20 Years-Unleashing the Power of HPC

SC1988

1988 Chair
George Michael
Orlando, FL

SUPERCOMPUTING '88

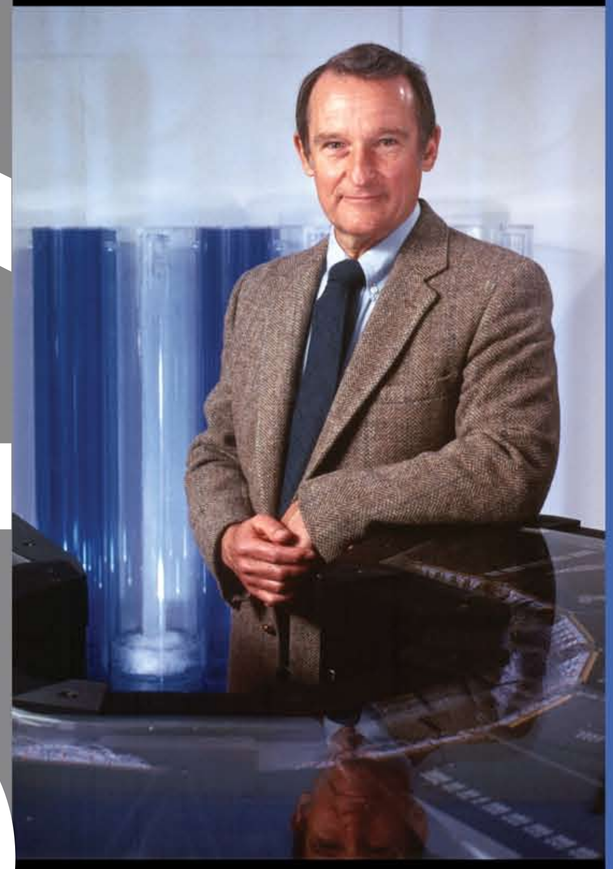

Seymour R. Cray

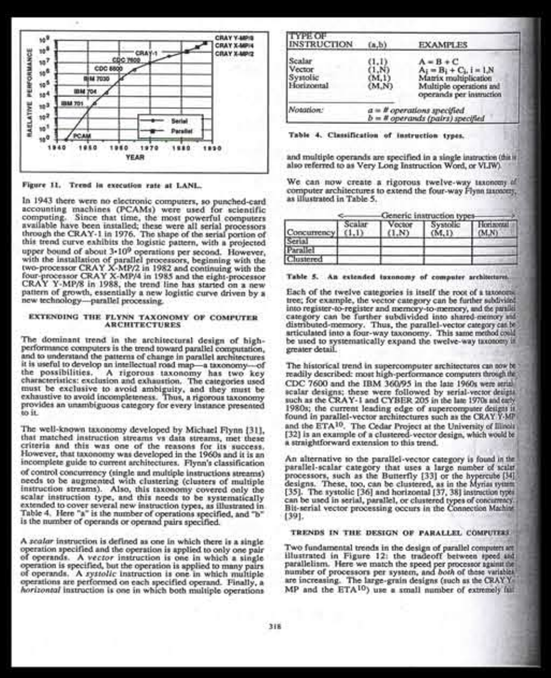

Figure 11. Trend in execution rate at LANL.

In 1943 there were no electronic computers, so punched-card accounting machines (PCAMs) were used for scientific computing. Since that time, the most powerful computers available have been installed; these were all serial processors through the CRAY-1 in 1976. The shape of the serial portion of this trend curve exhibits the logistic pattern, with a projected upper bound of about 3×10^9 operations per second. However, with the installation of parallel processors, beginning with the two-processor CRAY X-MP/2 in 1982 and continuing with the four-processor CRAY X-MP/4 in 1985 and the eight-processor CRAY Y-MP/8 in 1988, the trend line has started on a new pattern of growth, essentially a new logistic curve driven by a new technology—parallel processing.

EXTENDING THE FLYNN TAXONOMY OF COMPUTER ARCHITECTURES

The dominant trend in the architectural design of high-performance computers is the trend toward parallel computation, and to understand the patterns of change in parallel architectures it is useful to develop an intellectual road map—a taxonomy—of the possibilities. A rigorous taxonomy has two key characteristics: exclusion and exhaustion. The categories used must be exclusive to avoid ambiguity, and they must be exhaustive to avoid incompleteness. Thus, a rigorous taxonomy provides an unambiguous category for every instance presented to it.

The well-known taxonomy developed by Michael Flynn [31], that matched instruction streams vs data streams, met these criteria and this was one of the reasons for its success. However, that taxonomy was developed in the 1960s and it is an incomplete guide to current architectures. Flynn's classification of control concurrency (single and multiple instructions streams) needs to be augmented with clustering (clusters of multiple instruction streams). Also, this taxonomy covered only the scalar instruction type, and this needs to be systematically extended to cover several new instruction types, as illustrated in Table 4. Here "a" is the number of operations specified, and "b" is the number of operands or operand pairs specified.

A *scalar* instruction is defined as one in which there is a single operation specified and the operation is applied to only one pair of operands. A *vector* instruction is one in which a single operation is specified, but the operation is applied to many pairs of operands. A *systolic* instruction is one in which multiple operations are performed on each specified operand. Finally, a *horizontal* instruction is one in which both multiple operations

TYPE OF INSTRUCTION	(a,b)	EXAMPLES
Scalar	(1,1)	$A = B + C$
Vector	(1,N)	$A_i = B_i + C_i$, $i = 1,N$
Systolic	(M,1)	Matrix multiplication
Horizontal	(M,N)	Multiple operations and operands per instruction
Notation:		*a = # operations specified*
		b = # operands (pairs) specified

Table 4. Classification of instruction types.

and multiple operands are specified in a single instruction (this is also referred to as Very Long Instruction Word, or VLIW).

We can now create a rigorous twelve-way taxonomy of computer architectures to extend the four-way Flynn taxonomy, as illustrated in Table 5.

	<—Generic instruction types—			
Concurrency	Scalar (1,1)	Vector (1,N)	Systolic (M,1)	Horizontal (M,N)
Serial				
Parallel				
Clustered				

Table 5. An extended taxonomy of computer architectures.

Each of the twelve categories is itself the root of a taxonomic tree; for example, the vector category can be further subdivided into register-to-register and memory-to-memory, and the parallel category can be further subdivided into shared-memory and distributed-memory. Thus, the parallel-vector category can be articulated into a four-way taxonomy. This same method could be used to systematically expand the twelve-way taxonomy in greater detail.

The historical trend in supercomputer architectures can now be readily described: most high-performance computers through the CDC 7600 and the IBM 360/95 in the late 1960s were serial-scalar designs; these were followed by serial-vector designs such as the CRAY-1 and CYBER 205 in the late 1970s and early 1980s; the current leading edge of supercomputer designs is found in parallel-vector architectures such as the CRAY Y-MP and the ETA10. The Cedar Project at the University of Illinois [32] is an example of a clustered-vector design, which would be a straightforward extension to this trend.

An alternative to the parallel-vector category is found in the parallel-scalar category that uses a large number of scalar processors, such as the Butterfly [33] or the hypercube [34] designs. These, too, can be clustered, as in the Myrias system [35]. The systolic [36] and horizontal [37, 38] instruction types can be used in serial, parallel, or clustered types of concurrency. Bit-serial vector processing occurs in the Connection Machine [39].

TRENDS IN THE DESIGN OF PARALLEL COMPUTERS

Two fundamental trends in the design of parallel computers are illustrated in Figure 12: the tradeoff between speed and parallelism. Here we match the speed per processor against the number of processors per system, and both of these variables are increasing. The large-grain designs (such as the CRAY Y-MP and the ETA10) use a small number of extremely fast

316

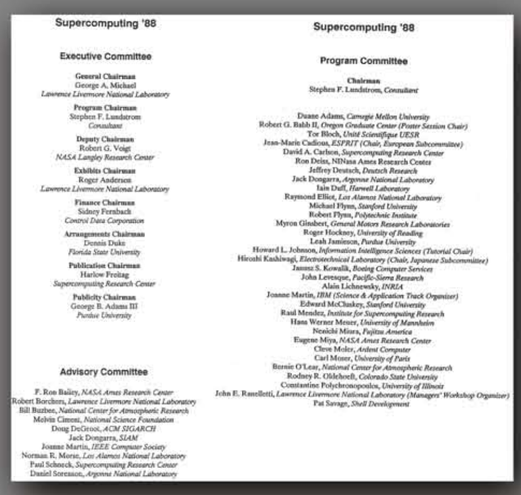

Supercomputing '88

Executive Committee

General Chairman
George A. Michael
Lawrence Livermore National Laboratory

Program Chairman
Stephen F. Lundstrom
Consultant

Deputy Chairman
Robert G. Voigt
NASA Langley Research Center

Exhibits Chairman
Roger Anderson
Lawrence Livermore National Laboratory

Finance Chairman
Sidney Fernbach
Control Data Corporation

Arrangements Chairman
Dennis Duke
Florida State University

Publication Chairman
Harlow Freitag
Supercomputing Research Center

Publicity Chairman
George B. Adams III
Purdue University

Advisory Committee

F. Ron Bailey, *NASA Ames Research Center*
Robert Borchers, *Lawrence Livermore National Laboratory*
Bill Buzbee, *National Center for Atmospheric Research*
Melvin Cimenti, *National Science Foundation*
Doug DeGroot, *ACM SIGARCH*
Jack Dongarra, *SIAM*
Joanne Martin, *IEEE Computer Society*
Norman R. Morse, *Los Alamos National Laboratory*
Paul Schneck, *Supercomputing Research Center*
Daniel Stevenson, *Argonne National Laboratory*

Supercomputing '88

Program Committee

Chairman
Stephen F. Lundstrom, *Consultant*

Duane Adams, *Carnegie Mellon University*
Robert G. Babb II, *Oregon Graduate Center (Poster Session Chair)*
Tor Bloch, *Unitil Scientifique URSR*
Jean-Marie Cadiou, *ESPRIT (Chair, European Subcommittee)*
David A. Carlson, *Supercomputing Research Center*
Ron Dror, *NASA Ames Research Center*
Jeffrey Deutsch, *Deutsch Research*
Jack Dongarra, *Argonne National Laboratory*
Iain Duff, *Harwell Laboratory*
Raymond Elliot, *Los Alamos National Laboratory*
Michael Flynn, *Stanford University*
Robert Flynn, *Polytechnic Institute*
Myron Ginsberg, *General Motors Research Laboratories*
Roger Hockney, *University of Reading*
Leah Jamieson, *Purdue University*
Howard L. Johnson, *Information Intelligence Sciences (Tutorial Chair)*
Hiroshi Kashiwagi, *Electrotechnical Laboratory (Chair, Japanese Subcommittee)*
Jasmin S. Kowalik, *Boeing Computer Services*
John Levesque, *Pacific-Sierra Research*
Alain Lichnewsky, *INRIA*
Joanne Martin, *IBM (Science & Application Track Organizer)*
Edward McCluskey, *Stanford University*
Raul Mendez, *Institute for Supercomputing Research*
Hans Werner Meuer, *University of Mannheim*
Noriaki Mura, *Fujitsu America*
Eugene Miya, *NASA Ames Research Center*
Clive Molex, *Ardent Computer*
Carl Moser, *University of Paris*
Bernie O'Lear, *National Center for Atmospheric Research*
Rodney R. Oldehoeft, *Colorado State University*
Constantine Polychronopoulos, *University of Illinois*
John E. Rastelletti, *Lawrence Livermore National Laboratory (Managers' Workshop Organizer)*
Pat Savage, *Shell Development*

1988

Notable Systems first mentioned this year in the proceedings:
- Cray X/MP and Y-MP
- Connection Machine CM-2
- Other Cray, Convex, Alliant, and IBM systems

Not mentioned in the proceedings after this year
- Amdahl VP-E
- CDC 7600, Cyber 700/800/990/995
- Apollo DN series
- Gould UTX
- Stellar GS
- Unisiys ISP

Noteworthy Architecture Topics
- Dataflow machines
- Vectorization

Notable Operating Systems
- CTSS
- IBM MVS and VM
- Unix

Research Machines
- Illinois CEDAR

Notable Programming Languages
- Lisp
- Lisp/Scheme
- Pascal
- Linda
- C
- Fortran

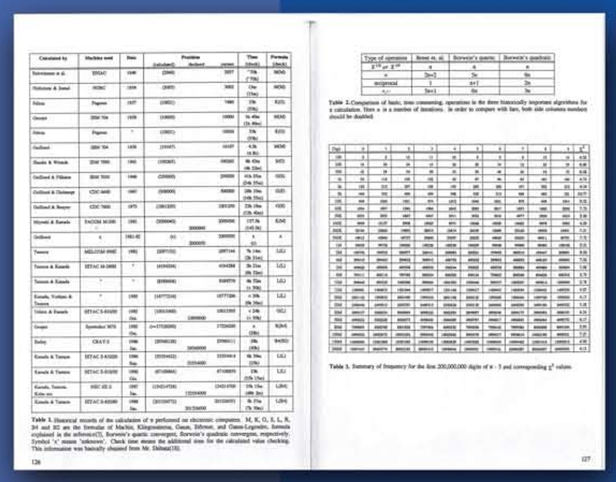

Table 3. Summary of frequency for the first 200,000,000 digits of π - 3 and corresponding g^k values.