

The reduced instruction set computer is an alternative to the general trend toward increasingly complex instruction sets. It executes most instructions in a single, short cycle.

A VLSI RISC

David A. Patterson and Carlo H. Séquin

University of California, Berkeley

A general trend in computers today is to increase the complexity of architectures commensurate with the increasing potential of implementation technologies, as exemplified by the complex successors of simpler machines. Compare, for example, the DEC VAX-11¹ to the PDP-11, the IBM System/38² to the System/3, and the Intel iAPX-432^{3,4} to the 8086. The complexity of this class of computers, which we call CISCs for complex instruction set computers, has some negative consequences: increased design time, increased design errors, and inconsistent implementations.⁵

Investigations of VLSI architectures indicate that the delay-power penalty of data transfers across chip boundaries and the still-limited resources (devices) available on a single chip are major design limitations. Even a million-transistor chip is insufficient if a whole computer has to be built from it.⁶ This raises the question of whether the extra hardware needed to implement a CISC is the best use of "scarce" resources.

The above findings led to the Reduced Instruction Set Computer Project. The purpose of the RISC Project is to explore alternatives to the general trend toward architectural complexity. The hypothesis is that by reducing the instruction set one can design a suitable VLSI architecture that uses scarce resources more effectively than a CISC. We also expect this approach to reduce design time, design errors, and the execution time of individual instructions.

An earlier version of this article, entitled "RISC I: A Reduced Instruction Set VLSI Computer," appeared in the *Proc. Eighth Int'l Symp. Computer Architecture*, May 1981, pp. 443-457.

Our initial version of such a computer is called RISC I. To meet our goals of simplicity and effective single-chip implementation, we somewhat artificially placed the following design constraints on the architecture:

- (1) *Execute one instruction per cycle.* RISC I instructions should be about as fast and no more complicated than microinstructions in current machines such as the PDP-11 or VAX.
- (2) *Make all instructions the same size.* This again simplifies implementation. We intentionally postponed attempts to reduce program size.
- (3) *Access memory only with load and store instructions; the rest operate between registers.* This restriction simplifies the design. The lack of complex addressing modes also makes it easier to restart instructions.
- (4) *Support high-level languages.* The degree of support is explained below. Our intent is to optimize the performance of RISC I for use with high-level languages.

RISC I supports 32-bit addresses, 8-, 16-, and 32-bit data, and several 32-bit registers. We intend to examine support for operating systems and floating-point calculations in the future.

It would appear that these constraints, based on our desire for simplicity and regularity, would result in a machine with substantially poorer code density, poorer performance, or both; but in spite of these constraints, the resulting architecture competes favorably with other

Recent developments

Since this article was submitted, we have received our first good silicon, and it looks like beginner's luck applies to VLSI. These chips correctly ran all diagnostic programs used to verify our original design. We (foolishly) created new diagnostics and uncovered a design error associated with the optional setting of condition codes on the load and shift instructions. Defying historical precedent for solving the problem by announcing a new architectural "feature," we decided to cover this minor error by modifying the RISC I assembler. (This was possible because ALU operations properly set all condition codes, whereas load and shift instructions do not set the negative condition bit. The patch consists of inserting an arithmetic test instruction when a conditional jump needs the N condition from a load or shift operation.)

The fastest of these chips runs all diagnostics at 1.5 MHz, or 2 μ sec per RISC I instruction. Several factors explain the difference between expected and measured performance. The chief one is inexperience; this was the first chip that any of us had built. A second is raw speed of transistors from this fabrication. Test structures ran about half the speed of other runs. The last stage of design involved connecting cells, and we concentrated our resources on logical

correctness rather than circuit speed. We recently reexamined the design and found four long clocked control lines that an analog circuit simulator predicts will limit the maximum clock speed to 4 MHz. Furthermore, many of our diagnostics can be run with a 3-MHz clock, suggesting that only a few RISC instructions are limiting performance. Finally, as we have still tested only 20 percent of the chips, we may well find faster RISC I's.

Even at 1.5 MHz and the assembler correction of the error, RISC I still runs programs faster than commercial microprocessors. RISC I was put onto a board with memory, I/O, and memory management on June 11, 1982, and ran its first program.

The bottom line of the RISC I effort is that students, as part of the graduate curriculum, designed and evaluated an architecture, learned Mead-Conway design, built new CAD tools, and tested their design. The end product, a 44,500-transistor integrated circuit, has one minor design error; it worked on the first good silicon and runs programs faster than commercial microprocessors.

More details can be found in an article in the September/October 1982 issue of *VLSI Design* entitled "Running RISCs."

microprocessors and minicomputers. This is due largely to a scheme of register organization we call overlapped register windows.

Support for high-level languages

Clearly, new architectures should be designed with the needs of high-level language programming in mind. It should not matter, however, whether a high-level language system is implemented mostly by hardware or mostly by software, provided the system hides any lower levels from the programmer.⁷ Given this framework, the role of the architect is to build a cost-effective system by deciding what pieces of the system should be in hardware and what pieces in software.

The selection of languages for consideration in RISC I was influenced by our environment; we chose "C" because of its large user community and, hence, considerable local expertise. Given the limited number of transistors that can be integrated into a single-chip computer, most of the pieces of a RISC high-level language system are in software, with hardware support for only the most time-consuming events.

To determine what constructs are used most frequently and, if possible, what constructs use the most time in average programs, we first looked at the frequency of classes of variables in high-level language programs. Data collected for Pascal and C are shown in Table 1.

The most important observation was that integer constants appeared almost as frequently as arrays or struc-

tures. What is not shown is that more than 80 percent of the scalars were local variables and more than 90 percent of the arrays or structures were global variables.

We also looked at the relative dynamic frequencies of high-level language statements for the same eight programs; average occurrences over one percent are shown in Table 2. This information does not tell what statements use the most time in the execution of typical programs. To answer that question, we have to look at the code produced by typical versions of each of these statements. A "typical" version of each statement was supplied by Wulf as part of his study into judging the quality of compilers.⁸ We used C compilers for the VAX, PDP-11, and 68000 to determine the average number of instructions and mem-

Table 1.
Dynamic percentage of operands in Pascal and C.

	P1	P2	P3	P4	C1	C2	C3	C4	AVERAGE
INTEGER CONSTANT	14	18	11	20	25	11	29	28	20 \pm 7
SCALAR	63	68	46	54	37	45	66	62	55 \pm 11
ARRAY/STRUCTURE	23	14	43	25	36	43	5	10	25 \pm 14

PROGRAM EXPLANATION

P1	COMP - a Pascal P-code style compiler
P2	MACRO - the macro expansion phase of the SCALD I design system
P3	PRINT - a prettyprinter for Pascal
P4	DIFF - a program that finds the differences between two files
C1	PCC - the portable C compiler for the VAX
C2	CIFPLOT - a program that plots VLSI mask layouts on a dot plotter
C3	NROFF - a text formatting program
C4	SORT - the Unix sorting program

Table 2.
Relative frequency of Pascal and C statements.

STATEMENTS*	P1	P2	P3	P4	AVERAGE	C1	C2	C3	C4	AVERAGE
ASSIGN	39	52	35	53	45±8	22	50	25	56	38±15
IF	35	30	36	16	29±8	59	31	61	22	43±17
CALL	15	14	16	15	15±1	6	17	9	16	12±5
WITH	2	0	5	13	5±5	2	2	3	5	3±1
LOOP	5	5	5	4	5±0	9	0	1	1	3±4
CASE	4	0	1	0	1±1	2	-	-	0	<1±1

*Because statements can be nested, we count each occurrence of a statement. Loop statements are counted once per execution rather than once per loop iteration. For example, if two IF statements and three assignment statements appear in a loop that iterates 5 times, we would count 26 statements with 15 assignments, 10 IF statements, and one loop. The WITH statement qualifies a record name.

ory references per statement. By multiplying the frequency of occurrence of each statement with the corresponding number of machine instructions and memory references, we obtain Table 3, which is ordered by memory references.

The data in Table 3 suggest that the procedure call/return is the most time-consuming operation in typical high-level language programs. These results corroborate studies by Lunde⁹ and Wichmann.¹⁰ The statistics on operands found in Table 1 emphasize the importance of local variables and constants. RISC I supports HLLs by enhancing performance of the most time-consuming features of typical HLL programs, as opposed to making the architecture "close" to a particular HLL; thus, RISC I attempts to handle local variables, constants, and procedure calls efficiently while leaving less frequent operations to instruction sequences or subroutines.

Basic architecture of RISC I

The RISC I architecture has 31 instructions, most of which do simple ALU and shift operations on registers. As shown in Table 4, they have been grouped into four

categories: arithmetic-logical, memory access, branch, and miscellaneous. Instructions, data, addresses, and registers are 32 bits. The execution time of a RISC I cycle is given by the time it takes to read and add two registers, and then store the result back into a register. Register 0, which always contains zero, allows us to synthesize a variety of operations and addressing modes.

Load and store instructions move data between registers and memory. Rather than lengthen the general cycle to permit a complete memory access, these instructions use two CPU cycles. There are eight variations of memory access instructions to accommodate sign-extended or zero-extended 8-bit, 16-bit, and 32-bit data. Although there appears to be only the index-plus-displacement addressing mode in data transfer instructions, absolute and register-indirect addressing can be synthesized using register 0 (see Table 5).

Branch instructions include call, return, conditional, and unconditional jump. The conditional instructions are the standard set used originally in the PDP-11 and found in most 16-bit microprocessors today. Most of the innovative features of RISC I are found in call, return, and jump; they will be discussed later.

Figure 1 shows the 32-bit format used by register-to-register instructions and memory access instructions. For register-to-register instructions, DEST selects one of the 32 registers as the destination of the result of the operation performed on the registers specified by SOURCE1 and SOURCE2. If IMM=0, the low-order five bits of SOURCE2 specify another register; if IMM=1, SOURCE2 expresses a sign-extended 13-bit constant. As mentioned above, the frequency of integer constants in HLL programs suggests architectural support, so immediate operands are available in every instruction. SCC determines whether or not the condition codes are set. Memory access instructions use SOURCE1 to specify the index register and SOURCE2 to specify the offset. One other format combines the last three fields to form a 19-bit PC-relative address and is used primarily by the branch instructions.

The examples in Table 6 show that many of the important VAX instructions can be synthesized from simple RISC I addressing modes and opcodes. Comparative measurements of benchmarks will demonstrate the effectiveness of the chosen instruction set.

Table 3.
Weighted relative frequency of HLL statements
(ordered by memory references).

STATEMENTS*	HLL (OCCURRENCE)		WEIGHTED (MACHINE INSTR.)		WEIGHTED (MEM. REF.)	
	P	C	P	C	P	C
CALL/RETURN	15±1	12±5	31±3	33±14	44±4	45±19
LOOPS	5±0	3±1	42±3	32±6	33±2	26±5
ASSIGN	45±5	38±15	13±2	13±5	14±2	15±6
IF	29±8	43±17	11±3	21±8	7±2	13±5
WITH	5±5	—	1±0	—	1±0	—
CASE	1±1	<1±1	1±1	1±1	1±1	1±1
GOTO	—	3±1	—	0±0	—	0±0

*For the CALL statement we counted passing parameters, saving/restoring general registers, and saving/restoring the program counter. The IF and CASE statements include instructions to evaluate expressions and to jump. For LOOP statements we count all the machine instructions executed during each iteration.

INST
ADD
ADDI
SUB
SUBI
SUBI
AND
OR
XOR
SLL
SRL
SRA
LDL
LDSI
LDSI
LDBI
LDBI
STL
STS
STB
JMP
JMP
CALI
CALI
RET
CALI
RETI
LDH
GTLI
GETI
PUTI
R
level
time
gran
ADDF
REGI
IMMI
INDE
ABSC
REG
Sept

Rd = DESTINATION REGISTER
 Rs = SOURCE REGISTER
 S2 = SOURCE 2

Table 4.
Assembly language definition for RISC I.

INSTRUCTION	OPERANDS	COMMENTS	
ADD	Rs, S2, Rd	Rd ← Rs + S2	integer add
ADDC	Rs, S2, Rd	Rd ← Rs + S2 + carry	add with carry
SUB	Rs, S2, Rd	Rd ← Rs - S2	integer subtract
SUBC	Rs, S2, Rd	Rd ← Rs - S2 - carry	subtract with carry
SUBR	Rs, S2, Rd	Rd ← S2 - Rs	integer subtract
SUBCR	Rs, S2, Rd	Rd ← S2 - Rs - carry	subtract with carry
AND	Rs, S2, Rd	Rd ← Rs & S2	logical AND
OR	Rs, S2, Rd	Rd ← Rs S2	logical OR
XOR	Rs, S2, Rd	Rd ← Rs xor S2	logical EXCLUSIVE OR
SLL	Rs, S2, Rd	Rd ← Rs shifted by S2	shift left
SRL	Rs, S2, Rd	Rd ← Rs shifted by S2	shift right logical
SRA	Rs, S2, Rd	Rd ← Rs shifted by S2	shift right arithmetic
LDL	(Rx)S2, Rd	Rd ← M[Rx + S2]	load long
LDSU	(Rx)S2, Rd	Rd ← M[Rx + S2]	load short unsigned
LDSS	(Rx)S2, Rd	Rd ← M[Rx + S2]	load short signed
LDBU	(Rx)S2, Rd	Rd ← M[Rx + S2]	load byte unsigned
LDBS	(Rx)S2, Rd	Rd ← M[Rx + S2]	load byte signed
STL	Rm, (Rx)S2	M[Rx + S2] ← Rm	store long
STS	Rm, (Rx)S2	M[Rx + S2] ← Rm	store short
STB	Rm, (Rx)S2	M[Rx + S2] ← Rm	store byte
JMP	COND, S2(Rx)	pc ← Rx + S2	conditional jump
JMPR	COND, Y	pc ← pc + Y	conditional relative
CALL	Rd, S2(Rx)	Rd ← pc, next pc ← Rx + S2, CWP ← CWP - 1	call and change window
CALLR	Rd, Y	Rd ← pc, next pc ← pc + Y, CWP ← CWP - 1	call relative and change window
RET	Rm, S2	pc ← Rm + S2, CWP ← CWP + 1	return and change window
CALLINT	Rd	Rd ← last pc; next CWP ← CWP - 1	disable interrupts
RETINT	Rm, S2	pc ← Rm + S2; next CWP ← CWP + 1	enable interrupts
LDHI	Rd, Y	Rd < 31:13 > ← Y; Rd < 12:0 > ← 0	load immediate high
GTLPC	Rd	Rd ← last pc	to restart delayed jump
GETPSW	Rd	Rd ← PSW	load status word
PUTPSW	Rm	PSW ← Rm	set status word

Register windows. Investigations into the use of high-level languages suggest that the procedure call is the most time-consuming operation in high-level language programs. Potentially, RISC programs may have even more

calls, because the complex instructions found in CISCs are subroutines in RISCs. Thus, the procedure call must be as fast as possible, perhaps no longer than a few jumps. Because of its register window scheme, RISC I approaches this goal and reduces data memory traffic.

Table 5.
Synthesizing VAX addressing modes.

ADDRESSING	VAX	RISC EQUIVALENT
REGISTER	Rx	Rx
IMMEDIATE	#LITERAL	S2 (13-BIT LITERAL)
INDEXED	Rx + DISPL	Rx + S2 (13-BIT DISPLACEMENT)
ABSOLUTE	@#ADDRESS	r0 + S2 (r0 = 0)
REG INDIRECT	(Rx)	Rx + 0

Using procedures involves two groups of time-consuming operations: saving or restoring registers on each call or return, and passing parameters and results to and from the procedure. The frequency of local scalar variables justifies the architectural support of placing locals in registers, and Baskett¹¹ and Sites¹² have proposed that




Figure 1. RISC I basic instruction format.

Table 6.
Synthesizing VAX instructions.

OPERATION	VAX		RISC I EQUIVALENT	
REG-REG MOVE	MOVL	Rm, Rn	ADD	R0, Rm, Rn (r0 ≡ 0)
COMPARE	CMPL	Rm, Rn	SUB	Rm, Rn, r0, {c}
COMPARE TO 0	TSTL	Rn	SUB	Rn, r0, r0, {c}
	TSTL	A	LDL	(r0)A, r0, {c}
CLEAR	CLRL	Rn	ADD	r0, r0, Rn
	CLRL	A	STL	r0, (r0)A
TWOS COMPLEMENT	MNEGL	Rm, Rn	SUB	r0, Rm, Rn
ONES COMPLEMENT	MCOML	Rm, Rn	XOR	Rm, # - 1, Rn
LOAD CONST	MOVL	\$N, Rm(N < 2 ¹²)	ADD	r0, #N, Rm
	MOVL	\$N, Rm(N ≥ 2 ¹²)	LDHI	#N < 31:13 >, Rm
INCREMENT	INCL	Rn	ADD	r0, #N < 12:0 >, Rm
DECREMENT	DECL	Rn	ADD	Rn, #1, Rn
			SUB	Rn, #1, Rn
CHECK INDEX BOUNDS, (A[0:U])	INDEX	Rm, #0, #U, #1, A, Rn;	SUB	Rm, #U, r0 {c};
TRAP IF ERROR, AND READ A[Rm]	MOVB	(Rn), Rp	JMP	lequ, OK; *
			CALL	error;
			OK: LDBU	(Rm)A, Rp

*This approach is better than the normal algorithm. We can think of an index as an unsigned integer since $0 \leq \text{index} \leq U$. A two's complement negative number ($1X...X$) is then a large unsigned number, so we only need make one unsigned test instead of two signed tests. Nonzero lower bounds are handled by subtracting the lower bound from the index, and multiple indices are handled by repeating the sequence and including a multiply and an add. This idea resulted from a discussion between Bill Joy, Peter Kessler, and George Taylor. Taylor coded the examples and found that on the VAX-11/780, the sequence of simple instructions was always faster than the index instruction. This optimization is found in the Unix C optimizer.

microprocessors keep multiple banks of registers on the chip to avoid register saving and restoring. A similar scheme was adopted by RISC I. Each procedure call allocates a new "window" of registers from the large register file for use by that procedure, and the return resets a pointer, restoring the old set. But some of the registers are not saved or restored on each procedure call; these registers (r0 through r9) are called global registers.

Furthermore, the sets of registers used by different procedures overlap, allowing parameters to be passed in registers. In other machines, parameters are usually passed on the stack, and the calling procedure uses a register (frame pointer) that points to the beginning of the parameters (and also the end of the locals). Thus, all references to parameters are indexed references to mem-




Figure 2. Naming within one virtual RISC I register window.

ory. Our approach partitions the set of window registers (10-31) into the three parts defined by their respective overlap. Every procedure sees the set of registers shown in Figure 2.

High registers 26 through 31 contain parameters passed from "above" the current procedure—that is, from the calling procedure. Local registers 16 through 25 are used for local scalar storage. Low registers 10 through 15 are used for temporaries and parameters passed to the procedure "below" the current procedure (the called procedure). On each procedure call a new set of registers, numbered 10-31, is allocated. The low registers of the "caller" become the high registers of the "callee" because of the hardware overlap between subsequent register windows. Thus, without moving information, parameters in registers 10-15 appear in registers 25-31 of the called window. Figure 3 illustrates this approach for the case where procedure A calls procedure B, which calls procedure C.

If the nesting depth is sufficiently large, all register windows will be used. RISC I handles a call overflow with a separate stack in memory. Overflow and underflow are handled with a trap to a software routine that adjusts that stack. Because this routine can save or restore several sets of registers, the overflow/underflow frequency is based on local variations in the depth of the stack rather than absolute depth. The effectiveness of this scheme depends on the relative frequency of overflows and underflows. Studies by Halbert and Kessler¹³ show that with eight register banks overflow will occur in less than one percent of the calls. This suggests that programs exhibit locality in the dynamic nesting of procedures, just as they exhibit locality in memory references.

Another problem with variables in registers occurs in referencing them with pointers, since this requires variable addresses. Because registers normally do not have memory addresses, we could let the compiler determine which variables have pointers and put these variables in

memory: slows acc by giving portion (compari: register. Because and stor feature v the use (leaves m This : addressi nested p variable: single pr environ Such a c ory add variable: dynamic dresses.

Delay just lon operatic store th mance t ecution ficulties machin: propriat niques a solutio effect u this as t

The d next ins instructio ranged RISC I; languag comple: mers of

Table normal the orde in RISC tion (T: now 10: jump co howeve sequenc while n appears quence followi not de equival

Arch chines importan:

Septem

memory, but this precludes separate compilation and slows access to these variables. RISC I solves that problem by giving addresses to the window registers. By reserving a portion of the address space, we can determine, with one comparison, whether a register address points to a CPU register or to one that has overflowed into memory. Because the only instructions accessing memory—load and store—already take an extra cycle, we can add this feature without reducing their performance. This permits the use of straightforward compiler technology and still leaves most of the variables in registers.

This addressing technique also solves the “up-level addressing” problem. Pascal and other languages allow nested procedure declarations, thereby creating a class of variables that are neither global variables nor local to a single procedure. Compilers keep track of each procedure environment using static and dynamic links or displays. Such a compiler for RISC I would also associate the memory address for the window of local variables. These variables would then be accessed by using the display or dynamic chains to find the corresponding memory addresses.

Delayed jump. The normal RISC I instruction cycle is just long enough to execute the following sequence of operations: read a register, do an ALU operation, and store the result back into a register. We increase performance by prefetching the next instruction during the execution of the current instruction. This introduces difficulties with branch instructions. Several high-end machines have elaborate techniques to prefetch the appropriate instruction after the branch,¹⁴ but these techniques are too complicated for a single-chip RISC. Our solution was to redefine jumps so that they do not take effect until after the *following* instructions; we refer to this as the delayed jump.

The delayed jump allows RISC I to always prefetch the next instruction during the execution of the current instruction. The machine language code is suitably arranged so that the desired results are obtained. Because RISC I is always intended to be programmed in high-level languages, we will not burden the programmer with this complexity; the “burden” will be carried by the programmers of the compiler, the optimizer, and the debugger.

Table 7 illustrates the delayed branch. Machines with normal jumps would execute the sequence in Table 7a in the order 100, 101, 102, 105, To get that same effect in RISC I, we would have to insert a no operation instruction (Table 7b). The sequence of instructions for RISC I is now 100, 101, 102, 103, 106, In the worst case, every jump could take two instructions. The RISC I compiler, however, includes an optimizer that tries to rearrange the sequence of instructions to do the equivalent operations while making use of the instruction slot where the NOP appears. As shown in Table 7c, the optimized RISC I sequence is 100, 101, 102, 105, Because the instruction following a jump is always executed and the jump at 101 is not dependent on the add at 102, this sequence is equivalent to the original program segment in Table 7a.

Architectural heritage. Since architects of new machines build on the work of others, we believe it is important to trace the genealogy of RISC I. Its earliest

ancestor is the 1951 Ferranti-Manchester MADM—the first machine with index registers—which also used a register to supply zero.¹⁵ Seymour Cray revived the idea in 1964 with the CDC-6400 and continued to use it in the CDC-7600 and the Cray 1. The delayed jump was first used in the Maniac I, which was completed just a year after the MADM, but we adopted the idea from micro-programmed control units, where delayed jumps are the norm.

The leading proponent of reduced instruction set computers for floating-point data is Cray. For the last 15 years, he has combined simple instruction sets with so-

Table 7.
Normal and delayed jumps.

ADDRESS	(a) NORMAL JUMP		(b) DELAYED JUMP		(c) OPTIMIZED DELAYED JUMP	
100	LOAD	X,A	LOAD	X,A	LOAD	X,A
101	ADD	1,A	ADD	1,A	JUMP	105
102	JUMP	105	JUMP	106	ADD	1,A
103	ADD	A,B	NOP		ADD	A,B
104	SUB	C,B	ADD	A,B	SUB	C,B
105	STORE	A,Z	SUB	C,B	STORE	A,Z
106			STORE	A,Z		




Figure 3. Use of three overlapped register windows.

phisticated pipelined implementations to create the most powerful floating-point engines in the world. While Cray concentrates on impressive floating-point rates at impressive costs, RISC I concentrates on improved performance at lower cost for integer programs written in HLLs.

A machine with similar goals that predates RISC I is the IBM 801. This project, led by John Cocke and George Radin, began in 1975 by reexamining the relationship between instruction sets, compilers, and operating systems. They pushed the state of the art of compiler technology and created an extremely fast, reduced-instruction-set ECL minicomputer. Alas, the architecture community was left to speculate on the truth of widely varying rumors about the technical details¹⁶ as well as the success or failure of the project.¹⁷ Fortunately, accurate information is beginning to emerge.¹⁸ It will be interesting to see the similarities between RISC I and the 801; one difference is that RISC I uses traditional compiler technology and the 801 uses a traditional register set.

In searching the annals of computer architecture we cannot find a clear reference to overlapped register windows. To our best knowledge, no machine uses the

scheme for fast, multiport registers in the CPU. Most modern machines support procedure call by having instructions that manage a portion of main memory as a stack to pass parameters and allocate locals. Theoretically, a cache should then make such a scheme as fast as the overlapped register windows. Registers are faster than caches because of the difference in speed between a small memory and a large memory, the difference in speed between a deterministic access and a probabilistic access, and the difference in speed between a nontranslated register access and a translated virtual memory access. Theoretically, hardware can overcome almost any obstacle, but it occasionally stumbles in implementation. The advantages of registers become apparent when we look at concrete realizations; as we shall see, procedure call/return on the VAX-11/780, using a software stack enhanced by a hardware cache, is about an order of magnitude slower than the overlapped register windows of RISC I.

There are a few machines that share features of RISC I's overlapped register window scheme. The BBN C/70, a recent machine, allocates a new set of registers on every procedure call, but it does not overlap register sets.




Figure 4. Photomicrograph of RISC I.

Table 8.
VLSI design metrics for Z8000, MC68000, iAPX-432, and RISC I.

	ZILOG Z8000	MOTOROLA 68000	43201	INTEL iAPX-432 43202	43203	RISC I
TOTAL DEVICES	17.5K	68K	110K	49K	60K	44K
TOTAL MINUS ROM	17.5K	37K	44K	49K	44K	44K
DRAWN DEVICES	3.5K	3.0K	5.6K	9.5K	5.7K	1.8K
REGULARIZATION FACTOR	5.0	12.1	7.9	5.2	7.7	25
SIZE OF CHIP (MILS)	238×251	246×281	318×323	366×313	358×326	406×305
(AREA IN MIL ²)	60K	69K	103K	115K	117K	124K
SIZE OF CONTROL (MIL ²)	37K	35K	67K	45K	47K	7K
PERCENT CONTROL	53%	50%	65%	39%	40%	6%
TIME TO FIRST SILICON (MONTHS)	30	30	33	33	21	19
DESIGN EFFORT (MAN MONTHS)	60	100	170	170	130	15
LAYOUT EFFORT (MAN MONTHS)	70	70	90	100	50	12

A popular architecture that comes close to RISC I is the Texas Instruments 990-9900 family. These machines allocate their general "registers" in memory, so adding the contents of one register to another results in three memory accesses. A single register points to the register work space; most of the machines allow the pointer to overlap work spaces. The latest generation of this family, the TI 99000, includes on-chip main memory, but the first models appear to still have slow register access.¹⁹ The machine that comes closest to the overlapped register windows is the Bell Labs MAC-8. The state of NMOS technology in 1975 precluded having a rich instruction set and a register file on the chip; the architects chose the rich instruction set. The main difference between the MAC-8 and TI 990 is that the Bell architects realized that overlapping the registers could improve the performance of the procedure call and provided instructions to specifically overlap the register windows in memory. It is our understanding that some C compilers used this feature. This machine was never implemented with on-chip registers, and the logical successor to this machine, the BellMAC-32, has abandoned this approach.

VLSI implementation

The transition from theoretical architecture to concrete circuits began on January 6, 1981. Mask descriptions were completed June 22, and we received first silicon on October 23. Figure 4 is a photomicrograph of RISC I. We followed the Mead-Conway design philosophy for NMOS with lambda at two microns and no buried contacts. This first version, RISC I "Gold" as it is known internally, implements the complete instruction set and six windows with a total of 78 registers. The only piece of the architecture not implemented is the mapping of registers into the memory address space.

We collected statistics on the design and layout of RISC I.²⁰ Table 8 compares these results to VLSI implementations of more complex architectures. The most visible impact of the reduced instruction set is the reduced control

area: control is only six percent of RISC I compared to 50 percent in others. RISC I is also more regular. Lattin defines the regularity factor as the total number of transistors (less those in ROM) divided by the number of individually drawn transistors.²¹ By this measure RISC I is two to five times more regular than the Z8000, 68000, or 432. The time from the first discussion of the RISC I architecture to the masks was 14 months—less than the development periods of other machines. This was due in part to the reduced instruction set and in part to the Berkeley CAD software, a good match for this style of VLSI design. The primary interface was Caesar, an excellent color graphics layout editor developed by Ousterhout.²²

Evaluation

Register windows. Two benchmarks—"puzzle" and "quicksort"—showed the effectiveness of window registers in reducing procedure cost. The two recursive programs behave quite differently. Quicksort has a large percentage of procedure calls. Puzzle has such a low density of calls that it is almost atypical for modern structured programs, but it does have a large nesting depth. In both cases, the window scheme proved to be beneficial. Table 9 shows the maximum depth of recursion, the number of register window overflows and underflows, and the total

Table 9.
Memory traffic due to call/return.

	CALLS PLUS RETURNS, % INSTRS	MAXIMUM NESTED DEPTH	RISC I OVERFLOWS + UNDERFLOWS	DATA MEMORY TRAFFIC RISC I, WORDS	VAX, WORDS
PUZZLE	43K 0.7%	20	124	8K 0.8%	444K 28.0%
QUICKSORT	111K 8.0%	10	64	4K 1.0%	696K 50.0%

number of words transferred between memory and the RISC CPU as a result of the overflows and underflows. It also shows the memory traffic due to saving and restoring registers in the VAX. For this simulation we assumed that half of the registers were saved on an overflow and half were restored on an underflow. We found that for RISC I, an average of 0.37 words are transferred to memory per procedure invocation for the puzzle program and 0.07 for quicksort. Note that half of the data memory references in quicksort are the result of the call/return overhead of the VAX.

Table 10 compares the average "cost" of the RISC procedure mechanism—measured in execution time, number of instructions executed, and data memory accesses—to that of three traditional machines. The data was collected by looking at the code generated by C compilers for these machines for procedure call and return statements, assuming that two parameters are passed and requiring that three registers be saved.

The window scheme also reduces off-chip memory accesses. In traditional machines, 30 to 50 percent of the instructions generally access data memory, but no more than 20 percent of the instructions are register-to-register.^{23,24} Because RISC I arithmetic and logical instructions cannot access memory, one might expect a higher percent of data transfer instructions. This is not the case. The static frequencies of RISC I instructions for nine typical C programs show that less than 20 percent of the instructions are loads and stores, but more than 50 percent are register-to-register. RISC I has successfully changed the allocation of variables from memory into

registers, thus minimizing the slower off-chip memory accesses. This demonstrates that complex addressing modes are not necessary to obtain an effective machine.

Delayed jump. The effectiveness of rearranging the code around jump instructions can be evaluated by counting the NOP instructions in a program. Static figures before optimization show that in typical C programs about 18 percent of the instructions are NOPs inserted after jump instructions. A simple peephole optimizer reduces this to about eight percent. The optimizer does well on unconditional branches (removing about 90 percent of the NOPs) but not so well with conditional branches (removing only about 20 percent of the NOPs). Note that these are the static numbers; the dynamic numbers can be worse.

This optimizer was improved to replace the NOP by the instruction at the target of a jump. This technique can be applied to conditional branches if the optimizer determines that the target instruction modifies temporary resources—for example, an instruction that only modifies the condition codes. In quicksort, this removes all NOPs except those that follow return instructions, dropping NOPs from 12 percent statically to three percent. The dynamic effectiveness of the delayed branch must now include the NOPs plus the instructions after conditional branches that need not be executed for a particular jump condition. The total percentages of either type of instruction are again program dependent, ranging from 4 to 22 percent.

Overall performance. Prototype versions of a RISC I compiler for C, optimizer, linker, assembler, and simulator were developed early in the project to predict the code size and performance of RISC I. The minicomputers and microprocessors chosen for this comparison are described in Table 11. We didn't have working hardware for either the 68000 or RISC I, so we used simulators to predict performance. The cycle time for the first RISC I prototype is expected to be 400 nsec to read and add two 32-bit registers, store the result in a register, and prefetch the next instruction. This estimate is both optimistic and pessimistic: optimistic in that it is unlikely that students can successfully build something that fast on their first

Table 10.
Procedure call/return overhead
(including parameter passing).

	EXECUTION TIME (μSECS)	INSTRUCTIONS EXECUTED	DATA MEMORY ACCESSES
VAX-11	26	5	19
PDP-11	22	19	15
68000	19	9	12
RISC I	2	6	0.2

Table 11.
Characteristics of six machines.

	MICROPROCESSORS—NMOS VLSI			MINICOMPUTERS—SHOTTKY TTL MSI		
	RISC I	68000	Z8002	VAX-11/780	PDP-11/70	C/70
YEAR OF INTRODUCTION	1981	1980	1979	1978	1975	1980
BASIC INSTRUCTIONS	31	61	110	248	65	40
GENERAL REGISTERS	32	15	14	13	6	8
ADDRESSING MODES	2	14	12	18	12	17
ADDRESS SIZE (BITS)	32	24	16	32	16	20
BASIC CLOCK FREQUENCY	7.5MHz	10MHz	6MHz	5MHz	7.5MHz	6.7MHz
REG. TO REG. ADD (μsec)	0.4	0.4	0.7	0.4	0.5	?
MODIFY INDEX, BRANCH IF ZERO (BRANCH TAKEN)	1.2	1.0	2.2	1.4	0.8	?

tr
te

pi
E
tc
th
a

Z
U
br
in
ha
ge

cc
th
si

—
B
E
F
H
K
I
A
PI
PI
RI
SI
TC

—
B
E
F
H
K
I
A
RE
PL
PL
SE
TC

—
Se

try, and pessimistic because an experienced IC design team could build a much faster machine.

We chose 11 C programs for the performance comparison. The first five programs are HLL versions of the EDN benchmarks.²⁵ The other C programs range from toy programs (e.g., towers of Hanoi) to programs from the Unix environment that are used every day (e.g., SED, a batch-oriented text editor).

The compilers used are quite similar: the VAX, C/70, Z8002, 68000, and RISC I C compilers are based on a Unix portable C compiler,²⁶ and the one for the PDP-11 is based on the Ritchie C compiler.²⁷ Experiments comparing the Ritchie and portable C compilers for the PDP-11 have shown that the average difference in the size of generated code is within one percent.²⁸

Tables 12 and 13 compare the relative performance and code size of these minicomputers and microprocessors on the 11 C programs. A surprising result is that, even though size optimization was virtually ignored, RISC I programs

are—at worst—a factor of two larger than programs for the other machines. To us, the most important figure of merit for a new architecture is execution time. Table 13 shows that RISC I executes C programs faster than currently available microprocessors—faster even than most minicomputers.

Discussion

The presentation of the RISC concept has led to many stimulating discussions. Listed below are frequently heard comments followed by a short discussion of that comment.

CISCs provide better support of HLLs since they include HLL primitives (CASE, CALL).

CISC architectures support HLLs by narrowing the gap between the semantics of the assembly language and

Table 12.
C benchmarks: RISC I program size (in bytes) and RISC I size ratio.

BENCHMARK	RISC I	PROGRAM SIZE RELATIVE TO RISC I				
		68000	Z8002	VAX-11/780	11/70	C/70
E—STRING SEARCH	144	.8	.9	.7	.8	.7
F—BIT TEST	120	1.2	1.5	1.2	1.4	1.0
H—LINKED LIST	176	.7	.8	1.2	1.7	.8
K—BIT MATRIX	288	1.1	1.3	1.0	1.3	1.1
I—QUICKSORT	992	.7	1.1	.9	1.1	.9
ACKERMAN(3,6)	144	—	2.1	.5	.6	.5
PUZZLE(SUBSCRIPT)	2736	—	.5	.5	.6	.6
PUZZLE(POINTER)	2796	.9	.5	.5	.5	.6
RECURSIVE QSORT	752	—	.8	.6	.8	.6
SED(BATCH EDITOR)	17720	—	1.0	.6	.5	.5
TOWERS HANOI(18)	96	—	2.5	.8	1.0	.7
AVERAGE		.9±.2	1.2±.6	.8±.3	.9±.4	.7±.2

Table 13.
C benchmarks: RISC I execution time (in milliseconds) and RISC I performance ratio.

BENCHMARK	RISC I	NUMBER OF TIMES SLOWER THAN RISC I				
		68000	Z8002	VAX-11/780	11/70	C/70
E—STRING SEARCH	.46	2.8	1.6	1.3	0.9	2.2
F—BIT TEST	.06	4.8	7.2	4.8	6.2	9.2
H—LINKED LIST	.10	1.6	2.4	1.2	1.9	2.5
K—BIT MATRIX	.43	4.0	5.2	3.0	4.0	9.3
I—QUICKSORT	50.4	4.1	5.2	3.0	3.6	5.8
ACKERMAN(3,6)	3200	—	2.8	1.6	1.6	—
RECURSIVE QSORT	800	—	5.9	2.3	3.2	1.3
PUZZLE(SUBSCRIPT)	4700	—	4.2	2.0	1.6	3.4
PUZZLE(POINTER)	3200	4.2	2.3	1.3	2.0	2.1
SED(BATCH EDITOR)	5100	—	4.4	1.1	1.1	2.6
TOWERS HANOI(18)	6800	—	4.2	1.8	2.3	1.6
AVERAGE		3.5±1.8	4.1±1.6	2.1±1.1	2.6±1.5	4.0±2.8

the semantics of an HLL. Support can also, however, be measured as the inverse of the "costs" of using typical HLL constructs on a particular machine. If the architect provides a feature that "looks" like the HLL construct but runs slowly, the compiler writer will omit the feature or, worse, the HLL programmer concerned with performance will avoid the construct. A recent study shows that CISCs penalize the use of HLLs far more than RISCs.²⁹

It is more difficult to write a compiler for a RISC than a CISC.

A recent paper by Wulf³⁰ helps explain why this is not true. He says that compiling is essentially a large "case analysis." The more ways there are to do something (more instructions and addressing modes), the more cases must be considered. The compiler writer must balance the speed of the compiler with his desire to get good code. In CISCs there may not be enough time to analyze the potential usage of all available instructions. Thus, Wulf recommends, "There should be precisely one way to do something, or all ways should be possible." In RISC we have taken the former approach. There are few choices; for example, if an operand is in memory, it must first be loaded into a register. Simple case analysis implies a simple compiler, even if more instructions must be generated in each case.

RISC I is tailored to C and will not work well with other HLLs.

Studies of other HLLs^{23,31} indicate that the most frequently executed operations are the same simple HLL constructs found in C, for which RISC I has been optimized. Unless an HLL significantly changes the way people program, we expect to see similar results. For languages that have unusual data types, such as Cobol, we need to find the simple operations that are used repeatedly in that environment and incorporate them into a RISC. Even if the RISC I architecture does not map Cobol efficiently, we believe this philosophy can lead to a RISC that does.

Comparisons of RISC I with the VAX are unfair in that the VAX provides a virtual address space; RISC I would be much slower if it had virtual memory.

To answer the question "How much slower?" we looked at solutions used by other microprocessors. National Semiconductor has announced the 16082, a memory management chip with an address cache that normally translates virtual address into physical addresses in 100 nsec.³² If we were to put this chip in a system with a RISC CPU, it would add another 100 nsec to every memory access. Memory is referenced every 400 nsec in RISC I, so such a combination would reduce RISC performance by 20 percent. Because 80 to 90 percent of memory references in RISC I are to instructions,¹ more sophisticated approaches, such as translating addresses only when crossing a page boundary, might limit performance reduction to only five percent. A final observation is that even if the addition of virtual memory doubled the cycle

time of NMOS RISC I, it would still be faster than most present-day microprocessors.

The good performance is due to the overlapped register windows; the reduced instruction set has nothing to do with it.

Certainly, a significant portion of the speed is due to the overlapped register windows of RISC I. A key point is that there would have been no room for register windows if control had not dropped from 50 to 6 percent. Furthermore, control is so simple in RISC that microprogramming is unnecessary; this eliminates the control loop as the limiting factor of the machine cycle, as is frequently the case in microprogrammed machines.

There is no difference between overlapped register windows and a data cache.

A cache is ineffective if it is too small. An effective data cache would require a much larger area than our register file, especially if it must provide the same number of ports as the register file. The more complicated virtual address translation and decoding would likely stretch the basic CPU cycle time. Finally, the more complicated cache control would have extended the design phase of RISC I.

RISC I represents a new style of computers that take less time to build yet provide higher performance. While traditional machines "support" HLLs with instructions that look like HLL constructs, this machine supports the use of HLLs with instructions that HLL compilers can use efficiently. The loss of complexity has not reduced RISC's functionality; the chosen subset, especially when combined with the register window scheme, emulates more complex machines. It also appears we can build such a single-chip computer much sooner and with less effort than traditional architectures.

As we go to press, we are just testing the RISC I chips. Unfortunately, the polysilicon layer was processed improperly, and we believe this accounts for the fact that the chips are only partially operational. We have not yet found any circuit design errors.

This research area is by no means closed. For example, an investigation of a RISC with two ALU operations per cycle and dual-port main memory has begun at Stanford,³³ and we are working on a new implementation with a denser register file and a more sophisticated timing scheme.³⁴ Some of the other topics to be investigated include the applicability of RISCs to other HLLs (e.g., Lisp, Cobol, Ada), the effectiveness of an operating system on RISC (e.g., Unix), the architecture of coprocessors for RISC (e.g., graphics, floating point), migration of software to RISC (e.g., a 370 emulator written in RISC machine language), and the implementation of RISC in other technologies (CMOS, TTL, ECL). This list is too big for one project; we hope to cooperate with industry and academia in exploring RISC architectures. ■

Acknowledgments

The RISC Project has been sustained by a large group of volunteers. We would like to thank all those in the Berkeley community who have helped push RISC from a concept to a chip. We would also like to give special thanks to a few.

John Ousterhout created, maintained, and revised Caesar, our principal design aid, and consistently provided useful technical and editorial advice. Lloyd Dickman was actively involved with the design of RISC during his sabbatical at Berkeley, supplying technical and managerial expertise. We also want to thank Richard Newton for dedicating his VLSI class to the RISC project.

The RISC research was investigated over a four-quarter sequence of graduate courses at Berkeley. Many have participated but a few contributed significantly. Manolis Katevenis did the initial block structure and the initial timing description and provided many important simplifications and ideas about the implementation and the architecture. Ralph Campbell wrote the initial C compiler, the optimizer, assembler, and linker. Yuval Tamir wrote a simulator, ran the benchmarks,³⁵ and provided many suggestions in the initial design of RISC I. Gary Corcoran wrote the initial ISPS description of RISC I. Jim Peek, Korbin Van Dyke, John Foderaro, Dan Fitzpatrick, and Zvi Peshkess were the principal VLSI designers of the first RISC I chip. Michael Arnold, Dan Fitzpatrick, John Foderaro, and Howard Landman all wrote CAD tools that were crucial to the VLSI implementation of RISC I. Peter Kessler helped derive the overlapped register windows and helped with the CAD software. Jim Beck and Bob Cmelik created the VLSI testing hardware and software. Bob Sherburne is currently working with Katevenis on a more efficient VLSI implementation of RISC. Earl Cohen and Neil Soiffer collected statistics on C programs, and Shafi Goldwasser collected similar statistics for Pascal.

We would also like to thank Korbin Van Dyke for his useful suggestions on improving this paper.

This research was funded in part by the Defense Advance Research Projects Agency, ARPA Order No. 3803, and monitored by the Naval Electronic System Command under Contract No. N00039-78-G-0013-0004. We would like to thank Duane Adams, Paul Losleben, and DARPA for providing the resources that allow universities to attempt projects involving high risk.

References

1. W. D. Strecker, "VAX-11/780: A Virtual Address Extension to the DEC PDP-11 Family," *AFIPS Conf. Proc.*, Vol. 47, 1978 NCC, pp. 967-980.
2. B. G. Utley et al., *IBM System/38 Technical Developments*, IBM GS80-0237, 1978.
3. P. Tyner, *iAPX-432 General Data Processor Architecture Reference Manual*, Order No. 171860-001, Intel, Santa Clara, Calif., 1981.
4. E. Organick, *A Programmer's View of the Intel 432 System*, McGraw-Hill, Hightstown, N.J., 1982.
5. D. A. Patterson and D. R. Ditzel, "The Case for the Reduced Instruction Set Computer," *Computer Architecture News*, Vol. 8, No. 6, Oct. 15, 1980, pp. 25-33.
6. D. A. Patterson and C. H. Séquin, "Design Considerations for Single-Chip Computers of the Future," *IEEE Trans. Computers*, Joint Special Issue on Microprocessors and Microcomputers, Vol. C-29, No. 2, pp. 108-116.
7. D. R. Ditzel and D. A. Patterson, "Retrospective on High-Level Language Computer Architecture," *Proc. Seventh Annual Int'l Symp. Computer Architecture*, May 6-8, 1980, pp. 97-104.
8. W. Wulf, private communication, Nov. 1980.
9. A. Lunde, "Empirical Evaluation of Some Features of Instruction Set Processor Architecture," *Comm. ACM*, Mar. 1977, Vol. 20, No. 3, pp. 143-153.
10. B. A. Wichmann, "Ackermann's Function: A Study in the Efficiency of Calling Procedures," *BIT*, Vol. 16, No. 1, Jan. 1976, pp. 103-110.
11. F. Baskett, "A VLSI Pascal Machine," public lecture, University of California, Berkeley, Fall 1978.
12. R. L. Sites, "How to Use 1000 Registers," *Caltech Conf. VLSI*, Jan. 1979.
13. D. Halbert and P. Kessler, *Windows of Overlapping Registers*, CS292R Final Reports, June 9, 1980.
14. D. Morris and R. N. Ibbett, *The MU-5 Computer System*, Springer-Verlag, New York, 1979.
15. F. C. Williams and T. Kilburn, "The University of Manchester Computing Machine," *Inaugural Conf. Manchester University Computer*, July 1951, pp. 5-11.
16. "Altering Computer Architecture is Way to Raise Throughput, Suggests IBM Researchers," *Electronics*, Vol. 49, No. 25, Dec. 23, 1976, pp. 30-31.
17. "IBM Mini a Radical Departure," *Datamation*, Vol. 25, No. 11, Oct. 79, pp. 53-55.
18. G. Radin, "The 801 Minicomputer," *Proc. Symp. Architectural Support for Programming Languages and Operating Systems*, Mar. 1-3, 1982.
19. R. V. Orlando and T. L. Anderson, "An Overview of the 9900 Microprocessor Family," *IEEE Micro*, Vol. 1, No. 3, Aug. 1981, pp. 38-42.
20. D. T. Fitzpatrick et al., "A RISCy Approach to VLSI," *VLSI Design*, Vol. 2, No. 4, Oct. 81, pp. 14-20.
21. W. W. Lattin et al., "A Methodology for VLSI Chip Design," *Lambda—The Magazine of VLSI Design*, Second Quarter 1981, pp. 34-44.
22. J. Ousterhout, "Caesar: An Interactive Editor for VLSI Circuits," *VLSI Design*, Vol. 2, No. 4, Nov. 1981, pp. 34-38.
23. W. C. Alexander and D. B. Wortman, "Static and Dynamic Characteristics of XPL Programs," *Computer*, Vol. 8, No. 11, Nov. 1975, pp. 41-46.
24. L. Shustek, "Analysis and Performance of Computer Instruction Sets," PhD Thesis, Stanford University, Jan. 1978.
25. R. G. Grappel and J. E. Hemmingsway, "A Tale of Four Microprocessors: Benchmarks Quantify Performance," *Electronic Design News*, Vol. 26, No. 7, Apr. 1, 1981, pp. 179-265.
26. S. C. Johnson, "A Portable Compiler: Theory and Practice," *Proc. Fifth Annual ACM Symp. Programming Languages*, Jan. 1978, pp. 97-104.
27. B. W. Kernighan and D. M. Ritchie, *The C Programming Language*, Prentice-Hall, Englewood Cliffs, N.J., 1978.
28. S. C. Johnson, private communication, Jan. 1981.

29. D. A. Patterson and R. S. Piepho, "RISC Assessment: A High-Level Language Experiment," *Proc. Ninth Int'l Symp. Computer Architecture*, Apr. 26-29, 1982, pp. 3-8. (Scheduled to appear in an upcoming issue of *IEEE Micro*.)
30. W. A. Wulf, "Compilers and Computer Architecture," *Computer*, Vol. 14, No. 7, July 1981, pp. 41-48.
31. D. R. Ditzel, "Program Measurements on a High-Level Language Computer," *Computer*, Vol. 13, No. 8, Aug. 1980, pp. 62-72.
32. Y. Lavi et al., "16-bit Microprocessor Enters Virtual Memory Domain," *Electronics*, Vol. 53, No. 9, Apr. 24, 1980, pp. 123-129.
33. J. Hennessy et al., "The MIPS Machine," *Digest of Papers Compton Spring 82*, Feb. 1982, pp. 2-7.
34. R. W. Sherburne et al., "Datapath Design for RISC," *Proc. Conf. Advanced Research in VLSI*, Jan. 25-27, 1982, pp. 53-62.
35. Y. Tamir, "Simulation and Performance Evaluation of the RISC Architecture," Electronics Research Laboratory Memorandum No. UCB/ERL M81/17, University of California, Berkeley, Mar. 1981.



David A. Patterson has been a member of the faculty in the Computer Science Division, Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, since 1977. He was named associate professor in 1981 and currently teaches computer architecture at the graduate and undergraduate levels. His research combines popular software, experimental architecture, and VLSI to create more effective computer systems.

Patterson spent the fall of 1979 on leave of absence at Digital Equipment Corporation developing microprogram design tools and reviewing computer designs. In the next academic year he developed courses that led to the design and implementation of RISC I, a 45,000-transistor microprocessor. In 1982 he received the Distinguished Teaching Award from the Berkeley division of the Academic Senate of the University of California. Patterson received a BA in mathematics and an MS and PhD in computer science from UCLA.



Carlo H. Séquin is a professor of computer science at the University of California, Berkeley. Since 1980, he has headed the Computer Science Division as associate chairman for computer sciences in the Department of Electrical Engineering and Computer Science. He joined the faculty in 1977. His research interests lie in the field of computer architecture and design tools for very large scale integrated systems. In particular, his research concerns multi-microprocessor computer networks, the mutual influence of advanced computer architecture and modern VLSI technology, and the implementation of special functions in silicon.

From 1970 to 1976, Séquin worked on the design and investigation of charge-coupled devices for imaging and signal processing applications at Bell Telephone Laboratories, Murray Hill, New Jersey. He has written many papers in that field and is an author of the first book on charge-transfer devices. Séquin received his PhD in experimental physics from the University of Basel, Switzerland, in 1969. He is a member of the ACM and the Swiss Physical Society, and a fellow of the IEEE.

COMPUTER ARCHITECTS:

Do you welcome the challenge of directing the future of large-scale computer systems? Are you motivated in a small, select group of talented colleagues? Do you enjoy working in a highly visible environment? Then consider Amdahl. Our business depends on skilled individuals who can advise our company officers on the development of architecture concepts and designs. Our Architecture Group is responsible for establishing product architecture directions and works closely with key designers and corporate planners to implement machines based on proprietary circuit technology.

OUR BUSINESS DEPENDS ON YOU.

If your background is in computer architecture, communications architecture, software systems design, processor organization, performance modeling or in the evaluation of hardware technologies, consider Amdahl. We offer the advantage of a career in a dynamic, entrepreneurial environment. The significance of the individual is the hallmark of Amdahl's progressive employee relations, and rewards are based on personal contribution.

Consider the potential for your career at Amdahl. Send your resume to *Elaine Patton, Amdahl Corporation, Dept. 9-186, P.O. Box 470, Sunnyvale, California 94086*. We are an equal opportunity employer through affirmative action.

amdahl
We offer an alternative.