2% for the 90% row means that 1.2% of servers were between 85% and 95% utilized.

Load	Performance	Watts	SPEC weightings	Weighted performance	Weighted watts	Figure 6.3 weightings	Weighted performance	Weighted watts
100%	2,889,020	662	9.09%	262,638	60	0.80%	22,206	5
90%	2,611,130	617	9.09%	237,375	56	1.20%	31,756	.8
80%	2,319,900	576	9.09%	210,900	52	1.50%	35,889	9
70%	2,031,260	533	9.09%	184,660	48	2.10%	42,491	11
60%	1,740,980	490	9.09%	158,271	45	5.10%	88,082	25
50%	1,448,810	451	9.09%	131,710	41	11.50%	166,335	52
40%	1,159,760	416	9.09%	105,433	38	19.10%	221,165	79
30%	869,077	382	9.09%	79,007	35	24.60%	213,929	94
20%	581,126	351	9.09%	52,830	32	15.30%	88,769	54
10%	290,762	308	9.09%	26,433	28	8.00%	23,198	25
0%	0	181	9.09%	0	16	10.90%	0	20
Total	15,941,825	4967		1,449,257	452		933,820	380
				ssj_ops/Watt	3210		ssj_ops/Watt	2454

Figure 6.4 SPECpower result from Figure 6.17 using the weightings from Figure 6.3 instead of even weightings.

HP shared-memory multiprocessor and a commodity HP server when running the TPC-C database benchmark. Unsurprisingly, Google buys low-end commodity servers.

Such WSC services also tend to develop their own software rather than buy third-party commercial software, in part to cope with the huge scale and in part to save money. For example, even on the best price-performance platform for TPC-C in 2011, including the cost of the Oracle database and Windows operating system doubles the cost of the Dell Poweredge 710 server. In contrast, Google runs Bigtable and the Linux operating system on its servers, for which it pays no licensing fees.

Given this review of the applications and systems software of a WSC, we are ready to look at the computer architecture of a WSC.

6.3

Computer Architecture of Warehouse-Scale Computers

Networks are the connective tissue that binds 50,000 servers together. Analogous to the memory hierarchy of Chapter 2, WSCs use a hierarchy of networks. Figure 6.5 shows one example. Ideally, the combined network would provide nearly the performance of a custom high-end switch for 50,000 servers at nearly the cost per port of a commodity switch designed for 50 servers. As we shall see in Section 6.6, the current solutions are far from that ideal, and networks for WSCs are an area of active exploration.

The 19-inch (48.26-cm) rack is still the standard framework to hold servers, despite this standard going back to railroad hardware from the 1930s. Servers are measured in the number of rack units (U) that they occupy in a rack. One U is 1.75 inches (4.45 cm) high, and that is the minimum space a server can occupy.

A 7-foot (213.36-cm) rack offers 48 U, so it's not a coincidence that the most popular switch for a rack is a 48-port Ethernet switch. This product has become a commodity that costs as little as \$30 per port for a 1 Gbit/sec Ethernet link in 2011 [Barroso and Hölzle 2009]. Note that the bandwidth within the rack is the same for each server, so it does not matter where the software places the sender and the receiver as long as they are within the same rack. This flexibility is ideal from a software perspective.

These switches typically offer two to eight uplinks, which leave the rack to go to the next higher switch in the network hierarchy. Thus, the bandwidth leaving the rack is 6 to 24 times smaller—48/8 to 48/2—than the bandwidth within the rack. This ratio is called *oversubscription*. Alas, large oversubscription means programmers must be aware of the performance consequences when placing senders and receivers in different racks. This increased software-scheduling burden is another argument for network switches designed specifically for the datacenter.

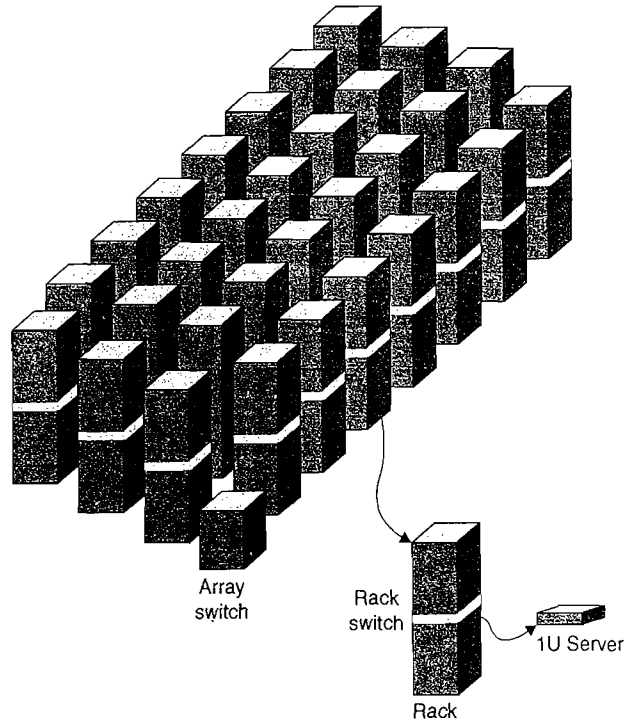


Figure 6.5 Hierarchy of switches in a WSC. (Based on Figure 1.2 of Barroso and Hölzle [2009].)

Storage

A natural design is to fill a rack with servers, minus whatever space you need for the commodity Ethernet rack switch. This design leaves open the question of where the storage is placed. From a hardware construction perspective, the simplest solution would be to include disks inside the server, and rely on Ethernet connectivity for access to information on the disks of remote servers. The alternative would be to use network attached storage (NAS), perhaps over a storage network like Infiniband. The NAS solution is generally more expensive per terabyte of storage, but it provides many features, including RAID techniques to improve dependability of the storage.

As you might expect from the philosophy expressed in the prior section, WSCs generally rely on local disks and provide storage software that handles connectivity and dependability. For example, GFS uses local disks and maintains at least three replicas to overcome dependability problems. This redundancy covers not just local disk failures, but also power failures to racks and to whole clusters. The eventual consistency flexibility of GFS lowers the cost of keeping replicas consistent, which also reduces the network bandwidth requirements of the storage

system. Local access patterns also mean high bandwidth to local storage, as we'll see shortly.

Beware that there is confusion about the term *cluster* when talking about the architecture of a WSC. Using the definition in Section 6.1, a WSC is just an extremely large cluster. In contrast, Barroso and Hölzle [2009] used the term cluster to mean the next-sized grouping of computers, in this case about 30 racks. In this chapter, to avoid confusion we will use the term *array* to mean a collection of racks, preserving the original meaning of the word cluster to mean anything from a collection of networked computers within a rack to an entire warehouse full of networked computers.

Array Switch

The switch that connects an array of racks is considerably more expensive than the 48-port commodity Ethernet switch. This cost is due in part because of the higher connectivity and in part because the bandwidth through the switch must be much higher to reduce the oversubscription problem. Barroso and Hölzle [2009] reported that a switch that has 10 times the *bisection bandwidth*—basically, the worst-case internal bandwidth—of a rack switch costs about 100 times as much. One reason is that the cost of switch bandwidth for n ports can grow as n^2 .

Another reason for the high costs is that these products offer high profit margins for the companies that produce them. They justify such prices in part by providing features such as packet inspection that are expensive because they must operate at very high rates. For example, network switches are major users of content-addressable memory chips and of field-programmable gate arrays (FPGAs), which help provide these features, but the chips themselves are expensive. While such features may be valuable for Internet settings, they are generally unused inside the datacenter.

WSC Memory Hierarchy

Figure 6.6 shows the latency, bandwidth, and capacity of memory hierarchy inside a WSC, and Figure 6.7 shows the same data visually. These figures are based on the following assumptions [Barroso and Hölzle 2009]:

	Local	Rack	Array
DRAM latency (microseconds)	0.1	100	300
Disk latency (microseconds)	10,000	11,000	12,000
DRAM bandwidth (MB/sec)	20,000	100	10
Disk bandwidth (MB/sec)	200	100	10
DRAM capacity (GB)	16	1040	31,200
Disk capacity (GB)	2000	160,000	4,800,000

Figure 6.6 Latency, bandwidth, and capacity of the memory hierarchy of a WSC [Barroso and Hölzle 2009]. Figure 6.7 plots this same information.




Figure 6.7 Graph of latency, bandwidth, and capacity of the memory hierarchy of a WSC for data in Figure 6.6 [Barroso and Hölzle 2009].

- Each server contains 16 GBytes of memory with a 100-nanosecond access time and transfers at 20 GBytes/sec and 2 terabytes of disk that offers a 10-millisecond access time and transfers at 200 MBytes/sec. There are two sockets per board, and they share one 1 Gbit/sec Ethernet port.
- Every pair of racks includes one rack switch and holds 80 2U servers (see Section 6.7). Networking software plus switch overhead increases the latency to DRAM to 100 microseconds and the disk access latency to 11 milliseconds. Thus, the total storage capacity of a rack is roughly 1 terabyte of DRAM and 160 terabytes of disk storage. The 1 Gbit/sec Ethernet limits the remote bandwidth to DRAM or disk within the rack to 100 MBytes/sec.
- The array switch can handle 30 racks, so storage capacity of an array goes up by a factor of 30: 30 terabytes of DRAM and 4.8 petabytes of disk. The array switch hardware and software increases latency to DRAM within an array to 500 microseconds and disk latency to 12 milliseconds. The bandwidth of the array switch limits the remote bandwidth to either array DRAM or array disk to 10 MBytes/sec.

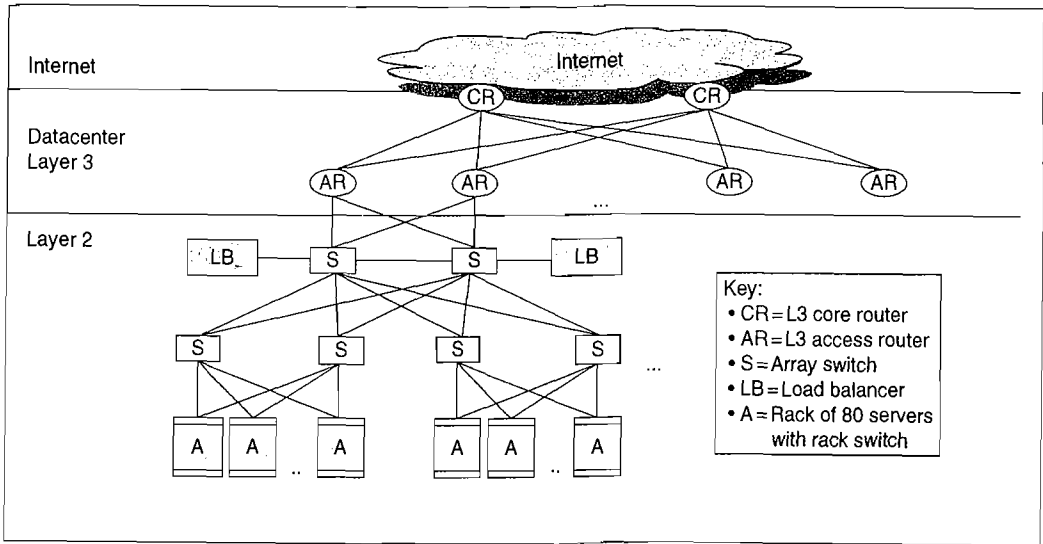


Figure 6.8 The Layer 3 network used to link arrays together and to the Internet [Greenberg et al. 2009]. Some WSCs use a separate *border router* to connect the Internet to the datacenter Layer 3 switches.

Figures 6.6 and 6.7 show that network overhead dramatically increases latency from local DRAM to rack DRAM and array DRAM, but both still have more than 10 times better latency than the local disk. The network collapses the difference in bandwidth between rack DRAM and rack disk and between array DRAM and array disk.

The WSC needs 20 arrays to reach 50,000 servers, so there is one more level of the networking hierarchy. Figure 6.8 shows the conventional Layer 3 routers to connect the arrays together and to the Internet.

Most applications fit on a single array within a WSC. Those that need more than one array use *sharding* or *partitioning*, meaning that the dataset is split into independent pieces and then distributed to different arrays. Operations on the whole dataset are sent to the servers hosting the pieces, and the results are coalesced by the client computer.

Example What is the average memory latency assuming that 90% of accesses are local to the server, 9% are outside the server but within the rack, and 1% are outside the rack but within the array?

Answer The average memory access time is

$$(90\% \times 0.1) + (9\% \times 100) + (1\% \times 300) = 0.09 + 9 + 3 = 12.09 \text{ microseconds}$$

or a factor of more than 120 slowdown versus 100% local accesses. Clearly, *locality of access within a server is vital for WSC performance.*

Example How long does it take to transfer 1000 MB between disks within the server, between servers in the rack, and between servers in different racks in the array? How much faster is it to transfer 1000 MB between DRAM in the three cases?

Answer A 1000 MB transfer between disks takes:

Within server = $1000/200 = 5$ seconds

Within rack = $1000/100 = 10$ seconds

Within array = $1000/10 = 100$ seconds

A memory-to-memory block transfer takes

Within server = $1000/20000 = 0.05$ seconds

Within rack = $1000/100 = 10$ seconds

Within array = $1000/10 = 100$ seconds

Thus, for block transfers outside a single server, it doesn't even matter whether the data are in memory or on disk since the rack switch and array switch are the bottlenecks. These performance limits affect the design of WSC software and inspire the need for higher performance switches (see Section 6.6).

Given the architecture of the IT equipment, we are now ready to see how to house, power, and cool it and to discuss the cost to build and operate the whole WSC, as compared to just the IT equipment within it.

6.4

Physical Infrastructure and Costs of Warehouse-Scale Computers

To build a WSC, you first need to build a warehouse. One of the first questions is where? Real estate agents emphasize location, but location for a WSC means proximity to Internet backbone optical fibers, low cost of electricity, and low risk from environmental disasters, such as earthquakes, floods, and hurricanes. For a company with many WSCs, another concern is finding a place geographically near a current or future population of Internet users, so as to reduce latency over the Internet. There are also many more mundane concerns, such as property tax rates.

Infrastructure costs for power distribution and cooling dwarf the construction costs of a WSC, so we concentrate on the former. Figures 6.9 and 6.10 show the power distribution and cooling infrastructure within a WSC.

Although there are many variations deployed, in North America electrical power typically goes through about five steps and four voltage changes on the way to the server, starting with the high-voltage lines at the utility tower of 115,000 volts:

1. The substation switches from 115,000 volts to medium-voltage lines of 13,200 volts, with an efficiency of 99.7%.

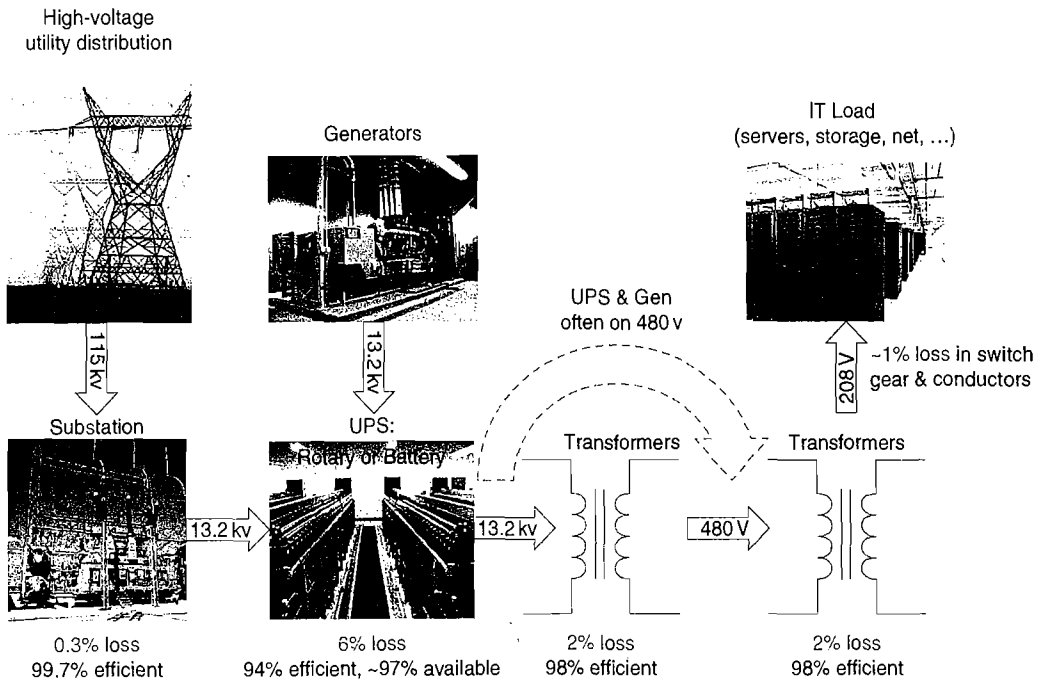


Figure 6.9 Power distribution and where losses occur. Note that the best improvement is 11%. (From Hamilton [2010].)

2. To prevent the whole WSC from going offline if power is lost, a WSC has an uninterruptible power supply (UPS), just as some servers do. In this case, it involves large diesel engines that can take over from the utility company in an emergency and batteries or flywheels to maintain power after the service is lost but before the diesel engines are ready. The generators and batteries can take up so much space that they are typically located in a separate room from the IT equipment. The UPS plays three roles: power conditioning (maintain proper voltage levels and other characteristics), holding the electrical load while the generators start and come on line, and holding the electrical load when switching back from the generators to the electrical utility. The efficiency of this very large UPS is 94%, so the facility loses 6% of the power by having a UPS. The WSC UPS can account for 7% to 12% of the cost of all the IT equipment.
3. Next in the system is a power distribution unit (PDU) that converts to low-voltage, internal, three-phase power at 480 volts. The conversion efficiency is 98%. A typical PDU handles 75 to 225 kilowatts of load, or about 10 racks.
4. There is yet another down step to two-phase power at 208 volts that servers can use, once again at 98% efficiency. (Inside the server, there are more steps to bring the voltage down to what chips can use; see Section 6.7.)

- The connectors, breakers, and electrical wiring to the server have a collective efficiency of 99%.

WSCs outside North America use different conversion values, but the overall design is similar.

Putting it all together, the efficiency of turning 115,000-volt power from the utility into 208-volt power that servers can use is 89%:

$$99.7\% \times 94\% \times 98\% \times 98\% \times 99\% = 89\%$$

This overall efficiency leaves only a little over 10% room for improvement, but as we shall see, engineers still try to make it better.

There is considerably more opportunity for improvement in the cooling infrastructure. The computer room air-conditioning (CRAC) unit cools the air in the server room using chilled water, similar to how a refrigerator removes heat by releasing it outside of the refrigerator. As a liquid absorbs heat, it evaporates. Conversely, when a liquid releases heat, it condenses. Air conditioners pump the liquid into coils under low pressure to evaporate and absorb heat, which is then sent to an external condenser where it is released. Thus, in a CRAC unit, fans push warm air past a set of coils filled with cold water and a pump moves the warmed water to the external chillers to be cooled down. The cool air for servers is typically between 64°F and 71°F (18°C and 22°C). Figure 6.10 shows the large collection of fans and water pumps that move air and water throughout the system.

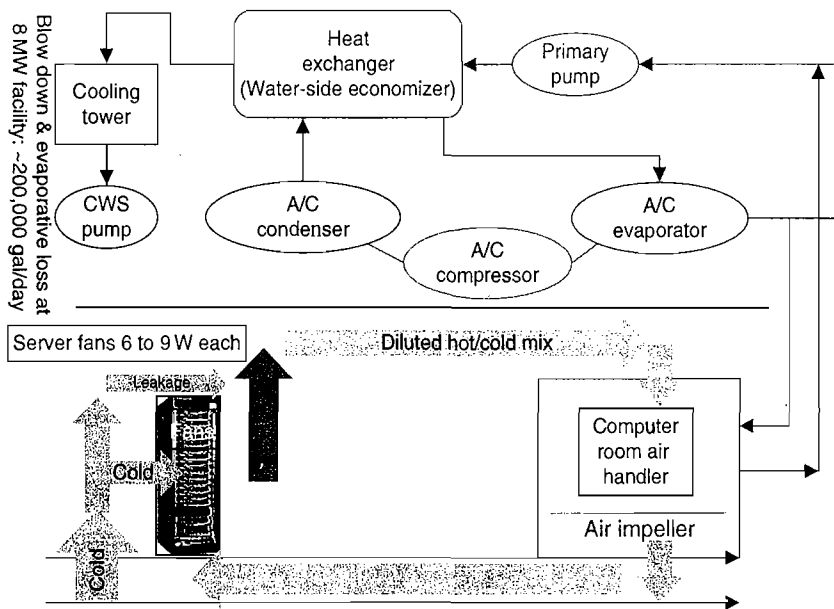


Figure 6.10 Mechanical design for cooling systems. CWS stands for circulating water system. (From Hamilton [2010].)

Clearly, one of the simplest ways to improve energy efficiency is simply to run the IT equipment at higher temperatures so that the air need not be cooled as much. Some WSCs run their equipment considerably above 71°F (22°C).

In addition to chillers, cooling towers are used in some datacenters to leverage the colder outside air to cool the water before it is sent to the chillers. The temperature that matters is called the *wet-bulb temperature*. The wet-bulb temperature is measured by blowing air on the bulb end of a thermometer that has water on it. It is the lowest temperature that can be achieved by evaporating water with air.

Warm water flows over a large surface in the tower, transferring heat to the outside air via evaporation and thereby cooling the water. This technique is called *airside economization*. An alternative is use cold water instead of cold air. Google's WSC in Belgium uses a water-to-water intercooler that takes cold water from an industrial canal to chill the warm water from inside the WSC.

Airflow is carefully planned for the IT equipment itself, with some designs even using airflow simulators. Efficient designs preserve the temperature of the cool air by reducing the chances of it mixing with hot air. For example, a WSC can have alternating aisles of hot air and cold air by orienting servers in opposite directions in alternating rows of racks so that hot exhaust blows in alternating directions.

In addition to energy losses, the cooling system also uses up a lot of water due to evaporation or to spills down sewer lines. For example, an 8 MW facility might use 70,000 to 200,000 gallons of water per day.

The relative power costs of cooling equipment to IT equipment in a typical datacenter [Barroso and Hölzle 2009] are as follows:

- Chillers account for 30% to 50% of the IT equipment power.
- CRAC accounts for 10% to 20% of the IT equipment power, due mostly to fans.

Surprisingly, it's not obvious to figure out how many servers a WSC can support after you subtract the overheads for power distribution and cooling. The so-called *nameplate power rating* from the server manufacturer is always conservative; it's the maximum power a server can draw. The first step then is to measure a single server under a variety of workloads to be deployed in the WSC. (Networking is typically about 5% of power consumption, so it can be ignored to start.)

To determine the number of servers for a WSC, the available power for IT could just be divided by the measured server power; however, this would again be too conservative according to Fan, Weber, and Barroso [2007]. They found that there is a significant gap between what thousands of servers could theoretically do in the worst case and what they will do in practice, since no real workloads will keep thousands of servers all simultaneously at their peaks. They found that they could safely oversubscribe the number of servers by as much as 40% based on the power of a single server. They recommended that WSC architects should do that to increase the average utilization of power within a WSC; however, they also suggested using extensive monitoring software along

with a safety mechanism that deschedules lower priority tasks in case the workload shifts.

Breaking down power usage inside the IT equipment itself, Barroso and Hölzle [2009] reported the following for a Google WSC deployed in 2007:

- 33% of power for processors
- 30% for DRAM
- 10% for disks
- 5% for networking
- 22% for other reasons (inside the server)

Measuring Efficiency of a WSC

A widely used, simple metric to evaluate the efficiency of a datacenter or a WSC is called *power utilization effectiveness* (or *PUE*):

$$\text{PUE} = (\text{Total facility power}) / (\text{IT equipment power})$$

Thus, PUE must be greater than or equal to 1, and the bigger the PUE the less efficient the WSC.

Greenberg et al. [2006] reported on the PUE of 19 datacenters and the portion of the overhead that went into the cooling infrastructure. Figure 6.11 shows what they found, sorted by PUE from most to least efficient. The median PUE is 1.69, with the cooling infrastructure using more than half as much power as the servers themselves—on average, 0.55 of the 1.69 is for cooling. Note that these are average PUEs, which can vary daily depending on workload and even external air temperature, as we shall see.

Since performance per dollar is the ultimate metric, we still need to measure performance. As Figure 6.7 above shows, bandwidth drops and latency increases depending on the distance to the data. In a WSC, the DRAM bandwidth within a server is 200 times larger than within a rack, which in turn is 10 times larger than within an array. Thus, there is another kind of locality to consider in the placement of data and programs within a WSC.

While designers of a WSC often focus on bandwidth, programmers developing applications on a WSC are also concerned with latency, since latency is visible to users. Users' satisfaction and productivity are tied to response time of a service. Several studies from the timesharing days report that user productivity is inversely proportional to time for an interaction, which was typically broken down into human entry time, system response time, and time for the person to think about the response before entering the next entry. The results of experiments showed that cutting system response time 30% shaved the time of an interaction by 70%. This implausible result is explained by human nature: People need less time to think when given a faster response, as they are less likely to get distracted and remain "on a roll."

Figure 6.12 shows the results of such an experiment for the Bing search engine, where delays of 50 ms to 2000 ms were inserted at the search server. As expected




Figure 6.11 Power utilization efficiency of 19 datacenters in 2006 [Greenberg et al. 2006]. The power for air conditioning (AC) and other uses (such as power distribution) is normalized to the power for the IT equipment in calculating the PUE. Thus, power for IT equipment must be 1.0 and AC varies from about 0.30 to 1.40 times the power of the IT equipment. Power for “other” varies from about 0.05 to 0.60 of the IT equipment.

from previous studies, time to next click roughly doubled the delay; that is, a 200 ms delay at the server led to a 500 ms increase in time to next click. Revenue dropped linearly with increasing delay, as did user satisfaction. A separate study on the Google search engine found that these effects lingered long after the 4-week experiment ended. Five weeks later, there were 0.1% fewer searchers per day for users who experienced 200 ms delays, and there were 0.2% fewer searches from users who experienced 400 ms delays. Given the amount of money made in search, even such small changes are disconcerting. In fact, the results were so negative that they ended the experiment prematurely [Schurman and Brutlag 2009].

Server delay (ms)	Increased time to next click (ms)	Queries/user	Any clicks/user	User satisfaction	Revenue/user
50	--	--	--	--	--
200	500	--	-0.3%	-0.4%	--
500	1200	--	-1.0%	-0.9%	-1.2%
1000	1900	-0.7%	-1.9%	-1.6%	-2.8%
2000	3100	-1.8%	-4.4%	-3.8%	-4.3%

Figure 6.12 Negative impact of delays at Bing search server on user behavior Schurman and Brutlag [2009].

Because of this extreme concern with satisfaction of all users of an Internet service, performance goals are typically specified that a high percentage of requests be below a latency threshold rather than just offer a target for the average latency. Such threshold goals are called *service level objectives (SLOs)* or *service level agreements (SLAs)*. An SLO might be that 99% of requests must be below 100 milliseconds. Thus, the designers of Amazon's Dynamo key-value storage system decided that, for services to offer good latency on top of Dynamo, their storage system had to deliver on its latency goal 99.9% of the time [DeCandia et al. 2007]. For example, one improvement of Dynamo helped the 99.9th percentile much more than the average case, which reflects their priorities.

Cost of a WSC

As mentioned in the introduction, unlike most architects, designers of WSCs worry about operational costs as well as the cost to build the WSC. Accounting labels the former costs as *operational expenditures (OPEX)* and the latter costs as *capital expenditures (CAPEX)*.

To put the cost of energy into perspective, Hamilton [2010] did a case study to estimate the costs of a WSC. He determined that the CAPEX of this 8 MW facility was \$88M, and that the roughly 46,000 servers and corresponding networking equipment added another \$79M to the CAPEX for the WSC. Figure 6.13 shows the rest of the assumptions for the case study.

We can now price the total cost of energy, since U.S. accounting rules allow us to convert CAPEX into OPEX. We can just amortize CAPEX as a fixed amount each month for the effective life of the equipment. Figure 6.14 breaks down the monthly OPEX for this case study. Note that the amortization rates differ significantly, from 10 years for the facility to 4 years for the networking equipment and 3 years for the servers. Hence, the WSC facility lasts a decade, but you need to replace the servers every 3 years and the networking equipment every 4 years. By amortizing the CAPEX, Hamilton came up with a monthly OPEX, including accounting for the cost of borrowing money (5% annually) to pay for the WSC. At \$3.8M, the monthly OPEX is about 2% of the CAPEX.

This figure allows us to calculate a handy guideline to keep in mind when making decisions about which components to use when being concerned about energy. The fully burdened cost of a watt per year in a WSC, including the cost of amortizing the power and cooling infrastructure, is

$$\frac{\text{Monthly cost of infrastructure} + \text{monthly cost of power}}{\text{Facility size in watts}} \times 12 = \frac{\$765\text{K} + \$475\text{K}}{8\text{M}} \times 12 = \$1.86$$

The cost is roughly \$2 per watt-year. Thus, to reduce costs by saving energy you shouldn't spend more than \$2 per watt-year (see Section 6.8).

Note that more than a third of OPEX is related to power, with that category trending up while server costs are trending down over time. The networking

Size of facility (critical load watts)	8,000,000
Average power usage (%)	80%
Power usage effectiveness	1.45
Cost of power (\$/kwh)	\$0.07
% Power and cooling infrastructure (% of total facility cost)	82%
CAPEX for facility (not including IT equipment)	\$88,000,000
Number of servers	45,978
Cost/server	\$1450
CAPEX for servers	\$66,700,000
Number of rack switches	1150
Cost/rack switch	\$4800
Number of array switches	22
Cost/array switch	\$300,000
Number of layer 3 switches	2
Cost/layer 3 switch	\$500,000
Number of border routers	2
Cost/border router	\$144,800
CAPEX for networking gear	\$12,810,000
Total CAPEX for WSC	\$167,510,000
Server amortization time	3 years
Networking amortization time	4 years
Facilities amortization time	10 years
Annual cost of money	5%

Figure 6.13 Case study for a WSC, based on Hamilton [2010], rounded to nearest \$5000. Internet bandwidth costs vary by application, so they are not included here. The remaining 18% of the CAPEX for the facility includes buying the property and the cost of construction of the building. We added people costs for security and facilities management in Figure 6.14, which were not part of the case study. Note that Hamilton's estimates were done before he joined Amazon, and they are not based on the WSC of a particular company.

equipment is significant at 8% of total OPEX and 19% of the server CAPEX, and networking equipment is not trending down as quickly as servers are. This difference is especially true for the switches in the networking hierarchy above the rack, which represent most of the networking costs (see Section 6.6). People costs for security and facilities management are just 2% of OPEX. Dividing the OPEX in Figure 6.14 by the number of servers and hours per month, the cost is about \$0.11 per server per hour.

Expense (% total)	Category	Monthly cost	Percent monthly cost
Amortized CAPEX (85%)	Servers	\$2,000,000	53%
	Networking equipment	\$290,000	8%
	Power and cooling infrastructure	\$765,000	20%
	Other infrastructure	\$170,000	4%
OPEX (15%)	Monthly power use	\$475,000	13%
	Monthly people salaries and benefits	\$85,000	2%
	Total OPEX	\$3,800,000	100%

Figure 6.14 Monthly OPEX for Figure 6.13, rounded to the nearest \$5000. Note that the 3-year amortization for servers means you need to purchase new servers every 3 years, whereas the facility is amortized for 10 years. Hence, the amortized capital costs for servers are about 3 times more than for the facility. People costs include 3 security guard positions continuously for 24 hours a day, 365 days a year, at \$20 per hour per person, and 1 facilities person for 24 hours a day, 365 days a year, at \$30 per hour. Benefits are 30% of salaries. This calculation doesn't include the cost of network bandwidth to the Internet, as it varies by application, nor vendor maintenance fees, as that varies by equipment and by negotiations.

Example The cost of electricity varies by region in the United States from \$0.03 to \$0.15 per kilowatt-hour. What is the impact on hourly server costs of these two extreme rates?

Answer We multiply the critical load of 8 MW by the PUE and by the average power usage from Figure 6.13 to calculate the average power usage:

$$8 \times 1.45 \times 80\% = 9.28 \text{ Megawatts}$$

The monthly cost for power then goes from \$475,000 in Figure 6.14 to \$205,000 at \$0.03 per kilowatt-hour and to \$1,015,000 at \$0.15 per kilowatt-hour. These changes in electricity cost change the hourly server costs from \$0.11 to \$0.10 and \$0.13, respectively.

Example What would happen to monthly costs if the amortization times were all made to be the same—say, 5 years? How does that change the hourly cost per server?

Answer The spreadsheet is available online at <http://mvdirona.com/jrh/TalksAndPapers/PerspectivesDataCenterCostAndPower.xls>. Changing the amortization time to 5 years changes the first four rows of Figure 6.14 to

Servers	\$1,260,000	37%
Networking equipment	\$242,000	7%
Power and cooling infrastructure	\$1,115,000	33%
Other infrastructure	\$245,000	7%

and the total monthly OPEX is \$3,422,000. If we replaced everything every 5 years, the cost would be \$0.103 per server hour, with more of the amortized costs now being for the facility rather than the servers, as in Figure 6.14.

The rate of \$0.11 per server per hour can be much less than the cost for many companies that own and operate their own (smaller) conventional datacenters. The cost advantage of WSCs led large Internet companies to offer computing as a utility where, like electricity, you pay only for what you use. Today, utility computing is better known as cloud computing.

6.5 Cloud Computing: The Return of Utility Computing

If computers of the kind I have advocated become the computers of the future, then computing may someday be organized as a public utility just as the telephone system is a public utility. . . . The computer utility could become the basis of a new and important industry.

John McCarthy

MIT centennial celebration (1961)

Driven by the demand of an increasing number of users, Internet companies such as Amazon, Google, and Microsoft built increasingly larger warehouse-scale computers from commodity components. This demand led to innovations in systems software to support operating at this scale, including Bigtable, Dynamo, GFS, and MapReduce. It also demanded improvement in operational techniques to deliver a service available at least 99.99% of the time despite component failures and security attacks. Examples of these techniques include failover, firewalls, virtual machines, and protection against distributed denial-of-service attacks. With the software and expertise providing the ability to scale and increasing customer demand that justified the investment, WSCs with 50,000 to 100,000 servers have become commonplace in 2011.

With increasing scale came increasing economies of scale. Based on a study in 2006 that compared a WSC with a datacenter with only 1000 servers, Hamilton [2010] reported the following advantages:

- *5.7 times reduction in storage costs*—It cost the WSC \$4.6 per GByte per year for disk storage versus \$26 per GByte for the datacenter.
- *7.1 times reduction in administrative costs*—The ratio of servers per administrator was over 1000 for the WSC versus just 140 for the datacenter.
- *7.3 times reduction in networking costs*—Internet bandwidth cost the WSC \$13 per Mbit/sec/month versus \$95 for the datacenter. Unsurprisingly, you can negotiate a much better price per Mbit/sec if you order 1000 Mbit/sec than if you order 10 Mbit/sec.

Another economy of scale comes during purchasing. The high level of purchasing leads to volume discount prices on the servers and networking gear. It also allows optimization of the supply chain. Dell, IBM, and SGI will deliver on new orders in a week to a WSC instead of 4 to 6 months. Short delivery time makes it much easier to grow the utility to match the demand.

Economies of scale also apply to operational costs. From the prior section, we saw that many datacenters operate with a PUE of 2.0. Large firms can justify hiring mechanical and power engineers to develop WSCs with lower PUEs, in the range of 1.2 (see Section 6.7).

Internet services need to be distributed to multiple WSCs for both dependability and to reduce latency, especially for international markets. All large firms use multiple WSCs for that reason. It's much more expensive for individual firms to create multiple, small datacenters around the world than a single datacenter in the corporate headquarters.

Finally, for the reasons presented in Section 6.1, servers in datacenters tend to be utilized only 10% to 20% of the time. By making WSCs available to the public, uncorrelated peaks between different customers can raise average utilization above 50%.

Thus, economies of scale for a WSC offer factors of 5 to 7 for several components of a WSC plus a few factors of 1.5 to 2 for the entire WSC.

While there are many cloud computing providers, we feature Amazon Web Services (AWS) in part because of its popularity and in part because of the low level and hence more flexible abstraction of their service. Google App Engine and Microsoft Azure raise the level of abstraction to managed runtimes and to offer automatic scaling services, which are a better match to some customers, but not as good a match as AWS to the material in this book.

Amazon Web Services

Utility computing goes back to commercial timesharing systems and even batch processing systems of the 1960s and 1970s, where companies only paid for a terminal and a phone line and then were billed based on how much computing they used. Many efforts since the end of timesharing then have tried to offer such pay as you go services, but they were often met with failure.

When Amazon started offering utility computing via the Amazon Simple Storage Service (Amazon S3) and then Amazon Elastic Computer Cloud (Amazon EC2) in 2006, it made some novel technical and business decisions:

- *Virtual Machines.* Building the WSC using x86-commodity computers running the Linux operating system and the Xen virtual machine solved several problems. First, it allowed Amazon to protect users from each other. Second, it simplified software distribution within a WSC, in that customers only need install an image and then AWS will automatically distribute it to all the instances being used. Third, the ability to kill a virtual machine reliably

makes it easy for Amazon and customers to control resource usage. Fourth, since Virtual Machines can limit the rate at which they use the physical processors, disks, and the network as well as the amount of main memory, that gave AWS multiple price points: the lowest price option by packing multiple virtual cores on a single server, the highest price option of exclusive access to all the machine resources, as well as several intermediary points. Fifth, Virtual Machines hide the identity of older hardware, allowing AWS to continue to sell time on older machines that might otherwise be unattractive to customers if they knew their age. Finally, Virtual Machines allow AWS to introduce new and faster hardware by either packing even more virtual cores per server or simply by offering instances that have higher performance per virtual core; virtualization means that offered performance need not be an integer multiple of the performance of the hardware.

- *Very low cost.* When AWS announced a rate of \$0.10 per hour per instance in 2006, it was a startlingly low amount. An instance is one Virtual Machine, and at \$0.10 per hour AWS allocated two instances per core on a multicore server. Hence, one EC2 computer unit is equivalent to a 1.0 to 1.2 GHz AMD Opteron or Intel Xeon of that era.
- *(Initial) reliance on open source software.* The availability of good-quality software that had no licensing problems or costs associated with running on hundreds or thousands of servers made utility computing much more economical for both Amazon and its customers. More recently, AWS started offering instances including commercial third-party software at higher prices.
- *No (initial) guarantee of service.* Amazon originally promised only best effort. The low cost was so attractive that many could live without a service guarantee. Today, AWS provides availability SLAs of up to 99.95% on services such as Amazon EC2 and Amazon S3. Additionally, Amazon S3 was designed for 99.99999999% durability by saving multiple replicas of each object across multiple locations. That is, the chances of permanently losing an object are one in 100 billion. AWS also provides a Service Health Dashboard that shows the current operational status of each of the AWS services in real time, so that AWS uptime and performance are fully transparent.
- *No contract required.* In part because the costs are so low, all that is necessary to start using EC2 is a credit card.

Figure 6.15 shows the hourly price of the many types of EC2 instances in 2011. In addition to computation, EC2 charges for long-term storage and for Internet traffic. (There is no cost for network traffic inside AWS regions.) Elastic Block Storage costs \$0.10 per GByte per month and \$0.10 per million I/O requests. Internet traffic costs \$0.10 per GByte going to EC2 and \$0.08 to \$0.15 per GByte leaving from EC2, depending on the volume. Putting this into historical perspective, for \$100 per month you can use the equivalent capacity of the sum of the capacities of all magnetic disks produced in 1960!

Instance	Per hour	Ratio to small	Compute units	Virtual cores	Compute units/core	Memory (GB)	Disk (GB)	Address size
Micro	\$0.020	0.5–2.0	0.5–2.0	1	0.5–2.0	0.6	EBS	32/64 bit
Standard Small	\$0.085	1.0	1.0	1	1.00	1.7	160	32 bit
Standard Large	\$0.340	4.0	4.0	2	2.00	7.5	850	64 bit
Standard Extra Large	\$0.680	8.0	8.0	4	2.00	15.0	1690	64 bit
High-Memory Extra Large	\$0.500	5.9	6.5	2	3.25	17.1	420	64 bit
High-Memory Double Extra Large	\$1.000	11.8	13.0	4	3.25	34.2	850	64 bit
High-Memory Quadruple Extra Large	\$2.000	23.5	26.0	8	3.25	68.4	1690	64 bit
High-CPU Medium	\$0.170	2.0	5.0	2	2.50	1.7	350	32 bit
High-CPU Extra Large	\$0.680	8.0	20.0	8	2.50	7.0	1690	64 bit
Cluster Quadruple Extra Large	\$1.600	18.8	33.5	8	4.20	23.0	1690	64 bit

Figure 6.15 Price and characteristics of on-demand EC2 instances in the United States in the Virginia region in January 2011. Micro Instances are the newest and cheapest category, and they offer short bursts of up to 2.0 compute units for just \$0.02 per hour. Customers report that Micro Instances average about 0.5 compute units. Cluster-Compute Instances in the last row, which AWS identifies as dedicated dual-socket Intel Xeon X5570 servers with four cores per socket running at 2.93 GHz, offer 10 Gigabit/sec networks. They are intended for HPC applications. AWS also offers Spot Instances at much less cost, where you set the price you are willing to pay and the number of instances you are willing to run, and then AWS will run them when the spot price drops below your level. They run until you stop them or the spot price exceeds your limit. One sample during the daytime in January 2011 found that the spot price was a factor of 2.3 to 3.1 lower, depending on the instance type. AWS also offers Reserved Instances for cases where customers know they will use most of the instance for a year. You pay a yearly fee per instance and then an hourly rate that is about 30% of column 1 to use it. If you used a Reserved Instance 100% for a whole year, the average cost per hour including amortization of the annual fee would be about 65% of the rate in the first column. The server equivalent to those in Figures 6.13 and 6.14 would be a Standard Extra Large or High-CPU Extra Large Instance, which we calculated to cost \$0.11 per hour.

Example Calculate the cost of running the average MapReduce jobs in Figure 6.2 on page 437 on EC2. Assume there are plenty of jobs, so there is no significant extra cost to round up so as to get an integer number of hours. Ignore the monthly storage costs, but include the cost of disk I/Os for AWS's Elastic Block Storage (EBS). Next calculate the cost per year to run all the MapReduce jobs.

Answer The first question is what is the right size instance to match the typical server at Google? Figure 6.21 on page 467 in Section 6.7 shows that in 2007 a typical Google server had four cores running at 2.2 GHz with 8 GB of memory. Since a single instance is one virtual core that is equivalent to a 1 to 1.2 GHz AMD Opteron, the closest match in Figure 6.15 is a High-CPU Extra Large with eight virtual cores and 7.0 GB of memory. For simplicity, we'll assume the average EBS storage access is 64 KB in order to calculate the number of I/Os.

	Aug-04	Mar-06	Sep-07	Sep-09
Average completion time (hours)	0.15	0.21	0.10	0.11
Average number of servers per job	157	268	394	488
Cost per hour of EC2 High-CPU XL instance	\$0.68	\$0.68	\$0.68	\$0.68
Average EC2 cost per MapReduce job	\$16.35	\$38.47	\$25.56	\$38.07
Average number of EBS I/O requests (millions)	2.34	5.80	3.26	3.19
EBS cost per million I/O requests	\$0.10	\$0.10	\$0.10	\$0.10
Average EBS I/O cost per MapReduce job	\$0.23	\$0.58	\$0.33	\$0.32
Average total cost per MapReduce job	\$16.58	\$39.05	\$25.89	\$38.39
Annual number of MapReduce jobs	29,000	171,000	2,217,000	3,467,000
Total cost of MapReduce jobs on EC2/EBS	\$480,910	\$6,678,011	\$57,394,985	\$133,107,414

Figure 6.16 Estimated cost if you ran the Google MapReduce workload (Figure 6.2) using 2011 prices for AWS ECS and EBS (Figure 6.15). Since we are using 2011 prices, these estimates are less accurate for earlier years than for the more recent ones.

Figure 6.16 calculates the average and total cost per year of running the Google MapReduce workload on EC2. The average 2009 MapReduce job would cost a little under \$40 on EC2, and the total workload for 2009 would cost \$133M on AWS. Note that EBS accesses are about 1% of total costs for these jobs.

Example Given that the costs of MapReduce jobs are growing and already exceed \$100M per year, imagine that your boss wants you to investigate ways to lower costs. Two potentially lower cost options are either AWS Reserved Instances or AWS Spot Instances. Which would you recommend?

Answer AWS Reserved Instances charge a fixed annual rate plus an hourly per-use rate. In 2011, the annual cost for the High-CPU Extra Large Instance is \$1820 and the hourly rate is \$0.24. Since we pay for the instances whether they are used or not, let's assume that the average utilization of Reserved Instances is 80%. Then the average price per hour becomes:

$$\frac{\frac{\text{Annual price}}{\text{Hours per year}} + \text{Hourly price}}{\text{Utilization}} = \frac{\frac{\$1820}{8760} + \$0.24}{80\%} = (0.21 + 0.24) \times 1.25 = \$0.56$$

Thus, the savings using Reserved Instances would be roughly 17% or \$23M for the 2009 MapReduce workload.

Sampling a few days in January 2011, the hourly cost of a High-CPU Extra Large Spot Instance averages \$0.235. Since that is the minimum price to bid to get one server, that cannot be the average cost since you usually want to run tasks to completion without being bumped. Let's assume you need to pay double the minimum price to run large MapReduce jobs to completion. The cost savings for Spot Instances for the 2009 workload would be roughly 31% or \$41M.

Thus, you tentatively recommend Spot Instances to your boss since there is less of an up-front commitment and they may potentially save more money. However, you tell your boss you need to try to run MapReduce jobs on Spot Instances to see what you actually end up paying to ensure that jobs run to completion and that there really are hundreds of High-CPU Extra Large Instances available to run these jobs daily.

In addition to the low cost and a pay-for-use model of utility computing, another strong attractor for cloud computing users is that the cloud computing providers take on the risks of over-provisioning or under-provisioning. Risk avoidance is a godsend for startup companies, as either mistake could be fatal. If too much of the precious investment is spent on servers before the product is ready for heavy use, the company could run out of money. If the service suddenly became popular, but there weren't enough servers to match the demand, the company could make a very bad impression with the potential new customers it desperately needs to grow.

The poster child for this scenario is FarmVille from Zynga, a social networking game on Facebook. Before FarmVille was announced, the largest social game was about 5 million daily players. FarmVille had 1 million players 4 days after launch and 10 million players after 60 days. After 270 days, it had 28 million daily players and 75 million monthly players. Because they were deployed on AWS, they were able to grow seamlessly with the number of users. Moreover, it sheds load based on customer demand.

More established companies are taking advantage of the scalability of the cloud, as well. In 2011, Netflix migrated its Web site and streaming video service from a conventional datacenter to AWS. Netflix's goal was to let users watch a movie on, say, their cell phone while commuting home and then seamlessly switch to their television when they arrive home to continue watching their movie where they left off. This effort involves batch processing to convert new movies to the myriad formats they need to deliver movies on cell phones, tablets, laptops, game consoles, and digital video recorders. These batch AWS jobs can take thousands of machines several weeks to complete the conversions. The transactional backend for streaming is done in AWS and the delivery of encoded files is done via Content Delivery Networks such as Akamai and Level 3. The online service is much less expensive than mailing DVDs, and the resulting low cost has made the new service popular. One study put Netflix as 30% of Internet download traffic in the United States during peak evening periods. (In contrast, YouTube was just 10% in the same 8 p.m. to 10 p.m. period.) In fact, the overall average is 22% of Internet traffic, making Netflix alone responsible for the largest portion of Internet traffic in North America. Despite accelerating growth rates in Netflix subscriber accounts, the growth rate of Netflix's datacenter has been halted, and all capacity expansion going forward has been done via AWS.

Cloud computing has made the benefits of WSC available to everyone. Cloud computing offers cost associativity with the illusion of infinite scalability at no extra cost to the user: 1000 servers for 1 hour cost no more than 1 server for

1000 hours. It is up to the cloud computing provider to ensure that there are enough servers, storage, and Internet bandwidth available to meet the demand. The optimized supply chain mentioned above, which drops time-to-delivery to a week for new computers, is a considerable aid in providing that illusion without bankrupting the provider. This transfer of risks, cost associativity, and pay-as-you-go pricing is a powerful argument for companies of varying sizes to use cloud computing.

Two crosscutting issues that shape the cost-performance of WSCs and hence cloud computing are the WSC network and the efficiency of the server hardware and software.

6.6

Crosscutting Issues

Net gear is the SUV of the datacenter.

James Hamilton (2009)

WSC Network as a Bottleneck

Section 6.4 showed that the networking gear above the rack switch is a significant fraction of the cost of a WSC. Fully configured, the list price of a 128-port 1 Gbit datacenter switch from Juniper (EX8216) is \$716,000 without optical interfaces and \$908,000 with them. (These list prices are heavily discounted, but they still cost more than 50 times as much as a rack switch did.) These switches also tend to be power hungry. For example, the EX8216 consumes about 19,200 watts, which is 500 to 1000 times more than a server in a WSC. Moreover, these large switches are manually configured and fragile at a large scale. Because of their price, it is difficult to afford more than dual redundancy in a WSC using these large switches, which limits the options for fault tolerance [Hamilton 2009].

However, the real impact on switches is how oversubscription affects the design of software and the placement of services and data within the WSC. The ideal WSC network would be a black box whose topology and bandwidth are uninteresting because there are no restrictions: You could run any workload in any place and optimize for server utilization rather than network traffic locality. The WSC network bottlenecks today constrain data placement, which in turn complicates WSC software. As this software is one of the most valuable assets of a WSC company, the cost of this added complexity can be significant.

For readers interested in learning more about switch design, Appendix F describes the issues involved in the design of interconnection networks. In addition, Thacker [2007] proposed borrowing networking technology from supercomputing to overcome the price and performance problems. Vahdat et al. [2010] did as well, and proposed a networking infrastructure that can scale to 100,000 ports and 1 petabit/sec of bisection bandwidth. A major benefit of these novel datacenter switches is to simplify the software challenges due to oversubscription.

Using Energy Efficiently Inside the Server

While PUE measures the efficiency of a WSC, it has nothing to say about what goes on inside the IT equipment itself. Thus, another source of electrical inefficiency not covered in Figure 6.9 is the power supply *inside* the server, which converts input of 208 volts or 110 volts to the voltages that chips and disks use, typically 3.3, 5, and 12 volts. The 12 volts are further stepped down to 1.2 to 1.8 volts on the board, depending on what the microprocessor and memory require. In 2007, many power supplies were 60% to 80% efficient, which meant there were greater losses inside the server than there were going through the many steps and voltage changes from the high-voltage lines at the utility tower to supply the low-voltage lines at the server. One reason is that they have to supply a range of voltages to the chips and the disks, since they have no idea what is on the motherboard. A second reason is that the power supply is often oversized in watts for what is on the board. Moreover, such power supplies are often at their worst efficiency at 25% load or less, even though as Figure 6.3 on page 440 shows, many WSC servers operate in that range. Computer motherboards also have voltage regulator modules (VRMs), and they can have relatively low efficiency as well.

To improve the state of the art, Figure 6.17 shows the Climate Savers Computing Initiative standards [2007] for rating power supplies and their goals over time. Note that the standard specifies requirements at 20% and 50% loading in addition to 100% loading.

In addition to the power supply, Barroso and Hölzle [2007] said the goal for the whole server should be *energy proportionality*; that is, servers should consume energy in proportion to the amount of work performed. Figure 6.18 shows how far we are from achieving that ideal goal using SPECpower, a server benchmark that measures energy used at different performance levels (Chapter 1). The energy proportional line is added to the actual power usage of the most efficient server for SPECpower as of July 2010. Most servers will not be that efficient; it was up to 2.5 times better than other systems benchmarked that year, and late in a benchmark competition systems are often configured in ways to win the benchmark that are not typical of systems in the field. For example, the best-rated SPECpower servers use solid-state disks whose capacity is smaller than main memory! Even so, this very efficient system still uses almost 30% of the full

Loading conditioning	Base	Bronze (June 2008)	Silver (June 2009)	Gold (June 2010)
20%	80%	82%	85%	87%
50%	80%	85%	88%	90%
100%	80%	82%	85%	87%

Figure 6.17 Efficiency ratings and goals for power supplies over time of the Climate Savers Computing Initiative. These ratings are for Multi-Output Power Supply Units, which refer to desktop and server power supplies in nonredundant systems. There is a slightly higher standard for single-output PSUs, which are typically used in redundant configurations (1U/2U single-, dual-, and four-socket and blade servers).




Figure 6.18 The best SPECpower results as of July 2010 versus the ideal energy proportional behavior. The system was the HP ProLiant SL2x170z G6, which uses a cluster of four dual-socket Intel Xeon L5640s with each socket having six cores running at 2.27 GHz. The system had 64 GB of DRAM and a tiny 60 GB SSD for secondary storage. (The fact that main memory is larger than disk capacity suggests that this system was tailored to this benchmark.) The software used was IBM Java Virtual Machine version 9 and Windows Server 2008, Enterprise Edition.

power when idle and almost 50% of full power at just 10% load. Thus, energy proportionality remains a lofty goal instead of a proud achievement.

Systems software is designed to use all of an available resource if it potentially improves performance, without concern for the energy implications. For example, operating systems use all of memory for program data or for file caches, despite the fact that much of the data will likely never be used. Software architects need to consider energy as well as performance in future designs [Carter and Rajamani 2010].

Example Using the data of the kind in Figure 6.18, what is the saving in power going from five servers at 10% utilization versus one server at 50% utilization?

Answer A single server at 10% load is 308 watts and at 50% load is 451 watts. The savings is then

$$5 \times 308 / 451 = (1540 / 451) \approx 3.4$$

or about a factor of 3.4. If we want to be good environmental stewards in our WSC, we must consolidate servers when utilizations drop, purchase servers that are more energy proportional, or find something else that is useful to run in periods of low activity.

Given the background from these six sections, we are now ready to appreciate the work of the Google WSC architects.

6.7

Putting It All Together: A Google Warehouse-Scale Computer

Since many companies with WSCs are competing vigorously in the marketplace, up until recently, they have been reluctant to share their latest innovations with the public (and each other). In 2009, Google described a state-of-the-art WSC as of 2005. Google graciously provided an update of the 2007 status of their WSC, making this section the most up-to-date description of a Google WSC [Clidas, Johnson, and Felderman 2010]. Even more recently, Facebook described their latest datacenter as part of <http://opencompute.org>.

Containers

Both Google and Microsoft have built WSCs using shipping containers. The idea of building a WSC from containers is to make WSC design modular. Each container is independent, and the only external connections are networking, power, and water. The containers in turn supply networking, power, and cooling to the servers placed inside them, so the job of the WSC is to supply networking, power, and cold water to the containers and to pump the resulting warm water to external cooling towers and chillers.

The Google WSC that we are looking at contains 45 40-foot-long containers in a 300-foot by 250-foot space, or 75,000 square feet (about 7000 square meters). To fit in the warehouse, 30 of the containers are stacked two high, or 15 pairs of stacked containers. Although the location was not revealed, it was built at the time that Google developed WSCs in The Dalles, Oregon, which provides a moderate climate and is near cheap hydroelectric power and Internet backbone fiber. This WSC offers 10 megawatts with a PUE of 1.23 over the prior 12 months. Of that 0.230 of PUE overhead, 85% goes to cooling losses (0.195 PUE) and 15% (0.035) goes to power losses. The system went live in November 2005, and this section describes its state as of 2007.

A Google container can handle up to 250 kilowatts. That means the container can handle 780 watts per square foot (0.09 square meters), or 133 watts per square foot across the entire 75,000-square-foot space with 40 containers. However, the containers in this WSC average just 222 kilowatts

Figure 6.19 is a cutaway drawing of a Google container. A container holds up to 1160 servers, so 45 containers have space for 52,200 servers. (This WSC has about 40,000 servers.) The servers are stacked 20 high in racks that form two long rows of 29 racks (also called *bays*) each, with one row on each side of the container. The rack switches are 48-port, 1 Gbit/sec Ethernet switches, which are placed in every other rack.




Figure 6.19 Google customizes a standard 1AAA container: 40 x 8 x 9.5 feet (12.2 x 2.4 x 2.9 meters). The servers are stacked up to 20 high in racks that form two long rows of 29 racks each, with one row on each side of the container. The cool aisle goes down the middle of the container, with the hot air return being on the outside. The hanging rack structure makes it easier to repair the cooling system without removing the servers. To allow people inside the container to repair components, it contains safety systems for fire detection and mist-based suppression, emergency egress and lighting, and emergency power shut-off. Containers also have many sensors: temperature, airflow pressure, air leak detection, and motion-sensing lighting. A video tour of the datacenter can be found at <http://www.google.com/corporate/green/datacenters/summit.html>. Microsoft, Yahoo!, and many others are now building modular datacenters based upon these ideas but they have stopped using ISO standard containers since the size is inconvenient.

Cooling and Power in the Google WSC

Figure 6.20 is a cross-section of the container that shows the airflow. The computer racks are attached to the ceiling of the container. The cooling is below a raised floor that blows into the aisle between the racks. Hot air is returned from behind the racks. The restricted space of the container prevents the mixing of hot and cold air, which improves cooling efficiency. Variable-speed fans are run at the lowest speed needed to cool the rack as opposed to a constant speed.

The “cold” air is kept 81°F (27°C), which is balmy compared to the temperatures in many conventional datacenters. One reason datacenters traditionally run so cold is not for the IT equipment, but so that hot spots within the datacenter don’t cause isolated problems. By carefully controlling airflow to prevent hot spots, the container can run at a much higher temperature.




Figure 6.20 Airflow within the container shown in Figure 6.19. This cross-section diagram shows two racks on each side of the container. Cold air blows into the aisle in the middle of the container and is then sucked into the servers. Warm air returns at the edges of the container. This design isolates cold and warm airflows.

External chillers have cutouts so that, if the weather is right, only the outdoor cooling towers need cool the water. The chillers are skipped if the temperature of the water leaving the cooling tower is 70°F (21°C) or lower.

Note that if it's too cold outside, the cooling towers need heaters to prevent ice from forming. One of the advantages of placing a WSC in The Dalles is that the annual wet-bulb temperature ranges from 15°F to 66°F (−9°C to 19°C) with an average of 41°F (5°C), so the chillers can often be turned off. In contrast,




Figure 6.21 Server for Google WSC. The power supply is on the left and the two disks are on the top. The two fans below the left disk cover the two sockets of the AMD Barcelona microprocessor, each with two cores, running at 2.2 GHz. The eight DIMMs in the lower right each hold 1 GB, giving a total of 8 GB. There is no extra sheet metal, as the servers are plugged into the battery and a separate plenum is in the rack for each server to help control the airflow. In part because of the height of the batteries, 20 servers fit in a rack.

Las Vegas, Nevada, ranges from -42°F to 62°F (-41°C to 17°C) with an average of 29°F (-2°C). In addition, having to cool only to 81°F (27°C) inside the container makes it much more likely that Mother Nature will be able to cool the water.

Figure 6.21 shows the server designed by Google for this WSC. To improve efficiency of the power supply, it only supplies 12 volts to the motherboard and the motherboard supplies just enough for the number of disks it has on the board. (Laptops power their disks similarly.) The server norm is to supply the many voltage levels needed by the disks and chips directly. This simplification means the 2007 power supply can run at 92% efficiency, going far above the Gold rating for power supplies in 2010 (Figure 6.17).

Google engineers realized that 12 volts meant that the UPS could simply be a standard battery on each shelf. Hence, rather than have a separate battery room, which Figure 6.9 shows as 94% efficient, each server has its own lead acid battery that is 99.99% efficient. This “distributed UPS” is deployed incrementally with each machine, which means there is no money or power spent on overcapacity. They use standard off-the-shelf UPS units to protect network switches.

What about saving power by using dynamic voltage-frequency scaling (DVFS), which Chapter 1 describes? DVFS was not deployed in this family of machines since the impact on latency was such that it was only feasible in very low activity regions for online workloads, and even in those cases the system-wide savings were very small. The complex management control loop needed to deploy it therefore could not be justified.




Figure 6.22 Power usage effectiveness (PUE) of 10 Google WSCs over time. Google A is the WSC described in this section. It is the highest line in Q3 '07 and Q2 '10. (From www.google.com/corporate/green/datacenters/measuring.htm.) Facebook recently announced a new datacenter that should deliver an impressive PUE of 1.07 (see <http://opencompute.org/>). The Prineville Oregon Facility has no air conditioning and no chilled water. It relies strictly on outside air, which is brought in one side of the building, filtered, cooled via misters, pumped across the IT equipment, and then sent out the building by exhaust fans. In addition, the servers use a custom power supply that allows the power distribution system to skip one of the voltage conversion steps in Figure 6.9.

One of the keys to achieving the PUE of 1.23 was to put measurement devices (called *current transformers*) in all circuits throughout the containers and elsewhere in the WSC to measure the actual power usage. These measurements allowed Google to tune the design of the WSC over time.

Google publishes the PUE of its WSCs each quarter. Figure 6.22 plots the PUE for 10 Google WSCs from the third quarter in 2007 to the second quarter in 2010; this section describes the WSC labeled Google A. Google E operates with a PUE of 1.16 with cooling being only 0.105, due to the higher operational temperatures and chiller cutouts. Power distribution is just 0.039, due to the distributed UPS and single voltage power supply. The best WSC result was 1.12, with Google A at 1.23. In April 2009, the trailing 12-month average weighted by usage across all datacenters was 1.19.

Servers in a Google WSC

The server in Figure 6.21 has two sockets, each containing a dual-core AMD Opteron processor running at 2.2 GHz. The photo shows eight DIMMs, and

these servers are typically deployed with 8 GB of DDR2 DRAM. A novel feature is that the memory bus is downclocked to 533 MHz from the standard 666 MHz since the slower bus has little impact on performance but a significant impact on power.

The baseline design has a single network interface card (NIC) for a 1 Gbit/sec Ethernet link. Although the photo in Figure 6.21 shows two SATA disk drives, the baseline server has just one. The peak power of the baseline is about 160 watts, and idle power is 85 watts.

This baseline node is supplemented to offer a storage (or “diskfull”) node. First, a second tray containing 10 SATA disks is connected to the server. To get one more disk, a second disk is placed into the empty spot on the motherboard, giving the storage node 12 SATA disks. Finally, since a storage node could saturate a single 1 Gbit/sec Ethernet link, a second Ethernet NIC was added. Peak power for a storage node is about 300 watts, and it idles at 198 watts.

Note that the storage node takes up two slots in the rack, which is one reason why Google deployed 40,000 instead of 52,200 servers in the 45 containers. In this facility, the ratio was about two compute nodes for every storage node, but that ratio varied widely across Google’s WSCs. Hence, Google A had about 190,000 disks in 2007, or an average of almost 5 disks per server.

Networking in a Google WSC

The 40,000 servers are divided into three arrays of more than 10,000 servers each. (Arrays are called *clusters* in Google terminology.) The 48-port rack switch uses 40 ports to connect to servers, leaving 8 for uplinks to the array switches.

Array switches are configured to support up to 480 1 Gbit/sec Ethernet links and a few 10 Gbit/sec ports. The 1 Gigabit ports are used to connect to the rack switches, as each rack switch has a single link to each of the array switches. The 10 Gbit/sec ports connect to each of two datacenter routers, which aggregate all array routers and provide connectivity to the outside world. The WSC uses two datacenter routers for dependability, so a single datacenter router failure does not take out the whole WSC.

The number of uplink ports used per rack switch varies from a minimum of 2 to a maximum of 8. In the dual-port case, rack switches operate at an oversubscription rate of 20:1. That is, there is 20 times the network bandwidth inside the switch as there was exiting the switch. Applications with significant traffic demands beyond a rack tended to suffer from poor network performance. Hence, the 8-port uplink design, which provided a lower oversubscription rate of just 5:1, was used for arrays with more demanding traffic requirements.

Monitoring and Repair in a Google WSC

For a single operator to be responsible for more than 1000 servers, you need an extensive monitoring infrastructure and some automation to help with routine events.

Google deploys monitoring software to track the health of all servers and networking gear. Diagnostics are running all the time. When a system fails, many of the possible problems have simple automated solutions. In this case, the next step is to reboot the system and then to try to reinstall software components. Thus, the procedure handles the majority of the failures.

Machines that fail these first steps are added to a queue of machines to be repaired. The diagnosis of the problem is placed into the queue along with the ID of the failed machine.

To amortize the cost of repair, failed machines are addressed in batches by repair technicians. When the diagnosis software is confident in its assessment, the part is immediately replaced without going through the manual diagnosis process. For example, if the diagnostic says disk 3 of a storage node is bad, the disk is replaced immediately. Failed machines with no diagnostic or with low-confidence diagnostics are examined manually.

The goal is to have less than 1% of all nodes in the manual repair queue at any one time. The average time in the repair queue is a week, even though it takes much less time for repair technician to fix it. The longer latency suggests the importance of repair throughput, which affects cost of operations. Note that the automated repairs of the first step take minutes for a reboot/reinstall to hours for running directed stress tests to make sure the machine is indeed operational.

These latencies do not take into account the time to idle the broken servers. The reason is that a big variable is the amount of state in the node. A stateless node takes much less time than a storage node whose data may need to be evacuated before it can be replaced.

Summary

As of 2007, Google had already demonstrated several innovations to improve the energy efficiency of its WSCs to deliver a PUE of 1.23 in Google A:

- In addition to providing an inexpensive shell to enclose servers, the modified shipping containers separate hot and cold air plenums, which helps reduce the variation in intake air temperature for servers. With less severe worst-case hot spots, cold air can be delivered at warmer temperatures.
- These containers also shrink the distance of the air circulation loop, which reduces energy to move air.
- Operating servers at higher temperatures means that air only has to be chilled to 81°F (27°C) instead of the traditional 64°F to 71°F (18°C to 22°C).
- A higher target cold air temperature helps put the facility more often within the range that can be sustained by evaporative cooling solutions (cooling towers), which are more energy efficient than traditional chillers.
- Deploying WSCs in temperate climates to allow use of evaporative cooling exclusively for portions of the year.
- Deploying extensive monitoring hardware and software to measure actual PUE versus designed PUE improves operational efficiency.

- Operating more servers than the worst-case scenario for the power distribution system would suggest, since it's statistically unlikely that thousands of servers would all be highly busy simultaneously, yet rely on the monitoring system to off-load work in the unlikely case that they did [Fan, Weber, and Barroso 2007] [Ranganathan et al. 2006]. PUE improves because the facility is operating closer to its fully designed capacity, where it is at its most efficient because the servers and cooling systems are not energy proportional. Such increased utilization reduces demand for new servers and new WSCs.
- Designing motherboards that only need a single 12-volt supply so that the UPS function could be supplied by standard batteries associated with each server instead of a battery room, thereby lowering costs and reducing one source of inefficiency of power distribution within a WSC.
- Carefully designing the server board itself to improve its energy efficiency. For example, underclocking the front-side bus on these microprocessors reduces energy usage with negligible performance impact. (Note that such optimizations do not impact PUE but do reduce overall WSC energy consumption.)

WSC design must have improved in the intervening years, as Google's best WSC has dropped the PUE from 1.23 for Google A to 1.12. Facebook announced in 2011 that they had driven PUE down to 1.07 in their new datacenter (see <http://opencompute.org/>). It will be interesting to see what innovations remain to improve further the WSC efficiency so that we are good guardians of our environment. Perhaps in the future we will even consider the energy cost to *manufacture* the equipment within a WSC [Chang et al. 2010].

6.8 Fallacies and Pitfalls

Despite WSC being less than a decade old, WSC architects like those at Google have already uncovered many pitfalls and fallacies about WSCs, often learned the hard way. As we said in the introduction, WSC architects are today's Seymour Crays.

Fallacy *Cloud computing providers are losing money.*

A popular question about cloud computing is whether it's profitable at these low prices.

Based on AWS pricing from Figure 6.15, we could charge \$0.68 per hour per server for computation. (The \$0.085 per hour price is for a Virtual Machine equivalent to one EC2 compute unit, not a full server.) If we could sell 50% of the server hours, that would generate \$0.34 of income per hour per server. (Note that customers pay no matter how little they use the servers they occupy, so selling 50% of the server hours doesn't necessarily mean that average server utilization is 50%.)

Another way to calculate income would be to use AWS *Reserved Instances*, where customers pay a yearly fee to reserve an instance and then a lower rate per

hour to use it. Combining the charges together, AWS would receive \$0.45 of income per hour per server for a full year.

If we could sell 750 GB per server for storage using AWS pricing, in addition to the computation income, that would generate another \$75 per month per server, or another \$0.10 per hour.

These numbers suggest an average income of \$0.44 per hour per server (via On-Demand Instances) to \$0.55 per hour (via Reserved Instances). From Figure 6.13, we calculated the cost per server as \$0.11 per hour for the WSC in Section 6.4. Although the costs in Figure 6.13 are estimates that are *not* based on actual AWS costs and the 50% sales for server processing and 750 GB utilization of per server storage are just examples, these assumptions suggest a gross margin of 75% to 80%. Assuming these calculations are reasonable, they suggest that cloud computing is profitable, especially for a service business.

Fallacy *Capital costs of the WSC facility are higher than for the servers that it houses.*

While a quick look at Figure 6.13 on page 453 might lead you to that conclusion, that glimpse ignores the length of amortization for each part of the full WSC. However, the facility lasts 10 to 15 years while the servers need to be repurchased every 3 or 4 years. Using the amortization times in Figure 6.13 of 10 years and 3 years, respectively, the capital expenditures over a decade are \$72M for the facility and $3.3 \times \$67\text{M}$, or \$221M, for servers. Thus, the capital costs for servers in a WSC over a decade are a factor of three higher than for the WSC facility.

Pitfall *Trying to save power with inactive low power modes versus active low power modes.*

Figure 6.3 on page 440 shows that the average utilization of servers is between 10% and 50%. Given the concern on operational costs of a WSC from Section 6.4, you would think low power modes would be a huge help.

As Chapter 1 mentions, you cannot access DRAMs or disks in these *inactive low power modes*, so you must return to fully active mode to read or write, no matter how low the rate. The pitfall is that the time and energy required to return to fully active mode make inactive low power modes less attractive. Figure 6.3 shows that almost all servers average at least 10% utilization, so you might expect long periods of low activity but not long periods of inactivity.

In contrast, processors still run in lower power modes at a small multiple of the regular rate, so *active low power modes* are much easier to use. Note that the time to move to fully active mode for processors is also measured in microseconds, so active low power modes also address the latency concerns about low power modes.

Pitfall *Using too wimpy a processor when trying to improve WSC cost-performance.*

Amdahl's law still applies to WSC, as there will be some serial work for each request, and that can increase request latency if it runs on a slow server [Hölzle 2010] [Lim et al. 2008]. If the serial work increases latency, then the cost of using a wimpy processor must include the software development costs to optimize the

code to return it to the lower latency. The larger number of threads of many slow servers can also be more difficult to schedule and load balance, and thus the variability in thread performance can lead to longer latencies. A 1 in 1000 chance of bad scheduling is probably not an issue with 10 tasks, but it is with 1000 tasks when you have to wait for the longest task. Many smaller servers can also lead to lower utilization, as it's clearly easier to schedule when there are fewer things to schedule. Finally, even some parallel algorithms get less efficient when the problem is partitioned too finely. The Google rule of thumb is currently to use the low-end range of server class computers [Barroso and Hölzle 2009].

As a concrete example, Reddi et al. [2010] compared embedded microprocessors (*Atom*) and server microprocessors (*Nehalem Xeon*) running the Bing search engine. They found that the latency of a query was about three times longer on *Atom* than on *Xeon*. Moreover, the *Xeon* was more robust. As load increases on *Xeon*, quality of service degrades gradually and modestly. *Atom* quickly violates its quality-of-service target as it tries to absorb additional load.

This behavior translates directly into search quality. Given the importance of latency to the user, as Figure 6.12 suggests, the Bing search engine uses multiple strategies to refine search results if the query latency has not yet exceeded a cut-off latency. The lower latency of the larger *Xeon* nodes means they can spend more time refining search results. Hence, even when the *Atom* had almost no load, it gave worse answers in 1% of the queries than *Xeon*. At normal loads, 2% of the answers were worse.

Fallacy *Given improvements in DRAM dependability and the fault tolerance of WSC systems software, you don't need to spend extra for ECC memory in a WSC.*

Since ECC adds 8 bits to every 64 bits of DRAM, potentially you could save a ninth of the DRAM costs by eliminating error-correcting code (ECC), especially since measurements of DRAM had claimed failure rates of 1000 to 5000 FIT (failures per billion hours of operation) per megabit [Tezzaron Semiconductor 2004].

Schroeder, Pinheiro, and Weber [2009] studied measurements of the DRAMs with ECC protection at the majority of Google's WSCs, which was surely many hundreds of thousands of servers, over a 2.5-year period. They found 15 to 25 times higher FIT rates than had been published, or 25,000 to 70,000 failures per megabit. Failures affected more than 8% of DIMMs, and the average DIMM had 4000 correctable errors and 0.2 uncorrectable errors per year. Measured at the server, about a third experienced DRAM errors each year, with an average of 22,000 correctable errors and 1 uncorrectable error per year. That is, for one-third of the servers, one memory error is corrected every 2.5 hours. Note that these systems used the more powerful chipkill codes rather than the simpler SECDED codes. If the simpler scheme had been used, the uncorrectable error rates would have been 4 to 10 times higher.

In a WSC that only had parity error protection, the servers would have to reboot for each memory parity error. If the reboot time were 5 minutes, one-third of the machines would spend 20% of their time rebooting! Such behavior would

lower the performance of the \$150M facility by about 6%. Moreover, these systems would suffer many uncorrectable errors without operators being notified that they occurred.

In the early years, Google used DRAM that didn't even have parity protection. In 2000, during testing before shipping the next release of the search index, it started suggesting random documents in response to test queries [Barroso and Hölzle 2009]. The reason was a stuck-at-zero fault in some DRAMs, which corrupted the new index. Google added consistency checks to detect such errors in the future. As WSC grew in size and as ECC DIMMs became more affordable, ECC became the standard in Google WSCs. ECC has the added benefit of making it much easier to find broken DIMMs during repair.

Such data suggest why the Fermi GPU (Chapter 4) adds ECC to its memory where its predecessors didn't even have parity protection. Moreover, these FIT rates cast doubts on efforts to use the Intel Atom processor in a WSC—due to its improved power efficiency—since the 2011 chip set does not support ECC DRAM.

Fallacy *Turning off hardware during periods of low activity improves cost-performance of a WSC.*

Figure 6.14 on page 454 shows that the cost of amortizing the power distribution and cooling infrastructure is 50% higher than the entire monthly power bill. Hence, while it certainly would save some money to compact workloads and turn off idle machines, even if you could save half the power it would only reduce the monthly operational bill by 7%. There would also be practical problems to overcome, since the extensive WSC monitoring infrastructure depends on being able to poke equipment and see it respond. Another advantage of energy proportionality and active low power modes is that they are compatible with the WSC monitoring infrastructure, which allows a single operator to be responsible for more than 1000 servers.

The conventional WSC wisdom is to run other valuable tasks during periods of low activity so as to recoup the investment in power distribution and cooling. A prime example is the batch MapReduce jobs that create indices for search. Another example of getting value from low utilization is spot pricing on AWS, which the caption in Figure 6.15 on page 458 describes. AWS users who are flexible about when their tasks are run can save a factor of 2.7 to 3 for computation by letting AWS schedule the tasks more flexibly using Spot Instances, such as when the WSC would otherwise have low utilization.

Fallacy *Replacing all disks with Flash memory will improve cost-performance of a WSC.*

Flash memory is much faster than disk for some WSC workloads, such as those doing many random reads and writes. For example, Facebook deployed Flash memory packaged as solid-state disks (SSDs) as a write-back cache called Flash-cache as part of its file system in its WSC, so that hot files stay in Flash and cold files stay on disk. However, since all performance improvements in a WSC must

be judged on cost-performance, before replacing all the disks with SSD the question is really I/Os per second per dollar and storage capacity per dollar. As we saw in Chapter 2, Flash memory costs at least 20 times more per GByte than magnetic disks: \$2.00/GByte versus \$0.09/Gbyte.

Narayanan et al. [2009] looked at migrating workloads from disk to SSD by simulating workload traces from small and large datacenters. Their conclusion was that SSDs were not cost effective for any of their workloads due to the low storage capacity per dollar. To reach the break-even point, Flash memory storage devices need to improve capacity per dollar by a factor of 3 to 3000, depending on the workload.

Even when you factor power into the equation, it's hard to justify replacing disk with Flash for data that are infrequently accessed. A one-terabyte disk uses about 10 watts of power, so, using the \$2 per watt-year rule of thumb from Section 6.4, the most you could save from reduced energy is \$20 a year per disk. However, the CAPEX cost in 2011 for a terabyte of storage is \$2000 for Flash and only \$90 for disk.

6.9 Concluding Remarks

Inheriting the title of building the world's biggest computers, computer architects of WSCs are designing the large part of the future IT that completes the mobile client. Many of us use WSCs many times a day, and the number of times per day and the number of people using WSCs will surely increase in the next decade. Already more than half of the nearly seven billion people on the planet have cell phones. As these devices become Internet ready, many more people from around the world will be able to benefit from WSCs.

Moreover, the economies of scale uncovered by WSC have realized the long dreamed of goal of computing as a utility. Cloud computing means anyone anywhere with good ideas and business models can tap thousands of servers to deliver their vision almost instantly. Of course, there are important obstacles that could limit the growth of cloud computing around standards, privacy, and the rate of growth of Internet bandwidth, but we foresee them being addressed so that cloud computing can flourish.

Given the increasing number of cores per chip (see Chapter 5), clusters will increase to include thousands of cores. We believe the technologies developed to run WSC will prove useful and trickle down to clusters, so that clusters will run the same virtual machines and systems software developed for WSC. One advantage would be easy support of "hybrid" datacenters, where the workload could easily be shipped to the cloud in a crunch and then shrink back afterwards to relying only on local computing.

Among the many attractive features of cloud computing is that it offers economic incentives for conservation. Whereas it is hard to convince cloud computing *providers* to turn off unused equipment to save energy given the

cost of the infrastructure investment, it is easy to convince cloud computing *users* to give up idle instances since they are paying for them whether or not they are doing anything useful. Similarly, charging by use encourages programmers to use computation, communication, and storage efficiently, which can be difficult to encourage without an understandable pricing scheme. The explicit pricing also makes it possible for researchers to evaluate innovations in cost-performance instead of just performance, since costs are now easily measured and believable. Finally, cloud computing means that researchers can evaluate their ideas at the scale of thousands of computers, which in the past only large companies could afford.

We believe that WSCs are changing the goals and principles of server design, just as the needs of mobile clients are changing the goals and principles of micro-processor design. Both are revolutionizing the software industry, as well. Performance per dollar and performance per joule drive both mobile client hardware and the WSC hardware, and parallelism is the key to delivering on those sets of goals.

Architects will play a vital role in both halves of this exciting future world. We look forward to seeing—and to using—what will come.

6.10 Historical Perspectives and References

Section L.8 (available online) covers the development of clusters that were the foundation of WSC and of utility computing. (Readers interested in learning more should start with Barroso and Hölzle [2009] and the blog postings and talks of James Hamilton at <http://perspectives.mvdirona.com>.)

Case Studies and Exercises by Parthasarathy Ranganathan

Case Study 1: Total Cost of Ownership Influencing Warehouse-Scale Computer Design Decisions

Concepts illustrated by this case study

- Total Cost of Ownership (TCO)
- Influence of Server Cost and Power on the Entire WSC
- Benefits and Drawbacks of Low-Power Servers

Total cost of ownership is an important metric for measuring the effectiveness of a warehouse-scale computer (WSC). TCO includes both the CAPEX and OPEX described in Section 6.4 and reflects the ownership cost of the entire datacenter to achieve a certain level of performance. In considering different servers, networks, and storage architectures, TCO is often the important comparison metric

used by datacenter owners to decide which options are best; however, TCO is a multidimensional computation that takes into account many different factors. The goal of this case study is to take a detailed look into WSCs, how different architectures influence TCO, and how TCO drives operator decisions. This case study will use the numbers from Figure 6.13 and Section 6.4, and assumes that the described WSC achieves the operator’s target level of performance. TCO is often used to compare different server options that have multiple dimensions. The exercises in this case study examine how such comparisons are made in the context of WSCs and the complexity involved in making the decisions.

- 6.1 [5/5/10] <6.2, 6.4> In this chapter, data-level parallelism has been discussed as a way for WSCs to achieve high performance on large problems. Conceivably, even greater performance can be obtained by using high-end servers; however, higher performance servers often come with a nonlinear price increase.
 - a. [5] <6.4> Assuming servers that are 10% faster at the same utilization, but 20% more expensive, what is the CAPEX for the WSC?
 - b. [5] <6.4> If those servers also use 15% more power, what is the OPEX?
 - c. [10] <6.2, 6.4> Given the speed improvement and power increase, what must the cost of the new servers be to be comparable to the original cluster? (*Hint:* Based on this TCO model, you may have to change the critical load of the facility.)
- 6.2 [5/10] <6.4, 6.8> To achieve a lower OPEX, one appealing alternative is to use low-power versions of servers to reduce the total electricity required to run the servers; however, similar to high-end servers, low-power versions of high-end components also have nonlinear trade-offs.
 - a. [5] <6.4, 6.8> If low-power server options offered 15% lower power at the same performance but are 20% more expensive, are they a good trade-off?
 - b. [10] <6.4, 6.8> At what cost do the servers become comparable to the original cluster? What if the price of electricity doubles?
- 6.3 [5/10/15] <6.4, 6.6> Servers that have different operating modes offer opportunities for dynamically running different configurations in the cluster to match workload usage. Use the data in Figure 6.23 for the power/performance modes for a given low-power server.
 - a. [5] <6.4, 6.6> If a server operator decided to save power costs by running all servers at medium performance, how many servers would be needed to achieve the same level of performance?

Mode	Performance	Power
High	100%	100%
Medium	75%	60%
Low	59%	38%

Figure 6.23 Power–performance modes for low-power servers.

- b. [10] <6.4, 6.6> What are the CAPEX and OPEX of such a configuration?
 - c. [15] <6.4, 6.6> If there was an alternative to purchase a server that is 20% cheaper but slower and uses less power, find the performance–power curve that provides a TCO comparable to the baseline server.
- 6.4 [Discussion] <6.4> Discuss the trade-offs and benefits of the two options in Exercise 6.3, assuming a constant workload being run on the servers.
- 6.5 [Discussion] <6.2, 6.4> Unlike high-performance computing (HPC) clusters, WSCs often experience significant workload fluctuation throughout the day. Discuss the trade-offs and benefits of the two options in Exercise 6.3, this time assuming a workload that varies.
- 6.6 [Discussion] <6.4, 6.7> The TCO model presented so far abstracts away a significant amount of lower level details. Discuss the impact of these abstractions to the overall accuracy of the TCO model. When are these abstractions safe to make? In what cases would greater detail provide significantly different answers?

Case Study 2: Resource Allocation in WSCs and TCO

Concepts illustrated by this case study

- Server and Power Provisioning within a WSC
- Time-Variance of Workloads
- Effects of Variance on TCO

Some of the key challenges to deploying efficient WSCs are provisioning resources properly and utilizing them to their fullest. This problem is complex due to the size of WSCs as well as the potential variance of the workloads being run. The exercises in this case study show how different uses of resources can affect TCO.

- 6.7 [5/5/10] <6.4> One of the challenges in provisioning a WSC is determining the proper power load, given the facility size. As described in the chapter, nameplate power is often a peak value that is rarely encountered.
- a. [5] <6.4> Estimate how the per-server TCO changes if the nameplate server power is 200 watts and the cost is \$3000.
 - b. [5] <6.4> Also consider a higher power, but cheaper option whose power is 300 watts and costs \$2000.
 - c. [10] <6.4> How does the per-server TCO change if the actual average power usage of the servers is only 70% of the nameplate power?
- 6.8 [15/10] <6.2, 6.4> One assumption in the TCO model is that the critical load of the facility is fixed, and the amount of servers fits that critical load. In reality, due to the variations of server power based on load, the critical power used by a facility can vary at any given time. Operators must initially provision the datacenter

based on its critical power resources and an estimate of how much power is used by the datacenter components.

- a. [15] <6.2, 6.4> Extend the TCO model to initially provision a WSC based on a server with a nameplate power of 300 watts, but also calculate the actual monthly critical power used and TCO assuming the server averages 40% utilization and 225 watts. How much capacity is left unused?
 - b. [10] <6.2, 6.4> Repeat this exercise with a 500-watt server that averages 20% utilization and 300 watts.
- 6.9 [10] <6.4, 6.5> WSCs are often used in an interactive manner with end users, as mentioned in Section 6.5. This interactive usage often leads to time-of-day fluctuations, with peaks correlating to specific time periods. For example, for Netflix rentals, there is a peak during the evening periods of 8 to 10 p.m.; the entirety of these time-of-day effects is significant. Compare the per-server TCO of a data-center with a capacity to match the utilization at 4 a.m. compared to 9 p.m.
- 6.10 [Discussion/15] <6.4, 6.5> Discuss some options to better utilize the excess servers during the off-peak hours or options to save costs. Given the interactive nature of WSCs, what are some of the challenges to aggressively reducing power usage?
- 6.11 [Discussion/25] <6.4, 6.6> Propose one possible way to improve TCO by focusing on reducing server power. What are the challenges to evaluating your proposal? Estimate the TCO improvements based on your proposal. What are advantages and drawbacks?

Exercises

- 6.12 [10/10/10] <6.1> One of the important enablers of WSC is ample request-level parallelism, in contrast to instruction or thread-level parallelism. This question explores the implication of different types of parallelism on computer architecture and system design.
- a. [10] <6.1> Discuss scenarios where improving the instruction- or thread-level parallelism would provide greater benefits than achievable through request-level parallelism.
 - b. [10] <6.1> What are the software design implications of increasing request-level parallelism?
 - c. [10] <6.1> What are potential drawbacks of increasing request-level parallelism?
- 6.13 [Discussion/15/15] <6.2> When a cloud computing service provider receives jobs consisting of multiple Virtual Machines (VMs) (e.g., a MapReduce job), many scheduling options exist. The VMs can be scheduled in a round-robin manner to spread across all available processors and servers or they can be consolidated to use as few processors as possible. Using these scheduling options, if a job with 24 VMs was submitted and 30 processors were available in the cloud (each able to run up to 3 VMs), round-robin would use 24 processors,

while consolidated scheduling would use 8 processors. The scheduler can also find available processor cores at different scopes: socket, server, rack, and an array of racks.

- a. [Discussion] <6.2> Assuming that the submitted jobs are all compute-heavy workloads, possibly with different memory bandwidth requirements, what are the pros and cons of round-robin versus consolidated scheduling in terms of power and cooling costs, performance, and reliability?
 - b. [15] <6.2> Assuming that the submitted jobs are all I/O-heavy workloads, what are the pros and cons of round-robin versus consolidated scheduling, at different scopes?
 - c. [15] <6.2> Assuming that the submitted jobs are network-heavy workloads, what are the pros and cons of round-robin versus consolidated scheduling, at different scopes?
- 6.14 [15/15/10/10] <6.2, 6.3> MapReduce enables large amounts of parallelism by having data-independent tasks run on multiple nodes, often using commodity hardware; however, there are limits to the level of parallelism. For example, for redundancy, MapReduce will write data blocks to multiple nodes, consuming disk and potentially network bandwidth. Assume a total dataset size of 300 GB, a network bandwidth of 1 Gb/sec, a 10 sec/GB map rate, and a 20 sec/GB reduce rate. Also assume that 30% of the data must be read from remote nodes, and each output file is written to two other nodes for redundancy. Use Figure 6.6 for all other parameters.
- a. [15] <6.2, 6.3> Assume that all nodes are in the same rack. What is the expected runtime with 5 nodes? 10 nodes? 100 nodes? 1000 nodes? Discuss the bottlenecks at each node size.
 - b. [15] <6.2, 6.3> Assume that there are 40 nodes per rack and that any remote read/write has an equal chance of going to any node. What is the expected runtime at 100 nodes? 1000 nodes?
 - c. [10] <6.2, 6.3> An important consideration is minimizing data movement as much as possible. Given the significant slowdown of going from local to rack to array accesses, software must be strongly optimized to maximize locality. Assume that there are 40 nodes per rack, and 1000 nodes are used in the MapReduce job. What is the runtime if remote accesses are within the same rack 20% of the time? 50% of the time? 80% of the time?
 - d. [10] <6.2, 6.3> Given the simple MapReduce program in Section 6.2, discuss some possible optimizations to maximize the locality of the workload.
- 6.15 [20/20/10/20/20/20] <6.2> WSC programmers often use data replication to overcome failures in the software. Hadoop HDFS, for example, employs three-way replication (one local copy, one remote copy in the rack, and one remote copy in a separate rack), but it's worth examining when such replication is needed.
- a. [20] <6.2> A Hadoop World 2010 attendee survey showed that over half of the Hadoop clusters had 10 nodes or less, with dataset sizes of 10 TB or less.

- Using the failure frequency data in Figure 6.1, what kind of availability does a 10-node Hadoop cluster have with one-, two-, and three-way replications?
- b. [20] <6.2> Assuming the failure data in Figure 6.1 and a 1000-node Hadoop cluster, what kind of availability does it have with one-, two-, and three-way replications?
 - c. [10] <6.2> The relative overhead of replication varies with the amount of data written per local compute hour. Calculate the amount of extra I/O traffic and network traffic (within and across rack) for a 1000-node Hadoop job that sorts 1 PB of data, where the intermediate results for data shuffling are written to the HDFS.
 - d. [20] <6.2> Using Figure 6.6, calculate the time overhead for two- and three-way replications. Using the failure rates shown in Figure 6.1, compare the expected execution times for no replication versus two- and three-way replications.
 - e. [20] <6.2> Now consider a database system applying replication on logs, assuming each transaction on average accesses the hard disk once and generates 1 KB of log data. Calculate the time overhead for two- and three-way replications. What if the transaction is executed in-memory and takes 10 μ s?
 - f. [20] <6.2> Now consider a database system with ACID consistency that requires two network round-trips for two-phase commitment. What is the time overhead for maintaining consistency as well as replications?
- 6.16 [15/15/20/15/] <6.1, 6.2, 6.8> Although request-level parallelism allows many machines to work on a single problem in parallel, thereby achieving greater overall performance, one of the challenges is avoiding dividing the problem too finely. If we look at this problem in the context of service level agreements (SLAs), using smaller problem sizes through greater partitioning can require increased effort to achieve the target SLA. Assume an SLA of 95% of queries respond at 0.5 sec or faster, and a parallel architecture similar to MapReduce that can launch multiple redundant jobs to achieve the same result. For the following questions, assume the query-response time curve shown in Figure 6.24. The curve shows the latency of response, based on the number of queries per second, for a baseline server as well as a “small” server that uses a slower processor model.
- a. [15] <6.1, 6.2, 6.8> How many servers are required to achieve that SLA, assuming that the WSC receives 30,000 queries per second, and the query-response time curve shown in Figure 6.24? How many “small” servers are required to achieve that SLA, given this response-time probability curve? Looking only at server costs, how much cheaper must the “wimpy” servers be than the normal servers to achieve a cost advantage for the target SLA?
 - b. [15] <6.1, 6.2, 6.8> Often “small” servers are also less reliable due to cheaper components. Using the numbers from Figure 6.1, assume that the number of events due to flaky machines and bad memories increases by 30%. How many “small” servers are required now? How much cheaper must those servers be than the standard servers?




Figure 6.24 Query-response time curve.

- c. [20] <6.1, 6.2, 6.8> Now assume a batch processing environment. The “small” servers provide 30% of the overall performance of the regular servers. Still assuming the reliability numbers from Exercise 6.15 part (b), how many “wimpy” nodes are required to provide the same expected throughput of a 2400-node array of standard servers, assuming perfect linear scaling of performance to node size and an average task length of 10 minutes per node? What if the scaling is 85%? 60%?
- d. [15] <6.1, 6.2, 6.8> Often the scaling is not a linear function, but instead a logarithmic function. A natural response may be instead to purchase larger nodes that have more computational power per node to minimize the array size. Discuss some of the trade-offs with this architecture.
- 6.17 [10/10/15] <6.3, 6.8> One trend in high-end servers is toward the inclusion of nonvolatile Flash memory in the memory hierarchy, either through solid-state disks (SSDs) or PCI Express-attached cards. Typical SSDs have a bandwidth of 250 MB/sec and latency of 75 μ s, whereas PCIe cards have a bandwidth of 600 MB/sec and latency of 35 μ s.
- a. [10] Take Figure 6.7 and include these points in the local server hierarchy. Assuming that identical performance scaling factors as DRAM are accessed at different hierarchy levels, how do these Flash memory devices compare when accessed across the rack? Across the array?
- b. [10] Discuss some software-based optimizations that can utilize the new level of the memory hierarchy.
- c. [25] Repeat part (a), instead assuming that each node has a 32 GB PCIe card that is able to cache 50% of all disk accesses.
- d. [15] As discussed in “Fallacies and Pitfalls” (Section 6.8), replacing all disks with SSDs is not necessarily a cost-effective strategy. Consider a WSC operator that uses it to provide cloud services. Discuss some scenarios where using SSDs or other Flash memory would make sense.

- 6.18 [20/20/Discussion] <6.3> *Memory Hierarchy*: Caching is heavily used in some WSC designs to reduce latency, and there are multiple caching options to satisfy varying access patterns and requirements.
- [20] Let's consider the design options for streaming rich media from the Web (e.g., Netflix). First we need to estimate the number of movies, number of encode formats per movie, and concurrent viewing users. In 2010, Netflix had 12,000 titles for online streaming, each title having at least four encode formats (at 500, 1000, 1600, and 2200 kbps). Let's assume that there are 100,000 concurrent viewers for the entire site, and an average movie is one hour long. Estimate the total storage capacity, I/O and network bandwidths, and video-streaming-related computation requirements.
 - [20] What are the access patterns and reference locality characteristics per user, per movie, and across all movies? (*Hint*: Random versus sequential, good versus poor temporal and spatial locality, relatively small versus large working set size.)
 - [Discussion] What movie storage options exist by using DRAM, SSD, and hard drives? Compare them in performance and TCO.
- 6.19 [10/20/20/Discussion/Discussion] <6.3> Consider a social networking Web site with 100 million active users posting updates about themselves (in text and pictures) as well as browsing and interacting with updates in their social networks. To provide low latency, Facebook and many other Web sites use memcached as a caching layer before the backend storage/database tiers.
- [10] Estimate the data generation and request rates per user and across the entire site.
 - [20] For the social networking Web site discussed here, how much DRAM is needed to host its working set? Using servers each having 96 GB DRAM, estimate how many local versus remote memory accesses are needed to generate a user's home page?
 - [20] Now consider two candidate memcached server designs, one using conventional Xeon processors and the other using smaller cores, such as Atom processors. Given that memcached requires large physical memory but has low CPU utilization, what are the pros and cons of these two designs?
 - [Discussion] Today's tight coupling between memory modules and processors often requires an increase in CPU socket count in order to provide large memory support. List other designs to provide large physical memory without proportionally increasing the number of sockets in a server. Compare them based on performance, power, costs, and reliability.
 - [Discussion] The same user's information can be stored in both the memcached and storage servers, and such servers can be physically hosted in different ways. Discuss the pros and cons of the following server layout in the WSC: (1) memcached collocated on the same storage server, (2) memcached and storage server on separate nodes in the same rack, or (3) memcached servers on the same racks and storage servers collocated on separate racks.

- 6.20 [5/5/10/10/Discussion/Discussion] <6.3, 6.6> *Datacenter Networking*: MapReduce and WSC are a powerful combination to tackle large-scale data processing; for example, Google in 2008 sorted one petabyte (1 PB) of records in a little more than 6 hours using 4000 servers and 48,000 hard drives.
- [5] Derive disk bandwidth from Figure 6.1 and associated text. How many seconds does it take to read the data into main memory and write the sorted results back?
 - [5] Assuming each server has two 1 Gb/sec Ethernet network interface cards (NICs) and the WSC switch infrastructure is oversubscribed by a factor of 4, how many seconds does it take to shuffle the entire dataset across 4000 servers?
 - [10] Assuming network transfer is the performance bottleneck for petabyte sort, can you estimate what oversubscription ratio Google has in their datacenter?
 - [10] Now let's examine the benefits of having 10 Gb/sec Ethernet without oversubscription—for example, using a 48-port 10 Gb/sec Ethernet (as used by the 2010 Indy sort benchmark winner TritonSort). How long does it take to shuffle the 1 PB of data?
 - [Discussion] Compare the two approaches here: (1) the massively scale-out approach with high network oversubscription ratio, and (2) a relatively small-scale system with a high-bandwidth network. What are their potential bottlenecks? What are their advantages and disadvantages, in terms of scalability and TCO?
 - [Discussion] Sort and many important scientific computing workloads are communication heavy, while many other workloads are not. List three example workloads that do not benefit from high-speed networking. What EC2 instances would you recommend to use for these two classes of workloads?
- 6.21 [10/25/Discussion] <6.4, 6.6> Because of the massive scale of WSCs, it is very important to properly allocate network resources based on the workloads that are expected to be run. Different allocations can have significant impacts on both the performance and total cost of ownership.
- [10] Using the numbers in the spreadsheet detailed in Figure 6.13, what is the oversubscription ratio at each access-layer switch? What is the impact on TCO if the oversubscription ratio is cut in half? What if it is doubled?
 - [25] Reducing the oversubscription ratio can potentially improve the performance if a workload is network-limited. Assume a MapReduce job that uses 120 servers and reads 5 TB of data. Assume the same ratio of read/intermediate/output data as in Figure 6.2, Sep-09, and use Figure 6.6 to define the bandwidths of the memory hierarchy. For data reading, assume that 50% of data is read from remote disks; of that, 80% is read from within the rack and 20% is read from within the array. For intermediate data and output data, assume that 30% of the data uses remote disks; of that, 90% is within the rack and 10% is within the array. What is the overall performance improvement

- when reducing the oversubscription ratio by half? What is the performance if it is doubled? Calculate the TCO in each case.
- c. [Discussion] We are seeing the trend to more cores per system. We are also seeing the increasing adoption of optical communication (with potentially higher bandwidth and improved energy efficiency). How do you think these and other emerging technology trends will affect the design of future WSCs?
- 6.22 [5/15/15/20/25] <6.5> *Realizing the Capability of Amazon Web Services*: Imagine you are the site operation and infrastructure manager of an Alexa.com top site and are considering using Amazon Web Services (AWS). What factors do you need to consider in determining whether to migrate to AWS, what services and instance types to use and how much cost could you save? You can use Alexa and site traffic information (e.g., Wikipedia provides page view stats) to estimate the amount of traffic received by a top site, or you can take concrete examples from the Web, such as the following example from DrupalCon San Francisco 2010: <http://2bits.com/sites/2bits.com/files/drupal-single-server-2.8-million-page-views-a-day.pdf>. The slides describe an Alexa #3400 site that receives 2.8 million page views per day, using a single server. The server has two quad-core Xeon 2.5 GHz processors with 8 GB DRAM and three 15 K RPM SAS hard drives in a RAID1 configuration, and it costs about \$400 per month. The site uses caching heavily, and the CPU utilization ranges from 50% to 250% (roughly 0.5 to 2.5 cores busy).
- a. [5] Looking at the available EC2 instances (<http://aws.amazon.com/ec2/instance-types/>), what instance types match or exceed the current server configuration?
- b. [15] Looking at the EC2 pricing information (<http://aws.amazon.com/ec2/pricing/>), select the most cost-efficient EC2 instances (combinations allowed) to host the site on AWS. What's the monthly cost for EC2?
- c. [15] Now add the costs for IP address and network traffic to the equation, and suppose the site transfers 100 GB/day in and out on the Internet. What's the monthly cost for the site now?
- d. [20] AWS also offers the Micro Instance for free for 1 year to new customers and 15 GB bandwidth each for traffic going in and out across AWS. Based on your estimation of peak and average traffic from your department Web server, can you host it for free on AWS?
- e. [25] A much larger site, Netflix.com, has also migrated their streaming and encoding infrastructure to AWS. Based on their service characteristics, what AWS services could be used by Netflix and for what purposes?
- 6.23 [Discussion/Discussion/20/20/Discussion] <6.4> Figure 6.12 shows the impact of user perceived response time on revenue, and motivates the need to achieve high-throughput while maintaining low latency.
- a. [Discussion] Taking Web search as an example, what are the possible ways of reducing query latency?

- b. [Discussion] What monitoring statistics can you collect to help understand where time is spent? How do you plan to implement such a monitoring tool?
 - c. [20] Assuming that the number of disk accesses per query follows a normal distribution, with an average of 2 and standard deviation of 3, what kind of disk access latency is needed to satisfy a latency SLA of 0.1 sec for 95% of the queries?
 - d. [20] In-memory caching can reduce the frequencies of long-latency events (e.g., accessing hard drives). Assuming a steady-state hit rate of 40%, hit latency of 0.05 sec, and miss latency of 0.2 sec, does caching help meet a latency SLA of 0.1 sec for 95% of the queries?
 - e. [Discussion] When can cached content become stale or even inconsistent? How often can this happen? How can you detect and invalidate such content?
- 6.24 [15/15/20] <6.4> The efficiency of typical power supply units (PSUs) varies as the load changes; for example, PSU efficiency can be about 80% at 40% load (e.g., output 40 watts from a 100-watt PSU), 75% when the load is between 20% and 40%, and 65% when the load is below 20%.
- a. [15] Assume a power-proportional server whose actual power is proportional to CPU utilization, with a utilization curve as shown in Figure 6.3. What is the average PSU efficiency?
 - b. [15] Suppose the server employs $2N$ redundancy for PSUs (i.e., doubles the number of PSUs) to ensure stable power when one PSU fails. What is the average PSU efficiency?
 - c. [20] Blade server vendors use a shared pool of PSUs not only to provide redundancy but also to dynamically match the number of PSUs to the server's actual power consumption. The HP c7000 enclosure uses up to six PSUs for a total of 16 servers. In this case, what is the average PSU efficiency for the enclosure of server with the same utilization curve?
- 6.25 [5/Discussion/10/15/Discussion/Discussion/Discussion] <6.4> *Power stranding* is a term used to refer to power capacity that is provisioned but not used in a data-center. Consider the data presented in Figure 6.25 [Fan, Weber, and Barroso 2007] for different groups of machines. (Note that what this paper calls a “cluster” is what we have referred to as an “array” in this chapter.)
- a. [5] What is the stranded power at (1) the rack level, (2) the power distribution unit level, and (3) the array (cluster) level? What are the trends with oversubscription of power capacity at larger groups of machines?
 - b. [Discussion] What do you think causes the differences between power stranding at different groups of machines?
 - c. [10] Consider an array-level collection of machines where the total machines never use more than 72% of the aggregate power (this is sometimes also referred to as the ratio between the peak-of-sum and sum-of-peaks usage). Using the cost model in the case study, compute the cost savings from comparing a datacenter provisioned for peak capacity and one provisioned for actual use.




Figure 6.25 Cumulative distribution function (CDF) of a real datacenter.

- d. [15] Assume that the datacenter designer chose to include additional servers at the array level to take advantage of the stranded power. Using the example configuration and assumptions in part (a), compute how many more servers can now be included in the warehouse-scale computer for the same total power provisioning.
 - e. [Discussion] What is needed to make the optimization of part (d) work in a real-world deployment? (*Hint*: Think about what needs to happen to cap power in the rare case when all the servers in the array are used at peak power.)
 - f. [Discussion] Two kinds of policies can be envisioned to manage power caps [Ranganathan et al. 2006]: (1) preemptive policies where power budgets are predetermined (“don’t assume you can use more power; ask before you do!”) or (2) reactive policies where power budgets are throttled in the event of a power budget violation (“use as much power as needed until told you can’t!”). Discuss the trade-offs between these approaches and when you would use each type.
 - g. [Discussion] What happens to the total stranded power if systems become more energy proportional (assume workloads similar to that of Figure 6.4)?
- 6.26 [5/20/Discussion] <6.4, 6.7> Section 6.7 discussed the use of per-server battery sources in the Google design. Let us examine the consequences of this design.
- a. [5] Assume that the use of a battery as a mini-server-level UPS is 99.99% efficient and eliminates the need for a facility-wide UPS that is only 92% efficient. Assume that substation switching is 99.7% efficient and that the efficiency for the PDU, step-down stages, and other electrical breakers are 98%, 98%, and 99%, respectively. Calculate the overall power infrastructure efficiency improvements from using a per-server battery backup.

- b. [20] Assume that the UPS is 10% of the cost of the IT equipment. Using the rest of the assumptions from the cost model in the case study, what is the break-even point for the costs of the battery (as a fraction of the cost of a single server) at which the total cost of ownership for a battery-based solution is better than that for a facility-wide UPS?
 - c. [Discussion] What are the other trade-offs between these two approaches? In particular, how do you think the manageability and failure model will change across these two different designs?
- 6.27 [5/5/Discussion] <6.4> For this exercise, consider a simplified equation for the total operational power of a WSC as follows: Total operational power = (1 + Cooling inefficiency multiplier) * IT equipment power.
- a. [5] Assume an 8 MW datacenter at 80% power usage, electricity costs of \$0.10 per kilowatt-hour, and a cooling-inefficiency multiplier of 0.8. Compare the cost savings from (1) an optimization that improves cooling efficiency by 20%, and (2) an optimization that improves the energy efficiency of the IT equipment by 20%.
 - b. [5] What is the percentage improvement in IT equipment energy efficiency needed to match the cost savings from a 20% improvement in cooling efficiency?
 - c. [Discussion/10] What conclusions can you draw about the relative importance of optimizations that focus on server energy efficiency and cooling energy efficiency?
- 6.28 [5/5/Discussion] <6.4> As discussed in this chapter, the cooling equipment in WSCs can themselves consume a lot of energy. Cooling costs can be lowered by proactively managing temperature. Temperature-aware workload placement is one optimization that has been proposed to manage temperature to reduce cooling costs. The idea is to identify the cooling profile of a given room and map the hotter systems to the cooler spots, so that at the WSC level the requirements for overall cooling are reduced.
- a. [5] The coefficient of performance (COP) of a CRAC unit is defined as the ratio of heat removed (Q) to the amount of work necessary (W) to remove that heat. The COP of a CRAC unit increases with the temperature of the air the CRAC unit pushes into the plenum. If air returns to the CRAC unit at 20 degrees Celsius and we remove 10KW of heat with a COP of 1.9, how much energy do we expend in the CRAC unit? If cooling the same volume of air, but now returning at 25 degrees Celsius, takes a COP of 3.1, how much energy do we expend in the CRAC unit now?
 - b. [5] Assume a workload distribution algorithm is able to match the hot workloads well with the cool spots to allow the computer room air-conditioning (CRAC) unit to be run at higher temperature to improve cooling efficiencies like in the exercise above. What is the power savings between the two cases described above?
 - c. [Discussion] Given the scale of WSC systems, power management can be a complex, multifaceted problem. Optimizations to improve energy efficiency

can be implemented in hardware and in software, at the system level, and at the cluster level for the IT equipment or the cooling equipment, etc. It is important to consider these interactions when designing an overall energy-efficiency solution for the WSC. Consider a consolidation algorithm that looks at server utilization and consolidates different workload classes on the same server to increase server utilization (this can potentially have the server operating at higher energy efficiency if the system is not energy proportional). How would this optimization interact with a concurrent algorithm that tried to use different power states (see ACPI, Advanced Configuration Power Interface, for some examples)? What other examples can you think of where multiple optimizations can potentially conflict with one another in a WSC? How would you solve this problem?

- 6.29 [5/10/15/20] <6.2> Energy proportionality (sometimes also referred to as energy scale-down) is the attribute of the system to consume no power when idle, but more importantly gradually consume more power in proportion to the activity level and work done. In this exercise, we will examine the sensitivity of energy consumption to different energy proportionality models. In the exercises below, unless otherwise mentioned, use the data in Figure 6.4 as the default.
- a. [5] A simple way to reason about energy proportionality is to assume linearity between activity and power usage. Using just the peak power and idle power data from Figure 6.4 and a linear interpolation, plot the energy-efficiency trends across varying activities. (Energy efficiency is expressed as performance per watt.) What happens if idle power (at 0% activity) is half of what is assumed in Figure 6.4? What happens if idle power is zero?
 - b. [10] Plot the energy-efficiency trends across varying activities, but use the data from column 3 of Figure 6.4 for power variation. Plot the energy efficiency assuming that the idle power (alone) is half of what is assumed in Figure 6.4. Compare these plots with the linear model in the previous exercise. What conclusions can you draw about the consequences of focusing purely on idle power alone?
 - c. [15] Assume the system utilization mix in column 7 of Figure 6.4. For simplicity, assume a discrete distribution across 1000 servers, with 109 servers at 0% utilization, 80 servers at 10% utilizations, etc. Compute the total performance and total energy for this workload mix using the assumptions in part (a) and part (b).
 - d. [20] One could potentially design a system that has a sublinear power versus load relationship in the region of load levels between 0% and 50%. This would have an energy-efficiency curve that peaks at lower utilizations (at the expense of higher utilizations). Create a new version of column 3 from Figure 6.4 that shows such an energy-efficiency curve. Assume the system utilization mix in column 7 of Figure 6.4. For simplicity, assume a discrete distribution across 1000 servers, with 109 servers at 0% utilization, 80 servers at 10% utilizations, etc. Compute the total performance and total energy for this workload mix.

Activity (%)	0	10	20	30	40	50	60	70	80	90	100
Power, case A (W)	181	308	351	382	416	451	490	533	576	617	662
Power, case B (W)	250	275	325	340	395	405	415	425	440	445	450

Figure 6.26 Power distribution for two servers.

Activity (%)	0	10	20	30	40	50	60	70	80	90	100
No. servers, case A and B	109	80	153	246	191	115	51	21	15	12	8
No. servers, case C	504	6	8	11	26	57	95	123	76	40	54

Figure 6.27 Utilization distributions across cluster, without and with consolidation.

- 6.30 [15/20/20] <6.2, 6.6> This exercise illustrates the interactions of energy proportionality models with optimizations such as server consolidation and energy-efficient server designs. Consider the scenarios shown in Figure 6.26 and Figure 6.27.
- [15] Consider two servers with the power distributions shown in Figure 6.26: case A (the server considered in Figure 6.4) and case B (a less energy-proportional but more energy-efficient server than case A). Assume the system utilization mix in column 7 of Figure 6.4. For simplicity, assume a discrete distribution across 1000 servers, with 109 servers at 0% utilization, 80 servers at 10% utilizations, etc., as shown in row 1 of Figure 6.27. Assume performance variation based on column 2 of Figure 6.4. Compare the total performance and total energy for this workload mix for the two server types.
 - [20] Consider a cluster of 1000 servers with data similar to the data shown in Figure 6.4 (and summarized in the first rows of Figures 6.26 and 6.27). What are the total performance and total energy for the workload mix with these assumptions? Now assume that we were able to consolidate the workloads to model the distribution shown in case C (second row of Figure 6.27). What are the total performance and total energy now? How does the total energy compare with a system that has a linear energy-proportional model with idle power of zero watts and peak power of 662 watts?
 - [20] Repeat part (b), but with the power model of server B, and compare with the results of part (a).
- 6.31 [10/Discussion] <6.2, 6.4, 6.6> *System-Level Energy Proportionality Trends:* Consider the following breakdowns of the power consumption of a server:
- CPU, 50%; memory, 23%; disks, 11%; networking/other, 16%
 CPU, 33%; memory, 30%; disks, 10%; networking/other, 27%
- [10] Assume a dynamic power range of 3.0× for the CPU (i.e., the power consumption of the CPU at idle is one-third that of its power consumption at

Tier 1	Single path for power and cooling distributions, without redundant components	99.0%
Tier 2	$(N + 1)$ redundancy = two power and cooling distribution paths	99.7%
Tier 3	$(N + 2)$ redundancy = three power and cooling distribution paths for uptime even during maintenance	99.98%
Tier 4	Two active power and cooling distribution paths, with redundant components in each path, to tolerate any single equipment failure without impacting the load	99.995%

Figure 6.28 Overview of data center tier classifications. (Adapted from Pitt Turner IV et al. [2008].)

- peak). Assume that the dynamic range of the memory systems, disks, and the networking/other categories above are respectively 2.0 \times , 1.3 \times , and 1.2 \times . What is the overall dynamic range for the total system for the two cases?
- b. [Discussion/10] What can you learn from the results of part (a)? How would we achieve better energy proportionality at the system level? (*Hint*: Energy proportionality at a system level cannot be achieved through CPU optimizations alone, but instead requires improvement across all components.)
- 6.32 [30] <6.4> Pitt Turner IV et al. [2008] presented a good overview of datacenter tier classifications. Tier classifications define site infrastructure performance. For simplicity, consider the key differences as shown in Figure 6.25 (adapted from Pitt Turner IV et al. [2008]). Using the TCO model in the case study, compare the cost implications of the different tiers shown.
- 6.33 [Discussion] <6.4> Based on the observations in Figure 6.13, what can you say qualitatively about the trade-offs between revenue loss from downtime and costs incurred for uptime?
- 6.34 [15/Discussion] <6.4> Some recent studies have defined a metric called TPUE, which stands for “true PUE” or “total PUE.” TPUE is defined as $PUE * SPUE$. PUE, the power utilization effectiveness, is defined in Section 6.4 as the ratio of the total facility power over the total IT equipment power. SPUE, or server PUE, is a new metric analogous to PUE, but instead applied to computing equipment, and is defined as the ratio of total server input power to its useful power, where useful power is defined as the power consumed by the electronic components directly involved in the computation: motherboard, disks, CPUs, DRAM, I/O cards, and so on. In other words, the SPUE metric captures inefficiencies associated with the power supplies, voltage regulators, and fans housed on a server.
- a. [15] <6.4> Consider a design that uses a higher supply temperature for the CRAC units. The efficiency of the CRAC unit is approximately a quadratic function of the temperature, and this design therefore improves the overall PUE, let’s assume by 7%. (Assume baseline PUE of 1.7.) However, the higher temperature at the server level triggers the on-board fan controller to

operate the fan at much higher speeds. The fan power is a cubic function of speed, and the increased fan speed leads to a degradation of SPUE. Assume a fan power model:

$$\text{Fan power} = 284 * ns * ns * ns - 75 * ns * ns,$$

where ns is the normalized fan speed = fan speed in rpm/18,000

and a baseline server power of 350 W. Compute the SPUE if the fan speed increases from (1) 10,000 rpm to 12,500 rpm and (2) 10,000 rpm to 18,000 rpm. Compare the PUE and TPUE in both these cases. (For simplicity, ignore the inefficiencies with power delivery in the SPUE model.)

- b. [Discussion] Part (a) illustrates that, while PUE is an excellent metric to capture the overhead of the facility, it does not capture the inefficiencies within the IT equipment itself. Can you identify another design where TPUE is potentially lower than PUE? (*Hint*: See Exercise 6.26.)
- 6.35 [Discussion/30/Discussion] <6.2> Two recently released benchmarks provide a good starting point for energy-efficiency accounting in servers—the SPECpower_ssj2008 benchmark (available at http://www.spec.org/power_ssj2008/) and the JouleSort metric (available at <http://sortbenchmark.org/>).
- a. [Discussion] <6.2> Look up the descriptions of the two benchmarks. How are they similar? How are they different? What would you do to improve these benchmarks to better address the goal of improving WSC energy efficiency?
 - b. [30] <6.2> JouleSort measures the total system energy to perform an out-of-core sort and attempts to derive a metric that enables the comparison of systems ranging from embedded devices to supercomputers. Look up the description of the JouleSort metric at <http://sortbenchmark.org>. Download a publicly available version of the sort algorithm and run it on different classes of machines—a laptop, a PC, a mobile phone, etc.—or with different configurations. What can you learn from the JouleSort ratings for different setups?
 - c. [Discussion] <6.2> Consider the system with the best JouleSort rating from your experiments above. How would you improve the energy efficiency? For example, try rewriting the sort code to improve the JouleSort rating.
- 6.36 [10/10/15] <6.1, 6.2> Figure 6.1 is a listing of outages in an array of servers. When dealing with the large scale of WSCs, it is important to balance cluster design and software architectures to achieve the required uptime without incurring significant costs. This question explores the implications of achieving availability through hardware only.
- a. [10] <6.1, 6.2> Assuming that an operator wishes to achieve 95% availability through server hardware improvements alone, how many events of each type would have to be reduced? For now, assume that individual server crashes are completely handled through redundant machines.
 - b. [10] <6.1, 6.2> How does the answer to part (a) change if the individual server crashes are handled by redundancy 50% of the time? 20% of the time? None of the time?

- c. [15] <6.1, 6.2> Discuss the importance of software redundancy to achieving a high level of availability. If a WSC operator considered buying machines that were cheaper, but 10% less reliable, what implications would that have on the software architecture? What are the challenges associated with software redundancy?
- 6.37 [15] <6.1, 6.8> Look up the current prices of standard DDR3 DRAM versus DDR3 DRAM that has error-correcting code (ECC). What is the increase in price per bit for achieving the higher reliability that ECC provides? Using the DRAM prices alone, and the data provided in Section 6.8, what is the uptime per dollar of a WSC with non-ECC versus ECC DRAM?
- 6.38 [5/Discussion] <6.1> *WSC Reliability and Manageability Concerns:*
- a. [5] Consider a cluster of servers costing \$2000 each. Assuming an annual failure rate of 5%, an average of an hour of service time per repair, and replacement parts requiring 10% of the system cost per failure, what is the annual maintenance cost per server? Assume an hourly rate of \$100 per hour for a service technician.
 - b. [Discussion] Comment on the differences between this manageability model versus that in a traditional enterprise datacenter with a large number of small or medium-sized applications each running on its own dedicated hardware infrastructure.

References

- Adve, S. V., and K. Gharachorloo [1996]. "Shared memory consistency models: A tutorial," *IEEE Computer* 29:12 (December), 66–76.
- Adve, S. V., and M. D. Hill [1990]. "Weak ordering—a new definition," *Proc. 17th Annual Int'l. Symposium on Computer Architecture (ISCA)*, May 28–31, 1990, Seattle, Wash., 2–14.
- Agarwal, A. [1987]. "Analysis of Cache Performance for Operating Systems and Multiprogramming," Ph.D. thesis, Tech. Rep. No. CSL-TR-87-332, Stanford University, Palo Alto, Calif.
- Agarwal, A. [1991]. "Limits on interconnection network performance," *IEEE Trans. on Parallel and Distributed Systems* 2:4 (April), 398–412.
- Agarwal, A., and S. D. Pudar [1993]. "Column-associative caches: A technique for reducing the miss rate of direct-mapped caches," *20th Annual Int'l. Symposium on Computer Architecture (ISCA)*, May 16–19, 1993, San Diego, Calif. Also appears in *Computer Architecture News* 21:2 (May), 179–190, 1993.
- Agarwal, A., R. Bianchini, D. Chaiken, K. Johnson, and D. Kranz [1995]. "The MIT Alewife machine: Architecture and performance," *Int'l. Symposium on Computer Architecture* (Denver, Colo.), June, 2–13.
- Agarwal, A., J. L. Hennessy, R. Simoni, and M. A. Horowitz [1988]. "An evaluation of directory schemes for cache coherence," *Proc. 15th Int'l. Symposium on Computer Architecture* (June), 280–289.
- Agarwal, A., J. Kubiawicz, D. Kranz, B.-H. Lim, D. Yeung, G. D'Souza, and M. Parkin [1993]. "Sparcle: An evolutionary processor design for large-scale multiprocessors," *IEEE Micro* 13 (June), 48–61.
- Agerwala, T., and J. Cocke [1987]. *High Performance Reduced Instruction Set Processors*, IBM Tech. Rep. RC12434, IBM, Armonk, N.Y.
- Akeley, K. and T. Jermoluk [1988]. "High-Performance Polygon Rendering," *Proc. 15th Annual Conf. on Computer Graphics and Interactive Techniques (SIGGRAPH 1988)*, August 1–5, 1988, Atlanta, Ga., 239–246.
- Alexander, W. G., and D. B. Wortman [1975]. "Static and dynamic characteristics of XPL programs," *IEEE Computer* 8:11 (November), 41–46.
- Alles, A. [1995]. "ATM Internetworking," White Paper (May), Cisco Systems, Inc., San Jose, Calif. (www.cisco.com/warp/public/614/12.html).
- Alliant. [1987]. *Alliant FX/Series: Product Summary*, Alliant Computer Systems Corp., Acton, Mass.
- Almasi, G. S., and A. Gottlieb [1989]. *Highly Parallel Computing*, Benjamin/Cummings, Redwood City, Calif.
- Alverson, G., R. Alverson, D. Callahan, B. Koblenz, A. Porterfield, and B. Smith [1992]. "Exploiting heterogeneous parallelism on a multithreaded multiprocessor," *Proc. ACM/IEEE Conf. on Supercomputing*, November 16–20, 1992, Minneapolis, Minn., 188–197.

- Amdahl, G. M. [1967]. "Validity of the single processor approach to achieving large scale computing capabilities," *Proc. AFIPS Spring Joint Computer Conf.*, April 18–20, 1967, Atlantic City, N.J., 483–485.
- Amdahl, G. M., G. A. Blaauw, and F. P. Brooks, Jr. [1964]. "Architecture of the IBM System 360," *IBM J. Research and Development* 8:2 (April), 87–101.
- Amza, C., A. L. Cox, S. Dwarkadas, P. Keleher, H. Lu, R. Rajamony, W. Yu, and W. Zwaenepoel [1996]. "Treadmarks: Shared memory computing on networks of workstations," *IEEE Computer* 29:2 (February), 18–28.
- Anderson, D. [2003]. "You don't know jack about disks," *Queue*, 1:4 (June), 20–30.
- Anderson, D., J. Dykes, and E. Riedel [2003]. "SCSI vs. ATA—More than an interface," *Proc. 2nd USENIX Conf. on File and Storage Technology (FAST '03)*, March 31–April 2, 2003, San Francisco.
- Anderson, D. W., F. J. Sparacio, and R. M. Tomasulo [1967]. "The IBM 360 Model 91: Processor philosophy and instruction handling," *IBM J. Research and Development* 11:1 (January), 8–24.
- Anderson, M. H. [1990]. "Strength (and safety) in numbers (RAID, disk storage technology)," *Byte* 15:13 (December), 337–339.
- Anderson, T. E., D. E. Culler, and D. Patterson [1995]. "A case for NOW (networks of workstations)," *IEEE Micro* 15:1 (February), 54–64.
- Ang, B., D. Chiou, D. Rosenband, M. Ehrlich, L. Rudolph, and Arvind [1998]. "StarT-Voyager: A flexible platform for exploring scalable SMP issues," *Proc. ACM/IEEE Conf. on Supercomputing*, November 7–13, 1998, Orlando, FL.
- Anjan, K. V., and T. M. Pinkston [1995]. "An efficient, fully-adaptive deadlock recovery scheme: Disha," *Proc. 22nd Annual Int'l. Symposium on Computer Architecture (ISCA)*, June 22–24, 1995, Santa Margherita, Italy.
- Anon. et al. [1985]. *A Measure of Transaction Processing Power*, Tandem Tech. Rep. TR85.2. Also appears in *Datamation* 31:7 (April), 112–118, 1985.
- Apache Hadoop. [2011]. <http://hadoop.apache.org>.
- Archibald, J., and J.-L. Baer [1986]. "Cache coherence protocols: Evaluation using a multiprocessor simulation model," *ACM Trans. on Computer Systems* 4:4 (November), 273–298.
- Armburst, M., A. Fox, R. Griffith, A. D. Joseph, R. Katz, A. Konwinski, G. Lee, D. Patterson, A. Rabkin, I. Stoica, and M. Zaharia [2009]. *Above the Clouds: A Berkeley View of Cloud Computing*, Tech. Rep. UCB/EECS-2009-28, University of California, Berkeley (<http://www.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-28.html>).
- Arpaci, R. H., D. E. Culler, A. Krishnamurthy, S. G. Steinberg, and K. Yelick [1995]. "Empirical evaluation of the CRAY-T3D: A compiler perspective," *22nd Annual Int'l. Symposium on Computer Architecture (ISCA)*, June 22–24, 1995, Santa Margherita, Italy.
- Asanovic, K. [1998]. "Vector Microprocessors," Ph.D. thesis, Computer Science Division, University of California, Berkeley.
- Associated Press. [2005]. "Gap Inc. shuts down two Internet stores for major overhaul," *USATODAY.com*, August 8, 2005.
- Atanasoff, J. V. [1940]. *Computing Machine for the Solution of Large Systems of Linear Equations*, Internal Report, Iowa State University, Ames.
- Atkins, M. [1991]. Performance and the i860 Microprocessor, *IEEE Micro*, 11:5 (September), 24–27, 72–78.
- Austin, T. M., and G. Sohi [1992]. "Dynamic dependency analysis of ordinary programs," *Proc. 19th Annual Int'l. Symposium on Computer Architecture (ISCA)*, May 19–21, 1992, Gold Coast, Australia, 342–351.

- Babbay, F., and A. Mendelson [1998]. "Using value prediction to increase the power of speculative execution hardware," *ACM Trans. on Computer Systems* 16:3 (August), 234–270.
- Baer, J.-L., and W.-H. Wang [1988]. "On the inclusion property for multi-level cache hierarchies," *Proc. 15th Annual Int'l. Symposium on Computer Architecture*, May 30–June 2, 1988, Honolulu, Hawaii, 73–80.
- Bailey, D. H., E. Barszcz, J. T. Barton, D. S. Browning, R. L. Carter, L. Dagum, R. A. Fatoohi, P. O. Frederickson, T. A. Lasinski, R. S. Schreiber, H. D. Simon, V. Venkatakrishnan, and S. K. Weeratunga [1991]. "The NAS parallel benchmarks," *Int'l. J. Supercomputing Applications* 5, 63–73.
- Bakoglu, H. B., G. F. Grohoski, L. E. Thatcher, J. A. Kaeli, C. R. Moore, D. P. Tattle, W. E. Male, W. R. Hardell, D. A. Hicks, M. Nguyen Phu, R. K. Montoye, W. T. Glover, and S. Dhawan [1989]. "IBM second-generation RISC processor organization," *Proc. IEEE Int'l. Conf. on Computer Design*, September 30–October 4, 1989, Rye, N.Y., 138–142.
- Balakrishnan, H., V. N. Padmanabhan, S. Seshan, and R. H. Katz [1997]. "A comparison of mechanisms for improving TCP performance over wireless links," *IEEE/ACM Trans. on Networking* 5:6 (December), 756–769.
- Ball, T., and J. Larus [1993]. "Branch prediction for free," *Proc. ACM SIGPLAN'93 Conference on Programming Language Design and Implementation (PLDI)*, June 23–25, 1993, Albuquerque, N.M., 300–313.
- Banerjee, U. [1979]. "Speedup of Ordinary Programs," Ph.D. thesis, Dept. of Computer Science, University of Illinois at Urbana-Champaign.
- Barham, P., B. Dragovic, K. Fraser, S. Hand, T. Harris, A. Ho, and R. Neugebauer [2003]. "Xen and the art of virtualization," *Proc. of the 19th ACM Symposium on Operating Systems Principles*, October 19–22, 2003, Bolton Landing, N.Y.
- Barroso, L. A. [2010]. "Warehouse Scale Computing [keynote address]," *Proc. ACM SIGMOD*, June 8–10, 2010, Indianapolis, Ind.
- Barroso, L. A., and U. Hölzle [2007]. "The case for energy-proportional computing," *IEEE Computer*, 40:12 (December), 33–37.
- Barroso, L. A., and U. Hölzle [2009]. *The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale Machines*, Morgan & Claypool, San Rafael, Calif.
- Barroso, L. A., K. Gharachorloo, and E. Bugnion [1998]. "Memory system characterization of commercial workloads," *Proc. 25th Annual Int'l. Symposium on Computer Architecture (ISCA)*, July 3–14, 1998, Barcelona, Spain, 3–14.
- Barton, R. S. [1961]. "A new approach to the functional design of a computer," *Proc. Western Joint Computer Conf.*, May 9–11, 1961, Los Angeles, Calif., 393–396.
- Bashe, C. J., W. Buchholz, G. V. Hawkins, J. L. Ingram, and N. Rochester [1981]. "The architecture of IBM's early computers," *IBM J. Research and Development* 25:5 (September), 363–375.
- Bashe, C. J., L. R. Johnson, J. H. Palmer, and E. W. Pugh [1986]. *IBM's Early Computers*, MIT Press, Cambridge, Mass.
- Baskett, F., and T. W. Keller [1977]. "An evaluation of the Cray-1 processor," in *High Speed Computer and Algorithm Organization*, D. J. Kuck, D. H. Lawrie, and A. H. Sameh, eds., Academic Press, San Diego, 71–84.
- Baskett, F., T. Jermoluk, and D. Solomon [1988]. "The 4D-MP graphics superworkstation: Computing + graphics = 40 MIPS + 40 MFLOPS and 10,000 lighted polygons per second," *Proc. IEEE COMPCON*, February 29–March 4, 1988, San Francisco, 468–471.
- BBN Laboratories. [1986]. *Butterfly Parallel Processor Overview*. Tech. Rep. 6148, BBN Laboratories, Cambridge, Mass.
- Bell, C. G. [1984]. "The mini and micro industries," *IEEE Computer* 17:10 (October), 14–30.

- Bell, C. G. [1985]. "Multis: A new class of multiprocessor computers," *Science* 228 (April 26), 462–467.
- Bell, C. G. [1989]. "The future of high performance computers in science and engineering," *Communications of the ACM* 32:9 (September), 1091–1101.
- Bell, G., and J. Gray [2001]. *Crays, Clusters and Centers*, Tech. Rep. MSR-TR-2001-76, Microsoft Research, Redmond, Wash.
- Bell, C. G., and J. Gray [2002]. "What's next in high performance computing?" *CACM* 45:2 (February), 91–95.
- Bell, C. G., and A. Newell [1971]. *Computer Structures: Readings and Examples*, McGraw-Hill, New York.
- Bell, C. G., and W. D. Strecker [1976]. "Computer structures: What have we learned from the PDP-11?," *Third Annual Int'l. Symposium on Computer Architecture (ISCA)*, January 19–21, 1976, Tampa, Fla., 1–14.
- Bell, C. G., and W. D. Strecker [1998]. "Computer structures: What have we learned from the PDP-11?" *25 Years of the International Symposia on Computer Architecture (Selected Papers)*, ACM, New York, 138–151.
- Bell, C. G., J. C. Mudge, and J. E. McNamara [1978]. *A DEC View of Computer Engineering*, Digital Press, Bedford, Mass.
- Bell, C. G., R. Cady, H. McFarland, B. DeLagi, J. O'Laughlin, R. Noonan, and W. Wulf [1970]. "A new architecture for mini-computers: The DEC PDP-11," *Proc. AFIPS Spring Joint Computer Conf.*, May 5–May 7, 1970, Atlantic City, N.J., 657–675.
- Benes, V. E. [1962]. "Rearrangeable three stage connecting networks," *Bell System Technical Journal* 41, 1481–1492.
- Bertozi, D., A. Jalabert, S. Murali, R. Tamhankar, S. Stergiou, L. Benini, and G. De Micheli [2005]. "NoC synthesis flow for customized domain specific multiprocessor systems-on-chip," *IEEE Trans. on Parallel and Distributed Systems* 16:2 (February), 113–130.
- Bhandarkar, D. P. [1995]. *Alpha Architecture and Implementations*, Digital Press, Newton, Mass.
- Bhandarkar, D. P., and D. W. Clark [1991]. "Performance from architecture: Comparing a RISC and a CISC with similar hardware organizations," *Proc. Fourth Int'l. Conf. on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, April 8–11, 1991, Palo Alto, Calif., 310–319.
- Bhandarkar, D. P., and J. Ding [1997]. "Performance characterization of the Pentium Pro processor," *Proc. Third Int'l. Symposium on High-Performance Computer Architecture*, February 1–February 5, 1997, San Antonio, Tex., 288–297.
- Bhuyan, L. N., and D. P. Agrawal [1984]. "Generalized hypercube and hyperbus structures for a computer network," *IEEE Trans. on Computers* 32:4 (April), 322–333.
- Bienia, C., S. Kumar, P. S. Jaswinder, and K. Li [2008]. *The Parsec Benchmark Suite: Characterization and Architectural Implications*, Tech. Rep. TR-811-08, Princeton University, Princeton, N.J.
- Bier, J. [1997]. "The Evolution of DSP Processors," presentation at University of California, Berkeley, November 14.
- Bird, S., A. Phansalkar, L. K. John, A. Mericas, and R. Indukuru [2007]. "Characterization of performance of SPEC CPU benchmarks on Intel's Core Microarchitecture based processor," *Proc. 2007 SPEC Benchmark Workshop*, January 21, 2007, Austin, Tex.
- Birman, M., A. Samuels, G. Chu, T. Chuk, L. Hu, J. McLeod, and J. Barnes [1990]. "Developing the WRL3170/3171 SPARC floating-point coprocessors," *IEEE Micro* 10:1, 55–64.
- Blackburn, M., R. Gamer, C. Hoffman, A. M. Khan, K. S. McKinley, R. Bentzur, A. Diwan, D. Feinberg, D. Frampton, S. Z. Guyer, M. Hirzel, A. Hosking, M. Jump, H. Lee, J. E. B. Moss, A. Phansalkar, D. Stefanovic, T. VanDrunen, D. von Dincklage,

- and B. Wiedermann [2006]. "The DaCapo benchmarks: Java benchmarking development and analysis," *ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA)*, October 22–26, 2006, 169–190.
- Blaum, M., J. Bruck, and A. Vardy [1996]. "MDS array codes with independent parity symbols," *IEEE Trans. on Information Theory*, IT-42 (March), 529–42.
- Blaum, M., J. Brady, J. Bruck, and J. Menon [1994]. "EVENODD: An optimal scheme for tolerating double disk failures in RAID architectures," *Proc. 21st Annual Int'l. Symposium on Computer Architecture (ISCA)*, April 18–21, 1994, Chicago, 245–254.
- Blaum, M., J. Brady, J. Bruck, and J. Menon [1995]. "EVENODD: An optimal scheme for tolerating double disk failures in RAID architectures," *IEEE Trans. on Computers* 44:2 (February), 192–202.
- Blaum, M., J. Brady, J., Bruck, J. Menon, and A. Vardy [2001]. "The EVENODD code and its generalization," in H. Jin, T. Cortes, and R. Buyya, eds., *High Performance Mass Storage and Parallel I/O: Technologies and Applications*, Wiley-IEEE, New York, 187–208.
- Bloch, E. [1959]. "The engineering design of the Stretch computer," *1959 Proceedings of the Eastern Joint Computer Conf.*, December 1–3, 1959, Boston, Mass., 48–59.
- Boddie, J. R. [2000]. "History of DSPs," www.lucent.com/micro/dsp/dsphist.html.
- Bolt, K. M. [2005]. "Amazon sees sales rise, profit fall," *Seattle Post-Intelligencer*, October 25 (http://seattlepi.nwsource.com/business/245943_techearn26.html).
- Bordawekar, R., U. Bondhugula, R. Rao [2010]. "Believe It or Not!: Multi-core CPUs can Match GPU Performance for a FLOP-Intensive Application!" *19th International Conference on Parallel Architecture and Compilation Techniques (PACT 2010)*. Vienna, Austria, September 11–15, 2010, 537–538.
- Borg, A., R. E. Kessler, and D. W. Wall [1990]. "Generation and analysis of very long address traces," *19th Annual Int'l. Symposium on Computer Architecture (ISCA)*, May 19–21, 1992, Gold Coast, Australia, 270–279.
- Bouknight, W. J., S. A. Deneberg, D. E. McIntyre, J. M. Randall, A. H. Sameh, and D. L. Slotnick [1972]. "The Illiac IV system," *Proc. IEEE* 60:4, 369–379. Also appears in D. P. Siewiorek, C. G. Bell, and A. Newell, *Computer Structures: Principles and Examples*, McGraw-Hill, New York, 1982, 306–316.
- Brady, J. T. [1986]. "A theory of productivity in the creative process," *IEEE CG&A* (May), 25–34.
- Brain, M. [2000]. "Inside a Digital Cell Phone," www.howstuffworks.com/inside-cellphone.htm.
- Brandt, M., J. Brooks, M. Cahir, T. Hewitt, E. Lopez-Pineda, and D. Sandness [2000]. *The Benchmarkers' Guide for Cray SVI Systems*. Cray Inc., Seattle, Wash.
- Brent, R. P., and H. T. Kung [1982]. "A regular layout for parallel adders," *IEEE Trans. on Computers* C-31, 260–264.
- Brewer, E. A., and B. C. Kuszmaul [1994]. "How to get good performance from the CM-5 data network," *Proc. Eighth Int'l. Parallel Processing Symposium*, April 26–27, 1994, Cancun, Mexico.
- Brin, S., and L. Page [1998]. "The anatomy of a large-scale hypertextual Web search engine," *Proc. 7th Int'l. World Wide Web Conf.*, April 14–18, 1998, Brisbane, Queensland, Australia, 107–117.
- Brown, A., and D. A. Patterson [2000]. "Towards maintainability, availability, and growth benchmarks: A case study of software RAID systems." *Proc. 2000 USENIX Annual Technical Conf.*, June 18–23, 2000, San Diego, Calif.
- Bucher, I. V., and A. H. Hayes [1980]. "I/O performance measurement on Cray-1 and CDC 7000 computers," *Proc. Computer Performance Evaluation Users Group, 16th Meeting*, NBS 500-65, 245–254.

- Bucher, I. Y. [1983]. "The computational speed of supercomputers," *Proc. Int'l. Conf. on Measuring and Modeling of Computer Systems (SIGMETRICS 1983)*, August 29–31, 1983, Minneapolis, Minn., 151–165.
- Bucholtz, W. [1962]. *Planning a Computer System: Project Stretch*, McGraw-Hill, New York.
- Burgess, N., and T. Williams [1995]. "Choices of operand truncation in the SRT division algorithm," *IEEE Trans. on Computers* 44:7, 933–938.
- Burkhardt III, H., S. Frank, B. Knobe, and J. Rothnie [1992]. *Overview of the KSRI Computer System*, Tech. Rep. KSR-TR-9202001, Kendall Square Research, Boston, Mass.
- Burks, A. W., H. H. Goldstine, and J. von Neumann [1946]. "Preliminary discussion of the logical design of an electronic computing instrument," Report to the U.S. Army Ordnance Department, p. 1; also appears in *Papers of John von Neumann*, W. Aspray and A. Burks, eds., MIT Press, Cambridge, Mass., and Tomash Publishers, Los Angeles, Calif., 1987, 97–146.
- Calder, B., G. Reinman, and D. M. Tullsen [1999]. "Selective value prediction," *Proc. 26th Annual Int'l. Symposium on Computer Architecture (ISCA)*, May 2–4, 1999, Atlanta, Ga.
- Calder, B., D. Grunwald, M. Jones, D. Lindsay, J. Martin, M. Mozer, and B. Zorn [1997]. "Evidence-based static branch prediction using machine learning," *ACM Trans. Program. Lang. Syst.* 19:1, 188–222.
- Callahan, D., J. Dongarra, and D. Levine [1988]. "Vectorizing compilers: A test suite and results," *Proc. ACM/IEEE Conf. on Supercomputing*, November 12–17, 1988, Orlando, Fla., 98–105.
- Cantin, J. F., and M. D. Hill [2001]. "Cache Performance for Selected SPEC CPU2000 Benchmarks," www.jfred.org/cache-data.html (June).
- Cantin, J. F., and M. D. Hill [2003]. "Cache Performance for SPEC CPU2000 Benchmarks, Version 3.0," www.cs.wisc.edu/multifacet/misc/spec2000cache-data/index.html.
- Carles, S. [2005]. "Amazon reports record Xmas season, top game picks," *Gamasutra*, December 27 (http://www.gamasutra.com/php-bin/news_index.php?story=7630.)
- Carter, J., and K. Rajamani [2010]. "Designing energy-efficient servers and data centers," *IEEE Computer* 43:7 (July), 76–78.
- Case, R. P., and A. Padeogs [1978]. "The architecture of the IBM System/370," *Communications of the ACM* 21:1, 73–96. Also appears in D. P. Siewiorek, C. G. Bell, and A. Newell, *Computer Structures: Principles and Examples*, McGraw-Hill, New York, 1982, 830–855.
- Censier, L., and P. Feautrier [1978]. "A new solution to coherence problems in multicache systems," *IEEE Trans. on Computers* C-27:12 (December), 1112–1118.
- Chandra, R., S. Devine, B. Verghese, A. Gupta, and M. Rosenblum [1994]. "Scheduling and page migration for multiprocessor compute servers," *Sixth Int'l. Conf. on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, October 4–7, 1994, San Jose, Calif., 12–24.
- Chang, F., J. Dean, S. Ghemawat, W. C. Hsieh, D. A. Wallach, M. Burrows, T. Chandra, A. Fikes, and R. E. Gruber [2006]. "Bigtable: A distributed storage system for structured data," *Proc. 7th USENIX Symposium on Operating Systems Design and Implementation (OSDI '06)*, November 6–8, 2006, Seattle, Wash.
- Chang, J., J. Meza, P. Ranganathan, C. Bash, and A. Shah [2010]. "Green server design: Beyond operational energy to sustainability," *Proc. Workshop on Power Aware Computing and Systems (HotPower '10)*, October 3, 2010, Vancouver, British Columbia.
- Chang, P. P., S. A. Mahlke, W. Y. Chen, N. J. Warter, and W. W. Hwu [1991]. "IMPACT: An architectural framework for multiple-instruction-issue processors," *18th Annual Int'l. Symposium on Computer Architecture (ISCA)*, May 27–30, 1991, Toronto, Canada, 266–275.

- Charlesworth, A. E. [1981]. "An approach to scientific array processing: The architecture design of the AP-120B/FPS-164 family," *Computer* 14:9 (September), 18–27.
- Charlesworth, A. [1998]. "Starfire: Extending the SMP envelope," *IEEE Micro* 18:1 (January/February), 39–49.
- Chen, P. M., and E. K. Lee [1995]. "Striping in a RAID level 5 disk array," *Proc. ACM SIGMETRICS Conf. on Measurement and Modeling of Computer Systems*, May 15–19, 1995, Ottawa, Canada, 136–145.
- Chen, P. M., G. A. Gibson, R. H. Katz, and D. A. Patterson [1990]. "An evaluation of redundant arrays of inexpensive disks using an Amdahl 5890," *Proc. ACM SIGMETRICS Conf. on Measurement and Modeling of Computer Systems*, May 22–25, 1990, Boulder, Colo.
- Chen, P. M., E. K. Lee, G. A. Gibson, R. H. Katz, and D. A. Patterson [1994]. "RAID: High-performance, reliable secondary storage," *ACM Computing Surveys* 26:2 (June), 145–188.
- Chen, S. [1983]. "Large-scale and high-speed multiprocessor system for scientific applications," *Proc. NATO Advanced Research Workshop on High-Speed Computing*, June 20–22, 1983, Jülich, West Germany. Also appears in K. Hwang, ed., "Superprocessors: Design and applications," *IEEE* (August), 602–609, 1984.
- Chen, T. C. [1980]. "Overlap and parallel processing," in H. Stone, ed., *Introduction to Computer Architecture*, Science Research Associates, Chicago, 427–486.
- Chow, F. C. [1983]. "A Portable Machine-Independent Global Optimizer—Design and Measurements," Ph.D. thesis, Stanford University, Palo Alto, Calif.
- Chrysos, G. Z., and J. S. Emer [1998]. "Memory dependence prediction using store sets," *Proc. 25th Annual Int'l. Symposium on Computer Architecture (ISCA)*, July 3–14, 1998, Barcelona, Spain, 142–153.
- Clark, B., T. Deshane, E. Dow, S. Evanchik, M. Finlayson, J. Herne, and J. Neefe Matthews [2004]. "Xen and the art of repeated research," *Proc. USENIX Annual Technical Conf.*, June 27–July 2, 2004, 135–144.
- Clark, D. W. [1983]. "Cache performance of the VAX-11/780," *ACM Trans. on Computer Systems* 1:1, 24–37.
- Clark, D. W. [1987]. "Pipelining and performance in the VAX 8800 processor," *Proc. Second Int'l. Conf. on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, October 5–8, 1987, Palo Alto, Calif., 173–177.
- Clark, D. W., and J. S. Emer [1985]. "Performance of the VAX-11/780 translation buffer: Simulation and measurement," *ACM Trans. on Computer Systems* 3:1 (February), 31–62.
- Clark, D., and H. Levy [1982]. "Measurement and analysis of instruction set use in the VAX-11/780," *Proc. Ninth Annual Int'l. Symposium on Computer Architecture (ISCA)*, April 26–29, 1982, Austin, Tex., 9–17.
- Clark, D., and W. D. Strecker [1980]. "Comments on 'the case for the reduced instruction set computer,'" *Computer Architecture News* 8:6 (October), 34–38.
- Clark, W. A. [1957]. "The Lincoln TX-2 computer development," *Proc. Western Joint Computer Conference*, February 26–28, 1957, Los Angeles, 143–145.
- Clidaras, J., C. Johnson, and B. Felderman [2010]. Private communication.
- Climate Savers Computing Initiative. [2007]. "Efficiency Specs," <http://www.climatesaverscomputing.org/>.
- Clos, C. [1953]. "A study of non-blocking switching networks," *Bell Systems Technical Journal* 32 (March), 406–424.
- Cody, W. J., J. T. Coonen, D. M. Gay, K. Hanson, D. Hough, W. Kahan, R. Karpinski, J. Palmer, F. N. Ris, and D. Stevenson [1984]. "A proposed radix- and word-length-independent standard for floating-point arithmetic," *IEEE Micro* 4:4, 86–100.

- Colwell, R. P., and R. Steck [1995]. "A 0.6 μm BiCMOS processor with dynamic execution." *Proc. of IEEE Int'l. Symposium on Solid State Circuits (ISSCC)*, February 15–17, 1995, San Francisco, 176–177.
- Colwell, R. P., R. P. Nix, J. J. O'Donnell, D. B. Papworth, and P. K. Rodman [1987]. "A VLIW architecture for a trace scheduling compiler," *Proc. Second Int'l. Conf. on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, October 5–8, 1987, Palo Alto, Calif., 180–192.
- Comer, D. [1993]. *Internetworking with TCP/IP*, 2nd ed., Prentice Hall, Englewood Cliffs, N.J.
- Compaq Computer Corporation. [1999]. *Compiler Writer's Guide for the Alpha 21264*, Order Number EC-RJ66A-TE, June, www1.support.compaq.com/alpha-tools/documentation/current/21264_EV67/ec-rj66a-te_comp_writ_gde_for_alpha21264.pdf.
- Conti, C., D. H. Gibson, and S. H. Pitkowsky [1968]. "Structural aspects of the System/360 Model 85. Part I. General organization," *IBM Systems J.* 7:1, 2–14.
- Coonen, J. [1984]. "Contributions to a Proposed Standard for Binary Floating-Point Arithmetic," Ph.D. thesis, University of California, Berkeley.
- Corbett, P., B. English, A. Goel, T. Gracanac, S. Kleiman, J. Leong, and S. Sankar [2004]. "Row-diagonal parity for double disk failure correction," *Proc. 3rd USENIX Conf. on File and Storage Technology (FAST '04)*, March 31–April 2, 2004, San Francisco.
- Crawford, J., and P. Gelsinger [1988]. *Programming the 80386*, Sybex Books, Alameda, Calif.
- Culler, D. E., J. P. Singh, and A. Gupta [1999]. *Parallel Computer Architecture: A Hardware/Software Approach*, Morgan Kaufmann, San Francisco.
- Curnow, H. J., and B. A. Wichmann [1976]. "A synthetic benchmark," *The Computer J.* 19:1, 43–49.
- Cvetanovic, Z., and R. E. Kessler [2000]. "Performance analysis of the Alpha 21264-based Compaq ES40 system," *Proc. 27th Annual Int'l. Symposium on Computer Architecture (ISCA)*, June 10–14, 2000, Vancouver, Canada, 192–202.
- Dally, W. J. [1990]. "Performance analysis of k -ary n -cube interconnection networks," *IEEE Trans. on Computers* 39:6 (June), 775–785.
- Dally, W. J. [1992]. "Virtual channel flow control," *IEEE Trans. on Parallel and Distributed Systems* 3:2 (March), 194–205.
- Dally, W. J. [1999]. "Interconnect limited VLSI architecture," *Proc. of the International Interconnect Technology Conference*, May 24–26, 1999, San Francisco.
- Dally, W. J., and C. I. Seitz [1986]. "The torus routing chip," *Distributed Computing* 1:4, 187–196.
- Dally, W. J., and B. Towles [2001]. "Route packets, not wires: On-chip interconnection networks," *Proc. 38th Design Automation Conference*, June 18–22, 2001, Las Vegas.
- Dally, W. J., and B. Towles [2003]. *Principles and Practices of Interconnection Networks*, Morgan Kaufmann, San Francisco.
- Darcy, J. D., and D. Gay [1996]. "FLECKmarks: Measuring floating point performance using a full IEEE compliant arithmetic benchmark," CS 252 class project, University of California, Berkeley (see [HTTP://CS.Berkeley.EDU/~darcy/Projects/cs252/](http://CS.Berkeley.EDU/~darcy/Projects/cs252/)).
- Darley, H. M. et al. [1989]. "Floating Point/Integer Processor with Divide and Square Root Functions," U.S. Patent 4,878,190, October 31.
- Davidson, E. S. [1971]. "The design and control of pipelined function generators," *Proc. IEEE Conf. on Systems, Networks, and Computers*, January 19–21, 1971, Oaxtepec, Mexico, 19–21.
- Davidson, E. S., A. T. Thomas, L. E. Shar, and J. H. Patel [1975]. "Effective control for pipelined processors," *Proc. IEEE COMPCON*, February 25–27, 1975, San Francisco, 181–184.

- Davie, B. S., L. L. Peterson, and D. Clark [1999]. *Computer Networks: A Systems Approach*, 2nd ed., Morgan Kaufmann, San Francisco.
- Dean, J. [2009]. "Designs, lessons and advice from building large distributed systems [keynote address]." *Proc. 3rd ACM SIGOPS Int'l. Workshop on Large-Scale Distributed Systems and Middleware. Co-located with the 22nd ACM Symposium on Operating Systems Principles*, October 11–14, 2009, Big Sky, Mont.
- Dean, J., and S. Ghemawat [2004]. "MapReduce: Simplified data processing on large clusters." In *Proc. Operating Systems Design and Implementation (OSDI)*, December 6–8, 2004, San Francisco, Calif., 137–150.
- Dean, J., and S. Ghemawat [2008]. "MapReduce: Simplified data processing on large clusters," *Communications of the ACM*, 51:1, 107–113.
- DeCandia, G., D. Hastorun, M. Jampani, G. Kakulapati, A. Lakshman, A. Pilchin, S. Sivasubramanian, P. Vosshall, and W. Vogels [2007]. "Dynamo: Amazon's highly available key-value store," *Proc. 21st ACM Symposium on Operating Systems Principles*, October 14–17, 2007, Stevenson, Wash.
- Dehnert, J. C., P. Y.-T. Hsu, and J. P. Bratt [1989]. "Overlapped loop support on the Cydra 5," *Proc. Third Int'l. Conf. on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, April 3–6, 1989, Boston, Mass., 26–39.
- Demmel, J. W., and X. Li [1994]. "Faster numerical algorithms via exception handling," *IEEE Trans. on Computers* 43:8, 983–992.
- Denehy, T. E., J. Bent, F. I. Popovici, A. C. Arpaci-Dusseau, and R. H. Arpaci-Dusseau [2004]. "Deconstructing storage arrays," *Proc. 11th Int'l. Conf. on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, October 7–13, 2004, Boston, Mass., 59–71.
- Desurvire, E. [1992]. "Lightwave communications: The fifth generation," *Scientific American* (International Edition) 266:1 (January), 96–103.
- Diep, T. A., C. Nelson, and J. P. Shen [1995]. "Performance evaluation of the PowerPC 620 microarchitecture," *Proc. 22nd Annual Int'l. Symposium on Computer Architecture (ISCA)*, June 22–24, 1995, Santa Margherita, Italy.
- Digital Semiconductor. [1996]. *Alpha Architecture Handbook, Version 3*, Digital Press, Maynard, Mass.
- Ditzel, D. R., and H. R. McLellan [1987]. "Branch folding in the CRISP microprocessor: Reducing the branch delay to zero," *Proc. 14th Annual Int'l. Symposium on Computer Architecture (ISCA)*, June 2–5, 1987, Pittsburgh, Penn., 2–7.
- Ditzel, D. R., and D. A. Patterson [1980]. "Retrospective on high-level language computer architecture," *Proc. Seventh Annual Int'l. Symposium on Computer Architecture (ISCA)*, May 6–8, 1980, La Baule, France, 97–104.
- Doherty, W. J., and R. P. Kelisky [1979]. "Managing VM/CMS systems for user effectiveness," *IBM Systems J.* 18:1, 143–166.
- Dongarra, J. J. [1986]. "A survey of high performance processors," *Proc. IEEE COMPCON*, March 3–6, 1986, San Francisco, 8–11.
- Dongarra, J., T. Sterling, H. Simon, and E. Strohmaier [2005]. "High-performance computing: Clusters, constellations, MPPs, and future directions," *Computing in Science & Engineering*, 7:2 (March/April), 51–59.
- Douceur, J. R., and W. J. Bolosky [1999]. "A large scale study of file-system contents," *Proc. ACM SIGMETRICS Conf. on Measurement and Modeling of Computer Systems*, May 1–9, 1999, Atlanta, Ga., 59–69.
- Douglas, J. [2005]. "Intel 8xx series and Paxville Xeon-MP microprocessors," paper presented at Hot Chips 17, August 14–16, 2005, Stanford University, Palo Alto, Calif.
- Duato, J. [1993]. "A new theory of deadlock-free adaptive routing in wormhole networks," *IEEE Trans. on Parallel and Distributed Systems* 4:12 (December) 1320–1331.

- Duato, J., and T. M. Pinkston [2001]. "A general theory for deadlock-free adaptive routing using a mixed set of resources," *IEEE Trans. on Parallel and Distributed Systems* 12:12 (December), 1219–1235.
- Duato, J., S. Yalamanchili, and L. Ni [2003]. *Interconnection Networks: An Engineering Approach*, 2nd printing, Morgan Kaufmann, San Francisco.
- Duato, J., I. Johnson, J. Flich, F. Naven, P. Garcia, and T. Nachiondo [2005a]. "A new scalable and cost-effective congestion management strategy for lossless multistage interconnection networks," *Proc. 11th Int'l. Symposium on High-Performance Computer Architecture*, February 12–16, 2005, San Francisco.
- Duato, J., O. Lysne, R. Pang, and T. M. Pinkston [2005b]. "Part I: A theory for deadlock-free dynamic reconfiguration of interconnection networks," *IEEE Trans. on Parallel and Distributed Systems* 16:5 (May), 412–427.
- Dubois, M., C. Scheurich, and F. Briggs [1988]. "Synchronization, coherence, and event ordering," *IEEE Computer* 21:2 (February), 9–21.
- Dunigan, W., K. Vetter, K. White, and P. Worley [2005]. "Performance evaluation of the Cray X1 distributed shared memory architecture," *IEEE Micro* January/February, 30–40.
- Eden, A., and T. Mudge [1998]. "The YAGS branch prediction scheme," *Proc. of the 31st Annual ACM/IEEE Int'l. Symposium on Microarchitecture*, November 30–December 2, 1998, Dallas, Tex., 69–80.
- Edmondson, J. H., P. I. Rubinfield, R. Preston, and V. Rajagopalan [1995]. "Superscalar instruction execution in the 21164 Alpha microprocessor," *IEEE Micro* 15:2, 33–43.
- Eggers, S. [1989]. "Simulation Analysis of Data Sharing in Shared Memory Multiprocessors," Ph.D. thesis, University of California, Berkeley.
- Elder, J., A. Gottlieb, C. K. Kruskal, K. P. McAuliffe, L. Randolph, M. Snir, P. Teller, and J. Wilson [1985]. "Issues related to MIMD shared-memory computers: The NYU Ultracomputer approach," *Proc. 12th Annual Int'l. Symposium on Computer Architecture (ISCA)*, June 17–19, 1985, Boston, Mass., 126–135.
- Ellis, J. R. [1986]. *Bulldog: A Compiler for VLIW Architectures*, MIT Press, Cambridge, Mass.
- Emer, J. S., and D. W. Clark [1984]. "A characterization of processor performance in the VAX-11/780," *Proc. 11th Annual Int'l. Symposium on Computer Architecture (ISCA)*, June 5–7, 1984, Ann Arbor, Mich., 301–310.
- Enriquez, P. [2001]. "What happened to my dial tone? A study of FCC service disruption reports," poster, *Richard Tapia Symposium on the Celebration of Diversity in Computing*, October 18–20, Houston, Tex.
- Erlichson, A., N. Nuckolls, G. Chesson, and J. L. Hennessy [1996]. "SoftFLASH: Analyzing the performance of clustered distributed virtual shared memory," *Proc. Seventh Int'l. Conf. on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, October 1–5, 1996, Cambridge, Mass., 210–220.
- Esmailzadeh, H., T. Cao, Y. Xi, S. M. Blackburn, and K. S. McKinley [2011]. "Looking Back on the Language and Hardware Revolution: Measured Power, Performance, and Scaling," *Proc. 16th Int'l. Conf. on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, March 5–11, 2011, Newport Beach, Calif.
- Evers, M., S. J. Patel, R. S. Chappell, and Y. N. Patt [1998]. "An analysis of correlation and predictability: What makes two-level branch predictors work," *Proc. 25th Annual Int'l. Symposium on Computer Architecture (ISCA)*, July 3–14, 1998, Barcelona, Spain, 52–61.
- Fabry, R. S. [1974]. "Capability based addressing," *Communications of the ACM* 17:7 (July), 403–412.

- Falsafi, B., and D. A. Wood [1997]. "Reactive NUMA: A design for unifying S-COMA and CC-NUMA," *Proc. 24th Annual Int'l. Symposium on Computer Architecture (ISCA)*, June 2-4, 1997, Denver, Colo., 229-240.
- Fan, X., W. Weber, and L. A. Barroso [2007]. "Power provisioning for a warehouse-sized computer," *Proc. 34th Annual Int'l. Symposium on Computer Architecture (ISCA)*, June 9-13, 2007, San Diego, Calif.
- Farkas, K. I., and N. P. Jouppi [1994]. "Complexity/performance trade-offs with non-blocking loads," *Proc. 21st Annual Int'l. Symposium on Computer Architecture (ISCA)*, April 18-21, 1994, Chicago.
- Farkas, K. I., N. P. Jouppi, and P. Chow [1995]. "How useful are non-blocking loads, stream buffers and speculative execution in multiple issue processors?," *Proc. First IEEE Symposium on High-Performance Computer Architecture*, January 22-25, 1995, Raleigh, N.C., 78-89.
- Farkas, K. I., P. Chow, N. P. Jouppi, and Z. Vranesic [1997]. "Memory-system design considerations for dynamically-scheduled processors," *Proc. 24th Annual Int'l. Symposium on Computer Architecture (ISCA)*, June 2-4, 1997, Denver, Colo., 133-143.
- Fazio, D. [1987]. "It's really much more fun building a supercomputer than it is simply inventing one," *Proc. IEEE COMPCON*, February 23-27, 1987, San Francisco, 102-105.
- Fisher, J. A. [1981]. "Trace scheduling: A technique for global microcode compaction," *IEEE Trans. on Computers* 30:7 (July), 478-490.
- Fisher, J. A. [1983]. "Very long instruction word architectures and ELI-512," *10th Annual Int'l. Symposium on Computer Architecture (ISCA)*, June 5-7, 1982, Stockholm, Sweden, 140-150.
- Fisher, J. A., and S. M. Freudenberger [1992]. "Predicting conditional branches from previous runs of a program," *Proc. Fifth Int'l. Conf. on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, October 12-15, 1992, Boston, Mass., 85-95.
- Fisher, J. A., and B. R. Rau [1993]. *Journal of Supercomputing*, January (special issue).
- Fisher, J. A., J. R. Ellis, J. C. Rutenberg, and A. Nicolau [1984]. "Parallel processing: A smart compiler and a dumb processor," *Proc. SIGPLAN Conf. on Compiler Construction*, June 17-22, 1984, Montreal, Canada, 11-16.
- Flemming, P. J., and J. J. Wallace [1986]. "How not to lie with statistics: The correct way to summarize benchmarks results," *Communications of the ACM* 29:3 (March), 218-221.
- Flynn, M. J. [1966]. "Very high-speed computing systems," *Proc. IEEE* 54:12 (December), 1901-1909.
- Forgie, J. W. [1957]. "The Lincoln TX-2 input-output system," *Proc. Western Joint Computer Conference* (February), Institute of Radio Engineers, Los Angeles, 156-160.
- Foster, C. C., and E. M. Riseman [1972]. "Percolation of code to enhance parallel dispatching and execution," *IEEE Trans. on Computers* C-21:12 (December), 1411-1415.
- Frank, S. J. [1984]. "Tightly coupled multiprocessor systems speed memory access time," *Electronics* 57:1 (January), 164-169.
- Freiman, C. V. [1961]. "Statistical analysis of certain binary division algorithms," *Proc. IRE* 49:1, 91-103.
- Friesenborg, S. E., and R. J. Wicks [1985]. *DASD Expectations: The 3380, 3380-23, and MVS/XA*, Tech. Bulletin GG22-9363-02, IBM Washington Systems Center, Gaithersburg, Md.
- Fuller, S. H., and W. E. Burr [1977]. "Measurement and evaluation of alternative computer architectures," *Computer* 10:10 (October), 24-35.
- Furber, S. B. [1996]. *ARM System Architecture*, Addison-Wesley, Harlow, England (see www.cs.man.ac.uk/amulet/publications/books/ARMSysArch).

- Gagliardi, U. O. [1973]. "Report of workshop 4—software-related advances in computer hardware," *Proc. Symposium on the High Cost of Software*, September 17–19, 1973, Monterey, Calif., 99–120.
- Gajski, D., D. Kuck, D. Lawrie, and A. Sameh [1983]. "CEDAR—a large scale multiprocessor," *Proc. Int'l. Conf. on Parallel Processing (ICPP)*, August, Columbus, Ohio, 524–529.
- Gallagher, D. M., W. Y. Chen, S. A. Mahlke, J. C. Gyllenhaal, and W. W. Hwu [1994]. "Dynamic memory disambiguation using the memory conflict buffer," *Proc. Sixth Int'l. Conf. on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, October 4–7, Santa Jose. Calif., 183–193.
- Galles, M. [1996]. "Scalable pipelined interconnect for distributed endpoint routing: The SGI SPIDER chip," *Proc. IEEE HOT Interconnects '96*, August 15–17, 1996, Stanford University, Palo Alto, Calif.
- Game, M., and A. Booker [1999]. "CodePack code compression for PowerPC processors," *MicroNews*, 5:1, www.chips.ibm.com/micronews/vol5_no1/codepack.html.
- Gao, Q. S. [1993]. "The Chinese remainder theorem and the prime memory system," *20th Annual Int'l. Symposium on Computer Architecture (ISCA)*, May 16–19, 1993, San Diego, Calif. (*Computer Architecture News* 21:2 (May), 337–340).
- Gap. [2005]. "Gap Inc. Reports Third Quarter Earnings," http://gapinc.com/public/documents/PR_Q405EarningsFeb2306.pdf.
- Gap. [2006]. "Gap Inc. Reports Fourth Quarter and Full Year Earnings," http://gapinc.com/public/documents/Q32005PressRelease_Final22.pdf.
- Garner, R., A. Agarwal, F. Briggs, E. Brown, D. Hough, B. Joy, S. Kleiman, S. Muchnick, M. Namjoo, D. Patterson, J. Pendleton, and R. Tuck [1988]. "Scalable processor architecture (SPARC)," *Proc. IEEE COMPCON*, February 29–March 4, 1988, San Francisco, 278–283.
- Gebis, J., and D. Patterson [2007]. "Embracing and extending 20th-century instruction set architectures," *IEEE Computer* 40:4 (April), 68–75.
- Gee, J. D., M. D. Hill, D. N. Pnevmatikatos, and A. J. Smith [1993]. "Cache performance of the SPEC92 benchmark suite," *IEEE Micro* 13:4 (August), 17–27.
- Gehringer, E. F., D. P. Siewiorek, and Z. Segall [1987]. *Parallel Processing: The Cm* Experience*, Digital Press, Bedford, Mass.
- Gharachorloo, K., A. Gupta, and J. L. Hennessy [1992]. "Hiding memory latency using dynamic scheduling in shared-memory multiprocessors," *Proc. 19th Annual Int'l. Symposium on Computer Architecture (ISCA)*, May 19–21, 1992, Gold Coast, Australia.
- Gharachorloo, K., D. Lenoski, J. Laudon, P. Gibbons, A. Gupta, and J. L. Hennessy [1990]. "Memory consistency and event ordering in scalable shared-memory multiprocessors," *Proc. 17th Annual Int'l. Symposium on Computer Architecture (ISCA)*, May 28–31, 1990, Seattle, Wash., 15–26.
- Ghemawat, S., H. Gobioff, and S.-T. Leung [2003]. "The Google file system," *Proc. 19th ACM Symposium on Operating Systems Principles*, October 19–22, 2003, Bolton Landing, N.Y.
- Gibson, D. H. [1967]. "Considerations in block-oriented systems design," *AFIPS Conf. Proc.* 30, 75–80.
- Gibson, G. A. [1992]. *Redundant Disk Arrays: Reliable, Parallel Secondary Storage*, ACM Distinguished Dissertation Series, MIT Press, Cambridge, Mass.
- Gibson, J. C. [1970]. "The Gibson mix," Rep. TR. 00.2043, IBM Systems Development Division, Poughkeepsie, N.Y. (research done in 1959).
- Gibson, J., R. Kunz, D. Ofelt, M. Horowitz, J. Hennessy, and M. Heinrich [2000]. "FLASH vs. (simulated) FLASH: Closing the simulation loop," *Proc. Ninth Int'l. Conf. on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, November 12–15, Cambridge, Mass., 49–58.

- Glass, C. J., and L. M. Ni [1992]. "The Turn Model for adaptive routing," *19th Annual Int'l. Symposium on Computer Architecture (ISCA)*, May 19–21, 1992, Gold Coast, Australia.
- Goldberg, D. [1991]. "What every computer scientist should know about floating-point arithmetic," *Computing Surveys* 23:1, 5–48.
- Goldberg, I. B. [1967]. "27 bits are not enough for 8-digit accuracy," *Communications of the ACM* 10:2, 105–106.
- Goldstein, S. [1987]. *Storage Performance—An Eight Year Outlook*, Tech. Rep. TR 03.308-1, Santa Teresa Laboratory, IBM Santa Teresa Laboratory, San Jose, Calif.
- Goldstine, H. H. [1972]. *The Computer: From Pascal to von Neumann*, Princeton University Press, Princeton, N.J.
- González, J., and A. González [1998]. "Limits of instruction level parallelism with data speculation," *Proc. Vector and Parallel Processing (VECPAR) Conf.*, June 21–23, 1998, Porto, Portugal, 585–598.
- Goodman, J. R. [1983]. "Using cache memory to reduce processor memory traffic," *Proc. 10th Annual Int'l. Symposium on Computer Architecture (ISCA)*, June 5–7, 1982, Stockholm, Sweden, 124–131.
- Goralski, W. [1997]. *SONET: A Guide to Synchronous Optical Network*, McGraw-Hill, New York.
- Gosling, J. B. [1980]. *Design of Arithmetic Units for Digital Computers*, Springer-Verlag, New York.
- Gray, J. [1990]. "A census of Tandem system availability between 1985 and 1990," *IEEE Trans. on Reliability*, 39:4 (October), 409–418.
- Gray, J. (ed.) [1993]. *The Benchmark Handbook for Database and Transaction Processing Systems*, 2nd ed., Morgan Kaufmann, San Francisco.
- Gray, J. [2006]. Sort benchmark home page, <http://sortbenchmark.org/>.
- Gray, J., and A. Reuter [1993]. *Transaction Processing: Concepts and Techniques*, Morgan Kaufmann, San Francisco.
- Gray, J., and D. P. Siewiorek [1991]. "High-availability computer systems," *Computer* 24:9 (September), 39–48.
- Gray, J., and C. van Ingen [2005]. *Empirical Measurements of Disk Failure Rates and Error Rates*, MSR-TR-2005-166, Microsoft Research, Redmond, Wash.
- Greenberg, A., N. Jain, S. Kandula, C. Kim, P. Lahiri, D. Maltz, P. Patel, and S. Sengupta [2009]. "VL2: A Scalable and Flexible Data Center Network," in *Proc. ACM SIGCOMM*, August 17–21, 2009, Barcelona, Spain.
- Grice, C., and M. Kanellos [2000]. "Cell phone industry at crossroads: Go high or low?," *CNET News*, August 31, technews.netscape.com/news/0-1004-201-2518386-0.html?tag=st.ne.1002.tgif.sf.
- Groe, J. B., and L. E. Larson [2000]. *CDMA Mobile Radio Design*, Artech House, Boston.
- Gunther, K. D. [1981]. "Prevention of deadlocks in packet-switched data transport systems," *IEEE Trans. on Communications* COM-29:4 (April), 512–524.
- Hagersten, E., and M. Koster [1998]. "WildFire: A scalable path for SMPs," *Proc. Fifth Int'l. Symposium on High-Performance Computer Architecture*, January 9–12, 1999, Orlando, Fla.
- Hagersten, E., A. Landin, and S. Haridi [1992]. "DDM—a cache-only memory architecture," *IEEE Computer* 25:9 (September), 44–54.
- Hamacher, V. C., Z. G. Vranesic, and S. G. Zaky [1984]. *Computer Organization*, 2nd ed., McGraw-Hill, New York.
- Hamilton, J. [2009]. "Data center networks are in my way," paper presented at the Stanford Clean Slate CTO Summit, October 23, 2009 (http://mvdirona.com/jrh/TalksAndPapers/JamesHamilton_CleanSlateCTO2009.pdf).

- Hamilton, J. [2010]. "Cloud computing economies of scale," paper presented at the *AWS Workshop on Genomics and Cloud Computing*, June 8, 2010, Seattle, Wash. (http://mvdirona.com/jrh/TalksAndPapers/JamesHamilton_GenomicsCloud20100608.pdf).
- Handy, J. [1993]. *The Cache Memory Book*, Academic Press, Boston.
- Hauck, E. A., and B. A. Dent [1968]. "Burroughs' B6500/B7500 stack mechanism," *Proc. AFIPS Spring Joint Computer Conf.*, April 30–May 2, 1968, Atlantic City, N.J., 245–251.
- Heald, R., K. Aingaran, C. Amir, M. Ang, M. Boland, A. Das, P. Dixit, G. Gouldsberry, J. Hart, T. Horel, W.-J. Hsu, J. Kaku, C. Kim, S. Kim, F. Klass, H. Kwan, R. Lo, H. McIntyre, A. Mehta, D. Murata, S. Nguyen, Y.-P. Pai, S. Patel, K. Shin, K. Tam, S. Vishwanthaiiah, J. Wu, G. Yee, and H. You [2000]. "Implementation of third-generation SPARC V9 64-b microprocessor," *ISSCC Digest of Technical Papers*, 412–413 and slide supplement.
- Heinrich, J. [1993]. *MIPS R4000 User's Manual*, Prentice Hall, Englewood Cliffs, N.J.
- Henly, M., and B. McNutt [1989]. *DASD I/O Characteristics: A Comparison of MVS to VM*, Tech. Rep. TR 02.1550 (May), IBM General Products Division, San Jose, Calif.
- Hennessy, J. [1984]. "VLSI processor architecture," *IEEE Trans. on Computers* C-33:11 (December), 1221–1246.
- Hennessy, J. [1985]. "VLSI RISC processors," *VLSI Systems Design* 6:10 (October), 22–32.
- Hennessy, J., N. Jouppi, F. Baskett, and J. Gill [1981]. "MIPS: A VLSI processor architecture," in *CMU Conference on VLSI Systems and Computations*, Computer Science Press, Rockville, Md.
- Hewlett-Packard. [1994]. *PA-RISC 2.0 Architecture Reference Manual*, 3rd ed., Hewlett-Packard, Palo Alto, Calif.
- Hewlett-Packard. [1998]. "HP's '5NINES:5MINUTES' Vision Extends Leadership and Redefines High Availability in Mission-Critical Environments," February 10, www.future.enterprisecomputing.hp.com/ta64/news/5nines_vision_pr.html.
- Hill, M. D. [1987]. "Aspects of Cache Memory and Instruction Buffer Performance," Ph.D. thesis, Tech. Rep. UCB/CSD 87/381, Computer Science Division, University of California, Berkeley.
- Hill, M. D. [1988]. "A case for direct mapped caches," *Computer* 21:12 (December), 25–40.
- Hill, M. D. [1998]. "Multiprocessors should support simple memory consistency models," *IEEE Computer* 31:8 (August), 28–34.
- Hillis, W. D. [1985]. *The Connection Multiprocessor*, MIT Press, Cambridge, Mass.
- Hillis, W. D. and G. L. Steele [1986]. "Data parallel algorithms," *Communications of the ACM* 29:12 (December), 1170–1183. (<http://doi.acm.org/10.1145/7902.7903>).
- Hinton, G., D. Sager, M. Upton, D. Boggs, D. Carmean, A. Kyker, and P. Roussel [2001]. "The microarchitecture of the Pentium 4 processor," *Intel Technology Journal*, February.
- Hintz, R. G., and D. P. Tate [1972]. "Control data STAR-100 processor design," *Proc. IEEE COMPCON*, September 12–14, 1972, San Francisco, 1–4.
- Hirata, H., K. Kimura, S. Nagamine, Y. Mochizuki, A. Nishimura, Y. Nakase, and T. Nishizawa [1992]. "An elementary processor architecture with simultaneous instruction issuing from multiple threads," *Proc. 19th Annual Int'l. Symposium on Computer Architecture (ISCA)*, May 19–21, 1992, Gold Coast, Australia, 136–145.
- Hitachi. [1997]. *SuperH RISC Engine SH7700 Series Programming Manual*, Hitachi, Santa Clara, Calif. (see www.halisp.hitachi.com/tech_prod/ and search for title).
- Ho, R., K. W. Mai, and M. A. Horowitz [2001]. "The future of wires," *Proc. of the IEEE* 89:4 (April), 490–504.
- Hoagland, A. S. [1963]. *Digital Magnetic Recording*, Wiley, New York.

- Hockney, R. W., and C. R. Jesshope [1988]. *Parallel Computers 2: Architectures, Programming and Algorithms*, Adam Hilger, Ltd., Bristol, England.
- Holland, J. H. [1959]. "A universal computer capable of executing an arbitrary number of subprograms simultaneously," *Proc. East Joint Computer Conf.* 16, 108–113.
- Holt, R. C. [1972]. "Some deadlock properties of computer systems," *ACM Computer Surveys* 4:3 (September), 179–196.
- Hopkins, M. [2000]. "A critical look at IA-64: Massive resources, massive ILP, but can it deliver?" *Microprocessor Report*, February.
- Hord, R. M. [1982]. *The Illiac-IV, The First Supercomputer*, Computer Science Press, Rockville, Md.
- Horel, T., and G. Lauterbach [1999]. "UltraSPARC-III: Designing third-generation 64-bit performance," *IEEE Micro* 19:3 (May–June), 73–85.
- Hospodor, A. D., and A. S. Hoagland [1993]. "The changing nature of disk controllers." *Proc. IEEE* 81:4 (April), 586–594.
- Hölzle, U. [2010]. "Brawny cores still beat wimpy cores, most of the time," *IEEE Micro* 30:4 (July/August).
- Hristea, C., D. Lenoski, and J. Keen [1997]. "Measuring memory hierarchy performance of cache-coherent multiprocessors using micro benchmarks," *Proc. ACM/IEEE Conf. on Supercomputing*, November 16–21, 1997, San Jose, Calif.
- Hsu, P. [1994]. "Designing the TFP microprocessor," *IEEE Micro* 18:2 (April), 2333.
- Huck, J. et al. [2000]. "Introducing the IA-64 Architecture" *IEEE Micro*, 20:5 (September–October), 12–23.
- Hughes, C. J., P. Kaul, S. V. Adve, R. Jain, C. Park, and J. Srinivasan [2001]. "Variability in the execution of multimedia applications and implications for architecture," *Proc. 28th Annual Int'l. Symposium on Computer Architecture (ISCA)*, June 30–July 4, 2001, Goteborg, Sweden, 254–265.
- Hwang, K. [1979]. *Computer Arithmetic: Principles, Architecture, and Design*, Wiley, New York.
- Hwang, K. [1993]. *Advanced Computer Architecture and Parallel Programming*, McGraw-Hill, New York.
- Hwu, W.-M., and Y. Patt [1986]. "HPSm, a high performance restricted data flow architecture having minimum functionality," *Proc. 13th Annual Int'l. Symposium on Computer Architecture (ISCA)*, June 2–5, 1986, Tokyo, 297–307.
- Hwu, W. W., S. A. Mahlke, W. Y. Chen, P. P. Chang, N. J. Warter, R. A. Bringmann, R. O. Ouellette, R. E. Hank, T. Kiyohara, G. E. Haab, J. G. Holm, and D. M. Lavery [1993]. "The superblock: An effective technique for VLIW and superscalar compilation," *J. Supercomputing* 7:1, 2 (March), 229–248.
- IBM. [1982]. *The Economic Value of Rapid Response Time*, GE20-0752-0, IBM, White Plains, N.Y., 11–82.
- IBM. [1990]. "The IBM RISC System/6000 processor" (collection of papers), *IBM J. Research and Development* 34:1 (January).
- IBM. [1994]. *The PowerPC Architecture*, Morgan Kaufmann, San Francisco.
- IBM. [2005]. "Blue Gene," *IBM J. Research and Development*, 49:2/3 (special issue).
- IEEE. [1985]. "IEEE standard for binary floating-point arithmetic," *SIGPLAN Notices* 22:2, 9–25.
- IEEE. [2005]. "Intel virtualization technology, computer," *IEEE Computer Society* 38:5 (May), 48–56.
- IEEE. 754-2008 Working Group. [2006]. "DRAFT Standard for Floating-Point Arithmetic 754-2008," <http://dx.doi.org/10.1109/IEEESTD.2008.4610935>.
- Imprimis Product Specification, 97209 Sabre Disk Drive IPI-2 Interface 1.2 GB*, Document No. 64402302, Imprimis, Dallas, Tex.

- InfiniBand Trade Association. [2001]. *InfiniBand Architecture Specifications Release 1.0.a*, www.infinibandta.org.
- Intel. [2001]. "Using MMX Instructions to Convert RGB to YUV Color Conversion," cedar.intel.com/cgi-bin/ids.dll/content/content.jsp?cntKey=Legacy::irtm_AP548_9996&cntType=IDS_EDITORIAL.
- Internet Retailer. [2005]. "The Gap launches a new site—after two weeks of downtime," *Internet® Retailer*, September 28, <http://www.internetretailer.com/2005/09/28/the-gap-launches-a-new-site-after-two-weeks-of-downtime>.
- Jain, R. [1991]. *The Art of Computer Systems Performance Analysis: Techniques for Experimental Design, Measurement, Simulation, and Modeling*, Wiley, New York.
- Jantsch, A., and H. Tenhunen (eds.) [2003]. *Networks on Chips*, Kluwer Academic Publishers, The Netherlands.
- Jimenez, D. A., and C. Lin [2002]. "Neural methods for dynamic branch prediction," *ACM Trans. on Computer Systems* 20:4 (November), 369–397.
- Johnson, M. [1990]. *Superscalar Microprocessor Design*, Prentice Hall, Englewood Cliffs, N.J.
- Jordan, H. F. [1983]. "Performance measurements on HEP—a pipelined MIMD computer," *Proc. 10th Annual Int'l. Symposium on Computer Architecture (ISCA)*, June 5–7, 1982, Stockholm, Sweden, 207–212.
- Jordan, K. E. [1987]. "Performance comparison of large-scale scientific processors: Scalar mainframes, mainframes with vector facilities, and supercomputers," *Computer* 20:3 (March), 10–23.
- Jouppi, N. P. [1990]. "Improving direct-mapped cache performance by the addition of a small fully-associative cache and prefetch buffers," *Proc. 17th Annual Int'l. Symposium on Computer Architecture (ISCA)*, May 28–31, 1990, Seattle, Wash., 364–373.
- Jouppi, N. P. [1998]. "Retrospective: Improving direct-mapped cache performance by the addition of a small fully-associative cache and prefetch buffers," *25 Years of the International Symposia on Computer Architecture (Selected Papers)*, ACM, New York, 71–73.
- Jouppi, N. P., and D. W. Wall [1989]. "Available instruction-level parallelism for superscalar and superpipelined processors," *Proc. Third Int'l. Conf. on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, April 3–6, 1989, Boston, 272–282.
- Jouppi, N. P., and S. J. E. Wilton [1994]. "Trade-offs in two-level on-chip caching," *Proc. 21st Annual Int'l. Symposium on Computer Architecture (ISCA)*, April 18–21, 1994, Chicago, 34–45.
- Kaeli, D. R., and P. G. Emma [1991]. "Branch history table prediction of moving target branches due to subroutine returns," *Proc. 18th Annual Int'l. Symposium on Computer Architecture (ISCA)*, May 27–30, 1991, Toronto, Canada, 34–42.
- Kahan, J. [1990]. "On the advantage of the 8087's stack," unpublished course notes, Computer Science Division, University of California, Berkeley.
- Kahan, W. [1968]. "7094-II system support for numerical analysis," *SHARE Secretarial Distribution* SSD-159, Department of Computer Science, University of Toronto.
- Kahaner, D. K. [1988]. "Benchmarks for 'real' programs," *SIAM News*, November.
- Kahn, R. E. [1972]. "Resource-sharing computer communication networks," *Proc. IEEE* 60:11 (November), 1397–1407.
- Kane, G. [1986]. *MIPS R2000 RISC Architecture*, Prentice Hall, Englewood Cliffs, N.J.
- Kane, G. [1996]. *PA-RISC 2.0 Architecture*, Prentice Hall, Upper Saddle River, N.J.
- Kane, G., and J. Heinrich [1992]. *MIPS RISC Architecture*, Prentice Hall, Englewood Cliffs, N.J.
- Katz, R. H., D. A. Patterson, and G. A. Gibson [1989]. "Disk system architectures for high performance computing," *Proc. IEEE* 77:12 (December), 1842–1858.

- Keckler, S. W., and W. J. Dally [1992]. "Processor coupling: Integrating compile time and runtime scheduling for parallelism," *Proc. 19th Annual Int'l. Symposium on Computer Architecture (ISCA)*, May 19–21, 1992, Gold Coast, Australia, 202–213.
- Keller, R. M. [1975]. "Look-ahead processors," *ACM Computing Surveys* 7:4 (December), 177–195.
- Keltcher, C. N., K. J. McGrath, A. Ahmed, and P. Conway [2003]. "The AMD Opteron processor for multiprocessor servers," *IEEE Micro* 23:2 (March–April), 66–76 ([dx.doi.org/10.1109/MM.2003.1191116](https://doi.org/10.1109/MM.2003.1191116)).
- Kembel, R. [2000]. "Fibre Channel: A comprehensive introduction," *Internet Week*, April.
- Kermani, P., and L. Kleinrock [1979]. "Virtual Cut-Through: A New Computer Communication Switching Technique," *Computer Networks* 3 (January), 267–286.
- Kessler, R. [1999]. "The Alpha 21264 microprocessor," *IEEE Micro* 19:2 (March/April) 24–36.
- Kilburn, T., D. B. G. Edwards, M. J. Lanigan, and F. H. Sumner [1962]. "One-level storage system," *IRE Trans. on Electronic Computers* EC-11 (April) 223–235. Also appears in D. P. Siewiorek, C. G. Bell, and A. Newell, *Computer Structures: Principles and Examples*, McGraw-Hill, New York, 1982, 135–148.
- Killian, E. [1991]. "MIPS R4000 technical overview—64 bits/100 MHz or bust," *Hot Chips III Symposium Record*, August 26–27, 1991, Stanford University, Palo Alto, Calif., 1.6–1.19.
- Kim, M. Y. [1986]. "Synchronized disk interleaving," *IEEE Trans. on Computers* C-35:11 (November), 978–988.
- Kissell, K. D. [1997]. "MIPS16: High-density for the embedded market," *Proc. Real Time Systems '97*, June 15, 1997, Las Vegas, Nev. (see www.sgi.com/MIPS/arch/MIPS16/MIPS16.whitepaper.pdf).
- Kitagawa, K., S. Tagaya, Y. Hagihara, and Y. Kanoh [2003]. "A hardware overview of SX-6 and SX-7 supercomputer," *NEC Research & Development J.* 44:1 (January), 2–7.
- Knuth, D. [1981]. *The Art of Computer Programming*, Vol. II, 2nd ed., Addison-Wesley, Reading, Mass.
- Kogge, P. M. [1981]. *The Architecture of Pipelined Computers*, McGraw-Hill, New York.
- Kohn, L., and S.-W. Fu [1989]. "A 1,000,000 transistor microprocessor," *Proc. of IEEE Int'l. Symposium on Solid State Circuits (ISSCC)*, February 15–17, 1989, New York, 54–55.
- Kohn, L., and N. Margulis [1989]. "Introducing the Intel i860 64-Bit Microprocessor," *IEEE Micro*, 9:4 (July), 15–30.
- Kontothanassis, L., G. Hunt, R. Stets, N. Hardavellas, M. Cierniak, S. Parthasarathy, W. Meira, S. Dwarkadas, and M. Scott [1997]. "VM-based shared memory on low-latency, remote-memory-access networks," *Proc. 24th Annual Int'l. Symposium on Computer Architecture (ISCA)*, June 2–4, 1997, Denver, Colo.
- Koren, I. [1989]. *Computer Arithmetic Algorithms*, Prentice Hall, Englewood Cliffs, N.J.
- Kozyrakakis, C. [2000]. "Vector IRAM: A media-oriented vector processor with embedded DRAM," paper presented at Hot Chips 12, August 13–15, 2000, Palo Alto, Calif., 13–15.
- Kozyrakakis, C., and D. Patterson, [2002]. "Vector vs. superscalar and VLIW architectures for embedded multimedia benchmarks," *Proc. 35th Annual Int'l. Symposium on Microarchitecture (MICRO-35)*, November 18–22, 2002, Istanbul, Turkey.
- Kroft, D. [1981]. "Lockup-free instruction fetch/prefetch cache organization," *Proc. Eighth Annual Int'l. Symposium on Computer Architecture (ISCA)*, May 12–14, 1981, Minneapolis, Minn., 81–87.
- Kroft, D. [1998]. "Retrospective: Lockup-free instruction fetch/prefetch cache organization," *25 Years of the International Symposia on Computer Architecture (Selected Papers)*, ACM, New York, 20–21.

- Kuck, D., P. P. Budnik, S.-C. Chen, D. H. Lawrie, R. A. Towle, R. E. Strebendt, E. W. Davis, Jr., J. Han, P. W. Kraska, and Y. Muraoka [1974]. "Measurements of parallelism in ordinary FORTRAN programs," *Computer* 7:1 (January), 37–46.
- Kuhn, D. R. [1997]. "Sources of failure in the public switched telephone network," *IEEE Computer* 30:4 (April), 31–36.
- Kumar, A. [1997]. "The HP PA-8000 RISC CPU," *IEEE Micro* 17:2 (March/April), 27–32.
- Kunimatsu, A., N. Ide, T. Sato, Y. Endo, H. Murakami, T. Kamei, M. Hirano, F. Ishihara, H. Tago, M. Oka, A. Ohba, T. Yutaka, T. Okada, and M. Suzuoki [2000]. "Vector unit architecture for emotion synthesis," *IEEE Micro* 20:2 (March–April), 40–47.
- Kunkel, S. R., and J. E. Smith [1986]. "Optimal pipelining in supercomputers," *Proc. 13th Annual Int'l. Symposium on Computer Architecture (ISCA)*, June 2–5, 1986, Tokyo, 404–414.
- Kurose, J. F., and K. W. Ross [2001]. *Computer Networking: A Top-Down Approach Featuring the Internet*, Addison-Wesley, Boston.
- Kuskin, J., D. Ofelt, M. Heinrich, J. Heinlein, R. Simoni, K. Gharachorloo, J. Chapin, D. Nakahira, J. Baxter, M. Horowitz, A. Gupta, M. Rosenblum, and J. L. Hennessy [1994]. "The Stanford FLASH multiprocessor," *Proc. 21st Annual Int'l. Symposium on Computer Architecture (ISCA)*, April 18–21, 1994, Chicago.
- Lam, M. [1988]. "Software pipelining: An effective scheduling technique for VLIW processors," *SIGPLAN Conf. on Programming Language Design and Implementation*, June 22–24, 1988, Atlanta, Ga., 318–328.
- Lam, M. S., and R. P. Wilson [1992]. "Limits of control flow on parallelism," *Proc. 19th Annual Int'l. Symposium on Computer Architecture (ISCA)*, May 19–21, 1992, Gold Coast, Australia, 46–57.
- Lam, M. S., E. E. Rothberg, and M. E. Wolf [1991]. "The cache performance and optimizations of blocked algorithms," *Proc. Fourth Int'l. Conf. on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, April 8–11, 1991, Santa Clara, Calif. (*SIGPLAN Notices* 26:4 (April), 63–74).
- Lambright, D. [2000]. "Experiences in measuring the reliability of a cache-based storage system," *Proc. of First Workshop on Industrial Experiences with Systems Software (WIESS 2000), Co-Located with the 4th Symposium on Operating Systems Design and Implementation (OSDI)*, October 22, 2000, San Diego, Calif.
- Lampert, L. [1979]. "How to make a multiprocessor computer that correctly executes multiprocess programs," *IEEE Trans. on Computers* C-28:9 (September), 241–248.
- Lang, W., J. M. Patel, and S. Shankar [2010]. "Wimpy node clusters: What about non-wimpy workloads?" *Proc. Sixth International Workshop on Data Management on New Hardware (DaMoN)*, June 7, Indianapolis, Ind.
- Laprie, J.-C. [1985]. "Dependable computing and fault tolerance: Concepts and terminology," *Proc. 15th Annual Int'l. Symposium on Fault-Tolerant Computing*, June 19–21, 1985, Ann Arbor, Mich., 2–11.
- Larson, E. R. [1973]. "Findings of fact, conclusions of law, and order for judgment," File No. 4-67, Civ. 138, *Honeywell v. Sperry-Rand and Illinois Scientific Development*, U.S. District Court for the State of Minnesota, Fourth Division (October 19).
- Laudon, J., and D. Lenoski [1997]. "The SGI Origin: A ccNUMA highly scalable server," *Proc. 24th Annual Int'l. Symposium on Computer Architecture (ISCA)*, June 2–4, 1997, Denver, Colo., 241–251.
- Laudon, J., A. Gupta, and M. Horowitz [1994]. "Interleaving: A multithreading technique targeting multiprocessors and workstations," *Proc. Sixth Int'l. Conf. on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, October 4–7, San Jose, Calif., 308–318.
- Lauterbach, G., and T. Horel [1999]. "UltraSPARC-III: Designing third generation 64-bit performance," *IEEE Micro* 19:3 (May/June).

- Lazowska, E. D., J. Zahorjan, G. S. Graham, and K. C. Sevcik [1984]. *Quantitative System Performance: Computer System Analysis Using Queueing Network Models*, Prentice Hall, Englewood Cliffs, N.J. (Although out of print, it is available online at www.cs.washington.edu/homes/lazowska/qsp/.)
- Lebeck, A. R., and D. A. Wood [1994]. "Cache profiling and the SPEC benchmarks: A case study," *Computer* 27:10 (October), 15–26.
- Lee, R. [1989]. "Precision architecture," *Computer* 22:1 (January), 78–91.
- Lee, W. V. et al. [2010]. "Debunking the 100X GPU vs. CPU myth: An evaluation of throughput computing on CPU and GPU," *Proc. 37th Annual Int'l. Symposium on Computer Architecture (ISCA)*, June 19–23, 2010, Saint-Malo, France.
- Leighton, F. T. [1992]. *Introduction to Parallel Algorithms and Architectures: Arrays, Trees, Hypercubes*, Morgan Kaufmann, San Francisco.
- Leiner, A. L. [1954]. "System specifications for the DYSEAC," *J. ACM* 1:2 (April), 57–81.
- Leiner, A. L., and S. N. Alexander [1954]. "System organization of the DYSEAC," *IRE Trans. of Electronic Computers* EC-3:1 (March), 1–10.
- Leiserson, C. E. [1985]. "Fat trees: Universal networks for hardware-efficient supercomputing," *IEEE Trans. on Computers* C-34:10 (October), 892–901.
- Lenoski, D., J. Laudon, K. Gharachorloo, A. Gupta, and J. L. Hennessy [1990]. "The Stanford DASH multiprocessor," *Proc. 17th Annual Int'l. Symposium on Computer Architecture (ISCA)*, May 28–31, 1990, Seattle, Wash., 148–159.
- Lenoski, D., J. Laudon, K. Gharachorloo, W.-D. Weber, A. Gupta, J. L. Hennessy, M. A. Horowitz, and M. Lam [1992]. "The Stanford DASH multiprocessor," *IEEE Computer* 25:3 (March), 63–79.
- Levy, H., and R. Eckhouse [1989]. *Computer Programming and Architecture: The VAX*, Digital Press, Boston.
- Li, K. [1988]. "IVY: A shared virtual memory system for parallel computing," *Proc. 1988 Int'l. Conf. on Parallel Processing*, Pennsylvania State University Press, University Park, Penn.
- Li, S., K. Chen, J. B. Brockman, and N. Jouppi [2011]. "Performance Impacts of Non-blocking Caches in Out-of-order Processors," HP Labs Tech Report HPL-2011-65 (full text available at <http://Library.hp.com/techpubs/2011/Hpl-2011-65.html>).
- Lim, K., P. Ranganathan, J. Chang, C. Patel, T. Mudge, and S. Reinhardt [2008]. "Understanding and designing new system architectures for emerging warehouse-computing environments," *Proc. 35th Annual Int'l. Symposium on Computer Architecture (ISCA)*, June 21–25, 2008, Beijing, China.
- Lincoln, N. R. [1982]. "Technology and design trade offs in the creation of a modern supercomputer," *IEEE Trans. on Computers* C-31:5 (May), 363–376.
- Lindholm, T., and F. Yellin [1999]. *The Java Virtual Machine Specification*, 2nd ed., Addison-Wesley, Reading, Mass. (also available online at java.sun.com/docs/books/vmspec/).
- Lipasti, M. H., and J. P. Shen [1996]. "Exceeding the dataflow limit via value prediction," *Proc. 29th Int'l. Symposium on Microarchitecture*, December 2–4, 1996, Paris, France.
- Lipasti, M. H., C. B. Wilkerson, and J. P. Shen [1996]. "Value locality and load value prediction," *Proc. Seventh Conf. on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, October 1–5, 1996, Cambridge, Mass., 138–147.
- Liptay, J. S. [1968]. "Structural aspects of the System/360 Model 85, Part II: The cache," *IBM Systems J.* 7:1, 15–21.
- Lo, J., L. Barroso, S. Eggers, K. Gharachorloo, H. Levy, and S. Parekh [1998]. "An analysis of database workload performance on simultaneous multithreaded processors," *Proc. 25th Annual Int'l. Symposium on Computer Architecture (ISCA)*, July 3–14, 1998, Barcelona, Spain, 39–50.

- Lo, J., S. Eggers, J. Emer, H. Levy, R. Stamm, and D. Tullsen [1997]. "Converting thread-level parallelism into instruction-level parallelism via simultaneous multithreading," *ACM Trans. on Computer Systems* 15:2 (August), 322–354.
- Lovett, T., and S. Thakkar [1988]. "The Symmetry multiprocessor system," *Proc. 1988 Int'l. Conf. of Parallel Processing*, University Park, Penn., 303–310.
- Lubeck, O., J. Moore, and R. Mendez [1985]. "A benchmark comparison of three supercomputers: Fujitsu VP-200, Hitachi S810/20, and Cray X-MP/2," *Computer* 18:12 (December), 10–24.
- Luk, C.-K., and T. C. Mowry [1999]. "Automatic compiler-inserted prefetching for pointer-based applications," *IEEE Trans. on Computers* 48:2 (February), 134–141.
- Lunde, A. [1977]. "Empirical evaluation of some features of instruction set processor architecture," *Communications of the ACM* 20:3 (March), 143–152.
- Luszczek, P., J. J. Dongarra, D. Koester, R. Rabenseifner, B. Lucas, J. Kepner, J. McCalpin, D. Bailey, and D. Takahashi [2005]. "Introduction to the HPC challenge benchmark suite," Lawrence Berkeley National Laboratory, Paper LBNL-57493 (April 25), repositories.cdlib.org/lbnl/LBNL-57493.
- Maberly, N. C. [1966]. *Mastering Speed Reading*, New American Library, New York.
- Magenheimer, D. J., L. Peters, K. W. Pettis, and D. Zuras [1988]. "Integer multiplication and division on the HP precision architecture," *IEEE Trans. on Computers* 37:8, 980–990.
- Mahlke, S. A., W. Y. Chen, W.-M. Hwu, B. R. Rau, and M. S. Schlansker [1992]. "Sentinel scheduling for VLIW and superscalar processors," *Proc. Fifth Int'l. Conf. on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, October 12–15, 1992, Boston, 238–247.
- Mahlke, S. A., R. E. Hank, J. E. McCormick, D. I. August, and W. W. Hwu [1995]. "A comparison of full and partial predicated execution support for ILP processors," *Proc. 22nd Annual Int'l. Symposium on Computer Architecture (ISCA)*, June 22–24, 1995, Santa Margherita, Italy, 138–149.
- Major, J. B. [1989]. "Are queuing models within the grasp of the unwashed?," *Proc. Int'l. Conf. on Management and Performance Evaluation of Computer Systems*, December 11–15, 1989, Reno, Nev., 831–839.
- Markstein, P. W. [1990]. "Computation of elementary functions on the IBM RISC System/6000 processor," *IBM J. Research and Development* 34:1, 111–119.
- Mathis, H. M., A. E. Mercias, J. D. McCalpin, R. J. Eickemeyer, and S. R. Kunkel [2005]. "Characterization of the multithreading (SMT) efficiency in Power5," *IBM J. Research and Development*, 49:4/5 (July/September), 555–564.
- McCalpin, J. [2005]. "STREAM: Sustainable Memory Bandwidth in High Performance Computers," www.cs.virginia.edu/stream/.
- McCalpin, J., D. Bailey, and D. Takahashi [2005]. *Introduction to the HPC Challenge Benchmark Suite*, Paper LBNL-57493 Lawrence Berkeley National Laboratory, University of California, Berkeley, repositories.cdlib.org/lbnl/LBNL-57493.
- McCormick, J., and A. Knies [2002]. "A brief analysis of the SPEC CPU2000 benchmarks on the Intel Itanium 2 processor," paper presented at Hot Chips 14, August 18–20, 2002, Stanford University, Palo Alto, Calif.
- McFarling, S. [1989]. "Program optimization for instruction caches," *Proc. Third Int'l. Conf. on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, April 3–6, 1989, Boston, 183–191.
- McFarling, S. [1993]. *Combining Branch Predictors*, WRL Technical Note TN-36, Digital Western Research Laboratory, Palo Alto, Calif.
- McFarling, S., and J. Hennessy [1986]. "Reducing the cost of branches," *Proc. 13th Annual Int'l. Symposium on Computer Architecture (ISCA)*, June 2–5, 1986, Tokyo, 396–403.

- McGhan, H., and M. O'Connor [1998]. "PicoJava: A direct execution engine for Java bytecode," *Computer* 31:10 (October), 22–30.
- McKeeman, W. M. [1967]. "Language directed computer design," *Proc. AFIPS Fall Joint Computer Conf.*, November 14–16, 1967, Washington, D.C., 413–417.
- McMahon, F. M. [1986]. "*The Livermore FORTRAN Kernels: A Computer Test of Numerical Performance Range*," Tech. Rep. UCRL-55745, Lawrence Livermore National Laboratory, University of California, Livermore.
- McNairy, C., and D. Soltis [2003]. "Itanium 2 processor microarchitecture," *IEEE Micro* 23:2 (March–April), 44–55.
- Mead, C., and L. Conway [1980]. *Introduction to VLSI Systems*, Addison-Wesley, Reading, Mass.
- Mellor-Crummey, J. M., and M. L. Scott [1991]. "Algorithms for scalable synchronization on shared-memory multiprocessors," *ACM Trans. on Computer Systems* 9:1 (February), 21–65.
- Menabrea, L. F. [1842]. "Sketch of the analytical engine invented by Charles Babbage," *Bibliothèque Universelle de Genève*, 82 (October).
- Menon, A., J. Renato Santos, Y. Turner, G. Janakiraman, and W. Zwaenepoel [2005]. "Diagnosing performance overheads in the xen virtual machine environment," *Proc. First ACM/USENIX Int'l. Conf. on Virtual Execution Environments*, June 11–12, 2005, Chicago, 13–23.
- Merlin, P. M., and P. J. Schweitzer [1980]. "Deadlock avoidance in store-and-forward networks. Part I. Store-and-forward deadlock," *IEEE Trans. on Communications* COM-28:3 (March), 345–354.
- Metcalfe, R. M. [1993]. "Computer/network interface design: Lessons from Arpanet and Ethernet," *IEEE J. on Selected Areas in Communications* 11:2 (February), 173–180.
- Metcalfe, R. M., and D. R. Boggs [1976]. "Ethernet: Distributed packet switching for local computer networks," *Communications of the ACM* 19:7 (July), 395–404.
- Metropolis, N., J. Howlett, and G. C. Rota (eds.) [1980]. *A History of Computing in the Twentieth Century*, Academic Press, New York.
- Meyer, R. A., and L. H. Seawright [1970]. A virtual machine time sharing system, *IBM Systems J.* 9:3, 199–218.
- Meyers, G. J. [1978]. "The evaluation of expressions in a storage-to-storage architecture," *Computer Architecture News* 7:3 (October), 20–23.
- Meyers, G. J. [1982]. *Advances in Computer Architecture*, 2nd ed., Wiley, New York.
- Micron. [2004]. "Calculating Memory System Power for DDR2," <http://download.micron.com/pdf/pubs/designline/dl1Q04.pdf>.
- Micron. [2006]. "The Micron[®] System-Power Calculator," <http://www.micron.com/systemcalc>.
- MIPS. [1997]. "MIPS16 Application Specific Extension Product Description," www.sgi.com/MIPS/arch/MIPS16/mips16.pdf.
- Miranker, G. S., J. Rubenstein, and J. Sanguinetti [1988]. "Squeezing a Cray-class super-computer into a single-user package," *Proc. IEEE COMPCON*, February 29–March 4, 1988, San Francisco, 452–456.
- Mitchell, D. [1989]. "The Transputer: The time is now," *Computer Design* (RISC suppl.), 40–41.
- Mitsubishi. [1996]. *Mitsubishi 32-Bit Single Chip Microcomputer M32R Family Software Manual*, Mitsubishi, Cypress, Calif.
- Miura, K., and K. Uchida [1983]. "FACOM vector processing system: VP100/200," *Proc. NATO Advanced Research Workshop on High-Speed Computing*, June 20–22, 1983, Jülich, West Germany. Also appears in K. Hwang, ed., "Superprocessors: Design and applications," *IEEE* (August 1984), 59–73.

- Miya, E. N. [1985]. "Multiprocessor/distributed processing bibliography," *Computer Architecture News* 13:1, 27–29.
- Montoye, R. K., E. Hokenek, and S. L. Runyon [1990]. "Design of the IBM RISC System/6000 floating-point execution," *IBM J. Research and Development* 34:1, 59–70.
- Moore, B., A. Padegs, R. Smith, and W. Bucholz [1987]. "Concepts of the System/370 vector architecture," *14th Annual Int'l. Symposium on Computer Architecture (ISCA)*, June 2–5, 1987, Pittsburgh, Penn., 282–292.
- Moore, G. E. [1965]. "Cramming more components onto integrated circuits," *Electronics*, 38:8 (April 19), 114–117.
- Morse, S., B. Ravenal, S. Mazor, and W. Pohlman [1980]. "Intel microprocessors—8080 to 8086," *Computer* 13:10 (October).
- Moshovos, A., and G. S. Sohi [1997]. "Streamlining inter-operation memory communication via data dependence prediction," *Proc. 30th Annual Int'l. Symposium on Microarchitecture*, December 1–3, Research Triangle Park, N.C., 235–245.
- Moshovos, A., S. Breach, T. N. Vijaykumar, and G. S. Sohi [1997]. "Dynamic speculation and synchronization of data dependences," *24th Annual Int'l. Symposium on Computer Architecture (ISCA)*, June 2–4, 1997, Denver, Colo.
- Moussouris, J., L. Crudele, D. Freitas, C. Hansen, E. Hudson, S. Przybylski, T. Riordan, and C. Rowen [1986]. "A CMOS RISC processor with integrated system functions," *Proc. IEEE COMPCON*, March 3–6, 1986, San Francisco, 191.
- Mowry, T. C., S. Lam, and A. Gupta [1992]. "Design and evaluation of a compiler algorithm for prefetching," *Proc. Fifth Int'l. Conf. on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, October 12–15, 1992, Boston (*SIGPLAN Notices* 27:9 (September), 62–73).
- MSN Money. [2005]. "Amazon Shares Tumble after Rally Fizzles," <http://moneycentral.msn.com/content/CNBCTV/Articles/Dispatches/P133695.asp>.
- Muchnick, S. S. [1988]. "Optimizing compilers for SPARC," *Sun Technology* 1:3 (Summer), 64–77.
- Mueller, M., L. C. Alves, W. Fischer, M. L. Fair, and I. Modi [1999]. "RAS strategy for IBM S/390 G5 and G6," *IBM J. Research and Development* 43:5–6 (September–November), 875–888.
- Mukherjee, S. S., C. Weaver, J. S. Emer, S. K. Reinhardt, and T. M. Austin [2003]. "Measuring architectural vulnerability factors," *IEEE Micro* 23:6, 70–75.
- Murphy, B., and T. Gent [1995]. "Measuring system and software reliability using an automated data collection process," *Quality and Reliability Engineering International* 11:5 (September–October), 341–353.
- Myer, T. H., and I. E. Sutherland [1968]. "On the design of display processors," *Communications of the ACM* 11:6 (June), 410–414.
- Narayanan, D., E. Thereska, A. Donnelly, S. Elnikety, and A. Rowstron [2009]. "Migrating server storage to SSDs: Analysis of trade-offs," *Proc. 4th ACM European Conf. on Computer Systems*, April 1–3, 2009, Nuremberg, Germany.
- National Research Council. [1997]. *The Evolution of Untethered Communications*, Computer Science and Telecommunications Board, National Academy Press, Washington, D.C.
- National Storage Industry Consortium. [1998]. "Tape Roadmap," www.nsic.org.
- Nelson, V. P. [1990]. "Fault-tolerant computing: Fundamental concepts," *Computer* 23:7 (July), 19–25.
- Ngai, T.-F., and M. J. Irwin [1985]. "Regular, area-time efficient carry-lookahead adders," *Proc. Seventh IEEE Symposium on Computer Arithmetic*, June 4–6, 1985, University of Illinois, Urbana, 9–15.

- Nicolau, A., and J. A. Fisher [1984]. "Measuring the parallelism available for very long instruction word architectures," *IEEE Trans. on Computers* C-33:11 (November), 968–976.
- Nikhil, R. S., G. M. Papadopoulos, and Arvind [1992]. "T: A multithreaded massively parallel architecture," *Proc. 19th Annual Int'l. Symposium on Computer Architecture (ISCA)*, May 19–21, 1992, Gold Coast, Australia, 156–167.
- Noordergraaf, L., and R. van der Pas [1999]. "Performance experiences on Sun's WildFire prototype," *Proc. ACM/IEEE Conf. on Supercomputing*, November 13–19, 1999, Portland, Ore.
- Nyberg, C. R., T. Barclay, Z. Cvetanovic, J. Gray, and D. Lomet [1994]. "AlphaSort: A RISC machine sort," *Proc. ACM SIGMOD*, May 24–27, 1994, Minneapolis, Minn.
- Oka, M., and M. Suzuoki [1999]. "Designing and programming the emotion engine," *IEEE Micro* 19:6 (November–December), 20–28.
- Okada, S., S. Okada, Y. Matsuda, T. Yamada, and A. Kobayashi [1999]. "System on a chip for digital still camera," *IEEE Trans. on Consumer Electronics* 45:3 (August), 584–590.
- Oliker, L., A. Canning, J. Carter, J. Shalf, and S. Ethier [2004]. "Scientific computations on modern parallel vector systems," *Proc. ACM/IEEE Conf. on Supercomputing*, November 6–12, 2004, Pittsburgh, Penn., 10.
- Pabst, T. [2000]. "Performance Showdown at 133 MHz FSB—The Best Platform for Coppermine," www6.tomshardware.com/mainboard/00q1/000302/.
- Padua, D., and M. Wolfe [1986]. "Advanced compiler optimizations for supercomputers," *Communications of the ACM* 29:12 (December), 1184–1201.
- Palacharla, S., and R. E. Kessler [1994]. "Evaluating stream buffers as a secondary cache replacement," *Proc. 21st Annual Int'l. Symposium on Computer Architecture (ISCA)*, April 18–21, 1994, Chicago, 24–33.
- Palmer, J., and S. Morse [1984]. *The 8087 Primer*, John Wiley & Sons, New York, 93.
- Pan, S.-T., K. So, and J. T. Rameh [1992]. "Improving the accuracy of dynamic branch prediction using branch correlation," *Proc. Fifth Int'l. Conf. on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, October 12–15, 1992, Boston, 76–84.
- Partridge, C. [1994]. *Gigabit Networking*, Addison-Wesley, Reading, Mass.
- Patterson, D. [1985]. "Reduced instruction set computers," *Communications of the ACM* 28:1 (January), 8–21.
- Patterson, D. [2004]. "Latency lags bandwidth," *Communications of the ACM* 47:10 (October), 71–75.
- Patterson, D. A., and D. R. Ditzel [1980]. "The case for the reduced instruction set computer," *Computer Architecture News* 8:6 (October), 25–33.
- Patterson, D. A., and J. L. Hennessy [2004]. *Computer Organization and Design: The Hardware/Software Interface*, 3rd ed., Morgan Kaufmann, San Francisco.
- Patterson, D. A., G. A. Gibson, and R. H. Katz [1987]. *A Case for Redundant Arrays of Inexpensive Disks (RAID)*, Tech. Rep. UCB/CSD 87/391, University of California, Berkeley. Also appeared in *Proc. ACM SIGMOD*, June 1–3, 1988, Chicago, 109–116.
- Patterson, D. A., P. Garrison, M. Hill, D. Lioupis, C. Nyberg, T. Sippel, and K. Van Dyke [1983]. "Architecture of a VLSI instruction cache for a RISC," *10th Annual Int'l. Conf. on Computer Architecture Conf. Proc.*, June 13–16, 1983, Stockholm, Sweden, 108–116.
- Pavan, P., R. Bez, P. Olivo, and E. Zanoni [1997]. "Flash memory cells—an overview," *Proc. IEEE* 85:8 (August), 1248–1271.
- Peh, L. S., and W. J. Dally [2001]. "A delay model and speculative architecture for pipeline routers," *Proc. 7th Int'l. Symposium on High-Performance Computer Architecture*, January 22–24, 2001, Monterrey, Mexico.

- Peng, V., S. Samudrala, and M. Gavrielov [1987]. "On the implementation of shifters, multipliers, and dividers in VLSI floating point units," *Proc. 8th IEEE Symposium on Computer Arithmetic*, May 19–21, 1987, Como, Italy, 95–102.
- Pfister, G. F. [1998]. *In Search of Clusters*, 2nd ed., Prentice Hall, Upper Saddle River, N.J.
- Pfister, G. F., W. C. Brantley, D. A. George, S. L. Harvey, W. J. Kleinfekder, K. P. McAuliffe, E. A. Melton, V. A. Norton, and J. Weiss [1985]. "The IBM research parallel processor prototype (RP3): Introduction and architecture," *Proc. 12th Annual Int'l. Symposium on Computer Architecture (ISCA)*, June 17–19, 1985, Boston, Mass., 764–771.
- Pineiro, E., W. D. Weber, and L. A. Barroso [2007]. "Failure trends in a large disk drive population," *Proc. 5th USENIX Conference on File and Storage Technologies (FAST '07)*, February 13–16, 2007, San Jose, Calif.
- Pinkston, T. M. [2004]. "Deadlock characterization and resolution in interconnection networks," in M. C. Zhu and M. P. Fanti, eds., *Deadlock Resolution in Computer-Integrated Systems*, CRC Press, Boca Raton, FL, 445–492.
- Pinkston, T. M., and J. Shin [2005]. "Trends toward on-chip networked microsystems," *Int'l. J. of High Performance Computing and Networking* 3:1, 3–18.
- Pinkston, T. M., and S. Warnakulasuriya [1997]. "On deadlocks in interconnection networks," *24th Annual Int'l. Symposium on Computer Architecture (ISCA)*, June 2–4, 1997, Denver, Colo.
- Pinkston, T. M., A. Benner, M. Krause, I. Robinson, and T. Sterling [2003]. "InfiniBand: The 'de facto' future standard for system and local area networks or just a scalable replacement for PCI buses?" *Cluster Computing* (special issue on communication architecture for clusters) 6:2 (April), 95–104.
- Postiff, M. A., D. A. Greene, G. S. Tyson, and T. N. Mudge [1999]. "The limits of instruction level parallelism in SPEC95 applications," *Computer Architecture News* 27:1 (March), 31–40.
- Przybylski, S. A. [1990]. *Cache Design: A Performance-Directed Approach*, Morgan Kaufmann, San Francisco.
- Przybylski, S. A., M. Horowitz, and J. L. Hennessy [1988]. "Performance trade-offs in cache design," *15th Annual Int'l. Symposium on Computer Architecture*, May 30–June 2, 1988, Honolulu, Hawaii, 290–298.
- Puente, V., R. Beivide, J. A. Gregorio, J. M. Prellezo, J. Duato, and C. Izu [1999]. "Adaptive bubble router: A design to improve performance in torus networks," *Proc. 28th Int'l. Conference on Parallel Processing*, September 21–24, 1999, Aizu-Wakamatsu, Fukushima, Japan.
- Radin, G. [1982]. "The 801 minicomputer," *Proc. Symposium Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, March 1–3, 1982, Palo Alto, Calif., 39–47.
- Rajesh Bordawekar, Uday Bondhugula, Ravi Rao: Believe it or not!: multi-core CPUs can match GPU performance for a FLOP-intensive application! 19th International Conference on Parallel Architecture and Compilation Techniques (PACT 2010), Vienna, Austria, September 11–15, 2010: 537–538.
- Ramamoorthy, C. V., and H. F. Li [1977]. "Pipeline architecture," *ACM Computing Surveys* 9:1 (March), 61–102.
- Ranganathan, P., P. Leech, D. Irwin, and J. Chase [2006]. "Ensemble-Level Power Management for Dense Blade Servers," *Proc. 33rd Annual Int'l. Symposium on Computer Architecture (ISCA)*, June 17–21, 2006, Boston, Mass., 66–77.
- Rau, B. R. [1994]. "Iterative modulo scheduling: An algorithm for software pipelining loops," *Proc. 27th Annual Int'l. Symposium on Microarchitecture*, November 30–December 2, 1994, San Jose, Calif., 63–74.

- Rau, B. R., C. D. Glaeser, and R. L. Picard [1982]. "Efficient code generation for horizontal architectures: Compiler techniques and architectural support," *Proc. Ninth Annual Int'l. Symposium on Computer Architecture (ISCA)*, April 26–29, 1982, Austin, Tex., 131–139.
- Rau, B. R., D. W. L. Yen, W. Yen, and R. A. Towle [1989]. "The Cydra 5 departmental supercomputer: Design philosophies, decisions, and trade-offs," *IEEE Computers* 22:1 (January), 12–34.
- Reddi, V. J., B. C. Lee, T. Chilimbi, and K. Vaid [2010]. "Web search using mobile cores: Quantifying and mitigating the price of efficiency," *Proc. 37th Annual Int'l. Symposium on Computer Architecture (ISCA)*, June 19–23, 2010, Saint-Malo, France.
- Redmond, K. C., and T. M. Smith [1980]. *Project Whirlwind—The History of a Pioneer Computer*, Digital Press, Boston.
- Reinhardt, S. K., J. R. Larus, and D. A. Wood [1994]. "Tempest and Typhoon: User-level shared memory," *21st Annual Int'l. Symposium on Computer Architecture (ISCA)*, April 18–21, 1994, Chicago, 325–336.
- Reinman, G., and N. P. Jouppi. [1999]. "Extensions to CACTI," research.compaq.com/wrl/people/jouppi/CACTI.html.
- Retzberg, R. D., W. R. Crowther, P. P. Carvey, and R. S. Towlinson [1990]. "The Monarch parallel processor hardware design," *IEEE Computer* 23:4 (April), 18–30.
- Riemens, A., K. A. Vissers, R. J. Schutten, F. W. Sijstermans, G. J. Hekstra, and G. D. La Hei [1999]. "Trimedia CPU64 application domain and benchmark suite," *Proc. IEEE Int'l. Conf. on Computer Design: VLSI in Computers and Processors (ICCD'99)*, October 10–13, 1999, Austin, Tex., 580–585.
- Riseman, E. M., and C. C. Foster [1972]. "Percolation of code to enhance parallel dispatching and execution," *IEEE Trans. on Computers* C-21:12 (December), 1411–1415.
- Robin, J., and C. Irvine [2000]. "Analysis of the Intel Pentium's ability to support a secure virtual machine monitor." *Proc. USENIX Security Symposium*, August 14–17, 2000, Denver, Colo.
- Robinson, B., and L. Blount [1986]. *The VM/HPO 3880-23 Performance Results*, IBM Tech. Bulletin GG66-0247-00, IBM Washington Systems Center, Gaithersburg, Md.
- Ropers, A., H. W. Lollman, and J. Wellhausen [1999]. *DSPstone: Texas Instruments TMS320C54x*, Tech. Rep. IB 315 1999/9-ISS-Version 0.9, Aachen University of Technology, Aachen, Germany (www.ert.rwth-aachen.de/Projekte/Tools/coal/dspstone_c54x/index.html).
- Rosenblum, M., S. A. Herrod, E. Witchel, and A. Gupta [1995]. "Complete computer simulation: The SimOS approach," in *IEEE Parallel and Distributed Technology* (now called *Concurrency*) 4:3, 34–43.
- Rowen, C., M. Johnson, and P. Ries [1988]. "The MIPS R3010 floating-point coprocessor," *IEEE Micro* 8:3 (June), 53–62.
- Russell, R. M. [1978]. "The Cray-1 processor system," *Communications of the ACM* 21:1 (January), 63–72.
- Rymarczyk, J. [1982]. "Coding guidelines for pipelined processors," *Proc. Symposium Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, March 1–3, 1982, Palo Alto, Calif., 12–19.
- Saavedra-Barrera, R. H. [1992]. "CPU Performance Evaluation and Execution Time Prediction Using Narrow Spectrum Benchmarking," Ph.D. dissertation, University of California, Berkeley.
- Salem, K., and H. Garcia-Molina [1986]. "Disk striping," *Proc. 2nd Int'l. IEEE Conf. on Data Engineering*, February 5–7, 1986, Washington, D.C., 249–259.

- Saltzer, J. H., D. P. Reed, and D. D. Clark [1984]. "End-to-end arguments in system design," *ACM Trans. on Computer Systems* 2:4 (November), 277–288.
- Samples, A. D., and P. N. Hilfinger [1988]. *Code Reorganization for Instruction Caches*, Tech. Rep. UCB/CSD 88/447, University of California, Berkeley.
- Santoro, M. R., G. Bewick, and M. A. Horowitz [1989]. "Rounding algorithms for IEEE multipliers," *Proc. Ninth IEEE Symposium on Computer Arithmetic*, September 6–8, Santa Monica, Calif., 176–183.
- Satran, J., D. Smith, K. Meth, C. Sapuntzakis, M. Wakeley, P. Von Stamwitz, R. Haagens, E. Zeidner, L. Dalle Ore, and Y. Klein [2001]. "iSCSI," IPS Working Group of IETF, Internet draft www.ietf.org/internet-drafts/draft-ietf-ips-iscsi-07.txt.
- Saulsbury, A., T. Wilkinson, J. Carter, and A. Landin [1995]. "An argument for Simple COMA," *Proc. First IEEE Symposium on High-Performance Computer Architectures*, January 22–25, 1995, Raleigh, N.C., 276–285.
- Schneck, P. B. [1987]. *Superprocessor Architecture*, Kluwer Academic Publishers, Norwell, Mass.
- Schroeder, B., and G. A. Gibson [2007]. "Understanding failures in petascale computers," *J. of Physics Conf. Series* 78(1), 188–198.
- Schroeder, B., E. Pinheiro, and W.-D. Weber [2009]. "DRAM errors in the wild: a large-scale field study," *Proc. Eleventh Int'l. Joint Conf. on Measurement and Modeling of Computer Systems (SIGMETRICS)*, June 15–19, 2009, Seattle, Wash.
- Schurman, E., and J. Brutlag [2009]. "The user and business impact of server delays," *Proc. Velocity: Web Performance and Operations Conf.*, June 22–24, 2009, San Jose, Calif.
- Schwartz, J. T. [1980]. "Ultracomputers," *ACM Trans. on Programming Languages and Systems* 4:2, 484–521.
- Scott, N. R. [1985]. *Computer Number Systems and Arithmetic*, Prentice Hall, Englewood Cliffs, N.J.
- Scott, S. L. [1996]. "Synchronization and communication in the T3E multiprocessor," *Seventh Int'l. Conf. on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, October 1–5, 1996, Cambridge, Mass.
- Scott, S. L., and J. Goodman [1994]. "The impact of pipelined channels on k -ary n -cube networks," *IEEE Trans. on Parallel and Distributed Systems* 5:1 (January), 1–16.
- Scott, S. L., and G. M. Thorson [1996]. "The Cray T3E network: Adaptive routing in a high performance 3D torus," *Proc. IEEE HOT Interconnects '96*, August 15–17, 1996, Stanford University, Palo Alto, Calif., 14–156.
- Scranton, R. A., D. A. Thompson, and D. W. Hunter [1983]. *The Access Time Myth*, Tech. Rep. RC 10197 (45223), IBM, Yorktown Heights, N.Y.
- Seagate. [2000]. *Seagate Cheetah 73 Family: ST173404LW/LWV/LC/LCV Product Manual*, Vol. 1, Seagate, Scotts Valley, Calif. (www.seagate.com/support/disc/manuals/scsi/29478b.pdf).
- Seitz, C. L. [1985]. "The Cosmic Cube (concurrent computing)," *Communications of the ACM* 28:1 (January), 22–33.
- Senior, J. M. [1993]. *Optical Fiber Communications: Principles and Practice*, 2nd ed., Prentice Hall, Hertfordshire, U.K.
- Sharangpani, H., and K. Arora [2000]. "Itanium Processor Microarchitecture," *IEEE Micro* 20:5 (September–October), 24–43.
- Shurkin, J. [1984]. *Engines of the Mind: A History of the Computer*, W. W. Norton, New York.
- Shustek, L. J. [1978]. "Analysis and Performance of Computer Instruction Sets," Ph.D. dissertation, Stanford University, Palo Alto, Calif.
- Silicon Graphics. [1996]. *MIPS V Instruction Set* (see http://www.sgi.com/MIPS/arch/ISA5/#MIPSV_idx).

- Singh, J. P., J. L. Hennessy, and A. Gupta [1993]. "Scaling parallel programs for multiprocessors: Methodology and examples," *Computer* 26:7 (July), 22–33.
- Sinharoy, B., R. N. Koala, J. M. Tandler, R. J. Eickemeyer, and J. B. Joyner [2005]. "POWER5 system microarchitecture," *IBM J. Research and Development*, 49:4–5, 505–521.
- Sites, R. [1979]. *Instruction Ordering for the CRAY-1 Computer*, Tech. Rep. 78-CS-023, Dept. of Computer Science, University of California, San Diego.
- Sites, R. L. (ed.) [1992]. *Alpha Architecture Reference Manual*, Digital Press, Burlington, Mass.
- Sites, R. L., and R. Witek, (eds.) [1995]. *Alpha Architecture Reference Manual*, 2nd ed., Digital Press, Newton, Mass.
- Skadron, K., and D. W. Clark [1997]. "Design issues and tradeoffs for write buffers," *Proc. Third Int'l. Symposium on High-Performance Computer Architecture*, February 1–5, 1997, San Antonio, Tex., 144–155.
- Skadron, K., P. S. Ahuja, M. Martonosi, and D. W. Clark [1999]. "Branch prediction, instruction-window size, and cache size: Performance tradeoffs and simulation techniques," *IEEE Trans. on Computers* 48:11 (November).
- Slater, R. [1987]. *Portraits in Silicon*, MIT Press, Cambridge, Mass.
- Slotnick, D. L., W. C. Borck, and R. C. McReynolds [1962]. "The Solomon computer," *Proc. AFIPS Fall Joint Computer Conf.*, December 4–6, 1962, Philadelphia, Penn., 97–107.
- Smith, A. J. [1982]. "Cache memories," *Computing Surveys* 14:3 (September), 473–530.
- Smith, A., and J. Lee [1984]. "Branch prediction strategies and branch-target buffer design," *Computer* 17:1 (January), 6–22.
- Smith, B. J. [1978]. "A pipelined, shared resource MIMD computer," *Proc. Int'l. Conf. on Parallel Processing (ICPP)*, August, Bellaire, Mich., 6–8.
- Smith, B. J. [1981]. "Architecture and applications of the HEP multiprocessor system," *Real-Time Signal Processing IV* 298 (August), 241–248.
- Smith, J. E. [1981]. "A study of branch prediction strategies," *Proc. Eighth Annual Int'l. Symposium on Computer Architecture (ISCA)*, May 12–14, 1981, Minneapolis, Minn., 135–148.
- Smith, J. E. [1984]. "Decoupled access/execute computer architectures," *ACM Trans. on Computer Systems* 2:4 (November), 289–308.
- Smith, J. E. [1988]. "Characterizing computer performance with a single number," *Communications of the ACM* 31:10 (October), 1202–1206.
- Smith, J. E. [1989]. "Dynamic instruction scheduling and the Astronautics ZS-1," *Computer* 22:7 (July), 21–35.
- Smith, J. E., and J. R. Goodman [1983]. "A study of instruction cache organizations and replacement policies," *Proc. 10th Annual Int'l. Symposium on Computer Architecture (ISCA)*, June 5–7, 1982, Stockholm, Sweden, 132–137.
- Smith, J. E., and A. R. Pleszkun [1988]. "Implementing precise interrupts in pipelined processors," *IEEE Trans. on Computers* 37:5 (May), 562–573. (This paper is based on an earlier paper that appeared in *Proc. 12th Annual Int'l. Symposium on Computer Architecture (ISCA)*, June 17–19, 1985, Boston, Mass.)
- Smith, J. E., G. E. Dermer, B. D. Vanderwam, S. D. Klinger, C. M. Rozewski, D. L. Fowler, K. R. Scidmore, and J. P. Laudon [1987]. "The ZS-1 central processor," *Proc. Second Int'l. Conf. on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, October 5–8, 1987, Palo Alto, Calif., 199–204.
- Smith, M. D., M. Horowitz, and M. S. Lam [1992]. "Efficient superscalar performance through boosting," *Proc. Fifth Int'l. Conf. on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, October 12–15, 1992, Boston, 248–259.

- Smith, M. D., M. Johnson, and M. A. Horowitz [1989]. "Limits on multiple instruction issue," *Proc. Third Int'l. Conf. on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, April 3–6, 1989, Boston, 290–302.
- Smotherman, M. [1989]. "A sequencing-based taxonomy of I/O systems and review of historical machines," *Computer Architecture News* 17:5 (September), 5–15. Reprinted in *Computer Architecture Readings*, M. D. Hill, N. P. Jouppi, and G. S. Sohi, eds., Morgan Kaufmann, San Francisco, 1999, 451–461.
- Sodani, A., and G. Sohi [1997]. "Dynamic instruction reuse," *Proc. 24th Annual Int'l. Symposium on Computer Architecture (ISCA)*, June 2–4, 1997, Denver, Colo.
- Sohi, G. S. [1990]. "Instruction issue logic for high-performance, interruptible, multiple functional unit, pipelined computers," *IEEE Trans. on Computers* 39:3 (March), 349–359.
- Sohi, G. S., and S. Vajapeyam [1989]. "Tradeoffs in instruction format design for horizontal architectures," *Proc. Third Int'l. Conf. on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, April 3–6, 1989, Boston, 15–25.
- Soundararajan, V., M. Heinrich, B. Vergheese, K. Gharachorloo, A. Gupta, and J. L. Hennessy [1998]. "Flexible use of memory for replication/migration in cache-coherent DSM multiprocessors," *Proc. 25th Annual Int'l. Symposium on Computer Architecture (ISCA)*, July 3–14, 1998, Barcelona, Spain, 342–355.
- SPEC. [1989]. *SPEC Benchmark Suite Release 1.0* (October 2).
- SPEC. [1994]. *SPEC Newsletter* (June).
- Sporer, M., F. H. Moss, and C. J. Mathais [1988]. "An introduction to the architecture of the Stellar Graphics supercomputer," *Proc. IEEE COMPCON*, February 29–March 4, 1988, San Francisco, 464.
- Spurgeon, C. [2001]. "Charles Spurgeon's Ethernet Web Site," wwwhost.ots.utexas.edu/ethernet/ethernet-home.html.
- Spurgeon, C. [2006]. "Charles Spurgeon's Ethernet Web SITE," www.ethermanage.com/ethernet/ethernet.html.
- Stenström, P., T. Joe, and A. Gupta [1992]. "Comparative performance evaluation of cache-coherent NUMA and COMA architectures," *Proc. 19th Annual Int'l. Symposium on Computer Architecture (ISCA)*, May 19–21, 1992, Gold Coast, Australia, 80–91.
- Sterling, T. [2001]. *Beowulf PC Cluster Computing with Windows and Beowulf PC Cluster Computing with Linux*, MIT Press, Cambridge, Mass.
- Stern, N. [1980]. "Who invented the first electronic digital computer?" *Annals of the History of Computing* 2:4 (October), 375–376.
- Stevens, W. R. [1994–1996]. *TCP/IP Illustrated* (three volumes), Addison-Wesley, Reading, Mass.
- Stokes, J. [2000]. "Sound and Vision: A Technical Overview of the Emotion Engine," arstechnica.com/reviews/1q00/playstation2/ee-1.html.
- Stone, H. [1991]. *High Performance Computers*, Addison-Wesley, New York.
- Strauss, W. [1998]. "DSP Strategies 2002," www.usadata.com/market_research/spr_05/spr_r127-005.htm.
- Strecker, W. D. [1976]. "Cache memories for the PDP-11?," *Proc. Third Annual Int'l. Symposium on Computer Architecture (ISCA)*, January 19–21, 1976, Tampa, Fla., 155–158.
- Strecker, W. D. [1978]. "VAX-11/780: A virtual address extension of the PDP-11 family," *Proc. AFIPS National Computer Conf.*, June 5–8, 1978, Anaheim, Calif., 47, 967–980.
- Sugumar, R. A., and S. G. Abraham [1993]. "Efficient simulation of caches under optimal replacement with applications to miss characterization," *Proc. ACM SIGMETRICS Conf. on Measurement and Modeling of Computer Systems*, May 17–21, 1993, Santa Clara, Calif., 24–35.

- Sun Microsystems. [1989]. *The SPARC Architectural Manual*, Version 8, Part No. 8001399-09, Sun Microsystems, Santa Clara, Calif.
- Sussenguth, E. [1999]. "IBM's ACS-1 Machine," *IEEE Computer* 22:11 (November).
- Swan, R. J., S. H. Fuller, and D. P. Siewiorek [1977]. "Cm*—a modular, multi-microprocessor," *Proc. AFIPS National Computing Conf.*, June 13–16, 1977, Dallas, Tex., 637–644.
- Swan, R. J., A. Bechtolsheim, K. W. Lai, and J. K. Ousterhout [1977]. "The implementation of the Cm* multi-microprocessor," *Proc. AFIPS National Computing Conf.*, June 13–16, 1977, Dallas, Tex., 645–654.
- Swartzlander, E. (ed.) [1990]. *Computer Arithmetic*, IEEE Computer Society Press, Los Alamitos, Calif.
- Takagi, N., H. Yasuura, and S. Yajima [1985]. "High-speed VLSI multiplication algorithm with a redundant binary addition tree," *IEEE Trans. on Computers* C-34:9, 789–796.
- Talagala, N. [2000]. "Characterizing Large Storage Systems: Error Behavior and Performance Benchmarks," Ph.D. dissertation, Computer Science Division, University of California, Berkeley.
- Talagala, N., and D. Patterson [1999]. *An Analysis of Error Behavior in a Large Storage System*, Tech. Report UCB/CSD-99-1042, Computer Science Division, University of California, Berkeley.
- Talagala, N., R. Arpaci-Dusseau, and D. Patterson [2000]. *Micro-Benchmark Based Extraction of Local and Global Disk Characteristics*, CSD-99-1063, Computer Science Division, University of California, Berkeley.
- Talagala, N., S. Asami, D. Patterson, R. Futernick, and D. Hart [2000]. "The art of massive storage: A case study of a Web image archive," *Computer* (November).
- Tamir, Y., and G. Frazier [1992]. "Dynamically-allocated multi-queue buffers for VLSI communication switches," *IEEE Trans. on Computers* 41:6 (June), 725–734.
- Tanenbaum, A. S. [1978]. "Implications of structured programming for machine architecture," *Communications of the ACM* 21:3 (March), 237–246.
- Tanenbaum, A. S. [1988]. *Computer Networks*, 2nd ed., Prentice Hall, Englewood Cliffs, N.J.
- Tang, C. K. [1976]. "Cache design in the tightly coupled multiprocessor system," *Proc. AFIPS National Computer Conf.*, June 7–10, 1976, New York, 749–753.
- Tanqueray, D. [2002]. "The Cray X1 and supercomputer road map," *Proc. 13th Daresbury Machine Evaluation Workshop*, December 11–12, 2002, Daresbury Laboratories, Daresbury, Cheshire, U.K.
- Tarjan, D., S. Thoziyoor, and N. Jouppi [2005]. "HPL Technical Report on CACTI 4.0," www.hpl.hp.com/techreports/2006/HPL=2006+86.html.
- Taylor, G. S. [1981]. "Compatible hardware for division and square root," *Proc. 5th IEEE Symposium on Computer Arithmetic*, May 18–19, 1981, University of Michigan, Ann Arbor, Mich., 127–134.
- Taylor, G. S. [1985]. "Radix 16 SRT dividers with overlapped quotient selection stages," *Proc. Seventh IEEE Symposium on Computer Arithmetic*, June 4–6, 1985, University of Illinois, Urbana, Ill., 64–71.
- Taylor, G., P. Hilfinger, J. Larus, D. Patterson, and B. Zorn [1986]. "Evaluation of the SPUR LISP architecture," *Proc. 13th Annual Int'l. Symposium on Computer Architecture (ISCA)*, June 2–5, 1986, Tokyo.
- Taylor, M. B., W. Lee, S. P. Amarasinghe, and A. Agarwal [2005]. "Scalar operand networks," *IEEE Trans. on Parallel and Distributed Systems* 16:2 (February), 145–162.
- Tendler, J. M., J. S. Dodson, J. S. Fields, Jr., H. Le, and B. Sinharoy [2002]. "Power4 system microarchitecture," *IBM J. Research and Development* 46:1, 5–26.
- Texas Instruments. [2000]. "History of Innovation: 1980s," www.ti.com/corp/docs/company/history/1980s.shtml.

- Tezzaron Semiconductor. [2004]. *Soft Errors in Electronic Memory*, White Paper, Tezzaron Semiconductor, Naperville, Ill. (http://www.tezzaron.com/about/papers/soft_errors_1_1_secure.pdf).
- Thacker, C. P., E. M. McCreight, B. W. Lampson, R. F. Sproull, and D. R. Boggs [1982]. "Alto: A personal computer," in D. P. Siewiorek, C. G. Bell, and A. Newell, eds., *Computer Structures: Principles and Examples*, McGraw-Hill, New York, 549–572.
- Thadhani, A. J. [1981]. "Interactive user productivity," *IBM Systems J.* 20:4, 407–423.
- Thekkath, R., A. P. Singh, J. P. Singh, S. John, and J. L. Hennessy [1997]. "An evaluation of a commercial CC-NUMA architecture—the CONVEX Exemplar SPP1200," *Proc. 11th Int'l. Parallel Processing Symposium (IPPS)*, April 1–7, 1997, Geneva, Switzerland.
- Thorlin, J. F. [1967]. "Code generation for PIE (parallel instruction execution) computers," *Proc. Spring Joint Computer Conf.*, April 18–20, 1967, Atlantic City, N.J., 27.
- Thornton, J. E. [1964]. "Parallel operation in the Control Data 6600," *Proc. AFIPS Fall Joint Computer Conf., Part II*, October 27–29, 1964, San Francisco, 26, 33–40.
- Thornton, J. E. [1970]. *Design of a Computer, the Control Data 6600*, Scott, Foresman, Glenview, Ill.
- Tjaden, G. S., and M. J. Flynn [1970]. "Detection and parallel execution of independent instructions," *IEEE Trans. on Computers* C-19:10 (October), 889–895.
- Tomasulo, R. M. [1967]. "An efficient algorithm for exploiting multiple arithmetic units," *IBM J. Research and Development* 11:1 (January), 25–33.
- Torrellas, J., A. Gupta, and J. Hennessy [1992]. "Characterizing the caching and synchronization performance of a multiprocessor operating system," *Proc. Fifth Int'l. Conf. on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, October 12–15, 1992, Boston (*SIGPLAN Notices* 27:9 (September), 162–174).
- Touma, W. R. [1993]. *The Dynamics of the Computer Industry: Modeling the Supply of Workstations and Their Components*, Kluwer Academic, Boston.
- Tuck, N., and D. Tullsen [2003]. "Initial observations of the simultaneous multithreading Pentium 4 processor," *Proc. 12th Int. Conf. on Parallel Architectures and Compilation Techniques (PACT'03)*, September 27–October 1, 2003, New Orleans, La., 26–34.
- Tullsen, D. M., S. J. Eggers, and H. M. Levy [1995]. "Simultaneous multithreading: Maximizing on-chip parallelism," *Proc. 22nd Annual Int'l. Symposium on Computer Architecture (ISCA)*, June 22–24, 1995, Santa Margherita, Italy, 392–403.
- Tullsen, D. M., S. J. Eggers, J. S. Emer, H. M. Levy, J. L. Lo, and R. L. Stamm [1996]. "Exploiting choice: Instruction fetch and issue on an implementable simultaneous multithreading processor," *Proc. 23rd Annual Int'l. Symposium on Computer Architecture (ISCA)*, May 22–24, 1996, Philadelphia, Penn., 191–202.
- Ungar, D., R. Blau, P. Foley, D. Samples, and D. Patterson [1984]. "Architecture of SOAR: Smalltalk on a RISC," *Proc. 11th Annual Int'l. Symposium on Computer Architecture (ISCA)*, June 5–7, 1984, Ann Arbor, Mich., 188–197.
- Unger, S. H. [1958]. "A computer oriented towards spatial problems," *Proc. Institute of Radio Engineers* 46:10 (October), 1744–1750.
- Vahdat, A., M. Al-Fares, N. Farrington, R. Niranjan Mysore, G. Porter, and S. Radhakrishnan [2010]. "Scale-Out Networking in the Data Center," *IEEE Micro* 30:4 (July/August), 29–41.
- Vaidya, A. S., A. Sivasubramaniam, and C. R. Das [1997]. "Performance benefits of virtual channels and adaptive routing: An application-driven study," *Proc. ACM/IEEE Conf. on Supercomputing*, November 16–21, 1997, San Jose, Calif.
- Vajapeyam, S. [1991]. "Instruction-Level Characterization of the Cray Y-MP Processor," Ph.D. thesis, Computer Sciences Department, University of Wisconsin-Madison.
- van Eijndhoven, J. T. J., F. W. Sijstermans, K. A. Vissers, E. J. D. Pol, M. I. A. Tromp, P. Struik, R. H. J. Bloks, P. van der Wolf, A. D. Pimentel, and H. P. E. Vranken [1999]. "Trimedia CPU64 architecture," *Proc. IEEE Int'l. Conf. on Computer Design: VLSI in Computers and Processors (ICCD'99)*, October 10–13, 1999, Austin, Tex., 586–592.

- Van Vleck, T. [2005]. "The IBM 360/67 and CP/CMS," <http://www.multicians.org/thvv/360-67.html>.
- von Eicken, T., D. E. Culler, S. C. Goldstein, and K. E. Schauer [1992]. "Active Messages: A mechanism for integrated communication and computation," *Proc. 19th Annual Int'l. Symposium on Computer Architecture (ISCA)*, May 19–21, 1992, Gold Coast, Australia.
- Waingold, E., M. Taylor, D. Srikrishna, V. Sarkar, W. Lee, V. Lee, J. Kim, M. Frank, P. Finch, R. Barua, J. Babb, S. Amarasinghe, and A. Agarwal [1997]. "Baring it all to software: Raw Machines," *IEEE Computer* 30 (September), 86–93.
- Wakerly, J. [1989]. *Microcomputer Architecture and Programming*, Wiley, New York.
- Wall, D. W. [1991]. "Limits of instruction-level parallelism," *Proc. Fourth Int'l. Conf. on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, April 8–11, 1991, Palo Alto, Calif., 248–259.
- Wall, D. W. [1993]. *Limits of Instruction-Level Parallelism*, Research Rep. 93/6, Western Research Laboratory, Digital Equipment Corp., Palo Alto, Calif.
- Walrand, J. [1991]. *Communication Networks: A First Course*, Aksen Associates/Irwin, Homewood, Ill.
- Wang, W.-H., J.-L. Baer, and H. M. Levy [1989]. "Organization and performance of a two-level virtual-real cache hierarchy," *Proc. 16th Annual Int'l. Symposium on Computer Architecture (ISCA)*, May 28–June 1, 1989, Jerusalem, 140–148.
- Watanabe, T. [1987]. "Architecture and performance of the NEC supercomputer SX system," *Parallel Computing* 5, 247–255.
- Waters, F. (ed.) [1986]. *IBM RT Personal Computer Technology*, SA 23-1057, IBM, Austin, Tex.
- Watson, W. J. [1972]. "The TI ASC—a highly modular and flexible super processor architecture," *Proc. AFIPS Fall Joint Computer Conf.*, December 5–7, 1972, Anaheim, Calif., 221–228.
- Weaver, D. L., and T. Germond [1994]. *The SPARC Architectural Manual*, Version 9, Prentice Hall, Englewood Cliffs, N.J.
- Weicker, R. P. [1984]. "Dhrystone: A synthetic systems programming benchmark," *Communications of the ACM* 27:10 (October), 1013–1030.
- Weiss, S., and J. E. Smith [1984]. "Instruction issue logic for pipelined supercomputers," *Proc. 11th Annual Int'l. Symposium on Computer Architecture (ISCA)*, June 5–7, 1984, Ann Arbor, Mich., 110–118.
- Weiss, S., and J. E. Smith [1987]. "A study of scalar compilation techniques for pipelined supercomputers," *Proc. Second Int'l. Conf. on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, October 5–8, 1987, Palo Alto, Calif., 105–109.
- Weiss, S., and J. E. Smith [1994]. *Power and PowerPC*, Morgan Kaufmann, San Francisco.
- Wendel, D., R. Kalla, J. Friedrich, J. Kahle, J. Leenstra, C. Lichtenau, B. Sinharoy, W. Starke, and V. Zyuban [2010]. "The Power7 processor SoC," *Proc. Int'l. Conf. on IC Design and Technology*, June 2–4, 2010, Grenoble, France, 71–73.
- Weste, N., and K. Eshraghian [1993]. *Principles of CMOS VLSI Design: A Systems Perspective*, 2nd ed., Addison-Wesley, Reading, Mass.
- Wiecek, C. [1982]. "A case study of the VAX 11 instruction set usage for compiler execution," *Proc. Symposium on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, March 1–3, 1982, Palo Alto, Calif., 177–184.
- Wilkes, M. [1965]. "Slave memories and dynamic storage allocation," *IEEE Trans. Electronic Computers* EC-14:2 (April), 270–271.

- Wilkes, M. V. [1982]. "Hardware support for memory protection: Capability implementations," *Proc. Symposium on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, March 1–3, 1982, Palo Alto, Calif., 107–116.
- Wilkes, M. V. [1985]. *Memoirs of a Computer Pioneer*, MIT Press, Cambridge, Mass.
- Wilkes, M. V. [1995]. *Computing Perspectives*, Morgan Kaufmann, San Francisco.
- Wilkes, M. V., D. J. Wheeler, and S. Gill [1951]. *The Preparation of Programs for an Electronic Digital Computer*, Addison-Wesley, Cambridge, Mass.
- Williams, S., A. Waterman, and D. Patterson [2009]. "Roofline: An insightful visual performance model for multicore architectures," *Communications of the ACM*, 52:4 (April), 65–76.
- Williams, T. E., M. Horowitz, R. L. Alverson, and T. S. Yang [1987]. "A self-timed chip for division," in P. Losleben, ed., *1987 Stanford Conference on Advanced Research in VLSI*, MIT Press, Cambridge, Mass.
- Wilson, A. W., Jr. [1987]. "Hierarchical cache/bus architecture for shared-memory multiprocessors," *Proc. 14th Annual Int'l. Symposium on Computer Architecture (ISCA)*, June 2–5, 1987, Pittsburgh, Penn., 244–252.
- Wilson, R. P., and M. S. Lam [1995]. "Efficient context-sensitive pointer analysis for C programs," *Proc. ACM SIGPLAN'95 Conf. on Programming Language Design and Implementation*, June 18–21, 1995, La Jolla, Calif., 1–12.
- Wolfe, A., and J. P. Shen [1991]. "A variable instruction stream extension to the VLIW architecture," *Proc. Fourth Int'l. Conf. on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, April 8–11, 1991, Palo Alto, Calif., 2–14.
- Wood, D. A., and M. D. Hill [1995]. "Cost-effective parallel computing," *IEEE Computer* 28:2 (February), 69–72.
- Wulf, W. [1981]. "Compilers and computer architecture," *Computer* 14:7 (July), 41–47.
- Wulf, W., and C. G. Bell [1972]. "C.mmp—A multi-mini-processor," *Proc. AFIPS Fall Joint Computer Conf.*, December 5–7, 1972, Anaheim, Calif., 765–777.
- Wulf, W., and S. P. Harbison [1978]. "Reflections in a pool of processors—an experience report on C.mmp/Hydra," *Proc. AFIPS National Computing Conf.* June 5–8, 1978, Anaheim, Calif., 939–951.
- Wulf, W. A., and S. A. McKee [1995]. "Hitting the memory wall: Implications of the obvious," *ACM SIGARCH Computer Architecture News*, 23:1 (March), 20–24.
- Wulf, W. A., R. Levin, and S. P. Harbison [1981]. *Hydra/C.mmp: An Experimental Computer System*, McGraw-Hill, New York.
- Yamamoto, W., M. J. Serrano, A. R. Talcott, R. C. Wood, and M. Nemirosky [1994]. "Performance estimation of multistreamed, superscalar processors," *Proc. 27th Annual Hawaii Int'l. Conf. on System Sciences*, January 4–7, 1994, Maui, 195–204.
- Yang, Y., and G. Mason [1991]. "Nonblocking broadcast switching networks," *IEEE Trans. on Computers* 40:9 (September), 1005–1015.
- Yeager, K. [1996]. "The MIPS R10000 superscalar microprocessor," *IEEE Micro* 16:2 (April), 28–40.
- Yeh, T., and Y. N. Patt [1993a]. "Alternative implementations of two-level adaptive branch prediction," *Proc. 19th Annual Int'l. Symposium on Computer Architecture (ISCA)*, May 19–21, 1992, Gold Coast, Australia, 124–134.
- Yeh, T., and Y. N. Patt [1993b]. "A comparison of dynamic branch predictors that use two levels of branch history," *Proc. 20th Annual Int'l. Symposium on Computer Architecture (ISCA)*, May 16–19, 1993, San Diego, Calif., 257–266.