

Computer Architecture: A Quantitative Approach, 3ed  
 John L. Hennessy and David A. Patterson  
 Errata for the 4th Printing

<b>Chapter</b>	<b>Page #</b>	<b>Description</b>
<b>Front Cover</b>	xxiv	Email address is "ca3bugs@mkp.com" The email address is listed on the 6th page inside the front cover as "cabugs@mkp.com"
	xxv	"two versions of this edition" --> "two versions of the second edition"
	xxviii	Emilio Salgueiro, Unysis --> Emilio Salgueiro, Unisys
1	19	"If alpha corresponds inversely to the number of masking levels" --> "If alpha corresponds <i>directly</i> to the number of masking levels"
1	28	last line: "11 integer benchmarks" --> "12 integer benchmarks"
1	30	figure 1.12 caption: "11 integer programs" --> "12 integer programs"
1	30	Figure 1.12: "perlmbk" --> "perlbmk"
1	32	Figure 1.13, row 6 (Telecommunications): "6" --> "5" (for number of kernels; the EEMBC Telecom suite has only 5 benchmarks)
1	32	Figure 1.13 caption: "consisting of 24 kernels" --> "consisting of 23 kernels"
1	37	Figure 1.6, caption "Programs" --> "Computers"
1	37	"Geometric means also have a nice property for two samples $X_i$ and $Y_i$ " --> "Geometric means also have a nice property for two samples $X$ and $Y$ of length $n$ : Geometric mean ( $X$ )/Geometric mean ( $Y$ ) = Geometric mean ( $(X_1/Y_1), (X_2/Y_2), \dots (X_n/Y_n)$ )"
1	37	In the formula for weightings: the summation from "i=1 to n of 1/Time_i" --> "from j=1 to n of 1/Time_j"
1	43	Line 18: change the numerator from "Instruction Count x Clock Cycle Time" to "Instruction Count x Cycles per Instruction"
1	44	Line 8 from top: "...number of instructions per clock for instruction I:" --> "...number of clocks per instruction for instruction I"
1	44	Example, Frequency of FP operations: Remove "(other than EPSQR)"
1	45	Answer, third equation: "1.625"--> "1.62"
1	52	Figure 1.21: (To concur with figure 1.22): Fujitsu PRIMEPOWER 20000 --> Fujitsu PRIMEPOWER 2000
1	52	Figure 1.21 caption, second line: "exSeries" --> "xSeries"
1	53	Figure 1.22 "Fujitsu PRIMEPOWER 20000" --> "Fujitsu PRIMEPOWER 2000"
1	53	Figure 1.23: "Transactions per minute (thouands)"-->"(thousands)"
1	54	Line 4 of first full paragraph, "AMD Elan" --> "AMD K6-2E+"
1	54	"PowerPC 650" --> "PowerPC 750"
1	56	Third full paragraph, line 3: "with a typical power consumption of 9.3 W" --> "with a typical power consumption of 9.6 W"
1	57	Figure 1.28, line 4: "how much faster a Pentium 4 at 1.7 GHz would be than a 1 Ghz Pentium III"-->"1.7 GHz Pentium III"
1	57	Figure 1.28, line 6: "approximation to how fast a P3 would run"-->"how fast such a Pentium III would run"
1	61	SPEC95 for eqntott should read "dropped," and the SPEC92 column should either contain a "modified" entry or be left blank
1	70	Line 10 in Commercial Developments section: "Computer Museum" --> "Computer History Museum" (they are two different places)
1	73	Add reference: Kembel, R. [2000]. "Fibre Channel: A Comprehensive Introduction." Internet Week, April 2000.
1	78	Exercise 1.8, part e: Digital 21064C --> Alpha 21264C (to match with Figure 1.34)
1	80	Exercise 1.13, part b: "What are the harmonic means (see Exercise 1.10 for the definition of harmonic mean) of the two sets of measurements"-->"what are the geometric means of the two sets of measurements"
1	80	Exercise 1.13, part b: "Which outlying data point affects the harmonic mean more"-->"Which outlying data point affects the geometric mean more"

1	80	Exercise 1.13, part c: "Which mean, arithmetic or harmonic"-->"Which mean, arithmetic or geometric"	
1	80	Exercise 1.13, part d: "How representative of the entire set for the arithmetic and harmonic mean statistics"-->"do the arithmetic and geometric mean statistics"	
1	82	1.16 remove the line "Only one enhancement is usable at a time"	
2	94	Figure 2.3, "Trimedia TM5200" --> "TM32"	
2	97	Figure 2.5, one of the misaligned examples includes a misaligned object that isn't shaded	
2	99	Delete last sentence in caption, Figure 2.7 (TeX is not a SPEC89 program)	
2	101	Line 3: "the general classes of integer operations in an instruction set" --> "the general classes of integer and floating-point operations in an instruction set"	
2	102	Figure 2.10 caption, add after last line: "Thus, 16 bits would capture about 80% and 8 bits about 50%."	
2	103	Last line of third full paragraph: "the use of novel addressing mode is sparse" --> "the use of novel addressing <i>modes</i> is sparse"	
2	103	line 13 from bottom: "used in our SPEC measurements" --> "used in our measurements" (TeX is not a SPEC program)	
2	103	"As the caption in Figure 2.10 suggests, these sizes would capture 50% to 80% of the immediates" --> this claim is not substantiated by the captions of the figure to which it refers	
2	104	Figure 2.11, line 9: "**ARx%"-->"*+ARx%"	
2	109	"Chapter 6 and Appendix F describe the full machines that pioneered these architectures" --> "Chapter 6 and Appendix G describe the computers that pioneered these architectures" (2 changes)	
2	111	Third full paragraph, line 11: addressees" --> "addresses"	
2	120	Line 6 of Summary section: "11 examples"-->"13 examples" (there are 10 examples in Appendix C and one each in Appendix D, E, and F	
2	127	The first sentence says "The missing elegance from these architectures," but the second sentence says "By making the vector width variable..." and then proceeds to explain how using a fixed vector length and variable vector width allows these SIMD instructions to "seamlessly switch between different data widths simply by increasing the number of elements per vector" where the vectors have a fixed length. The last sentence on page 127 gives an example using a variable vector length, which would agree with the first sentence but not the second. Suggest: Delete the first sentence; then change last sentence to "One generation might have 64 32-bit elements per vector register, and the next have 32 64-bit elements" in order to maintain the fixed vector length.	
2	130	Line 19: Delete "(see Chapter 1)" because Chapter 1 doesn't discuss the popularity of MIPS	
2	131	Line 7: "and popular in some programs" --> "and are popular in some programs"	
2	132	Figure 2.27: "Conditional branch instructions (rs is register, rd unused)" --> "Conditional branch instructions (rs is register 1, rd is register 2)" or "Conditional branch instructions (rs is register 1, rd is register 2)"	
2	133	Figure 2.28, "LH R1,40(R3)", "Meaning": there is an error of using 'Mem[41+Regs[[R3]]', since 41 has never appeared before.	
2	133	Figure 2.28, Row SW, Column Meaning: Add the subscript "32..63" to "Regt[R3]" since it is not clear which 32 bits are being stored	
2	134	Figure 2.29: "DSL" should be replaced with "SLT" to be consistent with Figure 2.31	
2	134	Line 4: "Where one source operand is R0" --> "where the source operand is R0"	
2	134	Figure 2.30, "JAL name": Regs[31]<-PC+4 --> Regs[31]<-PC+8	
2	135	Figure 2.30, "JALR R2": Regs[31]<-PC+4 --> Regs[31]<-PC+8	
2	136	Line 2 in MIPS Instruction Set Usage section: "SPECint92" --> "SPECint2000"	
2	136	Line 5 in Another View: delete "infinite"	
2	137	"SUB.D, SUB.S, ADD.PS" --> "SUB.D, SUB.S, SUB.PS" (See also correction to the back inside cover)	
2	137	Line 9: "from/to" --> "to/from" and "to/from" --> "from/to" (to correspond with the ordering of the instructions MFCO and MTCO)--see also, the same correction on the back inside cover	
2	137	Line 23: "BEQ, BNE Branch GPR equal/not equal" --> "BEQ, BNE Branch GPRs equal/not equal"	
2	138	Figure 2.32: "perl" --> "perlbmk"	
2	140	Figure 2.34: "perl" --> "perlbmk"	
2	140	Figure 2.34, "add/sun" --> "add/sub"	

2	141	Figure 2.35, Video Processing row: remove comma between "dynamic noise" and "reduction"
2	144	Figure 2.38, "perl" --> "perlbmk"
2	145	Figure 2.40 "N real update" --> "N real updates" and "N complex update" --> "N complex updates" (see reference given for DSPstone)
2	145	Figure 2.40: "IIR n biquad" --> "IIR n biquads"
2	145	Figure 2.40 caption, line 2: "using the 14 DSPstone kernels" --> "using 13 of the 14 DSPstone kernels" (the FFT-unput-scaled benchmark is not used)
2	145	Figure 2.40 caption: "using the 6 EEMBC Telecom kernel" --> "using the 5 EEMBC Telecom kernels"
2	148	Line 13 from bottom: "seven popular load-store computers" --> "nine popular load-store computers" (The nine computers are: Alpha, PA-RISC, PowerPC, SPARC, ARM, Thumb, SuperH, M32R, and MIPS16.)
2	154	Figure 2.41: "matrix" --> "matrix300"
2	155	Line 11: "infinite" --> "continuous"
2	163	Figure 2.42: Change "Figure 2.8" to "Figures 2.8 and 2.20"
2	163	Fourth line from bottom in table: add "18" as the number of offset magnitude bits
2	165	Second line from the bottom of 2.11: "other" --> "other logical"
3	175	Third code block, line 3: R1 != zero --> R1 != R2 (same as first code block)
3	184	"Invented by Robert Tomosulo ... and introduces Register Renaming to minimize WAW and RAW hazards" --> "to minimize WAW and WAR hazards"
3	188	Line 26: "six load buffers"-->"five load buffers" to be consistent with Figure 3.2
3	188	Line 27: "11 registers"-->"10 registers" to be consistent with Figure 3.2
3	192	Second line from bottom: "DADDUI"-->"DADDIU"
3	195	Line 20: "We really only need keep the..." --> "We really only need to keep the..."
3	200	Bottom of page, under heading "Correlating Branch Predictors": "DSUBUI"-->"DADDIU"; "#2"-->"#2"
3	201	L1: "DSUBUI"-->"DADDIU"; "#2"-->"#2"
3	207	Caption to Figure 3.16: "The counter is incremented whenever the 'predicted' predictor is correct and the other predictor is incorrect, and it is decremented in the reverse situation"-->"The counter is incremented whenever predictor 2 is correct and predictor 1 is incorrect, and it is decremented in the reverse situation." (As it is worded, when predictor 1 is "predicted" correctly, the counter may cause the new prediction to be predictor 2.
3	210	2nd paragraph, line 3: "the entry must be for this instruction"-->"the predictive entry must be matched to this instruction"
3	211	Line 7 from bottom: "Chapter 1"-->"Appendix A"
3	211	Line 6 in Example: "Assume that 60% of the branches are taken"--this assumption is not used so it can be removed
3	215	First line in section 3.6: "previous two sections"-->"previous three sections" (the previous two sections do not talk about data hazards)
3	215	Line 5: "the previous"-->"previous" (the immediately preceding section talks about control dependences)
3	216	Line 6: "discussed in this subsection"-->"discussed in this and the last subsection"
3	217	Line 4: "...instructions preceding that one in the --> preceding that one in the
3	222	Figure 3.25, second row from bottom, second column (Instructions): "DAADIU"-->"DADDIU"
3	233	Figure 3.32, section "FP operations and stores," "Action or bookkeeping", Line 3: ROB[b] --> ROB[h]
3	233	Figure 3.32, section "FP operations" Line 1: "RegisterStat[r].Qi=b; --> "RegisterStat[r].Reorder=b;"
3	233	Figure 3.32, section "Loads" Line 1: "RegisterStat[r].Qi=b;" --> "RegisterStat[r].Reorder=b;"
3	233	Figure 3.32, section "Write result all but store," "Action or bookkeeping," Line 1: RS[r].Reorder --> RS[r].Dest
3	233	Figure 3.32, section "Store": "ROB[h].Value" --> "ROB[h].Destination
3	233	Figure 3.32, section "Commit", Line 6: Address --> Destination
3	233	Figure 3.32, section "Commit", Line 11: "if RegisterStat[d].Qi==h" --> "if RegisterStat[d].Reorder==h"
3	233	Figure 3.32 caption, line 2: "r is the reservation station allocated, and b is the assigned ROB entry"-->"r is the reservation station allocated, b is the assigned ROB entry, and h is the head entry of the ROB."

3	235	4th line of code segment: Change "DADDIU R1, R1, #4" to "DADDIU R1, R1, #8" (the code segment loads and stores double words)
3	236	Figure 3.33, rows 4, 9, and 14: Change "DADDIU R1, R1, #4" to "DADDIU R1, R1, #8" (the code segment loads and stores double words)
3	236	Figure 3.33: Change one occurrence of "L.D" in each to "LD"
3	236	Figure 3.33, last line of table: "BNZ"-->"BNE"
3	237	Figure 3.34: Change one occurrence of "L.D" in each to "LD"
3	237	Figure 3.34, rows 4, 9, and 14: Change "DADDIU R1, R1, #4" to "DADDIU R1, R1, #8" (the code segment loads and stores double words)
3	266	Figure 3.53: The dark shading for L2 in the legend does not match the very light shading for L2 misses in the bar graph
3	266	Third line from bottom: "accessing"-->"assessing"
3	267	Line 5 from bottom, remove "varying from 1.0 to 1.75"
3	268	Figure 3.55: The caption for the x-axis of the figure is "Percentage of instructions that do not commit" but the numbers along the x-axis seem to be fractions of instructions rather than percentages of instructions
3	269	Line 1: "There are seven integer execution units in NetBurst versus five in P6"-->"There are more integer execution units in NetBurst versus the P6."
3	270	Figure 3.57: The line for actual CPI should pass through the point for apsi
3	270	Line 5: "two floating operations"-->"two floating point operations"
3	271	Line 5: "1.5 for vortex"-->"1.4 for vortex" [since $(700/1.4)*1.7$ is 850 while $(700/1.5)*1.7$ is only 793]
3	274	Line 19: "Of course, this fallacy is nothing more than a restatement of a pitfall from Chapter 2 about comparing processors using only one part of the performance equation"-- no such pitfall in Chapter 2; possibly "a restatement of a pitfall from Chapter 1 about comparing processors using only clock rate or the performance of a single benchmark suite"?
3	274	2nd Pitfall, line 2: "It had a 1994 clock rate of 60Mhz"--> "In 1994, the clock rate was 60Mhz"
3	275	Figure 3.59: Replace the benchmark "earsu2cor" on the x-axis with the benchmark "su2cor"
3	283	"The Engineering Design of the Stretch Computer" appeared in "1959 Proceedings of the Eastern Joint Computer Conference," not in "Proceedings of the Fall Joint Computer Conference"--this error also appears in Appendix A, see below.
3	286	Delete entry for Riseman, E.M. and C.C. Foster (repeated from page 285)
3	286	8th reference on page: "Postiff, M.A. D.A. Greene, G.S. Tyson, and T.N. Mudge [1992]"; date should be 1999.
3	291	"Initially, R1=0 and F0 contains a"--> either italicize "a" or add the words, "a floating point number," to avoid confusion
3	291	3.6, code segment, lines 6 and 7: "DADDUI" --> "DADDIU"
3	291	3.6, code segment, line 8: "DSGTUI"-->"DSGTIU"
3	291	Exercise 3.6: After changes from posted errata are made, the code will initialize R1 to X, but compare against a bound of #796. The correct bound to compare is X+796, as R1 will start at X. Or you can use R1 just as iteration count and use a separate register for indexing into X[]; this would sidestep potential student questions about constant sizing in the DSGTUI instruction. If you count R1 DOWN, you can keep the number of instructions constant for the exercise.
3	296	3.21, line 1: "the speculative Tomasulo processor shown in Figure 3.28 on page 225"-->"the speculative Tomasulo processor shown in Figure 3.29 on page 228"
4	311	2nd paragraph, line 3: "causes a decrease in the instruction miss rate." -> causes an increase
4	311	In the assembly code, the gray arrow that shows the dependency between the fifth and sixth line of instructions (ADD.D F8, F6, F2 and S.D. F8, - 8(R1); drop DADDUI & BNE)--> the arrow <i>should</i> point show the link between the two references to FB in both instructions, it <i>should not</i> point from F6 in the fifth line to F8 in the sixth line
4	315	"denoting that the instructions, since they contained several instructions, were very wide" --> "denoting that the instructions, since they <i>encode several operations</i> , were very wide"
4	330	in answer, at the bottom of the page; for Iteration I+2, "F4, 0(R1)" should not be in bold
4	341	Paragraph under example, line 2: (B<0) {A=-B;} else {A=B;} --> (B<0) {A=-B;} else {A=B;};
4	356	"to think of the opcode as being 4 bits + the M, F, I, B, L + X designation"--> "to think of the opcode as being 4 bits <i>plus</i> " for clarity
4	363	"and to simply instruction dispatch"--> to <i>simplify</i> instruction dispatch
4	366	All boxes in the legend for Figure 4.20 are equally shaded

4	366	In legend, TM1300 optimized code-size lines should be indicated with square data points! NEC VR5000 should be circular (based on text's comments about relative code size)
4	371	Three lines from the bottom of the page: "For example, for the first machine in Figure 4.24"-->"For example, in August 2001, for the first processor in Figure 4.24"
4	373	"the IBM Power4, which contains two Power3 processors and an integrated second-level cache"-->"which contains two POWER processors" or "which contains two processors"
4	378	Delete reference to Riseman, E.M. and C.C. Foster (repeated)
5	397	"We get the same answer as on page 395"-->"We get the same answer as on page 395, showing equivalence of the two equations"
5	427	Second line from bottom: 1 KB --> 4 KB
5	427	Last line: "1 + (15.05% x 82) = 13.374" --> "1 + (8.57% x 82) = 8.027"
5	430	3rd line from top: "1.00 + (0.133 x 25)" --> "1.00 + (0.098 x 25)"
5	468	Figure 5.37, In the line L1 cache tag of the figure, it should say "28" for 28 bits, not 43 bits.
5	469	"the cache index would also shrink by n bits" --> "the cache index would also shrink by $\log_2 n$ bits"
5	492	Line 5: MB/sec --> GB/sec
5	520	Exercise 5.19, part d--> the second group of workloads does not add up to 100%
6	582	first line: "unlike in a snooping scheme"
6	602	Figure 6.42: The calls to "barrier" do not have actual parameters. There is a mismatched parenthesis on the line "if (tree[mynode].count==k) {"--the closing "}" should go before "unlock (...)"
6	626	"Each node consists of pairs of MIPS R1000 processors sharing a single memory module"--> "MIPS R10000"
7	695	The "Not Read" signed should be named "Not Write"
7	723	Last equation: subscript "server" in "timeserver"
7	727	2nd line from bottom: exponentially random request arrival --> exponentially random service time
7	784	Exercise 7.19: "Redo the example that starts on page 728, but this time assume the distribution of disk service times has a squared coefficient..." --> "Redo the first example on page 728 but this time assume the distribution of disk service times has a coefficient..."
8	801	Figure 8.10: Cummulative percentage --> Cumulative percentage
<b>Appendices</b>		
A	4	Line 8: "concepts are significantly similar that we will not need to distinguish the exact architecture"-->"concepts are significantly similar that they will apply to any RISC."
A	4	Last line: "DADDUI" --> "DADDIU"
A	6	Line 2: "assume" --> "explore later"
A	6	Add a new line before "5. Write-back cycle"
A	9	Line 2 from bottom: "where the source and destination may be directly adjacent" --> "where the source and destination may <i>not</i> be directly adjacent"
A	9	"In the case of a pipelined processor, the pipeline registers also play the key role of carrying intermediate results from one stage to another where the source and destination may be directly adjacent" --> "where the source and destination are not in the same pipestage"
A	9	Mid-page, reference to "RF" --> "ID"
A	10	Line 6: "ID/IF, IF/RF, RF/EX, EX/MEM, MEM/WB" --> "IF/ID, ID/EX, EX/MEM, MEM/WB"
A	15	Figure A.5: Add to the end of the figure caption, "Note that this figure assumes that instructions i+1 and i+2 are not memory references."
A	16	Line 17: "ADD" --> "DADD"
A	16	Line 19: "SUB" --> "DSUB"

A	16	Line 22: "SUB" --> "DSUB"	
A	16	Line 23: "ADD" --> "DADD"	
A	16	Line 36: "SUB" --> "DSUB"	
A	23	Figure A.12: lines 3, 4 and 5 are not well aligned on the two top lines.	
A	28	Insert Text line 2: "...all the instructions. We initially used a less aggressive implementation of a branch instruction. We show how to implement the more aggressive version at the end of this section.	
A	28	Line 12: "BEQ with RO" --> "BEQ with R0"	
A	28	Line 13: ".branch we consider.) --> ...branch we consider.	
A	32	Figure A.19, "Stage EX", "Load or Store Instruction", Line 1: <-- ID/EX.IR to EX/MEM.IR to ID/EX.IR;	change EX/MEM.IR
A	32	Figure A.19 caption, line 4: "from one or two sources" --> "from one of two sources"	
A	39	Figure A.25, IF, line 3: "(IF/ID.IR16)16##IF/ID.IR16..31##00)"-->"sign-extend (IF/ID.IR [immediate field] <<2)"	
A	39	Figure A.25, ID, line 1: "Regs[IF/ID.IR6..10]"-->"Regs [rs]"; "Regs [IF/ID.IR11..15]"-->"Regs [rt]"	
A	39	Figure A.25, ID, line 3: "(IF/ID.IR16)16##IF/ID.IR16..31)"-->"sign-extend (IF/ID.IR [immediate field])"	
A	39	Figure A.25 caption, line 5: "ID/EX register" --> "ID stage"	
A	50	Figure A.31, Add one more small box on the right to the shaded boxes labelled DIV on the bottom of the figure	
A	51	Figure A.33, "ADD.D F2, F0, F8": The "WB" cycle is missing for Instruction at Cycle 17.	
A	51	Figure A.33 caption, line 4: "SD" --> "S.D".	
A	53	Lines 1 and 4: "LD" --> "L.D"	
A	39	Line 11: "high CPI processor" --> "low CPI processor"	
A	52	Figure A.34 caption: "L.D." --> "L.D".	
A	56	Line 6 in Performance of a MIPS FP Pipeline section: replace "compare" with "convert"	
A	58	Figure A.36: The numbers in the chart are correct, but the bars in the chart are of incorrect lengths. For example, doduc's bars with 0.07 and 0.08 are shorter than mdljdp's bar of 0.10, and hydro2d has a bar of length 0.22 that is extremely short.	
A	59	Figure A.38: Change "ADDD" to "DADD"	
A	61	"Exercise 4.8 asks you to explore..."-->"Exercise A.8 asks to explore..."	
A	63	Figure A-44 The first issues position fro ADD instruction: "U S+AA+S R+S" --> "U S+AA+R R+S"	
A	64	Caption for Figure A.46, line 4: "cycle 28 will be stalled until cycle 34" with "--> "cycle 28 will be stalled until cycle 36"	
A	68	Third full paragraph: "To allow us to begin executing the SUB.D in the above example" --> "To allow an instruction to begin execution as soon as its operands are available, even if a predecessor is stalled,"	
A	68	Third full paragraph: "We can still check for structural hazards when we issue the instruction; thus, we still use in-order instruction issue." --> "We decode and issue instructions in order."	
A	77	Lines 14 and 16 in Fallacies and Pitfalls section: "LD" --> "L.D" to be consistent with code segment	
A	78	Pitfall: "Evaluating a compile time schedule on the basis of unoptimized code"-->"Evaluating dynamic or static scheduling on the basis of unoptimized code"	
A	78	Pitfall, line 8: "To fairly evaluate a scheduler"-->"To fairly evaluate a compile-time scheduler or runtime dynamic scheduling"	
A	80	"The Engineering Design of the Stretch Computer" appeared in "1959 Proceedings of the Eastern Joint Computer Conference," not in "Proceedings of the Fall Joint Computer Conference"	
A	81	Exercise 1: "SD 0 (R2), R1; store R1 at address 0 + R2" ----> "SD R1, 0 (R2) ; store R1 at address 0 + R2"	
A	81	Exercise A.1.a "Use a pipeline timing chart like Figure A.6" --> "like Figure A.5"	
A	81	Exercise A.1.b "Show the timing of this instruction sequence for the RISC pipeline with normal forwarding" --> "with full forwarding"	
A	81	Exercise A.1.b "Use a pipeline timing chart like Figure A.6" --> "like Figure A.5"	
A	81	Problem A.2: change two occurrences of "DADDUI" to "DADDIU"	
A	84	Exercise A.5.f "Show all control hazard types by example" --> "Show all control hazards by example"	
A	85	Code segment: Replace "MULT.D" with "MUL.D"; change 2 occurrences of "DADDUI" to "DADDIU"; change "DSGTUI" to "SGTIU" or change the code to use "SLTIU" instead of "SGTIU" since "SGTIU" is not introduced in the text	

<b>A</b>	86	A.12 b: Replace "SGTI" with "SGTIU"	
<b>A</b>	86	A.12 b and c: Replace two occurrences of "SAXPY" with "DAXPY"	
<b>B</b>	3	Figure B-2: The Harmonic mean for C should be 2.0, not 5.0	
<b>B</b>	4	Answer to 1.13 d: The means and medians are calculated for the wrong sets. Arithmetic mean C should be Arithmetic mean D, Median D should be Median C, and so on.	
<b>B</b>	6	Answer to 1.18b, line 5: "1000" --> "100"	
<b>B</b>	7	Equation for MFLOPS normalized, second line: "287" --> "287 x 10^6"	
<b>B</b>	11	Solution to 3.2, third code fragment: change first "S.D" to "SD" to be compatible with the exercise	
<b>B</b>	11	Solution to 3.2: how can an output dependence exist since R2 is an integer register and F2 is a floating-point register?	
<b>B</b>	13	Line 7 from bottom: "discussion of BTBs in Section 3.4 of the text"-->"discussion of BTBs in Section 3.5 of the text"	
<b>B</b>	14	Line 5: "DLX" --> "MIPS" (DLX not introduced in this edition of the text)	
<b>B</b>	19	Problem 4.12: "Because one is a factor of two, the GCD test indicates that there is a dependence in the code." This implies that the GCD test can indicate that a loop has a dependence (that it is not parallel). The GCD test can only indicate that a loop does not have a dependence or, equivalently, that it is parallel. At best, it can only say that a loop <i>may</i> have a dependence (that is, that it <i>may</i> not be parallel).	
<b>B</b>	19	Problem 4.12: One being a factor of two is the same as saying that $[2 \bmod 1 = 0]$ . Working backwards, these values come from $[\gcd(2, 100) \bmod (d-b) = 0]$ . But this is not the GCD test--therefore, saying that one is a factor of two is unrelated to the GCD test.	
<b>B</b>	19	Problem 4.12: The conclusion is that there is a dependence in the code, or the loop may not be parallel. Correction: The loop <i>is</i> parallel for all indices since the right hand values are always odd and the left hand values are always even, and it is parallel according to the GCD test (since $1 \neq 0$ ).	
<b>B</b>	36	Figure B.18 and in the text: Three occurrences of "DADDUI" should be changed to "DADDIU"	
<b>B</b>	37	Figure B.19: Two occurrences of "DADDUI" should be changed to "DADDIU"	
<b>B</b>	40	"Pipeline stalls $real = (1*1\%)+(2*9\%)+(1*6\%) = 0.24$ " --> " $= 0.25$ "	
<b>G</b>	9	"much like to an assembly line"--> "much like an assembly line" or "similar to an assembly line"	
<b>R</b>	11	Add reference: Kembel, R. [2000]. "Fibre Channel: A Comprehensive Introduction." Internet Week, April 2000.	
<b>Back Cover</b>		Line 9, MFCO, MTCO: "Copy from/to GPR to/from a special register"-->"Copy from/to a special register to/from GPR"	
		"SUB.D, SUB.S, ADD.PS" --> "SUB.D, SUB.S, SUB.PS"	
	Hardware Description Notation	Notation column, row 2: "M" --> "Mem"	
		Example column, row 2: Change "Regs [R1] <- M[x];" to "Regs[R1] <- Mem[x];"	
		Example column, row 3: Change "M[y] <- 16 M[x]" to "Mem[y] <- 16 Mem [x]"	
		Example column, row 5: Change "Regs[R3]24..31 <- M[x];" to "Regs [R3]24..31 <- Mem[x];"	
		Example column, row 7: Change "Regs [R3] <- 024 ## M[x]; F2 ## F3 <- 64 M[x];" to "Regs [R3] <- 024 ## Mem [x]; F2 ## F# <-64 Mem[x];"	
		Meaning column, row 2: Replace five occurrences of "M" with "Mem"	
		row ##, column example: Delete example "F2##F3 <- 64 Mem[x]"	
		Switch the last two rows (since "&" is used to mean bitwise--and in the second-to-last row, and it is not introduced until the last row. "&" has only been introduced prior to the second-to-last row as the operator to get the address of a variable)	
		Meaning column: "the transferred bytes are M[i], M[i+1], M[i+2], and M[i=3]" --> "and M [i+3]"	
		Subset of the Instructions in MIPS64: "BEQ, BNE Branch <GPR> equal/not equal" --> ">GPRs<"	

	Events on Every Pipe Stage of the MIPS Pipeline	"Events on every Pip Stage...", "Stage EX", "Load or Store Instruction", Line 1: change EX/MEM.IR <-- ID/EX.IR to EX/MEM.IR to ID/EX.IR;	
--	---	--	--