

Computer Organization and Design: The Hardware/Software Interface
 David A. Patterson and John L. Hennessy
 Errata list as of 9/01/14

Chapter	Page #	Description	Correction
Preface	xviii	On page xviii, 9th line, the text says "Appendix C on logic design..."	The Basics of Logic Design has been moved to Appendix B .
Preface	xviii	On page xviii, 14th line, the text says "this firth edition"	Should read "this fifth edition"
Preface	xix	On page xix, the first subheading says "Changes for the Fifth Edition"	Subheading should read "Instructor Support"
1	10		
1	12	In the "Performance via Prediction" section, the first sentence incorrectly states that "the final great idea is prediction."	The sentence should read: "...the next great idea is prediction."
1	40	On page 40, the 2nd line of the caption for FIGURE 1.16, refers to "25 years"	Caption should say " 30 years" to be consistent with statement in the first sentence of Section 1.7.
1	55	Exercise 1.6 asks "which implementation is faster?"	For clarity, replace "which implementation is faster?" with " Which is faster: P1 or P2? "
1	58	Exercise 1.12.4 states: "Assume that 40% of the instructions executed on both P1 and P2 are floating-point instructions. Find the MFLOPS figures for the programs ."	P1 and P2 are processors. Therefore, "programs" should read " processors ".
1	59	On page 59, third line from the bottom (in the "Answers to Check Yourself"), it reads" §1.6, page 33: 1. a: both, b: latency, c: neither. 7 seconds.	There are two sub-questions, 1 and 2, for this problem. Therefore there should be a "2." preceding "7 seconds": §1.6, page 33: 1. a: both, b: latency, c: neither. 2. 7 seconds.
1	59	On page 59, Question 1.14 reads: "Assume a program requires the execution of 50 x 106 FP instructions, 110 x 106 INT instructions, 80 x 106 L/S instructions, and 16 x 106 branch instructions."	The "6" in each string of "106" should be superscript: "Assume a program requires the execution of 50 x 10 ⁶ FP instructions, 110 x 10 ⁶ INT instructions, 80 x 10 ⁶ L/S instructions, and 16 x 10 ⁶ branch instructions."
2	128	In the "Answer" section of the "Linking Object Files" example, the paragraphs above and below the table are identical.	Second instance of the paragraph should be removed.
2	135	In the section " The Procedure sort " the code example: The inner for loop does (j=i-1; j>=0 && v[j]>v[j+1]; j =1), though the latter should surely be j+=1 (or j++, j = j+1) - since otherwise the same comparison is made every time, for all i (the outer loop), so it may as well execute once. I'd suggest having made that change, that it's more readable to check v[j] > v[i], since j is defined i-1, odd to then use i-1+1.	[Needs confirmation]

2	164	Exercise 2.2: Problem statement reads "For the following MIPS assembly instructions above..."	The problem statement should read "For the following MIPS assembly instructions below ..."
2	165	Exercise 2.5: Problem statement reads "For the MIPS assembly instructions in Exercise 2.4, rewrite the assembly code to minimize the number of MIPS instructions..."	The problem statement should read "For the MIPS assembly instructions in Exercise 2.4, rewrite the assembly code to minimize the number of MIPS instructions..."
2	168	Exercises 2.18.1 and 2.18.2: Problem statement reads "How this would this affect the size of ..."	The problem statement should read "How would this affect the size of ..."
2	171	Exercise 2.31: At the end of the code the closing parentheses ")" is missing.	There should be one more line that has only ")", with it leftmost aligned. Similar to code in Question 2.34
2	171	Exercise 2.34: Relative time rating and reference section are missing.	Relative time and reference section should read: 2.34 [5] \$2.8
2	172	Exercise 2.37, 8th line reads: "50ten, 52ten, 0ten"	The 8th line should read ("ten" in subscript): $50_{ten}, 52_{ten}, 0_{ten}$
2	173	Exercise 2.42 reads "If the current value of the PC is 0x1FFFF000, ..." The value string "F" and "E" are mixed.	All the <code>FS</code> should be capital, ie: "0x1 FFFF 000"

3	206	On page 206, the 4th line reads: "Since $127 \geq +4 \geq -126$, there is no overflow or underflow."	The "+4" should be "-4", since the exponent of the final result above is "-4". Therefore the line should read: "Since $127 \geq -4 \geq -126$, there is no overflow or underflow."
3	223	Figure 3.19, the 4th and 5th lines, in the "Data transfer" column, read: "VLD(1,2,3,4)" and "VST(1,2,3,4)" The punctuation between "3" and "4" is a period (.)	The 4th and 5th lines should read: VLD(1,2,3,4) and VST(1,2,3,4)
3	225	The 8th line reads "instruction to perform two 64-bit floating-point multiplies." However, the 9th and 11th lines show these instructions: "addpd" and "vaddpd" Therefore, the text and the example instructions do not correspond.	The 8th line should read: "instruction to perform two 64-bit floating-point additions "
3	225	Elaboration section, 3rd line reads: <pre>vaddpd %ymm0, %ymm1, %ymm4 # %ymm4 = %ymm1 + %ymm2</pre> The variable names in the comment portion do not correspond to those in the real expression part.	Comment portion should read: <pre># %ymm4 = %ymm0 + %ymm1</pre>
3	226	Figure 3.21 caption reads: "Unoptimized C version..." Figure 3.22 caption reads: "...the <u>optimized</u> C code in Figure 3.21"	Figure 3.22 caption should read: "...the <u>unoptimized</u> C code in Figure 3.21"
3	226	Figure 3.22, 11th line reads: <pre>11. add \$0x1, %r11d # register %r11 = %r11 + 1</pre>	In the 11th line, "%r11d" should be " %r11 ": <pre>11. add \$0x1, %r11 # register %r11 = %r11 + 1</pre>
	229	On page 229, 6th line reads: "Recall that a binary number c , where x_i means the i th bit, represents the number"	The 6th line should read: "Recall that a binary number x , where x_i means the i th bit, represents the number"

3	229	The first equation reads: $\dots + (x^3 \times 2^3) + (x^2 \times 2^2) + (x1 \times 2^1) + (x0 \times 2^0)$	The first equation should read: $\dots + (x3 \times 2^3) + (x2 \times 2^2) + (x1 \times 2^1) + (x0 \times 2^0)$
3	229	The 8th line reads: "Shifting the bits of c right by n bits would seem to be the same as dividing by..."	The 8th line should read: "Shifting the bits of x right by n bits would seem to be the same as dividing by..."
3	229	The 9th line reads: " $2n \dots$ "	The 9th line should read: " $2^n \dots$ "
3	240	Exercises 3.35, 3.36, 3.37 read "... (3.41796875 $10^{-3} \dots$)"	Exercises 3.35, 3.36, 3.37 should read "... (3.41796875 $\times 10^{-3} \dots$)"
3	241	Exercises 3.43, 3.44, 3.45, 3.46. In Exercise 3.43, "value 1/3" is indicated. In the other exercises, it is not.	In exercises 3.45 and 3.46 should be re-worded to include "the fraction 1/3"
3	241	Exercise 3.47, 6th line (just below the equation) reads: "Assume you are to write an optimized implementation this code ..."	Exercise 3.47, 6th line (just below the equation) should read: "Assume you are to write an optimized implementation of this code ..."
4	277	On page 277, the 7th line from the bottom (a sub-heading) reads: Hazards	The sub-heading should read: Structural Hazards

4	282	<p>On page 282, the 1st sentence of the ANSWER section reads:</p> <p>"Figure 3.27 in Chapter 3 shows that branches are 17% of the instructions executed in SPECint2006."</p>	<p>The 1st sentence of the ANSWER section should read:</p> <p>"Figure 3.28 in Chapter 3 shows that branches are 17% of the instructions executed in SPECint2006."</p>
4	286	<p>On page 286, last paragraph of Section 4.5 is:</p> <p>"Instruction sets can either simplify or make life harder for pipeline designers, who must already cope with structural, control, and data hazards. Branch prediction and forwarding help make a computer fast while still getting the right answers."</p>	<p>The last paragraph of Section 4.5 should be included in the "BIG Picture" section that appears directly above it.</p>
4	291	<p>Figure 4.36, upper left, the "PC" box is not highlighted.</p>	<p>The "PC" box should be highlighted, since it is active.</p>
4	291	<p>In Figures 4.36, 4.38, and 4.40, there are two copies of the data path in each figure. In the top copy, the MUX after the MEM/WB pipelined registers has its 0,1 inputs swapped; in the bottom copy of the data path, the 0,1 inputs to this MUX are correct.</p>	<p>In the top data path, the MUX after the MEM/WB pipelined registers should read 1,0.</p>
4	328	<p>On page 328, 4 lines above EXAMPLE, reads:</p> <p>"In reality, we save the address +4, so the exception handling the software routine must first subtract 4 from the saved value."</p>	<p>The sentence should read (delete "the software"):</p> <p>"In reality, we save the address +4, so the exception handling routine must first subtract 4 from the saved value."</p>
4	338	<p>On page 338, the definition in the upper left margin reads:</p> <p>loop unrolling A technique to get more performance from loops that access arrays, in which multiple copies of the loop body are made and instructions from different iterations are"</p>	<p>The definition should read ("scheduled together" is missing from the conclusion of the sentence):</p> <p>loop unrolling A technique to get more performance from loops that access arrays, in which multiple copies of the loop body are made and instructions from different iterations are scheduled together."</p>
4	344	<p>Figure 4.73 states that the Intel Core i5 Nehalem has "1" Cores/Chip.</p>	<p>The Cores/Chip for the Core i5 Nehalem should read "2-4"</p>
4	348	<p>On page 348, Figure 4.77, bottom right corner of figure, in the rectangle 1 line up from bottom right, reads:</p> <p>"256 KB unified L2 cache (eight-way)"</p>	<p>The text should read ("L2" instead of "l2"):</p> <p>"256 KB unified L2 cache (eight-way)"</p>
4	351	<p>On page 351, 6 and 5th lines from bottom, reads:</p> <p>"We surround each intrinsic with a simple for loop that 4 iterations (lines 9, 14, and 20) ..."</p>	<p>The 6 and 5th lines from bottom should read ("loop with" rather than "loop that"):</p> <p>"We surround each intrinsic with a simple for loop with 4 iterations (lines 9, 14, and 20) ..."</p>
4	353	<p>On page 352, 9 to 7th lines from bottom, read:</p> <p>"Optimizations for subword parallelism and instruction level parallelism result in an overall speedup of 8.8 versus the unoptimized DGEMM in Figure 3.21."</p>	<p>In Figure 4.82, the GFLOPS values for unoptimized and AVX+unroll are 1.7 and 14.6, respectively. The speedup rate between them is therefore about 8.59.</p> <p>The sentence should read:</p> <p>"Optimizations for subword parallelism and instruction level parallelism result in an overall speedup of 8.59 versus the unoptimized DGEMM in Figure 3.21."</p>

4	353	<p>On page 353, Figure 4.81 shows:</p> <pre> 9 vmulpd (%rax), %ymm0, %ymm5 # Parallel mul %ymm1, 4 A ... 11 vmulpd 0x20(%rax), %ymm0, %ymm5 # Parallel mul %ymm1, 4 A elements ... 13 vmulpd 0x40(%rax), %ymm0, %ymm5 # Parallel mul %ymm1, 4 A elements 14 vmulpd 0x60(%rax), %ymm0, %ymm0 # Parallel mul %ymm1, 4 A elements </pre> <p>In these lines, the suffix number for the variable %ymm in the comment part differs from that in the command part.</p>	<p>The comments for these 4 lines should read</p> <pre> # Parallel mul %ymm0, 4 A elements </pre>
4	353	<p>On page 352, Figure 4.81 shows:</p> <pre> 19 jne 68 <dgemm+0x68> # jump if not %r8 != %rax </pre>	<p>"not" and "!=" are redundant. The line should read:</p> <pre> 19 jne 68 <dgemm+0x68> # jump if %r8 %rax </pre>
4	369	<p>On page 369, Question 4.19 reads:</p> <p>4.19.1 [10] <§§4.3, 4.6, 4.14></p>	<p>The reference section 4.14 in this question should read 4.15.</p> <p>The same applies to sub-questions that follow.</p>
5	386	<p>On page 386, lower part, in the table without Figure No.:</p> <p>In the "Hit or miss in cache" column, reference is made to (5.6b), (5.6c), and so on.</p>	<p>Figure 5.9 is the actual reference. Therefore (5.6b), (5.6c), etc. should be (5.9b), (5.9c), etc.</p>
5	414	<p>On page 414, Figure 5.20, Names of Arrays:</p> <p>In the body description, C, A, and B are used as the names of arrays. However, in this figure, x, y, or z appears above the top left corner of each square.</p>	<p>Array names should be: x to C, y to A, and z to B</p>
5	414	<p>On page 414, Figure 5.20, 3rd line of caption reads:</p> <p>Compared to Figure 5.21, elements ...</p>	<p>The 3rd line should read:</p> <p>Compared to Figure 5.22, elements ...</p>
5	414	<p>On page 414, Figure 5.20, last line of caption reads:</p> <p>... new elements of x.</p>	<p>The last line of the caption should read:</p> <p>... new elements of C.</p>

5	414	<p>On page 415, Figure 5.22, Names of Arrays</p> <p>In the body description, C, A, and B are used as the names of arrays. However, in this figure, x, y, or z appears above the top left corner of each square.</p>	<p>Array names should be: x to C, y to A, and z to B</p>
5	422	<p>On page 422, 2nd line from end of ANSWER reads:</p> <p>Parity bits 2 and 10 are incorrect.</p>	<p>In the EXAMPLE description, parity bit 8 is said to have an error. Therefore, "10" should be "8".</p> <p>The text should read: "Parity bits 2 and 8 are incorrect"</p>
5	462	<p>On page 462, 4th line reads:</p> <p>32-byte addresses</p>	<p>The 4th line should read:</p> <p>"32-bit addresses".</p>
5	472	<p>On page 472, Figure 5.44, heading of right most column, reads:</p> <p>Intel Nehalem</p>	<p>The heading of right most column should read:</p> <p>Intel Core i7</p>
5	489	<p>Page 489, Question 5.9.3</p> <p>The relative time rating and the reference section are missing.</p>	<p>The relative time rating and the reference section are the same as questions 5.9.1 and 5.9.2, ie:</p> <p>[5] <5.5></p>
6	533	<p>On page 533, 5th line in the code in the "Hardware/Software Interface" reads:</p> <pre>EmitIntermediate(w, "1"); // Produce list of all words reduce(String key, Iterator values):</pre>	<p>The 5th line in the code in the "Hardware/Software Interface" should read (line break inserted after "words"):</p> <pre>EmitIntermediate(w, "1"); // Produce list of all words reduce(String key, Iterator values):</pre>
6	533	<p>On page 533, 2nd line from bottom in the "Hardware/Software Interface" reads:</p> <pre>result += ParseInt(v); // get integer from key-value pair</pre>	<p>In the first paragraph on the next page, <code>ParseInt()</code> is explained to get the number of occurrences per word in all documents. To avoid confusion, change "integer" to this descriptive expression:</p> <pre>result += ParseInt(v); // get integer value from key-value pair // value represents number of occurrences per word in all documents</pre>

