CATEGORY: COMPUTATIONAL PHYSICS - CP16

POSTER P5266 | CONTACT NAME Jeffrey Kelling: j.kelling@hzdr.de

GPU TECHNOLOGY CONFERENCE

Handling Domain Decomposition in Massively Parallel Implementations of Stochastic Lattice Models

Jeffrey Kelling[1], Géza Ódor[2], Sibylle Gemming[1,3]

[1]Institute of Ion Beam Physics and Materials Research, Helmholtz-Zentrum Dresden-Rossendorf, Germany

[2]MTA TTK MFA Research Institute for Natural Sciences, Budapest, Hungary

[3]Institute of Physics, TU-Chemnitz, Germany

DRESDEN concept | HZDR HELMHOLTZ ZENTRUM DRESDEN ROSSENDORF

- nanopatterning of materials from molecular electronics to photovoltaics
- need large scale atomistic simulations to understand self-organization
- only stochastic models can bridge the gap from nano to micro
- GPGPU enables simulations of micron-sized volumes, billions of atoms and studies of the long-time evolution of systems
- random site-selection is essential but can be harmed by domain decomposition

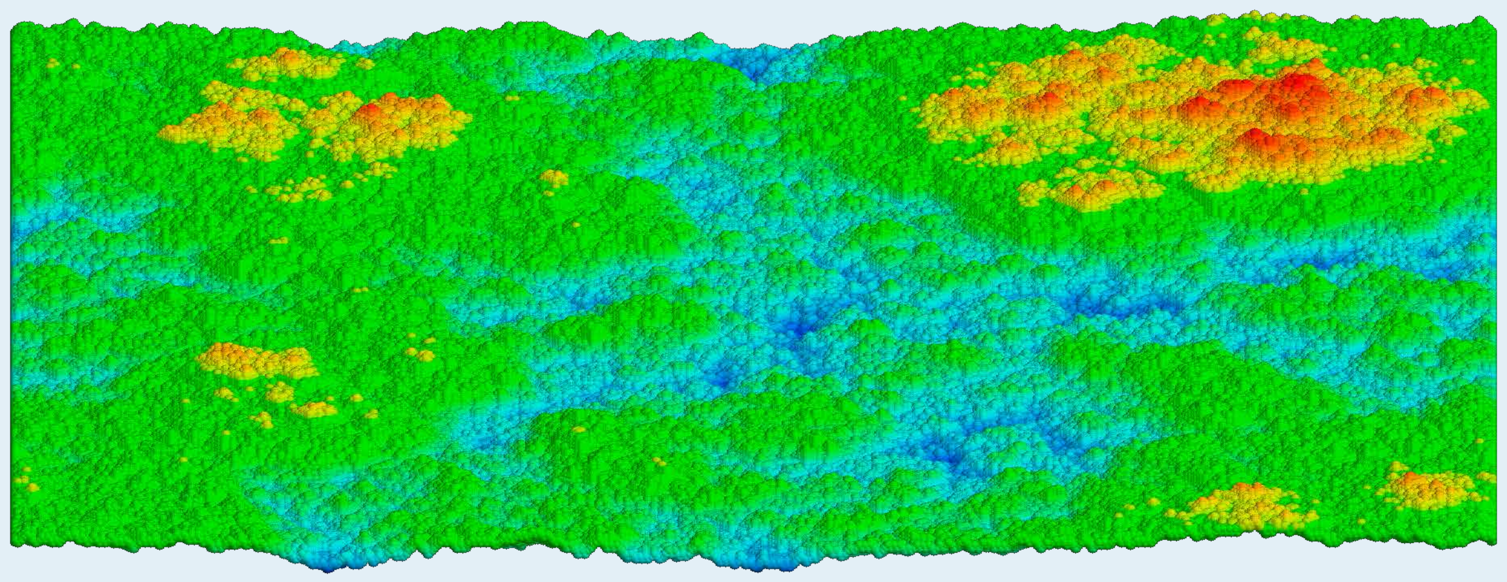

KPZ surface in the steady state

- Kardar–Parisi–Zhang (KPZ) universality class [1]
- can describe *growth processes*, directed polymers in random media, randomly stirred fluids, dissipative transport and magnetic flux lines in superconductors.

$$\mathrm{d}_t h(\mathbf{x},t) = \underbrace{v}_{\text{mean growth vel.}} + \underbrace{\sigma_2 \nabla^2 h(\mathbf{x},t)}_{\text{surface tension}} + \underbrace{\lambda\left[\nabla h(\mathbf{x},t)\right]^2}_{\text{local growth vel.}} + \underbrace{\eta(\mathbf{x},t)}_{\text{noise}}$$

KPZ stochastic differential equation

Roof-Top Model

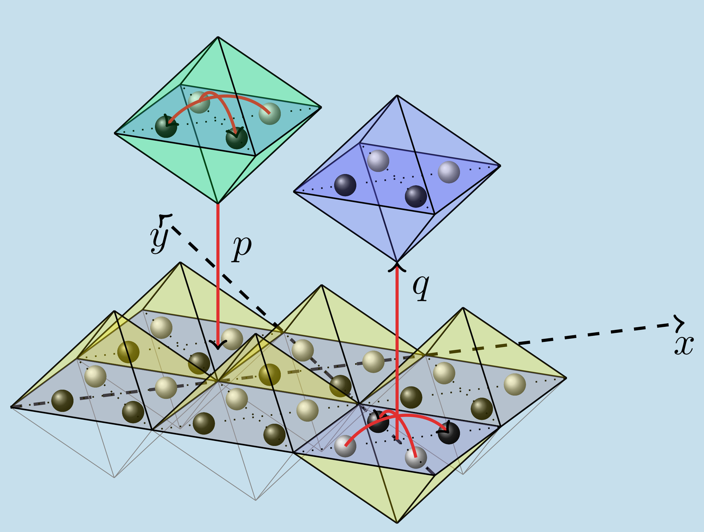

$2+1$D roof-top model—octahedron model

- stochastic model for surface growth [2, 3]
- 2D square lattice (octahedron model)
- carry out *random* depositions (probability p) or removals (prob. q)
- time t measured in sweeps of the lattice—Monte-Carlo-Steps (MCS)
- shows KPZ universality for $q = 0$
- bit-coded up- or down-slopes in between lattice-sites

Parallel Implementation

- CUDA implementation enables runs with $2^{17} \times 2^{17} \approx 16 \times 10^9$ lattice-sites [4]
  + self-averaging system ⇒ low noise in observables
- *aim:* GPU implementation should preserve random site-selection
  ⇒ *domain decomposition* to distribute work among multiprocessors (blocks)
  see right-hand side ↗
- cells loaded into shared memory, updated collectively by threads
  − single-hit double tiling scheme with fixed borders
  − each thread maintains independent RNG state (TinyMT [5])

Domain Decomposition

- simulation-cell is split into non-interacting cells
- cells are treated independently for short periods of time, effectively with fixed boundary conditions
  ⇒ introducing small errors that must be attenuated
  − larger active domains give better results
- *Dead Border Decomposition:* tiles, sites interacting with neighboring cells are inactive, origin randomized in intervals
- *Double Tiling Decompostion:* tiles split into 4 subtiles, only one set of subtiles active at a time, origin randomized

Auto-correlation of slopes during growth
$C(t,s) = \langle \phi(t)\phi(s) \rangle - \langle \phi(t) \rangle \langle \phi(s) \rangle$ for times $t > s$

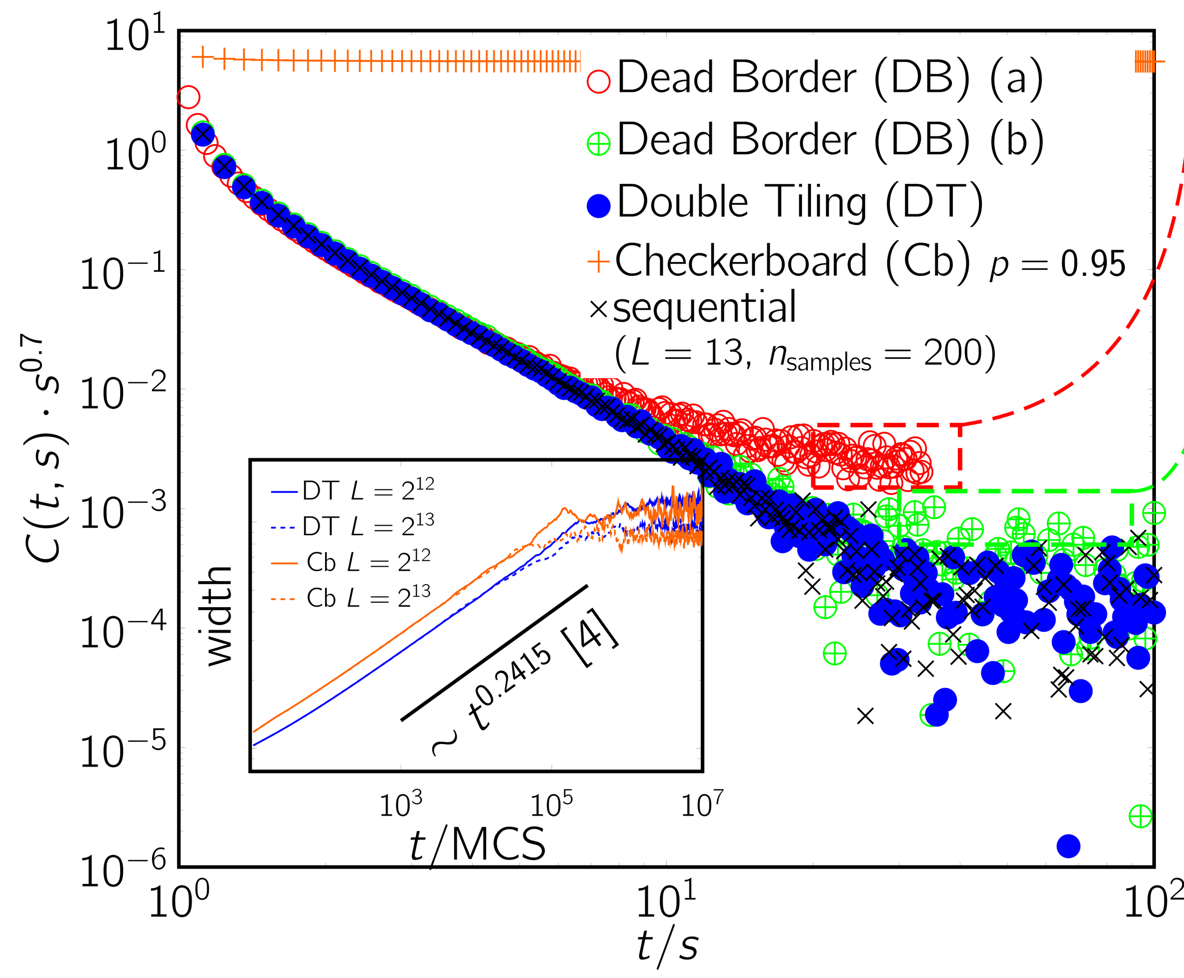

Benchmarks

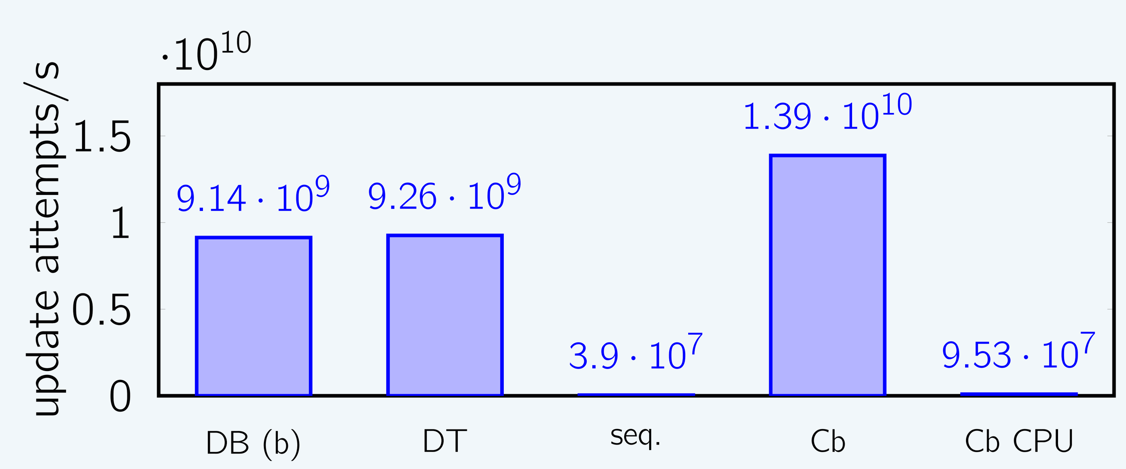

- sequential code: Intel Core i7-4930K @3.4 GHz
- CUDA code: NVIDIA GTX Titan Black

Conclusions

⇒ enabling long-time studies and large-scale simulations

- no correlations using double tiling with randomized origin
  − speedup ∼ 238× vs. sequential code
- straightforwardly adaptable for multi-GPU

Dead Border Decomposition

(a) smallest possible border for the Roof-Top Model: one lattice site to the right and bottom (see figure)
- borders can only be moved to word boundaries for encoding/performance reasons (chunks of 4×4 lattice sites)

+ fast, good scaling results
− restricted border-movement leads to correlations

(b) remove restriction for borders by using wider borders (four sites)

+ removes correlation
− bad signal-to-noise–ratio compared to sequential code … due to added disorder from crossing of wide borders

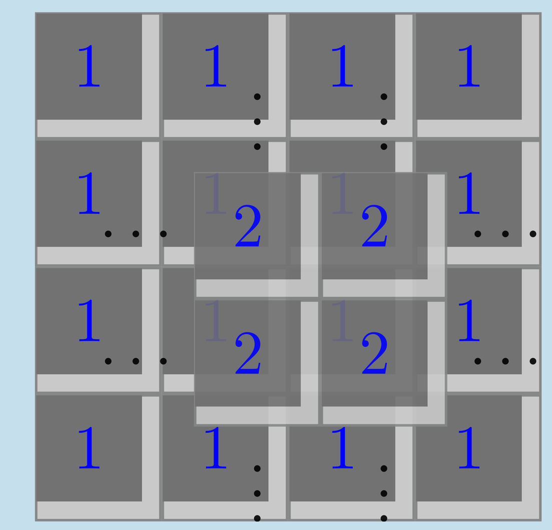

Dead-border decomposition: Light-gray areas are domain borders with inactive sites.

Double Tiling Decomposition

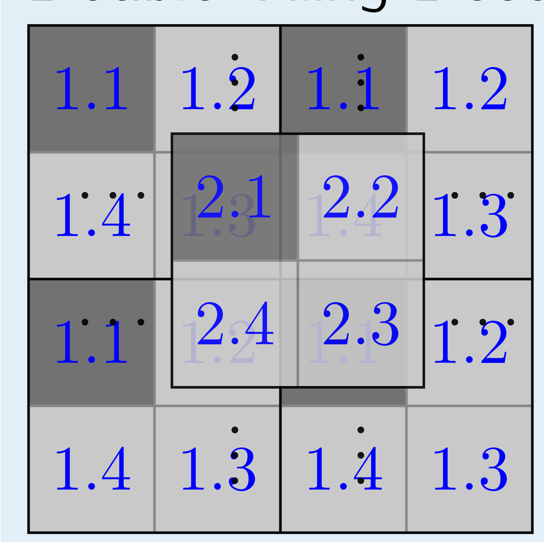

Double-tiling decomposition: Light-gray sub-cells are inactive.

- four non-interacting sets of sub-cells are updated in random order
- active regions overlap with inactive cells
  ⇒ works without shifting the origin for different models like Kinetic Metropolis Lattice Monte-Carlo [6]
- randomly shift origin of decomposition after every sweep to avoid auto-correlation error

+ noise en par with sequential code
  ⇒ *capability to perform large-scale aging studies*

Even/Odd site Checkerboard

- even sites do not interact with each other in single update

⇒ update even sites, then odd sites—often used for Ising model [7]

- replacing random site selection by finite update probability $p < 1$

+ correct growth results, very fast
− updates correlated, correlation depends on p

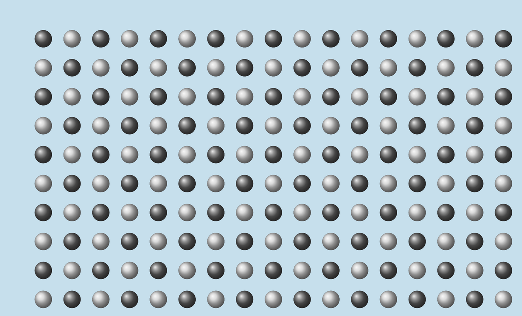

Lattice sites colored by parity.

Acknowledgements

We thank M. Bussmann for providing access to additional GPUs, the computing center at HZDR for support and the CCoE Dresden and the Helmholtz Initiative and Networking Funds for funding.

[1] Kardar, M., Parisi, G. & Zhang, Y.-C. *Phys. Rev. Lett.* **56**, 889–892 (1986).

[2] Plischke, M., Rácz, Z. & Liu, D. *Phys. Rev. B* **35**, 3485–3495 (1987).

[3] Ódor, G., Liedke, B. & Heinig, K.-H. *Phys. Rev. E* **79** (2009).

[4] Kelling, J. & Ódor, G. *Phys. Rev. E* **84**, 061150 (2011).

[5] TinyMT. URL http://www.math.sci.hiroshima-u.ac.jp/~%20m-mat/MT/TINYMT/index.html.

[6] Kelling, J., Ódor, G., Nagy, M. F., Schulz, H. & Heinig, K. *The European Physical Journal - Special Topics* **210**, 175–187 (2012).

[7] Preis, T. *et al. J. Comp. Phys.* **228**, 4468–4477 (2009).

[8] Kelling, J., Ódor, G. & Gemming, S. *in preparation* (2015).