


# Practical Parallel Programming Intensive

SciNet HPC Consortium

9–13 May 2011

# Welcome to the intensive parallel programming course!



# Part I

## The Course

# Main goal of the course

... to enable young researchers already experienced in scientific computing to leave with the knowledge necessary to begin writing the parallel codes needed for their research.

The course will be a mix of lectures and immediate feedback on practical assignments, designed to ensure that students leave with significant experience in both OpenMP and MPI, two of the standards for parallel computing today.

So there'll be a lot of typing and programming to help build skills with OpenMP and MPI.

We will use C and Fortran. We'll assume that you already know one of them, but not both.

# Schedule

Mon May 9	AM	Intro to Parallel Computing, SciNet resources
	PM	OpenMP I <i>+hands on</i>
Tue May 10	AM	OpenMP II <i>+hands on</i>
	PM	MPI I <i>+hands on</i>
Wed May 11	AM	MPI II <i>+hands on</i>
	PM	Explicit PDEs: Hydrodynamics <i>+hands on</i>
Thu May 12	AM	Particle Methods: N-body <i>+hands on</i>
	PM	GPU Programming <i>+hands on</i>
Fri May 13	AM	Performance tools & Best practices

# Strongly recommended books

(not provided by us)

- 1 B. Chapman, G. Jost and R. van der Pas, *Using OpenMP: Portable Shared Memory Parallel Programming* (MIT Press, Cambridge 2008)
- 2 W. Gropp, E. Lusk, A. Skjellum, *Using MPI: Portable Parallel Programming with the Message-Passing Interface*, second edition (MIT Press, Cambridge 1999).

## Part II

# Introduction to Parallel Programming

# Why Parallel Programming?



## 1 Faster

There's a limit to how fast 1 computer can compute.

**So use more computers!**



# Why Parallel Programming?



- 1 **Faster**  
There's a limit to how fast 1 computer can compute.
- 2 **Bigger**  
There's a limit to how much memory, disk, etc, can be put on 1 computer.

**So use more computers!**

# Why Parallel Programming?



- 1 Faster**  
There's a limit to how fast 1 computer can compute.
- 2 Bigger**  
There's a limit to how much memory, disk, etc, can be put on 1 computer.
- 3 More**  
Want to do the same thing that was done on 1 computer, but *thousands of times*.

**So use more computers!**

# Why is it necessary?

- Modern experiments and observations yield vastly more data to be processed than in the past.
- As more computing resources become available (SciNet), the bar for cutting edge simulations is raised.
- Science that before could not even be done becomes reachable.

# Why is it necessary?

- Modern experiments and observations yield vastly more data to be processed than in the past.
- As more computing resources become available (SciNet), the bar for cutting edge simulations is raised.
- Science that before could not even be done becomes reachable.

However:


# Why is it necessary?

- Modern experiments and observations yield vastly more data to be processed than in the past.
- As more computing resources become available (SciNet), the bar for cutting edge simulations is raised.
- Science that before could not even be done becomes reachable.

However:

- Advances in clock speeds, bigger and faster memory and disks have been lagging as compared to e.g. 10 years ago.  
*Can no longer “just wait a year” and get a better computer.*
- So more computing resources here means: more cores running **concurrently**.
- *Even most laptops now have 2 or more cpus.*
- So parallel computing is necessary.


# Wait, what about Moore's Law?



(source: Transistor Count and Moore's Law - 2008.svg, by Wgsimon, wikipedia)

# Wait, what about Moore's Law?


## CPU Transistor Counts 1971-2008 & Moore's Law



## Moore's law

*... describes a long-term trend in the history of computing hardware. The number of transistors that can be placed inexpensively on an integrated circuit doubles approximately every two years.*


(source: Moore's law, wikipedia)



(source: Transistor Count and Moore's Law - 2008.svg, by Wgsimon, wikipedia)

# Wait, what about Moore's Law?


CPU Transistor Counts 1971-2008 & Moore's Law



## Moore's law

*... describes a long-term trend in the history of computing hardware. The number of transistors that can be placed inexpensively on an integrated circuit doubles approximately every two years.*

(source: Moore's law, wikipedia)



## But...

- Moore's Law didn't promise us clock speed.
- More transistors but getting hard to push clock speed up. Power density is limiting factor.
- So more cores at fixed clock speed.

(source: Transistor Count and Moore's Law - 2008.svg, by Wgsimon, wikipedia)



# Concurrency


- Must have something to do for all these cores.
- Find parts of the program that can be done independently, and therefore concurrently.
- There must be many such parts.
- The order of execution should not matter either.
- **Data dependencies limit concurrency.**



(source: <http://flickr.com/photos/splorp>)

# Parameter study: best case scenario

- Aim is to get results from a model as a parameter varies.
- Can run the serial program on each processor at the same time.
- Get “more” done.




# Throughput


- How many tasks can you do per time unit?

$$\text{throughput} = \mathbf{H} = \frac{\mathbf{N}}{\mathbf{T}}$$

- Maximizing  $\mathbf{H}$  means that you can do as much as possible.
- Independent tasks: using  $\mathbf{P}$  processors increases  $\mathbf{H}$  by a factor  $\mathbf{P}$ .



vs.




# Scaling — Throughput

- How a problem's throughput scales as processor number increases (“strong scaling”).
- In this case, linear scaling:

$$H \propto P$$

- This is **Perfect scaling**.




# Scaling – Time

- How a problem's timing scales as processor number increases.
- Measured by the time to do one unit. In this case, inverse linear scaling:

$$T \propto 1/P$$

- Again this is the ideal case, or “embarrassingly parallel”.




# Scaling – Time

- How a problem's timing scales as processor number increases.
- Measured by the time to do one unit. In this case, inverse linear scaling:

$$T \propto 1/P$$

- Again this is the ideal case, or “embarrassingly parallel”.




# Scaling – Speedup

- How much faster the problem is solved as processor number increases.
- Measured by the serial time divided by the parallel time

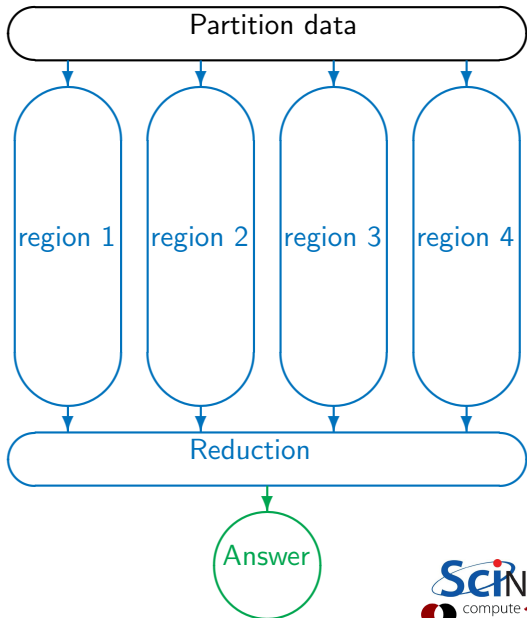
$$S = \frac{T_{\text{serial}}}{T(P)} \propto P$$

- For embarrassingly parallel applications: Linear speed up.




# Non-ideal cases

- Say we want to integrate some tabulated experimental data.
- Integration can be split up, so different regions are summed by each processor.
- Non-ideal:
  - ▶ First need to get data to processor
  - ▶ And at the end bring together all the sums: **“reduction”**






# Non-ideal cases




# Non-ideal cases

**Parallel region**  $\Rightarrow$   
Perfectly Parallel  
(for large **N**)



# Non-ideal cases




# Non-ideal cases

Parallel overhead  $\Rightarrow$

Parallel region  $\Rightarrow$

Perfectly Parallel  
(for large  $N$ )

Serial portion  $\Rightarrow$




# Non-ideal cases

Parallel overhead  $\Rightarrow$

Parallel region  $\Rightarrow$

Perfectly Parallel  
(for large  $N$ )

Serial portion  $\Rightarrow$



Suppose non-parallel part const:  $T_s$


# Amdahl's law

Speed-up (without parallel overhead):

$$S = \frac{NT_1 + T_s}{\frac{NT_1}{P} + T_s}$$

or, calling  $f = T_s / (T_s + NT_1)$  the serial fraction,

$$S = \frac{1}{f + (1 - f)/P}$$



(for  $f = 5\%$ )


# Amdahl's law

Speed-up (without parallel overhead):

$$S = \frac{NT_1 + T_s}{\frac{NT_1}{P} + T_s}$$

or, calling  $f = T_s / (T_s + NT_1)$  the serial fraction,

$$S = \frac{1}{f + (1 - f)/P} \quad \begin{matrix} P \rightarrow \infty \\ \longrightarrow \end{matrix} \quad \frac{1}{f}$$



(for  $f = 5\%$ )


# Amdahl's law

Speed-up (without parallel overhead):

$$S = \frac{NT_1 + T_s}{\frac{NT_1}{P} + T_s}$$

or, calling  $f = T_s / (T_s + NT_1)$  the serial fraction,

$$S = \frac{1}{f + (1 - f)/P} \quad \begin{matrix} P \rightarrow \infty \\ \longrightarrow \end{matrix} \quad \frac{1}{f}$$



Serial part dominates asymptotically.

Speed-up limited, no matter size of **P**.

(for  $f = 5\%$ )




# Amdahl's law

Speed-up (without parallel overhead):

$$S = \frac{NT_1 + T_s}{\frac{NT_1}{P} + T_s}$$

or, calling  $f = T_s / (T_s + NT_1)$  the serial fraction,

$$S = \frac{1}{f + (1 - f)/P} \quad \begin{matrix} P \rightarrow \infty \\ \longrightarrow \end{matrix} \quad \frac{1}{f}$$



Serial part dominates asymptotically.

Speed-up limited, no matter size of **P**.

And this is the overly optimistic case!

(for  $f = 5\%$ )

# Scaling efficiency

Speed-up compared to ideal factor **P**:

$$\text{Efficiency} = \frac{S}{P}$$

This will invariably fall off for larger **P** except for embarrassing parallel problems.

$$\text{Efficiency} \sim \frac{1}{fP} \xrightarrow{P \rightarrow \infty} 0$$



You cannot get 100% efficiency in any non-trivial problem.

All you can aim for here is to make the efficiency as least low as possible. Sometimes, that can mean running on less processors, but more problems at the same time.

# Timing example

- Say 100s in integration cost
- 5s in reduction
- Neglect communication cost
- What happens as we vary number of processors **P**?

$$\text{Time} = (100\text{s})/P + 5$$





# Throughput example

$$H(P) = \frac{N}{\text{Time}(P)}$$

- Say we are doing  $k$  at the same time, on  $P$  processors total.

$$H_k(P) = \frac{kN}{\text{Time}(P/k)}$$

Say  $N = 100$ :



## Big Lesson #1

Always keep throughput in mind: if you have several runs, running more of them at the same time on less processors per run is often advantageous.

# Less ideal case of Amdahl's law

We assumed reduction is constant.  
But it will in fact increase with  $P$ ,  
from sum of results of all processors

$$T_s \approx PT_1$$

Serial fraction now a function of  $P$ :

$$f(P) = \frac{P}{N}$$

Amdahl:

$$S(P) = \frac{1}{f(P) + [1 - f(P)]/P}$$

# Less ideal case of Amdahl's law

We assumed reduction is constant.  
But it will in fact increase with  $P$ ,  
from sum of results of all processors

$$T_s \approx PT_1$$

Serial fraction now a function of  $P$ :

$$f(P) = \frac{P}{N}$$

Amdahl:

$$S(P) = \frac{1}{f(P) + [1 - f(P)]/P}$$

Example:  $N = 100$ ,  $T_1 = 1s \dots$

# Less ideal case of Amdahl's law

We assumed reduction is constant.  
But it will in fact increase with  $P$ ,  
from sum of results of all processors

$$T_s \approx PT_1$$


Serial fraction now a function of  $P$ :

$$f(P) = \frac{P}{N}$$

Amdahl:

$$S(P) = \frac{1}{f(P) + [1 - f(P)]/P}$$

Example:  $N = 100$ ,  $T_1 = 1s \dots$






# Trying to beat Amdahl's law #1

## Scale up!

The larger **N**, the smaller  
the serial fraction:

$$f(P) = \frac{P}{N}$$




# Trying to beat Amdahl's law #1

## Scale up!

The larger  $N$ , the smaller the serial fraction:

$$f(P) = \frac{P}{N}$$



Weak scaling: Increase problem size while increasing  $P$

$$\text{Time}_{\text{weak}}(P) = \text{Time}(N = n \times P, P)$$


Good weak scaling means this time approaches a constant for large  $P$ .

# Trying to beat Amdahl's law #1

## Scale up!

The larger **N**, the smaller the serial fraction:

$$f(P) = \frac{P}{N}$$



Weak scaling: Increase problem size while increasing **P**

$$\text{Time}_{\text{weak}}(P) = \text{Time}(N = n \times P, P)$$


Good weak scaling means this time approaches a constant for large **P**.

## Gustafson's Law

Any large enough problem can be efficiently parallelized (Efficiency  $\rightarrow$  1).

# Trying to beat Amdahl's law #2

Parallel overhead  $\Rightarrow$



Parallel region  $\Rightarrow$


Perfectly Parallel  
(for large **N**)

Serial portion  $\Rightarrow$



# Trying to beat Amdahl's law #2

Parallel overhead  $\Rightarrow$



Parallel region  $\Rightarrow$

Perfectly Parallel  
(for large  $N$ )




Serial portion  $\Rightarrow$

Rewrite




# Trying to beat Amdahl's law #2

Parallel overhead  $\Rightarrow$




Parallel region  $\Rightarrow$

Perfectly Parallel  
(for large  $N$ )




Serial portion  $\Rightarrow$

Rewrite




# Trying to beat Amdahl's law #2

Parallel overhead  $\Rightarrow$



Parallel region  $\Rightarrow$


Perfectly Parallel  
(for large N)



Serial portion  $\Rightarrow$

Rewrite

$\propto^2 \log P$



# Trying to beat Amdahl's law #2

'Serial' fraction now different function of  $P$ :

$$f(P) = \frac{2 \log P}{N}$$

Amdahl:

$$S(P) = \frac{1}{f(P) + [1 - f(P)]/P}$$



# Trying to beat Amdahl's law #2

'Serial' fraction now different function of  $P$ :

$$f(P) = \frac{2 \log P}{N}$$

Amdahl:

$$S(P) = \frac{1}{f(P) + [1 - f(P)]/P}$$

Example:  $N = 100$ ,  $T_1 = 1s \dots$

# Trying to beat Amdahl's law #2


'Serial' fraction now different function of  $P$ :

$$f(P) = \frac{2 \log P}{N}$$

Amdahl:

$$S(P) = \frac{1}{f(P) + [1 - f(P)]/P}$$

Example:  $N = 100$ ,  $T_1 = 1s \dots$



# Trying to beat Amdahl's law #2

## Weak scaling

$$\text{Time}_{\text{weak}}(\mathbf{P}) = \text{Time}(\mathbf{N} = n \times \mathbf{P}, \mathbf{P})$$

Should approach constant for large  $\mathbf{P}$ .

Let's see...

# Trying to beat Amdahl's law #2


## Weak scaling

$$\text{Time}_{\text{weak}}(P) = \text{Time}(N = n \times P, P)$$

Should approach constant for large  $P$ .

Let's see...

Not quite!



# Trying to beat Amdahl's law #2

## Weak scaling


$$\text{Time}_{\text{weak}}(P) = \text{Time}(N = n \times P, P)$$

Should approach constant for large  $P$ .

Let's see...

Not quite!

But much better than before.



# Trying to beat Amdahl's law #2

## Weak scaling


$$\text{Time}_{\text{weak}}(P) = \text{Time}(N = n \times P, P)$$

Should approach constant for large  $P$ .

Let's see...

Not quite!

But much better than before.



# Trying to beat Amdahl's law #2

## Weak scaling

$$\text{Time}_{\text{weak}}(\mathbf{P}) = \text{Time}(\mathbf{N} = n \times \mathbf{P}, \mathbf{P})$$

Should approach constant for large  $\mathbf{P}$ .

Let's see...


Not quite!

But much better than before.

Gustafson?

It turns out that Gustafson's law assumes that the serial cost does not change with  $\mathbf{P}$ .

Here that grows logarithmically with  $\mathbf{P}$ , and this is reflected in the weak scaling.



# Trying to beat Amdahl's law #2

## Weak scaling

$$\text{Time}_{\text{weak}}(\mathbf{P}) = \text{Time}(\mathbf{N} = n \times \mathbf{P}, \mathbf{P})$$

Should approach constant for large  $\mathbf{P}$ .

Let's see...


Not quite!

But much better than before.

Gustafson?

It turns out that Gustafson's law assumes that the serial cost does not change with  $\mathbf{P}$ .

Here that grows logarithmically with  $\mathbf{P}$ , and this is reflected in the weak scaling.



Really not that bad.

*and other algorithms can do better.*




## Big Lesson #2

Optimal Serial Algorithm for your problem may not be the  $P \rightarrow 1$  limit of your optimal parallel algorithm.


# Synchronization

- Most problems are not purely concurrent.
- Some level of synchronization or exchange of information is needed between tasks.
- While synchronizing, nothing else happens: increases Amdahl's  $f$ .
- And synchronizations are themselves costly.



# Load balancing

- The division of calculations among the processors may not be equal.
- Some processors would already be done, while others are still going.
- Effectively using less than  $P$  processors: This reduces the efficiency.
- Aim for load balanced algorithms.





# Locality

- So far we neglected communication costs.
- But communication costs are more expensive than computation!
- To minimize communication to computation ratio:
  - \* Keep the data where it is needed.
  - \* Make sure as little data as possible is to be communicated.
  - \* Make shared data as local to the right processors as possible.
- Local data means less need for syncs, or smaller-scale syncs.
- Local syncs can alleviate load balancing issues.

# Locality

- So far we neglected communication costs.
- But communication costs are more expensive than computation!
- To minimize communication to computation ratio:
  - \* Keep the data where it is needed.
  - \* Make sure as little data as possible is to be communicated.
  - \* Make shared data as local to the right processors as possible.
- Local data means less need for syncs, or smaller-scale syncs.
- Local syncs can alleviate load balancing issues.

## Example (PDE Domain decomposition)



## Big Lesson #3

Parallel algorithm design is about finding as much concurrency as possible, and arranging it in a way that maximizes locality.

# Parallel Computers



## Top500.org:

List of the worlds  
500 largest  
supercomputers.  
Updated every 6  
months,

Info on  
architecture, etc.

[Home](#) [Lists](#) [November 2010](#)

### TOP500 List - November 2010 (1-100)

$R_{max}$  and  $R_{peak}$  values are in TFlops. For more details about other fields, check the [TOP500 description](#).

Power data in KW for entire system

[next](#)

Rank	Site	Computer/Year Vendor	Cores	Rmax	Rpeak	Power
1	National Supercomputing Center in Tianjin China	Tianhe-1A - NUDT TH MPP, X5670 2.93Ghz 6C, NVIDIA GPU, FT-1000 8C / 2010 NUDT	186368	2566.00	4701.00	4040.00
2	DOE/SC/Oak Ridge National Laboratory United States	Jaguar - Cray XT5-HE Opteron 6-core 2.6 GHz / 2009 Cray Inc.	224162	1759.00	2331.00	6950.60
3	National Supercomputing Centre in Shenzhen (NSCS) China	Nebulae - Dawning TC3600 Blade, Intel X5650, NVidia Tesla C2050 GPU / 2010 Dawning	120640	1271.00	2984.30	2580.00
4	GSIC Center, Tokyo Institute of Technology Japan	TSUBAME 2.0 - HP ProLiant SL390s G7 Xeon 6C X5670, Nvidia GPU, Linux/Windows / 2010 NEC/HP	73278	1192.00	2287.63	1398.61
5	DOE/SC/LBNL/NERSC United States	Hopper - Cray XE6 12-core 2.1 GHz / 2010 Cray Inc.	153408	1054.00	1288.63	2910.00
6	Commissariat a l'Energie Atomique (CEA) France	Tera-100 - Bull bullx super-node S6010/S6030 / 2010 Bull SA	138368	1050.00	1254.55	4590.00

# Supercomputer architectures

- Clusters, or, **distributed memory machines**  
In essence a bunch of desktops linked together by a network (“interconnect”). Easy and cheap.
- Multi-core machines, or, **shared memory machines**  
These can see the same memory. Limited number of cores, typically, and much more \$\$\$.
- Vector machines.  
These were the early supercomputers, and could do the same operation on a large number of numbers at the same time.  
Very \$\$\$\$\$\$, especially at scale.  
These days, most chips have some low-level, small size vectorization, but you rarely need to worry about it (compiler should do this).


Most supercomputers are a hybrid combo of these different architectures.



# Distributed Memory: Clusters

Simplest type of parallel computer to build

- Take existing powerful standalone computers
- And network them




(source: <http://flickr.com/photos/eurleif/>)

# Distributed Memory: Clusters

Each node is independent!

Parallel code consists of programs running on separate computers, communicating with each other.

Could be entirely different programs.



# Distributed Memory: Clusters

Each node is independent!


Parallel code consists of programs running on separate computers, communicating with each other.

Could be entirely different programs.

Each node has own memory!

Whenever it needs data from another region, requests it from that CPU.

Usual model: “message passing”



# Clusters+Message Passing

## Hardware:


Easy to build

(Harder to build well)

Can build larger and larger clusters relatively easily

## Software:

Every communication has to be hand-coded:  
hard to program



# Cluster Communication Cost

	Latency	Bandwidth
GigE	10 $\mu$ s (10,000 ns)	1 Gb/s ( 60 ns/double)
Infiniband	2 $\mu$ s (2,000 ns)	2-10 Gb/s ( 10 ns /double)

Processor speed:  $O(\text{GFLOP}) \sim \text{few ns or less.}$


# Shared Memory

One large bank of memory, different computing cores acting on it. All 'see' same data.

Any coordination done through memory

Could use message passing, but no need.

Each code is assigned a **thread of execution** of a single program that acts on the data.



# Threads versus Processes

## Threads:

Threads of execution within one process, with access to the same memory etc.

## Processes:

Independent tasks with their own memory and resources

```
ljdursi@gpc-f102n08
File Edit View Terminal Tabs Help
top - 17:27:34 up 2 days, 1:40, 1 user, load average: 1.81, 0.56, 0.20
Tasks: 142 total, 3 running, 139 sleeping, 0 stopped, 0 zombie
Cpu(s): 95.9%us, 3.0%sy, 0.0%ni, 0.0%id, 0.0%wa, 0.1%hi, 1.0%si, 0.0%st
Mem: 16411872k total, 2778368k used, 13633504k free, 256k buffers
Swap: 0k total, 0k used, 0k free, 2265652k cached

  PID USER      PR  NI  VIRT  RES  SHR  S %CPU  %MEM    TIME+  COMMAND
 18121 ljdursi   25   0 89536 1076 840  R 779.0  0.0   0:29.01 diffusion-omp
17193 root      15   0 35300 2580 60  S 15.0  0.0   0:01.57 pbs_mon
17192 root      15   0 35300 3216 696  R  6.0  0.0   0:00.48 pbs_mon
   1 root      15   0 10344  740 612  S  0.0  0.0   0:01.45 init
   2 root      RT  -5   0   0   0   0  S  0.0  0.0   0:00.00 migration/0
   3 root      34  19   0   0   0  S  0.0  0.0   0:00.00 ksoftirqd/0
   4 root      RT  -5   0   0   0   0  S  0.0  0.0   0:00.00 watchdog/0
   5 root      RT  -5   0   0   0   0  S  0.0  0.0   0:00.01 migration/1
   6 root      34  19   0   0   0  S  0.0  0.0   0:00.01 ksoftirqd/1
   7 root      RT  -5   0   0   0   0  S  0.0  0.0   0:00.00 watchdog/1
   8 root      RT  -5   0   0   0   0  S  0.0  0.0   0:00.00 migration/2
   9 root      34  19   0   0   0  S  0.0  0.0   0:00.00 ksoftirqd/2
  10 root      RT  -5   0   0   0   0  S  0.0  0.0   0:00.00 watchdog/2
  11 root      RT  -5   0   0   0   0  S  0.0  0.0   0:00.00 migration/3
```


```
ljdursi@gpc-f102n08
File Edit View Terminal Tabs Help
top - 17:33:58 up 2 days, 1:47, 1 user, load average: 0.80, 0.31, 0.17
Tasks: 150 total, 9 running, 141 sleeping, 0 stopped, 0 zombie
Cpu(s):100.0%us, 0.0%sy, 0.0%ni, 0.0%id, 0.0%wa, 0.0%hi, 0.0%si, 0.0%st
Mem: 16411872k total, 2801172k used, 13610700k free, 256k buffers
Swap: 0k total, 0k used, 0k free, 2268568k cached

  PID USER      PR  NI  VIRT  RES  SHR  S %CPU  %MEM    TIME+  COMMAND
18393 ljdursi   25   0 187m 5504 3484  R 100.2  0.0   0:05.45 diffusion-mpi
18395 ljdursi   25   0 187m 5512 3492  R 100.2  0.0   0:05.46 diffusion-mpi
18397 ljdursi   25   0 187m 5508 3488  R 100.2  0.0   0:05.46 diffusion-mpi
18392 ljdursi   25   0 187m 5580 3556  R 99.9  0.0   0:05.40 diffusion-mpi
18394 ljdursi   25   0 187m 5504 3488  R 99.9  0.0   0:05.45 diffusion-mpi
18396 ljdursi   25   0 187m 5512 3492  R 99.9  0.0   0:05.45 diffusion-mpi
18398 ljdursi   25   0 187m 5500 3480  R 99.9  0.0   0:05.43 diffusion-mpi
18399 ljdursi   25   0 187m 5512 3492  R 99.9  0.0   0:05.46 diffusion-mpi
   1 root      15   0 10344  740 612  S  0.0  0.0   0:01.45 init
   2 root      RT  -5   0   0   0   0  S  0.0  0.0   0:00.00 migration/0
   3 root      34  19   0   0   0  S  0.0  0.0   0:00.00 ksoftirqd/0
   4 root      RT  -5   0   0   0   0  S  0.0  0.0   0:00.00 watchdog/0
   5 root      RT  -5   0   0   0   0  S  0.0  0.0   0:00.01 migration/1
   6 root      34  19   0   0   0  S  0.0  0.0   0:00.01 ksoftirqd/1
```

# Shared Memory: NUMA

## Non-Uniform Memory Access


- Each core typically has some memory of its own.
- Cores have cache too.
- Keeping this memory coherent is extremely challenging.





# Coherency

- The different levels of memory imply multiple copies of some regions
- Multiple cores mean can update unpredictably
- Very expensive hardware
- Hard to scale up to lots of processors, very \$\$\$
- Very simple to program!!




# Shared Memory Communication Cost

	Latency	Bandwidth
GigE	10 $\mu$ s (10,000 ns)	1 Gb/s ( 60 ns/double)
Infiniband	2 $\mu$ s (2,000 ns)	2-10 Gb/s ( 10 ns /double)
NUMA (shared memory)	0.1 $\mu$ s (100 ns)	10-20 Gb/s ( 4 ns /double)

Processor speed:  $O(\text{GFLOP}) \sim \text{few ns or less.}$

# Hybrid Architectures

- Multicore machines linked together with an interconnect
- Many cores have modest vector capabilities.
- Machines with GPU: GPU is multi-core, but the amount of shared memory is limited.



We will focus on the aspects that affect the programmer:

- Shared memory: OpenMP
- Distributed memory: MPI
- Graphics computing: CUDA, OpenCL

## Big Lesson #4

The best approach to parallelizing your problem will depend on both details of your problem and of the hardware available.

# Resources at SciNet

## 1. General Purpose Cluster (GPC)



# Resources at SciNet

## 1. General Purpose Cluster (GPC)

- 3780 nodes with two 2.53GHz quad-core Intel Xeon 5500 (*Nehalem*) x86-64 processors (30240 cores total)
- 16GB RAM per node
- Gigabit ethernet network on all nodes for management, disk I/O, boot, etc.
- InfiniBand network on 1/4 of the nodes only used for job communication
- 306 TFlops
- #16 on the June 2009 *TOP500* supercomputer sites
- #1 in Canada

# Resources at SciNet

## 2. Tightly Coupled System (TCS)



# Resources at SciNet

## 2. Tightly Coupled System (TCS)

- 104 nodes of 16 dual-core 4.7GHz Power 6 processors.
- 128GB RAM per node
- Interconnected by full non-blocking InfiniBand
- 62 TFlops
- #80 on the June 2009 *TOP500* supercomputer sites
- #3 in Canada



# Resources at SciNet


## 2. Tightly Coupled System (TCS)

- 104 nodes of 16 dual-core 4.7GHz Power 6 processors.
- 128GB RAM per node
- Interconnected by full non-blocking InfiniBand
- 62 TFlops
- #80 on the June 2009 *TOP500* supercomputer sites
- #3 in Canada

Access disabled by default. For access, email us explaining the nature of your work. Your application should scale well to over 32 procs.

# Resources at SciNet

## 3. Accelerator Research Cluster (ARC)



## 3. Accelerator Research Cluster (ARC)

8 GPU devel nodes and 4 NVIDIA Tesla M2070. Per node:

- 2 × quad-core Intel Xeon X5550 2.67GHz
- 48 GB RAM
- Interconnected by DDR InfiniBand
- 2 × GPUs with CUDA capability 2.0 (Fermi) each with 448 CUDA cores @ 1.15GHz and 6 GB of RAM.

Max. computing power CPUs: 683.52 GFlops

Max. computing power GPUs: 4.12 TFlops (single prec)

2.06 TFlops (double prec)

*Access disabled by default.*

# Part III

## Review of C

# C review: Basics

- C was designed for (unix) system programming.
- C has a very small base.
- Most functionality is in (standard) libraries.
- We will use C99.

# C review: Basics

- C was designed for (unix) system programming.
- C has a very small base.
- Most functionality is in (standard) libraries.
- We will use C99.

## Example (Basic C program)

```
#include <stdio.h>
// include stdio.h to print
int main() // always called first
{ // braces delimit code block
    printf("Hello world.\n");
    // function call to print
    // line ends with a semicolon
    return 0;
    // optional return value to shell
}
```

# C review: Basics

- C was designed for (unix) system programming.
- C has a very small base.
- Most functionality is in (standard) libraries.
- We will use C99.

## Example (Basic C program)

```
#include <stdio.h>
// include stdio.h to print
int main() // always called first
{ // braces delimit code block
    printf("Hello world.\n");
    // function call to print
    // line ends with a semicolon
    return 0;
    // optional return value to shell
}
```

```
$ gcc -o hello hello.c -std=c99
$
```

$\left. \begin{array}{l} -O2 \\ -Os \\ -O3 \\ -Ofast \end{array} \right\}$

# C review: Basics

- C was designed for (unix) system programming.
- C has a very small base.
- Most functionality is in (standard) libraries.
- We will use C99.

## Example (Basic C program)

```
#include <stdio.h>
// include stdio.h to print
int main() // always called first
{ // braces delimit code block
  printf("Hello world.\n");
  // function call to print
  // line ends with a semicolon
  return 0;
  // optional return value to shell
}
```

```
$ gcc -o hello hello.c -std=c99
```

```
$ ./hello
Hello world.
$
```

```
{-02
  -0s
  -03
  -0fast
```



# C review: Basics

- C was designed for (unix) system programming.
- C has a very small base.
- Most functionality is in (standard) libraries.
- We will use C99.

## Example (Basic C program)

```
#include <stdio.h>
// include stdio.h to print
int main() // always called first
{ // braces delimit code block
  printf("Hello world.\n");
  // function call to print
  // line ends with a semicolon
  return 0;
  // optional return value to shell
}
```

```
$ gcc -o hello hello.c -std=c99
$ ./hello
Hello world.
$ echo $?
0
$ █
```

$\left. \begin{array}{l} -O2 \\ -Os \\ -O3 \\ -Ofast \end{array} \right\}$

# C review: Functions

## Function declaration (prototype)

```
returntype name(argument-spec);
```

## Function definition

```
returntype name(argument-spec) {  
  statements  
}
```

## Function call

```
var=name(arguments);  
f(name(arguments));
```

## Procedures

Procedures are functions with return-type `void` ; called without assignment.

# C review: Variables

Define a variable with

```
type name [= value];
```

- *type* may be a
  - \* built-in type:
    - floating point type:  
`float, double, long double`
    - integer type:  
`short, [unsigned] int, [unsigned] long int, [unsigned] long long int`
    - character or string of characters:  
`char, char*`
  - \* array, pointer
  - \* structure, enumerated type, union
- Variable declarations and code may be mixed in C99.
- Variables can be initialized to a *value* when declared.

**Any non-initialized variable is not set to zero, but has a random value!**

# C review: Loops

```
for (initialization; condition; increment) {  
    statements  
}
```

```
while (condition) {  
    statements  
}
```

You can use `break` to exit the loop.

# C review: Loops

```
for (initialization; condition; increment) {  
    statements  
}
```

```
while (condition) {  
    statements  
}
```

You can use `break` to exit the loop.

## Example

```
#include <stdio.h>  
int main() {  
    for (int i=1; i<=10; i++)  
        printf("%d ",i);  
    // note the omitted braces  
    printf("\n");  
}
```

# C review: Loops

```
for (initialization; condition; increment) {  
    statements  
}
```

```
while (condition) {  
    statements  
}
```

You can use `break` to exit the loop.

## Example

```
#include <stdio.h>  
int main() {  
    for (int i=1; i<=10; i++)  
        printf("%d ",i);  
    // note the omitted braces  
    printf("\n");  
}
```

```
$ gcc -o count count.c -O2 -std=c99  
$ ./count  
1 2 3 4 5 5 6 7 8 9 10  
$ █
```

# C review: Pointers

```
type *name;
```

# C review: Pointers

```
type *name;
```

## Example (Pointer assignment)

```
#include <stdio.h>
int main() {
    int a=7,b=5;
    int *ptr=&a;
    a = 13;
    b = *ptr;
    printf("b=%d\n",b);
}
```

```
$ gcc -o ptrex ptrex.c -O2 -std=c99
$ ./ptrex
b=13
$ █
```



# C review: Pointers

```
type *name;
```

## Example (Pointer assignment)

```
#include <stdio.h>
int main() {
    int a=7,b=5;
    int *ptr=&a;
    a = 13;
    b = *ptr;
    printf("b=%d\n",b);
}
```

```
$ gcc -o ptrex ptrex.c -O2 -std=c99
$ ./ptrex
b=13
$ █
```

## Example (Pass by reference)

```
void inc(int *i) { (*i)++; }
int main() {
    int j=10;
    inc(&j);
    return j;
}
```

```
$ gcc -o passref passref.c -O2 -std=c99
$ ./passref
$ echo $?
11
$ █
```

# C review: Automatic arrays

```
type name[number];
```

- *name* is equivalent to a pointer to the first element.
- Usage *name[i]*. Equivalent to  $*(name+i)$ .
- C arrays are zero-based.

# C review: Automatic arrays

```
type name[number];
```

- *name* is equivalent to a pointer to the first element.
- Usage *name[i]*. Equivalent to  $*(name+i)$ .
- C arrays are zero-based.

## Example

```
#include <stdio.h>
int main() {
    int a[10]={1,2,3,4,5,6,7,8,9,11};
    int sum=0;
    for (int i=0; i<10; i++)
        sum += a[i];
    printf("sum=%d\n,sum);
}
```

```
$ gcc -o autoarr autoarr.c -O2 -std=c99
$ ./autoarr
56
$ █
```

# C review: Automatic arrays

```
type name[number];
```

- *name* is equivalent to a pointer to the first element.
- Usage *name[i]*. Equivalent to  $*(name+i)$ .
- C arrays are zero-based.

## Example

```
#include <stdio.h>
int main() {
    int a[10]={1,2,3,4,5,6,7,8,9,11};
    int sum=0;
    for (int i=0; i<10; i++)
        sum += a[i];
    printf("sum=%d\n",sum);
}
```

```
$ gcc -o autoarr autoarr.c -O2 -std=c99
$ ./autoarr
56
$ █
```

## Gotcha:

- There's a compiler dependent limit on *number*.
- C standard only says at least 65535 bytes.

# C review: Dynamically allocated arrays

Requires header file:

```
#include <stdlib.h>
```

Defined as a pointer to memory:

```
type *name;
```

Allocated by a function call:

```
name=malloc(number*sizeof(type));
```

Usages:

```
a=name[number];
```

Deallocated by a function call:

```
free(name);
```

- System function call can access all available memory.
- Can check if allocation failed ( $name == 0$ ).
- Can control when memory is given back.

# C review: Dynamically allocated arrays

Example

# C review: Dynamically allocated arrays

## Example

```
#include <stdlib.h>
#include <stdio.h>
void printarr(int n, int *a) {
    for (int i=0;i<n;i++)
        printf("%d ", a[i]);
    printf("\n");
}
int main(){
    int n=100;
    int *b=malloc(n*sizeof(*b));
    for (int i=0;i<n;i++)
        b[i]=i*i;
    printarr(n,b);
    free(b);
}
```

# C review: Dynamically allocated arrays

## Example

```
#include <stdlib.h>
#include <stdio.h>
void printarr(int n, int *a) {
    for (int i=0;i<n;i++)
        printf("%d ", a[i]);
    printf("\n");
}
int main(){
    int n=100;
    int *b=malloc(n*sizeof(*b));
    for (int i=0;i<n;i++)
        b[i]=i*i;
    printarr(n,b);
    free(b);
}
```

```
$ gcc -o dynarr dynarr.c -O2 -std=c99
$ ./dynarr
0 1 4 9 16 25 36 49 64 81 100 121 144
169 196 225 256 289 324 361 400 441
484 529 576 625 676 729 784 841 900
961 1024 1089 1156 1225 1296 1369 1444
1521 1600 1681 1764 1849 1936 2025
2116 2209 2304 2401 2500 2601 2704
2809 2916 3025 3136 3249 3364 3481
3600 3721 3844 3969 4096 4225 4356
4489 4624 4761 4900 5041 5184 5329
5476 5625 5776 5929 6084 6241 6400
6561 6724 6889 7056 7225 7396 7569
7744 7921 8100 8281 8464 8649 8836
9025 9216 9409 9604 9801
$ █
```



# C review: Structs = collections of other variables

```
struct name {  
    type1 name1;  
    type2 name2;  
    ...  
};
```

# C review: Structs = collections of other variables

```
struct name {  
    type1 name1;  
    type2 name2;  
    ...  
};
```

## Example

```
#include <string.h>  
#include <stdio.h>  
struct Info {  
    char name[100];  
    unsigned int age;  
};  
int main() {  
    struct Info my;  
    my.age=38;  
    strcpy(my.name, "Ramses");  
    printf("%d %s\\", my.age, my.name);  
}
```

# C review: Structs = collections of other variables

```
struct name {  
    type1 name1;  
    type2 name2;  
    ...  
};
```

## Example

```
#include <string.h>  
#include <stdio.h>  
struct Info {  
    char name[100];  
    unsigned int age;  
};  
int main() {  
    struct Info my;  
    my.age=38;  
    strcpy(my.name, "Ramses");  
    printf("%d %s\\", my.age, my.name);  
}
```

```
$ gcc -o info info.c -O2 -std=c99  
$ ./info  
Ramses 38  
$ █
```

# C review: Conditionals

```
if (condition) {  
    statements  
} else if (other condition) {  
    statements  
} else {  
    statements  
}
```

## Example

# C review: Conditionals

```
if (condition) {  
    statements  
} else if (other condition) {  
    statements  
} else {  
    statements  
}
```

## Example

```
int main(){  
    int n=20;  
    int *b= malloc(n*sizeof(*b));  
    if (b==0)  
        return 1; //error  
    else {  
        for (int i=0;i<n;i++)  
            b[i]=i*i;  
        printarr(n,b);  
        free(b);  
    }  
}
```

# C review: Conditionals

```
if (condition) {  
    statements  
} else if (other condition) {  
    statements  
} else {  
    statements  
}
```

## Example

```
int main(){  
    int n=20;  
    int *b= malloc(n*sizeof(*b));  
    if (b==0)  
        return 1; //error  
    else {  
        for (int i=0;i<n;i++)  
            b[i]=i*i;  
        printarr(n,b);  
        free(b);  
    }  
}
```

```
$ gcc -o ifm ifm.c -O2 -std=c99  
$ ./ifm  
0 1 4 9 16 25 36 49 64 81 100 121  
144 169 196 225 256 289 324 361  
$ █
```

# C review: Multidimensional arrays

```
#include <stdlib.h>
#include <assert.h>
float **matrix(long n,long m) {
    float **a=malloc(n*sizeof(*a));
    assert(a); // check if a not null
    a[0]=malloc(n*m*sizeof(**a));
    assert(a[0]); // check if a[0] not null
    for (long i=1; i<n; i++)
        a[i]=&data[i*m];
    return a;
}
void free_matrix(float **a) {
    free(a[0]);
    free(a);
}
void fill(long n,long m,float **a,float v){
    for (long i=0; i<n; i++)
        for (long j=0; j<m; j++)
            a[i][j]=v;
}
```

# C review: Multidimensional arrays

```
#include <stdlib.h>
#include <assert.h>
float **matrix(long n,long m) {
    float **a=malloc(n*sizeof(*a));
    assert(a); // check if a not null
    a[0]=malloc(n*m*sizeof(**a));
    assert(a[0]); // check if a[0] not null
    for (long i=1; i<n; i++)
        a[i]=&data[i*m];
    return a;
}
void free_matrix(float **a) {
    free(a[0]);
    free(a);
}
void fill(long n,long m,float **a,float v){
    for (long i=0; i<n; i++)
        for (long j=0; j<m; j++)
            a[i][j]=v;
}
```

a




# C review: Multidimensional arrays

```
#include <stdlib.h>
#include <assert.h>

float **matrix(long n,long m) {
    float **a=malloc(n*sizeof(*a));
    assert(a); // check if a not null
    a[0]=malloc(n*m*sizeof(**a));
    assert(a[0]); // check if a[0] not null
    for (long i=1; i<n; i++)
        a[i]=&data[i*m];
    return a;
}

void free_matrix(float **a) {
    free(a[0]);
    free(a);
}

void fill(long n,long m,float **a,float v){
    for (long i=0; i<n; i++)
        for (long j=0; j<m; j++)
            a[i][j]=v;
}
```




# C review: Multidimensional arrays

```
#include <stdlib.h>
#include <assert.h>

float **matrix(long n,long m) {
    float **a=malloc(n*sizeof(*a));
    assert(a); // check if a not null
    a[0]=malloc(n*m*sizeof(**a));
    assert(a[0]); // check if a[0] not null
    for (long i=1; i<n; i++)
        a[i]=&data[i*m];
    return a;
}

void free_matrix(float **a) {
    free(a[0]);
    free(a);
}

void fill(long n,long m,float **a,float v){
    for (long i=0; i<n; i++)
        for (long j=0; j<m; j++)
            a[i][j]=v;
}
```




# C review: Multidimensional arrays

```
#include <stdlib.h>
#include <assert.h>

float **matrix(long n,long m) {
    float **a=malloc(n*sizeof(*a));
    assert(a); // check if a not null
    a[0]=malloc(n*m*sizeof(**a));
    assert(a[0]); // check if a[0] not null
    for (long i=1; i<n; i++)
        a[i]=&data[i*m];
    return a;
}

void free_matrix(float **a) {
    free(a[0]);
    free(a);
}

void fill(long n,long m,float **a,float v){
    for (long i=0; i<n; i++)
        for (long j=0; j<m; j++)
            a[i][j]=v;
}
```




# C review: Multidimensional arrays

```
#include <stdlib.h>
#include <assert.h>

float **matrix(long n,long m) {
    float **a=malloc(n*sizeof(*a));
    assert(a); // check if a not null
    a[0]=malloc(n*m*sizeof(**a));
    assert(a[0]); // check if a[0] not null
    for (long i=1; i<n; i++)
        a[i]=&data[i*m];
    return a;
}

void free_matrix(float **a) {
    free(a[0]);
    free(a);
}

void fill(long n,long m,float **a,float v){
    for (long i=0; i<n; i++)
        for (long j=0; j<m; j++)
            a[i][j]=v;
}
```




# C review: Multidimensional arrays

```
#include <stdlib.h>
#include <assert.h>

float **matrix(long n,long m) {
    float **a=malloc(n*sizeof(*a));
    assert(a); // check if a not null
    a[0]=malloc(n*m*sizeof(**a));
    assert(a[0]); // check if a[0] not null
    for (long i=1; i<n; i++)
        a[i]=&data[i*m];
    return a;
}

void free_matrix(float **a) {
    free(a[0]);
    free(a);
}

void fill(long n,long m,float **a,float v){
    for (long i=0; i<n; i++)
        for (long j=0; j<m; j++)
            a[i][j]=v;
}
```




# C review: Multidimensional arrays

```
#include <stdlib.h>
#include <assert.h>

float **matrix(long n,long m) {
    float **a=malloc(n*sizeof(*a));
    assert(a); // check if a not null
    a[0]=malloc(n*m*sizeof(**a));
    assert(a[0]); // check if a[0] not null
    for (long i=1; i<n; i++)
        a[i]=&data[i*m];
    return a;
}

void free_matrix(float **a) {
    free(a[0]);
    free(a);
}

void fill(long n,long m,float **a,float v){
    for (long i=0; i<n; i++)
        for (long j=0; j<m; j++)
            a[i][j]=v;
}
```




# C review: Multidimensional arrays

```
#include <stdlib.h>
#include <assert.h>

float **matrix(long n,long m) {
    float **a=malloc(n*sizeof(*a));
    assert(a); // check if a not null
    a[0]=malloc(n*m*sizeof(**a));
    assert(a[0]); // check if a[0] not null
    for (long i=1; i<n; i++)
        a[i]=&data[i*m];
    return a;
}

void free_matrix(float **a) {
    free(a[0]);
    free(a);
}

void fill(long n,long m,float **a,float v){
    for (long i=0; i<n; i++)
        for (long j=0; j<m; j++)
            a[i][j]=v;
}
```



# C review: Libraries

## Usage

- Put an include line in the source code, e.g.

```
#include <stdio.h>
#include <omp.h>
#include "mpi.h"
```

- Include the libraries at link time using `-l[libname]`.  
Implicit for most standard libraries, with `mpicc` and `gcc -fopenmp`.



# C review: Libraries

## Usage

- Put an include line in the source code, e.g.

```
#include <stdio.h>
#include <omp.h>
#include "mpi.h"
```

- Include the libraries at link time using `-l[libname]`.  
Implicit for most standard libraries, with `mpicc` and `gcc -fopenmp`.

## Common standard libraries

- `stdio.h`: input/output, e.g., `printf` and `fwrite`
- `stdlib.h`: memory, e.g. `malloc`
- `string.h`: strings, memory copies, e.g. `strcpy`
- `math.h`: special functions, e.g. `sqrt`.


When using `math`, you need to link with `-lm`.



# Compilation:

Building with make

# Compilation workflow



# Compiling with make

## Single source file

```
# This file is called makefile
CC      = gcc
CFLAGS  = -std=c99 -O2
LDFLAGS = -lm
main:   main.c
        $(CC) $(CFLAGS) $(LDFLAGS) $^ -o $@
```

# Compiling with make

## Single source file

```
# This file is called makefile
CC      = gcc
CFLAGS  = -std=c99 -O2
LDFLAGS = -lm
main:   main.c
        $(CC) $(CFLAGS) $(LDFLAGS) $^ -o $@
```

## Multiple source file application

```
CC      = gcc
CFLAGS  = -std=c99 -O2
LDFLAGS = -lm
main:   main.o mylib.o
        $(CC) $(LDFLAGS) $^ -o $@
main.o: main.c mylib.h
mylib.o: mylib.h mylib.c
clean:
        rm -f main.o mylib.o
```

# Compiling with make

When typing `make` at command line:

- Checks if `main.c` or `mylib.c` or `mylib.h` were changed.
- If so, invokes corresponding rules for object files.
- Only compiles changed code files: faster recompilation.
- Parallel make:

```
$ make -j 3
```

# Compiling with make

When typing `make` at command line:

- Checks if `main.c` or `mylib.c` or `mylib.h` were changed.
- If so, invokes corresponding rules for object files.
- Only compiles changed code files: faster recompilation.
- Parallel make:

```
$ make -j 3
```

## Gotcha

- Make does not detect changes in compiler, or in system.
- But `.o` files are system/compiler dependent, so need to be recompiled.
- Always specify a “clean” rule in the makefile, so that moving from one system or compiler to another, you can do a fresh rebuild:

```
$ make clean  
$ make
```

# Before we start with OpenMP...

```
$ ssh -X login.scinet.utoronto.ca
$ ssh -X gpc01
$ qsub -X -I -l nodes=1:ppn=8,walltime=9:00:00,os=centos53develibA
```

to get a dedicated development node (ensure this works).

```
$ cp -R ~ljdursi/intro-ppp ~/
$ source ~/intro-ppp/setup
$ cd ~/intro-ppp/gettingstarted
$ make omp_hello_world
$ ./omp_hello_world

$ make mpi_hello_world
$ mpirun -np 8 omp_hello_world
```




# Part V

## Introduction to OpenMP

# OpenMP

- For shared memory systems.
- Add parallelism to functioning serial code.
- <http://openmp.org>



 **Subscribe to the News Feed**

---

**» OpenMP Specifications**

- » About OpenMP
- » Compilers
- » Resources
- » Discussion Forum

---

**Events**

- » Multicore Expo '11: May 2-5 at the McEnery Convention Center in San Jose, California in booth #2206
- » WOMP 2011 Call For Papers (pdf) - 7th International Workshop on OpenMP, June 13-15, 2011, Chicago USA

---

**Input Register**

Alert the OpenMP.org webmaster about new products, events, or updates and we'll post it here.

[webmaster@openmp.org](mailto:webmaster@openmp.org)

---

**Search OpenMP.org**

---

**Archives**

- » March 2011
- » February 2011
- » January 2011
- » October 2010
- » July 2010
- » May 2010
- » June 2009
- » April 2009
- » May 2008

## The OpenMP News

---

» OpenMP at Multicore Expo '11 - May 2-5 - San Jose, CA



Look for OpenMP exhibiting at the Multicore Expo, May 2-5 at the McEnery Convention Center in San Jose, California in booth #2206.

The key objectives of the Multicore Technical Conference and Expo are to identify emerging challenges faced by designers, to suggest potential solutions or review actual designs, and to aid embedded designers in their continuing engineering education. Co-located with the Embedded Systems Conference, the agenda for this event promises to be interesting, with tracks on Multicore Debugging, Multicore Frameworks, Parallel Technologies, Software Design, and more.

Visit us in our booth talk about the latest in OpenMP, get answers to your questions, learn about release 3.1 of the OpenMP API, and pick up the latest OpenMP reference card! For registration and other information visit <http://www.multicore-expo.com/>

Posted on March 11, 2011

---

» Parallel Programming in Computational Engineering and Science PPECES 2011



Seminar/Workshop: March 21 - March 25, 2011 Aachen, Germany <http://www.rz.rwth-aachen.de/ppces>

This event is now over, but the course material is available on the Seminar website.

This year's seminar will include a special introduction session on Monday to present the new HPC-cluster to be delivered by IUL. During the remainder of the week, we will cover Serial Programming, Tuning, Debugging and Processor Architectures (Tuesday), Shared Memory Programming with OpenMP (Wednesday), Message Passing with MPI (Thursday) and GPGPU Programming on Friday. Some of these lectures will feature hands-on sessions.

Attendees should be comfortable with C/C++ or Fortran programming and interested in learning more about the technical details of application tuning and parallelization on their favored platform (Windows or Linux). The presentations will be given in English.

Dierker May (RWTH), Thomas Warschko (IUL), Herbert Cornelius (IUL), Jean-Pierre Parcerisa (Bull), Christian Bischof (RWTH) and Felix Wolf (German Research School for Simulation Science) for our Monday event. The remainder of the week will be covered by Rued von der Pae (Oracle), Michael Wolfe (PGI) and speakers of the HPC Team of the RWTH Aachen University.

The seminar is free. Allocation is on a first come, first served basis, as we are limited in capacity. Please register separately for any session you intend to participate. Go to:

## The OpenMP API

supports multi-platform shared-memory parallel programming in C/C++ and Fortran. OpenMP is a portable, scalable model with a simple and flexible interface for developing parallel applications on platforms from the desktop to the supercomputer.

[» Read about OpenMP.org](#)

---

**Get**

- » OpenMP specs

---

**Use**

- » OpenMP Compilers

---

**Learn**



- » Using OpenMP -- the book
- » Using OpenMP -- the examples
- » Using OpenMP -- the forum
- » Wikpedia
- » OpenMP Tutorial
- » More Resources

---

**Discuss**

- » User Forum

Ask the experts and get answers to questions about OpenMP

# OpenMP

- For shared memory systems.
  - Add parallelism to functioning serial code.
  - <http://openmp.org>
- 
- Compiler, run-time environment does a lot of work for us
  - Divides up work
  - But we have to tell it how to use variables, where to run in parallel, ...
  - Mark parallel regions.
  - Works by adding compiler directives to code.
- Invisible to non-openmp compilers.**

## OpenMP™

THE OPENMP API SPECIFICATION FOR PARALLEL PROGRAMMING

### OpenMP News

• OpenMP at Multicore Expo '11 - May 2-5 - San Jose, CA

**11th ANNUAL multicore TECHNICAL CONFERENCE & EXPO**  
May 2-5, 2011  
McEnery Convention Center  
San Jose, California

Look for OpenMP exhibiting at the Multicore Expo, May 2-5 at the McEnery Convention Center in San Jose, California in booth #2206.

The key objectives of the Multicore Technical Conference and Expo are to identify emerging challenges faced by designers, to suggest potential solutions or review actual designs, and to aid embedded designers in their continuing engineering education. Co-located with the Embedded Systems Conference, the agenda for this event promises to be interesting, with tracks on Multicore Debugging, Multicore Frameworks, Parallel Technologies, Software Design, and more.

Visit us in our booth talk about the latest in OpenMP, get answers to your questions, learn about release 3.1 of the OpenMP API, and pick up the latest OpenMP reference card! For registration and other information visit <http://www.multicore-expo.com/>

Posted on March 11, 2011

• Parallel Programming in Computational Engineering and Science PPECES 2011

Seminar/Workshop: March 21 - March 25, 2011  
Aachen, Germany <http://www.rz.rwth-aachen.de/ppces>

This event is now over, but the course material is available on the Seminar website.

This year's seminar will include a special introduction session on Monday to present the new HPC-cluster to be delivered by ILL. During the remainder of the week, we will cover Serial Programming, Tuning, Debugging and Processor Architectures (Tuesday), Shared Memory Programming with OpenMP (Wednesday), Message Passing with MPI (Thursday) and GPGPU Programming on Friday. Some of these lectures will feature hands-on sessions.

Attendees should be comfortable with C/C++ or Fortran programming and interested in learning more about the technical details of application tuning and parallelization on their favored platform (Windows or Linux). The presentations will be given in English.

Dealer an May (RWTH), Thomas Warschko (ILL), Herbert Cornelius (Intel), Jean-Pierre Pansier (Bull), Christian Bischof (RWTH) and Felix Wolf (German Research School for Simulation Science) for our Monday event. The remainder of the week will be covered by Rüdiger van der Pas (Oracle), Michael Wolfe (PGI) and speakers of the HPC Team of the RWTH Aachen University.

The seminar is free. Allocation is on a first come, first served basis, as we are limited in capacity. Please register separately for any session you intend to participate. Go to:

### The OpenMP API

supports multi-platform shared-memory parallel programming in C/C++ and Fortran. OpenMP is a portable, scalable model with a simple and flexible interface for developing parallel applications on platforms from the desktop to the supercomputer.

• Read about OpenMP.org

Get

- OpenMP specs

Use

- OpenMP Compilers

Learn

- Using OpenMP - the book
- Using OpenMP - the examples
- Using OpenMP - the forum
- Wikipedia
- OpenMP Tutorial
- More Resources

Discuss

- User Forum

Ask the experts and get answers to questions about OpenMP

Subscribe to the News Feed

- OpenMP Specifications
- About OpenMP
- Compilers
- Resources
- Discussion Forum

Events

- Multicore Expo '11: May 2-5 at the McEnery Convention Center in San Jose, California in booth #2206
- WOMP 2011 Call For Papers (pdf) - 7th International Workshop on OpenMP, June 13-15, 2011, Chicago USA

Input Register

Alert the OpenMP.org webmaster about new products, events, or updates and we'll post it here.

[webmaster@openmp.org](mailto:webmaster@openmp.org)

Search OpenMP.org

Search

Archives

- March 2011
- February 2011
- January 2011
- October 2010
- July 2010
- May 2010
- June 2009
- April 2009
- May 2008

# OpenMP basic operations

## In code:

- In C, you add lines starting with `#pragma omp`. This parallelizes the subsequent code block.
- In Fortran, you add lines starting with `!$omp`. An `!$omp end ...` is needed to close the parallel region.
- These lines are skipped (for C, sometimes with a warning) by compilers that do not support OpenMP.

## When compiling:

- To turn on OpenMP support in gcc and gfortran, add the `-fopenmp` flag to the compilation (and link!) commands.

## When running:

- The environment variable `OMP_NUM_THREADS` determines how many threads will be started in an OpenMP parallel block.

# OpenMP example

C:

```
#include <stdio.h>
#include <omp.h>
int main() {
    printf("At start of program\n");
    #pragma omp parallel
    {
        printf("Hello world from thread %d!\n", omp_get_thread_num());
    }
}
```

F90:

```
program omp_hello_world
use omp_lib
implicit none
print *, 'At start of program'
!$omp parallel
    print *, 'Hello world from thread ', omp_get_thread_num(), '!'
!$omp end parallel
end program omp_hello_world
```

# OpenMP example

```
$ gcc -std=c99 -O2 -o omp-hello-world omp-hello-world.c -fopenmp
or
$ gfortran -O2 -o omp-hello-world omp-hello-world.f90 -fopenmp

$ export OMP_NUM_THREADS=8
$ ./omp-hello-world
...
$ export OMP_NUM_THREADS=1
$ ./omp-hello-world
...
$ export OMP_NUM_THREADS=32
$ ./omp-hello-world
...
```

Let's see what happens...

# OpenMP example

```
$ gcc -o omp-hello-world omp-hello-world.c -fopenmp
$ export OMP_NUM_THREADS=8
$ ./omp-hello-world
At start of program
Hello, world, from thread 0!
Hello, world, from thread 6!
Hello, world, from thread 5!
Hello, world, from thread 4!
Hello, world, from thread 2!
Hello, world, from thread 1!
Hello, world, from thread 7!
Hello, world, from thread 3!
$ export OMP_NUM_THREADS=1
$ ./omp-hello-world
At start of program
Hello, world, from thread 0!
$ export OMP_NUM_THREADS=32
$ ./omp-hello-world
At start of program
Hello, world, from thread 11!
Hello, world, from thread 1!
Hello, world, from thread 16!
```

# So what happened precisely?

- OMP\_NUM\_THREADS threads were launched.
- Each prints “Hello, world ...”;
- In seemingly random order.
- Only one “At start of program”.

```
$ gcc -o omp-hello-world omp-hello-world.o
$ export OMP_NUM_THREADS=8
$ ./omp-hello-world
At start of program
Hello, world, from thread 0!
Hello, world, from thread 6!
Hello, world, from thread 5!
Hello, world, from thread 4!
Hello, world, from thread 2!
Hello, world, from thread 1!
Hello, world, from thread 7!
Hello, world, from thread 3!
$ export OMP_NUM_THREADS=1
$ ./omp-hello-world
At start of program
Hello, world, from thread 0!
$ export OMP_NUM_THREADS=32
$ ./omp-hello-world
At start of program
Hello, world, from thread 11!
Hello, world, from thread 1!
Hello, world, from thread 16!
```



# So what happened precisely?

```
#include <stdio.h>
#include <omp.h>
int main() {
    printf("At start of program\n");
    #pragma omp parallel
    {
        printf("Hello world from thread %d!\n", omp_get_thread_num());
    }
}
```

```
program omp_hello_world
use omp_lib
implicit none
print *, 'At start of program'
!$omp parallel
    print *, 'Hello world from thread ', omp_get_thread_num(), '!'
!$omp end parallel
end program omp_hello_world
```

# So what happened precisely?

```
#include <stdio.h>
#include <omp.h>
int main() {
    printf("At start of program\n");
    #pragma omp parallel
    {
        printf("Hello world from thread %d!\n", omp_get_thread_num());
    }
}
```

Program starts normally (single thread)

```
program omp_hello_world
use omp_lib
implicit none
print *, 'At start of program'
!$omp parallel
    print *, 'Hello world from thread ', omp_get_thread_num(), '!'
!$omp end parallel
end program omp_hello_world
```

# So what happened precisely?

```
#include <stdio.h>
#include <omp.h>
int main() {
    printf("At start of program\n");
    #pragma omp parallel
    {
        printf("Hello world from thread %d!\n", omp_get_thread_num());
    }
}
```

At start of parallel section, launching  
OMP\_NUM\_THREADS threads,  
Each executes the same code!

```
program omp_hello_world
use omp_lib
implicit none
print *, 'At start of program'
!$omp parallel
    print *, 'Hello world from thread ', omp_get_thread_num(), '!'
!$omp end parallel
end program omp_hello_world
```

# So what happened precisely?

```
#include <stdio.h>
#include <omp.h>
int main() {
    printf("At start of program\n");
    #pragma omp parallel
    {
        printf("Hello world from thread %d!\n", omp_get_thread_num());
    }
}
```


At end of parallel section,  
threads join back up,  
Execution continues serially.

```
program omp_hello_world
use omp_lib
implicit none
print *, 'At start of program'
!$omp parallel
    print *, 'Hello world from thread ', omp_get_thread_num(), '!'
!$omp end parallel
end program omp_hello_world
```

# So what happened precisely?

```
#include <stdio.h>
#include <omp.h>
int main() {
    printf("At start of program\n");
    #pragma omp parallel
    {
        printf("Hello world from thread %d!\n", omp_get_thread_num());
    }
}
```

Special function to find number of current thread (first=0).



```
program omp_hello_world
use omp_lib
implicit none
print *, 'At start of program'
!$omp parallel
    print *, 'Hello world from thread ', omp_get_thread_num(), '!'
!$omp end parallel
end program omp_hello_world
```

# OpenMP functions (from omp.h/omp\_lib)

```
#include <stdio.h>
#include <omp.h>
int main() {
    printf("At start of program\n");
    #pragma omp parallel
    {
        printf("Hello world from thread %d of %d!\n",
            omp_get_thread_num(),
            omp_get_num_threads());
    }
}
```

omp\_get\_num\_threads() called by all threads.

Let's see if we can fix that...

# OpenMP functions (from omp.h/omp\_lib)

```
#include <stdio.h>
#include <omp.h>
int main() {
    printf("At start of program\n");
    #pragma omp parallel
    {
        printf("Hello world from thread %d!\n",
            omp_get_thread_num());
    }
    printf("There were %d threads.\n", omp_get_num_threads());
}
```

What do you think, will this work?

# OpenMP functions (from omp.h/omp\_lib)

```
#include <stdio.h>
#include <omp.h>
int main() {
    printf("At start of program\n");
    #pragma omp parallel
    {
        printf("Hello world from thread %d!\n",
            omp_get_thread_num());
    }
    printf("There were %d threads.\n", omp_get_num_threads());
}
```

What do you think, will this work?

No:

Says 1 thread only!



# OpenMP functions (from omp.h/omp\_lib)

```
#include <stdio.h>
#include <omp.h>
int main() {
    printf("At start of program\n");
    #pragma omp parallel
    {
        printf("Hello world from thread %d!\n",
            omp_get_thread_num());
    }
    printf("There were %d threads.\n", omp_get_num_threads());
}
```

What do you think, will this work?

No:

Says 1 thread only!

Why?

Because that is true outside the parallel region!

Need to get the value from the parallel region somehow.

# Variables in OpenMP

Variables in parallel regions are a bit tricky.

```
#include <stdio.h>
#include <omp.h>
int main() {
    int mythread, nthreads;
    #pragma omp parallel default(none) shared(nthreads) private(mythread)
    {
        mythread = omp_get_thread_num();
        if (mythread == 0)
            nthreads = omp_get_num_threads();
    }
    printf("There were %d threads.\n", nthreads);
}
```

# Variables in OpenMP

Variables in parallel regions are a bit tricky.

```
#include <stdio.h>
#include <omp.h>
int main() {
    int mythread, nthreads;
    #pragma omp parallel default(none) shared(nthreads) private(mythread)
    {
        mythread = omp_get_thread_num();
        if (mythread == 0)
            nthreads = omp_get_num_threads();
    }
    printf("There were %d threads.\n", nthreads);
}
```

Variable declarations  
How used in parallel region

# Variables in OpenMP

Variables in parallel regions are a bit tricky.

```
#include <stdio.h>
#include <omp.h>
int main() {
    int mythread, nthreads;
    #pragma omp parallel default(none) shared(nthreads) private(mythread)
    {
        mythread = omp_get_thread_num();
        if (mythread == 0)
            nthreads = omp_get_num_threads();
    }
    printf("There were %d threads.\n", nthreads);
}
```

Variable declarations  
How used in parallel region

- default(none) can save you hours of debugging!
- shared: each thread sees it and can modify (be careful!).  
**Preserves value.**
- private: each thread gets its own copy, invisible for others  
**Initial and final value undefined!**  
(Advanced: firstprivate, lastprivate – copy in/out.)

# Variables in OpenMP

Variables in parallel regions are a bit tricky.

```
#include <stdio.h>
#include <omp.h>
int main() {
    int mythread, nthreads;
    #pragma omp parallel default(none) shared(nthreads) private(mythread)
    {
        mythread = omp_get_thread_num();
        if (mythread == 0)
            nthreads = omp_get_num_threads();
    }
    printf("There were %d threads.\n", nthreads);
}
```

- Program runs, launches threads.
- Each thread gets copy of mythread.
- Only thread 0 writes to nthreads.

# Variables in OpenMP

Variables in parallel regions are a bit tricky.

```
#include <stdio.h>
#include <omp.h>
int main() {
    int mythread, nthreads;
    #pragma omp parallel default(none) shared(nthreads) private(mythread)
    {
        mythread = omp_get_thread_num();
        if (mythread == 0)
            nthreads = omp_get_num_threads();
    }
    printf("There were %d threads.\n", nthreads);
}
```

- Program runs, launches threads.
- Each thread gets copy of mythread.
- Only thread 0 writes to nthreads.
- Good idea to declare mythread locally!  
(avoids many bugs)

# Variables in OpenMP

Variables in parallel regions are a bit tricky.

```
#include <stdio.h>
#include <omp.h>
int main() {
    int nthreads;
    #pragma omp parallel default(none) shared(nthreads)
    {
        int mythread = omp_get_thread_num();
        if (mythread == 0)
            nthreads = omp_get_num_threads();
    }
    printf("There were %d threads.\n", nthreads);
}
```

- Program runs, launches threads.
- Each thread gets copy of mythread.
- Only thread 0 writes to nthreads.
- Good idea to declare mythread locally!  
(avoids many bugs)

# Variables in OpenMP - Fortran version

```
program omp_vars
use omp_lib
implicit none
integer :: mythread, nthreads
!$omp parallel default(none) private(mythread) shared(nthreads)
  mythread = omp_get_thread_num()
  if (mythread == 0) then
    nthreads = omp_get_num_threads()
  endif
!$omp end parallel
print *, 'Number of threads was ', nthreads, '.'
end program omp_vars
```



# Single Execution in OpenMP

```
#include <stdio.h>
#include <omp.h>
int main() {
    int nthreads;
    #pragma omp parallel default(none) shared(nthreads) {
        int mythread = omp_get_thread_num();
        if (mythread == 0)
            nthreads = omp_get_num_threads();
    }
    printf("There were %d threads.\n", nthreads);
}
```

- Do we care that it's thread 0 in particular that updates nthreads?
- Often, we just want the first thread to go through, do not care which one.

# Single Execution in OpenMP

```
#include <stdio.h>
#include <omp.h>
int main() {
    int nthreads;
    #pragma omp parallel default(none) shared(nthreads)
    #pragma omp single
        nthreads = omp_get_num_threads();
    printf("There were %d threads.\n", nthreads);
}
```

```
program omp_vars
use omp_lib
implicit none
integer :: nthreads
!$omp parallel default(none) shared(nthreads)
!$omp single
    nthreads = omp_get_num_threads()
!$omp end single
!$omp end parallel
print *, 'Number of threads was ', nthreads, '.'
end program omp_vars
```

# Loops in OpenMP

Take one of your openmp programs and add a loop.

# Loops in OpenMP

Take one of your openmp programs and add a loop.

```
#include <stdio.h>
#include <omp.h>
int main() {
    int i, mythread;
    #pragma omp parallel default(none) \
        XXXX(i) XXXX(mythread)
    {
        mythread = omp_get_thread_num();
        for (i=0; i<16; i++)
            printf("Thread %d gets i=%d\n",
                mythread, i);
    }
}
```

```
program omp_loop
use omp_lib
implicit none
integer :: i, mythread
!$omp parallel default(none) &
!$omp XXXX(i) XXXX(mythread)
    mythread = omp_get_thread_num()
    do i=1,16
        print *, 'thread ', mythread, &
            ' gets i=', i
    enddo
!$omp end parallel
end program omp_loop
```

# Loops in OpenMP

Take one of your openmp programs and add a loop.

```
#include <stdio.h>
#include <omp.h>
int main() {
    int i, mythread;
    #pragma omp parallel default(none) \
        XXXX(i) XXXX(mythread)
    {
        mythread = omp_get_thread_num();
        for (i=0; i<16; i++)
            printf("Thread %d gets i=%d\n",
                mythread, i);
    }
}
```

```
program omp_loop
use omp_lib
implicit none
integer :: i, mythread
!$omp parallel default(none) &
!$omp XXXX(i) XXXX(mythread)
    mythread = omp_get_thread_num()
    do i=1,16
        print *, 'thread ', mythread, &
            ' gets i=', i
    enddo
!$omp end parallel
end program omp_loop
```

What would you imagine this does when run with e.g.  
OMP\_NUM\_THREADS=8?

# Worksharing constructs in OpenMP

- We don't generally want tasks to do exactly the same thing.
- Want to partition a problem into pieces, each thread works on a piece.
- Most scientific programming full of work-heavy loops.
- OpenMP has a worksharing construct: `omp for` (or `omp do`).

# Worksharing constructs in OpenMP

- We don't generally want tasks to do exactly the same thing.
- Want to partition a problem into pieces, each thread works on a piece.
- Most scientific programming full of work-heavy loops.
- OpenMP has a worksharing construct: `omp for` (or `omp do`).

```
#include <stdio.h>
#include <omp.h>
int main() {
    int i, mythread;
    #pragma omp parallel default(none) XXXX(i) XXXX(mythread)
    {
        mythread = omp_get_thread_num();
        #pragma omp for
        for (i=0; i<16; i++)
            printf("Thread %d gets i=%d\n",mythread,i);
    }
}
```

# Fortran version

```
program omp_loop
use omp_lib
implicit none
integer :: i, mythread
!$omp parallel default(none) XXXX(i) XXXX(mythread)
  mythread = omp_get_thread_num()
  !$omp do
  do i=1,16
    print *, 'thread ', mythread, ' gets i=', i
  enddo
  !$omp end do
!$omp end parallel
end program omp_loop
```



# Worksharing constructs in OpenMP

- `omp for/omp do` construct breaks up the iterations by thread.
- If doesn't divide evenly, does the best it can.
- Allows easy breaking up of work!
- Advanced: can break up work of arbitrary blocks of code with `omp task` construct.

```
$ ./omp_loop
thread 3 gets i= 6
thread 3 gets i= 7
thread 4 gets i= 8
thread 4 gets i= 9
thread 5 gets i= 10
thread 5 gets i= 11
thread 6 gets i= 12
thread 6 gets i= 13
thread 1 gets i= 2
thread 1 gets i= 3
thread 0 gets i= 0
thread 0 gets i= 1
thread 2 gets i= 4
thread 2 gets i= 5
thread 7 gets i= 14
thread 7 gets i= 15
$
```

# Less trivial example: DAXPY

- multiply a vector by a scalar, add a vector.
- (a X plus Y, in double precision)
- Implement this, first serially, then with OpenMP
- daxpy.c or daxpy.f90
- make daxpy or make fdaxpy

$$z = ax + y$$

## Warning

This is a common linear algebra construct that you really shouldn't implement yourself. Various so-called BLAS implementations will do a much better job than you. But good for illustration.

```

#include <stdio.h>
#include "pca_utils.h"
void daxpy(int n, double a, double *x, double *y, double *z) {
    for (int i=0; i<n; i++) {
        x[i] = (double)i*(double)i;
        y[i] = ((double)i+1.)*((double)i-1.);
    }
    for (int i=0; i<n; i++)
        z[i] += a * x[i] + y[i];
}
int main() {
    int n=1e7;
    double *x = vector(n);
    double *y = vector(n);
    double *z = vector(n);
    double a = 5./3.;
    pca_time tt;
    tick(&tt);
    daxpy(n,a,x,y,z);
    tock(&tt);
    free(z);
    free(y);
    free(x);
}

```

```
#include <stdio.h>
#include "pca_utils.h"
void daxpy(int n, double a, double *x, double *y, double *z) {
    for (int i=0; i<n; i++) {
        x[i] = (double)i*(double)i;
        y[i] = ((double)i+1.)*((double)i-1.);
    }
    for (int i=0; i<n; i++)
        z[i] += a * x[i] + y[i];
}
int main() {
    int n=1e7;
    double *x = vector(n);
    double *y = vector(n);
    double *z = vector(n);
    double a = 5./3.;
    pca_time tt;
    tick(&tt);
    daxpy(n,a,x,y,z);
    tock(&tt);
    free(z);
    free(y);
    free(x);
}
```

Utilities for this course; `double` is a numerical type which can be set to single or double precision

```
#include <stdio.h>
#include "pca_utils.h"
void daxpy(int n, double a, double *x, double *y, double *z) {
    for (int i=0; i<n; i++) {
        x[i] = (double)i*(double)i;
        y[i] = ((double)i+1.)*((double)i-1.);
    }
    for (int i=0; i<n; i++)
        z[i] += a * x[i] + y[i];
}
int main() {
    int n=1e7;
    double *x = vector(n);
    double *y = vector(n);
    double *z = vector(n);
    double a = 5./3.;
    pca_time tt;
    tick(&tt);
    daxpy(n,a,x,y,z);
    tock(&tt);
    free(z);
    free(y);
    free(x);
}
```

← Fill arrays with calculated values.

```
#include <stdio.h>
#include "pca_utils.h"
void daxpy(int n, double a, double *x, double *y, double *z) {
    for (int i=0; i<n; i++) {
        x[i] = (double)i*(double)i;
        y[i] = ((double)i+1.)*((double)i-1.);
    }
    for (int i=0; i<n; i++)
        z[i] += a * x[i] + y[i];
}
int main() {
    int n=1e7;
    double *x = vector(n);
    double *y = vector(n);
    double *z = vector(n);
    double a = 5./3.;
    pca_time tt;
    tick(&tt);
    daxpy(n,a,x,y,z);
    tock(&tt);
    free(z);
    free(y);
    free(x);
}
```

Do calculation.



```
#include <stdio.h>
#include "pca_utils.h"
void daxpy(int n, double a, double *x, double *y, double *z) {
    for (int i=0; i<n; i++) {
        x[i] = (double)i*(double)i;
        y[i] = ((double)i+1.)*((double)i-1.);
    }
    for (int i=0; i<n; i++)
        z[i] += a * x[i] + y[i];
}
int main() {
    int n=1e7;
    double *x = vector(n);
    double *y = vector(n);
    double *z = vector(n);
    double a = 5./3.;
    pca_time tt;
    tick(&tt);
    daxpy(n,a,x,y,z);
    tock(&tt);
    free(z);
    free(y);
    free(x);
}
```

← Driver (setup, call, timing).

```
#include <stdio.h>
#include "pca_utils.h"
void daxpy(int n, double a, double *x, double *y, double *z) {
    for (int i=0; i<n; i++) {
        x[i] = (double)i*(double)i;
        y[i] = ((double)i+1.)*((double)i-1.);
    }
    for (int i=0; i<n; i++)
        z[i] += a * x[i] + y[i];
}
int main() {
    int n=1e7;
    double *x = vector(n);
    double *y = vector(n);
    double *z = vector(n);
    double a = 5./3.;
    pca_time tt;
    tick(&tt);
    daxpy(n,a,x,y,z);
    tock(&tt);
    free(z);
    free(y);
    free(x);
}
```

HANDS-ON: Try OpenMPing...



HANDS-ON 1:  
Parallelize daxpy with OpenMP.  
Also do the scaling analysis!