Productivity Portability Performance for Heterogeneous Systems

Michael Wolfe
The Portland Group, Inc.

August 2012

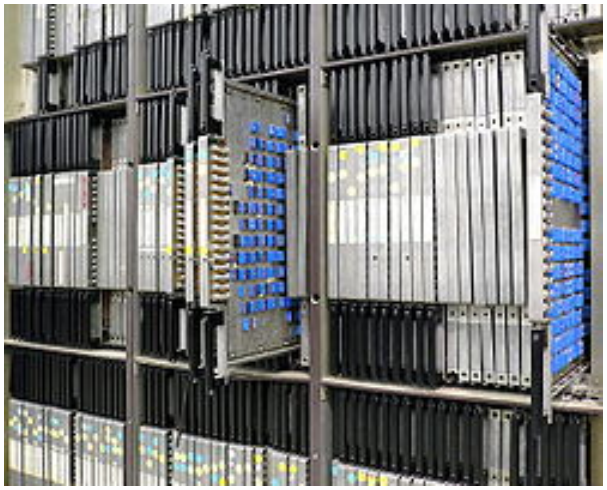

Illiac IV

Star-100

TI ASC

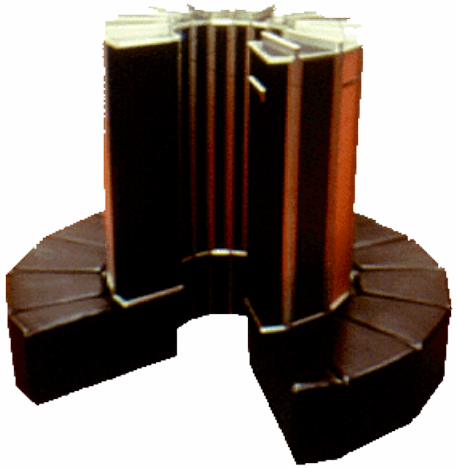

Cray 1

Illiac IV

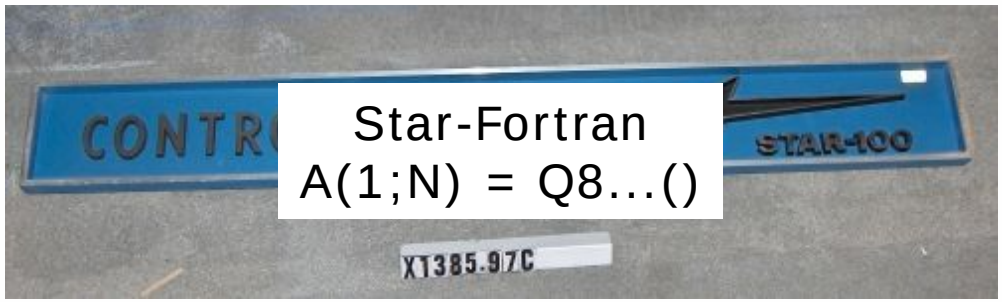

Star-100

TI ASC

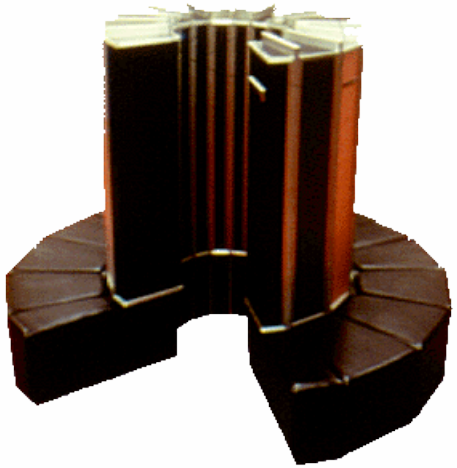

Cray 1

Sequent Balance

Encore Multimax

Cray 2

ETA 10

Sequent directives
Threading library

Sequent Balance

Encore Multimax

POSIX Threads
OpenMP
#pragma omp parallel

Cray Microtasking

Cray 2

ETA 10

Cosmic Cube

Intel iPSC/2

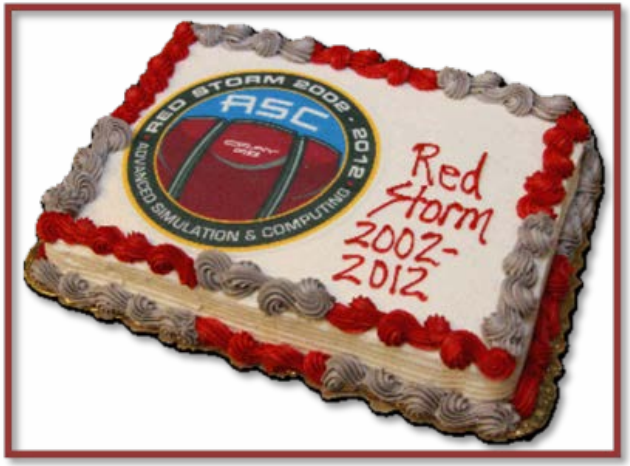

Cray Red Storm

IBM Blue Gene

Cosmic Cube

Intel iPSC/2

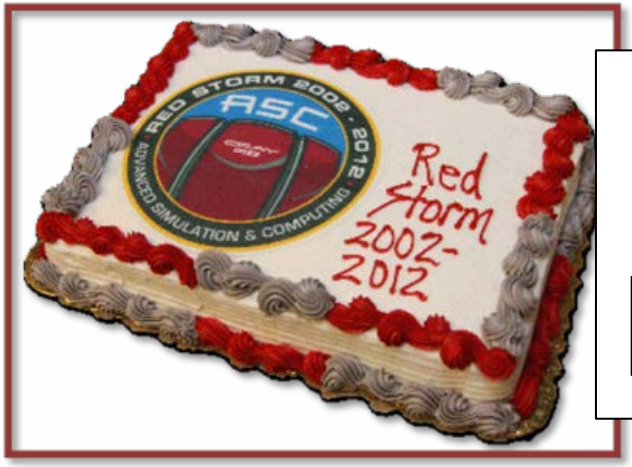

Cray Red Storm

MPI
HPF

IBM Blue Gene

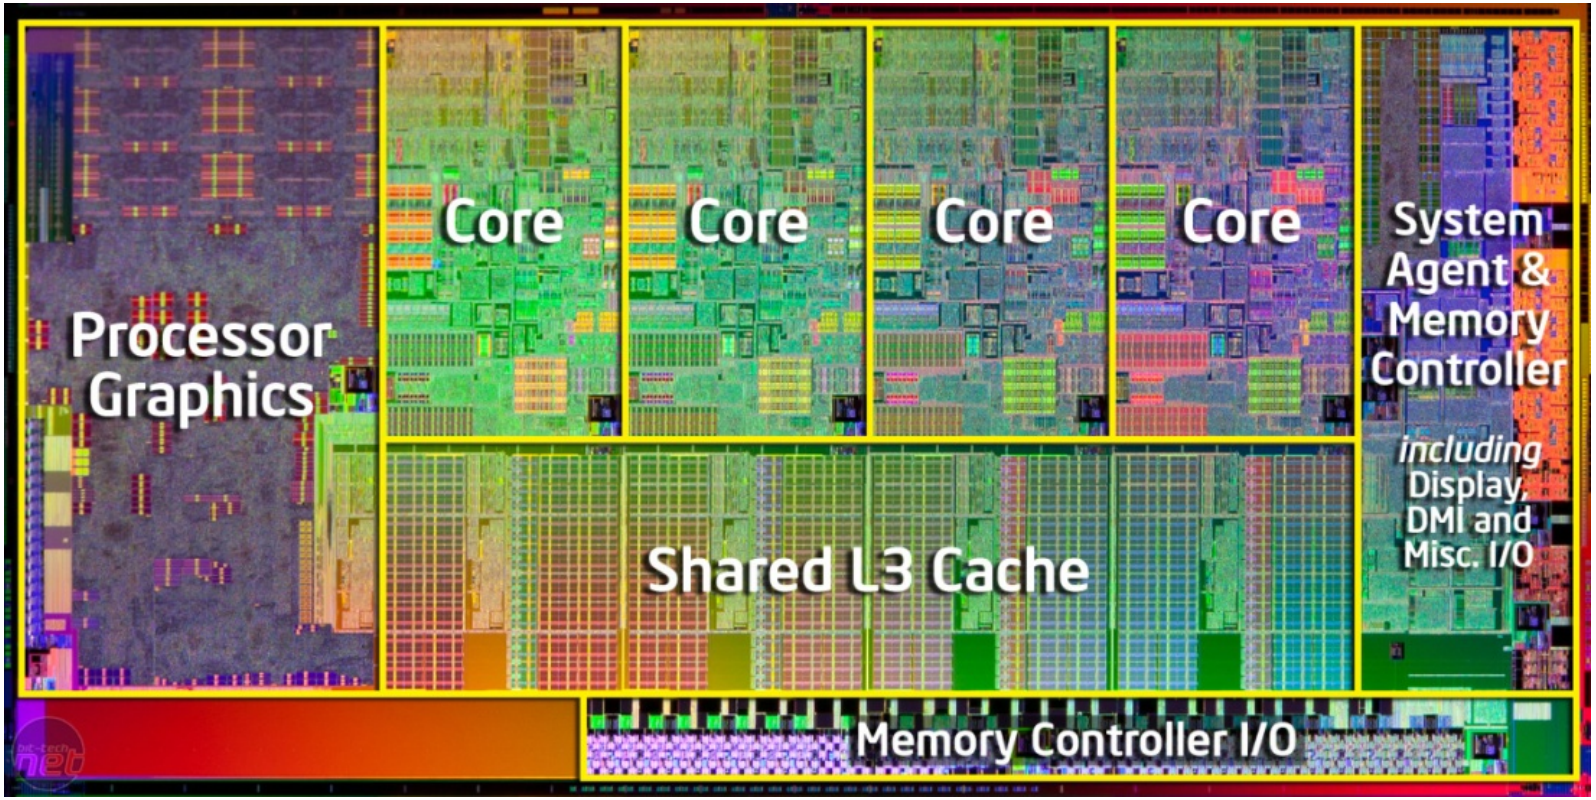

Intel Sandybridge

NVIDIA Tesla

AMD Firestream

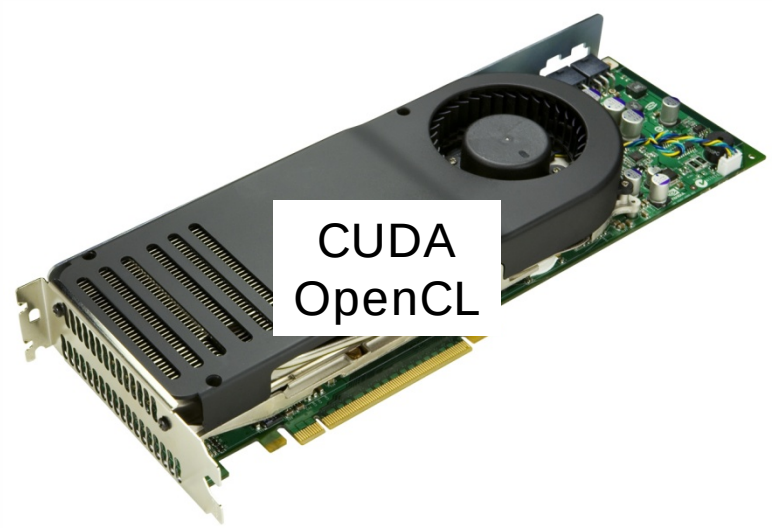

NVIDIA Tesla

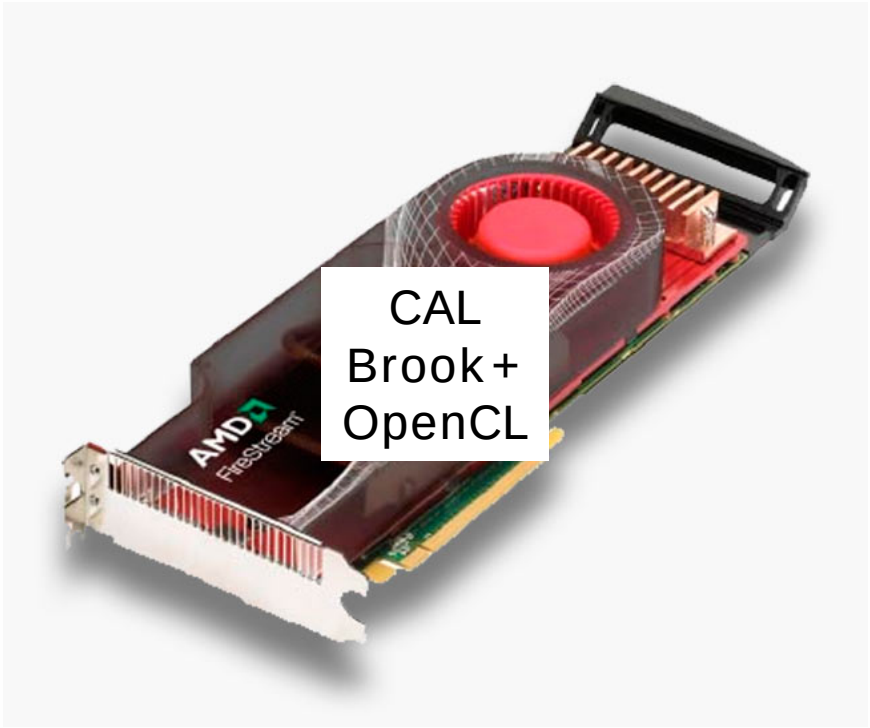

AMD Firestream

OpenACC Directives

```
#pragma acc kernels loop
  for( i = 1; i < n-1; ++i ){
    for( j = 1; j < n-1; ++j ){
      a[j+i*n] = 0.25*(b[j-1+i*n] + b[j+(i-1)*n] +
                       b[j+1+i*n] + b[j+(i+1)*n]);
    }
  }
```

OpenACC Directives

```
void iterate( float* a, float* b, int n ){
  int i, j;
#pragma acc kernels loop
  for( i = 1; i < n-1; ++i ){
    for( j = 1; j < n-1; ++j ){
      a[j+i*n] = 0.25*(b[j-1+i*n] + b[j+(i-1)*n] +
                       b[j+1+i*n] + b[j+(i+1)*n]);
    }
  }
}
...
#pragma acc data copy(a[0:n*n],b[0:n*n])
{
  for( iter = 0; ... ){
    iterate( a, b, n );
    tmp=a;a=b;b=tmp;
  }
}
```

Successful Models

- Vectorizing Compilers
  - local analysis suffices
  - programmer feedback to guide program changes
  - portable across a wide range of target machines
- OpenMP
  - modest changes to the program
  - automatic analysis could not compete
  - portable across a wide range of target machines
- MPI
  - required programmers to recode for locality
  - exposes cost of communication
  - portable across a wide range of target machines

Goals and Requirements

- **Virtualize** as much as possible
  - no hard constants in the program
- **Expose** the bottlenecks
  - communication is always slow
- **Reprogramming** is inevitable
  - minimize, guide, focus
- **Portability** across targets is the key
  - tomorrow's machine is not isomorphic to today's