Chapter 1 Fundamentals of Quantitative Design and Analysis

© 2019 Elsevier Inc. All rights reserved.



Figure 1.1 Growth in processor performance over 40 years. This chart plots program performance relative to the VAX 11/780 as measured by the SPEC integer benchmarks (see Section 1.8). Prior to the mid-1980s, growth in processor performance was largely technology-driven and averaged about 22% per year, or doubling performance every 3.5 years. The increase in growth to about 52% starting in 1986, or doubling every 2 years, is attributable to more advanced architectural and organizational ideas typified in RISC architectures. By 2003 this growth led to a difference in performance of an approximate factor of 25 versus the performance that would have occurred if it had continued at the 22% rate. In 2003 the limits of power due to the end of Dennard scaling and the available instruction-level parallelism slowed uniprocessor performance to 23% per year until 2011, or doubling every 3.5 years. (The fastest SPECintbase performance since 2007 has had automatic parallelization turned on, so uniprocessor speed is harder to gauge. These results are limited to single-chip systems with usually four cores per chip.) From 2011 to 2015, the annual improvement was less than 12%, or doubling every 8 years in part due to the limits of parallelism of Amdahl's Law. Since 2015, with the end of Moore's Law, improvement has been just 3.5% per year, or doubling every 20 years! Performance for floating-point-oriented calculations follows the same trends, but typically has 1% to 2% higher annual growth in each shaded region. Figure 1.11 on page 27 shows the improvement in clock rates for these same eras. Because SPEC has changed over the years, performance of newer machines is estimated by a scaling factor that relates the performance for different versions of SPEC: SPEC89, SPEC92, SPEC95, SPEC2000, and SPEC2006. There are too few results for SPEC2017 to plot yet.

Feature	Personal mobile device (PMD)	Desktop	Server	Clusters/warehouse- scale computer	Internet of things/ embedded
Price of system	\$100-\$1000	\$300-\$2500	\$5000-\$10,000,000	\$100,000-\$200,000,000	\$10-\$100,000
Price of microprocessor	\$10-\$100	\$50-\$500	\$200-\$2000	\$50-\$250	\$0.01-\$100
Critical system design issues	Cost, energy, media performance, responsiveness	Price- performance, energy, graphics performance	Throughput, availability, scalability, energy	Price-performance, throughput, energy proportionality	Price, energy, application- specific performance

Figure 1.2 A summary of the five mainstream computing classes and their system characteristics. Sales in 2015 included about 1.6 billion PMDs (90% cell phones), 275 million desktop PCs, and 15 million servers. The total number of embedded processors sold was nearly 19 billion. In total, 14.8 billion ARM-technology-based chips were shipped in 2015. Note the wide range in system price for servers and embedded systems, which go from USB keys to network routers. For servers, this range arises from the need for very large-scale multiprocessor systems for high-end transaction processing.

Application		Annual losses with downtime of						
	Cost of downtime per hour	1% (87.6 h/year)	0.5% (43.8 h/year)	0.1% (8.8 h/year)				
Brokerage service	\$4,000,000	\$350,400,000	\$175,200,000	\$35,000,000				
Energy	\$1,750,000	\$153,300,000	\$76,700,000	\$15,300,000				
Telecom	\$1,250,000	\$109,500,000	\$54,800,000	\$11,000,000				
Manufacturing	\$1,000,000	\$87,600,000	\$43,800,000	\$8,800,000				
Retail	\$650,000	\$56,900,000	\$28,500,000	\$5,700,000				
Health care	\$400,000	\$35,000,000	\$17,500,000	\$3,500,000				
Media	\$50,000	\$4,400,000	\$2,200,000	\$400,000				

Figure 1.3 Costs rounded to nearest \$100,000 of an unavailable system are shown by analyzing the cost of downtime (in terms of immediately lost revenue), assuming three different levels of availability, and that downtime is distributed uniformly. These data are from Landstrom (2014) and were collected and analyzed by Contingency Planning Research.

Register	Name	Use	Saver
x0	zero	The constant value 0	N.A.
x1	ra	Return address	Caller
x2	sp	Stack pointer	Callee
х3	gp	Global pointer	_
×4	tp	Thread pointer	_
x5-x7	t0-t2	Temporaries	Caller
x8	s0/fp	Saved register/frame pointer	Callee
x9	s1	Saved register	Callee
x10-x11	a0-a1	Function arguments/return values	Caller
x12-x17	a2-a7	Function arguments	Caller
x18-x27	s2-s11	Saved registers	Callee
x28-x31	t3-t6	Temporaries	Caller
f0-f7	ft0-ft7	FP temporaries	Caller
f8-f9	fs0-fs1	FP saved registers	Callee
f10-f11	fa0-fa1	FP function arguments/return values	Caller
f12-f17	fa2-fa7	FP function arguments	Caller
f18-f27	fs2-fs11	FP saved registers	Callee
f28-f31	ft8-ft11	FP temporaries	Caller

Figure 1.4 RISC-V registers, names, usage, and calling conventions. In addition to the 32 general-purpose registers (x0–x31), RISC-V has 32 floating-point registers (f0–f31) that can hold either a 32-bit single-precision number or a 64-bit double-precision number. The registers that are preserved across a procedure call are labeled "Callee" saved.

Instruction type/opcode	Instruction meaning
Data transfers	Move data between registers and memory, or between the integer and FP or special registers; only memory address mode is 12-bit displacement+contents of a GPR
lb,lbu,sb	Load byte, load byte unsigned, store byte (to/from integer registers)
lh,lhu,sh	Load half word, load half word unsigned, store half word (to/from integer registers)
lw,lwu,sw	Load word, load word unsigned, store word (to/from integer registers)
ld,sd	Load double word, store double word (to/from integer registers)
flw, fld, fsw, fsd	Load SP float, load DP float, store SP float, store DP float
fmvx, fmv.x	Copy from/to integer register to/from floating-point register; ""=S for single- precision, D for double-precision
csrrw,csrrwi,csrrs, csrrsi,csrrc,csrrci	Read counters and write status registers, which include counters: clock cycles, time, instructions retired
Arithmetic/logical	Operations on integer or logical data in GPRs
add,addi,addw,addiw	Add, add immediate (all immediates are 12 bits), add 32-bits only & sign-extend to 64 bits, add immediate 32-bits only
sub, subw	Subtract, subtract 32-bits only
mul,mulw,mulh,mulhsu, mulhu	Multiply, multiply 32-bits only, multiply upper half, multiply upper half signed- unsigned, multiply upper half unsigned
div,divu,rem,remu	Divide, divide unsigned, remainder, remainder unsigned
divw,divuw,remw,remuw	Divide and remainder: as previously, but divide only lower 32-bits, producing 32-bit sign-extended result
and, andi	And, and immediate
or,ori,xor,xori	Or, or immediate, exclusive or, exclusive or immediate
lui	Load upper immediate; loads bits 31-12 of register with immediate, then sign-extends
auipc	Adds immediate in bits $31-12$ with zeros in lower bits to PC; used with JALR to transfer control to any 32-bit address
sll,slli,srl,srli,sra, srai	Shifts: shift left logical, right logical, right arithmetic; both variable and immediate forms
sllw,slliw,srlw,srliw, sraw,sraiw	Shifts: as previously, but shift lower 32-bits, producing 32-bit sign-extended result
slt,slti,sltu,sltiu	Set less than, set less than immediate, signed and unsigned
Control	Conditional branches and jumps; PC-relative or through register
beq, bne, blt, bge, bltu, bgeu	Branch GPR equal/not equal; less than; greater than or equal, signed and unsigned
jal,jalr	Jump and link: save PC+4, target is PC-relative (JAL) or a register (JALR); if specify $\times 0$ as destination register, then acts as a simple jump
ecall	Make a request to the supporting execution environment, which is usually an OS
ebreak	Debuggers used to cause control to be transferred back to a debugging environment
fence,fence.i	Synchronize threads to guarantee ordering of memory accesses; synchronize instructions and data for stores to instruction memory

Figure 1.5 Subset of the instructions in RISC-V. RISC-V has a base set of instructions (R64I) and offers optional extensions: multiply-divide (RVM), single-precision floating point (RVF), double-precision floating point (RVD). This figure includes RVM and the next one shows RVF and RVD. Appendix A gives much more detail on RISC-V.

Instruction type/opcode	Instruction meaning
Floating point	FP operations on DP and SP formats
fadd.d, fadd.s	Add DP, SP numbers
fsub.d,fsub.s	Subtract DP, SP numbers
fmul.d,fmul.s	Multiply DP, SP floating point
fmadd.d,fmadd.s,fnmadd.d, fnmadd.s	Multiply-add DP, SP numbers; negative multiply-add DP, SP numbers
fmsub.d,fmsub.s,fnmsub.d, fnmsub.s	Multiply-sub DP, SP numbers; negative multiply-sub DP, SP numbers
fdiv.d,fdiv.s	Divide DP, SP floating point
fsqrt.d,fsqrt.s	Square root DP, SP floating point
fmax.d,fmax.s,fmin.d, fmin.s	Maximum and minimum DP, SP floating point
fcvt, fcvtu, fcvtu	Convert instructions: $FCVT$. x. y converts from type x to type y, where x and y are L (64-bit integer), W (32-bit integer), D (DP), or S (SP). Integers can be unsigned (U)
feq,flt,fle	Floating-point compare between floating-point registers and record the Boolean result in integer register; " $_$ " = S for single-precision, D for double-precision
fclass.d, fclass.s	Writes to integer register a 10-bit mask that indicates the class of the floating-point number $(-\infty, +\infty, -0, +0, NaN,)$
fsgnj,fsgnjn, fsgnjx	Sign-injection instructions that changes only the sign bit: copy sign bit from other source, the oppositive of sign bit of other source, XOR of the 2 sign bits

Figure 1.6 Floating point instructions for RISC-V. RISC-V has a base set of instructions (R64I) and offers optional extensions for single-precision floating point (RVF) and double-precision floating point (RVD). SP = single precision; DP = double precision.

31	25 24	4 2	20 19	15	14 12 ⁻	11	7	6	0
funct7		rs2	rs	1	funct3	rd		opcode	R-type
									-
imm [11:0]		rs	1	funct3	rd		opcode	I-type
									-
imm [11:5]		rs2	rs	1	funct3	imm [4:0]		opcode	S-type
									-
imm [12] imm [1	0:5]	rs2	rs	1	funct3	imm [4:1 11]		opcode	B-type
									-
imm [31:12]					rd		opcode	U-type	
									1
imn	n [20	10:1 11 1	9:12]			rd		opcode	J-type

Figure 1.7 The base RISC-V instruction set architecture formats. All instructions are 32 bits long. The R format is for integer register-to-register operations, such as ADD, SUB, and so on. The I format is for loads and immediate operations, such as LD and ADDI. The B format is for branches and the J format is for jumps and link. The S format is for stores. Having a separate format for stores allows the three register specifiers (rd, rs1, rs2) to always be in the same location in all formats. The U format is for the wide immediate instructions (LUI, AUIPC).

Functional requirements	Typical features required or supported
Application area	Target of computer
Personal mobile device	Real-time performance for a range of tasks, including interactive performance for graphics, video, and audio; energy efficiency (Chapters 2–5 and 7; Appendix A)
General-purpose desktop	Balanced performance for a range of tasks, including interactive performance for graphics, video, and audio (Chapters 2–5; Appendix A)
Servers	Support for databases and transaction processing; enhancements for reliability and availability; support for scalability (Chapters 2, 5, and 7; Appendices A, D, and F)
Clusters/warehouse-scale computers	Throughput performance for many independent tasks; error correction for memory; energy proportionality (Chapters 2, 6, and 7; Appendix F)
Internet of things/embedded computing	Often requires special support for graphics or video (or other application-specific extension); power limitations and power control may be required; real-time constraints (Chapters 2, 3, 5, and 7; Appendices A and E)
Level of software compatibility	Determines amount of existing software for computer
At programming language	Most flexible for designer; need new compiler (Chapters 3, 5, and 7; Appendix A)
Object code or binary compatible	Instruction set architecture is completely defined—little flexibility—but no investment needed in software or porting programs (Appendix A)
Operating system requirements	Necessary features to support chosen OS (Chapter 2; Appendix B)
Size of address space	Very important feature (Chapter 2); may limit applications
Memory management	Required for modern OS; may be paged or segmented (Chapter 2)
Protection	Different OS and application needs: page versus segment; virtual machines (Chapter 2)
Standards	Certain standards may be required by marketplace
Floating point	Format and arithmetic: IEEE 754 standard (Appendix J), special arithmetic for graphics or signal processing
I/O interfaces	For I/O devices: Serial ATA, Serial Attached SCSI, PCI Express (Appendices D and F)
Operating systems	UNIX, Windows, Linux, CISCO IOS
Networks	Support required for different networks: Ethernet, Infiniband (Appendix F)
Programming languages	Languages (ANSI C, C++, Java, Fortran) affect instruction set (Appendix A)

Figure 1.8 Summary of some of the most important functional requirements an architect faces. The left-hand column describes the class of requirement, while the right-hand column gives specific examples. The right-hand column also contains references to chapters and appendices that deal with the specific issues.



Figure 1.9 Log-log plot of bandwidth and latency milestones in Figure 1.10 relative to the first milestone. Note that latency improved 8–91 ×, while bandwidth improved about 400–32,000 ×. Except for networking, we note that there were modest improvements in latency and bandwidth in the other three technologies in the six years since the last edition: 0%–23% in latency and 23%–70% in bandwidth. Updated from Patterson, D., 2004. Latency lags bandwidth. Commun. ACM 47 (10), 71–75.

Microprocessor	16-Bit address/ bus, microcoded	32-Bit address/ bus, microcoded	5-Stage pipeline, on-chip I & D caches, FPU	2-Way superscalar, 64-bit bus	Out-of-order 3-way superscalar	Out-of-order superpipelined, on-chip L2 cache	Multicore OOO 4-way on chip L3 cache, Turbo
Product	Intel 80286	Intel 80386	Intel 80486	Intel Pentium	Intel Pentium Pro	Intel Pentium 4	Intel Core i7
Year	1982	1985	1989	1993	1997	2001	2015
Die size (mm ²)	47	43	81	90	308	217	122
Transistors	134,000	275,000	1,200,000	3,100,000	5,500,000	42,000,000	1,750,000,000
Processors/chip	1	1	1	1	1	1	4
Pins	68	132	168	273	387	423	1400
Latency (clocks)	6	5	5	5	10	22	14
Bus width (bits)	16	32	32	64	64	64	196
Clock rate (MHz)	12.5	16	25	66	200	1500	4000
Bandwidth (MIPS)	2	6	25	132	600	4500	64,000
Latency (ns)	320	313	200	76	50	15	4
Memory module	DRAM	Page mode DRAM	Fast page mode DRAM	Fast page mode DRAM	Synchronous DRAM	Double data rate SDRAM	DDR4 SDRAM
Module width (bits)	16	16	32	64	64	64	64
Year	1980	1983	1986	1993	1997	2000	2016
Mbits/DRAM chip	0.06	0.25	1	16	64	256	4096
Die size (mm ²)	35	45	70	130	170	204	50
Pins/DRAM chip	16	16	18	20	54	66	134
Bandwidth (MBytes/s)	13	40	160	267	640	1600	27,000
Latency (ns)	225	170	125	75	62	52	30
Local area network	Ethernet	Fast Ethernet	Gigabit Ethernet	10 Gigabit Ethernet	100 Gigabit Ethernet	400 Gigabit Ethernet	
IEEE standard	802.3	803.3u	802.3ab	802.3ac	802.3ba	802.3bs	
Year	1978	1995	1999	2003	2010	2017	
Bandwidth (Mbits/seconds)	10	100	1000	10,000	100,000	400,000	
Latency (µs)	3000	500	340	190	100	60	
Hard disk	3600 RPM	5400 RPM	7200 RPM	10,000 RPM	15,000 RPM	15,000 RPM	
Product	CDC WrenI 94145-36	Seagate ST41600	Seagate ST15150	Seagate ST39102	Seagate ST373453	Seagate ST600MX0062	
Year	1983	1990	1994	1998	2003	2016	
Capacity (GB)	0.03	1.4	4.3	9.1	73.4	600	
Disk form factor	5.25 in.	5.25 in.	3.5 in.	3.5 in.	3.5 in.	3.5 in.	
Media diameter	5.25 in.	5.25 in.	3.5 in.	3.0 in.	2.5 in.	2.5 in.	
Interface	ST-412	SCSI	SCSI	SCSI	SCSI	SAS	
Bandwidth (MBytes/s)	0.6	4	9	24	86	250	
Latency (ms)	48.3	17.1	12.7	8.8	5.7	3.6	

Figure 1.10 Performance milestones over 25–40 years for microprocessors, memory, networks, and disks. The microprocessor milestones are several generations of IA-32 processors, going from a 16-bit bus, microcoded 80286 to a 64-bit bus, multicore, out-of-order execution, superpipelined Core i7. Memory module milestones go from 16-bit-wide, plain DRAM to 64-bit-wide double data rate version 3 synchronous DRAM. Ethernet advanced from 10 Mbits/s to 400 Gbits/s. Disk milestones are based on rotation speed, improving from 3600 to 15,000 RPM. Each case is best-case bandwidth, and latency is the time for a simple operation assuming no contention.

Updated from Patterson, D., 2004. Latency lags bandwidth. Commun. ACM 47 (10), 71–75.



Figure 1.11 Growth in clock rate of microprocessors in Figure 1.1. Between 1978 and 1986, the clock rate improved less than 15% per year while performance improved by 22% per year. During the "renaissance period" of 52% performance improvement per year between 1986 and 2003, clock rates shot up almost 40% per year. Since then, the clock rate has been nearly flat, growing at less than 2% per year, while single processor performance improved recently at just 3.5% per year.



Figure 1.12 Energy savings for a server using an AMD Opteron microprocessor, 8 GB of DRAM, and one ATA disk. At 1.8 GHz, the server can handle at most up to two-thirds of the workload without causing service-level violations, and at 1 GHz, it can safely handle only one-third of the workload (Figure 5.11 in Barroso and Hölzle, 2009).

		F	Relativ	ve energ	jy cost			Rela	tive a	irea co	st
Operation:	Energy (pJ)						Area (µm²)				
8b Add	0.03						36				
16b Add	0.05						67				
32b Add	0.1						137				
16b FB Add	0.4						1360				
32b FB Add	0.9		-				4184				
8b Mult	0.2						282				
32b Mult	3.1		-				3495				
16b FB Mult	1.1						1640				
32b FB Mult	3.7		_				7700				
32b SRAM Read (8KB)	5						N/A				
32b DRAM Read	640						N/A]			
		1	10	100	1000	10000)	1 .	10	100	1000

Energy numbers are from Mark Horowitz *Computing's Energy problem (and what we can do about it)*. ISSCC 2014 Area numbers are from synthesized result using Design compiler under TSMC 45nm tech node. FP units used DesignWare Library.

Figure 1.13 Comparison of the energy and die area of arithmetic operations and energy cost of accesses to SRAM and DRAM. [Azizi][Dally]. Area is for TSMC 45 nm technology node.



Figure 1.14 Photograph of an Intel Skylake microprocessor die, which is evaluated in Chapter 4.



Figure 1.15 The components of the microprocessor die in Figure 1.14 are labeled with their functions.



Figure 1.16 This 200 mm diameter wafer of RISC-V dies was designed by SiFive. It has two types of RISC-V dies using an older, larger processing line. An FE310 die is 2.65 mm × 2.72 mm and an SiFive test die that is 2.89 mm × 2.72 mm. The wafer contains 1846 of the former and 1866 of the latter, totaling 3712 chips.

	SPEC2017	SPEC2006	SPEC2000	SPEC95	SPEC92	SPEC89
GNU C compiler	•					_ gcc
Perl interpreter	-			- perl]	espresso
Route planning	4		- mcf		-	li
General data compression	XZ		bzip2		compress	eqntott
Discrete Event simulation - computer network	-	 omnetpp 	vortex	go	sc	
XML to HTML conversion via XSLT	-	 xalancbmk 	gzip	ijpeg		-
Video compression	X264	h264ref	eon	m88ksim		
Artificial Intelligence: alpha-beta tree search (Chess)	deepsjeng	sjeng	twolf			
Artificial Intelligence: Monte Carlo tree search (Go)	leela	gobmk	vortex			
Artificial Intelligence: recursive solution generator (Sudoku)	exchange2	astar	vpr			
		hmmer	crafty			
		libquantum	parser			
Explosion modeling	-	— bwaves				fpppp
Physics: relativity		cactuBSSN				tomcatv
Molecular dynamics	-	namd]	doduc
Ray tracing	•	povray				nasa7
Fluid dynamics	-	Ibm				spice
Weather forecasting	•	wrf			swim	matrix300
Biomedical imaging: optical tomography with finite elements	parest	gamess	1	apsi	hydro2d	
3D rendering and animation	blender			mgrid	su2cor	
Atmosphere modeling	cam4	milc	wupwise	applu	wave5	
Image manipulation	imagick	zeusmp	apply	turb3d	J	
Molecular dynamics	nab	gromacs	galgel			
Computational Electromagnetics	fotonik3d	leslie3d	mesa			
Regional ocean modeling	roms	dealll	art			
		soplex	equake			
		calculix	racerec			
		GemsFDTD	lucas			
		tonto	fma3d			
		sphinx3	sixtrack			

Benchmark name by SPEC generation

Figure 1.17 SPEC2017 programs and the evolution of the SPEC benchmarks over time, with integer programs above the line and floating-point programs below the line. Of the 10 SPEC2017 integer programs, 5 are written in C, 4 in C++., and 1 in Fortran. For the floating-point programs, the split is 3 in Fortran, 2 in C++, 2 in C, and 6 in mixed C, C++, and Fortran. The figure shows all 82 of the programs in the 1989, 1992, 1995, 2000, 2006, and 2017 releases. Gcc is the senior citizen of the group. Only 3 integer programs and 3 floating-point programs survived three or more generations. Although a few are carried over from generation to generation, the version of the program changes and either the input or the size of the benchmark is often expanded to increase its running time and to avoid perturbation in measurement or domination of the execution time by some factor other than CPU time. The benchmark descriptions on the left are for SPEC2017 only and do not apply to earlier versions. Programs in the same row from different generations of SPEC are generally not related; for example, fpppp is not a CFD code like bwaves.

Category	Name	Measures performance of			
Cloud	Cloud_IaaS 2016	Cloud using NoSQL database transaction and K-Means clustering using map/reduce			
CPU	CPU2017	Compute-intensive integer and floating-point workloads			
	SPECviewperf [®] 12	3D graphics in systems running OpenGL and Direct X			
Graphics and workstation performance	SPECwpc V2.0	Workstations running professional apps under the Windows OS			
	SPECapcSM for 3ds Max 2015 TM	3D graphics running the proprietary Autodesk 3ds Max 2015 app			
	SPECapcSM for Maya [®] 2012	3D graphics running the proprietary Autodesk 3ds Max 2012 app			
	SPECapcSM for PTC Creo 3.0	3D graphics running the proprietary PTC Creo 3.0 app			
	SPECapcSM for Siemens NX 9.0 and 10.0	3D graphics running the proprietary Siemens NX 9.0 or 10.0 app			
	SPECapcSM for SolidWorks 2015	3D graphics of systems running the proprietary SolidWorks 2015 CAD/CAM app			
	ACCEL	Accelerator and host CPU running parallel applications using OpenCL and OpenACC			
High performance computing	MPI2007	MPI-parallel, floating-point, compute-intensive programs running on clusters and SMPs			
	OMP2012	Parallel apps running OpenMP			
Java client/server	SPECjbb2015	Java servers			
Power	SPECpower_ssj2008	Power of volume server class computers running SPECjbb2015			
Solution File	SFS2014	File server throughput and response time			
Server (SFS)	SPECsfs2008	File servers utilizing the NFSv3 and CIFS protocols			
Virtualization	SPECvirt_sc2013	Datacenter servers used in virtualized server consolidation			

Figure 1.18 Active benchmarks from SPEC as of 2017.

Benchmarks	Sun Ultra Enterprise 2 time (seconds)	AMD A10- 6800K time (seconds)	SPEC 2006Cint ratio	Intel Xeon E5-2690 time (seconds)	SPEC 2006Cint ratio	AMD/Intel times (seconds)	Intel/AMD SPEC ratios
perlbench	9770	401	24.36	261	37.43	1.54	1.54
bzip2	9650	505	19.11	422	22.87	1.20	1.20
gcc	8050	490	16.43	227	35.46	2.16	2.16
mcf	9120	249	36.63	153	59.61	1.63	1.63
gobmk	10,490	418	25.10	382	27.46	1.09	1.09
hmmer	9330	182	51.26	120	77.75	1.52	1.52
sjeng	12,100	517	23.40	383	31.59	1.35	1.35
libquantum	20,720	84	246.08	3	7295.77	29.65	29.65
h264ref	22,130	611	36.22	425	52.07	1.44	1.44
omnetpp	6250	313	19.97	153	40.85	2.05	2.05
astar	7020	303	23.17	209	33.59	1.45	1.45
xalancbmk	6900	215	32.09	98	70.41	2.19	2.19
Geometric mean			31.91		63.72	2.00	2.00

Figure 1.19 SPEC2006Cint execution times (in seconds) for the Sun Ultra 5—the reference computer of SPEC2006—and execution times and SPECRatios for the AMD A10 and Intel Xeon E5-2690. The final two columns show the ratios of execution times and SPEC ratios. This figure demonstrates the irrelevance of the reference computer in relative performance. The ratio of the execution times is identical to the ratio of the SPEC ratios, and the ratio of the geometric means (63.7231.91/20.86 = 2.00) is identical to the geometric mean of the ratios (2.00). Section 1.11 discusses libquantum, whose performance is orders of magnitude higher than the other SPEC benchmarks.

	Syster	tem 1 System 2		System 3		
Component		Cost (% Cost)		Cost (% Cost)		Cost (% Cost)
Base server	PowerEdge R710	\$653 (7%)	PowerEdge R815	\$1437 (15%)	PowerEdge R815	\$1437 (11%)
Power supply	570 W		1100 W		1100 W	
Processor	Xeon X5670	\$3738 (40%)	Opteron 6174	\$2679 (29%)	Opteron 6174	\$5358 (42%)
Clock rate	2.93 GHz		2.20 GHz		2.20 GHz	
Total cores	12		24		48	
Sockets	2		2		4	
Cores/socket	6		12		12	
DRAM	12 GB	\$484 (5%)	16 GB	\$693 (7%)	32 GB	\$1386 (11%)
Ethernet Inter.	Dual 1-Gbit	\$199 (2%)	Dual 1-Gbit	\$199 (2%)	Dual 1-Gbit	\$199 (2%)
Disk	50 GB SSD	\$1279 (14%)	50 GB SSD	\$1279 (14%)	50 GB SSD	\$1279 (10%)
Windows OS		\$2999 (32%)		\$2999 (33%)		\$2999 (24%)
Total		\$9352 (100%)		\$9286 (100%)		\$12,658 (100%)
Max ssj_ops	910,978		926,676		1,840,450	
Max ssj_ops/\$	97		100		145	

Figure 1.20 Three Dell PowerEdge servers being measured and their prices as of July 2016. We calculated the cost of the processors by subtracting the cost of a second processor. Similarly, we calculated the overall cost of memory by seeing what the cost of extra memory was. Hence the base cost of the server is adjusted by removing the estimated cost of the default processor and memory. Chapter 5 describes how these multisocket systems are connected together, and Chapter 6 describes how clusters are connected together.



Figure 1.21 Power-performance of the three servers in Figure 1.20. Ssj_ops/watt values are on the left axis, with the three columns associated with it, and watts are on the right axis, with the three lines associated with it. The horizontal axis shows the target workload, as it varies from 100% to Active Idle. The single node R630 has the best ssj_ops/watt at each workload level, but R730 consumes the lowest power at each level.



Figure 1.22 Predictions of logic transistor dimensions from two editions of the ITRS report. These reports started in 2001, but 2015 will be the last edition, as the group has disbanded because of waning interest. The only companies that can produce state-of-the-art logic chips today are GlobalFoundaries, Intel, Samsung, and TSMC, whereas there were 19 when the first ITRS report was released. With only four companies left, sharing of plans was too hard to sustain. From IEEE Spectrum, July 2016, "Transistors will stop shrinking in 2021, Moore's Law Roadmap Predicts," by Rachel Courtland.



Figure 1.23 Relative bandwidth for microprocessors, networks, memory, and disks over time, based on data in Figure 1.10.



Figure 1.24 Percentage of peak performance for four programs on four multiprocessors scaled to 64 processors. The Earth Simulator and X1 are vector processors (see Chapter 4 and Appendix G). Not only did they deliver a higher fraction of peak performance, but they also had the highest peak performance and the lowest clock rates. Except for the Paratec program, the Power 4 and Itanium 2 systems delivered between 5% and 10% of their peak. From Oliker, L., Canning, A., Carter, J., Shalf, J., Ethier, S., 2004. Scientific computations on modern parallel vector systems. In: Proc. ACM/IEEE Conf. on Supercomputing, November 6–12, 2004, Pittsburgh, Penn., p. 10.

Appendix	Title
A	Instruction Set Principles
В	Review of Memory Hierarchies
С	Pipelining: Basic and Intermediate Concepts
D	Storage Systems
Е	Embedded Systems
F	Interconnection Networks
G	Vector Processors in More Depth
Н	Hardware and Software for VLIW and EPIC
Ι	Large-Scale Multiprocessors and Scientific Applications
J	Computer Arithmetic
K	Survey of Instruction Set Architectures
L	Advanced Concepts on Address Translation
М	Historical Perspectives and References

Figure 1.25 List of appendices.

Chip	Die Size (mm ²)	Estimated defect rate (per cm ²)	N	Manufacturing size (nm)	Transistors (billion)	Cores
BlueDragon	180	0.03	12	10	7.5	4
RedDragon	120	0.04	14	7	7.5	4
Phoenix ⁸	200	0.04	14	7	12	8

Figure 1.26 Manufacturing cost factors for several hypothetical current and future processors.

System	Chip	TDP	Idle power	Busy power
General-purpose	Haswell E5-2699 v3	504 W	159 W	455 W
Graphics processor	NVIDIA K80	1838 W	357 W	991 W
Custom ASIC	TPU	861 W	290 W	384 W

Figure 1.27 Hardware characteristics for general-purpose processor, graphical processing unit-based or custom ASIC-based system, including measured power (cite ISCA paper).

Guetarra		Throughput			% Max IPS		
System	Cnip	Α	В	с	Α	В	с
General-purpose	Haswell E5-2699 v3	5482	13,194	12,000	42%	100%	90%
Graphics processor	NVIDIA K80	13,461	36,465	15,000	37%	100%	40%
Custom ASIC	TPU	225,000	280,000	2000	80%	100%	1%

Figure 1.28 Performance characteristics for general-purpose processor, graphical processing unit-based or custom ASIC-based system on two neural-net workloads (cite ISCA paper). Workloads A and B are from published results. Workload C is a fictional, more general-purpose application.