Scientific Applications

Jing Jiang, Jayanth Gummaraju,
Rohit Gupta

Outline

- Application Study
  - Vortex
- Architectural Issues
- Benchmarks

Applications

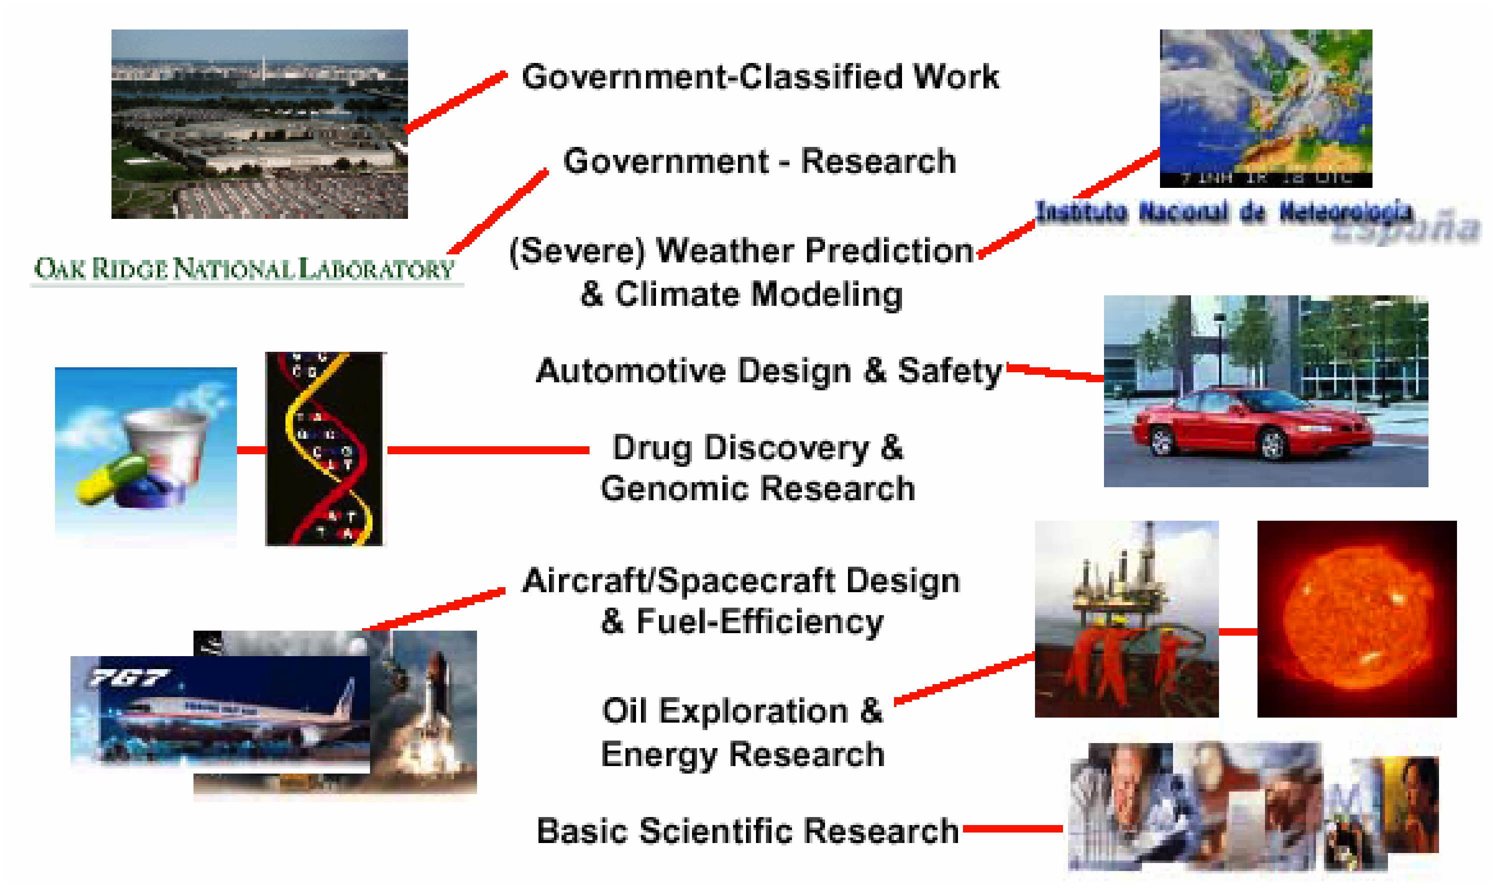

taken from Cray Research website: www.cray.com

Example Application: Vortex

- N-body Simulation
- O (N^2) Interactions
- Each Processor has N/P bodies
- Binary Tree Reduction
- In this example:
  - 4096 bodies
  - 100 stages

Code Size

APPLICATION	LANG	SCL	MTYPE	PROCS
CLIMATE	Fortran	80K	Delta	256
SEMI	Fortran	50K	Delta	512
MOLECULE	Fortran	1K	nCUBE 2	512
RENDER	C	2K	Delta	32
EXFLOW	C & Fortran	12K	Delta	512
QCD	C	2K	nCUBE 1	256
VORTEX	Fortran	1K	nCUBE 2	64
REACT	C & Fortran	42K	Delta	512

Memory Requirements

APPL	DATA	DATA/PROC	CODE/PROC	OS/PROC	% USED
CLIMATE	1750MB	7168KB	4096KB	4096KB	94
SEMI	1000MB	2048KB	NA	4096KB	NA
MOLECULE	1000MB	2048KB	200KB	200KB	60
RENDER	280MB	8960KB	260KB	4096KB	81
EXFLOW	732MB	1464KB	720KB	4096KB	38
QCD	17MB	70KB	98KB	100KB	52
VORTEX	0MB	3KB	492KB	200KB	17
REACT	536MB	1072KB	432KB	4096KB	34
AVE	665MB	2854KB	900KB	4096KB	54
MAX	1750MB	8960KB	4096KB	4096KB	94

Processing Requirements

APPLICATION	FLOAT OPS.	FLOAT OPS/PROC	PRECISION
CLIMATE	2970G	12200M	32
SEMI	10000G	20000M	64
MOLECULE	1000G	2000M	32
RENDER	24G	768M	32
EXFLOW	3994G	7987M	32
QCD	119G	474M	32
VORTEX	42G	677M	32
REACT	27648G	55296M	64

I/O Requirements

APPLICATION	INPUT	OUTPUT	VOLUME	VOL/MFLOP	DISK	TAPE
CLIMATE	1MB	1500MB	1500MB	517B	10MB	1500MB
SEMI	10MB	100MB	1000MB	100B	1000MB	0MB
MOLECULE	0MB	0MB	0MB	0B	0MB	0MB
RENDER	180MB	28MB	208MB	8858B	208MB	0MB
EXFLOW	0MB	1MB	1MB	0B	1MB	0MB
QCD	0MB	6MB	6MB	52B	6MB	0MB
VORTEX	0MB	0MB	0MB	3B	0MB	0MB
REACT	0MB	160MB	3400MB	126B	1600MB	0MB
AVE	24MB	224MB	764MB	1207B	353MB	187MB
MAX	180MB	1500MB	3400MB	8858B	1600MB	1500MB

Communication Requirements

APPLICATION	VOLUME	VOL/MFLOP	COUNT	COUNT/MFLOP	AVE SIZE
CLIMATE	965GB	325KB	1956M	660.0	505B
SEMI	120GB	12KB	15M	1.5	8192B
MOLECULE	1956GB	1956KB	44M	44.0	45568B
RENDER	2GB	98KB	0M	0.2	512000B
EXFLOW	562GB	144KB	256M	65.6	2248B
QCD	7GB	57KB	94M	810.0	72B
VORTEX	1GB	35KB	1M	29.4	1245B
REACT	132GB	5KB	12M	0.4	11264B
AVE	468GB	329KB	297M	201.4	72637B
MAX	1956GB	1956KB	1956M	810.0	512000B

General Characteristics

- Number Crunching Applications – typically have high arithmetic:memory operations.
- Large data sets – working set is also typically large but depends on application.
- Typically low temporal locality
- Depending on regularity of application – can have high spatial locality

Parallelism

- Lots of DLP, TLP, & ILP
- DLP
  - Same operation performed on all bodies
- TLP
  - Convert DLP to TLP
  - More flexibility compared to DLP
- ILP
  - Parallelism within "threads"
- Example: Vortex
  - Mostly DLP

Architectural Issues

Performance Trends

Scaling faster than Moore's law!

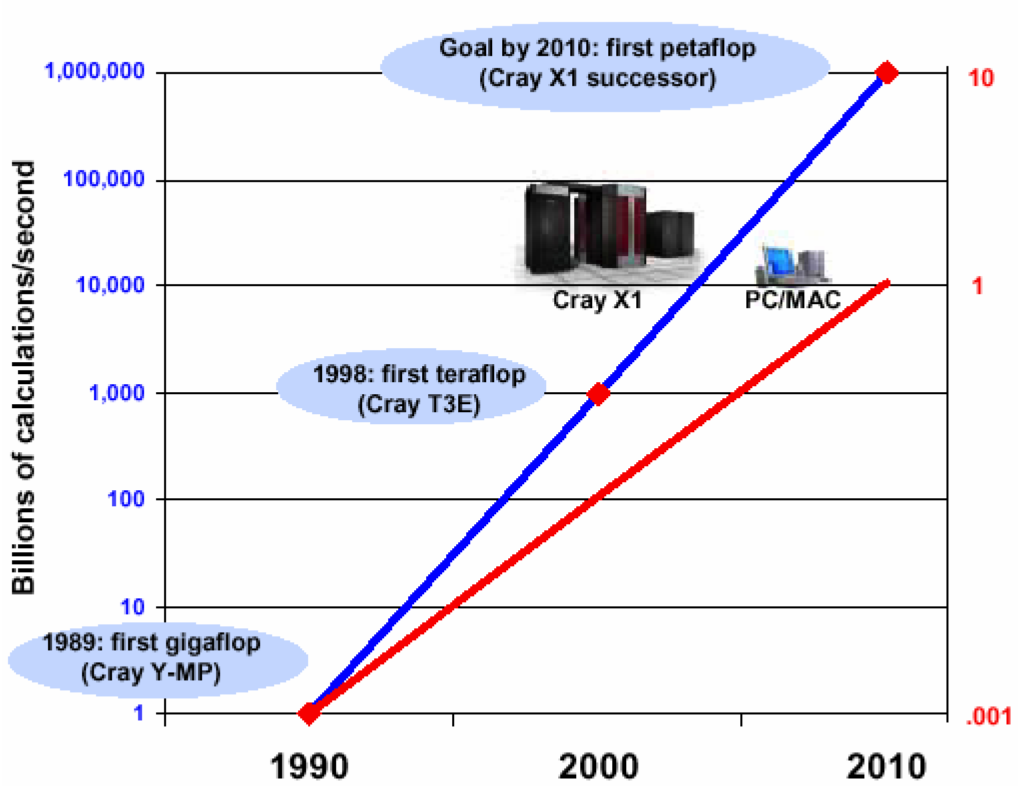

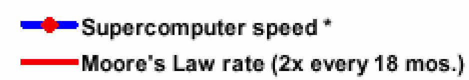

Year	Supercomputer speed: microprocessor speed
2010	100,000:1
2000	10,000:1
1990	1,000:1

taken from Cray Research website: www.cray.com

Processing Requirements

- Two approaches to achieve computational capacity:
  - Cluster Systems: typically 100s-1000s of processors
  - Stream/Vector Systems: fewer custom designed highly powerful processors

Interconnection Networks

- Both BW and latency important

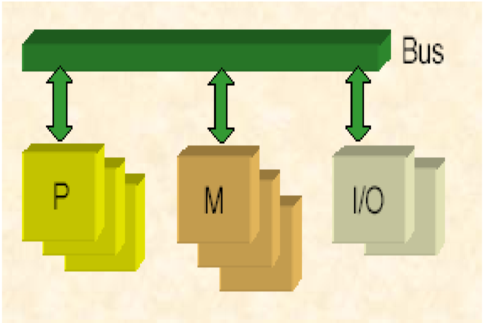

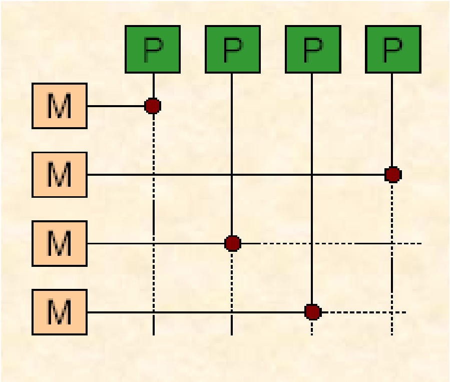

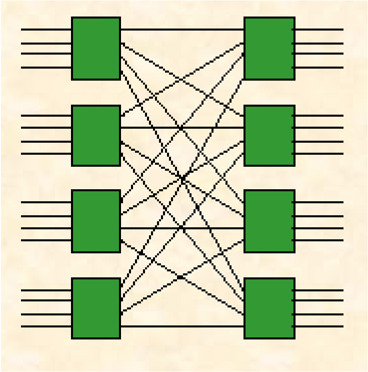

Bus:
•Need arbitration protocol
•Only one device at a time.

Crossbar Switch:
•All processors & memories connected
•O(N2) – doesn't scale well

Multistage Switch:
•N X N switches built from smaller switches
•E.g. 16X16 built from 2 stages of 4X4

Architectural Issues in Vortex

Data Memory	O (N)
FLOPS	O (N^2 + P)
I/O Volume	O (N)
Communication Volume	O (NP + P^2)
Communication Count	O (P^2)

Current Design Challenges

- System Performance-to-Cost ratio: millions of dollars to build
  - Custom vs. Cluster systems
- Programming model not very intuitive
- I/O Scalability
- Power

Benchmarks

DLAB suite –measuring performance of distributed & resource sharing systems on scientific applications

Performance Type	Measuring Benchmark
Floating Point	LFK, Linpack, Nbench-byte-2.1, EuroBen-V3.9, NPB2.3-serial
Integer Arithmetic	Nbench-byte-2.1, NPB2.3-serial (IS)
Memory Subsystem	Stream, Nbench-Byte-2.1, Stream_OpenMp
Communications	PMB-MPI1, Eff_bandwidth, NPB2.3-parallel
Sample Compact Applications	NPB2.3-Parallel (CFD), Parallel_Chem (Chemistry)
Full Applications	Angus (CFD), SBLI (CFD), DLPOLY-2.13 (Molecular Dynamics)

taken from CCLRC UK website: http://www.cse.clrc.ac.uk

Benchmarks

- Two important performance measures:
  - Peak Performance – dependent on maximum computation capacity. eg. Linpack
  - Sustained Performance – depends on overall system architecture (interconnects, memory BW)
- DLAB measures this and other characteristics

Memory interconnects

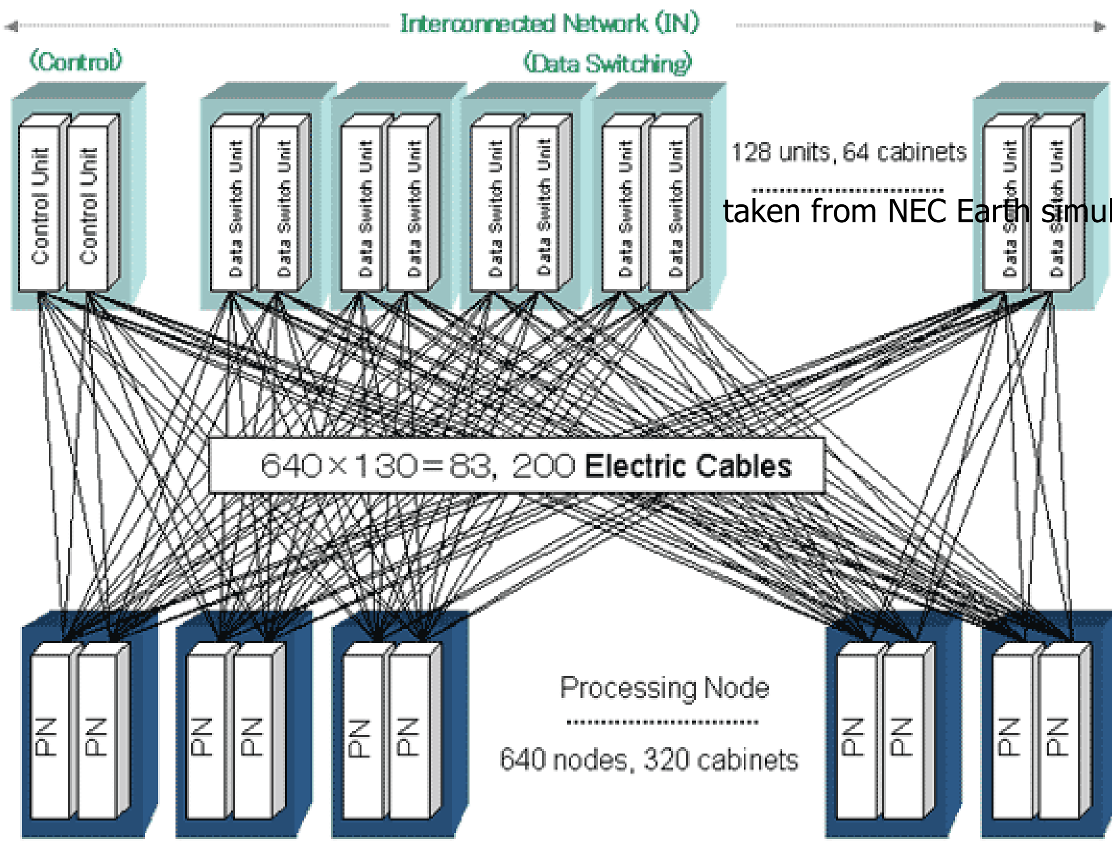

taken from NEC Earth simulator website: www.nec.co.jp