

# Parallel Simulation of Dendritic Growth On Unstructured Grids

Andreas Schäfer, Julian Hammer, Dietmar Fey

Friedrich-Alexander-Universität Erlangen-Nürnberg

IA<sup>3</sup> Workshop @SC11, Nov. 13th, 2011


# Outline

- 1 What and why?
- 2 Specialized Parallelization
- 3 Stencilized Parallelization

# Outlook

- 1 What and why?
- 2 Specialized Parallelization
- 3 Stencilized Parallelization

# Simulation of Dendritic Growth in Al/Cu




Microscope



Simulation

- two classes of models:
  - 1 cellular automata (our approach)
  - 2 phase field method  
(Peta-scale Phase-Field Simulation for Dendritic Solidification on the TSUBAME 2.0 Supercomputer)
- meshfree (no regular grid)

# Simulation Model

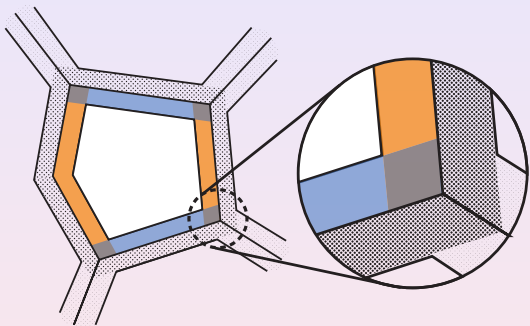


- black: solid cells
- green: liquid cells
- squares: particles on phase boundary


# Outlook

- 1 What and why?
- 2 Specialized Parallelization**
- 3 Stencilized Parallelization

- 1 decompose graph via ParMETIS
- 2 loop
  - 1 sync ghostzones
  - 2 update
  - 3 (output)




# Communication Graph



- $10^6$  cells, 10 MPI processes
- 74k ghost cells
- 2 GB/step



# Communication Graph




- $10^6$  cells, 100 MPI processes
- 475k ghost cells
- 11 GB/step

# Outlook

- 1 What and why?
- 2 Specialized Parallelization
- 3 Stencilized Parallelization**

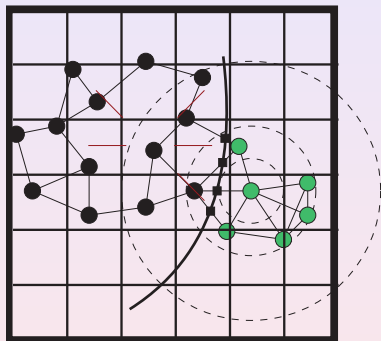
# Stencilization

- superimpose grid on irregular graph
- place cells into container cells
- **physically equivalent**
- reuse existing library:  
LibGeoDecomp
  - overlapping comm. & calc.
  - hybrid parallelization

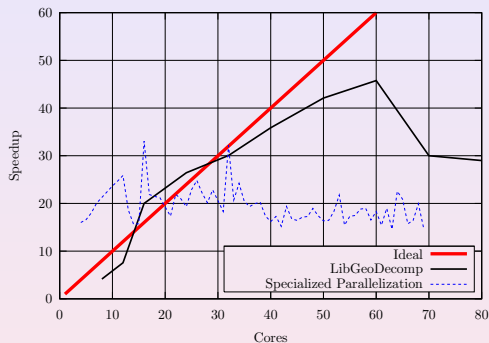


# Stencilization

- superimpose grid on irregular graph
- place cells into container cells
- **physically equivalent**
- reuse existing library:  
LibGeoDecomp
  - overlapping comm. & calc.
  - hybrid parallelization



# Evaluation: Speedup



- testbed: 28 IBM LS21 blades (Opteron dual-cores)
- 10 Gb InfiniBand

# Overlapping Communication and Calculation

- Myth #1: It's as easy as calling `MPI_Isend()`
- Myth #2: It's not possible at all

- 1 `MPI_Isend()`
- 2 loop
  - 1 `MPI_Test()`
  - 2 `work()`
- 3 `MPI_Wait()`

# Overlapping Communication and Calculation

- Myth #1: It's as easy as calling `MPI_Isend()`
- Myth #2: It's not possible at all

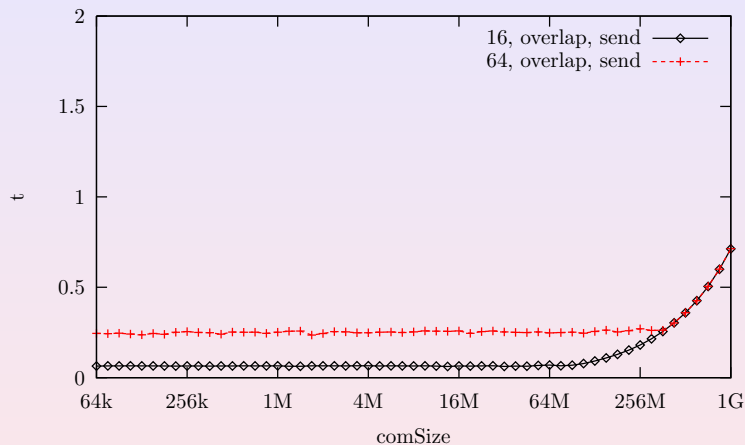
```
1 MPI_Isend()  
2 loop  
  1 MPI_Test()  
  2 work()  
3 MPI_Wait()
```

# Overlapping Communication and Calculation

- Myth #1: It's as easy as calling `MPI_Isend()`
  - Myth #2: It's not possible at all
- 1 `MPI_Isend()`
  - 2 loop
    - 1 `MPI_Test()`
    - 2 `work()`
  - 3 `MPI_Wait()`



# Overlapping Micro Benchmark



- Open MPI + InfiniBand

# Conclusion


- communication-bound model
- *stencilization* surprisingly efficient
  - reduces number of neighbors
  - but model changes may be substantial
- use MPI+OpenMP to reduce memory traffic
- asynchronous communication by repeatedly poking MPI



<http://www.libgeodecomp.org>


# Backup

# Improving Data Locality for Communication

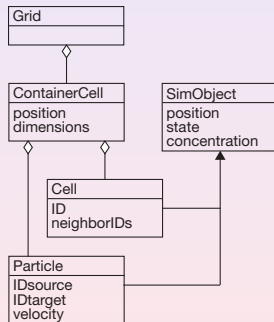


# Efficient Memory Layout


## Original Layout




## Optimized Layout



# Simulation of Dendritic Growth in Al/Cu



Microscope



Simulation

- model courtesy of
  - Department of Metallic Materials, FSU Jena, Germany
- name derived from greek  $\delta\epsilon\nu\delta\rho\nu$  (dendron)
- **not supercooled**