The APE Experience

Nicola Cabibbo

Università di Roma “La Sapienza”
INFN — Sezione di Roma

apeNEXT: Computational Challenges
and First Physics Results

Birth and early growth of LQCD

- K. Wilson 1974: Introduces LQCD.
- M. Creutz 1979: C. Rebbi 1980 First Monte-Carlo simulations.
- H. Hamber and G. Parisi 1981, D. Weingarten 1982: Quenched approximation.
- N. C., G. Martinelli, R. Petronzio 1983: Simulation of weak interactions.

The need for computer power!

> In the next Section we outline the method of computation. In section 3 we describe the numerical results based on 14 link configurations in a 10^3x20 lattice, for which the Wilson quark propagators were already available /6/. This analysis is

Cray vs. Dedicated Supercomputers

The top commercial machine was the Cray: $\approx$ 1 Gflops for $\approx$ 20 G$. Available in Italy at CINECA ($\approx$ 300 KLire/hour), at CEA in Paris, etc. Difficult or expensive access to University groups.

The alternative was offered by home-brew machines. The shining example was the CERN/SLAC "3081/E Emulator" project. Emulator farms were the forerunner of modern PC-Clusters, and were widely used, e.g. in LEP experiments.

APE was conceived in october 1984 with the aim to be as powerful as a contemporary Cray (1 Gflops) but at a fraction of the price.

The "3081/E Emulator" was a great inspiration for the inception of the APE project; parts of the Emulator machines — the integer board, the crates, etc. — were used in the first APE.

APE over the years.

Other important projects of the 80's:

- Columbia University machine – N. H. Christ
- QCD-PACS in Japan
- GF11 at IBM – D. Weingarten

Il team del primo APE

P. Bacilieri
INFN-CNAF, Bologna, Italy

S. Cabasino, A. Frighi, F. Marzano, N. Matone, P. S. Paolucci, S. Petrarca, G. Salina
INFN, Sezione di Roma, Italy

N. Cabibbo, E. Marinari, G. Parisi
Dipartimento di Fisica, II Univerita' di Roma "Tor Vergata"; INFN, Sezione di Roma, Italy

F. Costantini, G. Fiorentini, S. Galeotti, D. Passuello, R. Tripiccione
Dipartimento di Fisica, Univerita' di Pisa; INFN, Sezione di Pisa, Italy

A. Fucci, R. Petronzio, F. Rapuano
CERN, Geneva, Switzerland

D. Pascoli, P. Rossi
Dipartimento di Fisica, Univerita' di Padova; INFN, Sezione di Padova, Italy

E.Remiddi
Dipartimento di Fisica, Univerita' di Bologna; INFN-CNAF, Bologna, Italy; INFN, Sezione di Bologna, Italy

R.Rusack
Rockefeller University, New York, U.S.A.

B.Tirozzi
Dipartimento di Matematica - Universita' "La Sapienza" Roma, Italy

---

Many younger people joined the group over the years, often as thesis students, probaly close to 100 by now.

The Processor of the first APE

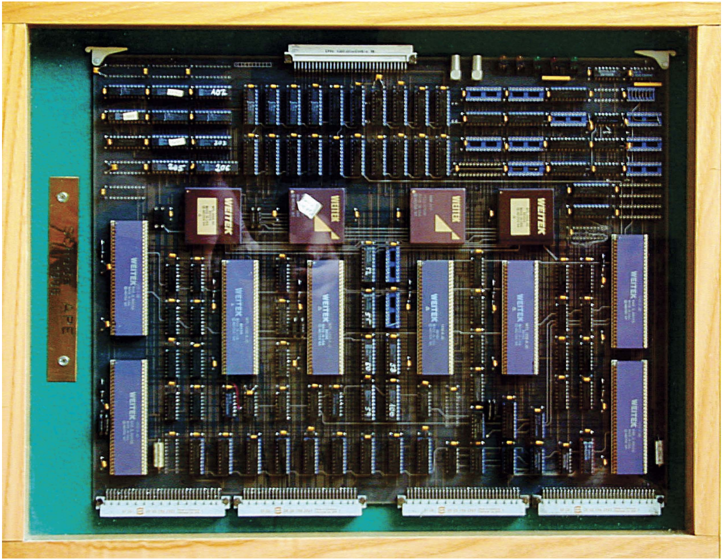

The Processor of the first APE

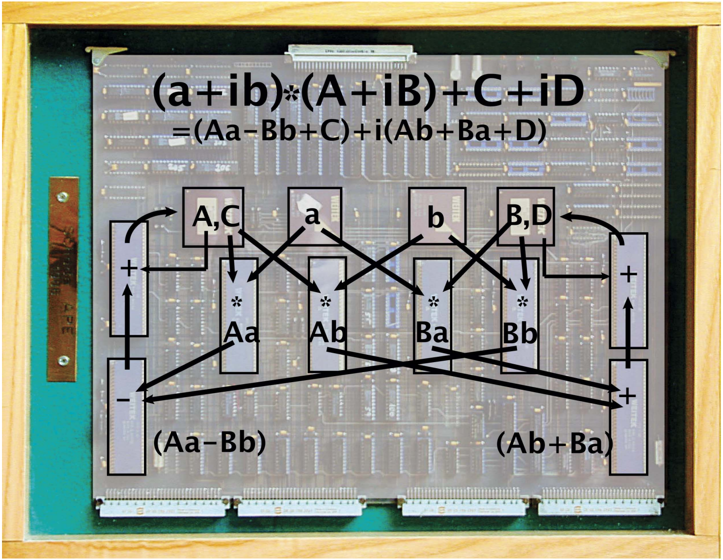

APE100 — Processor on Chip

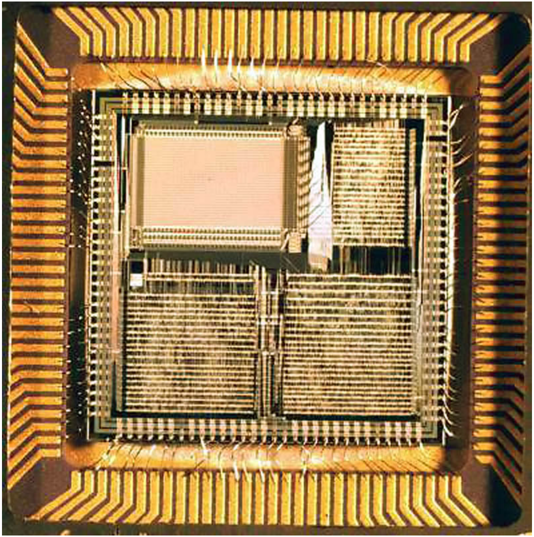

VLIW — Very Long Instruction Word

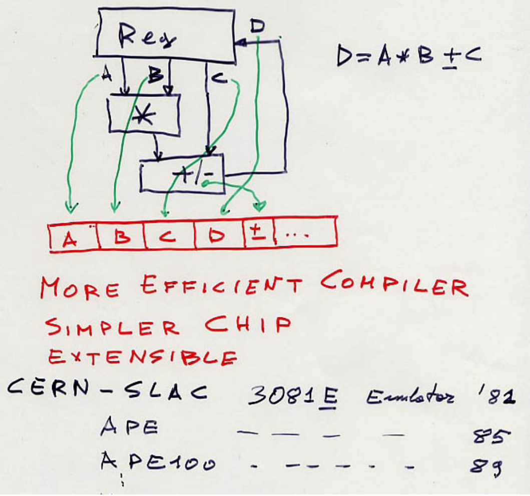

The Very Long Instruction Word structure was borrowed from the 3081/E design.

VLIW simplifies the processor structure — no (or minimal) instruction decoding — and significantly reduces power consumption.

The Pipeline and the Normal Operation

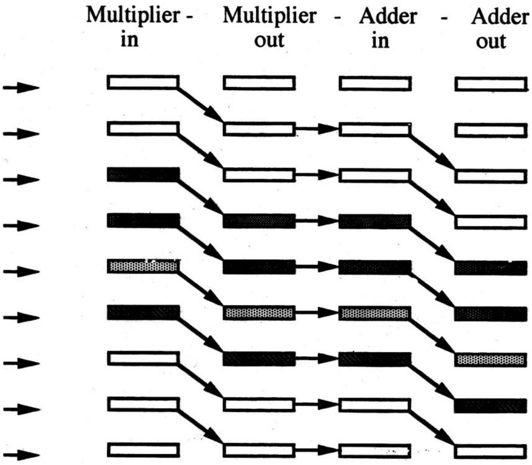

Not providing separate add and multiply instructions actually *improves* the processor efficiency, as it helps filling the pipeline.

SIMD — Single Instruction Multiple Data

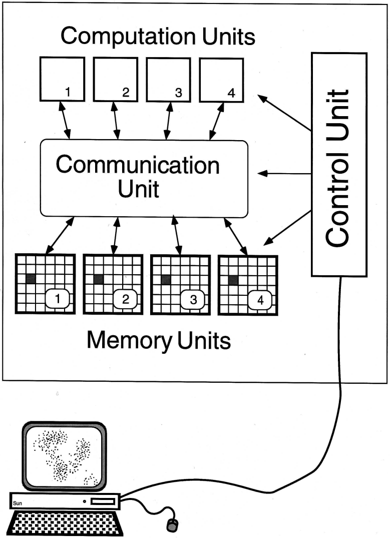

J&T: Aritmetica

•Pipelined "normal" complessa a*b+c (8 flops) per cycle

4 multiplier

4 adder/sub

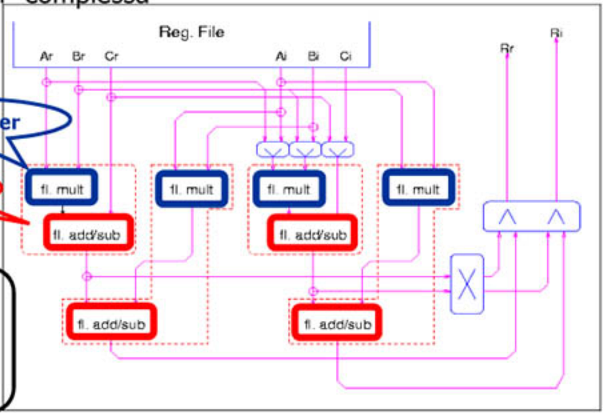

A 200 MHz (fully piped) = 1.6 GFlops

Saving Energy and Saving Space

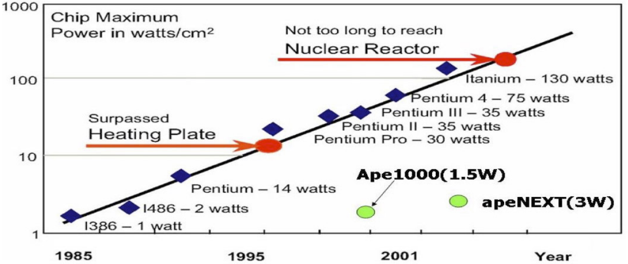

Machine	RLX TM5600	RLX TM5800	Avalon	ASCI Red	ASCI White	apeNEXT
Performance (Gflops)	21.4	3.3	17.6	600	2500	3670
Power (kilowatts)	5.2	0.52	18.0	1200	2000	80
Perf/Power (Mflops/watt)	4.12	6.35	0.978	0.5	1.25	46

Table 4. Performance-Power Ratio for Five Parallel-Computing Systems

Machine	RLX TM5600	RLX TM5800	Avalon	ASCI Red	ASCI White	apeNEXT
Performance (Gflops)	21.4	3.3	17.6	600	2500	3670
Area (feet^2)	6	6	120	1600	9920	72
Perf/Power (Mflops/feet^2)	3500	550	150	375	252	50972

Table 5. Performance-Space Ratio for Five Parallel-Computing Systems

APE in Numbers

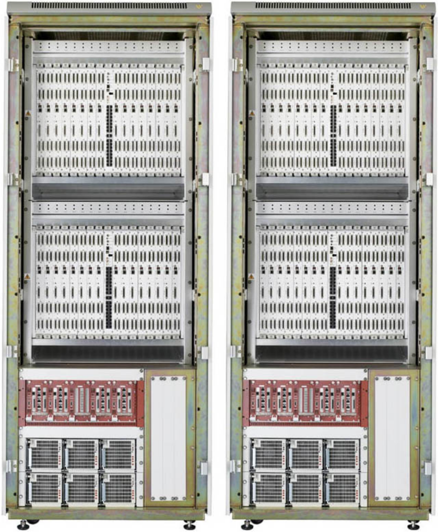

Apemille (2000):

	Italy	1365 GF
	Germany	650 GF
	UK	65 GF
	France	16 GF
Total		2 TF

apeNEXT (2005):
Development costs = 2000 k€uro
1100 k€uro VLSI NRE
250 k€uro non-VLSI NRE
650 k€uro prototype procurement
Manpower = 20 man/year
Mass production cost ~ 0.5 €uro/Mflops
Installations:

	Italy	10.6 TF
	Germany	8.0 TF
	France	1.6 TF
Total		20.2 TF

A Future for APE?

The question should be turned around:

Is there a physics problem which is worth pursuing
and requires 10 or 100 times the computational power of apeNEXT?

If the answer is yes, the APE way is probably
still today the best way to do it.

The question is still open,
but in the mean time, let us get the best of apeNEXT.

Designing APE

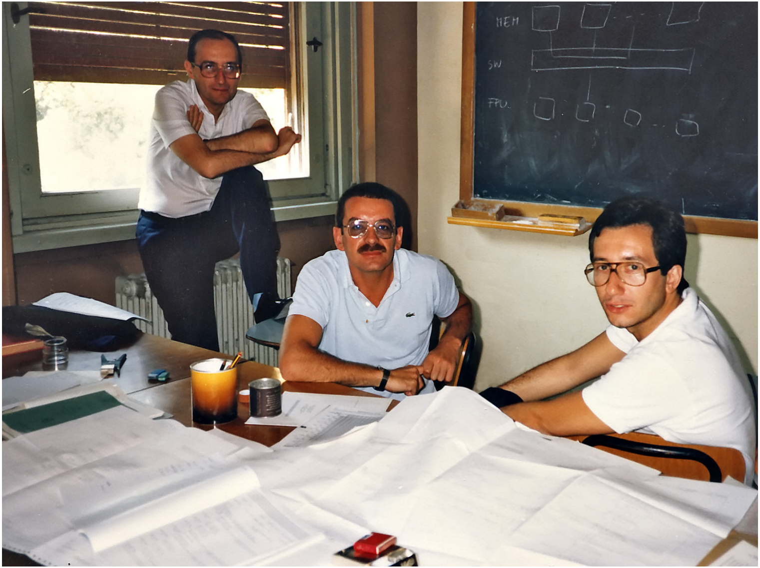

Getting it Right

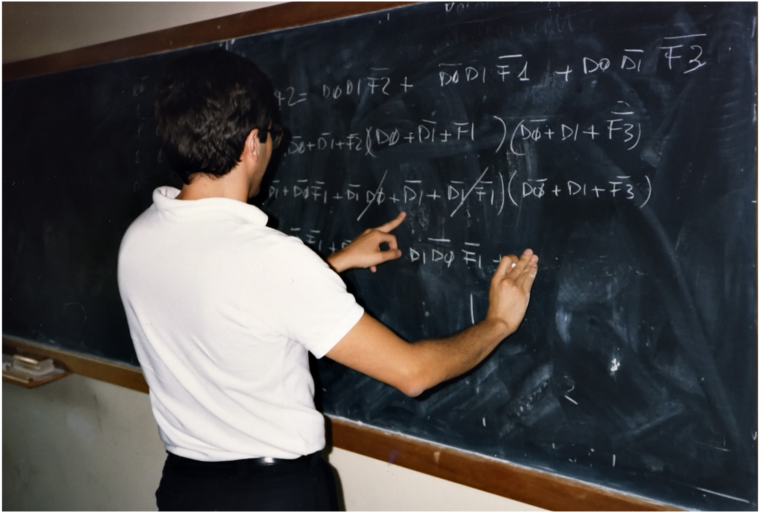

Building APE

Proud of the first APE

The four processor APE

The End