<section-header><section-header><section-header><section-header><section-header><section-header><section-header><section-header><text><text><text><text><image>


```
__shared__ int deltaE[THREADS];
deltaE[n] = ie;
for(int stride = THREADS>>1; stride > 0; stride >>= 1) {
   __syncthreads();
   if(n < stride) deltaE[n] += deltaE[n+stride];
}
if(n == 0) result[blockIdx.y*GRIDL+blockIdx.x] += deltaE[0];
```

Ising model: Measurements

Consider Metropolis kernel for the 2D Ising model discussed before:

GPU code v3 - kernel

}

```
__global__ void metro_checkerboard_three(spin_t *s, int *ranvec, int offset)
```

```
int n = blockDim.x*blockIdx.x + threadIdx.x;
int cur = blockDim.x*blockIdx.x + threadIdx.x + offset*(N/2);
int north = cur + (1-2*offset)*(N/2);
int east = ((north+1)%L) ? north + 1 : north-L+1;
int west = (north%L) ? north - 1 : north+L-1;
int south = (n - (1-2*offset)*L + N/2)%(N/2) + (1-offset)*(N/2);
int ide = s[cur]*(s[west]+s[north]+s[east]+s[south]);
if(fabs(RAN(ranvec[n])*4.656612e-10f) < tex1Dfetch(boltzT, ide+2*DIM)) {
s[cur] = -s[cur];
```

How can measurements of the internal energy, say, be incorporated?

 \Rightarrow local changes can be tracked

M. Weigel (Coventry/Mainz)

advanced simulations

01/11/2012 3 / 34



7

Ising Spin glass

Recall Hamiltonian:

$$\mathcal{H} = -\sum_{\langle i,j\rangle} J_{ij} s_i s_j$$

where J_{ij} are quenched random variables. For reasonable equilibrium results, average over thousands of realizations is necessary.

- same domain decomposition (checkerboard)
- slightly bigger effort due to non-constant couplings
- higher performance due to larger independence?
- very simple to combine with parallel tempering

Spin glass: performance



M. Weigel (Coventry/Mainz)

advanced simulations

01/11/2012 6 / 34

01/11/2012 0

Spin glasses: continued

Seems to work well with

- 15 ns per spin flip on CPU
- 70 ps per spin flip on GPU

but not better than ferromagnetic Ising model.

Further improvement: use multi-spin coding

- Synchronous multi-spin coding: different spins in a single configurations in one word
- Asynchronous multi-spin coding: spins from different realizations in one word

```
\Rightarrow brings us down to about 2 ps per spin flip
```

```
Local updates
Implementation
for(int i = 0; i < SWEEPS_LOCAL; ++i) {</pre>
  float r = RAN(ranvecS[n]);
  if(r < boltzD[4]) sS(x1,y) = -sS(x1,y);
  else {
   if(r < boltzD[2]) {
     ido = p1 | p2 | p3 | p4;
     sS(x1,y) = ido \cap sS(x1,y);
   } else {
     ido1 = p1 & p2; ido2 = p1 ^ p2;
     ido3 = p3 & p4; ido4 = p3 ^ p4;
     ido = ido1 | ido3 | (ido2 & ido4);
     sS(x1,y) = ido ^ sS(x1,y);
   3
  3
  __syncthreads();
  r = RAN(ranvecS[n]);
  if(r < boltzD[4]) sS(x2,y) = -sS(x2,y);
  else f
   if(r < boltzD[2]) {
    ido = p1 | p2 | p3 | p4;
     sS(x2,y) = ido ^ sS(x2,y);
   } else {
     ido1 = p1 & p2; ido2 = p1 ^ p2;
     ido3 = p3 & p4; ido4 = p3 ^ p4;
ido = ido1 | ido3 | (ido2 & ido4);
     sS(x2,y) = ido \cap sS(x2,y);
  3
   _syncthreads();
```

Spin glasses: continued

Seems to work well with

- 15 ns per spin flip on CPU
- ${\rm \circ}~70~{\rm ps}$ per spin flip on GPU

but not better than ferromagnetic Ising model.

Further improvement: use multi-spin coding

- Synchronous multi-spin coding: different spins in a single configurations in one word
- Asynchronous multi-spin coding: spins from different realizations in one word
- \Rightarrow brings us down to about $2~\mathrm{ps}$ per spin flip

M. Weigel (Coventry/Mainz)

advanced simulations

Janus

JANUS, a modular massively parallel and reconfigurable FPGA-based computing system.

Local updates

	JAN	US	PC			
MODEL	Algorithm	Max size	perfs	AMSC	SMSC	NO MSC
3D Ising EA	Metropolis	96 ³	16 ps	$45 \times$	$190 \times$	
3D Ising EA	Heat Bath	96 ³	16 ps	$60 \times$		
Q = 4 3D Glassy Potts	Metropolis	16^{3}	64 ps	$1250 \times$	$1900 \times$	
Q = 4 3D disordered Potts	Metropolis	88 ³	32 ps	$125 \times$		$1800 \times$
$Q = 4, C_m = 4$ random graph	Metropolis	24000	2.5 ns	$2.4 \times$		$10 \times$

Costs:

- ${\ensuremath{\, \circ }}$ Janus: 256 units, total cost about 700,000 Euros
- Same performance with GPU: 64 PCs (2000 Euros) with 2 GTX 295 cards (500 Euros) \Rightarrow 200,000 Euros
- Same performance with CPU only (assuming a speedup of ~ 50): 800 blade servers with two dual Quadcore sub-units (3500 Euros) $\Rightarrow 2,800,000$ Euros

JANUS, a modular massively parallel and reconfigurable FPGA-based computing system.



Critical configuration

Janus



01/11/2012

10 / 34

Cluster algorithms

Beat critical slowing down

with local spin flips.

Need to

Spatial correlations close to a continuous

• have the right geometric properties,

i.e., fractal dimensions etc.

phase transition impede decorrelation

• update non-local variables

• of characteristic size, i.e., percolating at transition, that

Cluster updates



Cluster algorithms

These are strong requirements only fulfilled for a few algorithms, most notably for the Potts, O(n) and related lattice models.

01/11/2012

14 / 34

M. Weigel (Coventry/Mainz)

advanced simulations

Cluster updates

Swendsen-Wang update



Cluster algorithms

Would need to use cluster algorithms for efficient equilibrium simulation of spin models at criticality:

- (1) Activate bonds between like spins with probability $p = 1 e^{-2\beta J}$.
- ② Construct (Swendsen-Wang) spin clusters from domains connected by active bonds.
- (3) Flip independent clusters with probability 1/2.
- ④ Goto 1.

Steps 1 and 3 are local \Rightarrow Can be efficiently ported to GPU. What about step 2? \Rightarrow Domain decomposition into tiles.



Swendsen-Wang update



Cluster updates

Label activation

one thread per site

thread updates two bonds

• low load, therefore updating more bonds per thread beneficial

M. Weigel (Coventry/Mainz)

advanced simulations

01/11/2012 18 / 34

Cluster updates

BFS code

__global__ void localBFS(int *inclus, bond_t *bond, int *next) int clus = 0, clus_leng, tested; int xmin = blockIdx.x*(L/gridDim.x), xmax = (blockIdx.x+1)*(L/gridDim.x); int ymin = blockIdx.y*(L/gridDim.y), ymax = (blockIdx.y+1)*(L/gridDim.y); int offset = ymin*L+xmin; for(int y = ymin; y < ymax; ++y) { for(int x = xmin; x < xmax; ++x) { if(inclus[y*L+x] <= clus) continue;</pre> clus_leng = tested = 0; clus = next[offset] = y*L+x; while(tested <= clus_leng) {</pre> int xx = next[offset+tested] % L; int yy = next[offset+tested] / L; int k1 = yy * L + (xx + 1);if(xx < xmax-1 && inclus[k1] > clus && bond[yy*L+xx]) { ++clus_leng; next[offset+clus_leng] = k1; inclus[k1] = clus; k1 = yy * L + (xx - 1);if(xx > xmin && inclus[k1] > clus && bond[k1]) { ++clus_leng; next[offset+clus_leng] = k1; inclus[k1] = clus; 3 // other directions ++tested; } } 3

BFS or Ants in the Labyrinth

56	57	58	59	60	61	62	63
48	49	50	51	52	53	54	55
40	41	42	43	36	-36	46	47
32	33	34	36-	-36-	-36	38	39
24	25	26	36	28	36	30	31
16	17	18	19	20	21	22	23
8	9	10	11	12	13	14	15
0	1	2	3	4	5	6	7

advanced simulations

only wave-front vectorization would be possible \Rightarrow many idle threads

Cluster updates

M. Weigel (Coventry/Mainz)

01/11/2012 19 / 34

Self-labeling



effort is $O(L^3)$ at the critical point, but can be vectorized with $O(L^2)$ threads

advanced simulations

Cluster updates

Self-labeling code



- guarantee no interference by other threads
- fast intrinsic implementation
- atomicAdd(), atomicMin(), atomicInc(), and others
- __syncthreads_or() synchronizes threads and returns true if *any* of the thread expression evaluates to true

M. Weigel (Coventry/Mainz)

advanced simulations

01/11/2012 22 / 34

Cluster updates

Union-find

56	57	58	59	60	61	62	63
48	41	41	51	52	53	54	55
40	32	41	41	44	45	46	47
32	32	34	30	30	30	38	39
24	25	26	27	30	30	13	31
16	17	18	19	20	21	13	23
8	9	10	11		13	13	15
0	1	2	3	4	5	6	7

tree structure with two optimizations:

balanced trees
 path compression

path compression

M. Weigel (Coventry/Mainz)

 \Rightarrow root finding and cluster union

essentially O(1) operations

Self-labeling code

self-labeling on tiles

```
do {
    changed = 0;
    if(bnd[me] && in[me] != in[right]) {
        in[me] = in[right] = min(in[me], in[right]);
        changed = 1;
    }
    if(bnd[me+BLOCKL*BLOCKL] && in[me] != in[up]) {
        in[me] = in[up] = min(in[me], in[up]);
        changed = 1;
    }
} while(__syncthreads_or(changed));
```

Cluster updates

- synchronization not strictly necessary here
- "non-deterministic" version is faster
- might need a couple of more iterations

M. Weigel (Coventry/Mainz)

advanced simulations

01/11/2012 22 / 34

Comparison



advanced simulations

Cluster updates

Cluster updates

Label consolidation



Cluster updates

Performance



Cluster updates Performance 400- relax - setup boundary 300 - label local bonds $\underset{\rm Su}{\rm diff} 200$ 100 0_{32} 64 128 25651210242048 L

Performance

512

L

2048

128

32

8192

Multicanonical simulations



② Choose multicanonical weights $\omega_1(E) = \hat{S}_0(E)$ for next simulation.





Multicanonical simulations

Choice of weights

To overcome barriers, we need to broaden P(E), in the extremal case to a constant distribution.

$$P_{\text{muca}}(E) = Z_{\text{muca}}^{-1} \Omega(E) / W(E) = Z_{\text{muca}}^{-1} e^{S(E) - \omega(E)} \stackrel{!}{=} \text{const}$$

where $S(E) = \ln \Omega(E)$ is the microcanonical entropy.

Under these assumptions, $W(E) = \Omega(E)$ is optimal, i.e., we again desire to estimate the density of states. This is not known a priori, so (again) use histogram estimator

 $\hat{\Omega}(E) = Z_{\text{muca}} \, \hat{H}_{\text{muca}}(E) / N \times e^{\omega(E)}.$

Canonical averages can be recovered at any time by reweighting:

$$\langle A \rangle_{K} = \frac{\sum_{E} A(E) P_{K}(E) / P_{\text{muca}}(E)}{\sum_{E} P_{K}(E) / P_{\text{muca}}(E)}$$

advanced simulations

M. Weigel (Coventry/Mainz)

01/11/2012 28 / 34

Multicanonical and Wang-Landau simulations Multicanonical simulations

Muca iteration

Determine muca weights/density of states iteratively:

- (1) Use, e.g., a K = 0 canonical simulation to get initial estimate $\hat{S}_0(E) = \ln \hat{\Omega}_0(E).$
- ② Choose multicanonical weights $\omega_1(E) = \hat{S}_0(E)$ for next simulation.

Iterate.



Multicanonical simulations





Multicanonical simulations

Muca iteration

Determine muca weights/density of states iteratively:

- (1) Use, e.g., a K = 0 canonical simulation to get initial estimate $\hat{S}_0(E) = \ln \hat{\Omega}_0(E)$.
- 2 Choose multicanonical weights $\omega_1(E) = \hat{S}_0(E)$ for next simulation.

Iterate.

Advantages:

- always in equilibrium
- arbitrary distributions possible
- system ideally performs an unbiased random walk in energy space → fast(er) dynamics



01/11/2012

28 / 34

M. Weigel (Coventry/Mainz)

Multicanonical and Wang-Landau simulations

advanced simulations

Multicanonical simulations



 $\omega_{i+1}(E) - \omega_i(E) = \phi,$

each time an energy E is seen.

 $(1) \\ (2) \\ (3) \\ (4) \\ (4) \\ (5) \\ (6) \\ (7)$

Multicanonical simulations



Multicanonical and Wang-Landau simulations

Multicanonical simulations

Wang-Landau simulations



Summarv

Performance

Benchmark results for various models considered:

			CPU	C1060	GTX 480	
System	Algorithm	L	ns/flip	ns/flip	ns/flip	speed-up
2D Ising	Metropolis	32	8.3	2.58	1.60	3/5
2D Ising	Metropolis	16 384	8.0	0.077	0.034	103/235
2D Ising	Metropolis, $k = 1$	16 384	8.0	0.292	0.133	28/60
3D Ising	Metropolis	512	14.0	0.13	0.067	107/209
2D Heisenberg	Metro. double	4096	183.7	4.66	1.94	39/95
2D Heisenberg	Metro. single	4096	183.2	0.74	0.50	248/366
2D Heisenberg	Metro. fast math	4096	183.2	0.30	0.18	611/1018
2D spin glass	Metropolis	32	14.6	0.15	0.070	97/209
2D spin glass	Metro. multi-spin	32	0.18	0.0075	0.0023	24/78
2D Ising	Swendsen-Wang	10240	77.4	_	2.97	-/26
2D Ising	multicanonical	64	42.1	_	0.33	-/128
2D Ising	Wang-Landau	64	43.6	_	0.94	-/46

Wang-Landau simulations



All the rest

What I haven't talked about:

- no-copy pinning of system memory
- asynchronous execution and data transfers: streams etc.

Summarv

- surface memory
- GPU-GPU communication
- C++ language features:
 - ${\scriptstyle \circ \ }$ C++ new and delete
 - virtual functions
 - $\bullet \ \ \text{inline} \ \mathsf{PTX}$
- Thrust library (similar to STL):
 - data structures: device_vector, host_vector, ...
 - ${\scriptstyle \circ }$ algorithms: sort, reduce , ...
- libraries: FFT, BLAS,
- other language bindings: pyCUDA, Copperhead, Fortran, ...
- OpenCL

Summary and outlook

This lecture

In this lecture we have covered some more advanced simulation techniques. As expected, highly (data) non-local algorithms give significantly smaller speed-ups. Still, these are sizeable.

Summarv

All lectures

Scientific computing with graphics processing units:

- tailoring to the hardware needed for really good performance
- special, but general enough to provide significant speed-ups for most problems
- GPU computing might be a fashion, but CPU computing goes the same way

Reading

For this lecture:

- M. Weigel, Phys. Rev. E 84, 036709 (2011) [arXiv:1105.5804].
- M. Weigel and T. Yavors'kii, Physics Procedia 15, 92 (2011) [arXiv:1107.5463].
- http://www.cond-mat.physik.uni-mainz.de/~weigel/GPU

advanced simulations

01/11/2012 33 / 34

Some advertisement: EPJST issue



http://epjst.epj.org/articles/epjst/abs/2012/10/contents/contents.html

M. Weigel (Coventry/Mainz)

advanced simulations

01/11/2012 34 / 34