Who needs High Performance Computing?

Weather forecasting

Molecular modeling

Plasma physics

Serious data mining

Bioinformatics and computational genomics

Simulation of airplanes in wind tunnel

Search engines and query processing

Cryptanalysis

Benchmarking

Numerous yardsticks are used …

Real programs

Spice (CAD software)

TeX (Text processing software)

Compiling N lines of a C-program

Kernels

Small samples extracting key pieces of real programs.

Livermore loops (A set of 24 Fortran DO-loops extracted from operational codes used at the Lawrence Livermore National Laboratory). They are a mixture of **vectorizable** and non-vectorizable loops.

(What is vectorization?)

LINPACK is a collection of Fortran subroutines that analyze and solve linear equations and linear least-squares problems. The package solves linear systems **(A.x = B)** whose matrices have a diverse range of properties

Loop Vectorization

Original loop

```
For (j = 0; j < N, j++) {
    A[j] = A[j] + B[j]
}
```

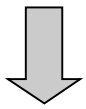

Vectorized loop

```
A[0:N] = A[0:N] + B[0:N]
```

{Runs more efficiently on a vector processor}

Non-vectorizable loop

```
For (j = 1; j < N, j++) {
    A[j] = A[j-1] + B[j]
}
```

Synthetic benchmarks

Dhrystone is a metric for determining the integer performance FORTRAN and C (and Java) codes

Whetstone measures the floating-point arithmetic performance of programs. Derived for the statistics of program behaviors in KDF9

Benchmark suites

Collection of benchmarks consisting of a variety of applications:

For example, **SPEC 92** consists of

Espresso	C-code minimizes Boolean function
Gcc	Gnu C-compiler
Spice2g6	Circuit simulation package
Ora	Ray tracing application
Nasa 7	7 kernels to do matrix op, FFT etc
Etc	
Etc	

Problems with benchmarking

Dual core processors

The latest buzz in the processor industry. AMD took the limelight with dual core AMD Opteron processors. It is two processor cores on one die. Opteron was designed with an extra HyperTransport link that simply means a faster connection between two cores. This does not mean that the chip itself is faster.

A dual core processor is between a single core processor and a dual processor system for architecture. A dual core processor will share some of the other hardware like the memory controller and bus, but everything else will be separate.

Special machines

Deep blue for chess playing

1st computer to win a chess game against Garry Kasparov

It is a massively parallel, **30-node**, RS/6000 SP-based computer system enhanced with 480 special purpose VLSI chess processors capable of evaluating 100,000,000 positions per second. Each node is an SMP with 2-4 processors

The IBM RS/6000 SP (1999)

Machine type	RISC-based distributed-memory cluster
Operating system	AIX (IBMs Unix variant)
Connection structure	Omega-switch
Compilers	Fortran 90, HPF, XL C, C++

Speed measured at 7.2 teraflops on a 8000-processor system

The architecture

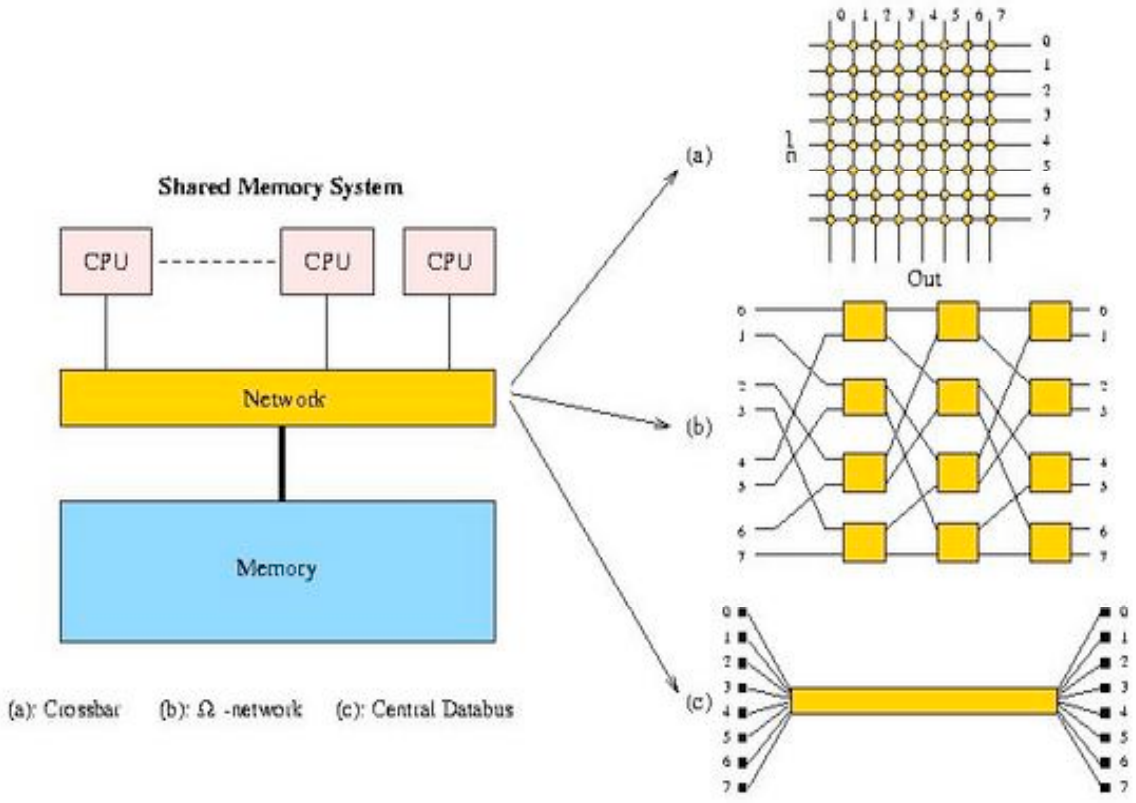

Blue Gene architecture

Blue Gene is an IBM Research project dedicated to the advancement of our understanding of important biological processes such as protein folding.

The full Blue Gene/L machine is being built with the Department of Energy's NNSA/Lawrence Livermore National Laboratory in California, and will have a peak speed of 360 Teraflops. Blue Gene/L occupies the **#1 position in the TOP500 supercomputer list** announced in November 2005 and IBM now offers a Blue Gene Solution. IBM and its collaborators are currently exploring a growing list of applications including hydrodynamics, quantum chemistry, molecular dynamics, climate modeling and financial modeling.

Deep crack 1998 broke DES in 56 hours