

The S-Machine, an Architecture for Symbolic Processing

Cristian Băleanu, Dan Tomescu

Ro-Micro, June 28, 2024

Context

- Computer architecture was a remarkably active field of research in the 1980s, primarily due to impressive advances in circuit technology. Much of the flurry of research of that time was related to investigating the relationship between architecture and programming languages, i.e. identifying the kind of hardware/software interfaces that would optimize the execution of (compiled) code written in various high level languages.
- Our interest in computer architecture dates from the late 1970s, when we were part of a Bucharest Politehnica University team that designed and built the production prototype of DIAGRAM, a graphics workstation produced and commercialized by FEPER. DIAGRAM was built around a Z80-based bi-processor machine featuring a hardware accelerator for graphics.

Projects (1)

- After the DIAGRAM project we joined **ICPE**, a research institute for electrical engineering, to do applied research work on innovative computer architectures. Our goal was to build, from the ground up, a complete solution for a real world problem that would be intractable using stock hardware.
- Some local market research brought us in front of such a customer, **Petromar**, a Romanian company doing off-shore oil drilling and exploration work in the Black Sea. Petromar's CEO had been impressed by work done in the West with expert systems for oil drilling, but commercial embargoes in force at that time (1985) did not allow him to consider the purchase of a turnkey solution. He accepted our project proposal to build an original solution around a novel architecture, the *S-machine*, that would be capable of running the expert system (Prolog) code within the real-time constraints imposed by their application.
- **The S-machine was designed to be used as an “intelligent” hardware accelerator connected to a general-purpose computer through a high-speed interface and fully integrated with the host operating system.** DEC PDP 11 running Unix and IBM PC were the two types of computers that could be connected to the S-machine. Code libraries written in C were designed and implemented to make the use of the S-machine transparent to programs running on host machines, i.e. the S-machine was seen as a regular API (application programming interface). While the hardware team was working on the implementation of the machine, a software simulator was being used to support the development of the Prolog programming environment. At the same time, modules for the expert system to be run on host machines could be implemented and tested in a simulated environment.

Projects (2)

- In 1989, while the Petromar project was in an advanced stage, our team started to work on a new job. ITCI, a research institute for computing technology, was interested to bid for a request for a Lisp machine coming through Comecon (the Eastern block trade organization) from a **Soviet research institute for the aerospace industry**. ITCI asked us to join forces with them to write a project proposal that was accepted by the customer. As a consequence, part of our team started to work with the ITCI team on this new project and a new version of the S-machine was about to be built. The Lisp programming environment was on a path very similar to that successfully taken to build the Prolog system: implementing a Lisp runtime environment on the S-machine and combining it with modules developed for conventional implementations.
- **1990 was a very eventful year in Eastern Europe**. As a consequence, progress on our projects slowed down considerably, before coming to a halt (e.g. Comecon was disbanded and took with it our contract with the Moscow institute, but not before successfully reaching the first milestone).
- In conclusion, our S-machine projects were a success while they lasted. We managed to successfully build a hardware machine having a novel architecture together with its associated software and our work was fully funded by real customers. However, our goal to implement the machine in VLSI technology and turn it into a commercial product had to be abandoned.

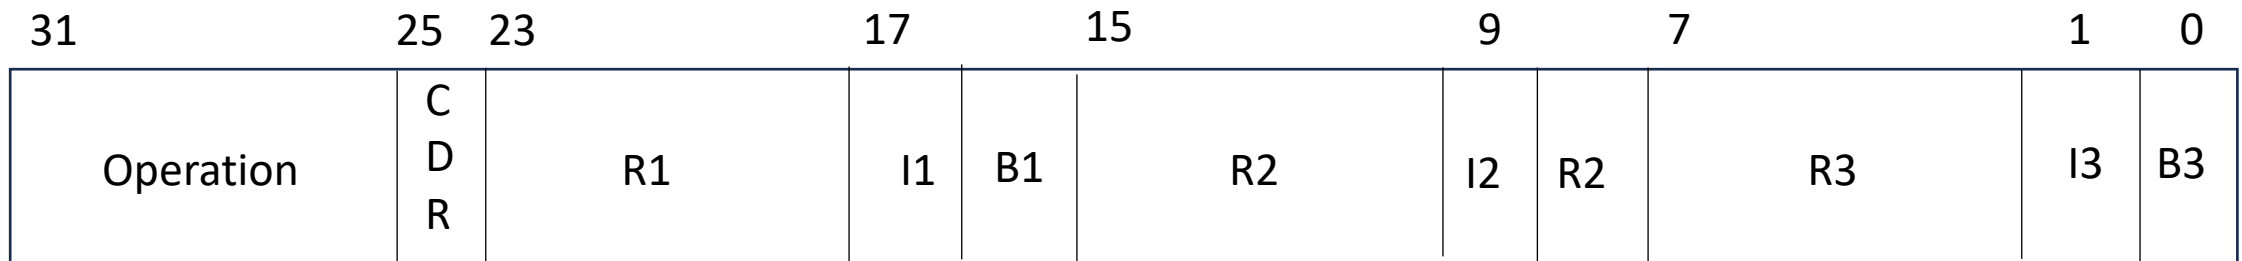
The S-Machine (1)

- The S-machine was a **32-bit computer with hardware support for symbolic languages** like Lisp and Prolog, but able to run with no time penalty compiled code for programs written in procedural languages like C. Lisp implementations were usually inefficient because the control of program execution on conventional machines is implicitly sequential and memories are linearly organized, while in Lisp both data structures (lists) and evaluation procedures are recursive. We did not want to restrict our machine to Lisp, but similar problems were posed by Prolog, so we chose to design a general-purpose architecture with a reduced set of instructions and hardware support for recursive control strategies and compact list representation of both data and programs.
- Given a list l , $\text{car}(l)$ is a function that selects the first element of l ; $\text{cdr}(l)$ is a function whose value is l without its first element (e.g. if $l = (a, b, (c, d))$, $\text{car}(l) = a$, $\text{cdr}(l) = (b, (c, d))$). Lists are represented usually by means of cells made up of two pointers each, CAR and CDR. However, statistical studies carried out on large Lisp programs showed that if lists were linearized, more than 98% of list CDRs pointed to the next cell. As a consequence, a technique called **CDR-coding** can be used to represent lists: every pointer P is represented in a separate memory word together with two additional bits encoding information about CDRs of list cells, i.e. “CAR is P , CDR is in the next word”, “CAR is P , CDR is nil (empty list)”, “CAR is P , CDR is the next word”, “the cell is relocated at P ”.

The S-Machine (2)

- Memory words are 32 bit long and contain a 24 bit pointer field (P), a two bit CDR field (used for CDR coding) and a six bit data type field identifying the type of data pointed to by P. When lists are used to store machine programs, the data type field is interpreted as an operation code, while the pointer field holds the operands. The representation of programs by means of lists implies that, unlike in the case of conventional computers, on the S-machine any program, regardless of the language it is written in, can be manipulated by any other program.
- While the logical address space of the S-machine is continuous, the first 16 Kwords were implemented in fast memory. The machine had **128 general-purpose registers, out of which 64 were global and 64 local**. The global register set takes up the first 64 words of fast memory, while the local register set specifies a window made up of any 64 consecutive memory words, with addresses calculated by adding an offset (derived from the register number) to the contents of a base register. The local register set can be used to hold local variables of functions and procedures: if the base register is appropriately set on procedure entry, then local variables are accessible via register references. For efficiency reasons, live contexts should reside in fast memory: it is up to compilers to generate code for dealing with fast memory management.

List representation of data/programs



The S-Machine (3)

- The S-machine has a set of **61 simple hardwired instructions** with a uniform representation. Instructions are register-to-register (except for jumps). Besides the usual data transfer, logical, and arithmetic instructions, two instructions implement the most frequently used list operations: *car ri, rj* and *cdr ri, rj*. A wide range of jump instructions, including 2^n jumps, some of them using masks to test selected bits of special purpose state registers, can be used for efficient dynamic data type checking, which is a vital part of the runtime support for symbolic languages. At the hardware implementation level, **most functional modules have a symmetric structure**, resulting in increased efficiency of program execution. Main memory is interleaved while the fast memory is made up of two identical banks containing the same information (two simultaneous read operations per memory cycle can therefore be executed, with write operations altering corresponding addresses from both banks). Efficiency gains brought about by the dual hardware structure include simultaneous access to left and right operands residing in global registers and local registers allocated in fast memory, the execution of parallel operations in the two sections of the processor, and fast transfers of compact data blocks. Prefetching of instructions is favoured by the compact list representation of programs.
- **The S-machine prototype was built in TTL-S and TTL-LS technology.** The processor is sequentially controlled by a 100 ns single phase clock. The fast memory has a 100 ns cycle and is built out of TTL-S chips, while the main memory is implemented in MOS DRAM technology and has a 350 ns cycle.