Directives Beyond Shared Memory

Dr. Michael K. Bane

HIGH END COMPUTE

http://highendcompute.co.uk

OMP Versions 1, 2, 3

- OpenMP formed as *the* standard for shared memory programming
  - Directives to set-up parallel regions (v1)
  - Directives to share the work (v1)
  - Directives for task-based (v3)

Tasks

- Quantum of independent work
  - "independent" as in *internal* work can proceed without any need for further input
- Then define the simulation as
  - Set of tasks
  - Dependency between tasks (eg DAG)
  - More of a *dataflow* approach
- Presuming an excellent task manager, then should get good throughput and speedup

Tasks: The OpenMP Way

- Create parallel region
- Have a single thread create the tasks
- Then the tasks launch (one per thread over all threads of parallel region)

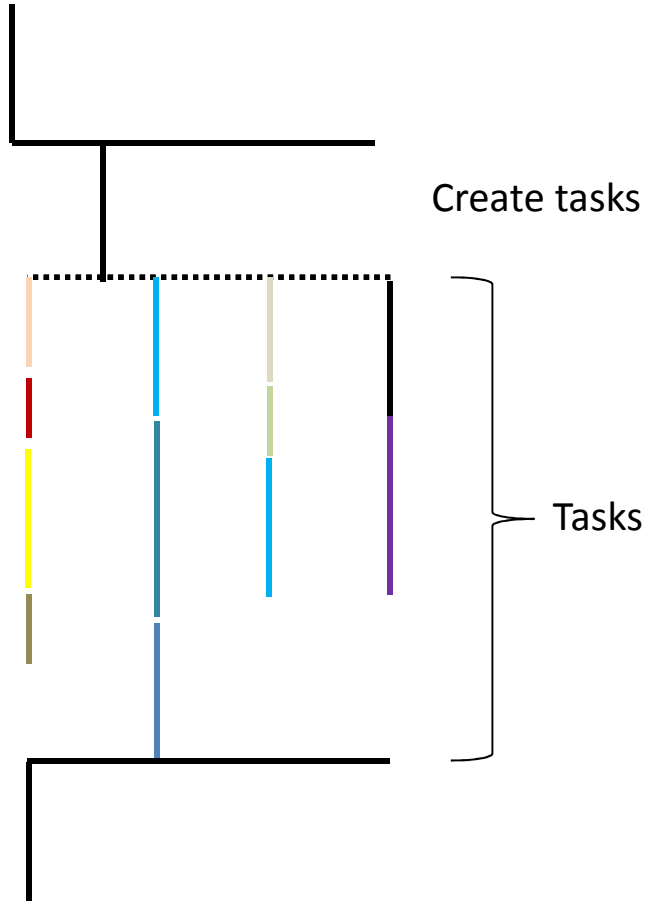

The Tasking Concept In OpenMP

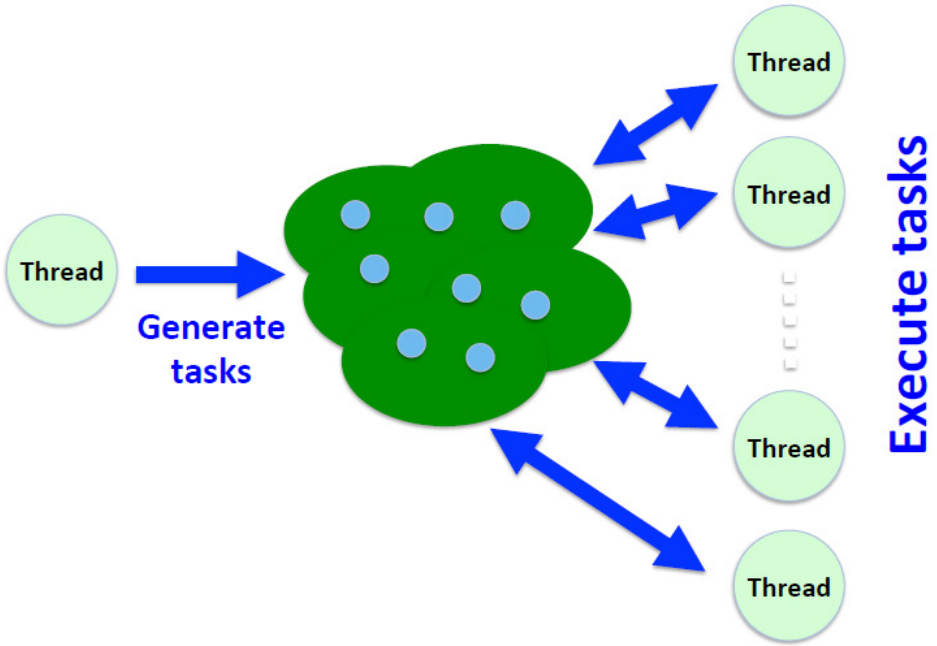

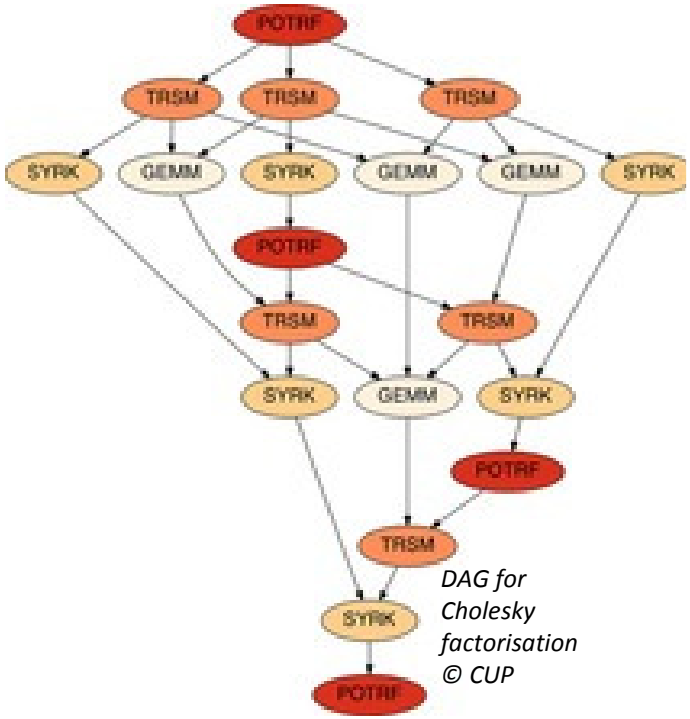

DAG for Cholesky factorisation © CUP

```
#pragma omp parallel
{
#pragma omp single
{
printf("A ");
#pragma omp task
{printf("car ");}
#pragma omp task
{printf("race ");}
#pragma omp taskwait
printf("is fun to watch ");
}
} // End of parallel region
```

2 tasks

Synchronisation (barrier) for tasks


```
$ cc -xopenmp -fast hello.c
$ export OMP_NUM_THREADS=2
$ ./a.out
A car race is fun to watch
$ ./a.out
A car race is fun to watch
$ ./a.out
A race car is fun to watch
```

What about Accelerators?

- OpenMP 4 introduced directives to offload work to a co-processor (GPU, KNC at end of PCI-e)
- OpenMP 4.5 refined & improved

- OpenACC
  - Directives based
  - Somewhere similar to OpenMP (liked by Intel)
  - Moves more quickly, but less vendors (loved by NVIDIA)

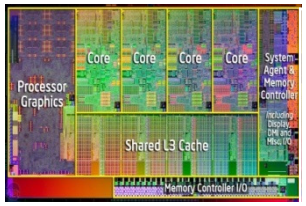

CPU

1 to maybe 64 cores, running at 2 to 3 GHz

High clock speed but general purpose

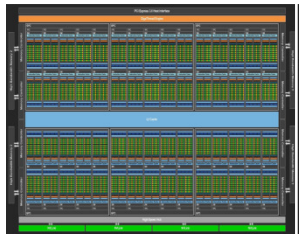

GPU

15 to 56 "streaming multiprocessors" (SMs), each with 64-128 "CUDA Cores". Base freq about 1 GHz

Very high throughput of vector arithmetic (particularly integer)

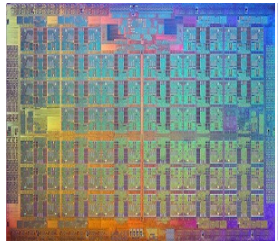

Xeon Phi

60-70 cores

Low grunt but general purpose cores

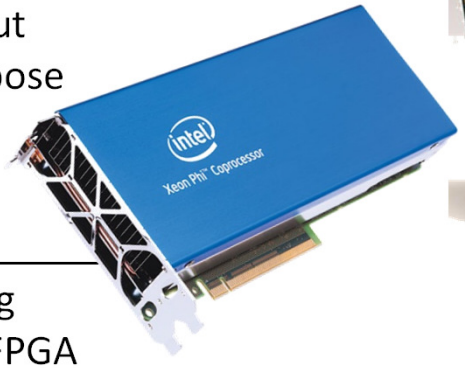

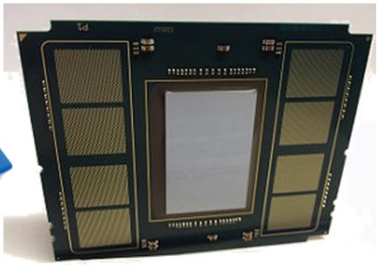

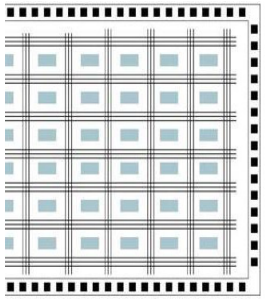

FPGA

Research paper now out showing OpenMP code being pushed to FPGA (without user doing intermediate steps)

ASIC

(out of the reach of us mere mortals!)

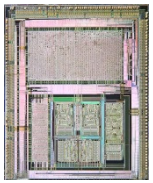

OpenMP Example

```
!$OMP PARALLEL DO
DO I=1, N
  Y(I) = A*X(I)*X(I) + B*X(I) + C
END DO
!$OMP END PARALLEL DO
```

CPU

```
!$OMP PARALLEL TARGET DEVICE(0) DO
DO I=1, N
  Y(I) = A*X(I)*X(I) + B*X(I) + C
END DO
!$OMP END PARALLEL DO
```

ACCELERATOR

TARGET is referring to a device (GPU or XPhi) for pushing the iterations of the DO loop. Impl dep how defined hw to DEVICE(n)

For GPU, most likely also want to use TEAMS DISTRIBUTE to make effective use of their Streaming Multiprocessors

```
!$OMP PARALLEL TARGET DO
DO I=1, N
  Y(I) = A*X(I)*X(I) + B*X(I) + C
END DO
!$OMP END PARALLEL DO
```

OpenMP 4

OpenACC Example

```
!$ACC PARALLEL LOOP
DO I=1, N
  Y(I) = A*X(I)*X(I) + B*X(I) + C
END DO
!$ACC END PARALLEL LOOP
```

!$ACC Parallel directive: put it on the accelerator

!$ACC Loop directive: spread iterations over threads

Improved efficiency by use of !$ACC Data directive (determine which IO (and can do async IO) to accelerator)

Want to know/try more…

- "Supplementary Materials directory
- HEC qwikLabs offer
  - Choice of labs that use GPUs in AWS cloud
  - Tokens for *you*
    indicate on feedback form