

Solutions for Efficient Memory Access for Cubic Lattices and Random Number Algorithms



SAPIENZA
UNIVERSITÀ DI ROMA



UNIVERSITEIT TWENTE.



GTC 2015

Speaker: Dr. Matteo Lulli
Prof. M. Bernaschi and Prof. G. Parisi

March the 19th, 2015

Outlook

- 1 Cubic stencils
- 2 PRNGs
- 3 Multi-GPU and MPI
- 4 Results
- 5 Conclusions & Perspectives

Phase Transitions in disordered systems

Phase Transitions in disordered systems

- Equilibrium Monte Carlo analysis works well for non-disordered systems
- Disordered systems are very hard to equilibrate

Phase Transitions in disordered systems

- Equilibrium Monte Carlo analysis works well for non-disordered systems
- Disordered systems are very hard to equilibrate
- A large number of disorder realizations, **samples**, is required
- 3D Ising spin glass at most $L = 40$ equilibrated so far (Janus, FPGA dedicated machine)

Phase Transitions in disordered systems

- Equilibrium Monte Carlo analysis works well for non-disordered systems
- Disordered systems are very hard to equilibrate
- A large number of disorder realizations, **samples**, is required
- 3D Ising spin glass at most $L = 40$ equilibrated so far (Janus, FPGA dedicated machine)
- Robust out-of-equilibrium methods for the study of phase transitions would be very useful

Phase Transitions in disordered systems

- Equilibrium Monte Carlo analysis works well for non-disordered systems
- Disordered systems are very hard to equilibrate
- A large number of disorder realizations, **samples**, is required
- 3D Ising spin glass at most $L = 40$ equilibrated so far (Janus, FPGA dedicated machine)
- Robust out-of-equilibrium methods for the study of phase transitions would be very useful

Why GPUs

Even with out-of-equilibrium methods usual CPUs are not sufficiently powerful in order to obtain good estimates

Outline for section 1

1 Cubic stencils

2 PRNGs

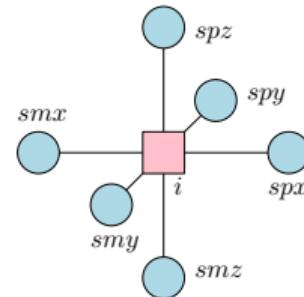
3 Multi-GPU and MPI

4 Results

5 Conclusions & Perspectives

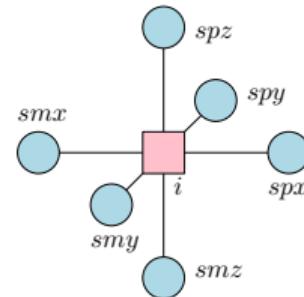
Standard checkerboard pattern

- Nearest-neighbours based problems in 3D



Standard checkerboard pattern

- Nearest-neighbours based problems in 3D
- Cubic lattice of linear size $L = 2n$: Checkerboard colouring
- Each lattice site has nearest neighbours of the other colour

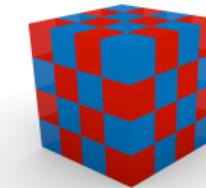
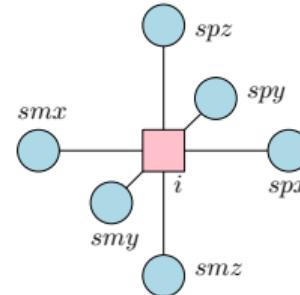


Standard checkerboard pattern

- Nearest-neighbours based problems in 3D
- Cubic lattice of linear size $L = 2n$: Checkerboard colouring
- Each lattice site has nearest neighbours of the other colour

Allocation choices

- Unified allocation: one array

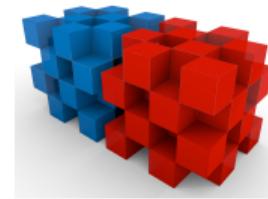
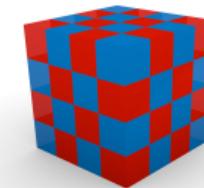
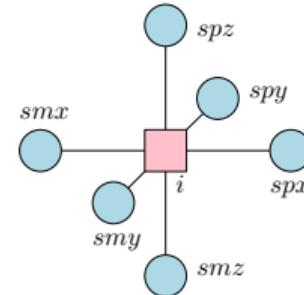


Standard checkerboard pattern

- Nearest-neighbours based problems in 3D
- Cubic lattice of linear size $L = 2n$: Checkerboard colouring
- Each lattice site has nearest neighbours of the other colour

Allocation choices

- Unified allocation: one array
- Separated allocation: two arrays

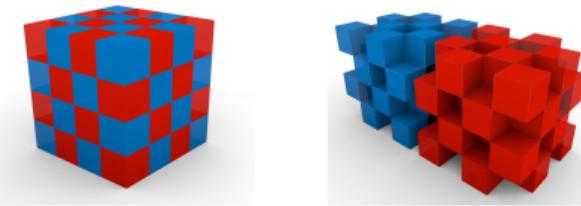
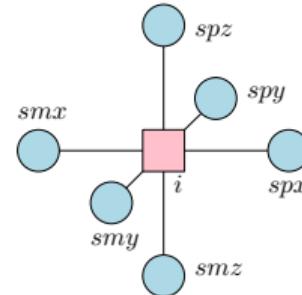


Standard checkerboard pattern

- Nearest-neighbours based problems in 3D
- Cubic lattice of linear size $L = 2n$: Checkerboard colouring
- Each lattice site has nearest neighbours of the other colour

Allocation choices

- Unified allocation: one array
- Separated allocation: two arrays

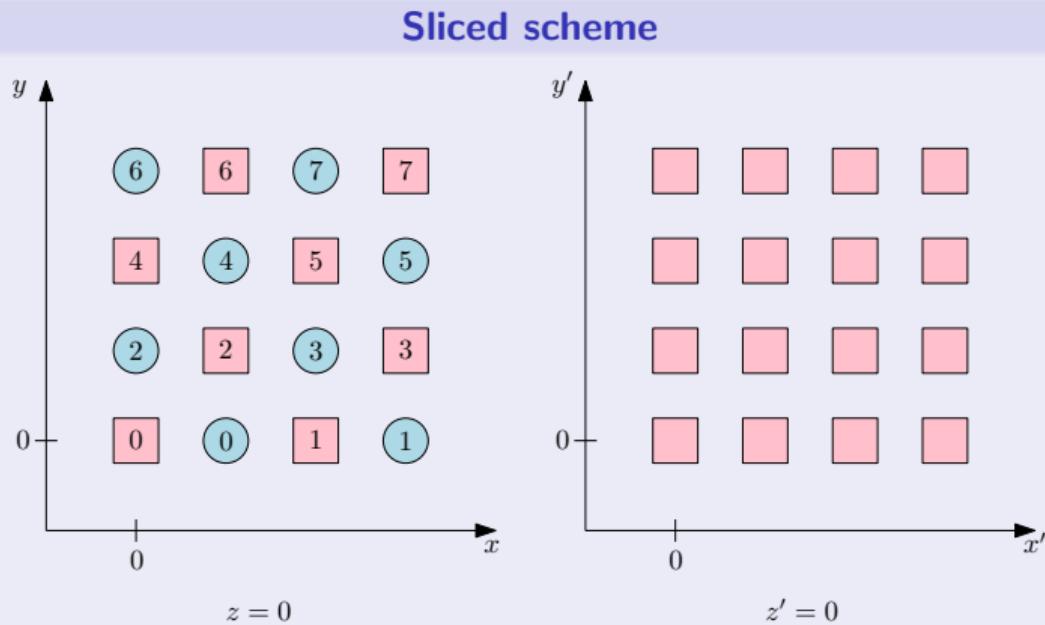


The parity, $(-1)^{x_i+y_i+z_i}$, of each lattice site has to be taken into account

Sliced checkerboard pattern: definition

It is always possible to **remap** the lattice sites

- Yavors'kii *et al.*, Heisenberg spin glass, 'snake-like' pattern, unified
- Ferrero *et al.*, Potts 2D, separated

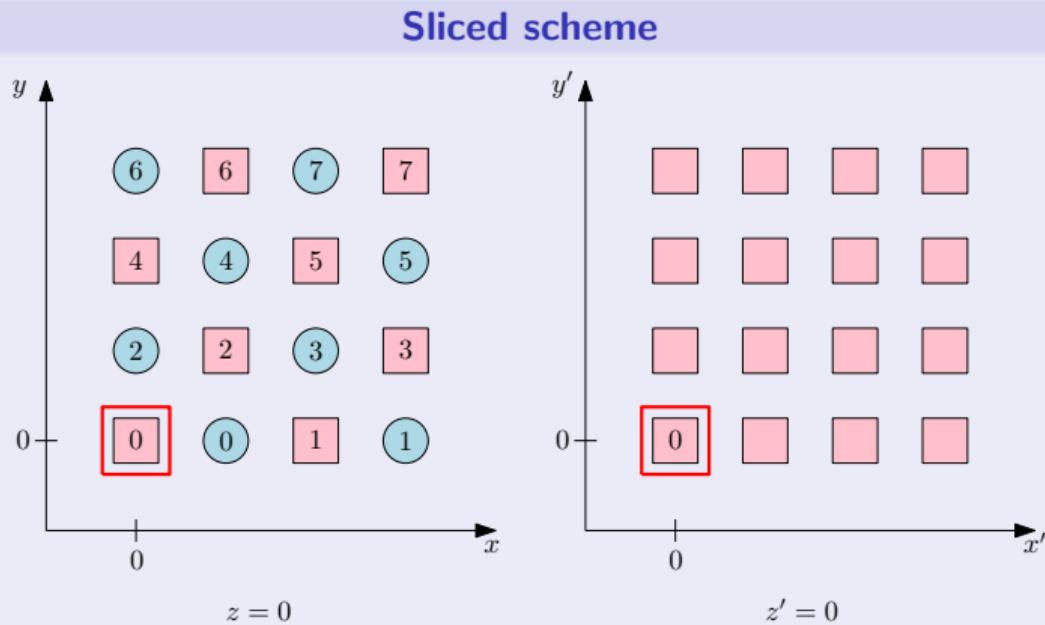


Periodic boundary conditions are necessary

Sliced checkerboard pattern: definition

It is always possible to **remap** the lattice sites

- Yavors'kii *et al.*, Heisenberg spin glass, 'snake-like' pattern, unified
- Ferrero *et al.*, Potts 2D, separated

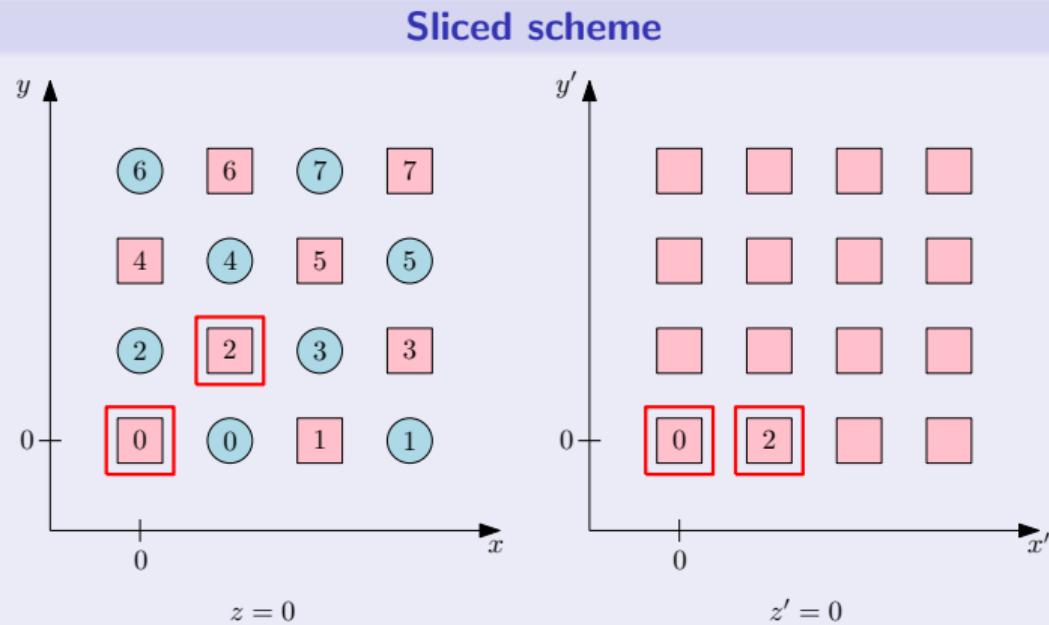


Periodic boundary conditions are necessary

Sliced checkerboard pattern: definition

It is always possible to **remap** the lattice sites

- Yavors'kii *et al.*, Heisenberg spin glass, 'snake-like' pattern, unified
- Ferrero *et al.*, Potts 2D, separated

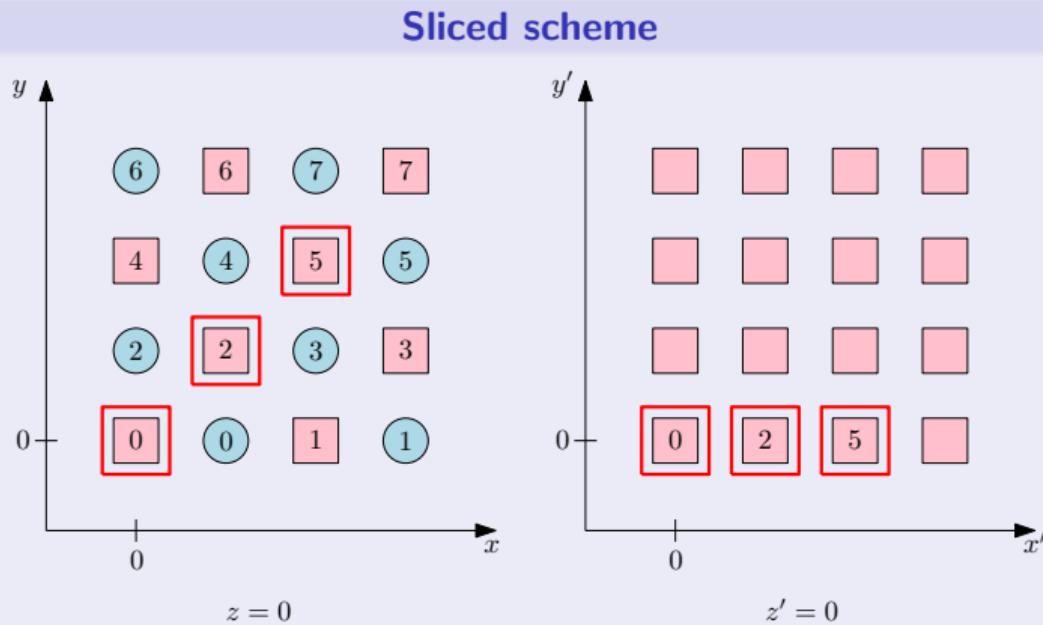


Periodic boundary conditions are necessary

Sliced checkerboard pattern: definition

It is always possible to **remap** the lattice sites

- Yavors'kii *et al.*, Heisenberg spin glass, 'snake-like' pattern, unified
- Ferrero *et al.*, Potts 2D, separated

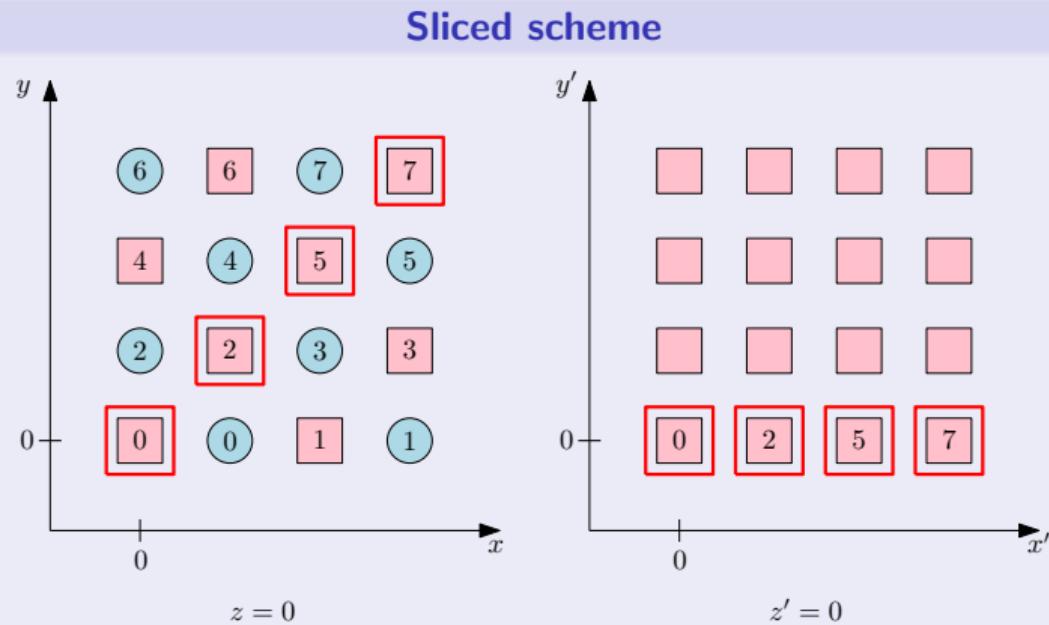


Periodic boundary conditions are necessary

Sliced checkerboard pattern: definition

It is always possible to **remap** the lattice sites

- Yavors'kii *et al.*, Heisenberg spin glass, 'snake-like' pattern, unified
- Ferrero *et al.*, Potts 2D, separated

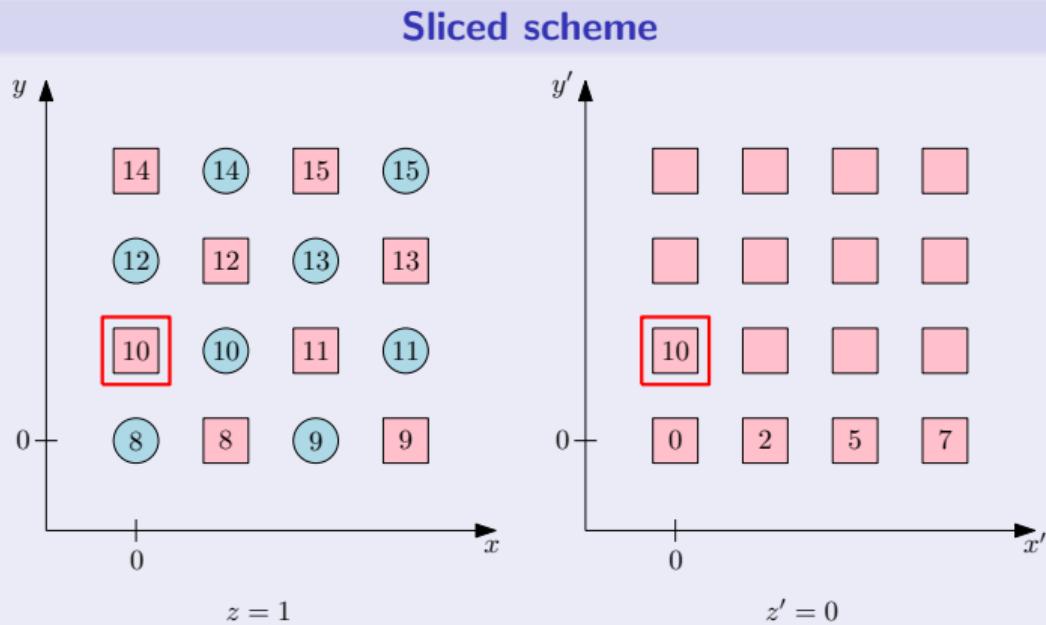


Periodic boundary conditions are necessary

Sliced checkerboard pattern: definition

It is always possible to **remap** the lattice sites

- Yavors'kii *et al.*, Heisenberg spin glass, 'snake-like' pattern, unified
- Ferrero *et al.*, Potts 2D, separated

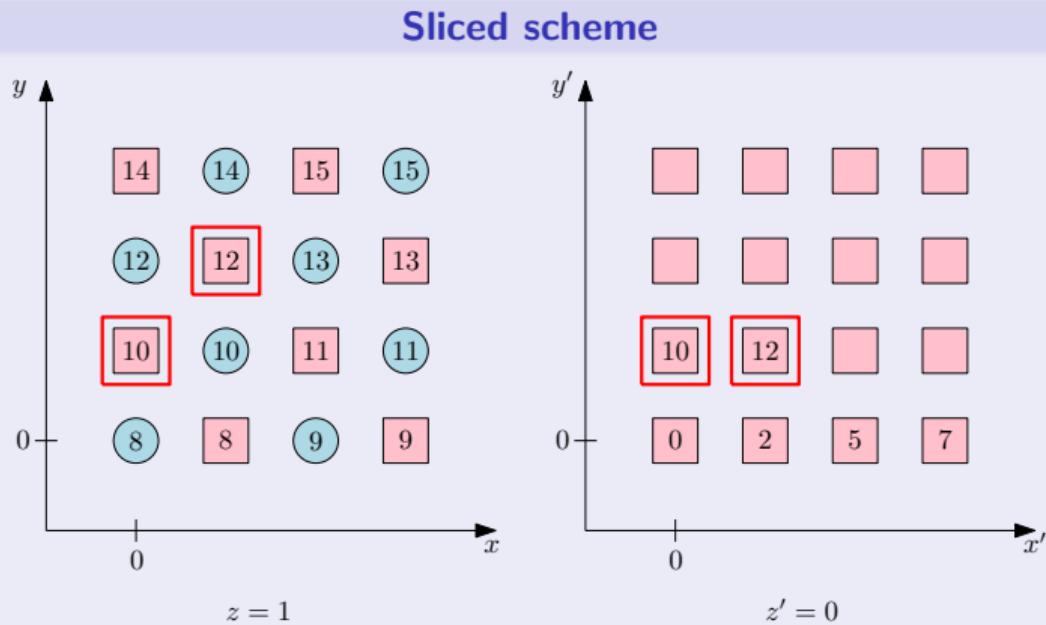


Periodic boundary conditions are necessary

Sliced checkerboard pattern: definition

It is always possible to **remap** the lattice sites

- Yavors'kii *et al.*, Heisenberg spin glass, 'snake-like' pattern, unified
- Ferrero *et al.*, Potts 2D, separated

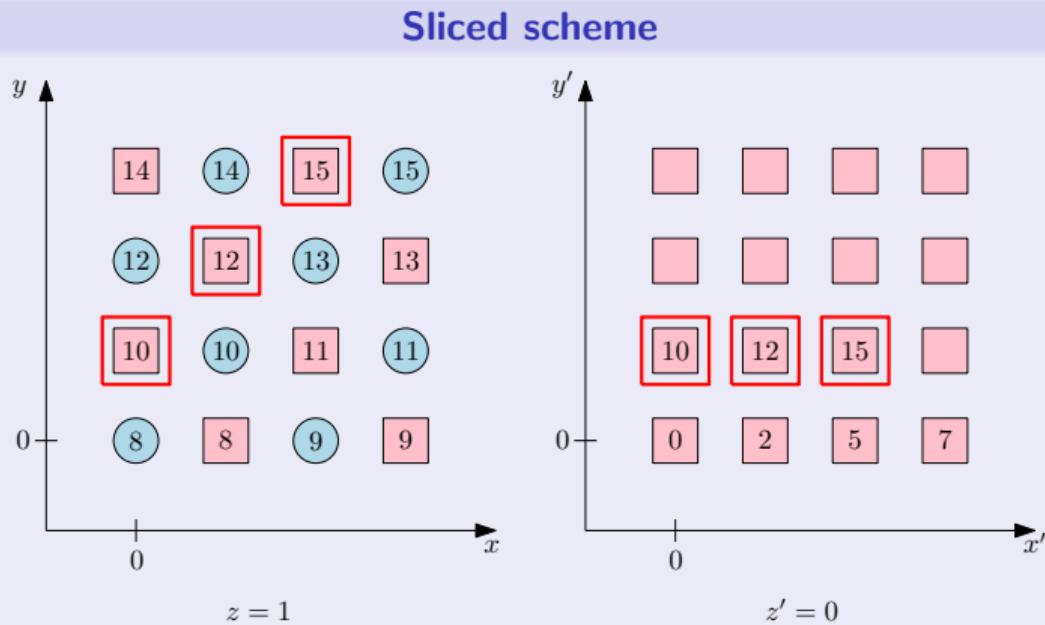


Periodic boundary conditions are necessary

Sliced checkerboard pattern: definition

It is always possible to **remap** the lattice sites

- Yavors'kii *et al.*, Heisenberg spin glass, 'snake-like' pattern, unified
- Ferrero *et al.*, Potts 2D, separated

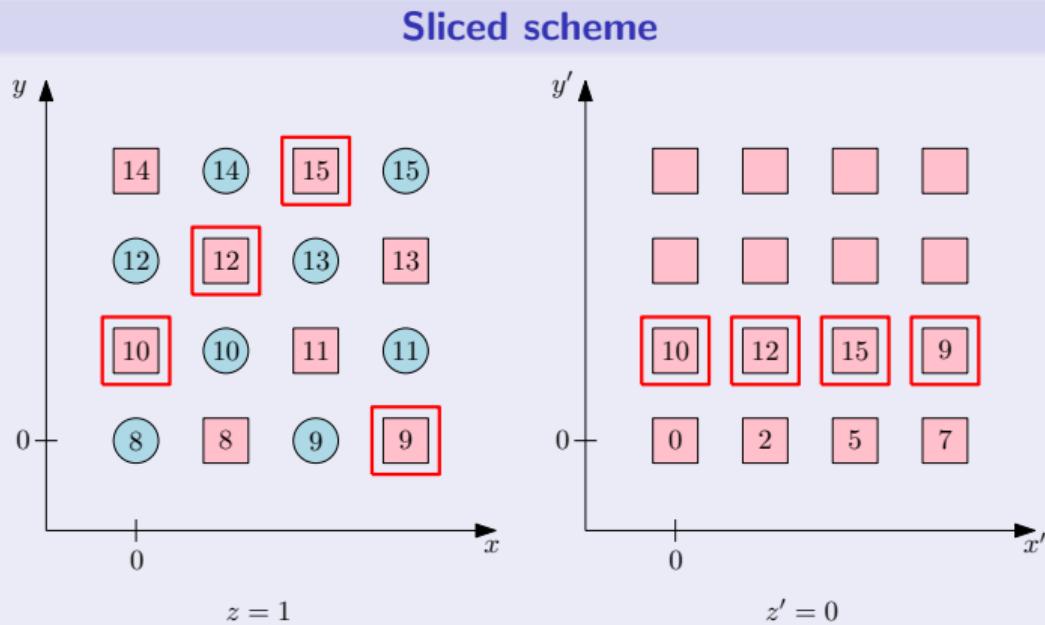


Periodic boundary conditions are necessary

Sliced checkerboard pattern: definition

It is always possible to **remap** the lattice sites

- Yavors'kii *et al.*, Heisenberg spin glass, 'snake-like' pattern, unified
- Ferrero *et al.*, Potts 2D, separated

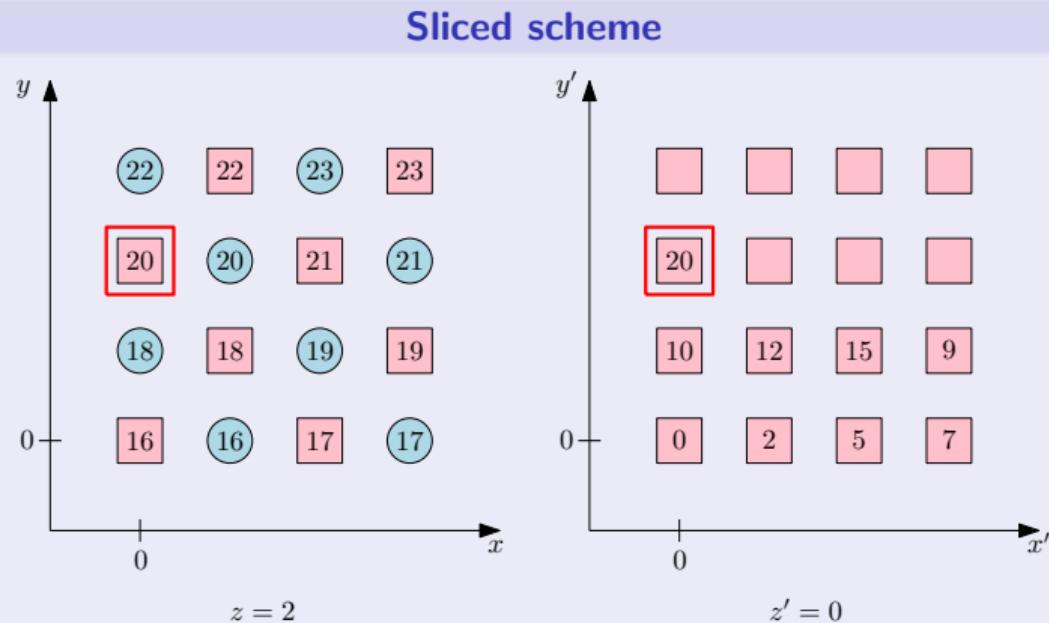


Periodic boundary conditions are necessary

Sliced checkerboard pattern: definition

It is always possible to **remap** the lattice sites

- Yavors'kii *et al.*, Heisenberg spin glass, 'snake-like' pattern, unified
- Ferrero *et al.*, Potts 2D, separated

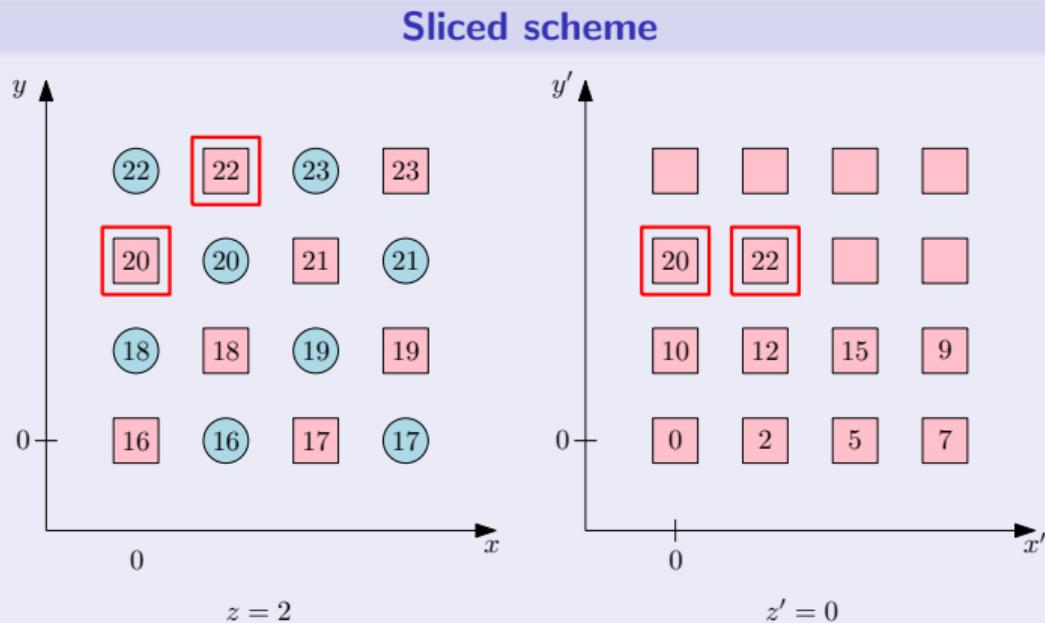


Periodic boundary conditions are necessary

Sliced checkerboard pattern: definition

It is always possible to **remap** the lattice sites

- Yavors'kii *et al.*, Heisenberg spin glass, 'snake-like' pattern, unified
- Ferrero *et al.*, Potts 2D, separated

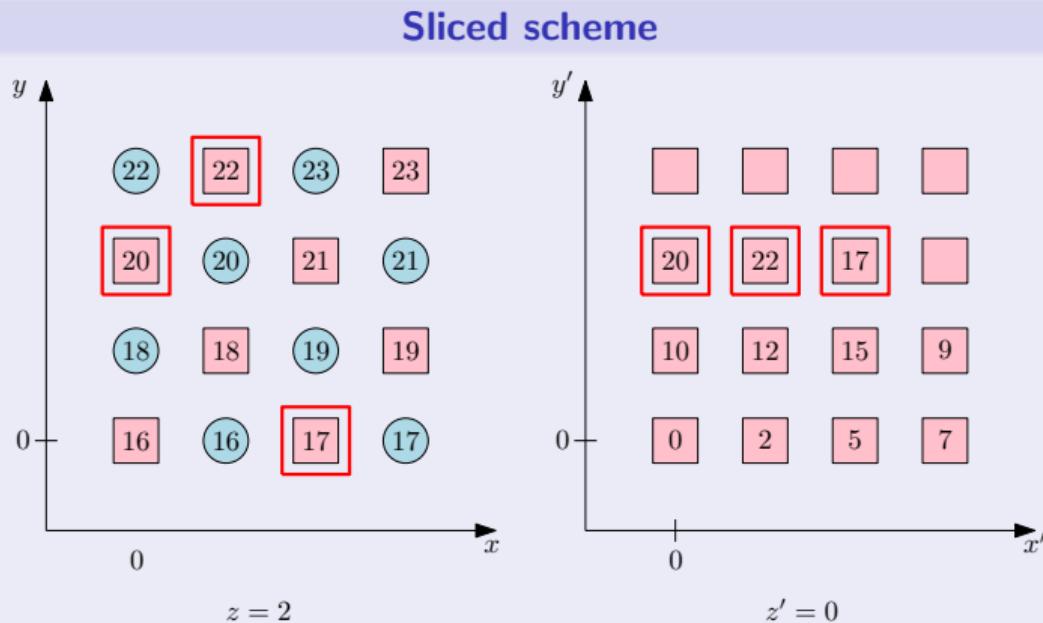


Periodic boundary conditions are necessary

Sliced checkerboard pattern: definition

It is always possible to **remap** the lattice sites

- Yavors'kii *et al.*, Heisenberg spin glass, 'snake-like' pattern, unified
- Ferrero *et al.*, Potts 2D, separated

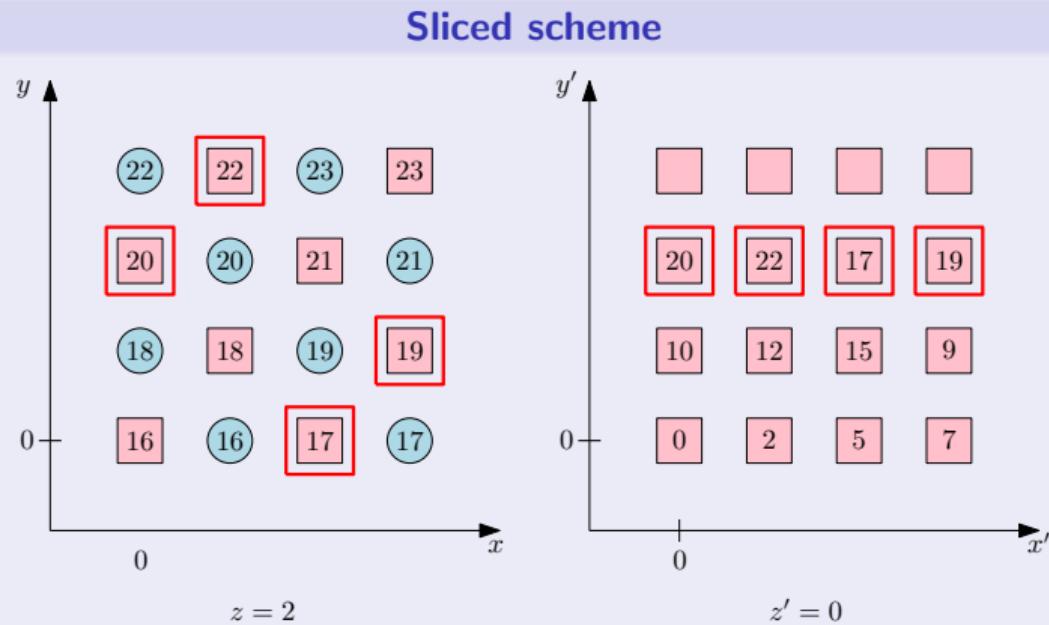


Periodic boundary conditions are necessary

Sliced checkerboard pattern: definition

It is always possible to **remap** the lattice sites

- Yavors'kii *et al.*, Heisenberg spin glass, 'snake-like' pattern, unified
- Ferrero *et al.*, Potts 2D, separated

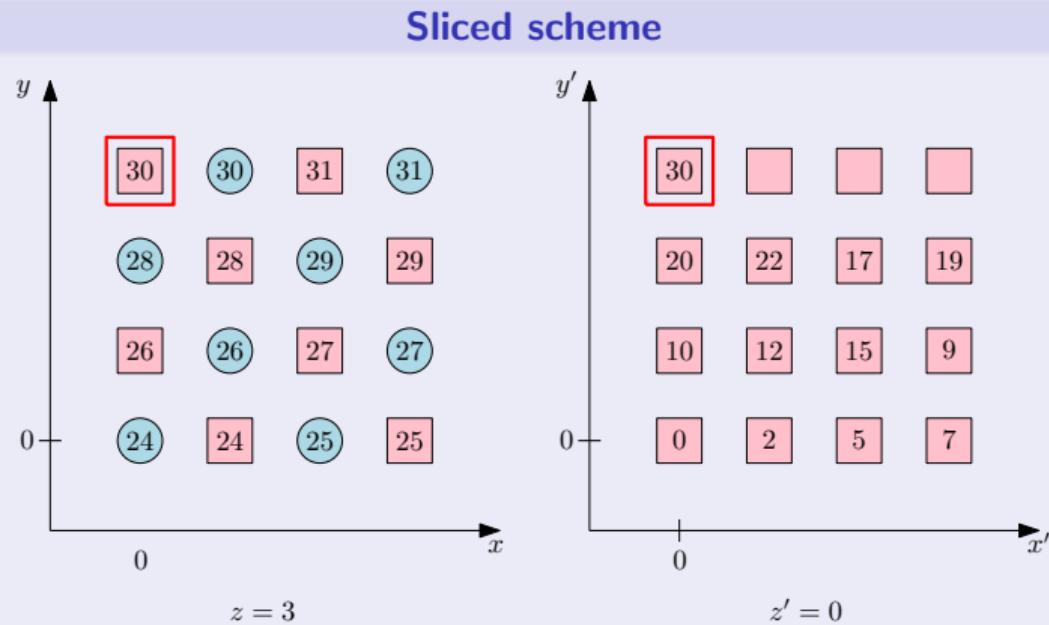


Periodic boundary conditions are necessary

Sliced checkerboard pattern: definition

It is always possible to **remap** the lattice sites

- Yavors'kii *et al.*, Heisenberg spin glass, 'snake-like' pattern, unified
- Ferrero *et al.*, Potts 2D, separated

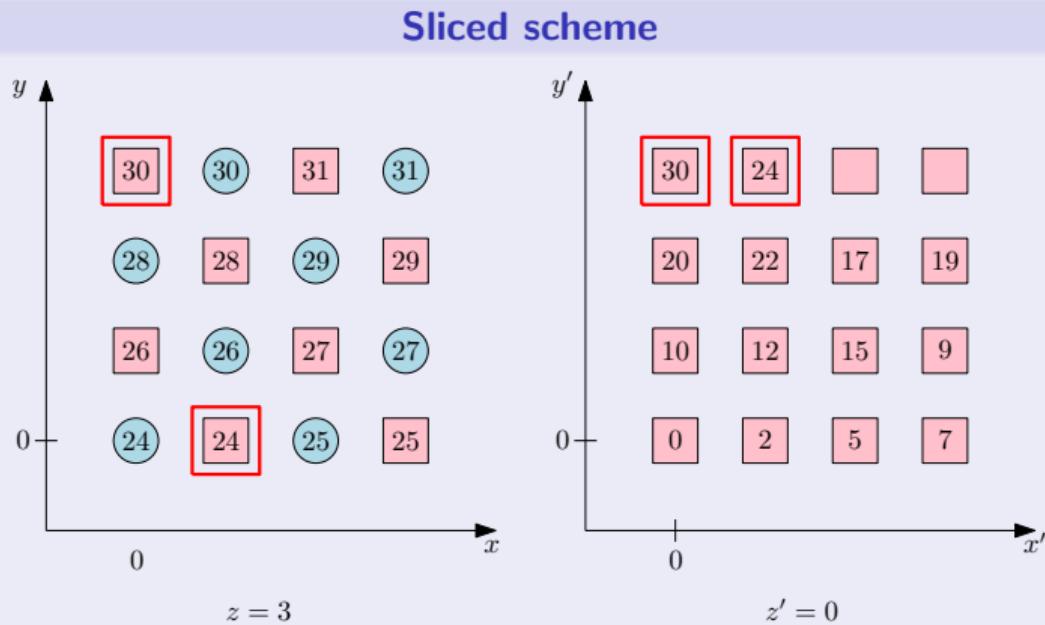


Periodic boundary conditions are necessary

Sliced checkerboard pattern: definition

It is always possible to **remap** the lattice sites

- Yavors'kii *et al.*, Heisenberg spin glass, 'snake-like' pattern, unified
- Ferrero *et al.*, Potts 2D, separated

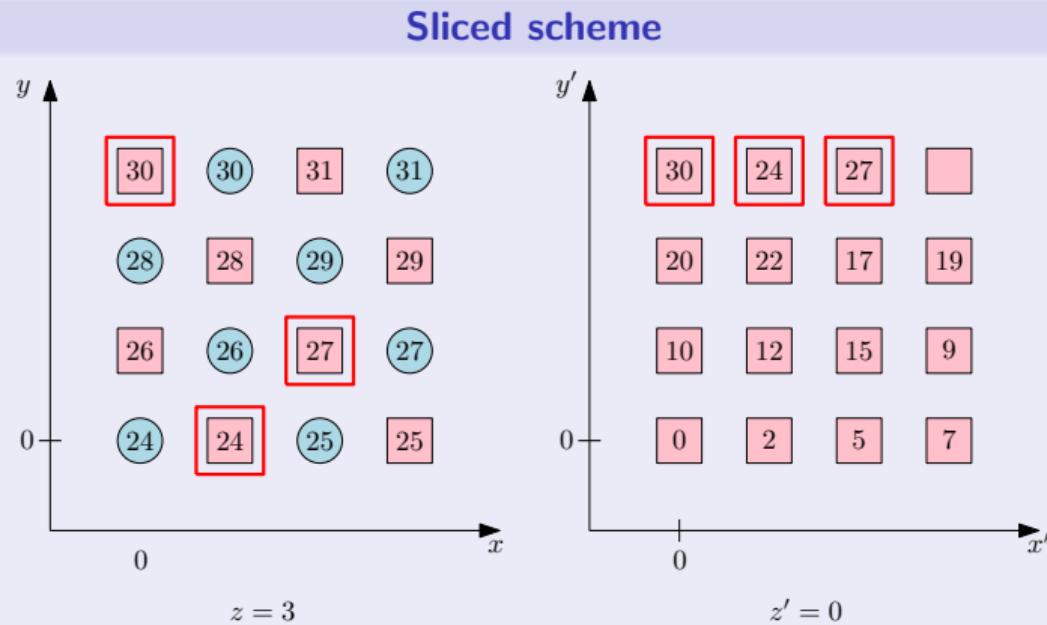


Periodic boundary conditions are necessary

Sliced checkerboard pattern: definition

It is always possible to **remap** the lattice sites

- Yavors'kii *et al.*, Heisenberg spin glass, 'snake-like' pattern, unified
- Ferrero *et al.*, Potts 2D, separated

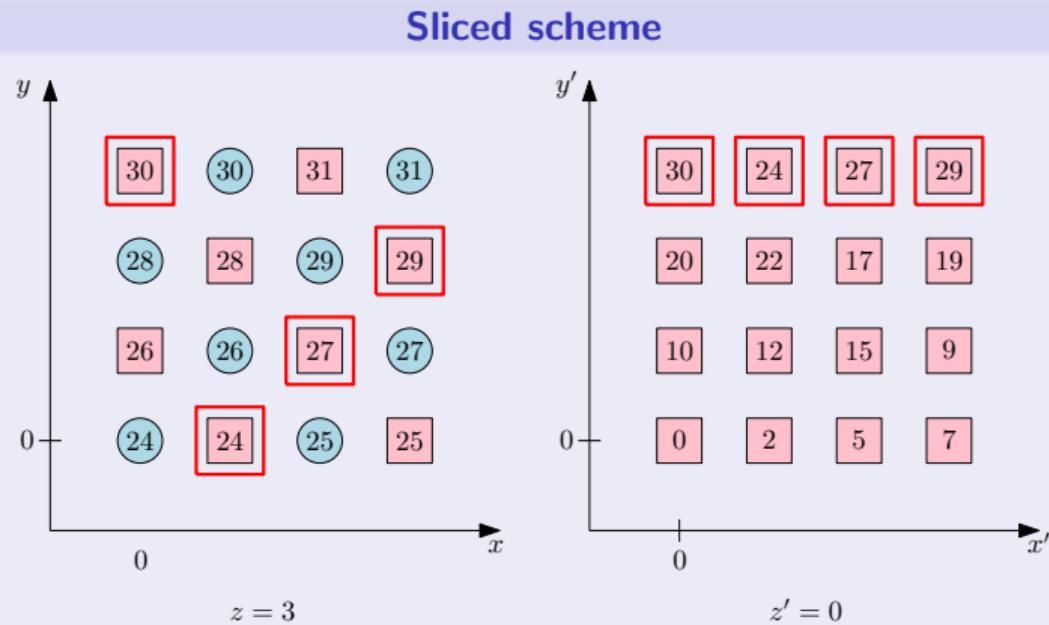


Periodic boundary conditions are necessary

Sliced checkerboard pattern: definition

It is always possible to **remap** the lattice sites

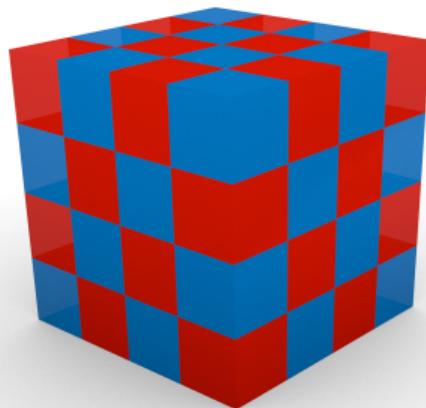
- Yavors'kii *et al.*, Heisenberg spin glass, 'snake-like' pattern, unified
- Ferrero *et al.*, Potts 2D, separated



Periodic boundary conditions are necessary

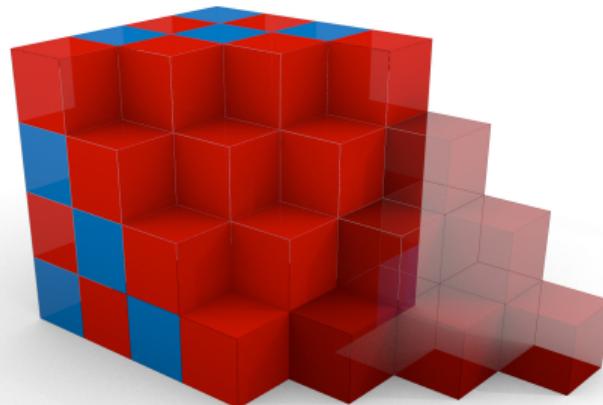
Sliced checkerboard pattern: geometric interpretation

Sites of planes orthogonal to $\vec{n} = (+1, -1, +1)$ are one-coloured.



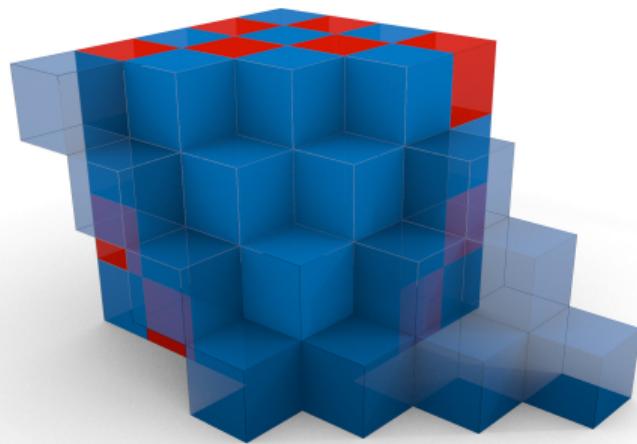
Sliced checkerboard pattern: geometric interpretation

Sites of planes orthogonal to $\vec{n} = (+1, -1, +1)$ are one-coloured.



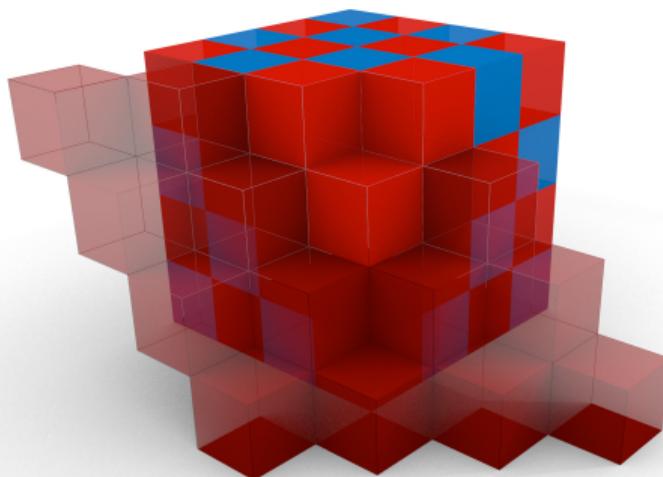
Sliced checkerboard pattern: geometric interpretation

Sites of planes orthogonal to $\vec{n} = (+1, -1, +1)$ are one-coloured.



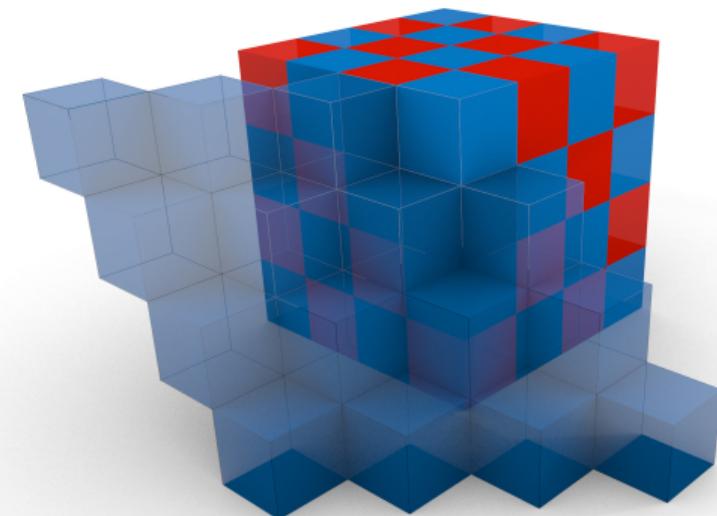
Sliced checkerboard pattern: geometric interpretation

Sites of planes orthogonal to $\vec{n} = (+1, -1, +1)$ are one-coloured.



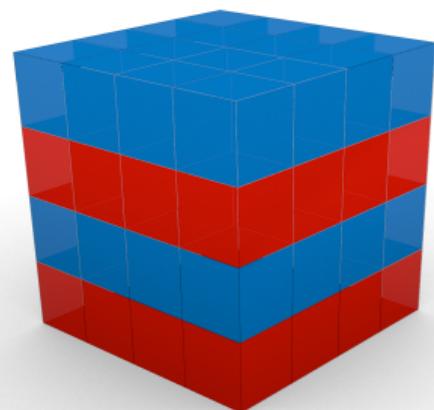
Sliced checkerboard pattern: geometric interpretation

Sites of planes orthogonal to $\vec{n} = (+1, -1, +1)$ are one-coloured.



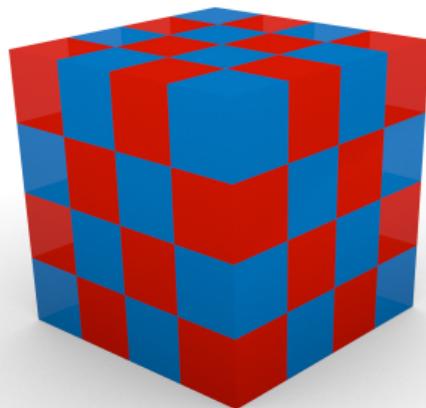
Sliced checkerboard pattern: geometric interpretation

Sites of planes orthogonal to $\vec{n} = (+1, -1, +1)$ are one-coloured.

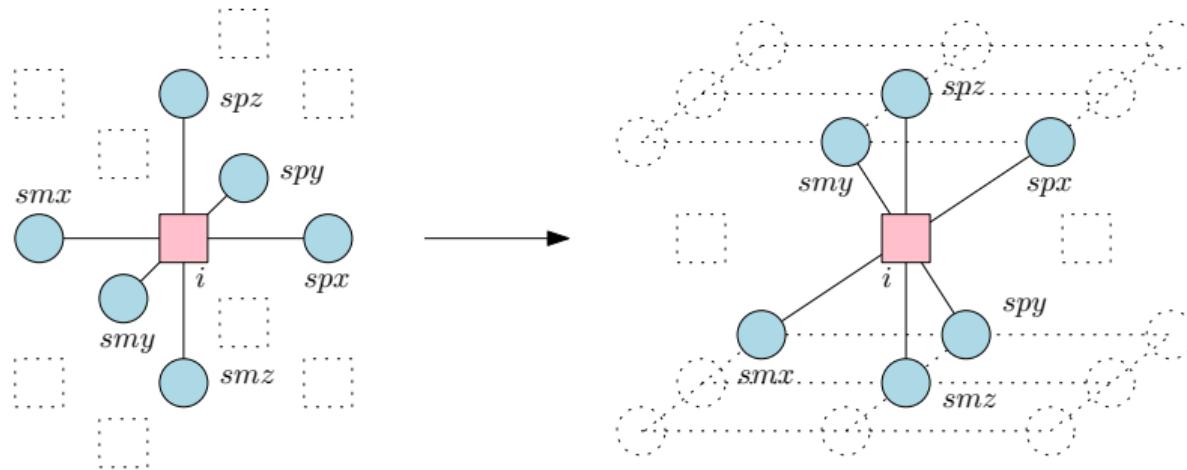


Sliced checkerboard pattern: geometric interpretation

Sites of planes orthogonal to $\vec{n} = (+1, -1, +1)$ are one-coloured.

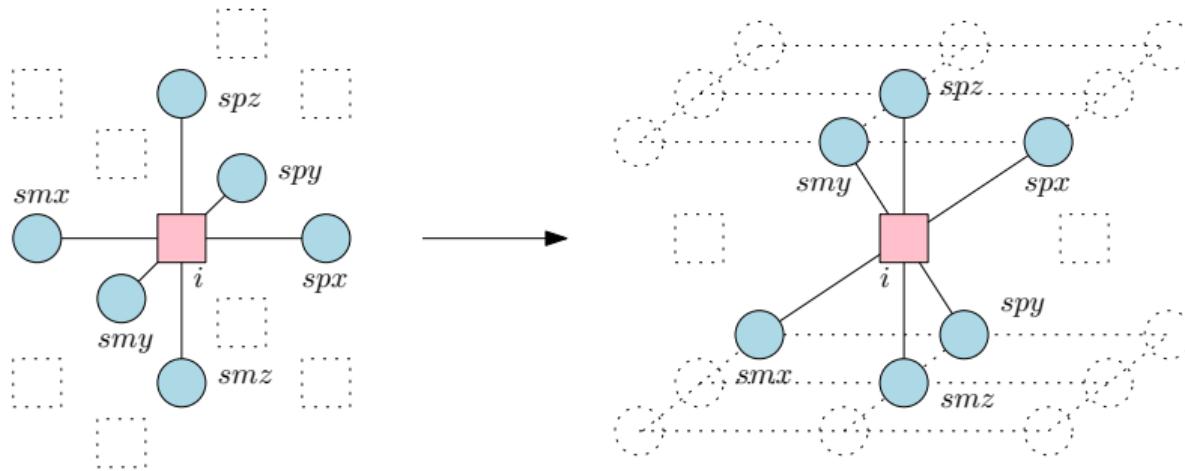


Sliced checkerboard pattern: cubic stencils



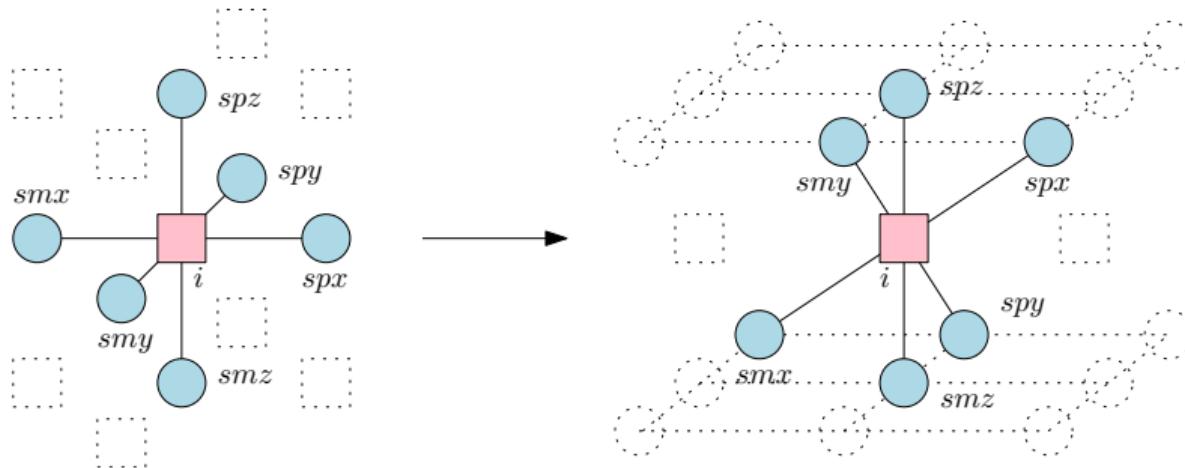
- Variables belonging to one slice are **decoupled**
- No sites parity involved

Sliced checkerboard pattern: cubic stencils



- Variables belonging to one slice are **decoupled**
- No sites parity involved
- The spz , smy , spx nearest neighbours belong to the **upper** slice
- The smz , spy , smx nearest neighbours belong to the **lower** slice

Sliced checkerboard pattern: cubic stencils



- Variables belonging to one slice are **decoupled**
- No sites parity involved
- The spz , smy , spx nearest neighbours belong to the **upper** slice
- The smz , spy , smx nearest neighbours belong to the **lower** slice
- The method can be generalized to any dimension

Outline for section 2

1 Cubic stencils

2 PRNGs

3 Multi-GPU and MPI

4 Results

5 Conclusions & Perspectives

Linear Congruential PRNG

- Lehmer-Park-Miller MINSTD

$$R_{n+1} = (16807 R_n) \bmod (2^{31} - 1),$$

- D. Carta implementation: no modulus and overflow handled with 32 bit integers only

MINSTD overflow handling

No 64 bits variables and no 'mod' involved

```
RNGT lo = 16807*(seed&0xffff);  
RNGT hi = 16807*(seed>>16);  
lo += (hi&0x7fff)<<16;  
lo += hi>>15;  
if(lo > 0x7fffffff) lo -= 0x7fffffff;  
return lo;
```

Linear Congruential PRNG

- Lehmer-Park-Miller MINSTD

$$R_{n+1} = (16807 R_n) \bmod (2^{31} - 1),$$

- D. Carta implementation: no modulus and overflow handled with 32 bit integers only
- Easy implementation: one word state
- Coalescence is obtained by definition

R[i]	0	1	2	3	4	5	...
	↑	↑	↑	↑	↑	↑	

MINSTD overflow handling

No 64 bits variables and no 'mod' involved

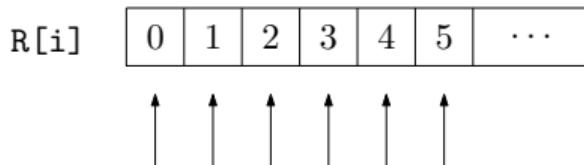
```
RNGT lo = 16807*(seed&0xffff);  
RNGT hi = 16807*(seed>>16);  
lo += (hi&0x7fff)<<16;  
lo += hi>>15;  
if(lo > 0x7fffffff) lo -= 0x7fffffff;  
return lo;
```

Linear Congruential PRNG

- Lehmer-Park-Miller MINSTD

$$R_{n+1} = (16807 R_n) \bmod (2^{31} - 1),$$

- D. Carta implementation: no modulus and overflow handled with 32 bit integers only
- Easy implementation: one word state
- Coalescence is obtained by definition



- More than one instance: read once, store once
- Low quality random numbers, but still useful in some cases

MINSTD overflow handling

No 64 bits variables and no 'mod' involved

```
RNGT lo = 16807*(seed&0xffff);  
RNGT hi = 16807*(seed>>16);  
lo += (hi&0x7fff)<<16;  
lo += hi>>15;  
if(lo > 0x7fffffff) lo -= 0x7fffffff;  
return lo;
```

Linear Congruential PRNG

- Lehmer-Park-Miller MINSTD

$$R_{n+1} = (16807 R_n) \bmod (2^{31} - 1),$$

- D. Carta implementation: no modulus and overflow handled with 32 bit integers only
- Easy implementation: one word state
- Coalescence is obtained by definition

R[i]	0	1	2	3	4	5	...
	↑	↑	↑	↑	↑	↑	

- More than one instance: read once, store once
- Low quality random numbers, but still useful in some cases

MINSTD overflow handling

No 64 bits variables and no 'mod' involved

```
RNGT lo = 16807*(seed&0xffff);  
RNGT hi = 16807*(seed>>16);  
lo += (hi&0x7fff)<<16;  
lo += hi>>15;  
if(lo > 0x7fffffff) lo -= 0x7fffffff;  
return lo;
```

Avoiding 'if' statement

We need to avoid branching

```
lo -= (((lo&0x80000000)>>31))&0x7fffffff;
```

Lagged-Fibonacci-like PRNGs: Parisi-Rapuano

Definition

- Modified lagged-Fibonacci PRNG

```
ira[i] = ira[i - 24] + ira[i - 55];  
R = ira[i]^ira[i - 61];
```

- At least 62 entries state

Lagged-Fibonacci-like PRNGs: Parisi-Rapuano

Definition

- Modified lagged-Fibonacci PRNG

```
ira[i] = ira[i - 24] + ira[i - 55];  
R = ira[i]^ira[i - 61];
```

- At least 62 entries state

GPU implementation

- Common practice: load one or more states in Shared Memory. Lags can be used for threads to work together.
- However, lags not suitable (Weigel 2012)
- Good trade-off between efficiency and quality

Lagged-Fibonacci-like PRNGs: Parisi-Rapuano

Definition

- Modified lagged-Fibonacci PRNG

```
ira[i] = ira[i - 24] + ira[i - 55];
R = ira[i]^ira[i - 61];
```

- At least 62 entries state

GPU implementation

- Common practice: load one or more states in Shared Memory. Lags can be used for threads to work together.
- However, lags not suitable (Weigel 2012)
- Good trade-off between efficiency and quality

Lagged-Fibonacci-like PRNGs: new memory access

- One state per thread with coalescing

```
seed = ira[(i - 24)*threads + threadId]
      + ira[(i - 55)*threads + threadId];
```

```
ira[i*threads + threadId] = seed;
```

```
seed ^= ira[(i - 61)*threads + threadId];
```

- The method is general and it can be applied to the well-known Mersenne-Twister (MT19937)
- Huge memory consumption $\propto N_{\text{threads}} \times N_{\text{state}}$
- Parisi-Rapuano: $N_{PR} = 62$
- Mersenne Twister MT19937 $N_{MT} = 624$

Benchmarks: host API

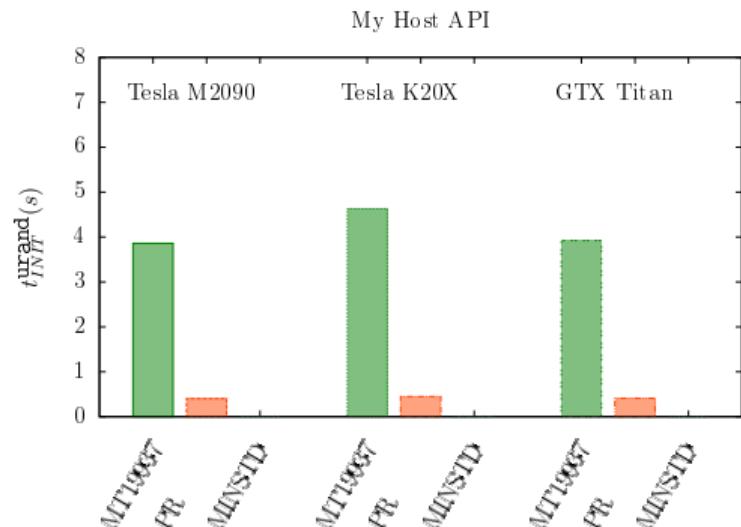
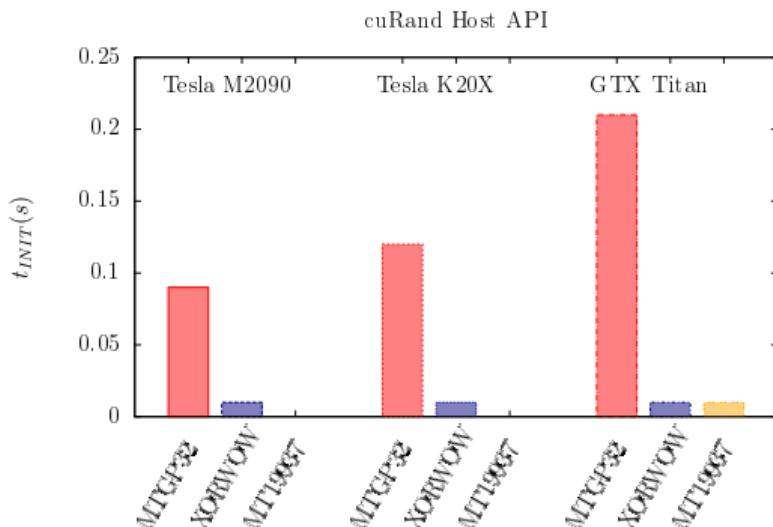
- PRAND benchmark results using cuRand host API
- Filling an array of 2^{29} single-precision floating point variables

Benchmarks: host API

- PRAND benchmark results using cuRand host API
- Filling an array of 2^{29} single-precision floating point variables
- t_{INIT} : time needed to initialize the PRNG
- t_{INIT}^{urand} : time needed to initialize the PRNG using /dev/urandom
- t_{GEN} : is the generation time

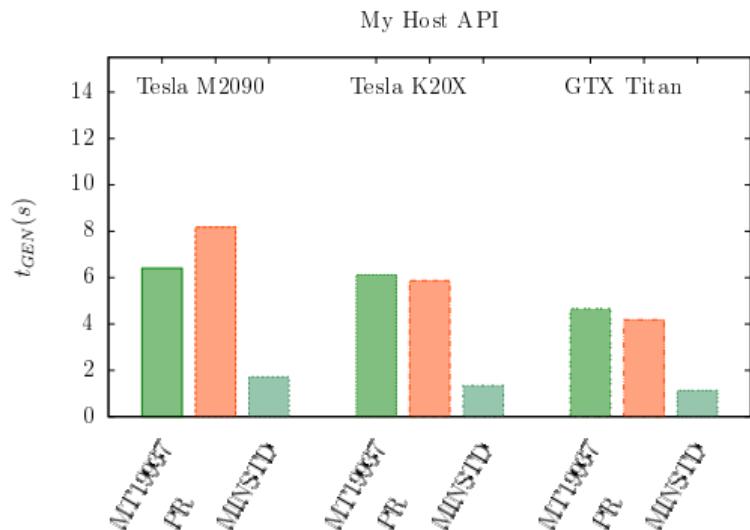
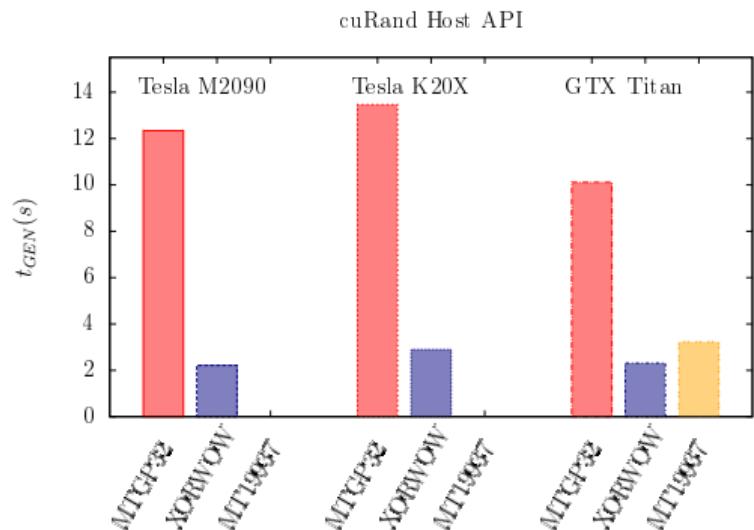
Benchmarks: host API

- PRAND benchmark results using cuRand host API
- Filling an array of 2^{29} single-precision floating point variables
- t_{INIT} : time needed to initialize the PRNG
- t_{INIT}^{urand} : time needed to initialize the PRNG using `/dev/urandom`
- t_{GEN} : is the generation time



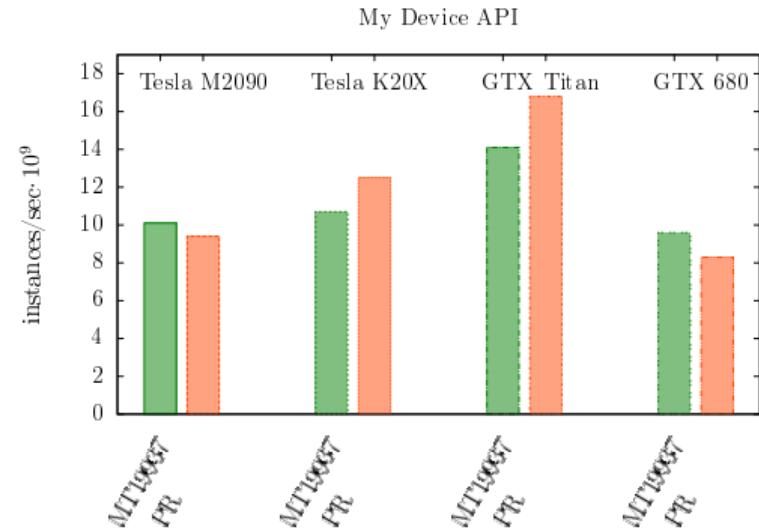
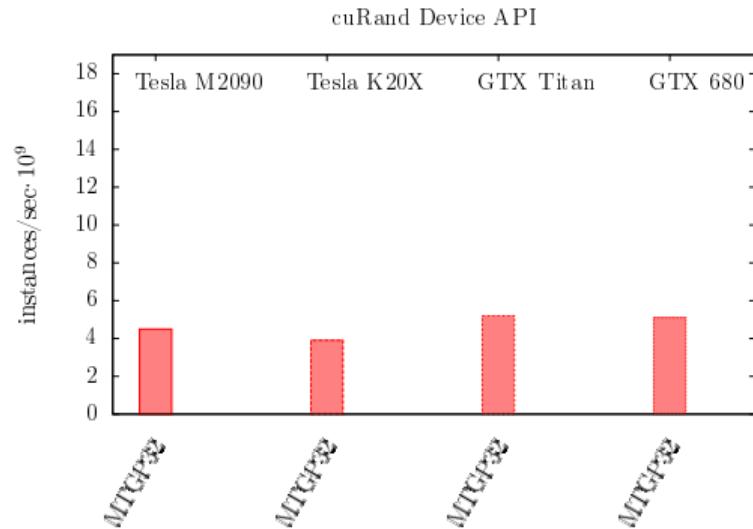
Benchmarks: host API

- PRAND benchmark results using cuRand host API
- Filling an array of 2^{29} single-precision floating point variables
- t_{INIT} : time needed to initialize the PRNG
- t_{INIT}^{urand} : time needed to initialize the PRNG using `/dev/urandom`
- t_{GEN} : is the generation time



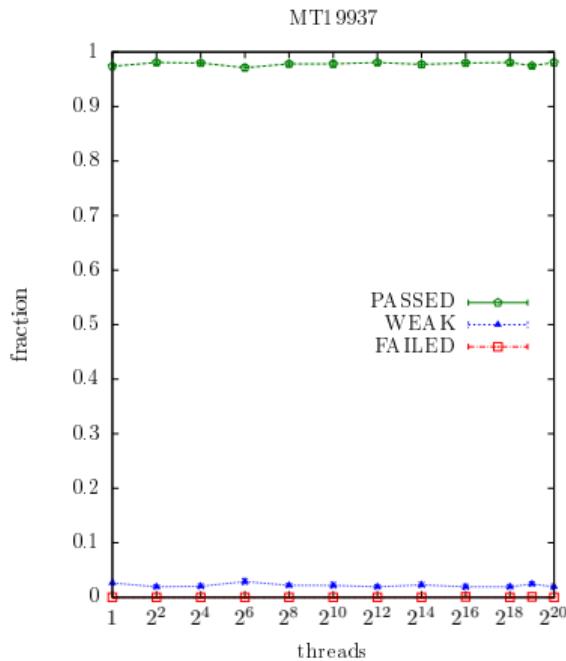
Benchmarks: device API

- Number of instances per second
- Launch configuration: 64 blocks of 256 threads
- Each thread produces 2^{15} instances, repeated 10 times



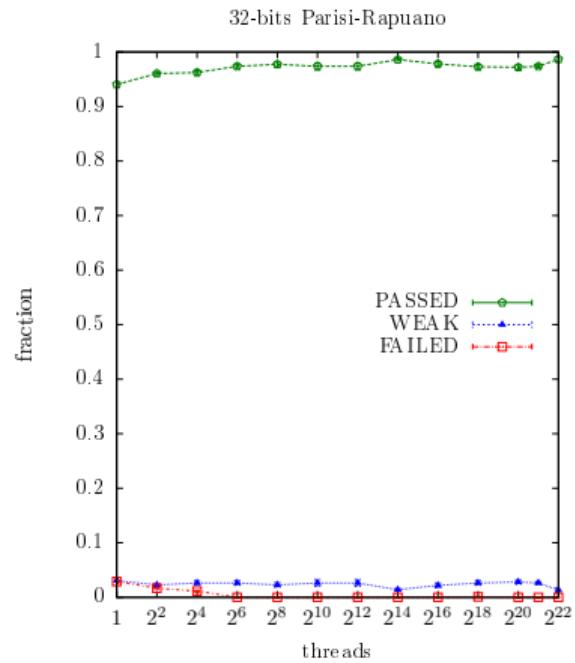
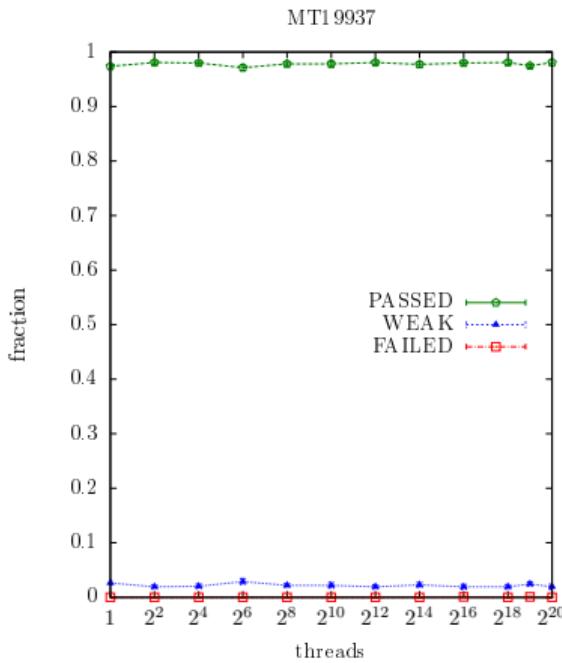
Dieharder Tests

- Let us test the generators initialized via `/dev/urandom` against **Dieharder**



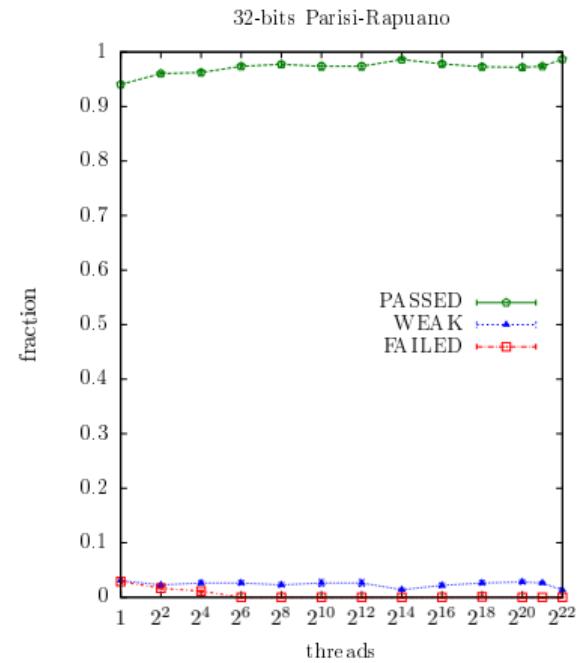
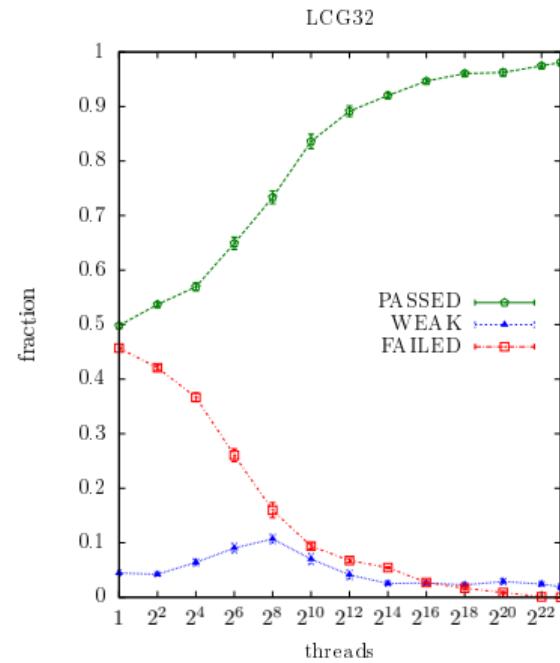
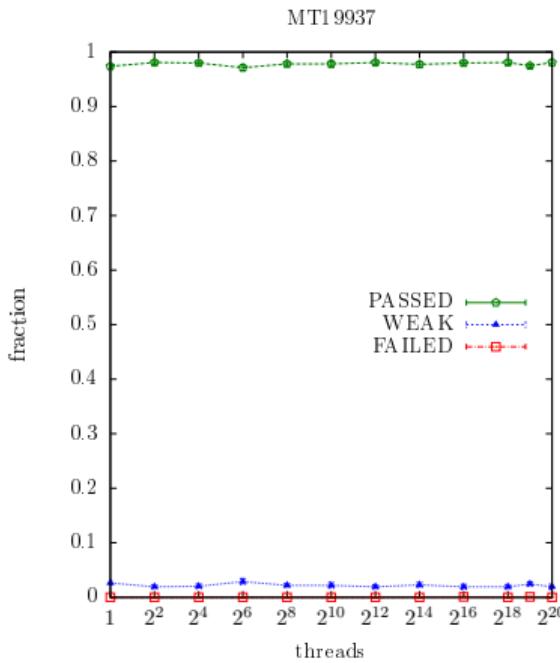
Dieharder Tests

- Let us test the generators initialized via `/dev/urandom` against **Dieharder**



Dieharder Tests

- Let us test the generators initialized via `/dev/urandom` against Dieharder



Outline for section 3

1 Cubic stencils

2 PRNGs

3 Multi-GPU and MPI

4 Results

5 Conclusions & Perspectives

CUDA streams

- Separated update: all **reds** and then all **blues**
- Communication handled by MPI
- z direction slicing of the system: bulk + boundary

CUDA streams

- Separated update: all **reds** and then all **blues**
- Communication handled by MPI
- z direction slicing of the system: bulk + boundary

Strategy

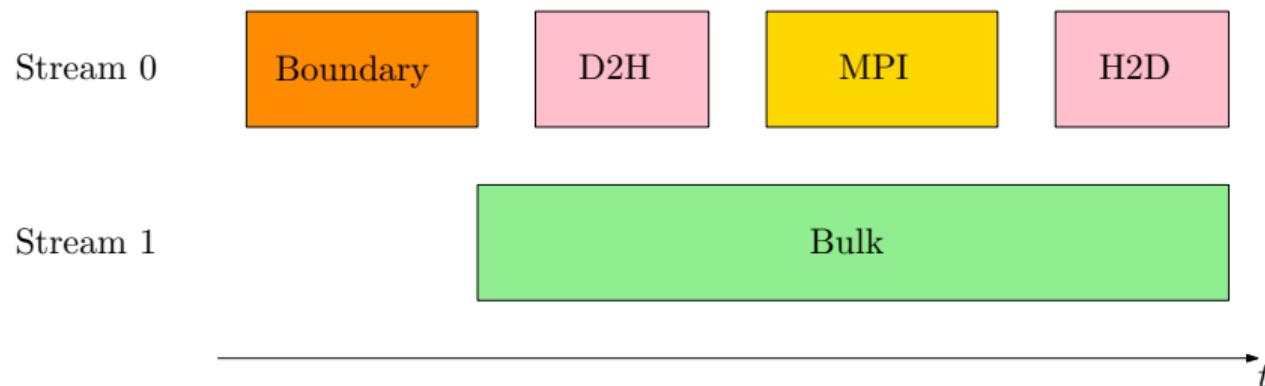
- Overlapping the update of the boundaries and their transfer with the update of the bulk
- One stream each

CUDA streams

- Separated update: all **reds** and then all **blues**
- Communication handled by MPI
- z direction slicing of the system: bulk + boundary

Strategy

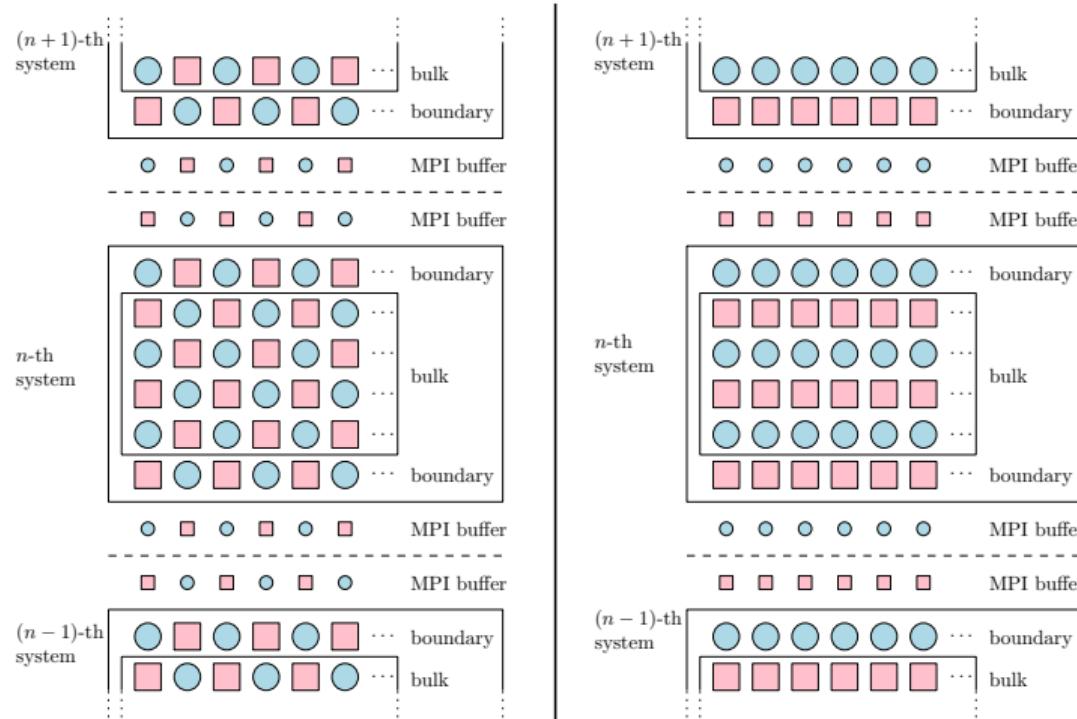
- Overlapping the update of the boundaries and their transfer with the update of the bulk
- One stream each



Sliced checkerboard and MPI

- Standard checkerboard boundaries are **two-coloured**
- Sliced checkerboard boundaries are **one-coloured**

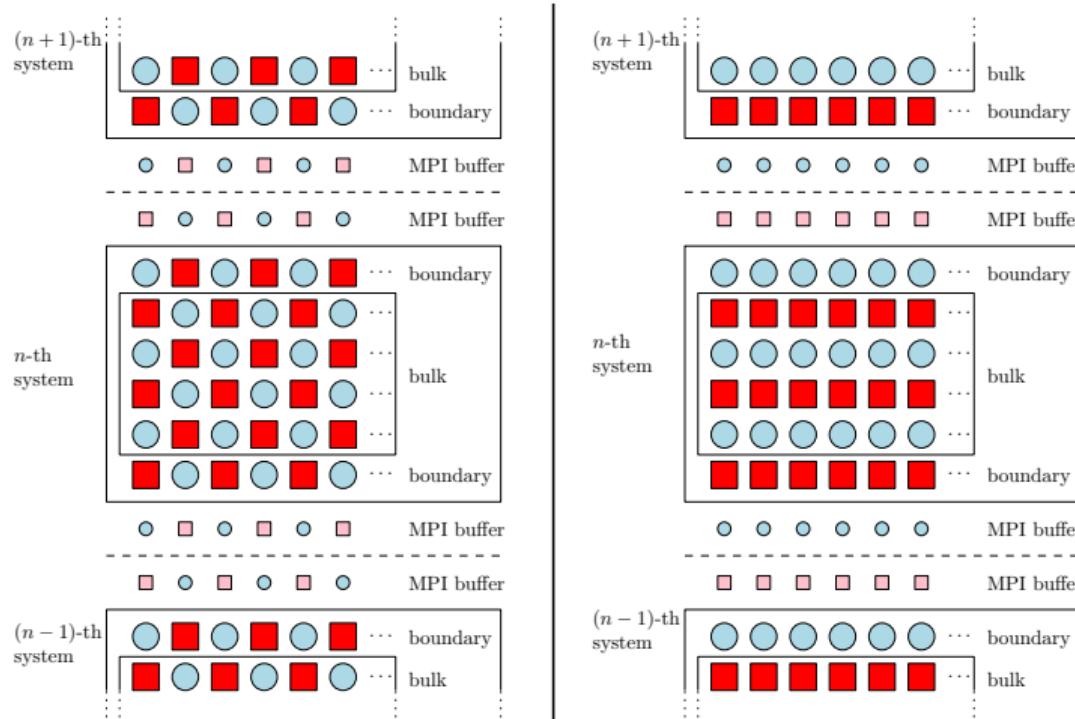
Sliced scheme: communication between nodes is **one-directional only**



Sliced checkerboard and MPI

- Standard checkerboard boundaries are **two-coloured**
- Sliced checkerboard boundaries are **one-coloured**

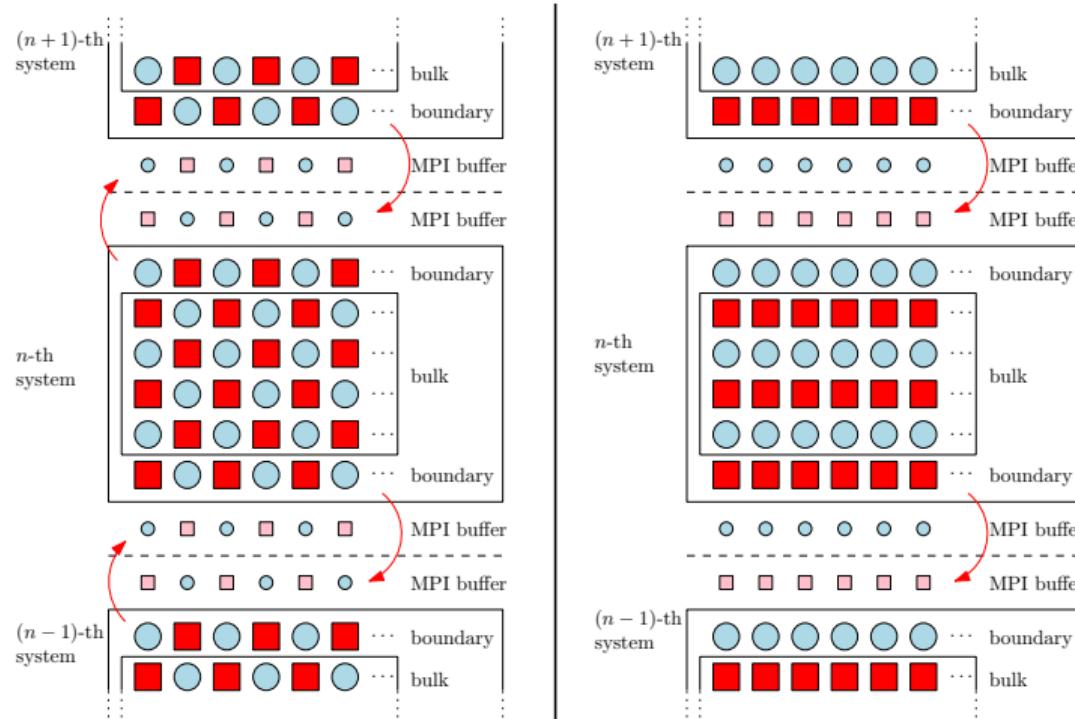
Sliced scheme: communication between nodes is **one-directional only**



Sliced checkerboard and MPI

- Standard checkerboard boundaries are **two-coloured**
- Sliced checkerboard boundaries are **one-coloured**

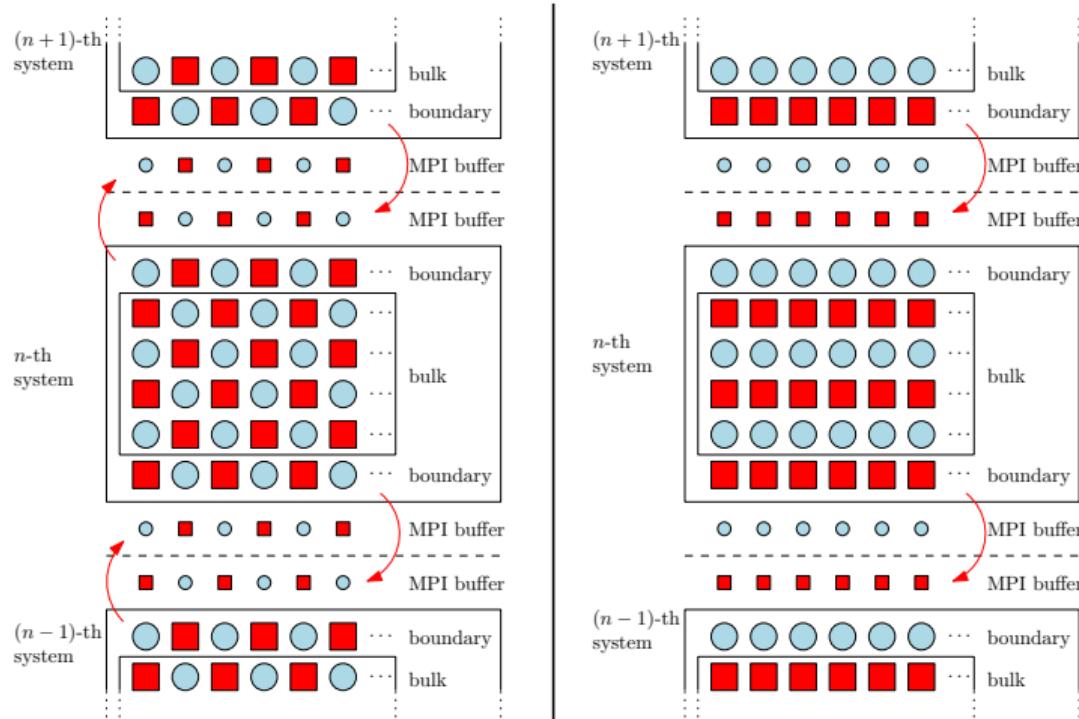
Sliced scheme: communication between nodes is **one-directional only**



Sliced checkerboard and MPI

- Standard checkerboard boundaries are **two-coloured**
- Sliced checkerboard boundaries are **one-coloured**

Sliced scheme: communication between nodes is **one-directional only**



Outline for section 4

1 Cubic stencils

2 PRNGs

3 Multi-GPU and MPI

4 Results

5 Conclusions & Perspectives

Three-dimensional Edwards-Anderson model

- Three-dimensional Spin Glass model

$$H = - \sum_{\langle ik \rangle} J_{ik} \sigma_i \sigma_k, \quad \sigma_i \in \{-1, +1\}, \quad J_{ik} \in \{-1, +1\}$$

- Bimodal disorder

$$P(J_{ik}) = \frac{1}{2} [\delta_K(J_{ik} + 1) + \delta_K(J_{ik} - 1)]$$

- Two-valued variables: one variable-one bit
- Asynchronous MultiSpin-Coding (AMSC) avoiding 'if' statements

$$J_{ik} = -1 \rightarrow 1$$

$$J_{ik} = +1 \rightarrow 0$$

$$\sigma_i = +1 \rightarrow 1$$

$$\sigma_i = -1 \rightarrow 0$$

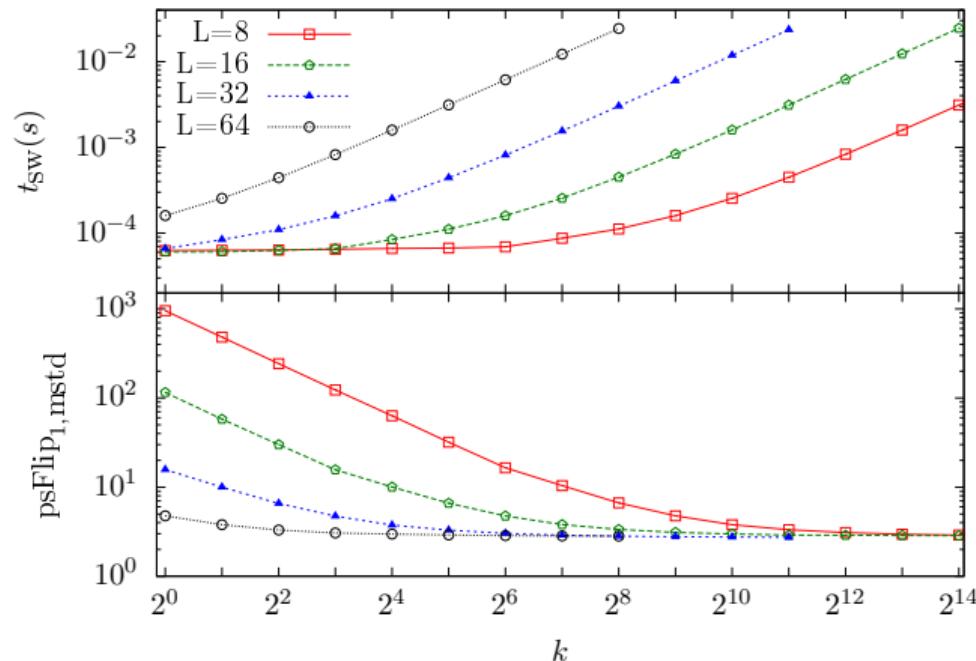
$$\mathbf{e}_{ik} = \sigma_i \hat{\sigma}_k \hat{J}_{ik}$$

$$\begin{array}{c} \mathbf{e}_{ik0} \\ \mathbf{e}_{ik1} \\ \mathbf{e}_{ik2} \\ \vdots \end{array} = \begin{array}{ccc} \hat{\sigma}_{i0} & \hat{\sigma}_{k0} & \hat{J}_{ik0} \\ \hat{\sigma}_{i1} & \hat{\sigma}_{k1} & \hat{J}_{ik1} \\ \hat{\sigma}_{i2} & \hat{\sigma}_{k2} & \hat{J}_{ik2} \\ \vdots & \vdots & \vdots \end{array}$$
$$\sum_{\langle ik \rangle} \begin{array}{c} \mathbf{e}_{ik0} \\ \mathbf{e}_{ik1} \\ \mathbf{e}_{ik2} \\ \vdots \end{array} = \begin{array}{ccc} \mathbf{s2}_{i0} & \mathbf{s1}_{i0} & \mathbf{s0}_{i0} \\ \mathbf{s2}_{i1} & \mathbf{s1}_{i1} & \mathbf{s0}_{i1} \\ \mathbf{s2}_{i2} & \mathbf{s1}_{i2} & \mathbf{s0}_{i2} \\ \vdots & \vdots & \vdots \end{array}$$

Single- & Multi-GPU Results

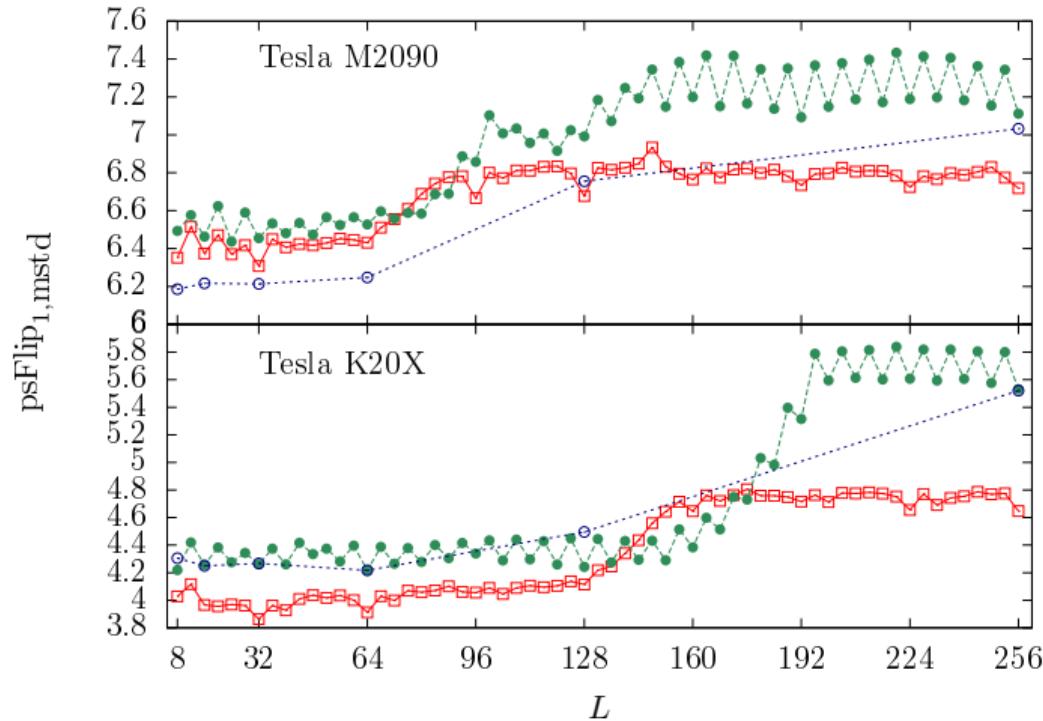
- The smaller the lattice the more disorder realizations are needed to attain saturation
- Let k be the number of coded disorder realizations and $\text{psFlip}_{n,\text{mstd}}$

$$\text{psFlip}_{n,\text{mstd}}(L, k) = t_{\text{sw}} \cdot n \cdot (32 \cdot k \cdot 4 \cdot L^3)^{-1},$$



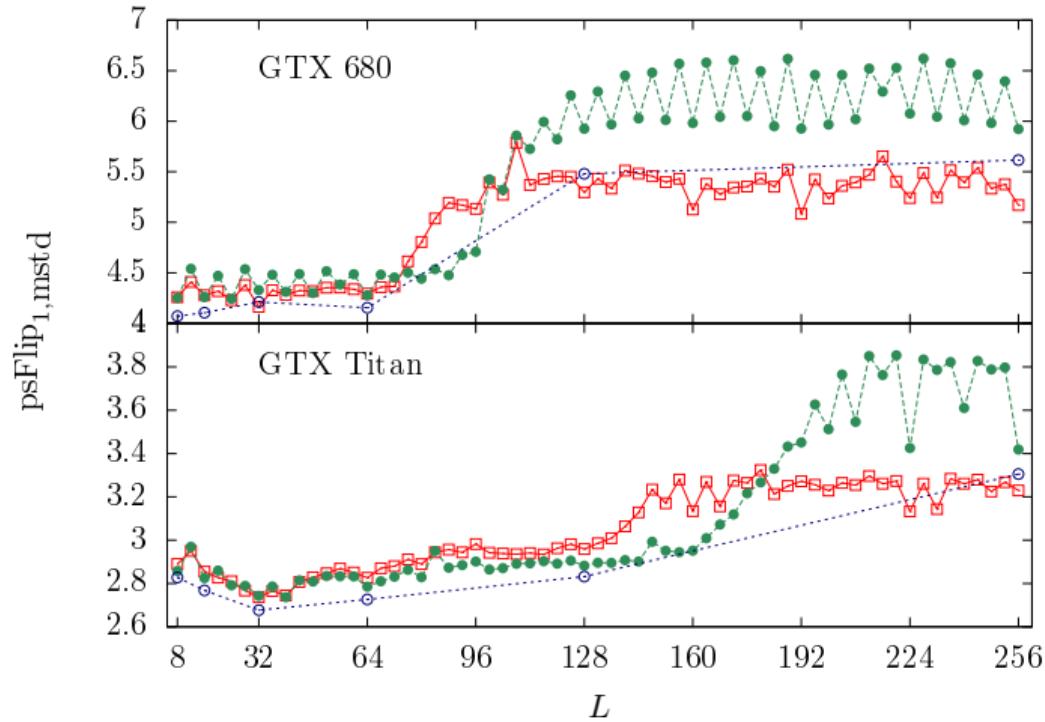
Single- & Multi-GPU Results

- Values for $\text{psFlip}_{1,\text{mstd}}$ for different values of $L = 4k$
- Sliced (red), Standard (green), Standard-Bitwise (blue)



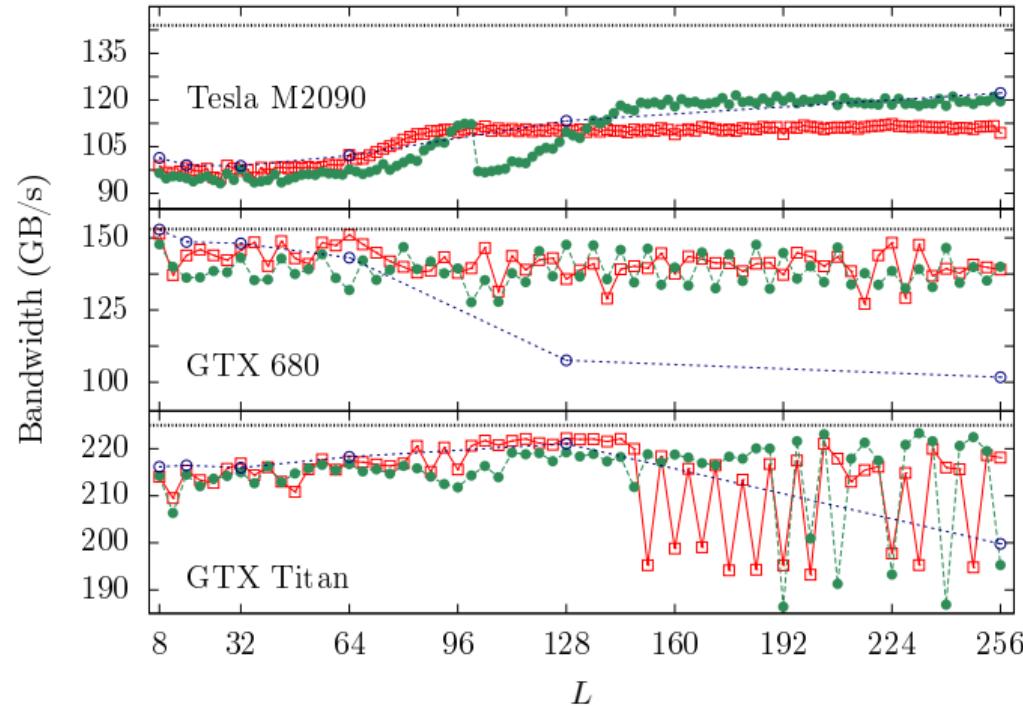
Single- & Multi-GPU Results

- Values for $\text{psFlip}_{1,\text{mstd}}$ for different values of $L = 4k$
- Sliced (red), Standard (green), Standard-Bitwise (blue)



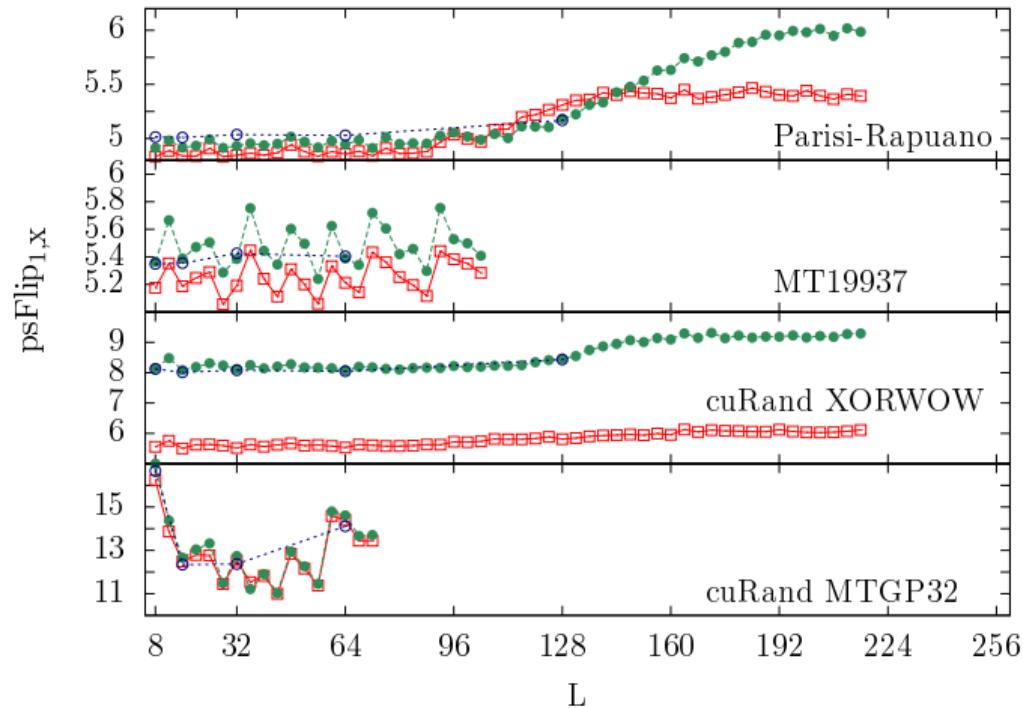
Single- & Multi-GPU Results

- Bandwidth for different values of $L = 4k$
- Sliced (red), Standard (green), Standard-Bitwise (blue)



Single- & Multi-GPU Results

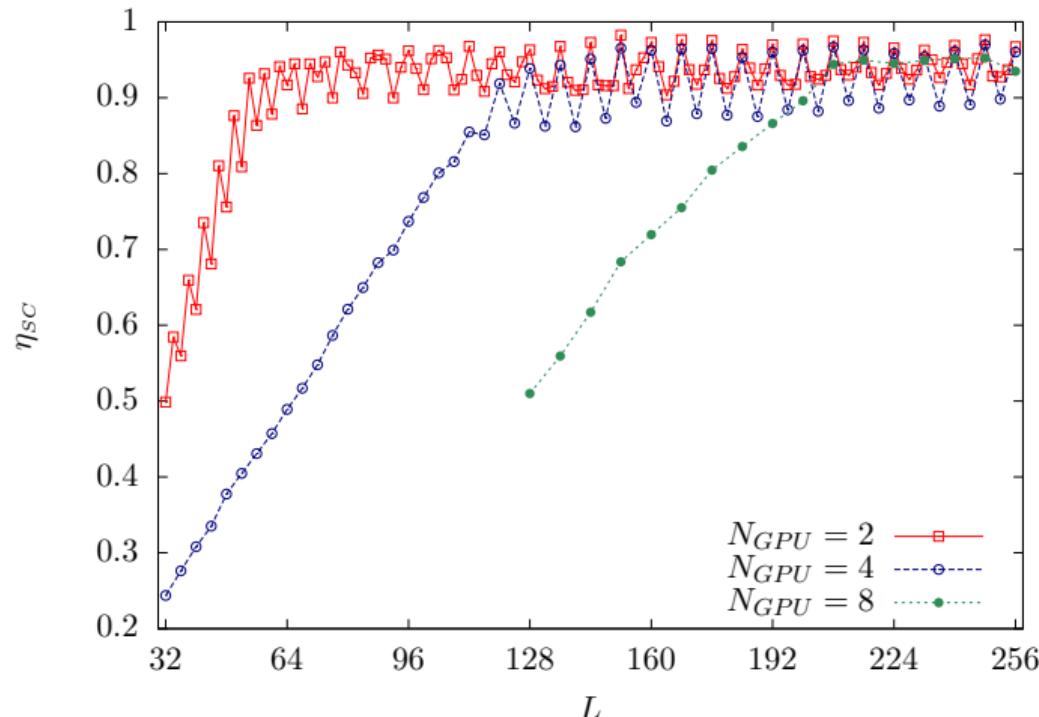
- Values for $\text{psFlip}_{1,X}$ for different random number generators and $L = 4k$
- Sliced (red), Standard (green), Standard-Bitwise (blue)



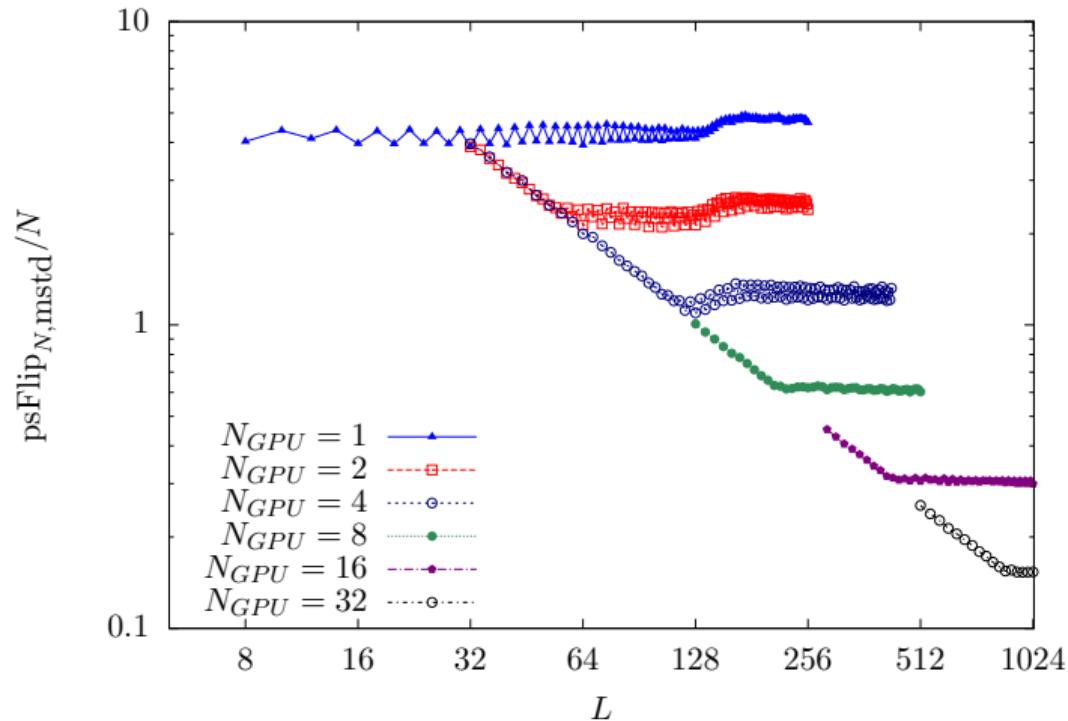
Single- & Multi-GPU Results (CSCS Piz Daint)

- strong-scaling efficiency

$$\eta_{SC} = \frac{\text{psFlip}_{1,X}}{\text{psFlip}_{N,X}},$$



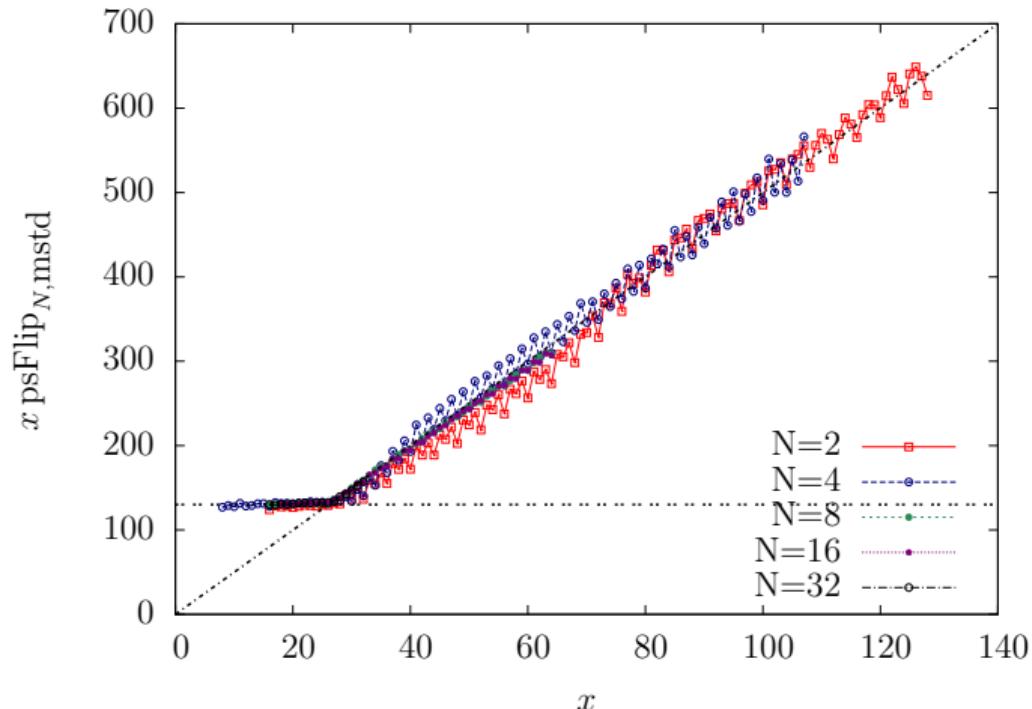
Single- & Multi-GPU Results (CSCS Piz Daint)



Single- & Multi-GPU Results (CSCS Piz Daint)

- Data scaling

$$\frac{L}{N} \text{psFlip}_{N,x}(L, k) \propto t_{\text{sw}} \left(\frac{L}{N} \right)^{-2}, \quad x = \frac{L}{N}.$$



Physics results

- This GPU implementation together with a **new** theoretical approach allowed us to obtain **cutting-edge** estimations for the critical parameters for the 3D Ising Spin Glass phase transition
- 3.1 years of one GTX Titan
- ~ 20 faster than high-end CPUs implementations and comparable with FPGA (Janus 2, 2013)

	T_c	ν	η	ω	z	Years
Janus 2013	1.1019(29)	2.562(42)	-0.3900(36)	1.12(10)		75 (FPGA)
Our Work¹	1.099(5)	2.47(10)	-0.39(1)	1.3(2)	6.80(15)	3.1 (GPU)
Haus. et al. 2008	1.109(10)	2.45(15)	-0.375(10)	1.0(1)		40 (CPU)

The use of GPUs was necessary and unavoidable in order to prove the validity and effectiveness of the new theoretic approach

¹M. Lulli, G. Parisi & A. Pelissetto in preparation

Outline for section 5

- 1 Cubic stencils
- 2 PRNGs
- 3 Multi-GPU and MPI
- 4 Results
- 5 Conclusions & Perspectives

Conclusions & Perspectives

Conclusions

- A new **sliced** cubic stencil access pattern which in some cases performs better than highly tuned solutions
- A new access pattern for lagged-Fibonacci-like PRNGs which performs better than analogues in CURAND ($\sim 2\times$)
- 3D Ising Spin Glass single GPU: performances are stable for a large interval (the largest so far) of L ; the MT19937 PRNG performs **only** $\sim 10\%$ slower than the PR
- The **sliced** multi-GPU version shows a very good strong scaling efficiency and competitive speeds for dynamic studies
- Together with a new off-equilibrium finite-size scaling approach we obtained the **2nd** most precise estimates for EA3D critical parameters so far

Perspectives

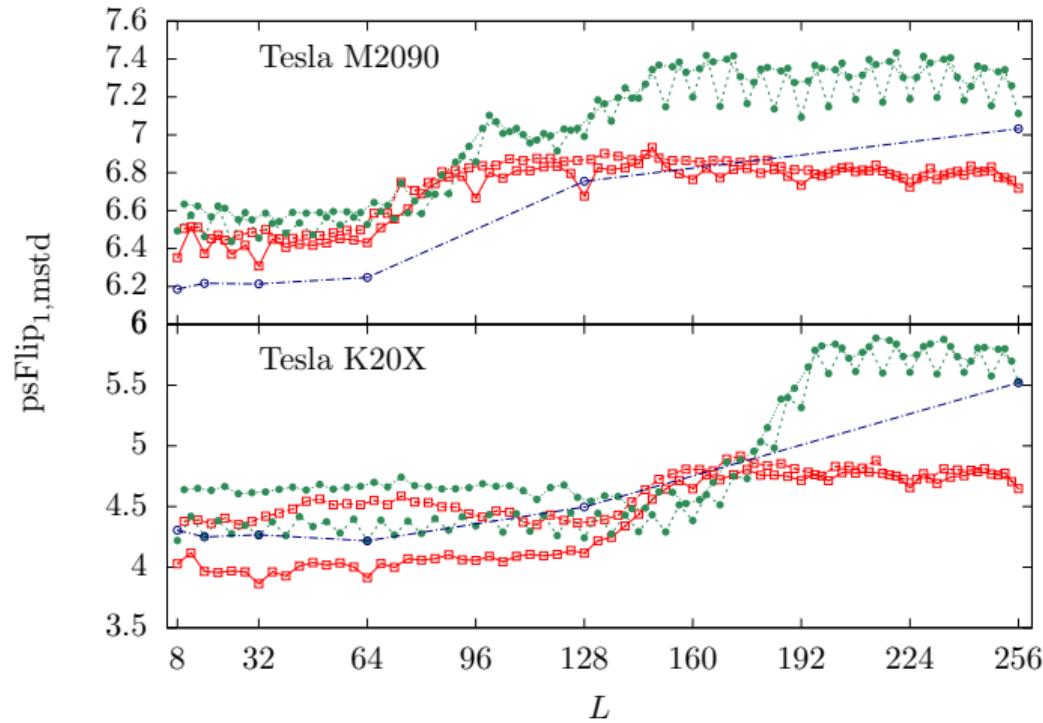
- Implement Parallel Tempering dynamics for equilibrium simulations



matteo.lulli@gmail.com
...and please give your feedback!

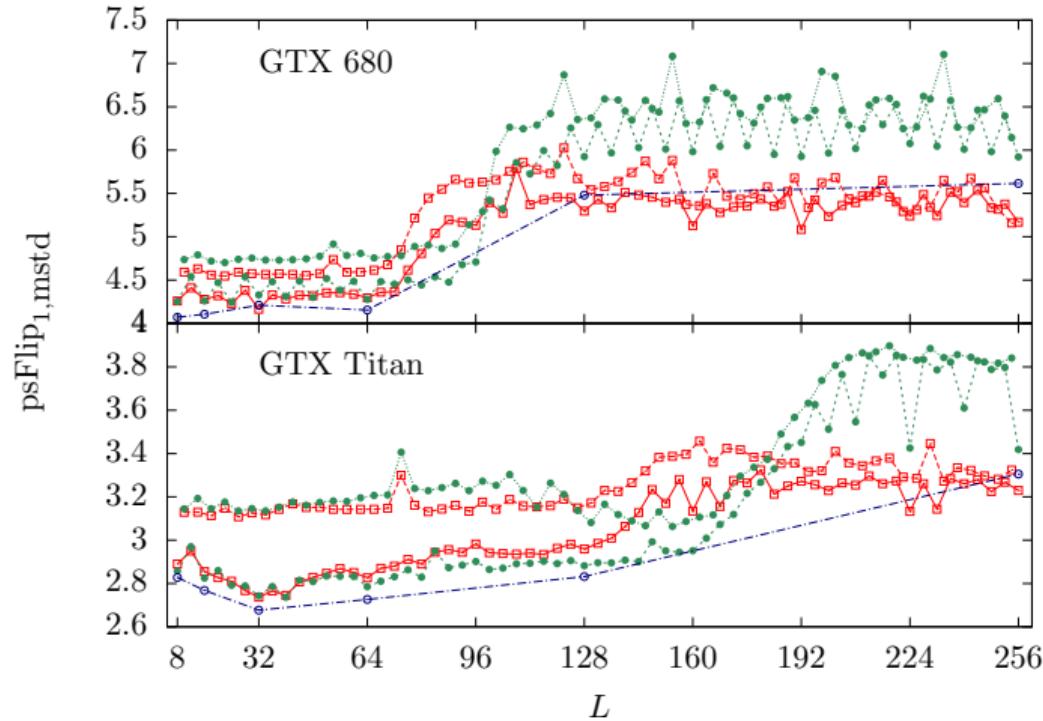
Even and Odd subsequences for $L = 2n$

- Values for $\text{psFlip}_{1,\text{mstd}}$ for different values of $L = 2n$
- Sliced (red), Standard (green), Standard-Bitwise (blue)



Even and Odd subsequences for $L = 2n$

- Values for $\text{psFlip}_{1,\text{mstd}}$ for different values of $L = 2n$
- Sliced (red), Standard (green), Standard-Bitwise (blue)



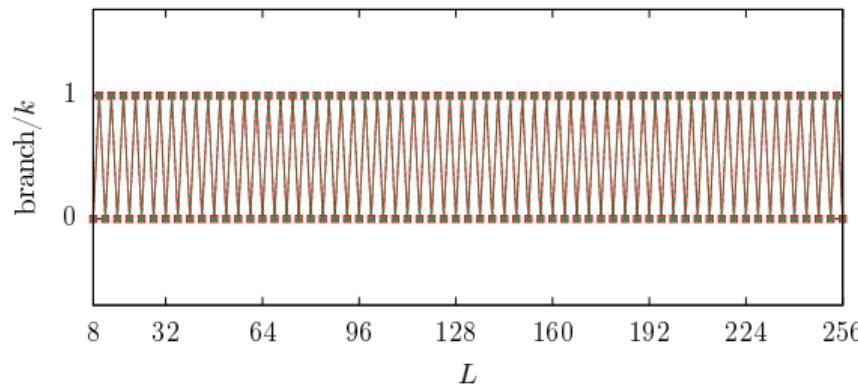
Even and Odd subsequences for $L = 2n$

What is happening?

$$L_0 = 2(2m), \quad \frac{V_0}{2} = \frac{(4m)^3}{2} = 32m^3, \quad (32m^3) \bmod(32) = 0, \quad \forall m \in \mathbb{N},$$

$$L_1 = 2(2m + 1), \quad \frac{V_1}{2} = \frac{[2(2m + 1)]^3}{2} = 4(2m + 1)^3, \quad [4(2m + 1)^3] \bmod(32) \neq 0, \quad \forall m \in \mathbb{N}.$$

- The **odd** subsequence is never commensurate to the actual warp size.
- For **every** sample there is **one** warp branching



More on MultiSpin-Coding

- Two-valued variables: one variable-one bit

More on MultiSpin-Coding

- Two-valued variables: one variable-one bit
- Asynchronous MultiSpin-Coding (AMSC)

More on MultiSpin-Coding

- Two-valued variables: one variable-one bit
- Asynchronous MultiSpin-Coding (AMSC)

$$J_{ik} = -1 \rightarrow 1$$

$$J_{ik} = +1 \rightarrow 0$$

$$\sigma_i = +1 \rightarrow 1$$

$$\sigma_i = -1 \rightarrow 0$$

$$e_{ik} = \sigma_i \hat{\wedge} \sigma_k \hat{\wedge} J_{ik}$$

$$\begin{array}{c|c|c|c} e_{ik0} & \sigma_{i0} & \sigma_{k0} & J_{ik0} \\ e_{ik1} & \sigma_{i1} & \sigma_{k1} & J_{ik1} \\ e_{ik2} & \sigma_{i2} & \sigma_{k2} & J_{ik2} \\ \vdots & \vdots & \vdots & \vdots \end{array} = \sum_{\langle ik \rangle} \begin{array}{c|c|c|c} e_{ik0} & s2_{i0} & s1_{i0} & s0_{i0} \\ e_{ik1} & s2_{i1} & s1_{i1} & s0_{i1} \\ e_{ik2} & s2_{i2} & s1_{i2} & s0_{i2} \\ \vdots & \vdots & \vdots & \vdots \end{array}$$

More on MultiSpin-Coding

- Two-valued variables: one variable-one bit
- Asynchronous MultiSpin-Coding (AMSC)

$$J_{ik} = -1 \rightarrow 1$$

$$J_{ik} = +1 \rightarrow 0$$

$$\sigma_i = +1 \rightarrow 1$$

$$\sigma_i = -1 \rightarrow 0$$

$$e_{ik} = \sigma_i \hat{\wedge} \sigma_k \hat{\wedge} J_{ik}$$

$$\begin{array}{c|c|c|c} e_{ik0} & \sigma_{i0} & \sigma_{k0} & J_{ik0} \\ e_{ik1} & \sigma_{i1} & \sigma_{k1} & J_{ik1} \\ e_{ik2} & \sigma_{i2} & \sigma_{k2} & J_{ik2} \\ \vdots & \vdots & \vdots & \vdots \end{array} = \sum_{\langle ik \rangle} \begin{array}{c|c|c|c} e_{ik0} & s2_{i0} & s1_{i0} & s0_{i0} \\ e_{ik1} & s2_{i1} & s1_{i1} & s0_{i1} \\ e_{ik2} & s2_{i2} & s1_{i2} & s0_{i2} \\ \vdots & \vdots & \vdots & \vdots \end{array}$$

Metropolis Dynamics - Avoiding 'if' statements

$$\Delta E = H[\{\sigma_{i \neq a}, -\sigma_a\}] - H[\{\sigma_{i \neq a}, \sigma_a\}] = -12, -8, -4, 0, 4, 8, 12.$$

The acceptance probability

$$P_{\text{flip}}(\Delta E) = \begin{cases} 1, & \Delta E \leq 0 \\ e^{-\beta \Delta E}, & \Delta E > 0 \end{cases}$$

More on MultiSpin-Coding

- Two-valued variables: one variable-one bit
- Asynchronous MultiSpin-Coding (AMSC)

$$J_{ik} = -1 \rightarrow 1$$

$$J_{ik} = +1 \rightarrow 0$$

$$\sigma_i = +1 \rightarrow 1$$

$$\sigma_i = -1 \rightarrow 0$$

$$e_{ik} = \sigma_i \hat{\wedge} \sigma_k \hat{\wedge} J_{ik}$$

$$\begin{array}{c|c|c|c} e_{ik0} & \sigma_{i0} & \sigma_{k0} & J_{ik0} \\ e_{ik1} & \sigma_{i1} & \sigma_{k1} & J_{ik1} \\ e_{ik2} & \sigma_{i2} & \sigma_{k2} & J_{ik2} \\ \vdots & \vdots & \vdots & \vdots \end{array} = \sum_{\langle ik \rangle} \begin{array}{c|c|c|c} e_{ik0} & s2_{i0} & s1_{i0} & s0_{i0} \\ e_{ik1} & s2_{i1} & s1_{i1} & s0_{i1} \\ e_{ik2} & s2_{i2} & s1_{i2} & s0_{i2} \\ \vdots & \vdots & \vdots & \vdots \end{array}$$

Metropolis Dynamics - Avoiding 'if' statements

$$(0, 0, 0) = 0 \rightarrow \Delta E = -12 \quad (1, 0, 0) = 4 \rightarrow \Delta E = 4$$

$$(0, 0, 1) = 1 \rightarrow \Delta E = -8 \quad (1, 0, 1) = 5 \rightarrow \Delta E = 8$$

$$(0, 1, 0) = 2 \rightarrow \Delta E = -4 \quad (1, 1, 0) = 6 \rightarrow \Delta E = 12$$

$$(0, 1, 1) = 3 \rightarrow \Delta E = 0$$

More on MultiSpin-Coding

- Two-valued variables: one variable-one bit
- Asynchronous MultiSpin-Coding (AMSC)

$$J_{ik} = -1 \rightarrow 1$$

$$J_{ik} = +1 \rightarrow 0$$

$$\sigma_i = +1 \rightarrow 1$$

$$\sigma_i = -1 \rightarrow 0$$

$$e_{ik} = \sigma_i \wedge \sigma_k \wedge J_{ik}$$

$$\begin{array}{c|c|c|c} e_{ik0} & \sigma_{i0} & \sigma_{k0} & J_{ik0} \\ e_{ik1} & \sigma_{i1} & \sigma_{k1} & J_{ik1} \\ e_{ik2} & \sigma_{i2} & \sigma_{k2} & J_{ik2} \\ \vdots & \vdots & \vdots & \vdots \end{array} = \sum_{\langle ik \rangle} \begin{array}{c|c|c|c} e_{ik0} & s2_{i0} & s1_{i0} & s0_{i0} \\ e_{ik1} & s2_{i1} & s1_{i1} & s0_{i1} \\ e_{ik2} & s2_{i2} & s1_{i2} & s0_{i2} \\ \vdots & \vdots & \vdots & \vdots \end{array}$$

Metropolis Dynamics - Avoiding 'if' statements

Normalized transition probabilities & random number comparison

$$R_{max} \exp(-\beta \Delta E) = \text{EXP12}, \text{EXP8}, \text{EXP4}$$

```
cond12 = -(R < EXP12); cond8 = -(R < EXP8); cond4 = -(R < EXP4);
```

Spin flip: `spin = mask ^ spin;`

```
mask = cond12 | (~sum2)
      | ((sum2 & (sum2^sum1)) & (cond8 | (cond4 & (~sum0))));
```

Tiny differences - Calligraphy

Sliced Addressing

```
x = i%d_L; y = (i/d_L)%d_L; z = i/d_A;  
  
smz = i + (SM(z - 1, d_hL) - z)*d_A;  
spy = smz + (SP(y + 1, d_L) - y)*d_L;  
smy = i + (SM(y - 1, d_L) - y)*d_L;  
smx = spy - x + SM(x - 1, d_L);  
spx = smy - x + SP(x + 1, d_L);
```

Standard Addressing

```
x = i%d_hL; y = (i/d_hL)%d_L; z = i/d_hA;  
par = (y^z)&1;  
  
spx = i - x + SP(x + 1 - (par^1), d_hL);  
smx = i - x + SM(x - 1 + par, d_hL);  
spy = i + (SP(y + 1, d_L) - y)*d_hL;  
smy = i + (SM(y - 1, d_L) - y)*d_hL;  
spz = i + (SP(z + 1, d_L) - z)*d_hA;  
smz = i + (SM(z - 1, d_L) - z)*d_hA;
```

For the **sliced scheme** $spz = i$, and **parity** is not involved