



High-Performance and Scalable Agent-Based Simulation With BioDynaMo

*Lukas Breitwieser, Ahmad Hesam, Fons Rademakers,
Juan Gómez Luna, and Onur Mutlu*

<https://arxiv.org/abs/2301.06984>, <https://doi.org/10.1145/3572848.3577480>

ACM SIGPLAN Annual Symposium on Principles and Practice of Parallel Programming
27 February 2023, Montreal Canada



ETH zürich

SAFARI

TUDelft

Key problem



Most existing agent-based simulation platforms exhibit **low performance**

Impact of low performance

Impact of low performance



Limits the model size and complexity

Impact of low performance



Limits the model size and complexity



Increases the development time

Impact of low performance



Limits the model size and complexity



Increases the development time



Limits the capability to optimize parameters
→ might lead to less optimal solution

Impact of low performance



Limits the model size and complexity



Increases the development time



Limits the capability to optimize parameters
→ might lead to less optimal solution



Increases costs

Our solution: BioDynaMo



BioDynaMo is a **modular and high-performance** agent-based simulation platform written in C++.

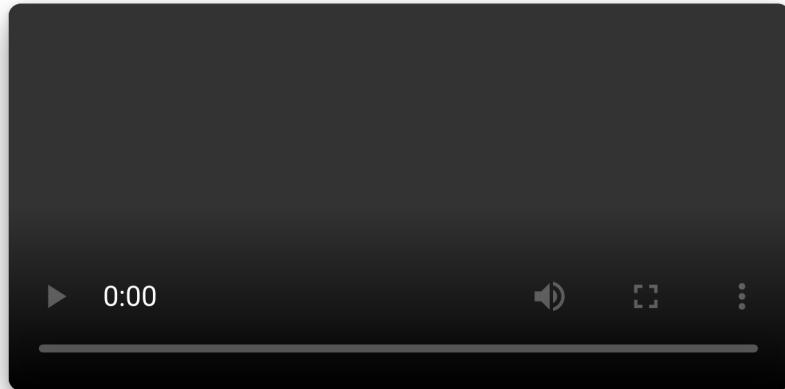
<https://biodynamo.org>



Background

Example agent-based model

Modeling a swarm of birds



Example agent-based model

Modeling a swarm of birds

Agent is a bird with the following attributes:

- position
- velocity



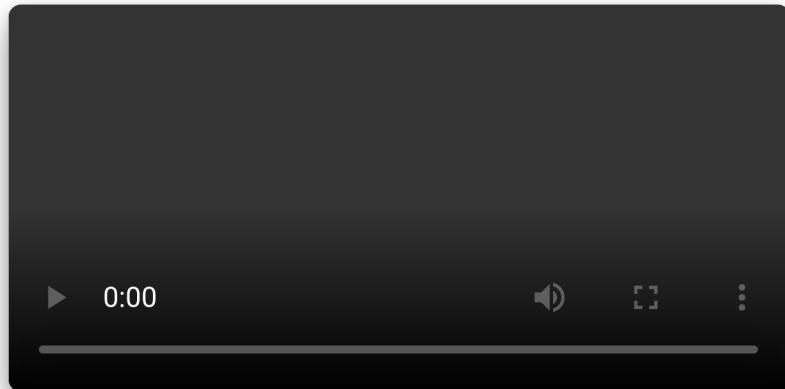
Example agent-based model

Modeling a swarm of birds

Agent is a bird with the following attributes:

- position
- velocity

and three behaviors



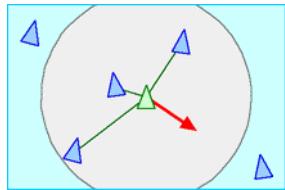
Example agent-based model

Modeling a swarm of birds

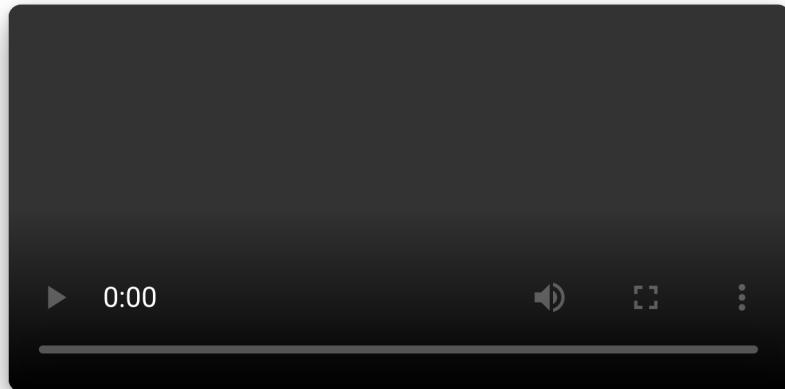
Agent is a bird with the following attributes:

- position
- velocity

and three behaviors



Separation



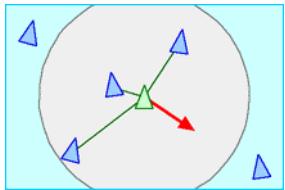
Example agent-based model

Modeling a swarm of birds

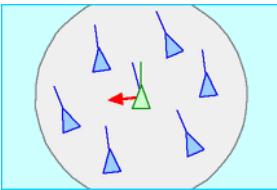
Agent is a bird with the following attributes:

- position
- velocity

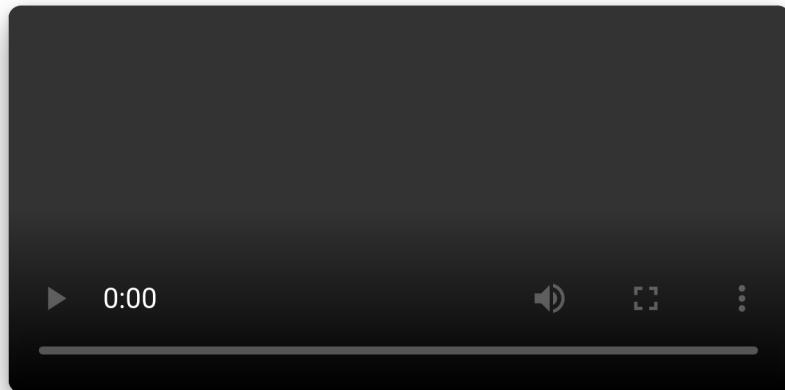
and three behaviors



Separation



Alignment



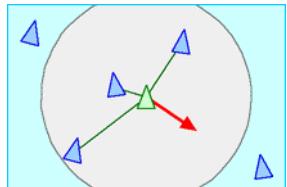
Example agent-based model

Modeling a swarm of birds

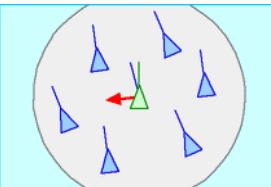
Agent is a bird with the following attributes:

- position
- velocity

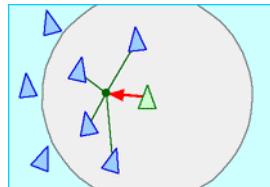
and three behaviors



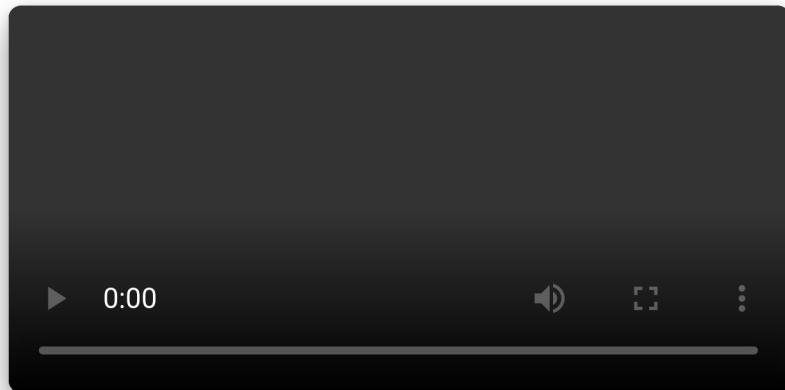
Separation



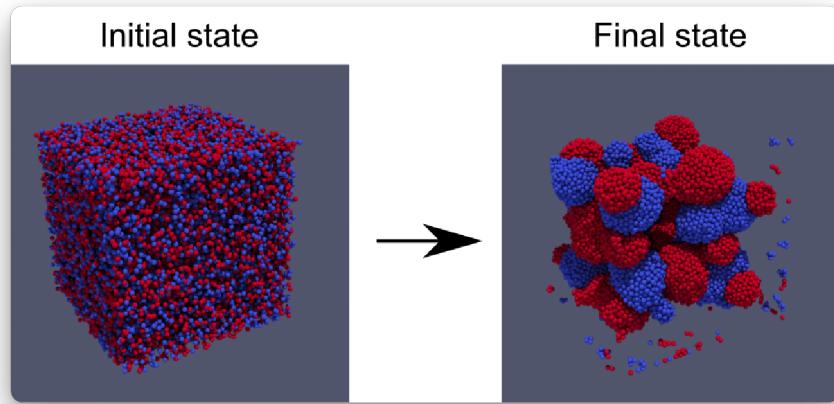
Alignment



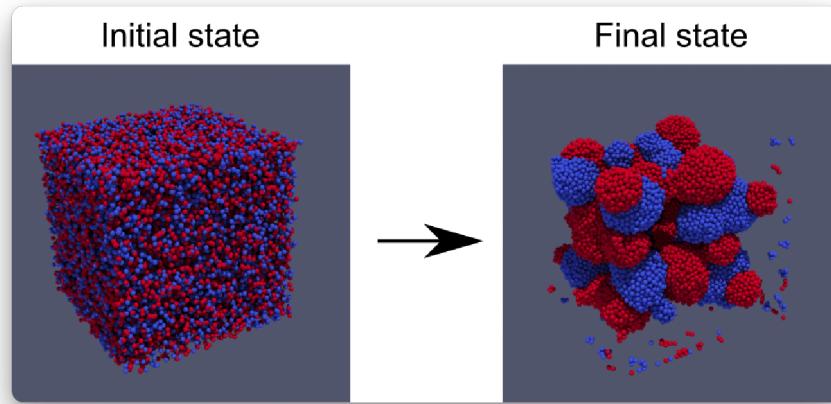
Cohesion



Cell sorting in BioDynaMo



Cell sorting in BioDynaMo

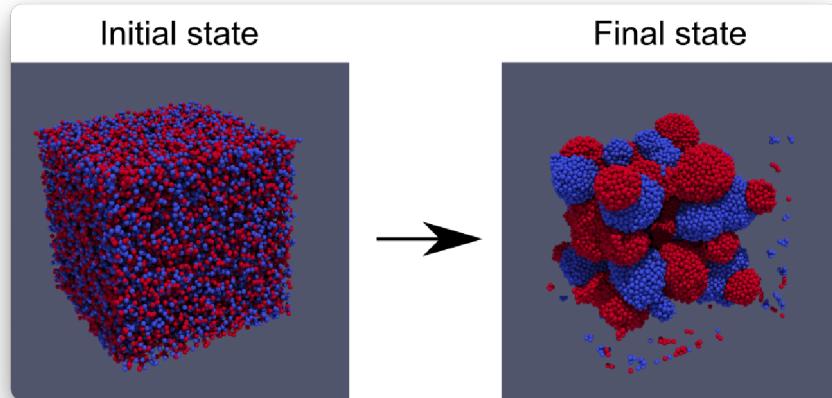


Agent

Spherical cell with type attribute (red and blue)

Behaviors

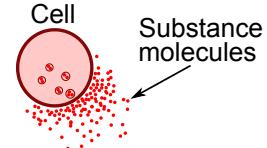
Cell sorting in BioDynaMo



Agent

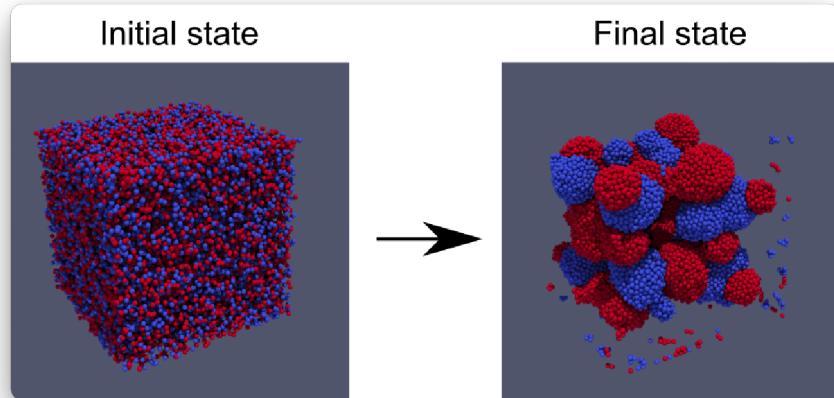
Spherical cell with type attribute (red and blue)

Behaviors



Substance
secretion

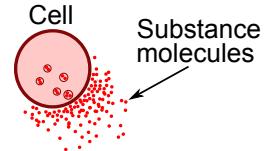
Cell sorting in BioDynaMo



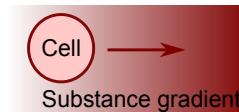
Agent

Spherical cell with type attribute (red and blue)

Behaviors

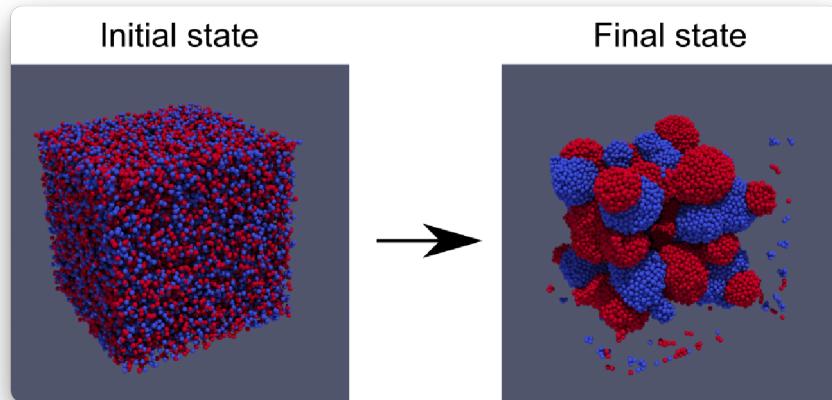


Substance
secretion



Chemotaxis

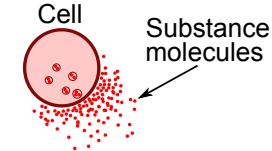
Cell sorting in BioDynaMo



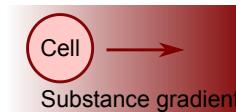
Agent

Spherical cell with type attribute (red and blue)

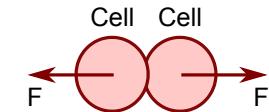
Behaviors



Substance secretion



Chemotaxis



Mechanical forces

Cell sorting implementation

```
1 // Model initialization
2 AddSubstance('purple');
3 AddSubstance('light-blue');
4 CreateAgentsRandom(number, 'red');
5 CreateAgentsRandom(number, 'blue');
6
7 // Simulation
8 for each iteration {
9     UpdateNeighborSearchIndex();
10
11     for each agent {
12         SecreteSubstance(agent);
13         FollowSubstanceGradient(agent);
14         CalculateForces(agent);
15     }
16
17     DiffuseSubstance('purple');
18     DiffuseSubstance('light-blue');
19
20     Visualize();
21 }
```

Cell sorting implementation

```
1 // Model initialization
2 AddSubstance('purple');
3 AddSubstance('light-blue');
4 CreateAgentsRandom(number, 'red');
5 CreateAgentsRandom(number, 'blue');
6
7 // Simulation
8 for each iteration {
9     UpdateNeighborSearchIndex();
10
11     for each agent {
12         SecreteSubstance(agent);
13         FollowSubstanceGradient(agent);
14         CalculateForces(agent);
15     }
16
17     DiffuseSubstance('purple');
18     DiffuseSubstance('light-blue');
19
20     Visualize();
21 }
```

Cell sorting implementation

```
1 // Model initialization
2 AddSubstance('purple');
3 AddSubstance('light-blue');
4 CreateAgentsRandom(number, 'red');
5 CreateAgentsRandom(number, 'blue');
6
7 // Simulation
8 for each iteration {
9     UpdateNeighborSearchIndex();
10
11    for each agent {
12        SecreteSubstance(agent);
13        FollowSubstanceGradient(agent);
14        CalculateForces(agent);
15    }
16
17    DiffuseSubstance('purple');
18    DiffuseSubstance('light-blue');
19
20    Visualize();
21 }
```

Cell sorting implementation

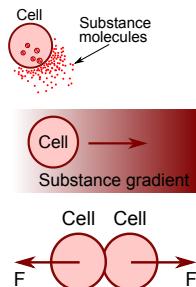
```
1 // Model initialization
2 AddSubstance('purple');
3 AddSubstance('light-blue');
4 CreateAgentsRandom(number, 'red');
5 CreateAgentsRandom(number, 'blue');
6
7 // Simulation
8 for each iteration {
9     UpdateNeighborSearchIndex();
10
11     for each agent {
12         SecreteSubstance(agent);
13         FollowSubstanceGradient(agent);
14         CalculateForces(agent);
15     }
16
17     DiffuseSubstance('purple');
18     DiffuseSubstance('light-blue');
19
20     Visualize();
21 }
```

Cell sorting implementation

```
1 // Model initialization
2 AddSubstance('purple');
3 AddSubstance('light-blue');
4 CreateAgentsRandom(number, 'red');
5 CreateAgentsRandom(number, 'blue');
6
7 // Simulation
8 for each iteration {
9     UpdateNeighborSearchIndex();
10
11     for each agent {
12         SecreteSubstance(agent);
13         FollowSubstanceGradient(agent);
14         CalculateForces(agent);
15     }
16
17     DiffuseSubstance('purple');
18     DiffuseSubstance('light-blue');
19
20     Visualize();
21 }
```

Cell sorting implementation

```
1 // Model initialization
2 AddSubstance('purple');
3 AddSubstance('light-blue');
4 CreateAgentsRandom(number, 'red');
5 CreateAgentsRandom(number, 'blue');
6
7 // Simulation
8 for each iteration {
9     UpdateNeighborSearchIndex();
10
11     for each agent {
12         SecreteSubstance(agent);
13         FollowSubstanceGradient(agent);
14         CalculateForces(agent);
15     }
16
17     DiffuseSubstance('purple');
18     DiffuseSubstance('light-blue');
19
20     Visualize();
21 }
```



Cell sorting implementation

```
1 // Model initialization
2 AddSubstance('purple');
3 AddSubstance('light-blue');
4 CreateAgentsRandom(number, 'red');
5 CreateAgentsRandom(number, 'blue');
6
7 // Simulation
8 for each iteration {
9     UpdateNeighborSearchIndex();
10
11     for each agent {
12         SecreteSubstance(agent);
13         FollowSubstanceGradient(agent);
14         CalculateForces(agent);
15     }
16
17     DiffuseSubstance('purple');
18     DiffuseSubstance('light-blue');
19
20     Visualize();
21 }
```

Cell sorting implementation

```
1 // Model initialization
2 AddSubstance('purple');
3 AddSubstance('light-blue');
4 CreateAgentsRandom(number, 'red');
5 CreateAgentsRandom(number, 'blue');
6
7 // Simulation
8 for each iteration {
9     UpdateNeighborSearchIndex();
10
11     for each agent {
12         SecreteSubstance(agent);
13         FollowSubstanceGradient(agent);
14         CalculateForces(agent);
15     }
16
17     DiffuseSubstance('purple');
18     DiffuseSubstance('light-blue');
19
20     Visualize();
21 }
```

Cell sorting implementation

```
1 // Model initialization
2 AddSubstance('purple');
3 AddSubstance('light-blue');
4 CreateAgentsRandom(number, 'red');
5 CreateAgentsRandom(number, 'blue');
6
7 // Simulation
8 for each iteration {
9     UpdateNeighborSearchIndex();
10
11     for each agent {
12         SecreteSubstance(agent);
13         FollowSubstanceGradient(agent);
14         CalculateForces(agent);
15     }
16
17     DiffuseSubstance('purple');
18     DiffuseSubstance('light-blue');
19
20     Visualize();
21 }
```

Fundamental operations

Cell sorting implementation

```
1 // Model initialization
2 AddSubstance('purple');
3 AddSubstance('light-blue');
4 CreateAgentsRandom(number, 'red');
5 CreateAgentsRandom(number, 'blue');
6
7 // Simulation
8 for each iteration {
9     UpdateNeighborSearchIndex();
10
11     for each agent {
12         SecreteSubstance(agent);
13         FollowSubstanceGradient(agent);
14         CalculateForces(agent);
15     }
16
17     DiffuseSubstance('purple');
18     DiffuseSubstance('light-blue');
19
20     Visualize();
21 }
```

Fundamental operations

- Iterate over agents and execute a function

Cell sorting implementation

```
1 // Model initialization
2 AddSubstance('purple');
3 AddSubstance('light-blue');
4 CreateAgentsRandom(number, 'red');
5 CreateAgentsRandom(number, 'blue');
6
7 // Simulation
8 for each iteration {
9     UpdateNeighborSearchIndex();
10
11     for each agent {
12         SecreteSubstance(agent);
13         FollowSubstanceGradient(agent);
14         CalculateForces(agent);
15     }
16
17     DiffuseSubstance('purple');
18     DiffuseSubstance('light-blue');
19
20     Visualize();
21 }
```

Fundamental operations

- Iterate over agents and execute a function
- Determine neighbors

Cell sorting implementation

```
1 // Model initialization
2 AddSubstance('purple');
3 AddSubstance('light-blue');
4 CreateAgentsRandom(number, 'red');
5 CreateAgentsRandom(number, 'blue');
6
7 // Simulation
8 for each iteration {
9     UpdateNeighborSearchIndex();
10
11    for each agent {
12        SecreteSubstance(agent);
13        FollowSubstanceGradient(agent);
14        CalculateForces(agent);
15    }
16
17    DiffuseSubstance('purple');
18    DiffuseSubstance('light-blue');
19
20    Visualize();
21 }
```

Fundamental operations

- Iterate over agents and execute a function
- Determine neighbors
- **Iterate over neighbors and execute a function**

Cell sorting implementation

```
1 // Model initialization
2 AddSubstance('purple');
3 AddSubstance('light-blue');
4 CreateAgentsRandom(number, 'red');
5 CreateAgentsRandom(number, 'blue');

6
7 // Simulation
8 for each iteration {
9     UpdateNeighborSearchIndex();

10    for each agent {
11        SecreteSubstance(agent);
12        FollowSubstanceGradient(agent);
13        CalculateForces(agent);
14    }

15    DiffuseSubstance('purple');
16    DiffuseSubstance('light-blue');

17    Visualize();
18 }
```

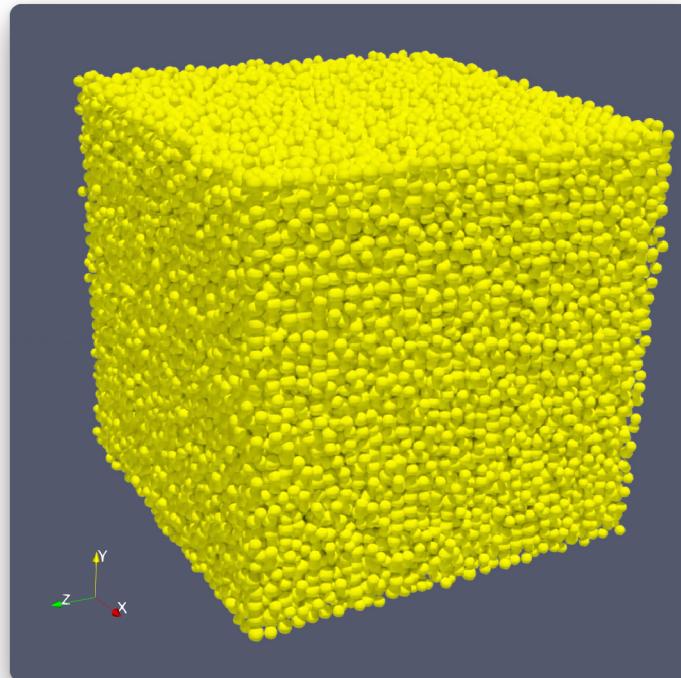
Fundamental operations

- Iterate over agents and execute a function
- Determine neighbors
- Iterate over neighbors and execute a function
- **Add and remove agents**

Benchmark simulations

The selected simulations cover a broad spectrum of performance related simulation characteristics

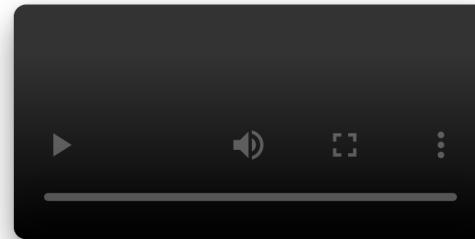
Characteristic	Cell proliferation	Cell clustering	Epidemiology	Neuroscience	Oncology
Create new agents during simulation	x			x	x
Delete agents during simulation					x
Agents modify neighbors				x	
Load imbalance			x	x	
Agents move randomly			x		x
Simulation uses diffusion	x			x	
Simulation has static regions				x	
Number of iterations	500	1000	1000	500	288
Number of agents (in millions)	12.6	2	10	9	10
Number of diffusion volumes	0	54m	0	65k	0



Benchmark simulations

The selected simulations cover a broad spectrum of performance related simulation characteristics

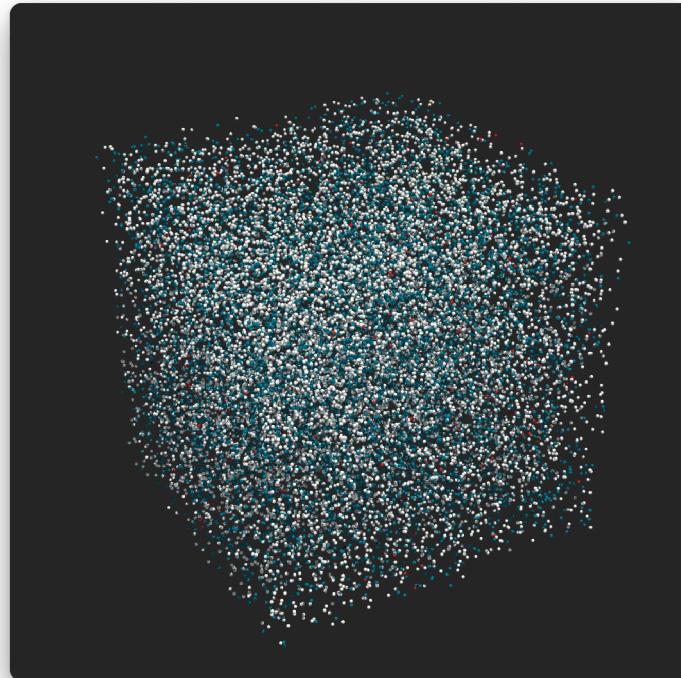
Characteristic	Cell proliferation	Cell clustering	Epidemiology	Neuroscience	Oncology
Create new agents during simulation	x			x	x
Delete agents during simulation				x	
Agents modify neighbors				x	
Load imbalance			x	x	
Agents move randomly		x			x
Simulation uses diffusion		x		x	
Simulation has static regions				x	
Number of iterations	500	1000	1000	500	288
Number of agents (in millions)	12.6	2	10	9	10
Number of diffusion volumes	0	54m	0	65k	0



Benchmark simulations

The selected simulations cover a broad spectrum of performance related simulation characteristics

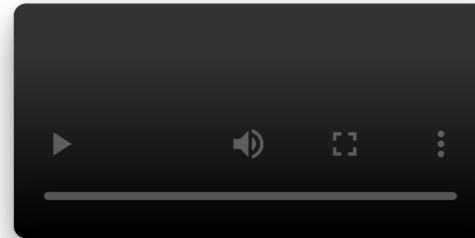
Characteristic	Cell proliferation	Cell clustering	Epidemiology	Neuroscience	Oncology
Create new agents during simulation	x			x	x
Delete agents during simulation					x
Agents modify neighbors			x	x	
Load imbalance		x	x		
Agents move randomly		x			x
Simulation uses diffusion	x			x	
Simulation has static regions				x	
Number of iterations	500	1000	1000	500	288
Number of agents (in millions)	12.6	2	10	9	10
Number of diffusion volumes	0	54m	0	65k	0



Benchmark simulations

The selected simulations cover a broad spectrum of performance related simulation characteristics

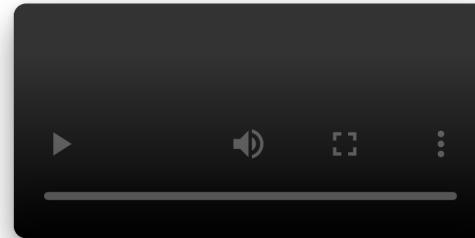
Characteristic	Cell proliferation	Cell clustering	Epidemiology	Neuroscience	Oncology
Create new agents during simulation	x		x	x	x
Delete agents during simulation				x	x
Agents modify neighbors			x	x	
Load imbalance		x	x	x	
Agents move randomly		x	x		x
Simulation uses diffusion	x		x	x	
Simulation has static regions			x	x	
Number of iterations	500	1000	1000	500	288
Number of agents (in millions)	12.6	2	10	9	10
Number of diffusion volumes	0	54m	0	65k	0



Benchmark simulations

The selected simulations cover a broad spectrum of performance related simulation characteristics

Characteristic	Cell proliferation	Cell clustering	Epidemiology	Neuroscience	Oncology
Create new agents during simulation	x			x	x
Delete agents during simulation				x	x
Agents modify neighbors			x	x	
Load imbalance		x	x		
Agents move randomly		x	x		x
Simulation uses diffusion	x		x		
Simulation has static regions			x		
Number of iterations	500	1000	1000	500	288
Number of agents (in millions)	12.6	2	10	9	10
Number of diffusion volumes	0	54m	0	65k	0



More details about the simulations and BioDynaMo's modeling features

<https://doi.org/10.1093/bioinformatics/btab649>

Bioinformatics, 38(2), 2022, 453–460
doi: 10.1093/bioinformatics/btab649
Advance Access Publication Date: 16 September 2021
Original Paper

OXFORD

Systems biology
BioDynaMo: a modular platform for high-performance agent-based simulation

Lukas Breitwieser  ^{1,2,*}, Ahmad Hesam ^{1,3,*}, Jean de Montigny ¹,
Vasileios Vavourakis ^{4,5}, Alexandros Iosif ⁴, Jack Jennings ⁶, Marcus Kaiser ^{6,7,8},
Marco Manca  ⁹, Alberto Di Meglio ¹, Zaid Al-Ars ³, Fons Rademakers ¹,
Onur Mutlu ^{2,10,*} and Roman Bauer ^{11,*}

¹CERN openlab, IT Department, CERN, Geneva 1211, Switzerland, ²Department of Computer Science, ETH Zurich, Zurich 8092, Switzerland, ³Department of Quantum & Computer Engineering, Delft University of Technology, Delft 2628CD, The Netherlands, ⁴Department of Mechanical & Manufacturing Engineering, University of Cyprus, Nicosia 2109, Cyprus, ⁵Department of Medical Physics & Biomedical Engineering, University College London, London WC1E 6BT, UK, ⁶School of Computing, Newcastle University, Newcastle upon Tyne NE4 5TG, UK, ⁷Department of Functional Neurosurgery, Ruijin Hospital, Shanghai Jiao Tong University School of Medicine, Shanghai 200025, China, ⁸Precision Imaging Beacon, School of Medicine, University of Nottingham, Nottingham NG7 2UH, UK, ⁹SCImPulse Foundation, Geleen 6162 BC, The Netherlands, ¹⁰Department of Information Technology and Electrical Engineering, ETH Zurich, Zurich 8092, Switzerland and ¹¹Department of Computer Science, University of Surrey, Guildford GU2 7XH, UK

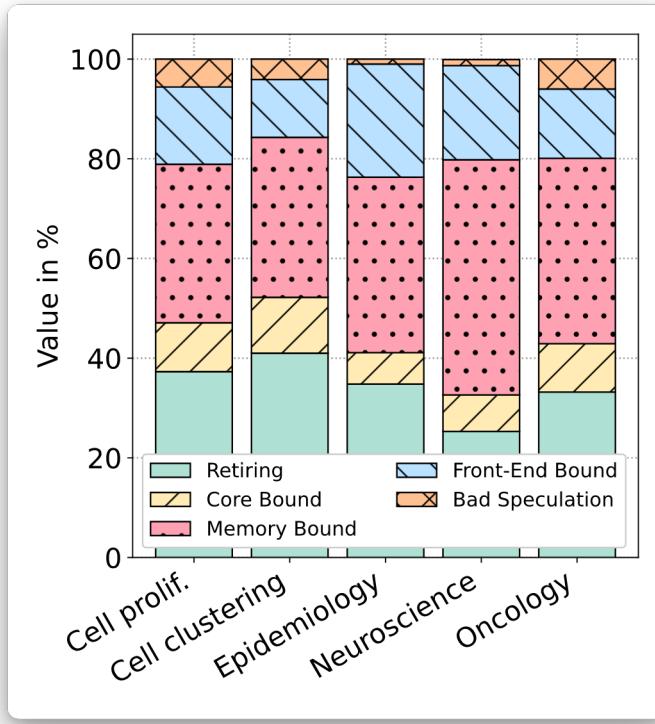
*To whom correspondence should be addressed.
Associate Editor: Jonathan Wren

Received on February 4, 2021; revised on September 2, 2021; editorial decision on September 3, 2021; accepted on September 13, 2021

Optimizations

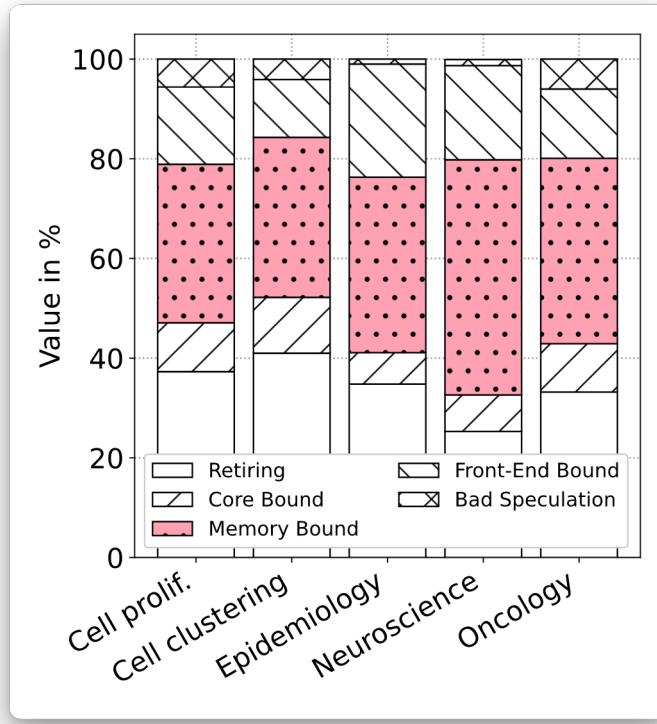
Key observation: Agent-based simulations are often **memory-bound**

Microarchitecture analysis for five simulations



Key observation: Agent-based simulations are often **memory-bound**

Microarchitecture analysis for five simulations



Key challenges and improvements

Improved fundamental operations:

- Determine neighbors
- Add and remove agents
- Iterate over agents and execute a function
- Iterate over neighbors and execute a function

Key challenges and improvements



I: Maximize parallelization

- Optimize radial neighbor search
- Parallelize the addition and removal of agents

Improved fundamental operations:

- Determine neighbors
- Add and remove agents
- Iterate over agents and execute a function
- Iterate over neighbors and execute a function

Key challenges and improvements



I: Maximize parallelization

- Optimize radial neighbor search
- Parallelize the addition and removal of agents



II: Memory layout

- NUMA-aware iteration
- Agent sorting and balancing
- Pool-based memory allocator

Improved fundamental operations:

- Determine neighbors
- Add and remove agents
- Iterate over agents and execute a function
- Iterate over neighbors and execute a function

Key challenges and improvements



I: Maximize parallelization

- Optimize radial neighbor search
- Parallelize the addition and removal of agents



II: Memory layout

- NUMA-aware iteration
- Agent sorting and balancing
- Pool-based memory allocator



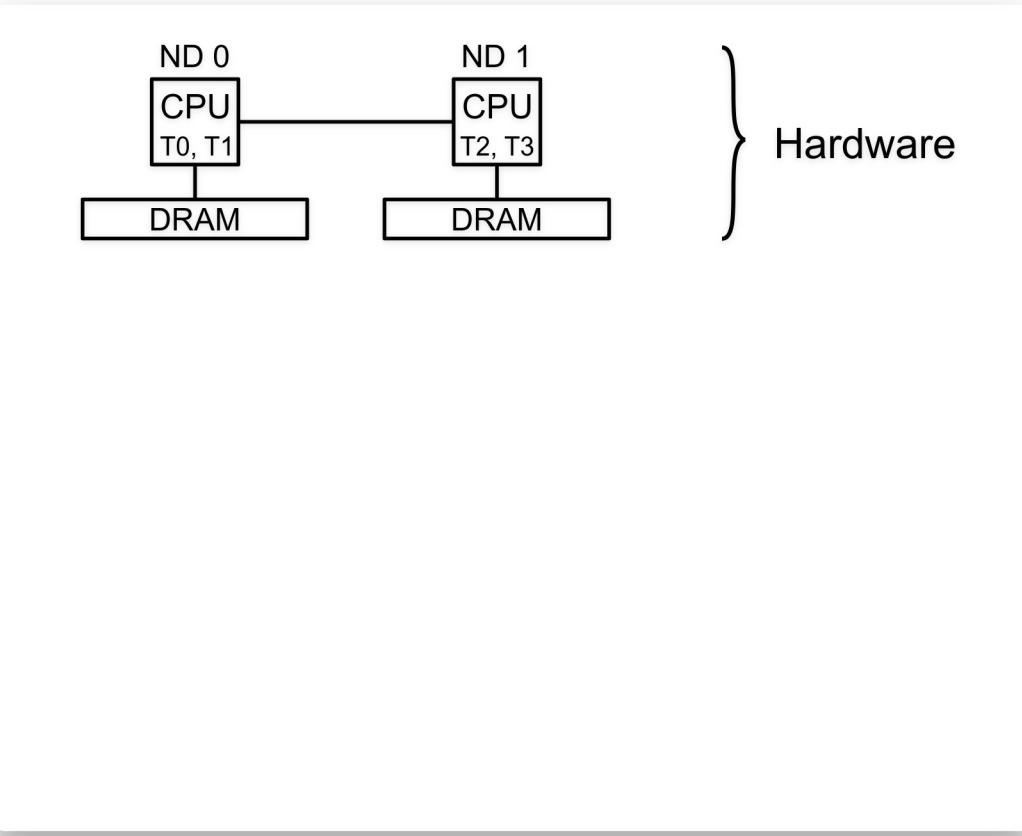
III: Avoid unnecessary work

- Pair-wise force calculation for static regions

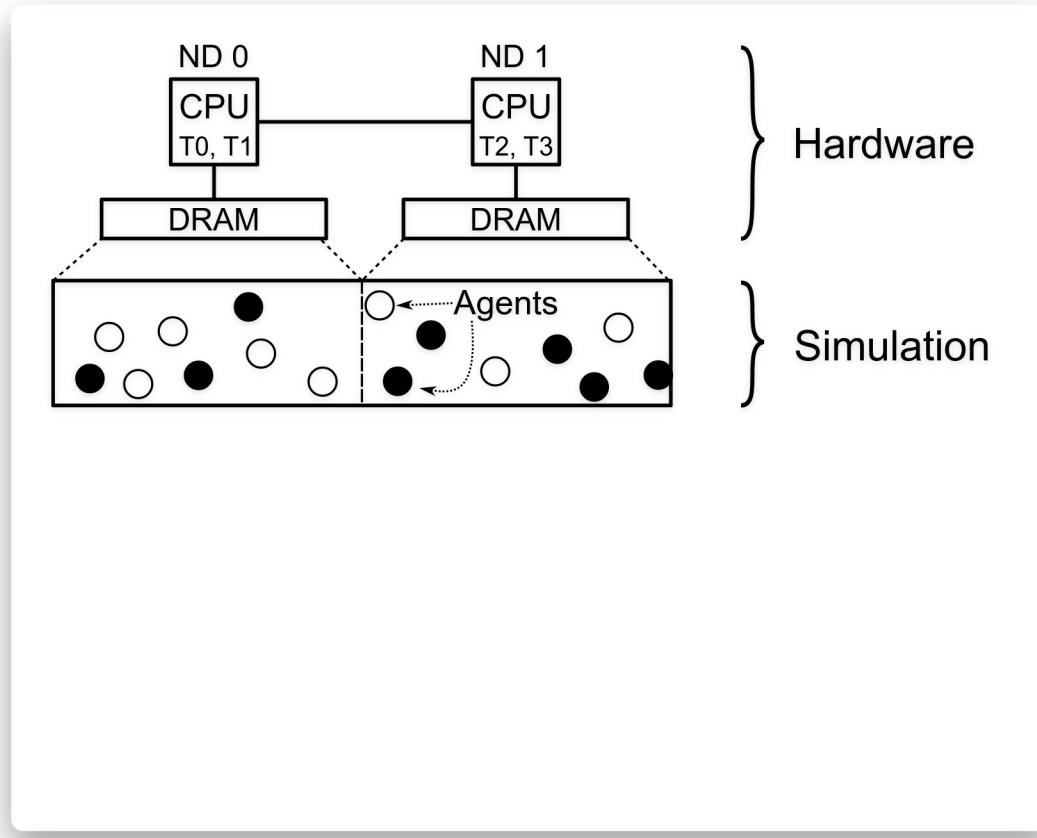
Improved fundamental operations:

- Determine neighbors
- Add and remove agents
- Iterate over agents and execute a function
- Iterate over neighbors and execute a function

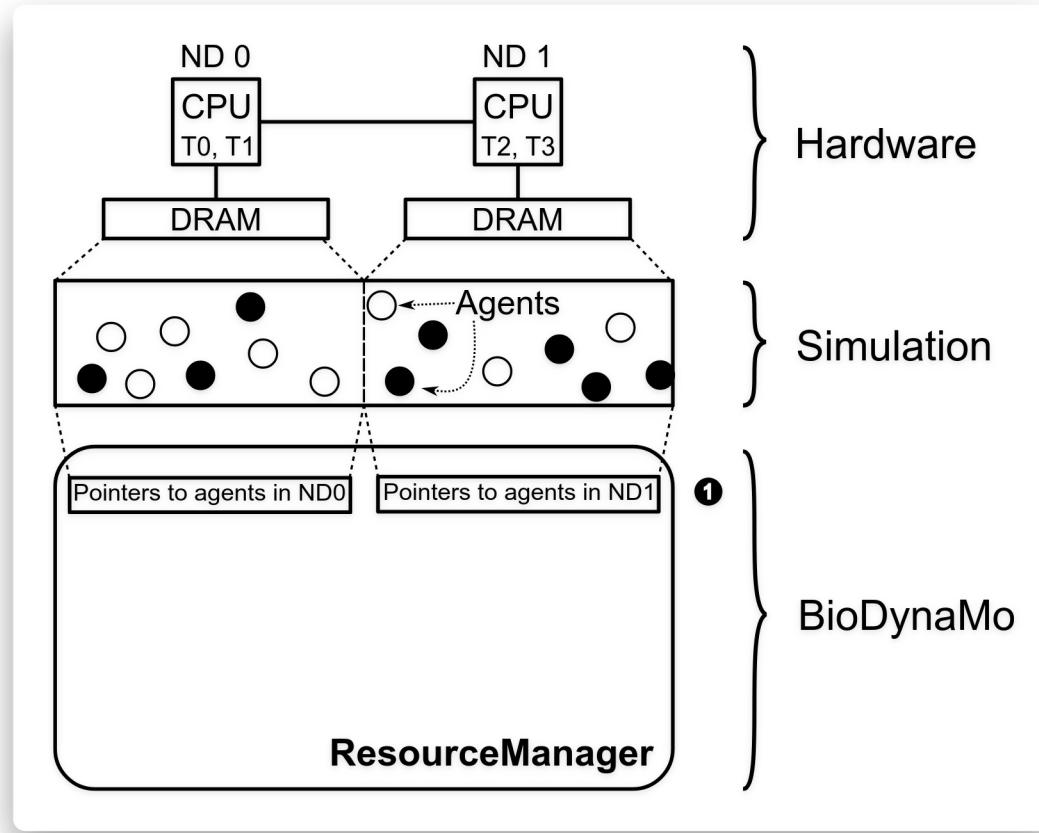
NUMA-aware iteration



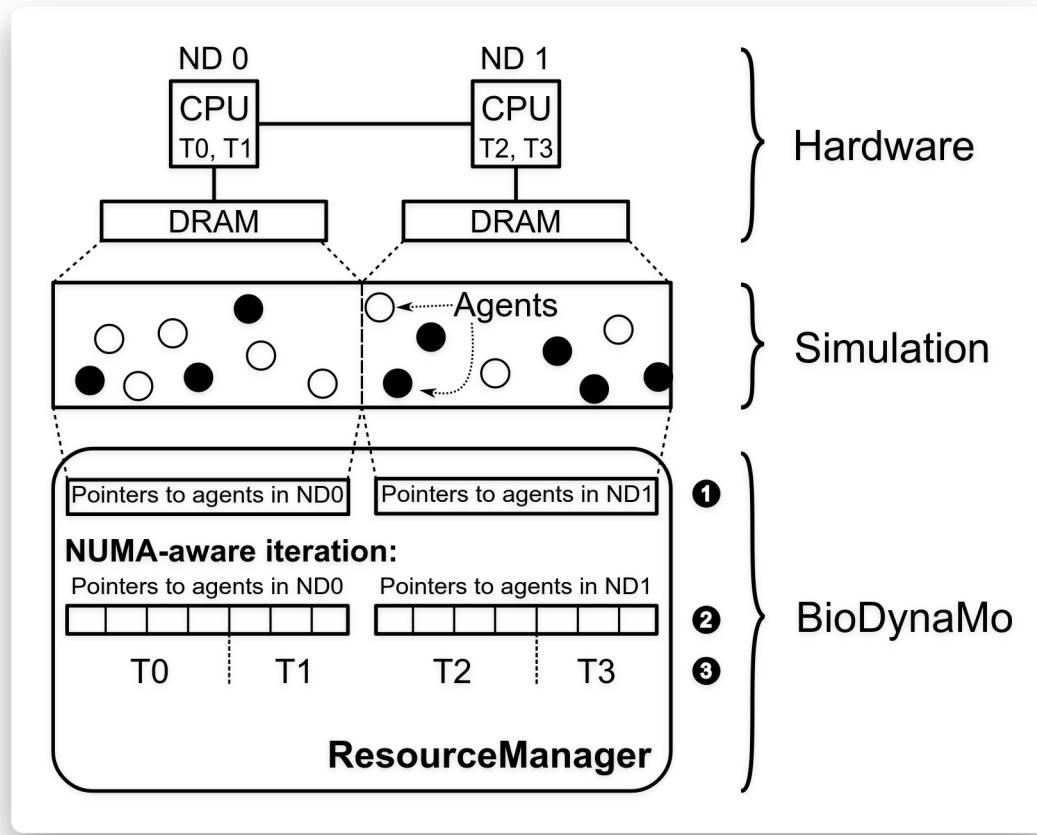
NUMA-aware iteration



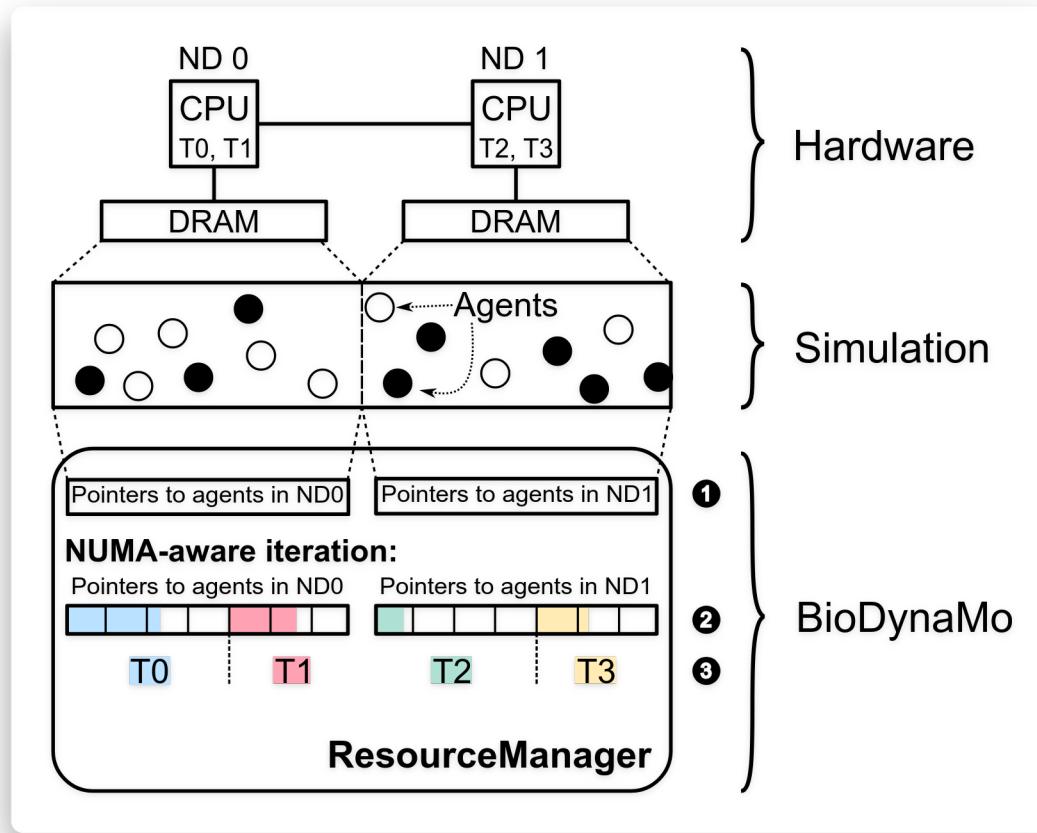
NUMA-aware iteration



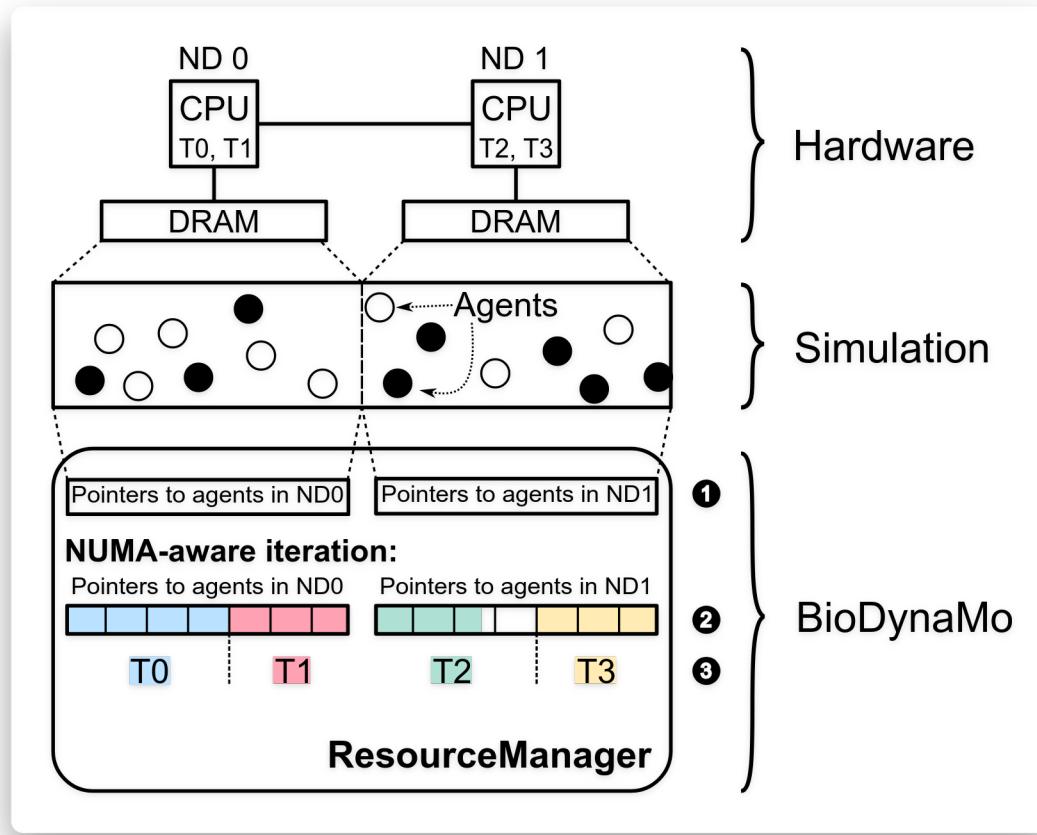
NUMA-aware iteration



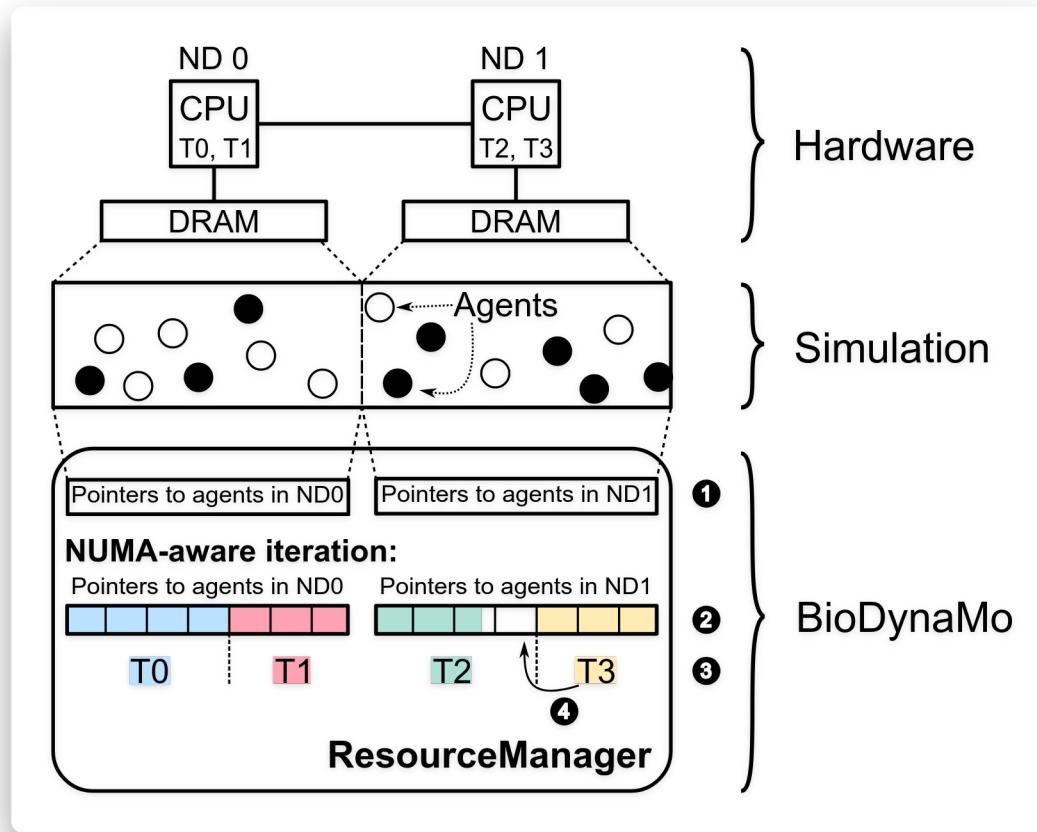
NUMA-aware iteration



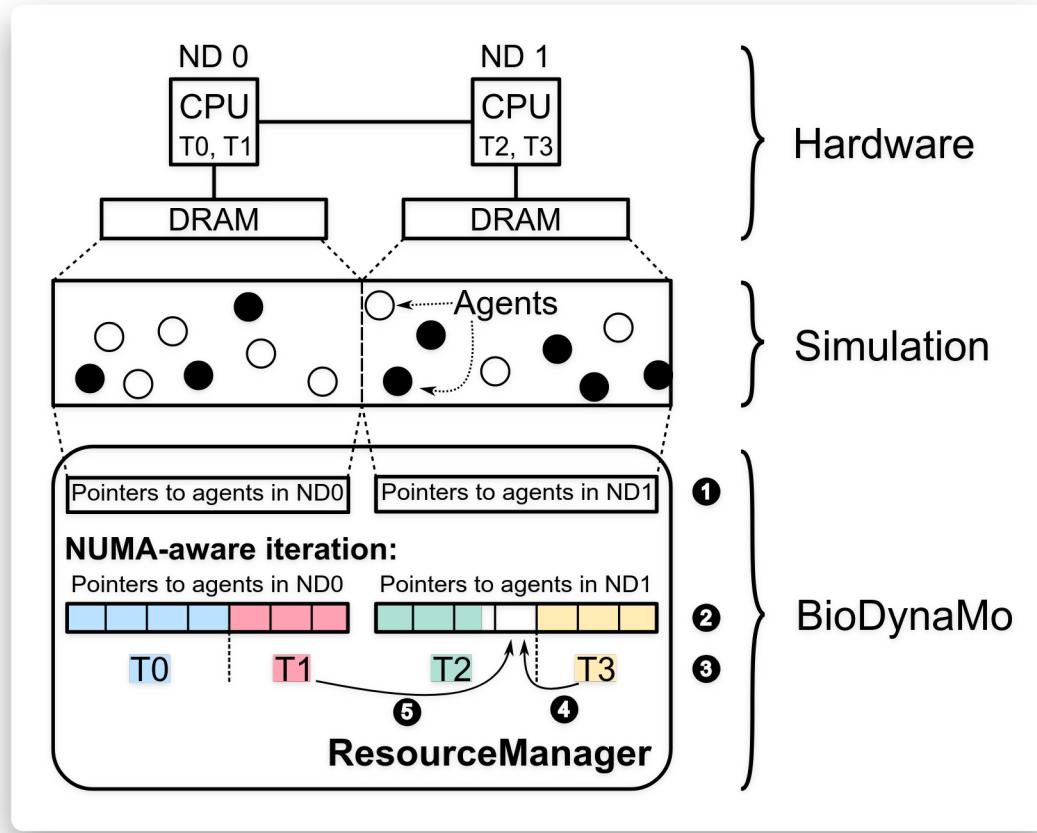
NUMA-aware iteration



NUMA-aware iteration



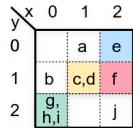
NUMA-aware iteration



NUMA-aware iteration

Agent sorting and balancing

A Agents in 3x3 grid



Agent sorting and balancing

A Agents in 3x3 grid

x	0	1	2
0	a	e	
1	b	c,d	f
2	g, h,i		j

B Grid box indices

x	0	1	2
0	0	1	2
1	3	4	5
2	6	7	8

Agent sorting and balancing

A Agents in 3x3 grid

x	0	1	2
0	a	e	
1	b	c,d	f
2	g,h,i		j

B Grid box indices

x	0	1	2
0	0	1	2
1	3	4	5
2	6	7	8

C Morton order of 4x4 grid

x	0	1	2	3
0	0	1	4	5
1	2	3	6	7
2	8	9	12	13
3	10	11	14	15

Agent sorting and balancing

A Agents in 3x3 grid

x	0	1	2
0	a	e	
1	b	c,d	f
2	g		j

B Grid box indices

x	0	1	2
0	0	1	2
1	3	4	5
2	6	7	8

C Morton order of 4x4 grid

x	0	1	2	3
0	0	1	4	5
1	2	3	6	7
2	8	9	12	13
3	10	11	14	15

Agent sorting and balancing

A Agents in 3x3 grid

x	0	1	2
0	a	e	
1	b	c,d	f
2	g,h,i		j

B Grid box indices

x	0	1	2
0	0	1	2
1	3	4	5
2	6	7	8

C Morton order of 4x4 grid

x	0	1	2	3
0	0	1	2	3
1	z	3	4	5
2	8	9	12	13
3	10	11	14	15

NUMA domain 0 \leftrightarrow NUMA domain 1



Agent sorting and balancing

A Agents in 3x3 grid

x	0	1	2
0	a	e	
1	b	c,d	f
2	g,h,i		j

B Grid box indices

x	0	1	2
0	0	1	2
1	3	4	5
2	6	7	8

C Morton order of 4x4 grid

x	0	1	2	3
0	0	1	2	3
1	3	4	5	6
2	8	9	12	13
3	10	11	14	15

NUMA domain 0 \leftrightarrow NUMA domain 1



Agent sorting and balancing

A Agents in 3x3 grid

x	0	1	2
0	a	e	
1	b	c,d	f
2	g,h,i		j

B Grid box indices

x	0	1	2
0	0	1	2
1	3	4	5
2	6	7	8

C Morton order of 4x4 grid

x	0	1	2	3
0	0	1		
1	2	3	4	5
2	6	7	8	9
3	10	11	12	13

NUMA domain 0 \leftrightarrow NUMA domain 1



Agent sorting and balancing

A Agents in 3x3 grid

x	0	1	2
0	a	e	
1	b	c,d	f
2	g,h,i		j

B Grid box indices

x	0	1	2
0	0	1	2
1	3	4	5
2	6	7	8

C Morton order of 4x4 grid

x	0	1	2	3
0	0	1		
1	2	3	4	5
2	6	7	8	9
3	10	11	12	13

NUMA domain 0 \leftrightarrow NUMA domain 1



Agent sorting and balancing

A Agents in 3x3 grid

x	0	1	2
0	a	e	
1	b	c,d	f
2	g,h,i		j

B Grid box indices

x	0	1	2
0	0	1	2
1	3	4	5
2	6	7	8

C Morton order of 4x4 grid

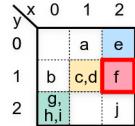
x	0	1	2	3
0	0	1	2	3
1	3	4	5	6
2	8	9	12	13
3	10	11	14	15

NUMA domain 0 \leftrightarrow NUMA domain 1

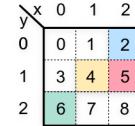


Agent sorting and balancing

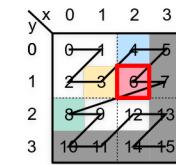
A Agents in 3x3 grid



B Grid box indices



C Morton order of 4x4 grid



NUMA domain 0 \leftrightarrow NUMA domain 1



Agent sorting and balancing

A Agents in 3x3 grid

x	0	1	2
0	a	e	
1	b	c,d	f
2	g,h,i		j

B Grid box indices

x	0	1	2
0	0	1	2
1	3	4	5
2	6	7	8

C Morton order of 4x4 grid

x	0	1	2	3
0	0	1		
1	2	3	4	5
2	6	7	8	9
3	10	11	12	13

NUMA domain 0 \leftrightarrow NUMA domain 1



Agent sorting and balancing

A Agents in 3x3 grid

x	0	1	2
0	a	e	
1	b	c,d	f
2	g,i		j

B Grid box indices

x	0	1	2
0	0	1	2
1	3	4	5
2	6	7	8

C Morton order of 4x4 grid

x	0	1	2	3
0	0	1		
1	z	3	6	
2	8	9	12	13
3	10	11	14	15

NUMA domain 0 \leftrightarrow NUMA domain 1



Agent sorting and balancing

A Agents in 3x3 grid

x	0	1	2
0	a	e	
1	b	c,d	f
2	g,h,i		j

B Grid box indices

x	0	1	2
0	0	1	2
1	3	4	5
2	6	7	8

C Morton order of 4x4 grid

x	0	1	2	3
0	0	1		
1	z	3	6	
2	8	9	12	13
3	10	11	14	15

NUMA domain 0 \leftrightarrow NUMA domain 1

a	b	c	d	e	f	g	h	i	j
---	---	---	---	---	---	---	---	---	---

Agent sorting and balancing

A Agents in 3x3 grid

x	0	1	2
0	a	e	
1	b	c,d	f
2	g,h,i		j

B Grid box indices

x	0	1	2
0	0	1	2
1	3	4	5
2	6	7	8

C Morton order of 4x4 grid

x	0	1	2	3
0	0	1	4	5
1	2	3	6	7
2	8	9	12	13
3	10	11	14	15

NUMA domain 0 \leftrightarrow NUMA domain 1



Agent sorting and balancing

A Agents in 3x3 grid

x	0	1	2
0	a	e	
1	b	c,d	f
2	g,h,i		j

B Grid box indices

x	0	1	2
0	0	1	2
1	3	4	5
2	6	7	8

C Morton order of 4x4 grid

x	0	1	2	3
0	0	1		
1	z	3	6	
2	8	9	12	13
3	10	11	14	15

NUMA domain 0 \leftrightarrow NUMA domain 1



Agent sorting and balancing

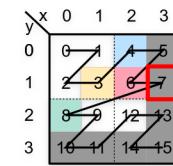
A Agents in 3x3 grid

x	0	1	2
0	a	e	
1	b	c,d	f
2	g,h,i		j

B Grid box indices

x	0	1	2
0	0	1	2
1	3	4	5
2	6	7	8

C Morton order of 4x4 grid



NUMA domain 0 \leftrightarrow NUMA domain 1



Agent sorting and balancing

A Agents in 3x3 grid

x	0	1	2
0	a	e	
1	b	c,d	f
2	g,h,i		j

B Grid box indices

x	0	1	2
0	0	1	2
1	3	4	5
2	6	7	8

C Morton order of 4x4 grid

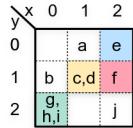
x	0	1	2	3
0	0	1		
1	z	3	6	
2	8	9	12	13
3	10	11	14	15

NUMA domain 0 \leftrightarrow NUMA domain 1

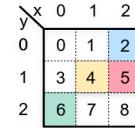
a	b	c	d	e	f	g	h	i	j
---	---	---	---	---	---	---	---	---	---

Agent sorting and balancing

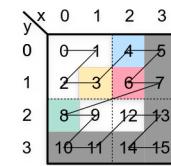
A Agents in 3x3 grid



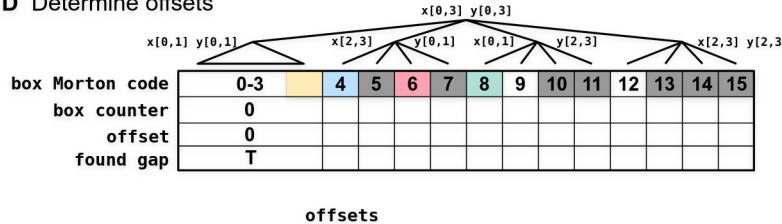
B Grid box indices



C Morton order of 4x4 grid



D Determine offsets



Agent sorting and balancing

A Agents in 3x3 grid

x	0	1	2
0	a	e	
1	b	c,d	f
2	g		j

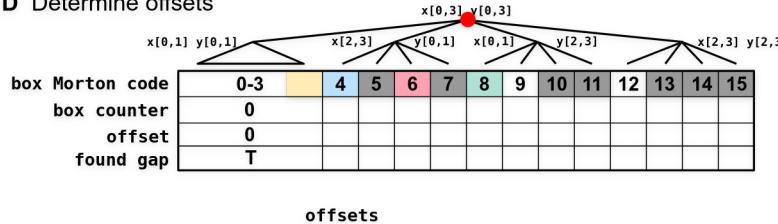
B Grid box indices

x	0	1	2
0	0	1	2
1	3	4	5
2	6	7	8

C Morton order of 4x4 grid

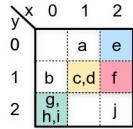
x	0	1	2	3
0	0	1	4	5
1	2	3	6	7
2	8	9	12	13
3	10	11	14	15

D Determine offsets

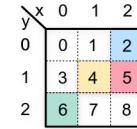


Agent sorting and balancing

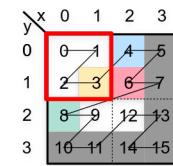
A Agents in 3x3 grid



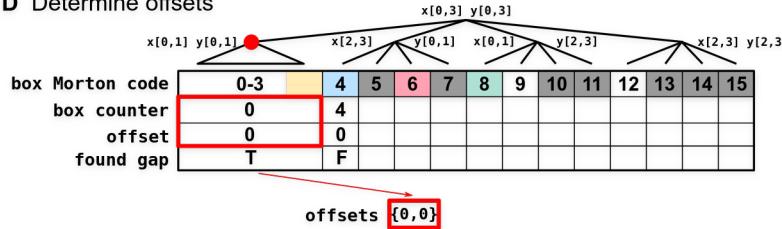
B Grid box indices



C Morton order of 4x4 grid

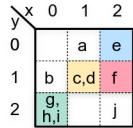


D Determine offsets

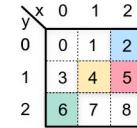


Agent sorting and balancing

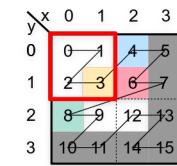
A Agents in 3x3 grid



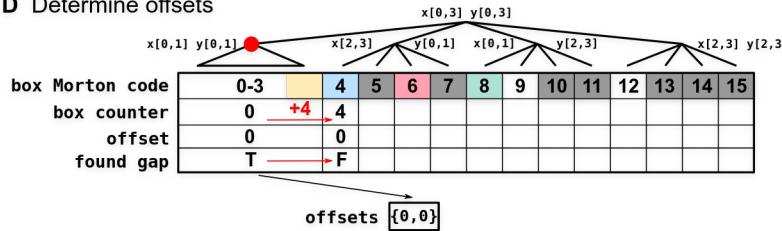
B Grid box indices



C Morton order of 4x4 grid



D Determine offsets



Agent sorting and balancing

A Agents in 3x3 grid

x	0	1	2
0	a	e	
1	b	c,d	f
2	g,h,i		j

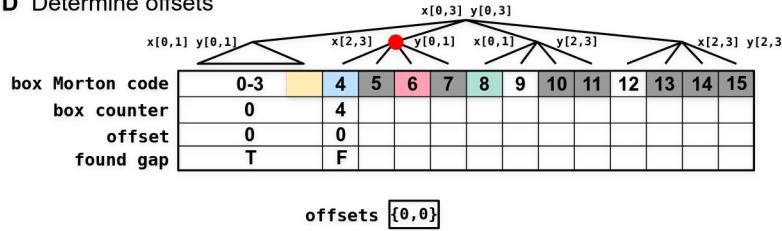
B Grid box indices

x	0	1	2
0	0	1	2
1	3	4	5
2	6	7	8

C Morton order of 4x4 grid

x	0	1	2	3
0	0	1	4	5
1	2	3	6	7
2	8	9	12	13
3	10	11	14	15

D Determine offsets



Agent sorting and balancing

A Agents in 3x3 grid

x	0	1	2
0	a	e	
1	b	c,d	f
2	g		j

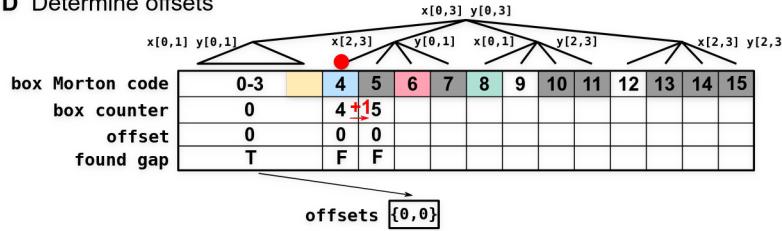
B Grid box indices

x	0	1	2
0	0	1	2
1	3	4	5
2	6	7	8

C Morton order of 4x4 grid

x	0	1	2	3
0	0	1	4	5
1	2	3	6	7
2	8	9	12	13
3	10	11	14	15

D Determine offsets



Agent sorting and balancing

A Agents in 3x3 grid

x	0	1	2
0	a	e	
1	b	c,d	f
2	g		j

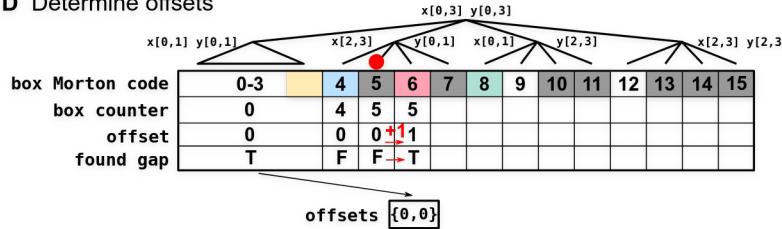
B Grid box indices

x	0	1	2
0	0	1	2
1	3	4	5
2	6	7	8

C Morton order of 4x4 grid

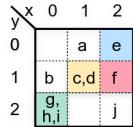
x	0	1	2	3
0	0	1	4	5
1	2	3	6	7
2	8	9	12	13
3	10	11	14	15

D Determine offsets

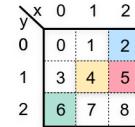


Agent sorting and balancing

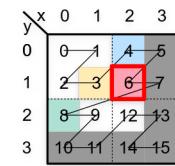
A Agents in 3x3 grid



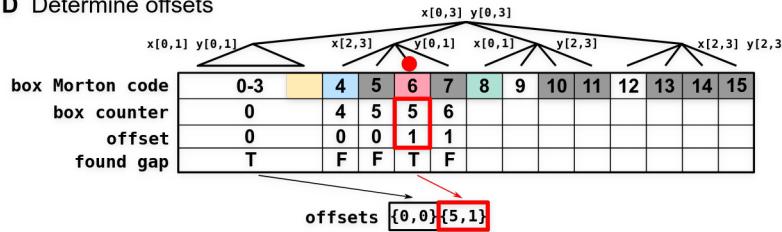
B Grid box indices



C Morton order of 4x4 grid



D Determine offsets



Agent sorting and balancing

A Agents in 3x3 grid

x	0	1	2
0	a	e	
1	b	c,d	f
2	g		j

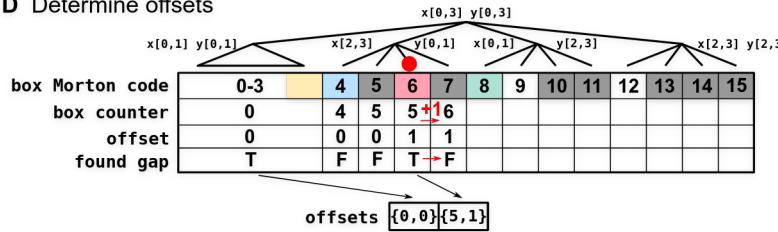
B Grid box indices

x	0	1	2
0	0	1	2
1	3	4	5
2	6	7	8

C Morton order of 4x4 grid

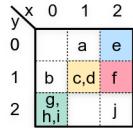
x	0	1	2	3
0	0	1	4	5
1	2	3	6	7
2	8	9	12	13
3	10	11	14	15

D Determine offsets

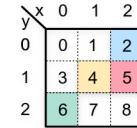


Agent sorting and balancing

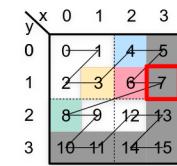
A Agents in 3x3 grid



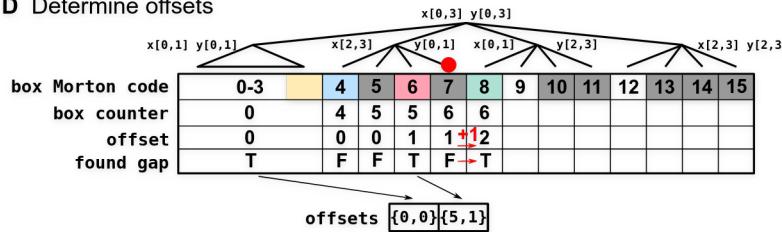
B Grid box indices



C Morton order of 4x4 grid

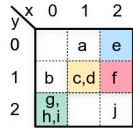


D Determine offsets

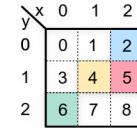


Agent sorting and balancing

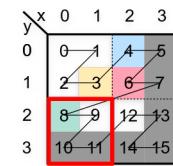
A Agents in 3x3 grid



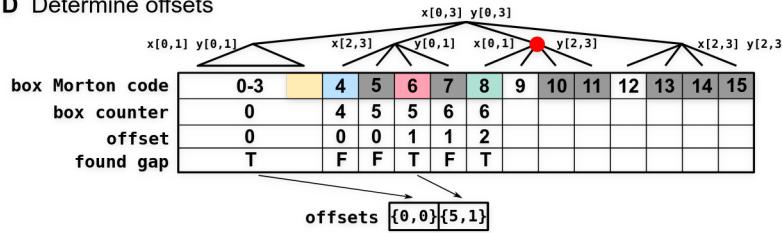
B Grid box indices



C Morton order of 4x4 grid

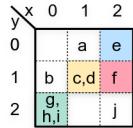


D Determine offsets

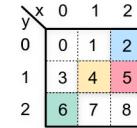


Agent sorting and balancing

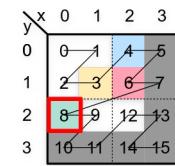
A Agents in 3x3 grid



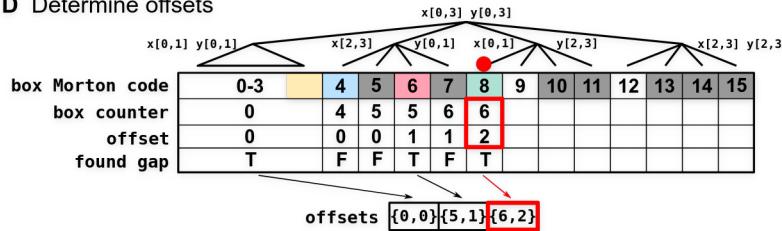
B Grid box indices



C Morton order of 4x4 grid

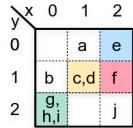


D Determine offsets

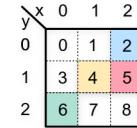


Agent sorting and balancing

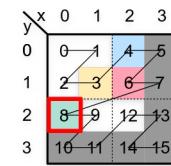
A Agents in 3x3 grid



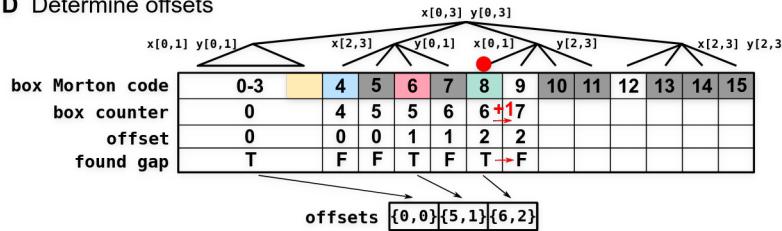
B Grid box indices



C Morton order of 4x4 grid

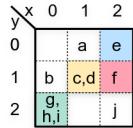


D Determine offsets

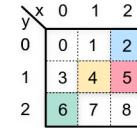


Agent sorting and balancing

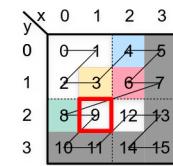
A Agents in 3x3 grid



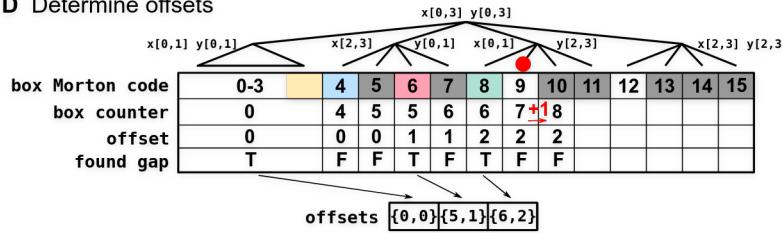
B Grid box indices



C Morton order of 4x4 grid



D Determine offsets



Agent sorting and balancing

A Agents in 3x3 grid

x	0	1	2
0	a	e	
1	b	c,d	f
2	g		j

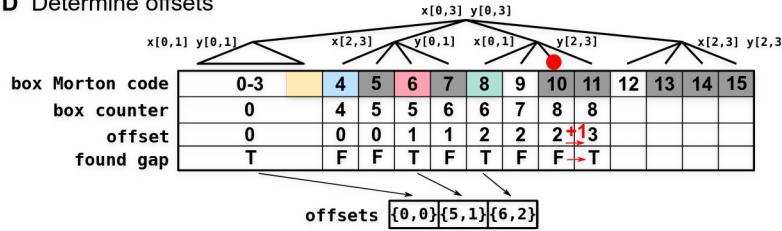
B Grid box indices

x	0	1	2
0	0	1	2
1	3	4	5
2	6	7	8

C Morton order of 4x4 grid

x	0	1	2	3
0	0	1	4	5
1	2	3	6	7
2	8	9	12	13
3	10	11	14	15

D Determine offsets



Agent sorting and balancing

A Agents in 3x3 grid

x	0	1	2
0	a	e	
1	b	c,d	f
2	g		j

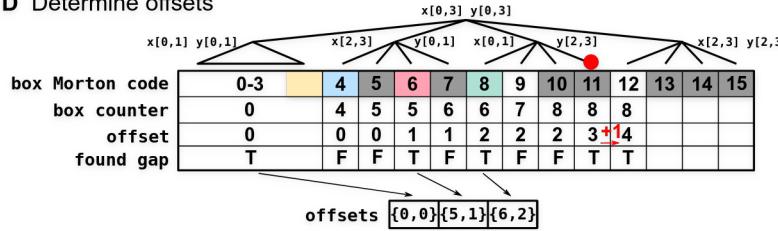
B Grid box indices

x	0	1	2
0	0	1	2
1	3	4	5
2	6	7	8

C Morton order of 4x4 grid

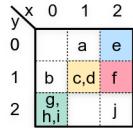
x	0	1	2	3
0	0	1	4	5
1	2	3	6	7
2	8	9	12	13
3	10	11	14	15

D Determine offsets

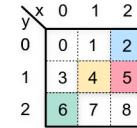


Agent sorting and balancing

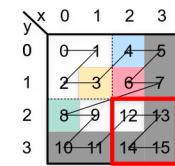
A Agents in 3x3 grid



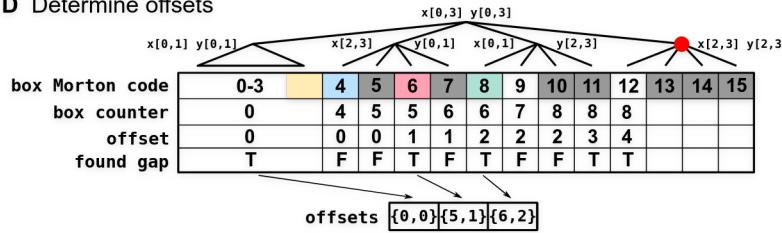
B Grid box indices



C Morton order of 4x4 grid



D Determine offsets



Agent sorting and balancing

A Agents in 3x3 grid

x	0	1	2
0	a	e	
1	b	c,d	f
2	g		j

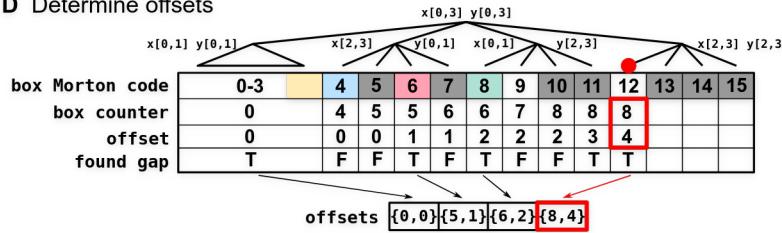
B Grid box indices

x	0	1	2
0	0	1	2
1	3	4	5
2	6	7	8

C Morton order of 4x4 grid

x	0	1	2	3
0	0	1	4	5
1	2	3	6	7
2	8	9	12	13
3	10	11	14	15

D Determine offsets



Agent sorting and balancing

A Agents in 3x3 grid

x	0	1	2
0	a	e	
1	b	c,d	f
2	g		j

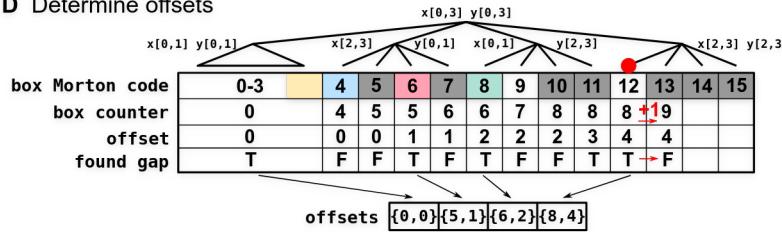
B Grid box indices

x	0	1	2
0	0	1	2
1	3	4	5
2	6	7	8

C Morton order of 4x4 grid

x	0	1	2	3
0	0	1	4	5
1	2	3	6	7
2	8	9	12	13
3	10	11	14	15

D Determine offsets



Agent sorting and balancing

A Agents in 3x3 grid

x	0	1	2
0	a	e	
1	b	c,d	f
2	g,h,i		j

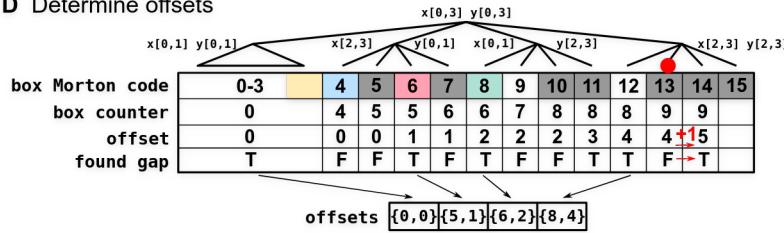
B Grid box indices

x	0	1	2
0	0	1	2
1	3	4	5
2	6	7	8

C Morton order of 4x4 grid

x	0	1	2	3
0	0	1	4	5
1	2	3	6	7
2	8	9	12	13
3	10	11	14	15

D Determine offsets



Agent sorting and balancing

A Agents in 3x3 grid

x	0	1	2
0	a	e	
1	b	c,d	f
2	g		j

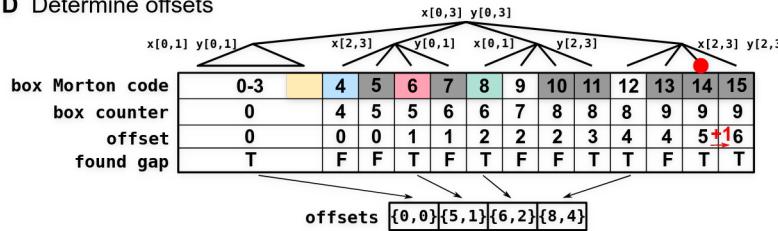
B Grid box indices

x	0	1	2
0	0	1	2
1	3	4	5
2	6	7	8

C Morton order of 4x4 grid

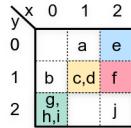
x	0	1	2	3
0	0	1	4	5
1	2	3	6	7
2	8	9	12	13
3	10	11	14	15

D Determine offsets

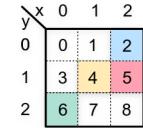


Agent sorting and balancing

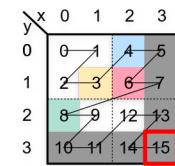
A Agents in 3x3 grid



B Grid box indices



C Morton order of 4x4 grid



D Determine offsets

	x[0,1] y[0,1]		x[2,3] y[0,1]		y[0,1] x[0,1]		y[2,3] x[2,3] y[2,3]	
box Morton code	0-3	4	5	6	7	8	9	10-15
box counter	0	4	5	5	6	6	7	8
offset	0	0	0	1	1	2	2	3
found gap	T	F	F	T	F	T	T	T

offsets {0,0}{5,1}{6,2}{8,4}

Agent sorting and balancing

A Agents in 3x3 grid

x	0	1	2
0	a	e	
1	b	c,d	f
2	g,h,i		j

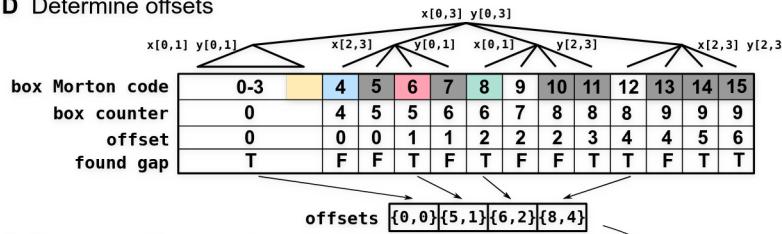
B Grid box indices

x	0	1	2
0	0	1	2
1	3	4	5
2	6	7	8

C Morton order of 4x4 grid

x	0	1	2	3
0	0	1	4	5
1	2	3	6	7
2	8	9	12	13
3	10	11	14	15

D Determine offsets



E Determine Morton order

index	0	1	2	3	4	5	6	7	8
+ offsets				0		1	2	4	
= Morton order	0	1	2	3	4	6	8	9	12

Agent sorting and balancing

A Agents in 3x3 grid

x	0	1	2
0	a	e	
1	b	c,d	f
2	g,h,i		j

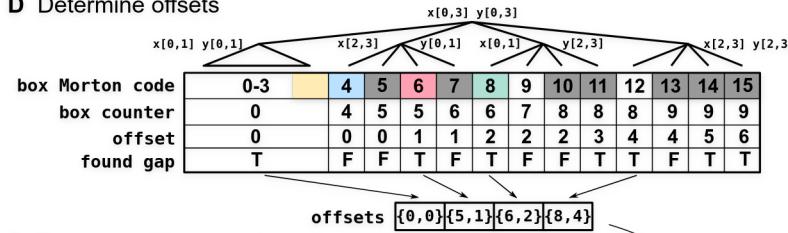
B Grid box indices

x	0	1	2
0	0	1	2
1	3	4	5
2	6	7	8

C Morton order of 4x4 grid

x	0	1	2	3
0	0	1	4	5
1	2	3	6	7
2	8	9	12	13
3	10	11	14	15

D Determine offsets

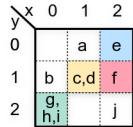


E Determine Morton order

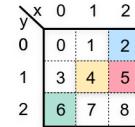
index	0	1	2	3	4	5	6	7	8
+ offsets					0		1	2	4
$= \text{Morton order}$									
0 1 2 3 4 6 8 9 12									

Agent sorting and balancing

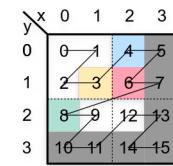
A Agents in 3x3 grid



B Grid box indices



C Morton order of 4x4 grid



D Determine offsets

		x[0,3] y[0,3]												
		x[0,1]	y[0,1]	x[2,3]	y[0,1]	x[0,1]	y[2,3]	x[2,3]	y[2,3]					
box	Morton code	0-3	4	5	6	7	8	9	10	11	12	13	14	15
box	counter	0	4	5	5	6	6	7	8	8	8	9	9	9
	offset	0	0	0	1	1	2	2	2	3	4	4	5	6
	found gap	T	F	F	T	F	T	F	T	T	F	T	T	

offsets {0,0}{5,1}{6,2}{8,4}

E Determine Morton order

index	0	1	2	3	4	5	6	7	8
+ offsets				0		1	2	4	
= Morton order									
0 1 2 3 4 6 8 9 12									

Agent sorting and balancing

A Agents in 3x3 grid

x	0	1	2
0	a	e	
1	b	c,d	f
2	g		j

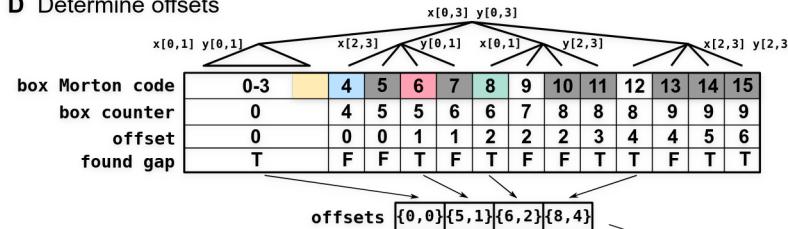
B Grid box indices

x	0	1	2
0	0	1	2
1	3	4	5
2	6	7	8

C Morton order of 4x4 grid

x	0	1	2	3
0	0	1	4	5
1	2	3	6	7
2	8	9	12	13
3	10	11	14	15

D Determine offsets



E Determine Morton order

index	0	1	2	3	4	5	6	7	8
+ offsets				0		1	2	4	
= Morton order									
0 1 2 3 4 6 8 9 12									

Agent sorting and balancing

A Agents in 3x3 grid

x	0	1	2
0	a	e	
1	b	c,d	f
2	g		j

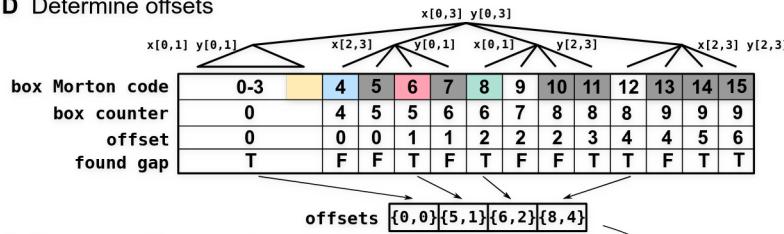
B Grid box indices

x	0	1	2
0	0	1	2
1	3	4	5
2	6	7	8

C Morton order of 4x4 grid

x	0	1	2	3
0	0	1	4	5
1	2	3	6	7
2	8	9	12	13
3	10	11	14	15

D Determine offsets



E Determine Morton order

index	0	1	2	3	4	5	6	7	8
+ offsets				0		1	2	4	
= Morton order	0	1	2	3	4	6	8	9	12

F Partition

agents per box	0	1	1	2	1	1	3	0	1
prefix sum									

Agent sorting and balancing

A Agents in 3x3 grid

x	0	1	2
0	a	e	
1	b	c,d	f
2	g		j

B Grid box indices

x	0	1	2
0	0	1	2
1	3	4	5
2	6	7	8

C Morton order of 4x4 grid

x	0	1	2	3
0	0	1	4	5
1	2	3	6	7
2	8	9	12	13
3	10	11	14	15

D Determine offsets

	x[0,3] y[0,3]						
	x[0,1]	y[0,1]	x[2,3]	y[0,1]	x[0,1]	y[2,3]	
box Morton code	0-3		4	5	6	7	8 9 10 11 12 13 14 15
box counter	0		4	5	5	6	6 7 8 8 8 9 9 9
offset	0		0	0	1	1	2 2 2 3 4 4 4 5 6
found gap	T		F	F	T	F	F T T F T T T T T T T

offsets {0,0}{5,1}{6,2}{8,4}

E Determine Morton order

index	0	1	2	3	4	5	6	7	8
+ offsets				0		1	2	4	
= Morton order	0	1	2	3	4	6	8	9	12

F Partition

agents per box	0	1	1	2	1	1	3	0	1
prefix sum									

Agent sorting and balancing

A Agents in 3x3 grid

x	0	1	2
0	a	e	
1	b	c,d	f
2	g		j
h,i			

B Grid box indices

x	0	1	2
0	0	1	2
1	3	4	5
2	6	7	8

C Morton order of 4x4 grid

x	0	1	2	3
0	0	1	4	5
1	2	3	6	7
2	8	9	12	13
3	10	11	14	15

D Determine offsets

	x[0,3] y[0,3]							
	x[0,1]	y[0,1]	x[2,3]	y[0,1]	x[0,1]	y[2,3]	x[2,3]	y[2,3]
box Morton code	0-3		4	5	6	7	8	9
box counter	0		4	5	5	6	6	7
offset	0		0	0	1	1	2	2
found gap	T		F	F	T	F	T	T

offsets {0,0}{5,1}{6,2}{8,4}

E Determine Morton order

index	0	1	2	3	4	5	6	7	8
+ offsets				0		1	2	4	
= Morton order	0	1	2	3	4	6	8	9	12

F Partition

agents per box	0	1	1	2	1	1	3	0	1
prefix sum									

Agent sorting and balancing

A Agents in 3x3 grid

x	0	1	2
0	a	e	
1	b	c,d	f
2	g,i		j

B Grid box indices

x	0	1	2
0	0	1	2
1	3	4	5
2	6	7	8

C Morton order of 4x4 grid

x	0	1	2	3
0	0	1	4	5
1	2	3	6	7
2	8	9	12	13
3	10	11	14	15

D Determine offsets

	x[0,3] y[0,3]		x[2,3] y[0,3]		y[0,1] x[0,1]		y[2,3] x[2,3]	
x[0,1] y[0,1]	0-3	4	5	6	7	8	9	10-15
box Morton code	0	4	5	5	6	6	7	8
box counter	0	4	5	5	6	6	7	9
offset	0	0	0	1	1	2	2	3
found gap	T	F	F	T	F	T	F	T

offsets {0,0}{5,1}{6,2}{8,4}

E Determine Morton order

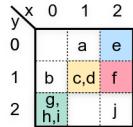
index	0	1	2	3	4	5	6	7	8
+ offsets				0		1	2	4	
= Morton order	0	1	2	3	4	6	8	9	12

F Partition

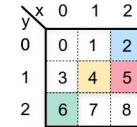
agents per box	0	1	1	2	1	1	3	0	1
prefix sum									

Agent sorting and balancing

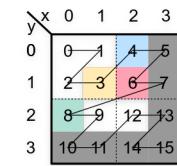
A Agents in 3x3 grid



B Grid box indices



C Morton order of 4x4 grid



D Determine offsets

		x[0,3] y[0,3]												
		x[0,1] y[0,1]	x[2,3]	y[0,1]	x[0,1]	y[2,3]	x[2,3]	y[2,3]	x[2,3]	y[0,3]				
box	Morton code	0-3	4	5	6	7	8	9	10	11	12	13	14	15
box	counter	0	4	5	5	6	6	7	8	8	8	9	9	9
offset	0	0	0	1	1	2	2	2	3	4	4	5	6	
found gap	T	F	F	T	F	T	F	F	T	T	F	T	T	

offsets {0,0}{5,1}{6,2}{8,4}

E Determine Morton order

index	0	1	2	3	4	5	6	7	8
+ offsets					0	1	2	4	
= Morton order	0	1	2	3	4	6	8	9	12

F Partition

agents per box	0	1	1	2	1	1	3	0	1
prefix sum	0	1	2	4	5	6	9	9	10

Agent sorting and balancing

A Agents in 3x3 grid

x	0	1	2
0	a	e	
1	b	c,d	f
2	g		j

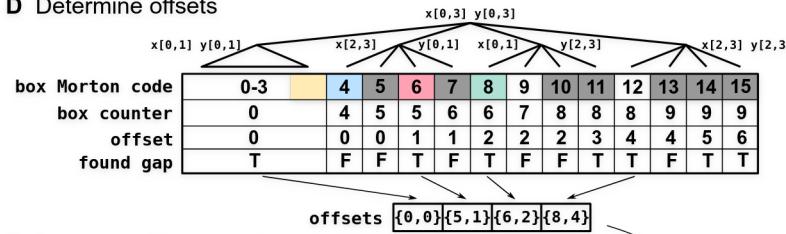
B Grid box indices

x	0	1	2
0	0	1	2
1	3	4	5
2	6	7	8

C Morton order of 4x4 grid

x	0	1	2	3
0	0	1	4	5
1	2	3	6	7
2	8	9	12	13
3	10	11	14	15

D Determine offsets



E Determine Morton order

index	0	1	2	3	4	5	6	7	8
+ offsets				0		1	2	4	
= Morton order	0	1	2	3	4	5	6	8	9 12

F Partition

agents per box	0	1	1	2	1	1	3	0	1
prefix sum	0	1	2	4	5	6	9	9	10

NUMA domain 0 ←→ NUMA domain 1

Agent sorting and balancing

A Agents in 3x3 grid

x	0	1	2
0	a	e	
1	b	c,d	f
2	g,h,i		j

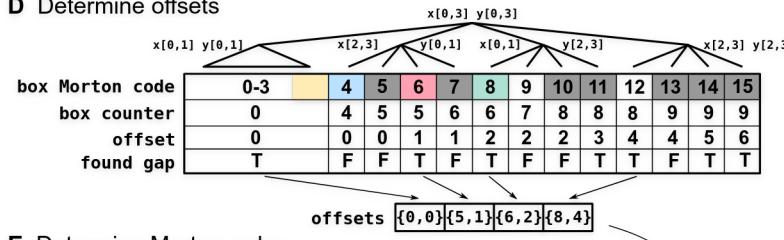
B Grid box indices

x	0	1	2
0	0	1	2
1	3	4	5
2	6	7	8

C Morton order of 4x4 grid

x	0	1	2	3
0	0	1	4	5
1	2	3	6	7
2	8	9	12	13
3	10	11	14	15

D Determine offsets



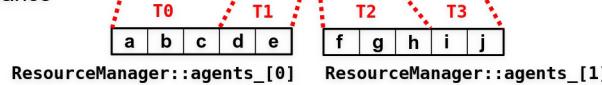
E Determine Morton order

index	0	1	2	3	4	5	6	7	8
+ offsets					0	1	2	4	
= Morton order	0	1	2	3	4	5	6	8	9 12

F Partition

agents per box	0	1	1	2	1	1	3	0	1
prefix sum	0	1	2	4	5	6	9	9	10

G Sort and balance



Agent sorting and balancing

A Agents in 3x3 grid

x	0	1	2
0	a	e	
1	b	c,d	f
2	g		j

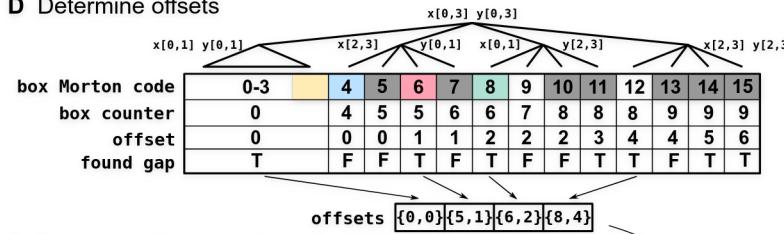
B Grid box indices

x	0	1	2
0	0	1	2
1	3	4	5
2	6	7	8

C Morton order of 4x4 grid

x	0	1	2	3
0	0	1	4	5
1	2	3	6	7
2	8	9	12	13
3	10	11	14	15

D Determine offsets



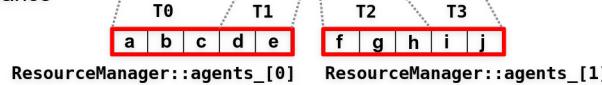
E Determine Morton order

index	0	1	2	3	4	5	6	7	8
+ offsets					0	1	2	4	
= Morton order	0	1	2	3	4	5	6	8	9 12

F Partition

agents per box	0	1	1	2	1	1	3	0	1
prefix sum	0	1	2	4	5	6	9	9	10

G Sort and balance



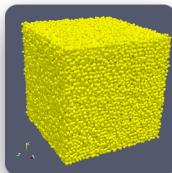
Evaluation

Benchmark hardware

System	Memory	CPU	OS
A	504 GB	Four Intel(R) Xeon(R) E7-8890 v3 CPUs @ 2.50GHz with a total of 72 physical cores, two threads per core and four NUMA domains.	CentOS
B	1008 GB		7.9.2009
C	62 GB	Two Intel(R) Xeon(R) E5-2683 v3 CPUs @ 2.00GHz with a total of 28 physical cores, two threads per core and two NUMA domains.	CentOS Stream 8

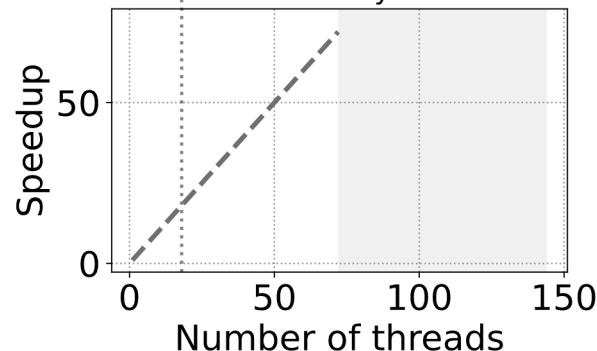
BioDynaMo scales well across NUMA nodes & large CPU core counts

Cell proliferation simulation

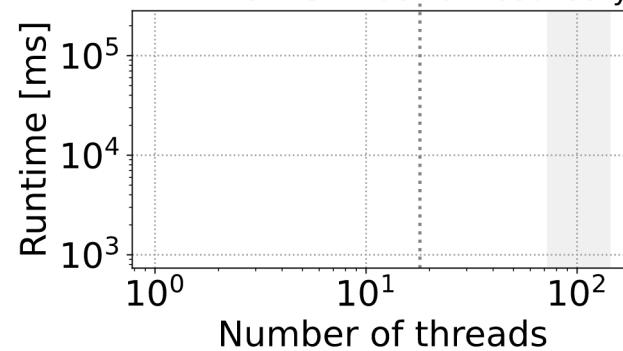


— Ideal speedup

First NUMA domain boundary

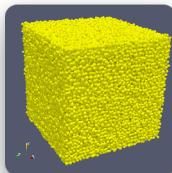


First NUMA domain boundary



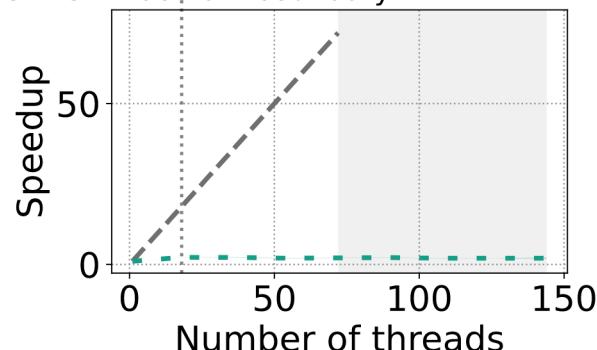
BioDynaMo scales well across NUMA nodes & large CPU core counts

Cell proliferation simulation

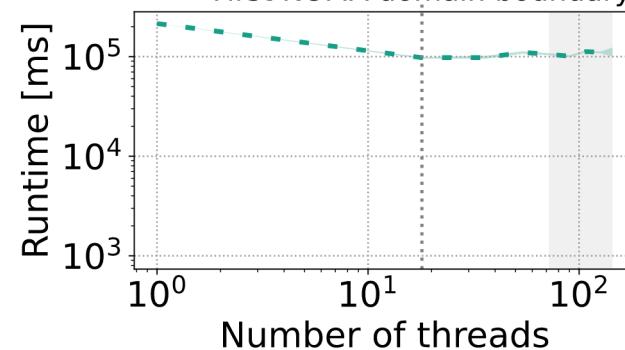


- Ideal speedup
- - BDM standard implementation

First NUMA domain boundary

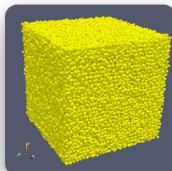


First NUMA domain boundary



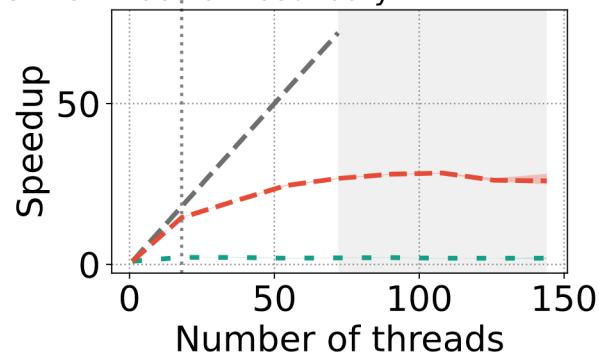
BioDynaMo scales well across NUMA nodes & large CPU core counts

Cell proliferation simulation

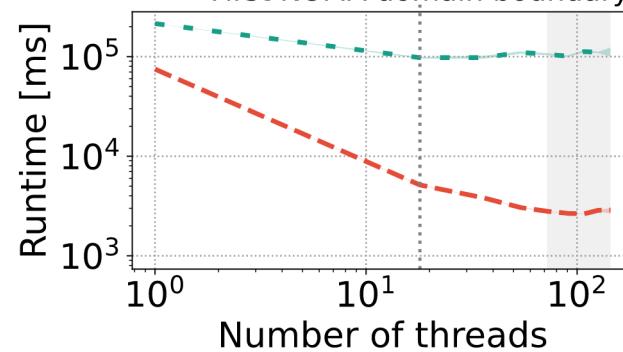


- Ideal speedup
- - BDM standard implementation
- - Plus BDM uniform grid

First NUMA domain boundary

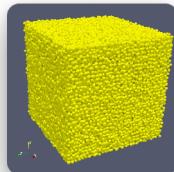


First NUMA domain boundary



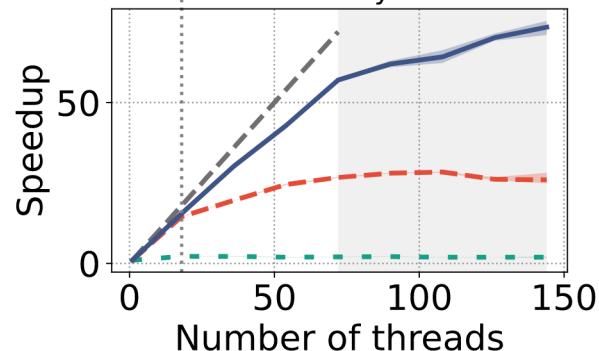
BioDynaMo scales well across NUMA nodes & large CPU core counts

Cell proliferation simulation

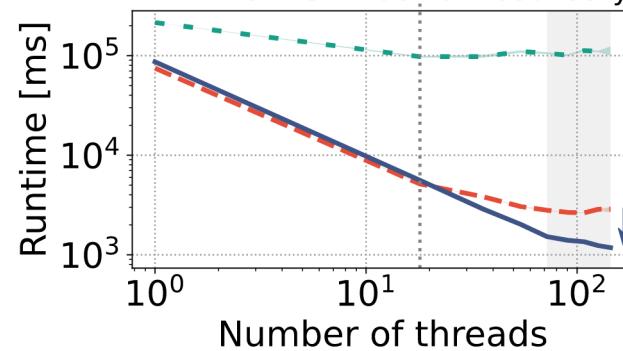


- Ideal speedup
- - BDM standard implementation
- - - Plus BDM uniform grid
- Plus memory improvements

First NUMA domain boundary

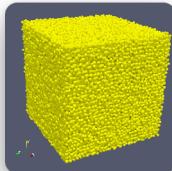


First NUMA domain boundary



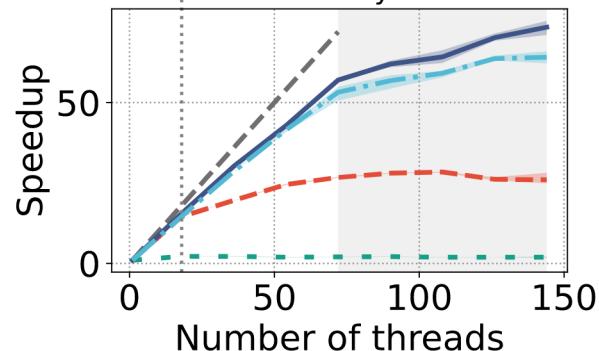
BioDynaMo scales well across NUMA nodes & large CPU core counts

Cell proliferation simulation

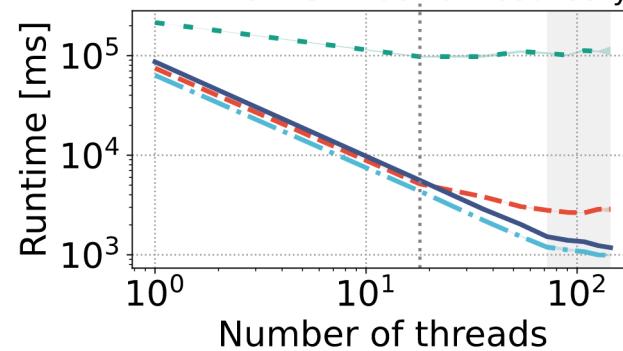


- Ideal speedup
- - BDM standard implementation
- - Plus BDM uniform grid
- Plus memory improvements
- - Plus extra memory

First NUMA domain boundary

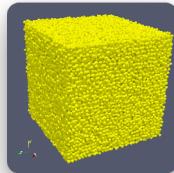


First NUMA domain boundary



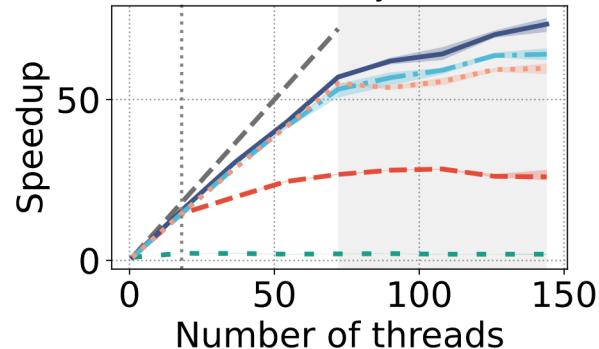
BioDynaMo scales well across NUMA nodes & large CPU core counts

Cell proliferation simulation

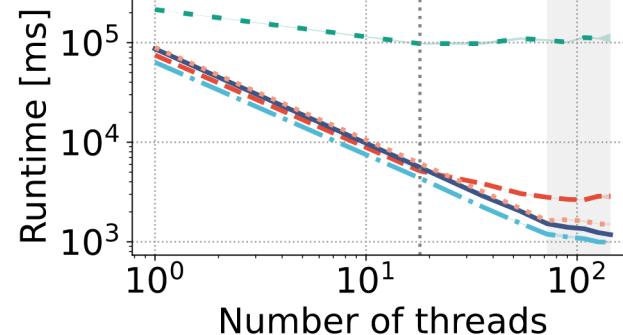


- Ideal speedup
- - - BDM standard implementation
- - - Plus BDM uniform grid
- Plus memory improvements
- - - Plus extra memory
- - - Plus static agents detection

First NUMA domain boundary

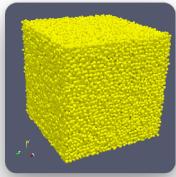


First NUMA domain boundary

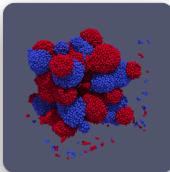


BioDynaMo scales well across NUMA nodes & large CPU core counts

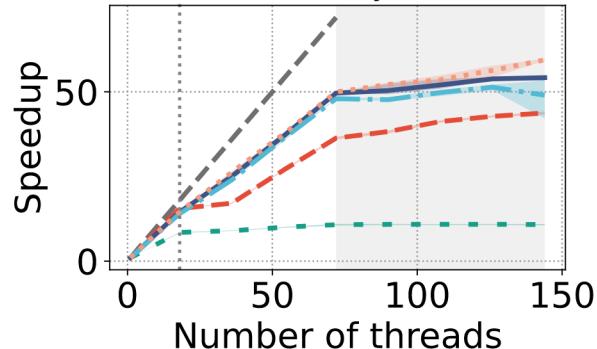
Cell proliferation simulation



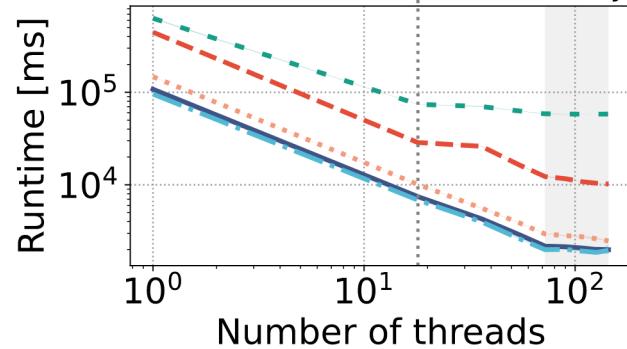
Cell clustering benefits from BioDynaMo's memory optimizations for any number of CPU cores



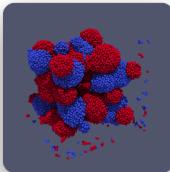
First NUMA domain boundary



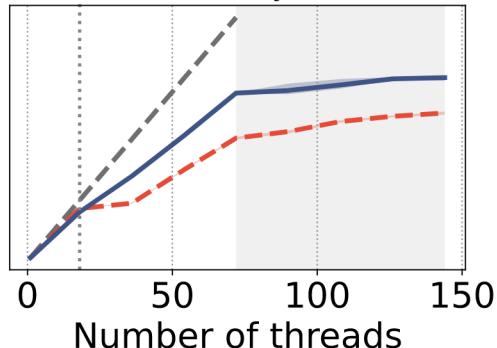
First NUMA domain boundary



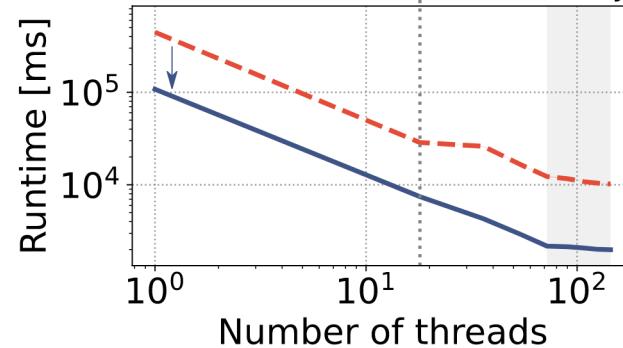
Cell clustering benefits from BioDynaMo's memory optimizations for any number of CPU cores



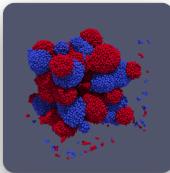
First NUMA domain boundary



First NUMA domain boundary

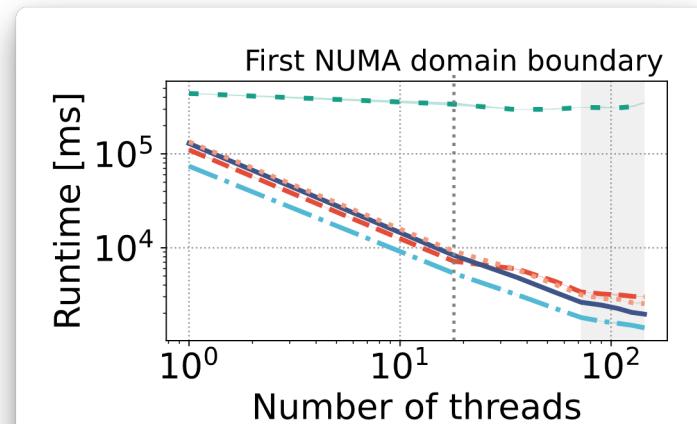
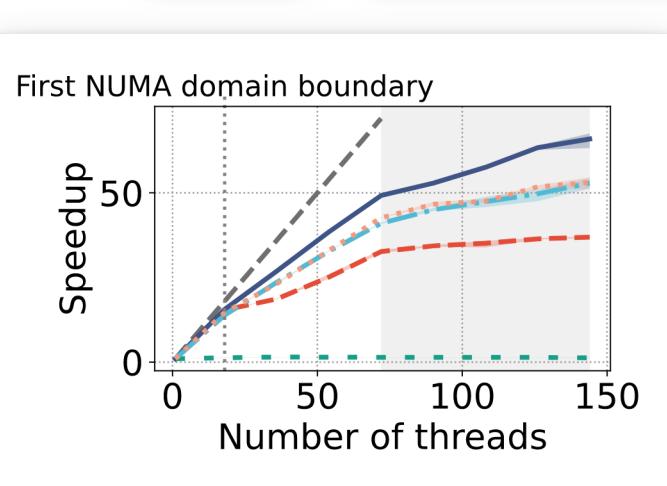
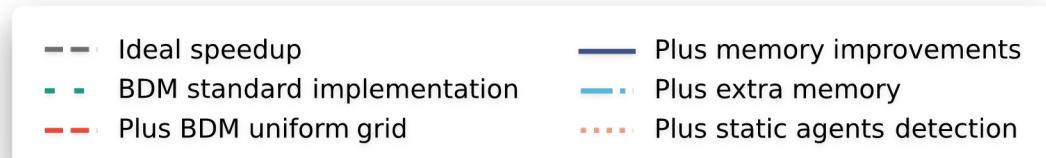


Cell clustering benefits from BioDynaMo's memory optimizations for any number of CPU cores

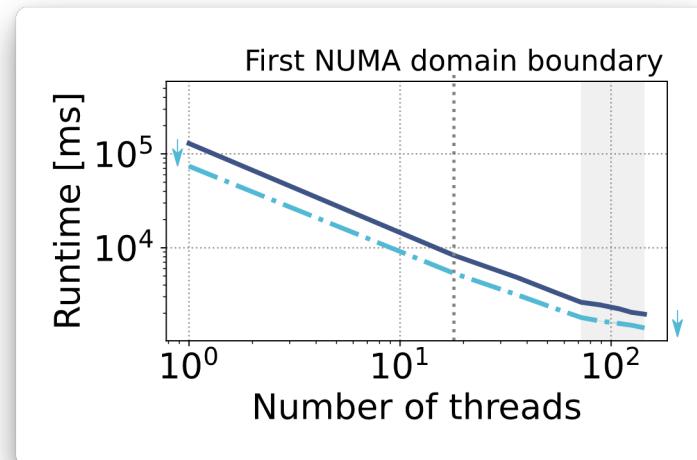
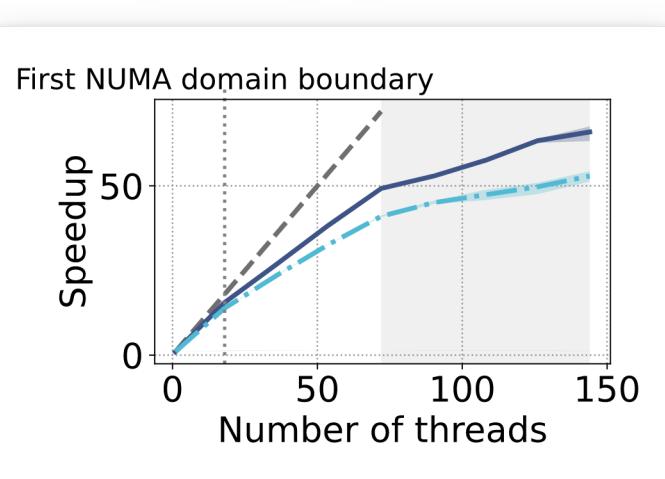
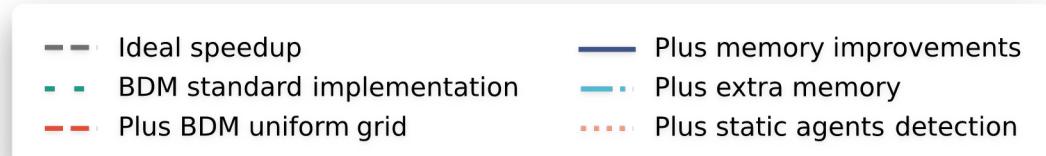


- Ideal speedup
- - BDM standard implementation
- - Plus BDM uniform grid
- Plus memory improvements
- Plus extra memory
- Plus static agents detection

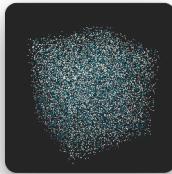
The epidemiology use case benefits from using additional memory



The epidemiology use case benefits from using additional memory



The epidemiology use case benefits from using additional memory

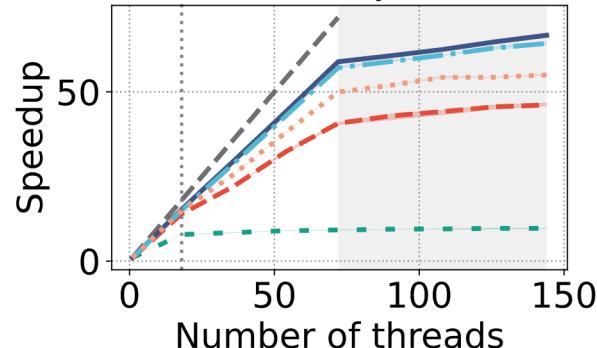


- Ideal speedup
- - BDM standard implementation
- - Plus BDM uniform grid
- Plus memory improvements
- Plus extra memory
- Plus static agents detection

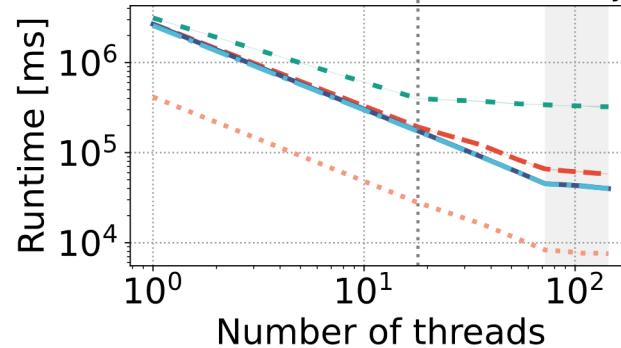
The neuroscience use case benefits substantially from static region detection



First NUMA domain boundary



First NUMA domain boundary



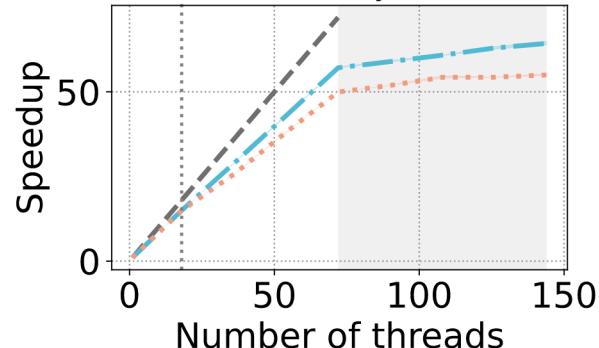
The neuroscience use case benefits substantially from static region detection



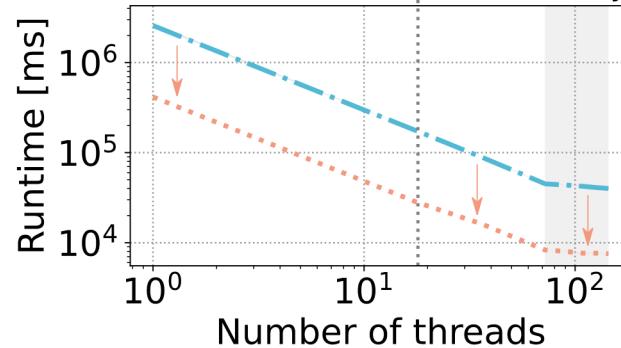
— Ideal speedup
- - BDM standard implementation
- - Plus BDM uniform grid

— Plus memory improvements
- - Plus extra memory
- - Plus static agents detection

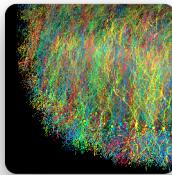
First NUMA domain boundary



First NUMA domain boundary



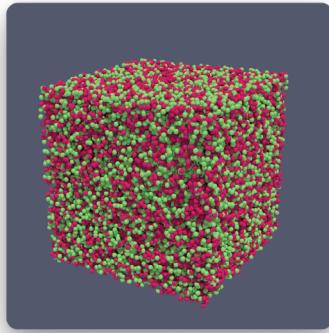
The neuroscience use case benefits substantially from static region detection



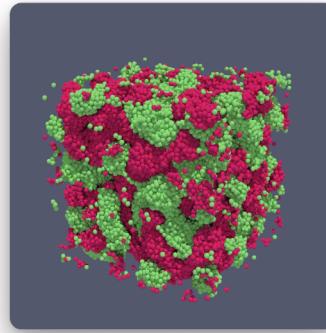
- Ideal speedup
- - BDM standard implementation
- - Plus BDM uniform grid
- Plus memory improvements
- Plus extra memory
- Plus static agents detection

BioDynaMo is up to $9\times$ more efficient than Biocellion

We implement Biocellions cell clustering simulation in BioDynaMo



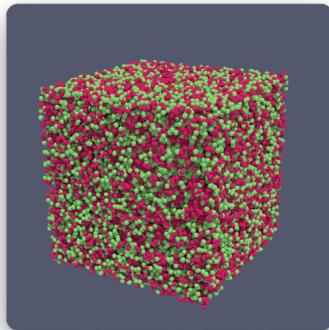
Initial state



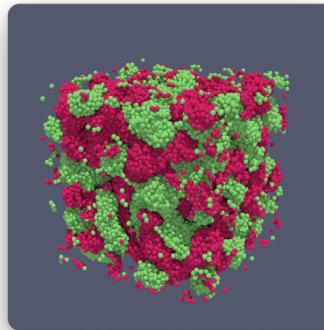
Final state

BioDynaMo is up to $9\times$ more efficient than Biocellion

We implement Biocellions cell clustering simulation in BioDynaMo



Initial state

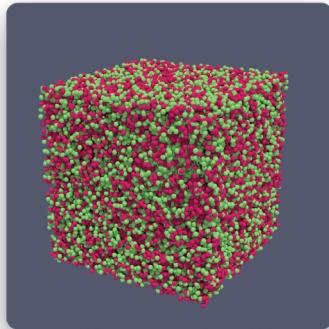


Final state

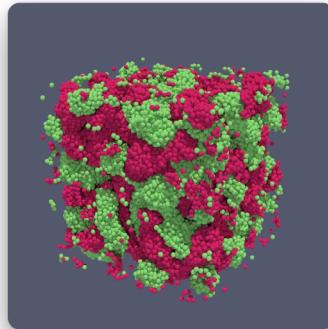
- Small-scale with **26.8 million** cells
 - 16 CPU cores (Biocellion & BioDynaMo)
 - Result: **BioDynaMo is $4.2\times$ faster**

BioDynaMo is up to $9\times$ more efficient than Biocellion

We implement Biocellions cell clustering simulation in BioDynaMo



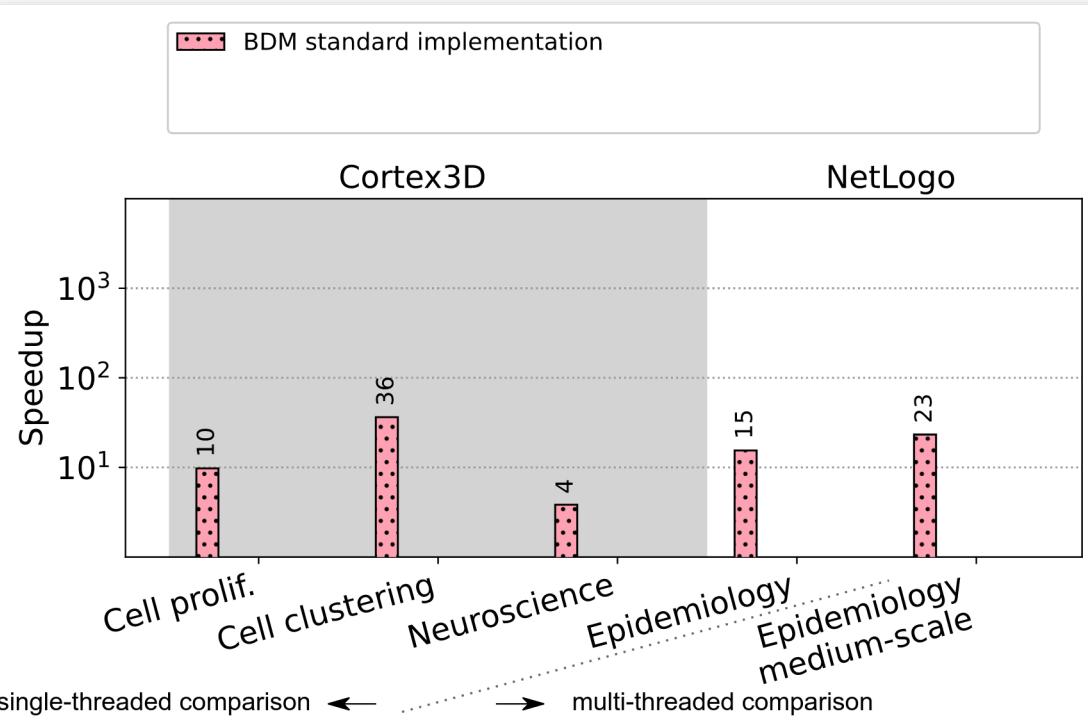
Initial state



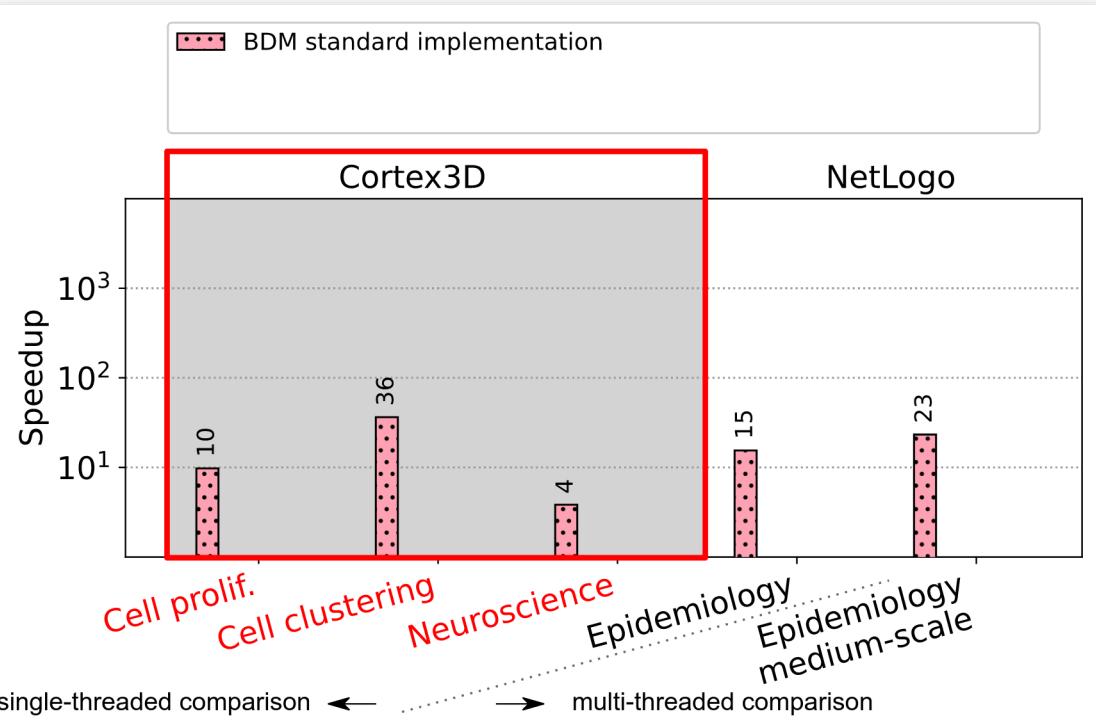
Final state

- Small-scale with **26.8 million cells**
 - 16 CPU cores (Biocellion & BioDynaMo)
 - Result: **BioDynaMo is $4.2\times$ faster**
- Large-scale with 281 million cells
 - Biocellion: 21 nodes with 672 CPU cores
 - BioDynaMo: 1 node with 72 CPU cores
 - Result: same runtime, but BioDynaMo uses **$9.3\times$ fewer CPU cores**

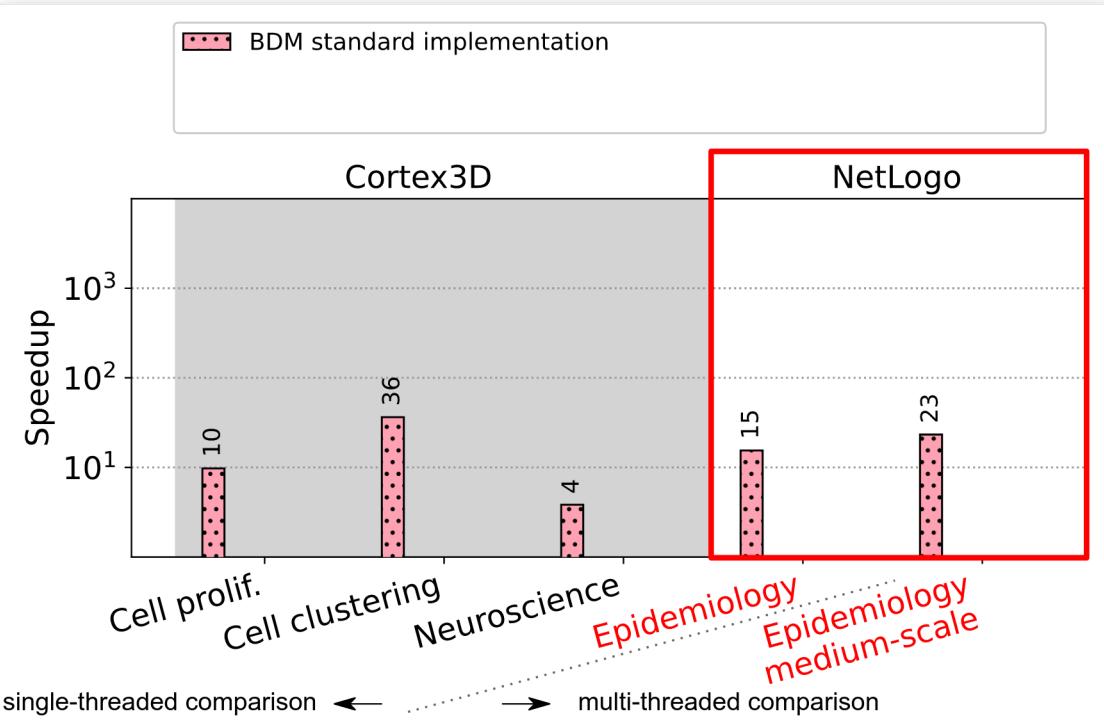
BioDynaMo is $1000\times$ faster than Cortex3D and NetLogo



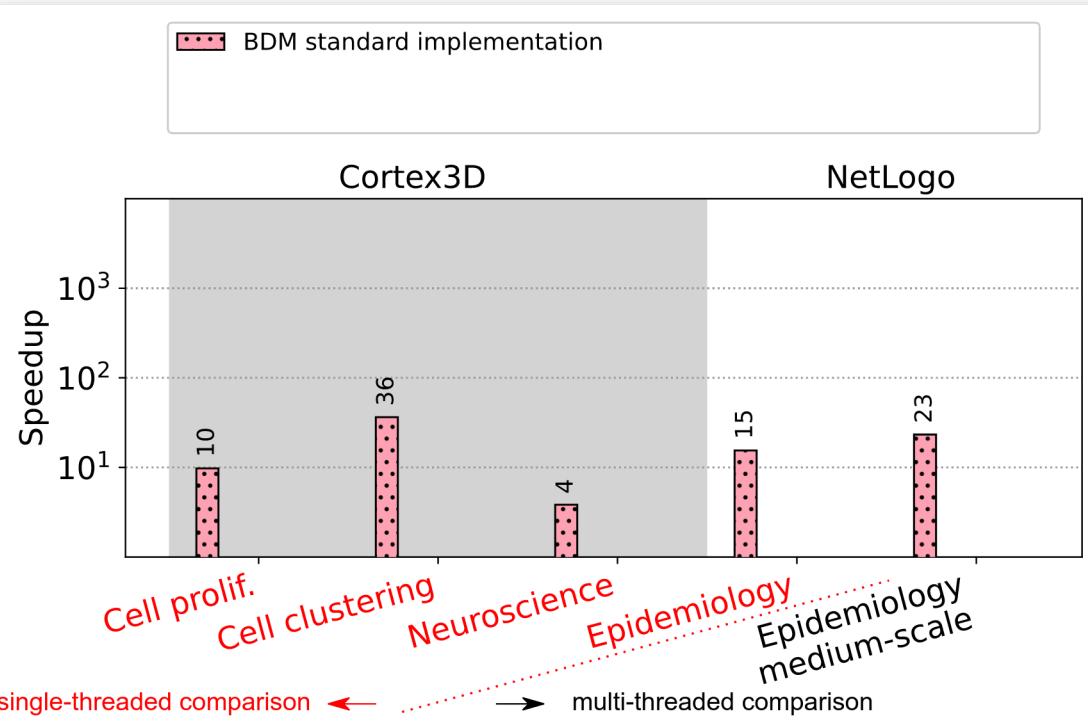
BioDynaMo is $1000\times$ faster than Cortex3D and NetLogo



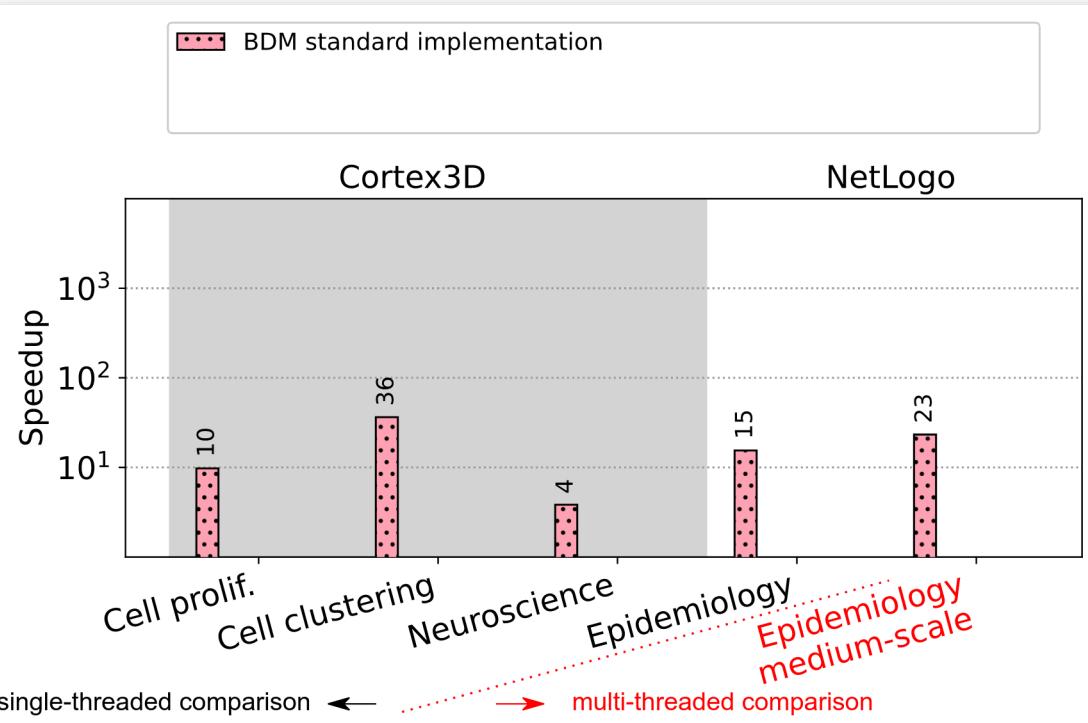
BioDynaMo is $1000\times$ faster than Cortex3D and NetLogo



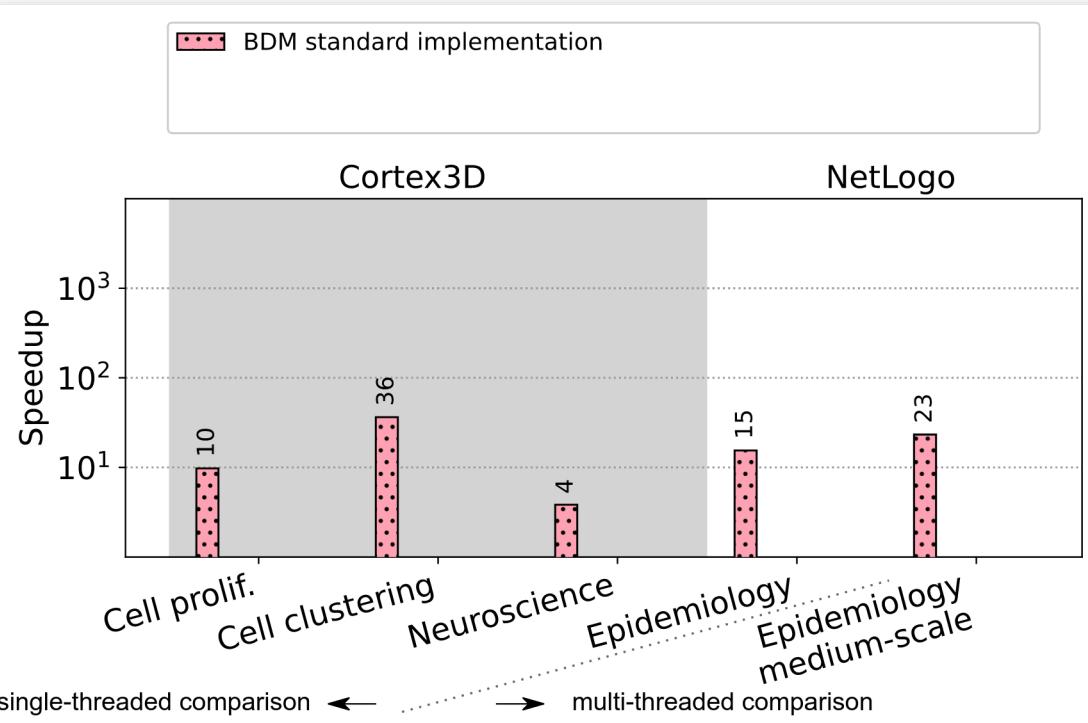
BioDynaMo is $1000\times$ faster than Cortex3D and NetLogo



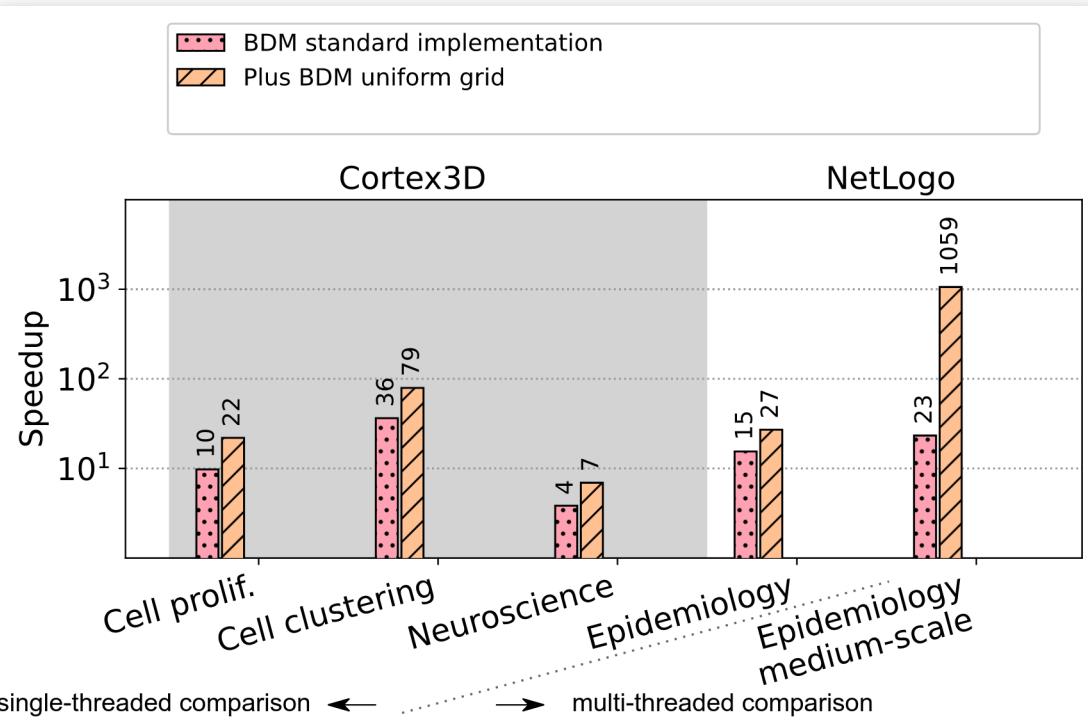
BioDynaMo is $1000\times$ faster than Cortex3D and NetLogo



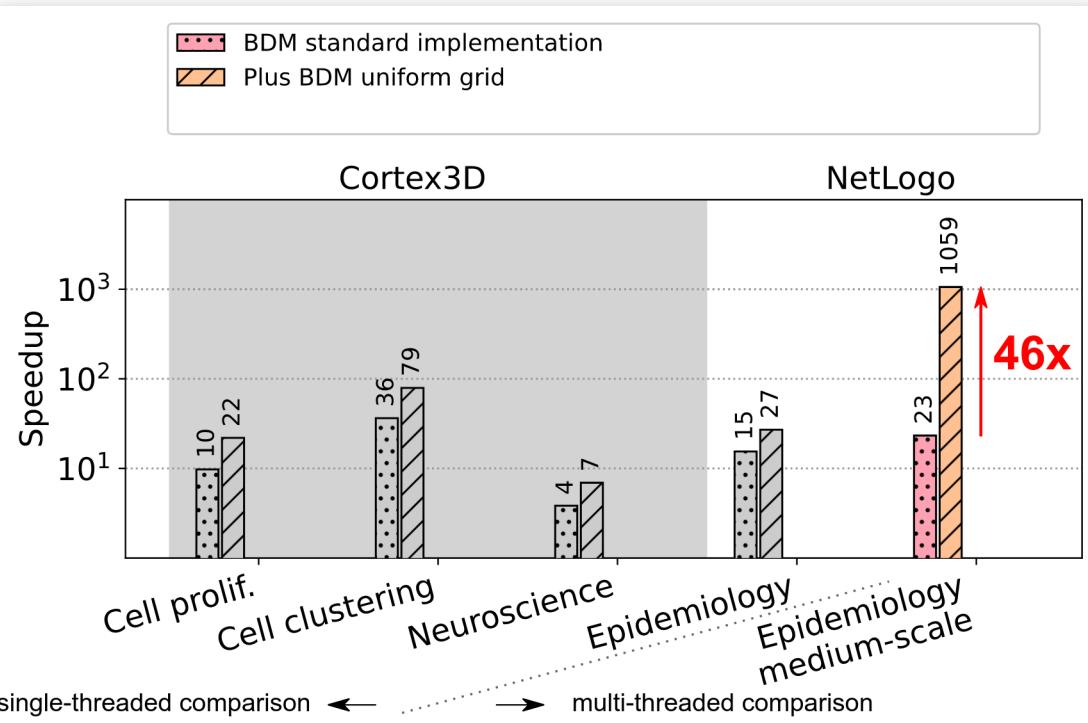
BioDynaMo is $1000\times$ faster than Cortex3D and NetLogo



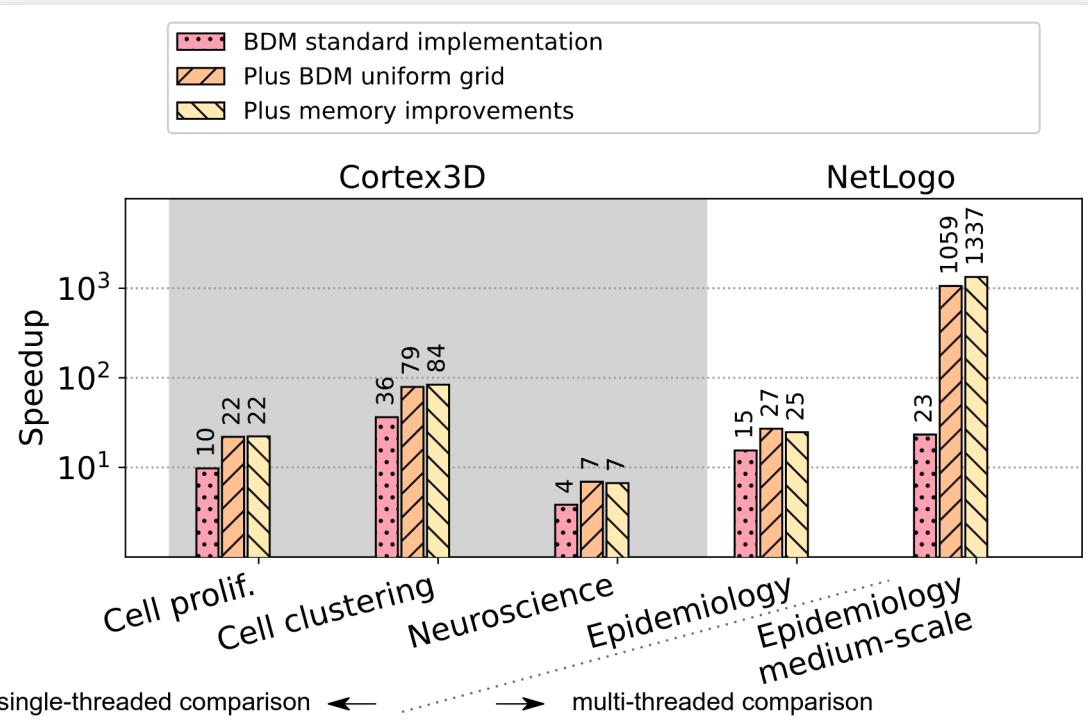
BioDynaMo is $1000\times$ faster than Cortex3D and NetLogo



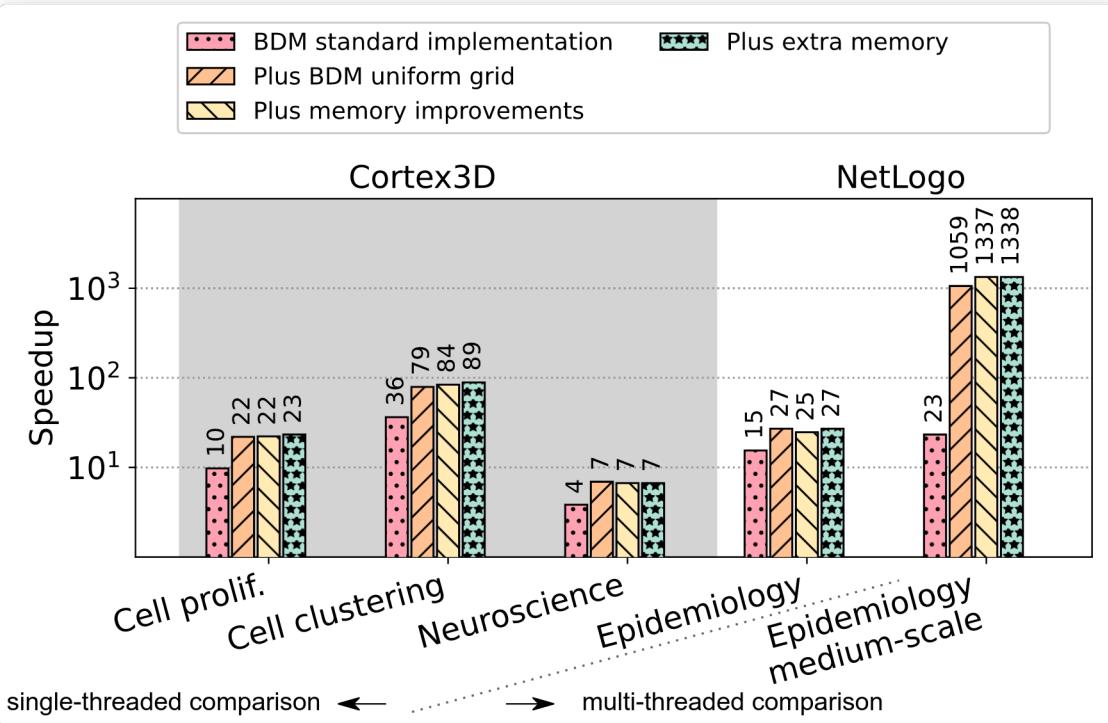
BioDynaMo is $1000\times$ faster than Cortex3D and NetLogo



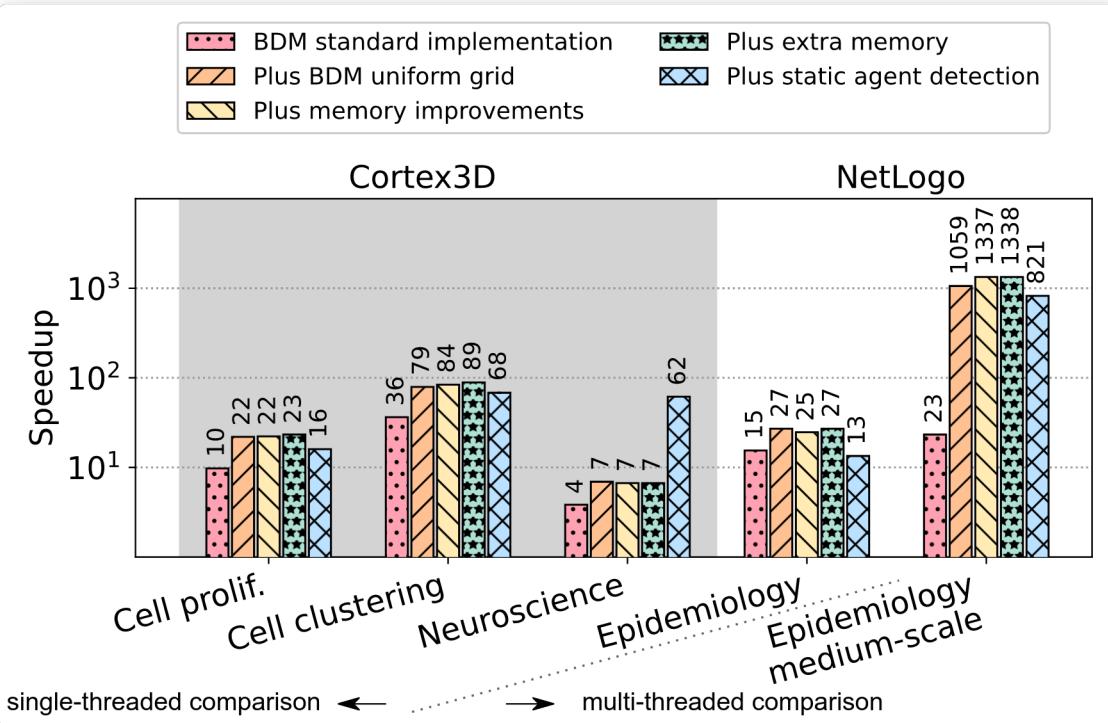
BioDynaMo is $1000\times$ faster than Cortex3D and NetLogo



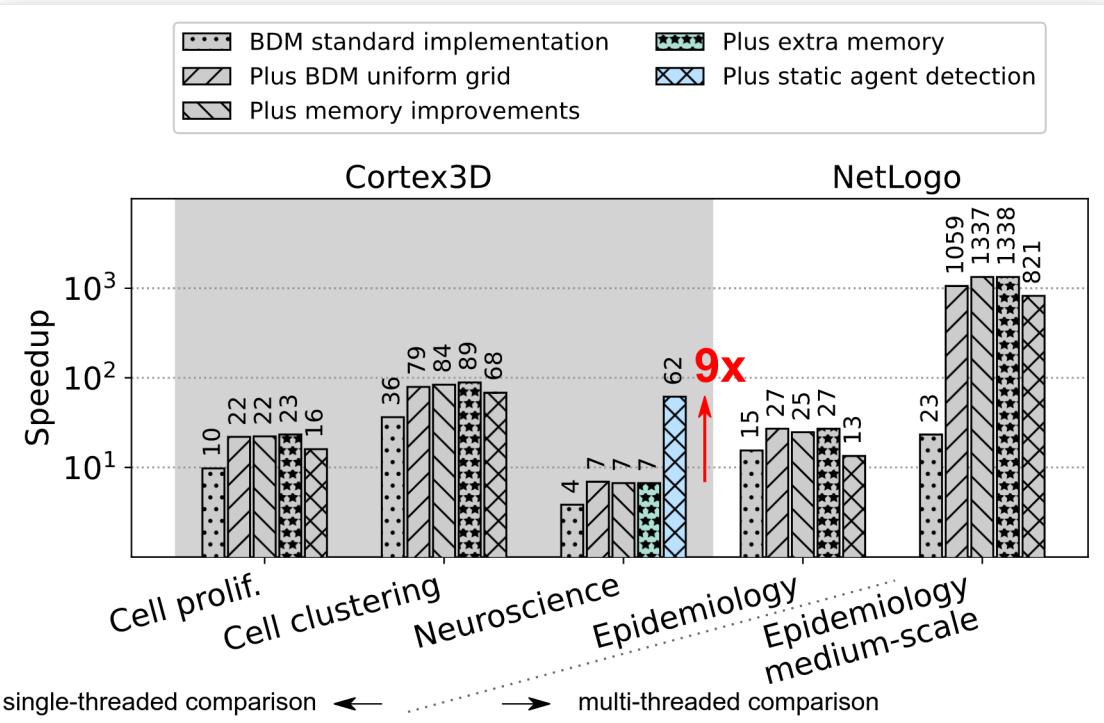
BioDynaMo is $1000\times$ faster than Cortex3D and NetLogo



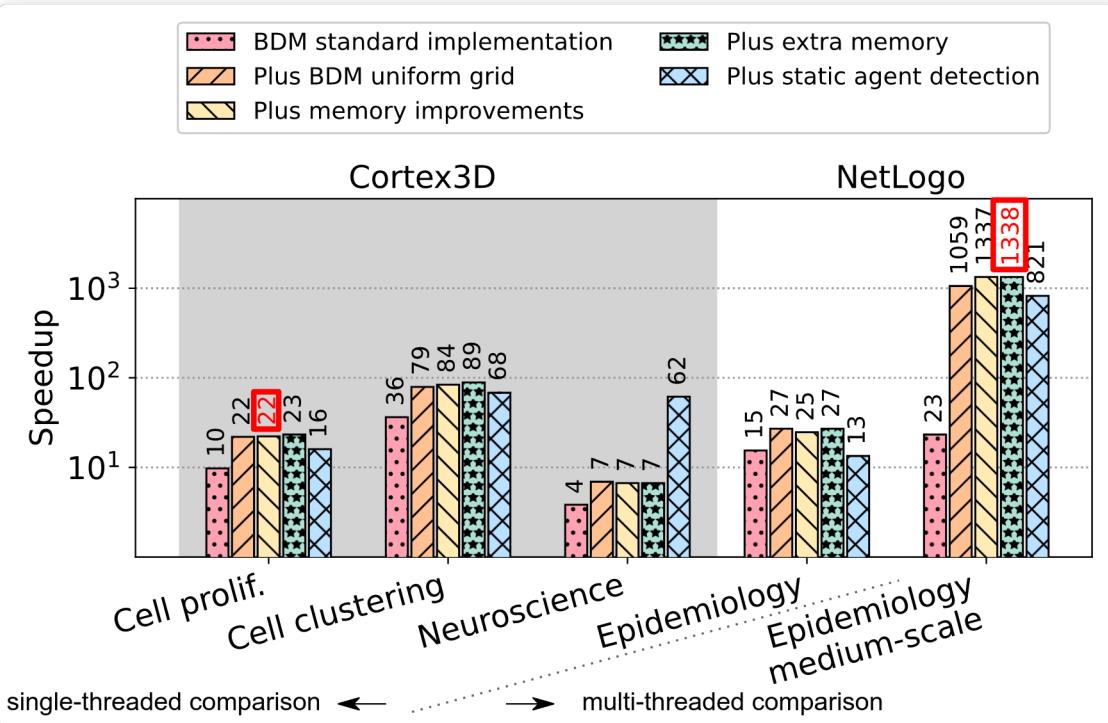
BioDynaMo is $1000\times$ faster than Cortex3D and NetLogo



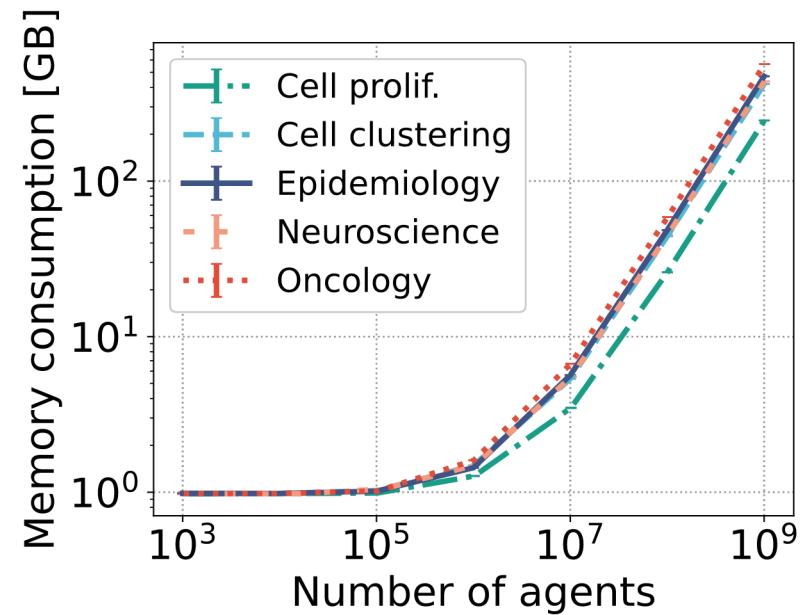
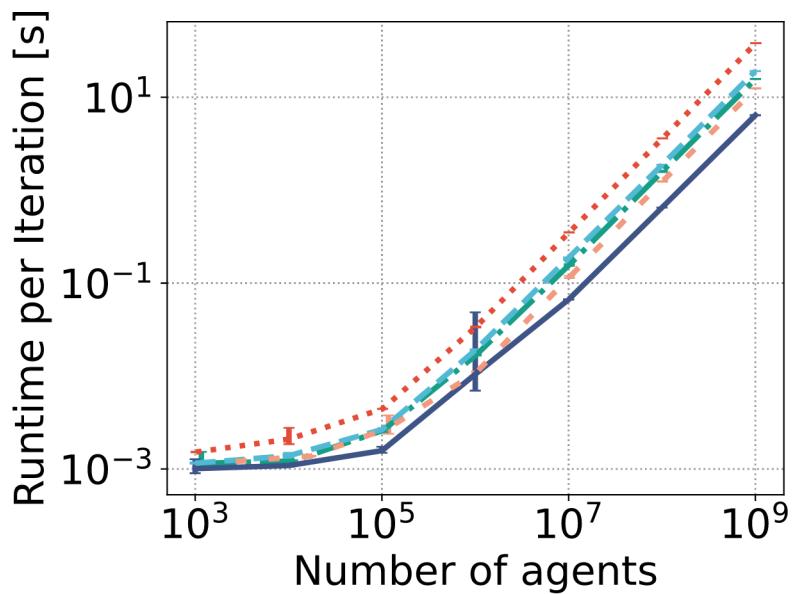
BioDynaMo is $1000\times$ faster than Cortex3D and NetLogo



BioDynaMo is $1000\times$ faster than Cortex3D and NetLogo



BioDynaMo's runtime and memory consumption increases linearly with model size



More in the paper

<https://arxiv.org/abs/2301.06984>, <https://doi.org/10.1145/3572848.3577480>



High-Performance and Scalable Agent-Based Simulation with BioDynaMo

Lukas Breitwieser*

CERN, Switzerland

ETH Zurich, Switzerland

Ahmad Hesam

Delft University of Technology,

The Netherlands

Fons Rademakers

CERN, Switzerland

Juan Gómez Luna

ETH Zurich, Switzerland

Onur Mutlu†

ETH Zurich, Switzerland

Abstract

Agent-based modeling plays an essential role in gaining insights into biology, sociology, economics, and other fields. However, many existing agent-based simulation platforms are not suitable for large-scale studies due to the low performance of the underlying simulation engines. To overcome this limitation, we present a novel high-performance simulation engine.

We identify three key challenges for which we present the following solutions. First, to maximize parallelization, we present an optimized grid to search for neighbors and parallelize the merging of thread-local results. Second, we reduce

1 Introduction

Agent-based modeling (ABM) allows to simulate complex dynamics in a wide range of research fields. ABM has been used to answer research questions in biology [30, 43, 73], sociology [18], economics [65], technology [50], business [56], and more fields [39]. Agents are individual entities that, among others, can represent subcellular structures to simulate the growth of a neuron, a cell to investigate cancer development, or a person to simulate the spread of infectious diseases [10]. The actions of an agent are defined through instances of class *behavior*. To stay with the examples from before, possible behaviors are neurite bifurcation, uncontrolled cell division, migration, etc.

Supplementary Materials

Provides additional evaluations and all files required to reproduce the results (  ):

- <https://doi.org/10.5281/zenodo.6463816>
- <https://github.com/CMU-SAFARI/BioDynaMo>

Supplementary Materials

Provides additional evaluations and all files required to reproduce the results (  ):

- <https://doi.org/10.5281/zenodo.6463816>
- <https://github.com/CMU-SAFARI/BioDynaMo>

```
1 # extract the provided code archive
2 mkdir reproduce
3 tar -xzf <path>/SF2-code.tar.gz -c reproduce
4 cd reproduce/bdm-paper-examples
5
6 # load the provided self-contained docker image
7 docker/load.sh <path>/SF3-bdm-publication-image.tar.gz
8
9 # execute the scripts
10 docker/run.sh ./run-main.sh
```

Supplementary Materials

Provides additional evaluations and all files required to reproduce the results (  ):

- <https://doi.org/10.5281/zenodo.6463816>
- <https://github.com/CMU-SAFARI/BioDynaMo>

```
1 # extract the provided code archive
2 mkdir reproduce
3 tar -xzf <path>/SF2-code.tar.gz -c reproduce
4 cd reproduce/bdm-paper-examples
5
6 # load the provided self-contained docker image
7 docker/load.sh <path>/SF3-bdm-publication-image.tar.gz
8
9 # execute the scripts
10 docker/run.sh ./run-main.sh
```

Supplementary Materials

Provides **additional evaluations** and all files required to **reproduce** the results (  ):

- <https://doi.org/10.5281/zenodo.6463816>
- <https://github.com/CMU-SAFARI/BioDynaMo>

```
1 # extract the provided code archive
2 mkdir reproduce
3 tar -xzf <path>/SF2-code.tar.gz -C reproduce
4 cd reproduce/bdm-paper-examples
5
6 # load the provided self-contained docker image
7 docker/load.sh <path>/SF3-bdm-publication-image.tar.gz
8
9 # execute the scripts
10 docker/run.sh ./run-main.sh
```

Supplementary Materials

Provides additional evaluations and all files required to reproduce the results (  ):

- <https://doi.org/10.5281/zenodo.6463816>
- <https://github.com/CMU-SAFARI/BioDynaMo>

```
1 # extract the provided code archive
2 mkdir reproduce
3 tar -xzf <path>/SF2-code.tar.gz -C reproduce
4 cd reproduce/bdm-paper-examples
5
6 # load the provided self-contained docker image
7 docker/load.sh <path>/SF3-bdm-publication-image.tar.gz
8
9 # execute the scripts
10 docker/run.sh ./run-main.sh
```

Supplementary Materials

Provides additional evaluations and all files required to reproduce the results (  ):

- <https://doi.org/10.5281/zenodo.6463816>
- <https://github.com/CMU-SAFARI/BioDynaMo>

```
1 # extract the provided code archive
2 mkdir reproduce
3 tar -xzf <path>/SF2-code.tar.gz -C reproduce
4 cd reproduce/bdm-paper-examples
5
6 # load the provided self-contained docker image
7 docker/load.sh <path>/SF3-bdm-publication-image.tar.gz
8
9 # execute the scripts
10 docker/run.sh ./run-main.sh
```

More details in file SF1-readme.pdf →

Instructions to Reproduce the Results in the Paper:
High-Performance and Scalable Agent-Based Simulation
with BioDynaMo

Lukas Breitwieser Ahmad Hesam Fons Rademakers
Juan Gómez Luna Onur Muthu

Contents

1	Introduction	2
2	Getting started	3
2.1	Extract code repositories	3
2.2	Install host machine prerequisites	3
2.3	Load Docker image	4
2.4	Verify successful setup	4
3	Documentation	5
3.1	The bdm-paper-examples repository	5
3.2	Mapping between simulation names in the paper and bdm-paper-examples	5
3.3	Docker image	6
3.4	Important information	6
3.5	Result directory structure	6
3.6	Long execution time	7
4	Reproducing results	8
4.1	Main results (Figure 5 (left) and 9–12 in the paper)	8
4.2	Comparison with Cortex3D and NetLogo (Figure 8 in the paper)	10
4.3	Runtime and space complexity (Figure 6 in the paper)	11
4.4	Workload profiling (Figure 5 (right) in the paper)	12
4.5	Biocellion comparison small (Figure 7 in the paper)	13
4.6	Biocellion comparison large (Figure 7 in the paper)	14
5	Reusing and repurposing this artifact	15
5.1	Add additional benchmarks	15
5.2	Evaluate the effectiveness of additional optimizations	17
5.3	Evaluate BioDynaMo’s performance for additional simulations	17
6	Contact	18

Supplementary Materials

Provides additional evaluations and all files required to reproduce the results (  ):

- <https://doi.org/10.5281/zenodo.6463816>
- <https://github.com/CMU-SAFARI/BioDynaMo>

```
1 # extract the provided code archive
2 mkdir reproduce
3 tar -xzf <path>/SF2-code.tar.gz -C reproduce
4 cd reproduce/bdm-paper-examples
5
6 # load the provided self-contained docker image
7 docker/load.sh <path>/SF3-bdm-publication-image.tar.gz
8
9 # execute the scripts
10 docker/run.sh ./run-main.sh
```

More details in file SF1-readme.pdf →



Instructions to Reproduce the Results in the Paper:
High-Performance and Scalable Agent-Based Simulation
with BioDynaMo

Lukas Breitwieser Ahmad Hesam Fons Rademakers
Juan Gómez Luna Onur Muthu

Contents

1	Introduction	2
2	Getting started	3
2.1	Extract code repositories	3
2.2	Install host machine prerequisites	3
2.3	Load Docker image	4
2.4	Verify successful setup	4
3	Documentation	5
3.1	The bdm-paper-examples repository	5
3.2	Mapping between simulation names in the paper and bdm-paper-examples	5
3.3	Docker image	6
3.4	Important information	6
3.5	Result directory structure	6
3.6	Long execution time	7
4	Reproducing results	8
4.1	Main results (Figure 5 (left) and 9–12 in the paper)	8
4.2	Comparison with Cortex3D and NetLogo (Figure 8 in the paper)	10
4.3	Runtime and space complexity (Figure 6 in the paper)	11
4.4	Workload profiling (Figure 5 (right) in the paper)	12
4.5	Biocellion comparison small (Figure 7 in the paper)	13
4.6	Biocellion comparison large (Figure 7 in the paper)	14
5	Reusing and repurposing this artifact	15
5.1	Add additional benchmarks	15
5.2	Evaluate the effectiveness of additional optimizations	17
5.3	Evaluate BioDynaMo’s performance for additional simulations	17
6	Contact	18

Summary

Our optimizations to maximize parallelism, improve the memory layout, and avoid unnecessary work are effective and give BioDynaMo the following performance characteristics.

Summary

Our optimizations to maximize parallelism, improve the memory layout, and avoid unnecessary work are effective and give BioDynaMo the following performance characteristics.

BioDynaMo scales well across a large number of CPU cores

Summary

Our optimizations to maximize parallelism, improve the memory layout, and avoid unnecessary work are effective and give BioDynaMo the following performance characteristics.

BioDynaMo scales well across a large number of CPU cores

Runtime and memory consumption increases linearly with the number of agents

Summary

Our optimizations to maximize parallelism, improve the memory layout, and avoid unnecessary work are effective and give BioDynaMo the following performance characteristics.

BioDynaMo scales well across a large number of CPU cores

Runtime and memory consumption increases linearly with the number of agents

BioDynaMo is more than $1000\times$ faster than Cortex3D and NetLogo and up to $9\times$ more efficient than Biocellion

Summary

Our optimizations to maximize parallelism, improve the memory layout, and avoid unnecessary work are effective and give BioDynaMo the following performance characteristics.

BioDynaMo scales well across a large number of CPU cores

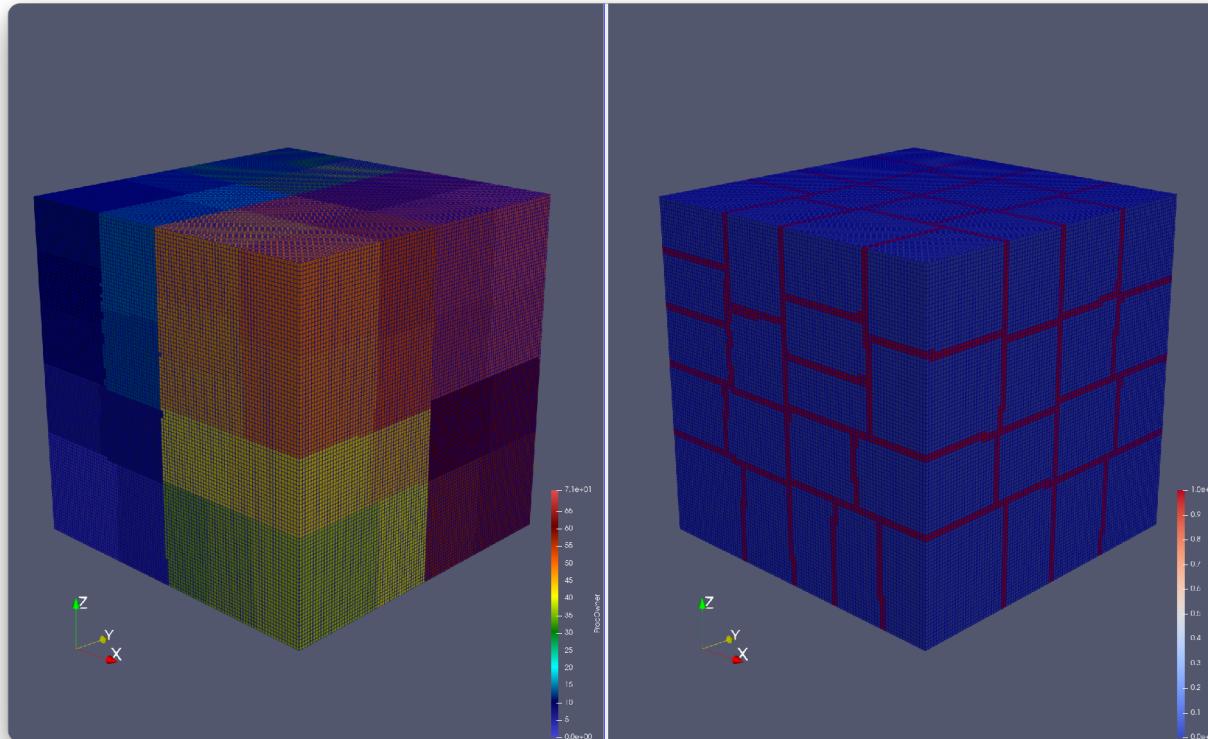
Runtime and memory consumption increases linearly with the number of agents

BioDynaMo is more than $1000\times$ faster than Cortex3D and NetLogo and up to $9\times$ more efficient than Biocellion

These performance characteristics enable simulations with billions of agents on a single server.

Future work

Distributed simulation engine



BioDynaMo is available under the permissive Apache 2.0 open source license

<https://github.com/BioDynaMo/biodynamo>

The screenshot shows the GitHub repository page for BioDynaMo. The top navigation bar includes 'Code' (selected), 'Issues 8', 'Pull requests 6', 'Discussions', 'Actions', 'Projects', 'Wiki', 'Security', 'Insights', and 'Settings'. Below the navigation is a search bar with dropdowns for 'master', '52 branches', and '8 tags'. Buttons for 'Go to file', 'Add file', and 'Code' are also present. The main content area displays a list of recent commits by TobiasDuswald:

Commit	Message	Date
.github	Introduce libgit2 to track used code for reproducibility (#309)	last month
benchmark	Update Copyright and extend GHA for Copyright	last year
cli	Avoid parsing of bdm.json with bdm test	10 months ago
cmake	Add Paraview sha's for macOS 13.2, also created the links to 13.2 ve...	5 days ago
demo	Replace rm->AddAgent with ctxt->AddAgent in demos (#309)	4 days ago
doc	Replace rm->AddAgent with ctxt->AddAgent in demos (#309)	4 days ago
etc	Pix several issues on macOS and allow building without paraview.	2 years ago
notebook	Make floating-point precision adjustable (#253)	6 months ago
paraview_plugin	Update copyright information	last year
src	Replace rm->AddAgent with ctxt->AddAgent in demos (#309)	4 days ago
test	Add PrintInfo member function to DiffusionGrid (#293)	2 months ago
third_party	Fix optimlib / armadillo dependency (#235)	2 months ago
util	Introduce libgit2 to track used code for reproducibility (#300)	last month
.clang-format	Generated new version by clang-format-3.9. Honors Google includ...	6 years ago
.clang-tidy	Fix clang-tidy errors	3 years ago
.clang-tidy-ignore	Improve automated code checks	5 years ago
.editorconfig	Add support for editors supporting http://editorconfig.org .	6 years ago
.gitignore	Introduce libgit2 to track used code for reproducibility (#300)	last month
.mailmap	Update mailmap file	4 months ago

To the right of the commit list is the 'About' section, which describes BioDynaMo as a high-performance and modular, agent-based simulation platform. It lists various tags: simulation, cancer, biology, high-performance, neuroscience, epidemiology, agent-based, large-scale, modular-design, agent-based-framework, and agent-based-modelling. Below the 'About' section are links to 'Readme', 'Apache-2.0 license', 'Code of conduct', '61 stars', '10 watching', and '41 forks'. The 'Releases' section shows one release: BioDynaMo 1.04 (Latest, released on Oct 5, 2022). The 'Contributors' section shows 22 contributors with their GitHub profiles.

Join us in pushing the limits of agent-based modeling

BioDynamo is an open, welcoming, and collaborative project: <https://biodynamo.org>

Besides the organizations behind this paper,



ETH zürich

TUDelft

BioDynamo is developed by the BioDynamo collaboration,



and further organizations.



More in the paper

<https://arxiv.org/abs/2301.06984>, <https://doi.org/10.1145/3572848.3577480>



High-Performance and Scalable Agent-Based Simulation with BioDynaMo

Lukas Breitwieser*, Ahmad Hesam, Fons Rademakers
CERN, Switzerland
Delft University of Technology,
The Netherlands
CERN, Switzerland

Juan Gómez Luna, Onur Mutlu†
ETH Zurich, Switzerland
ETH Zurich, Switzerland

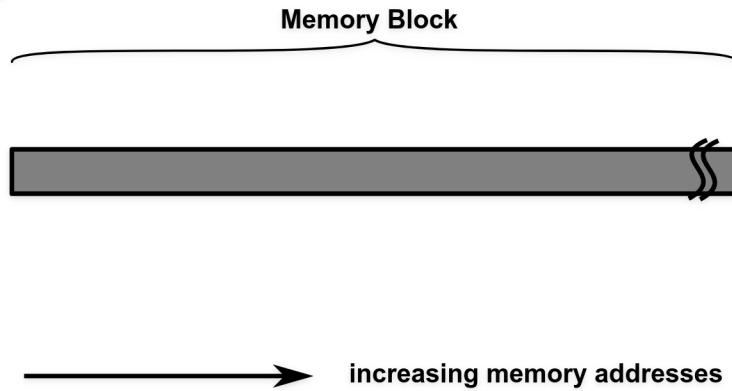
Abstract
Agent-based modeling plays an essential role in gaining insights into biology, sociology, economics, and other fields. However, many existing agent-based simulation platforms are not suitable for large-scale studies due to the low performance of the underlying simulation engines. To overcome this limitation, we present a novel high-performance simulation engine.
We identify three key challenges for which we present the following solutions. First, to maximize parallelization, we present an optimized grid to search for neighbors and parallelize the merging of thread-local results. Second, we reduce the memory access latency with a NUMA-aware agent iterator, agent sorting with a space-filling curve, and a custom heap memory allocator. Third, we present a mechanism to omit the collision force calculation under certain conditions.
Our evaluation shows an order of magnitude improvement over Biocellion, three orders of magnitude speedup over Cortex3D and NetLogo, and the ability to simulate 1.72 billion agents on a single server.
Supplementary Materials, including instructions to reproduce the results, are available at: <https://doi.org/10.5281/zendo.6463816>



Appendix

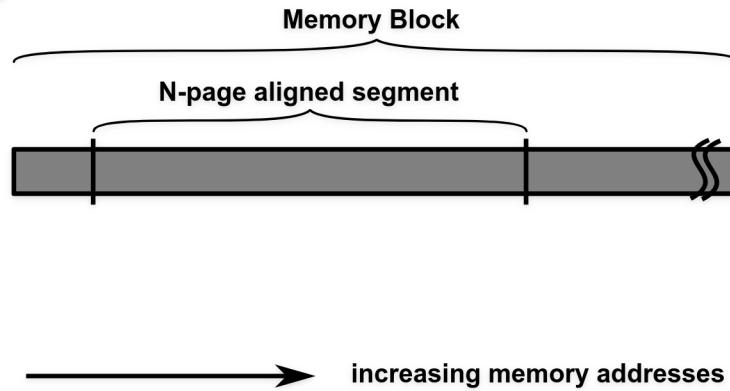
Memory allocation

A Layout



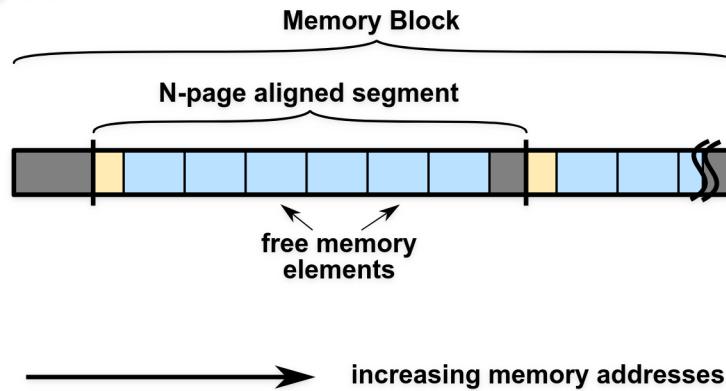
Memory allocation

A Layout



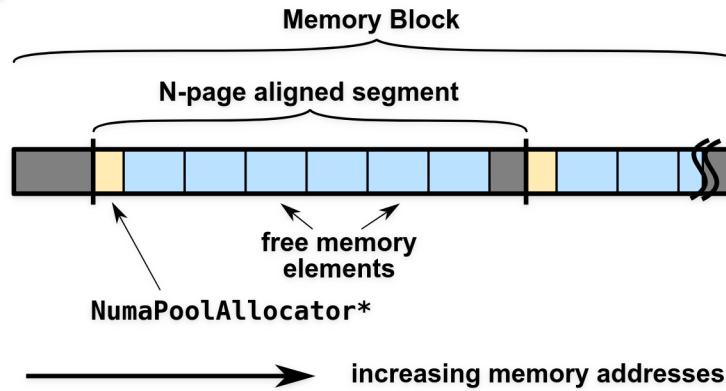
Memory allocation

A Layout



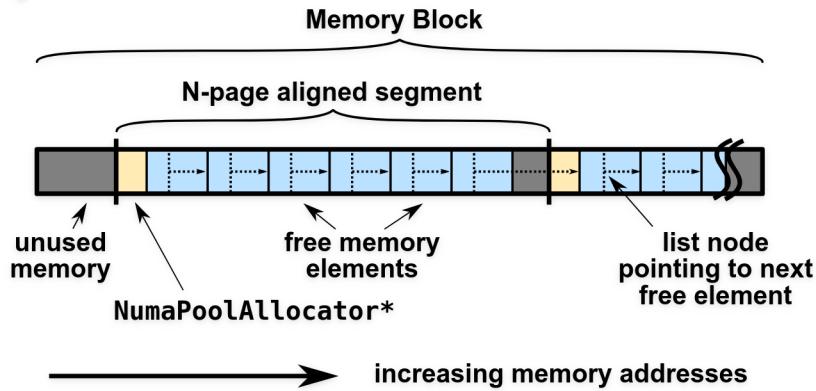
Memory allocation

A Layout



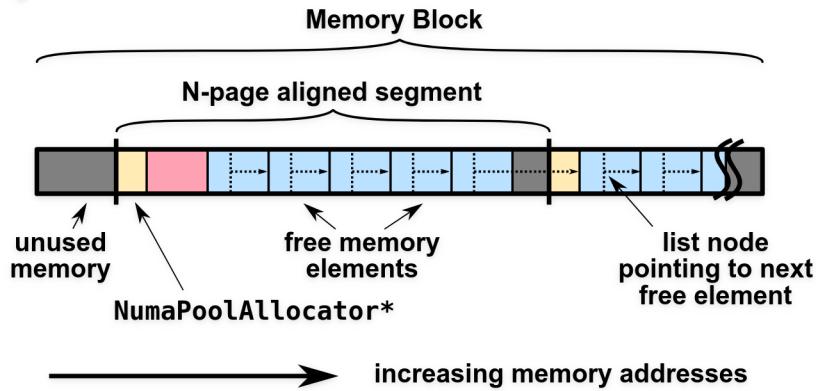
Memory allocation

A Layout



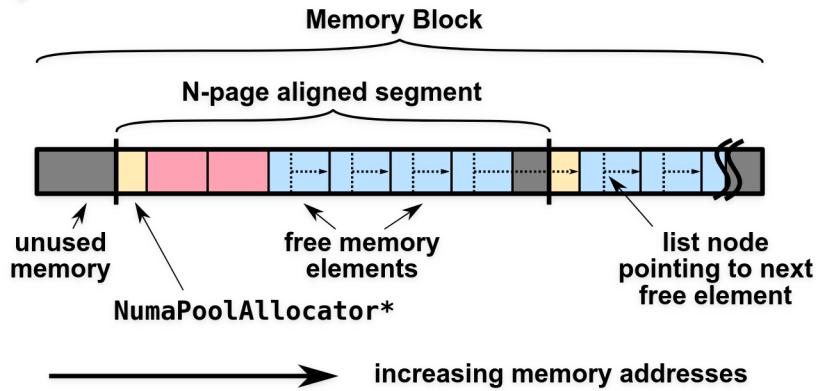
Memory allocation

A Layout



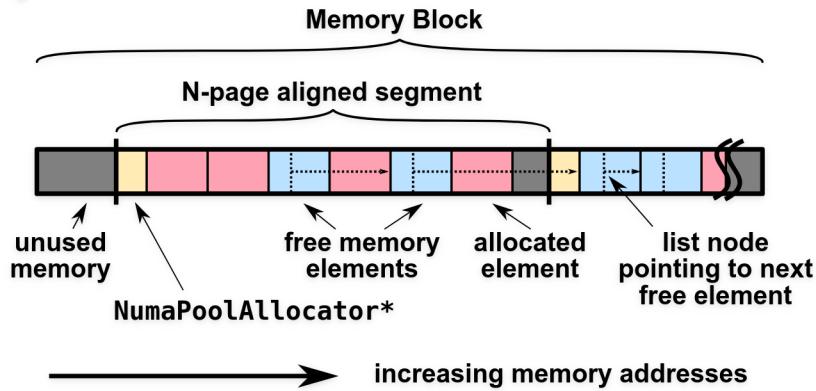
Memory allocation

A Layout



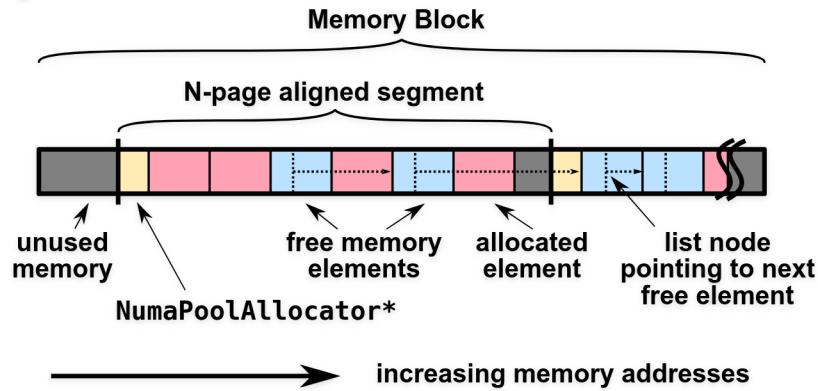
Memory allocation

A Layout

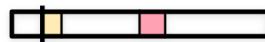


Memory allocation

A Layout

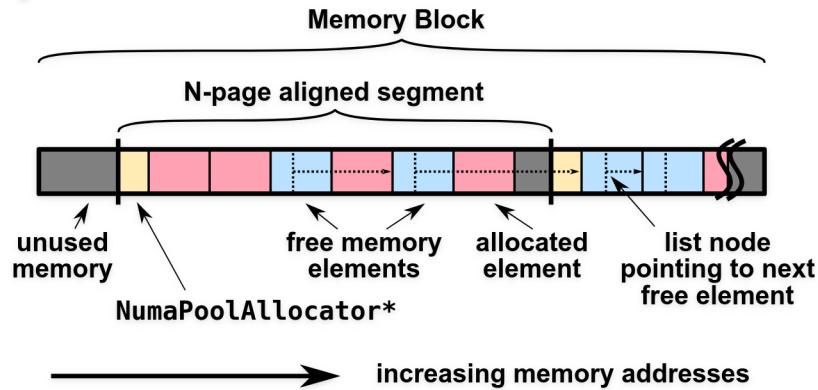


B Deallocation

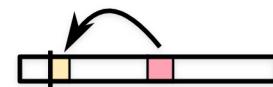


Memory allocation

A Layout



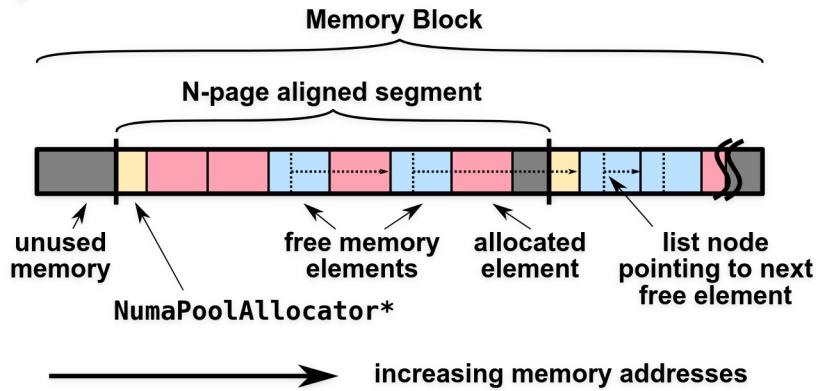
B Deallocation



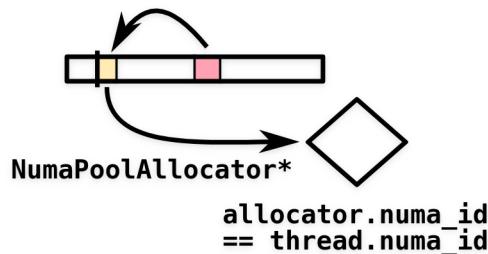
NumaPoolAllocator*

Memory allocation

A Layout

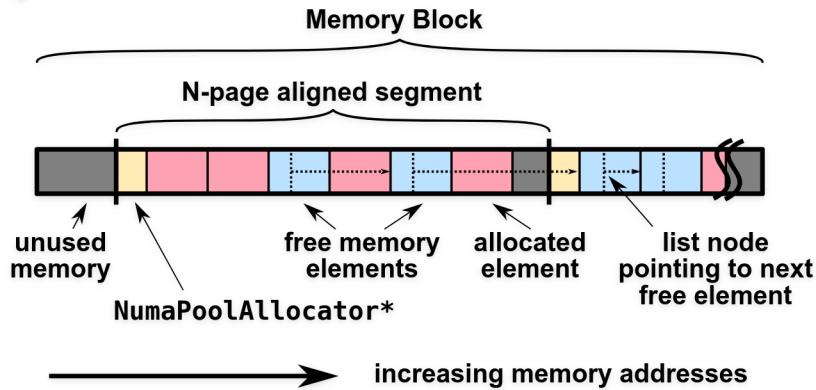


B Deallocation

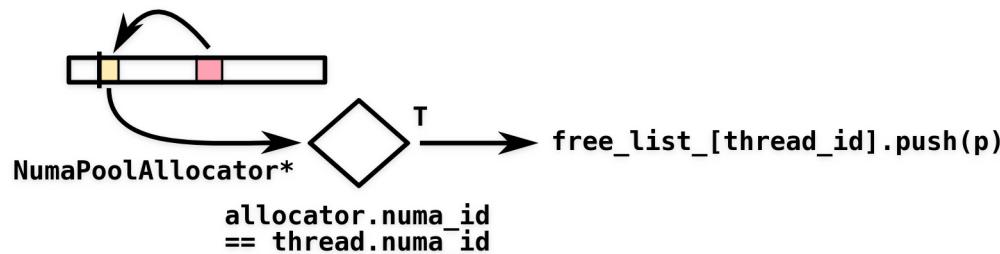


Memory allocation

A Layout

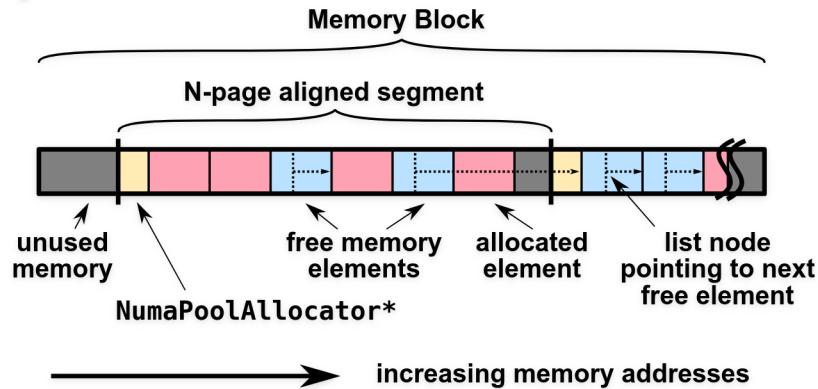


B Deallocation

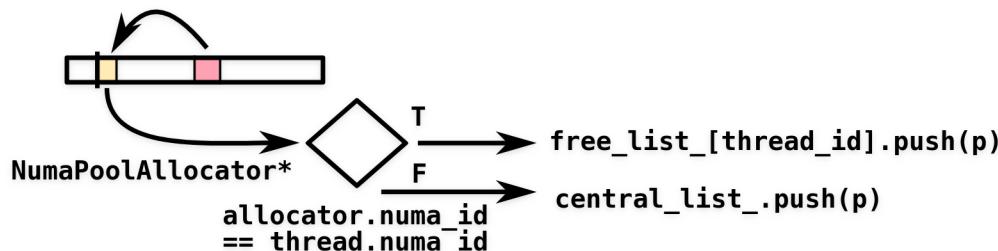


Memory allocation

A Layout



B Deallocation



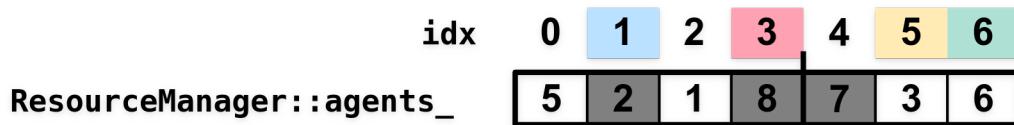
Memory allocation

Parallel agent removal

Removed agents:

Thread 0: {2, 8}

Thread 1: {7}



1. Initialize

to_right	<table border="1"><tr><td>∞</td><td>∞</td><td>∞</td></tr></table>	∞	∞	∞	<table border="1"><tr><td>0</td><td>0</td><td>0</td></tr></table>	0	0	0	not_to_left
∞	∞	∞							
0	0	0							

2. Fill

to_right	<table border="1"><tr><td>1</td><td>3</td><td>∞</td></tr></table>	1	3	∞	<table border="1"><tr><td>1</td><td>0</td><td>0</td></tr></table>	1	0	0	not_to_left
1	3	∞							
1	0	0							

3. Reorder

to_right	<table border="1"><tr><td>1</td><td>3</td><td>∞</td></tr></table>	1	3	∞	<table border="1"><tr><td>5</td><td>0</td><td>6</td></tr></table>	5	0	6	to_left
1	3	∞							
5	0	6							

#swaps

<table border="1"><tr><td>2</td><td>0</td></tr></table>	2	0	<table border="1"><tr><td>1</td><td>1</td></tr></table>	1	1
2	0				
1	1				

4. Swap

to_right	<table border="1"><tr><td>1</td><td>3</td><td>∞</td></tr></table>	1	3	∞	<table border="1"><tr><td>5</td><td>0</td><td>6</td></tr></table>	5	0	6	to_left
1	3	∞							
5	0	6							

Thread 0

Thread 1

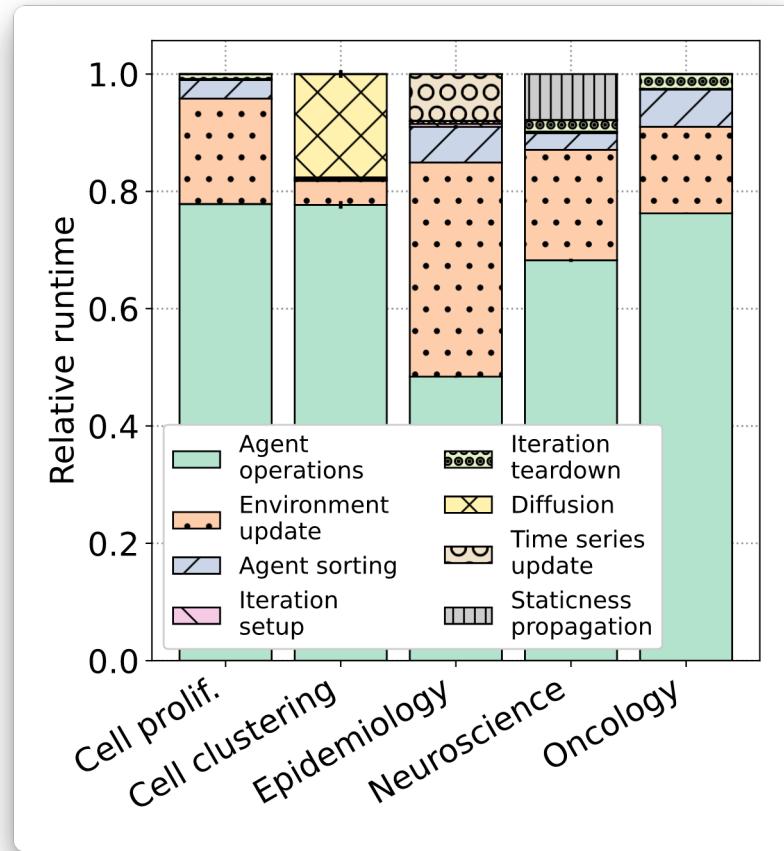
5. Resize

5	3	1	6
---	---	---	---

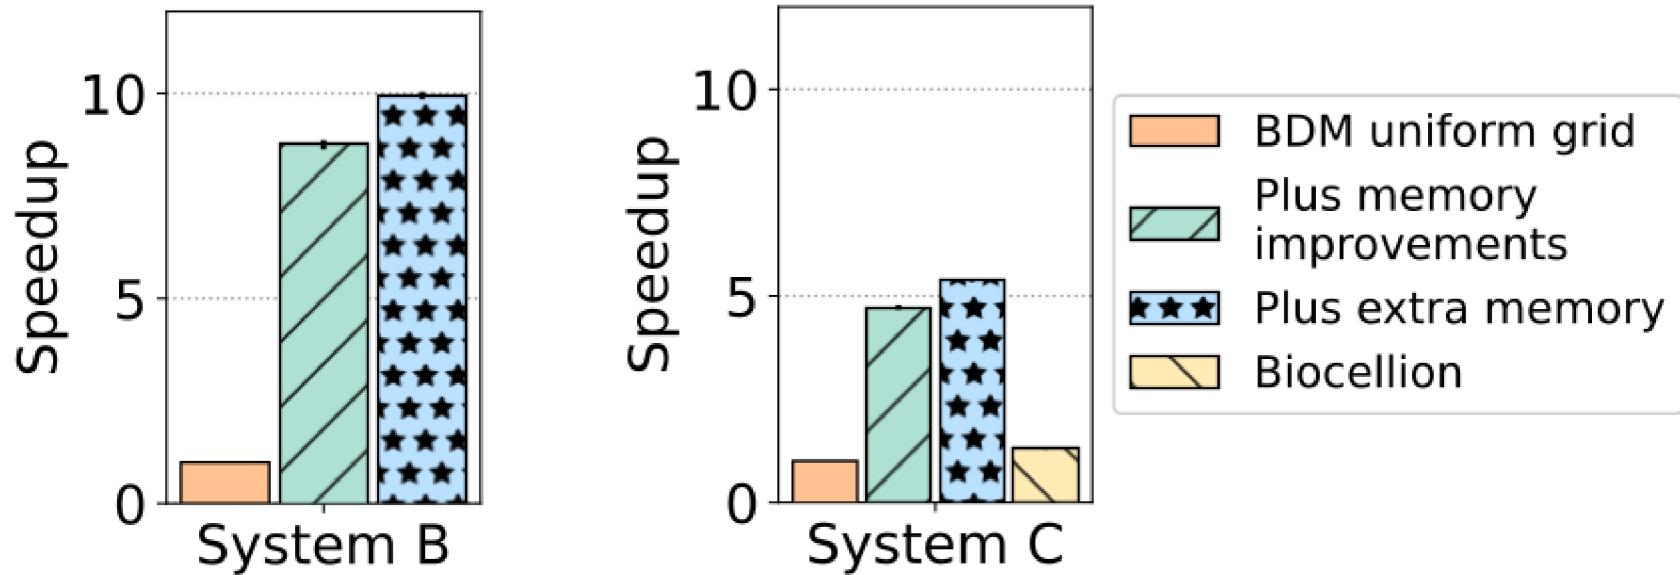
Performance-related simulation characteristics

Characteristic	Cell proliferation	Cell clustering	Epidemiology	Neuroscience	Oncology
Create new agents during simulation	x			x	x
Delete agents during simulation					x
Agents modify neighbors			x	x	x
Load imbalance			x	x	
Agents move randomly			x		x
Simulation uses diffusion	x		x		
Simulation has static regions			x		
Number of iterations	500	1000	1000	500	288
Number of agents (in millions)	12.6	2	10	9	10
Number of diffusion volumes	0	54m	0	65k	0

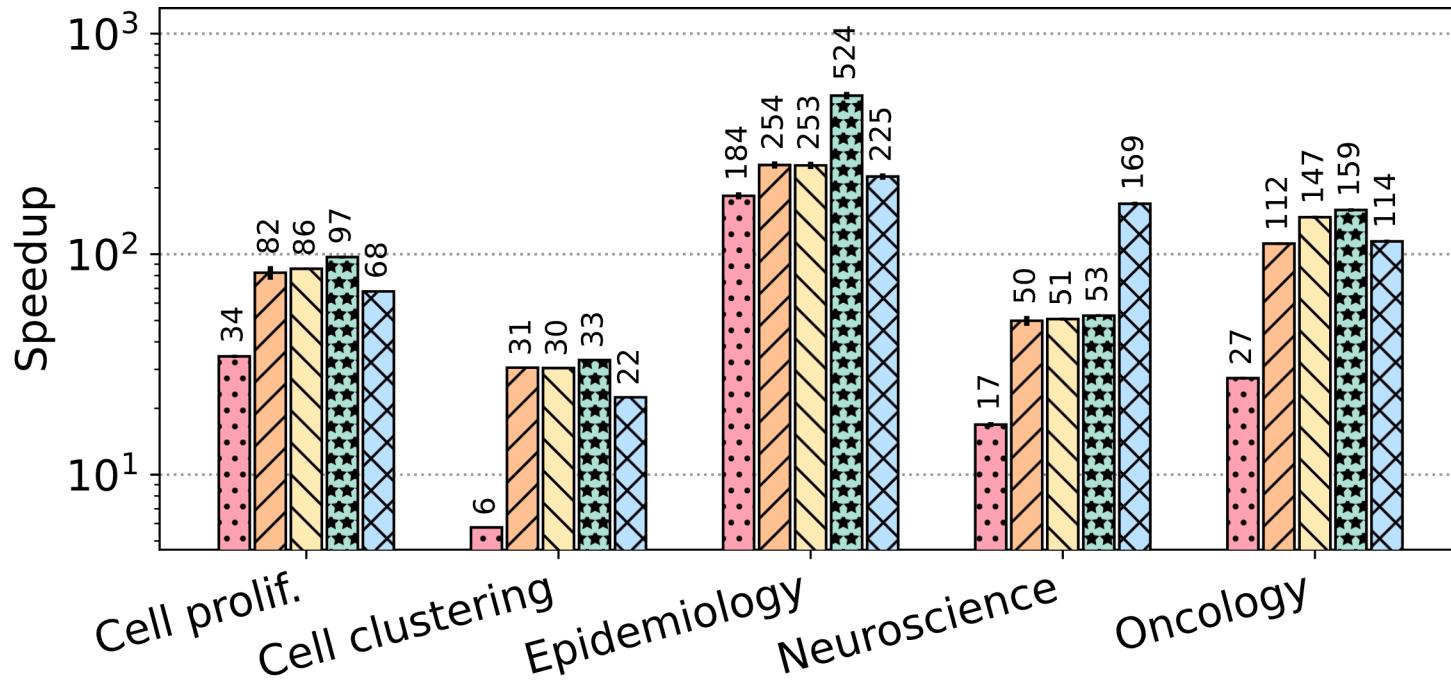
Operation runtime breakdown



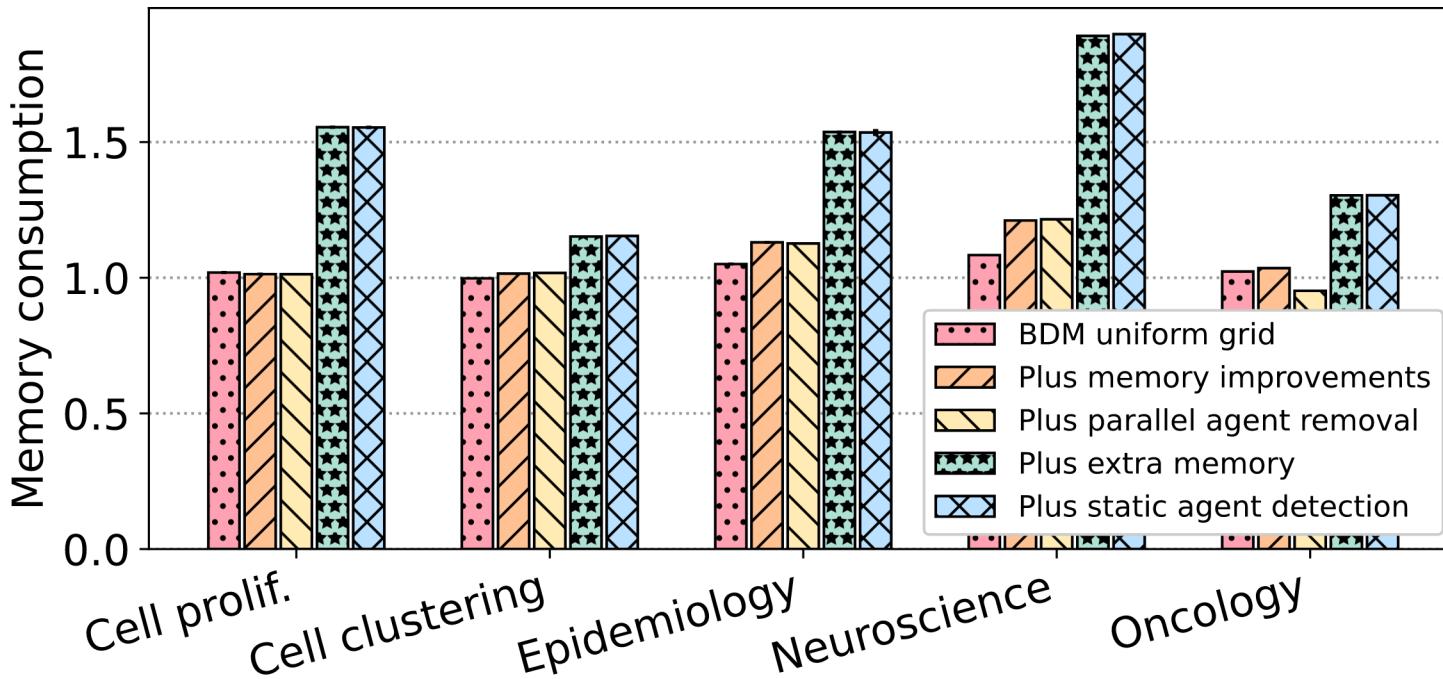
BioDynaMo / Biocellion comparison analysis



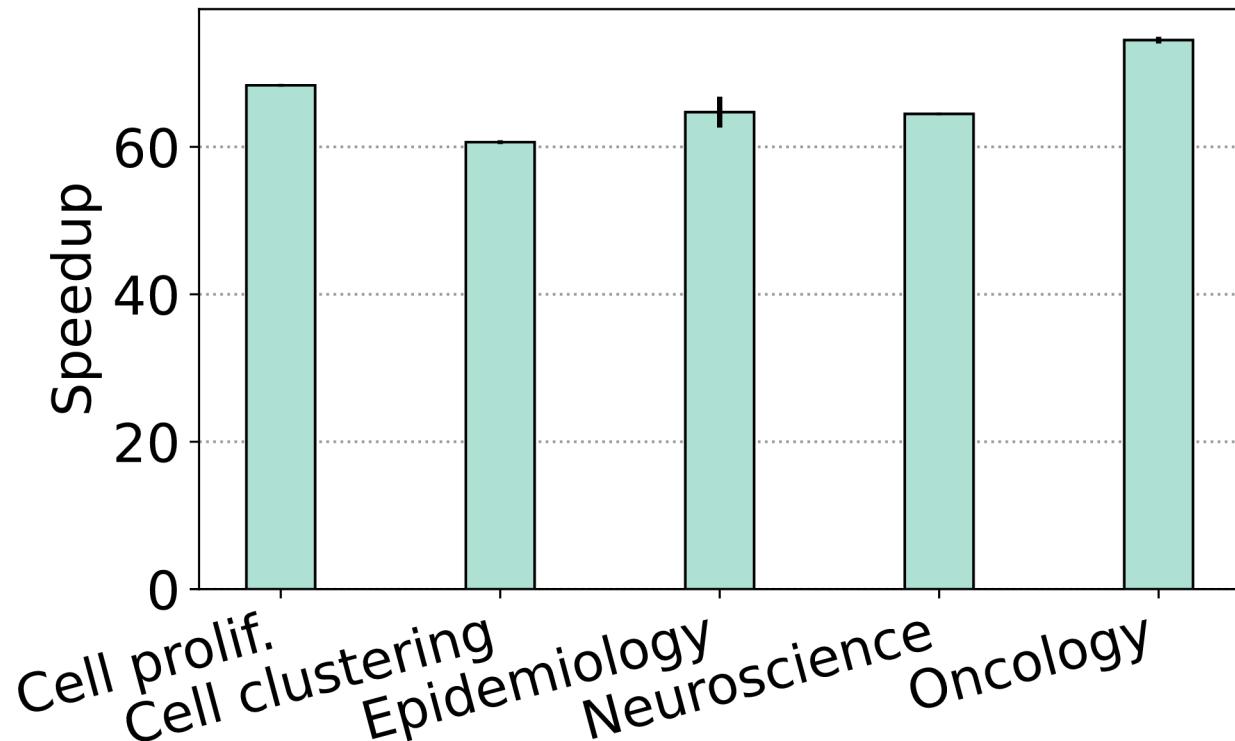
Optimization overview for larger scale simulations



Optimization overview for larger scale simulations

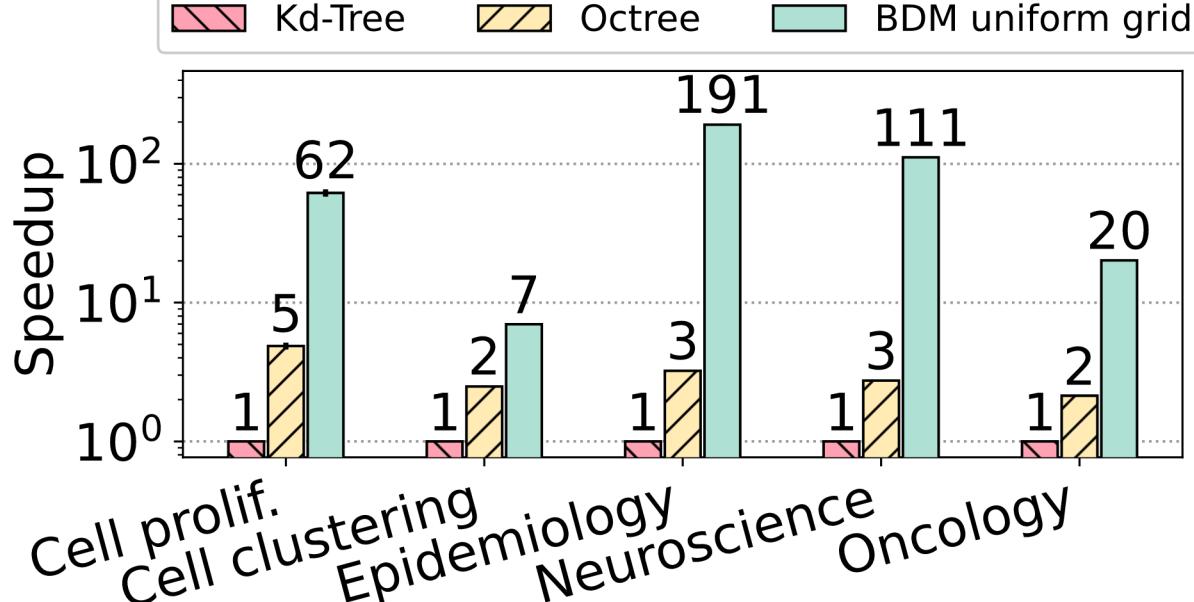


Full simulation scalability



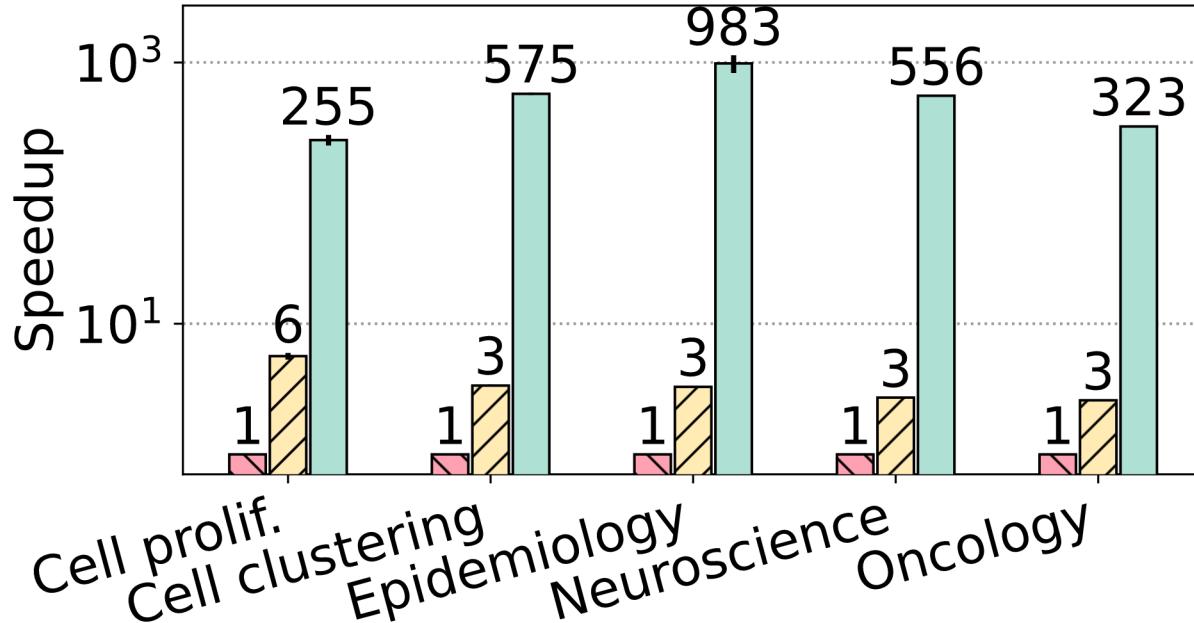
Radial neighbor search comparison

Whole simulation (144 threads and 4 NUMA domains)



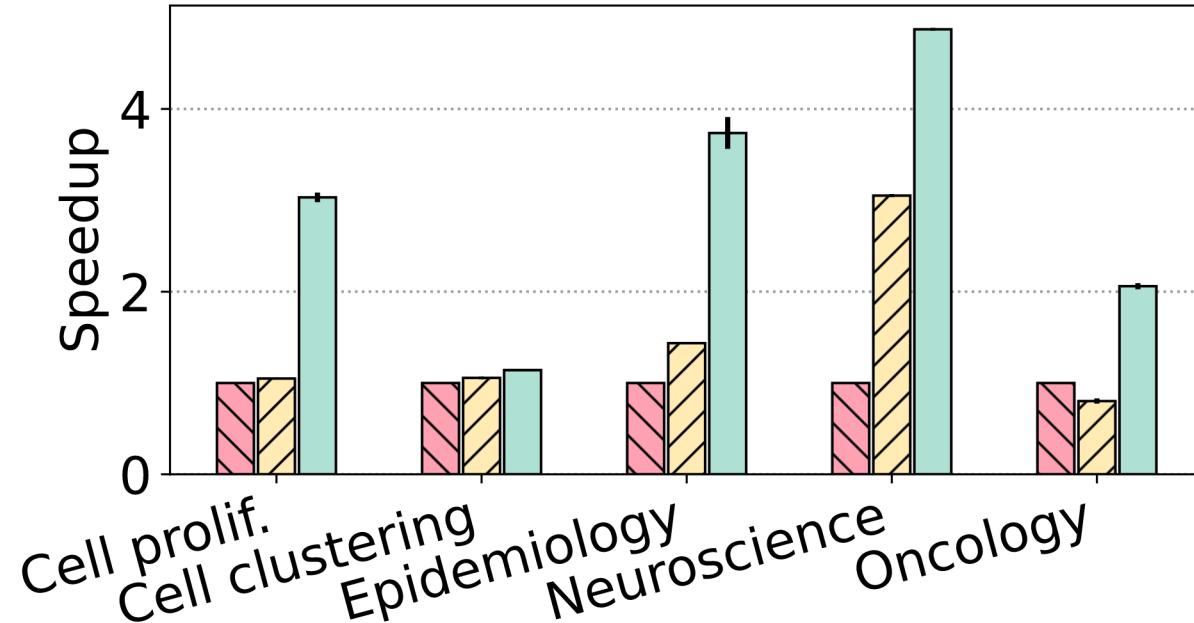
Radial neighbor search comparison

Build time (144 threads and 4 NUMA domains)



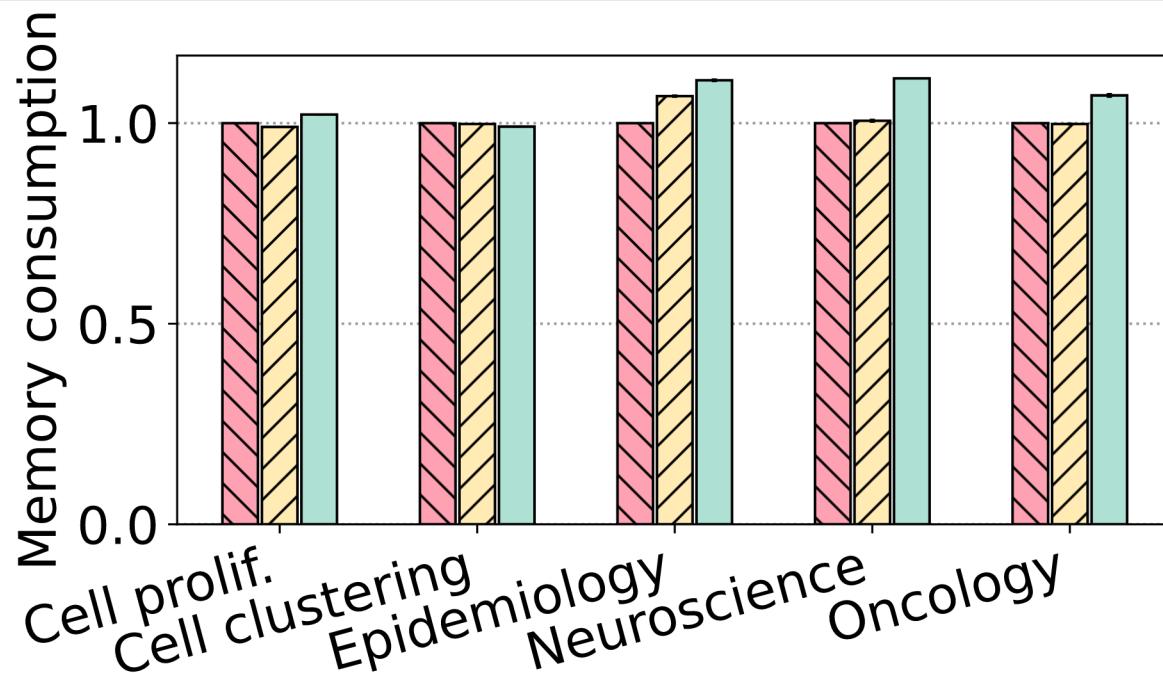
Radial neighbor search comparison

Search time indirect (144 threads and 4 NUMA domains)



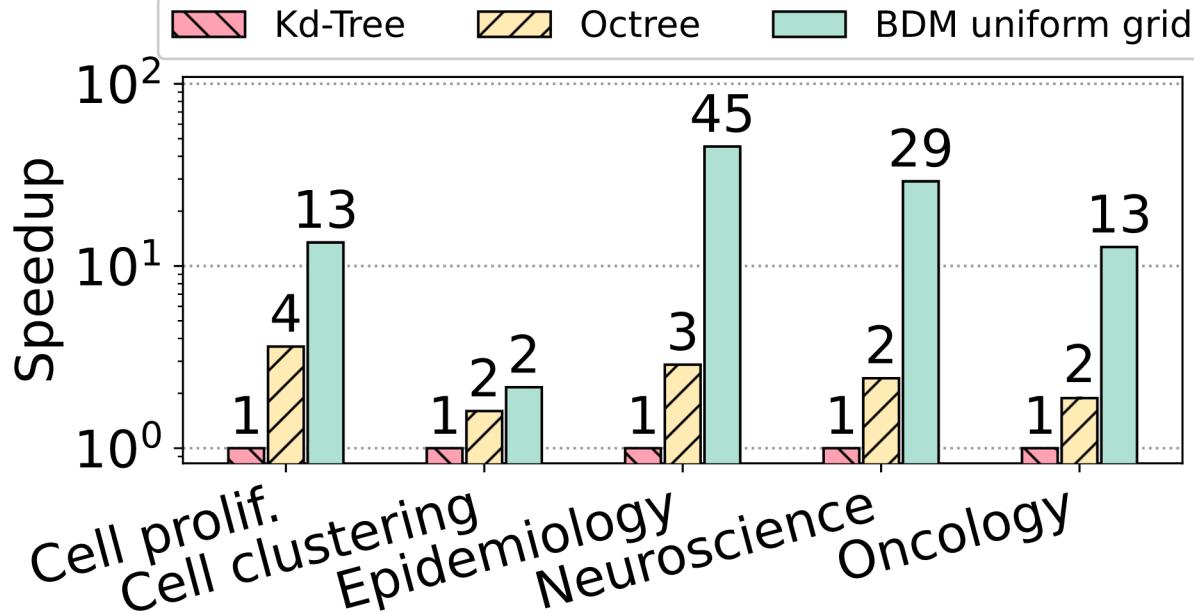
Radial neighbor search comparison

Memory consumption (144 threads and 4 NUMA domains)



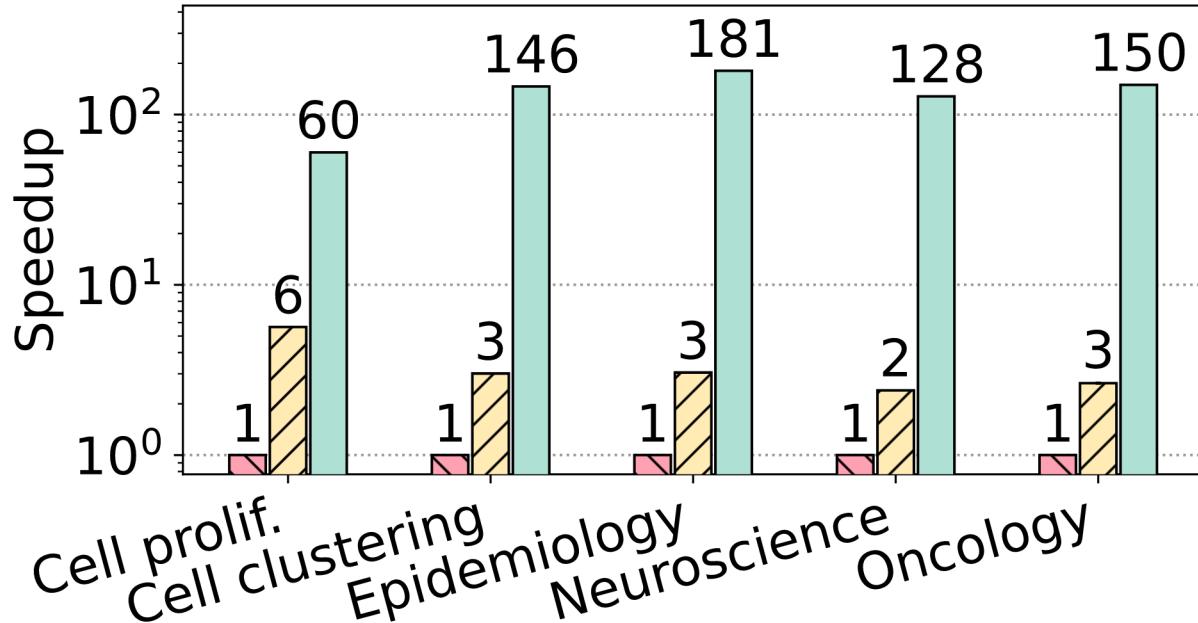
Radial neighbor search comparison

Whole simulation (18 threads and one NUMA domain)



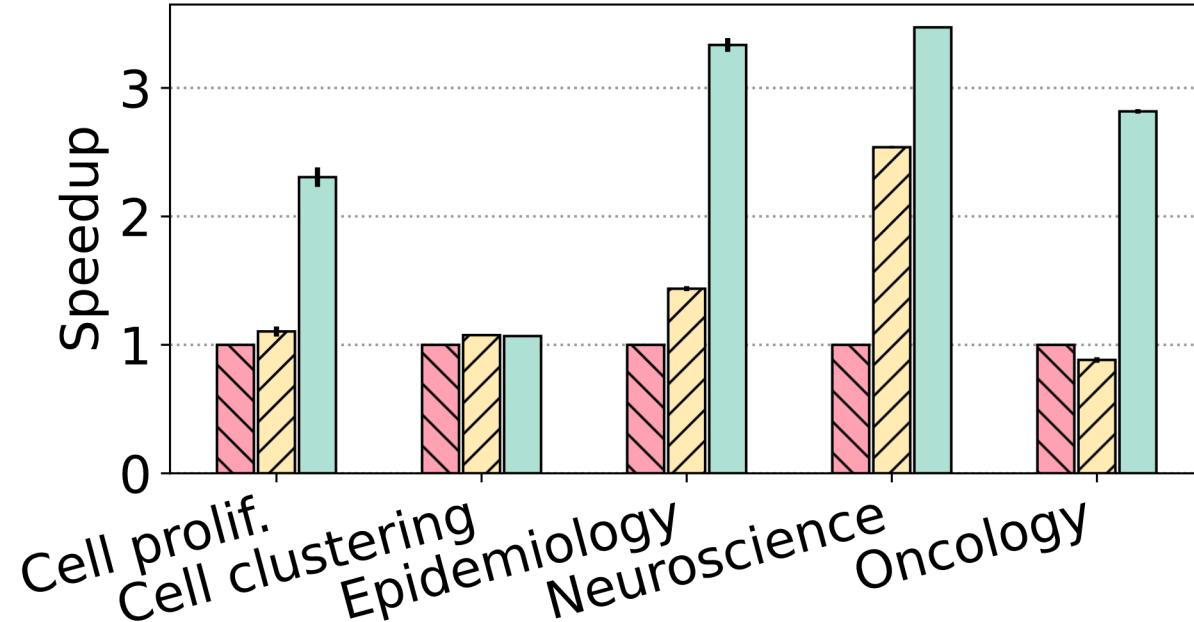
Radial neighbor search comparison

Build time (18 threads and one NUMA domain)



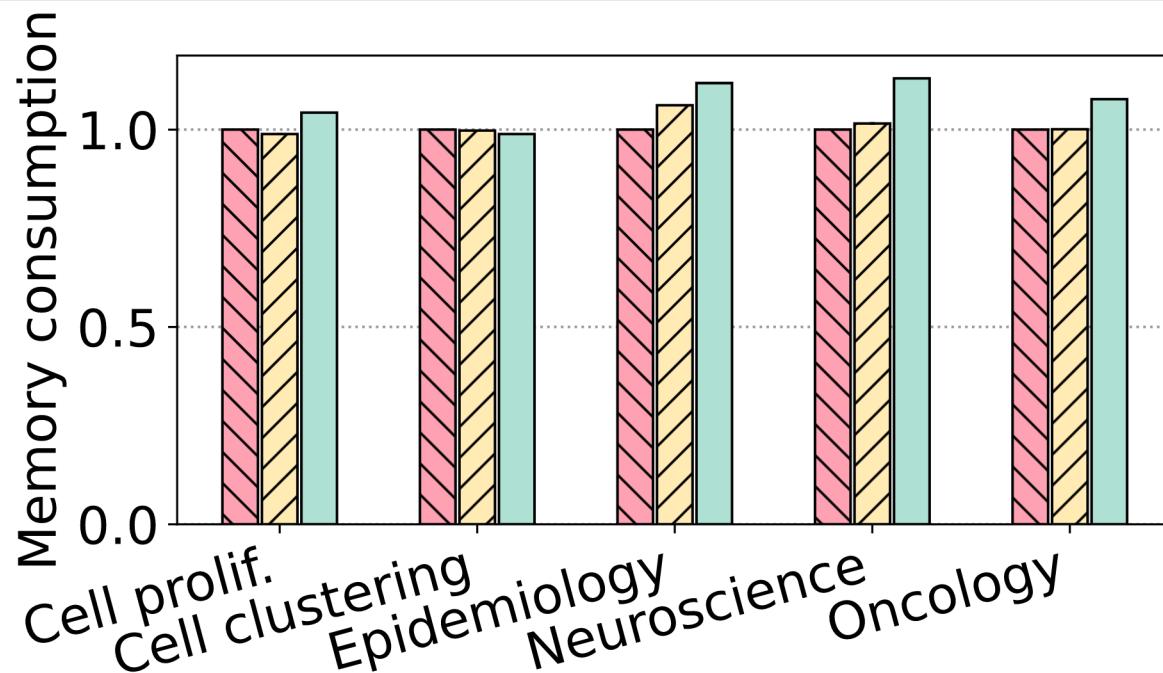
Radial neighbor search comparison

Search time indirect (18 threads and one NUMA domain)

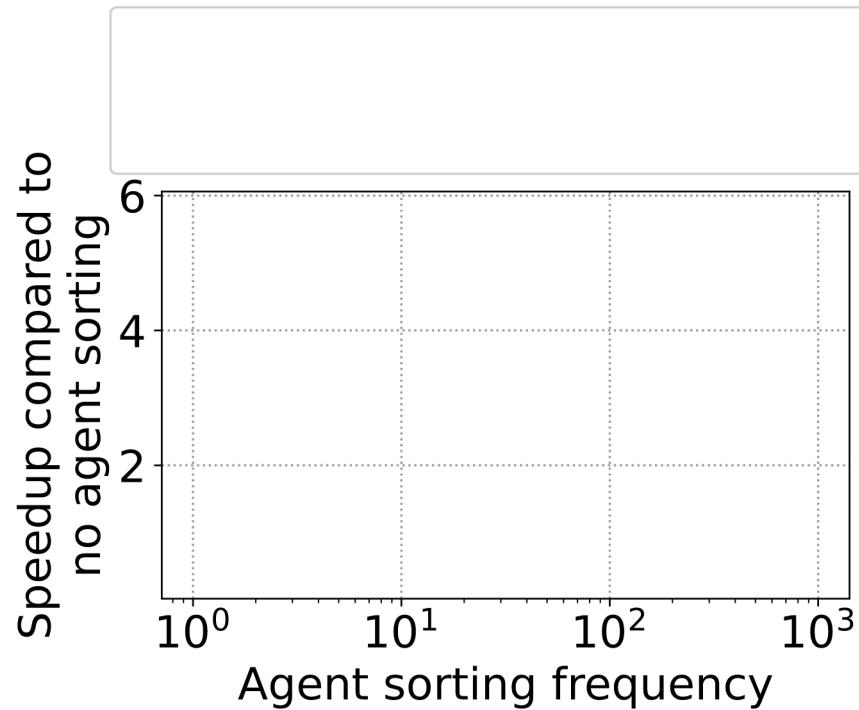


Radial neighbor search comparison

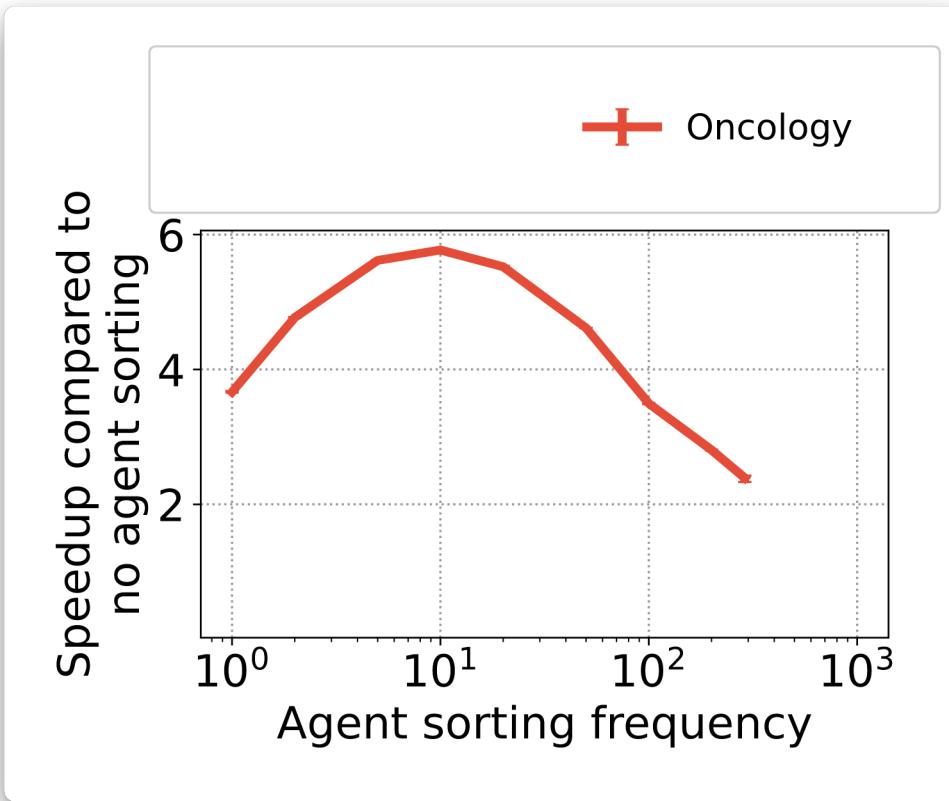
Memory consumption (18 threads and one NUMA domain)



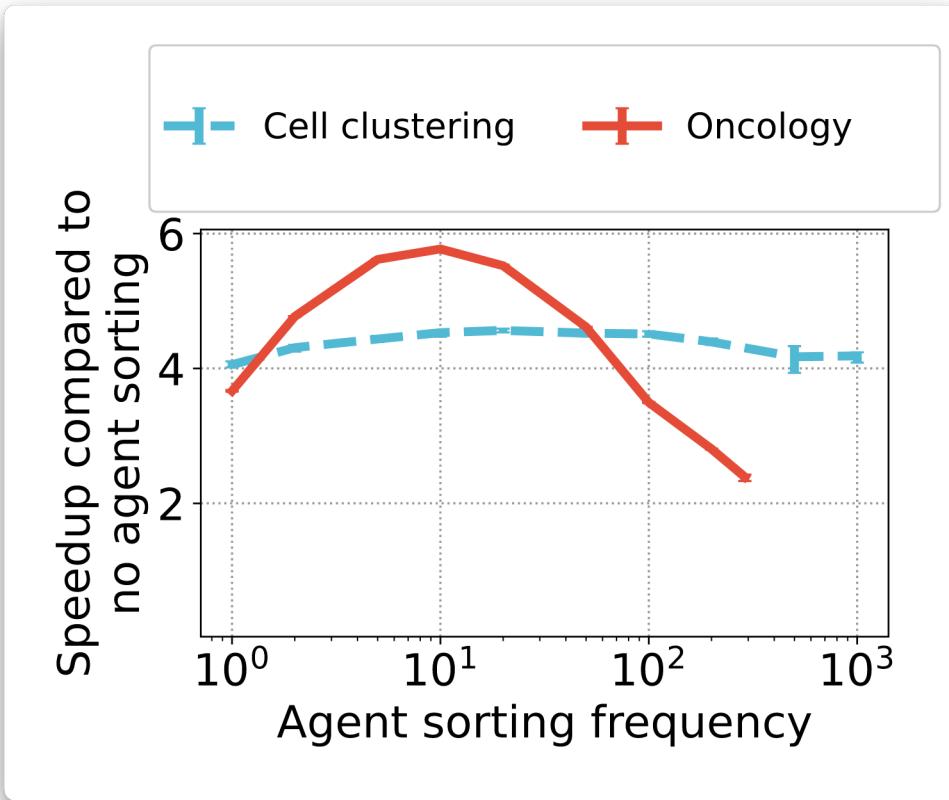
Agent sorting speeds up simulations up to $6\times$



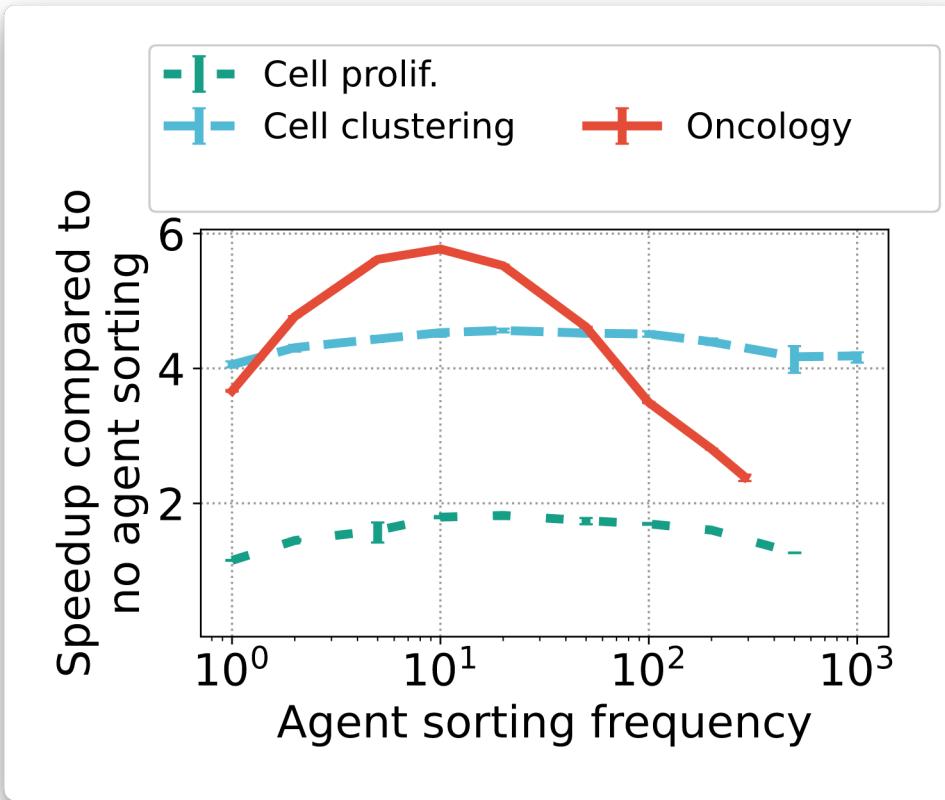
Agent sorting speeds up simulations up to 6×



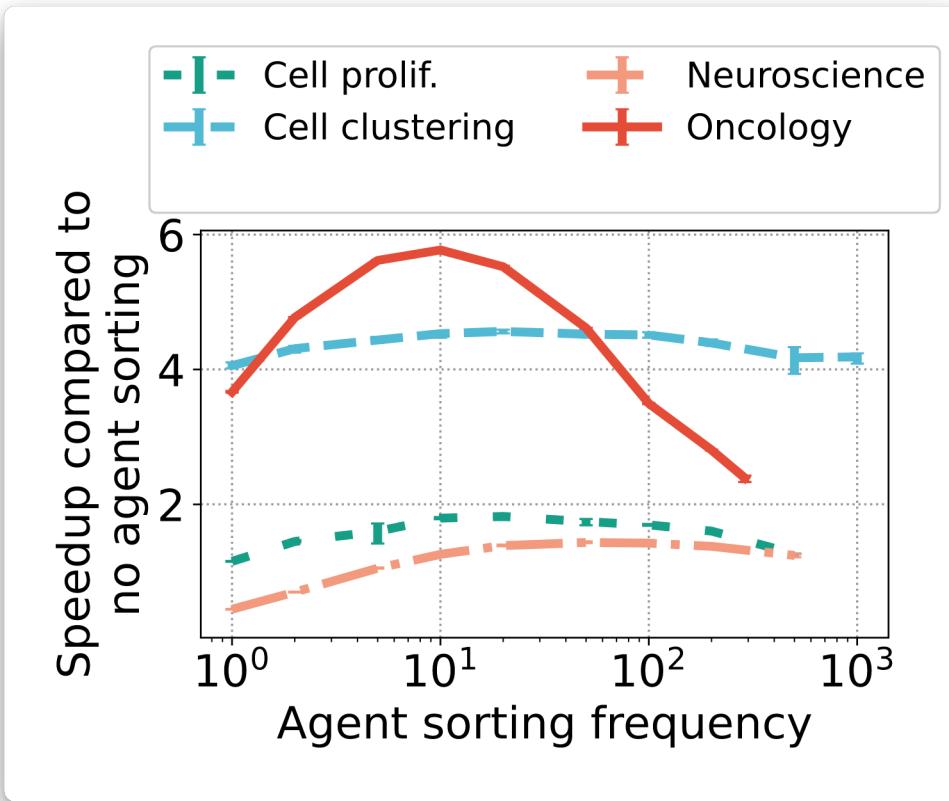
Agent sorting speeds up simulations up to $6\times$



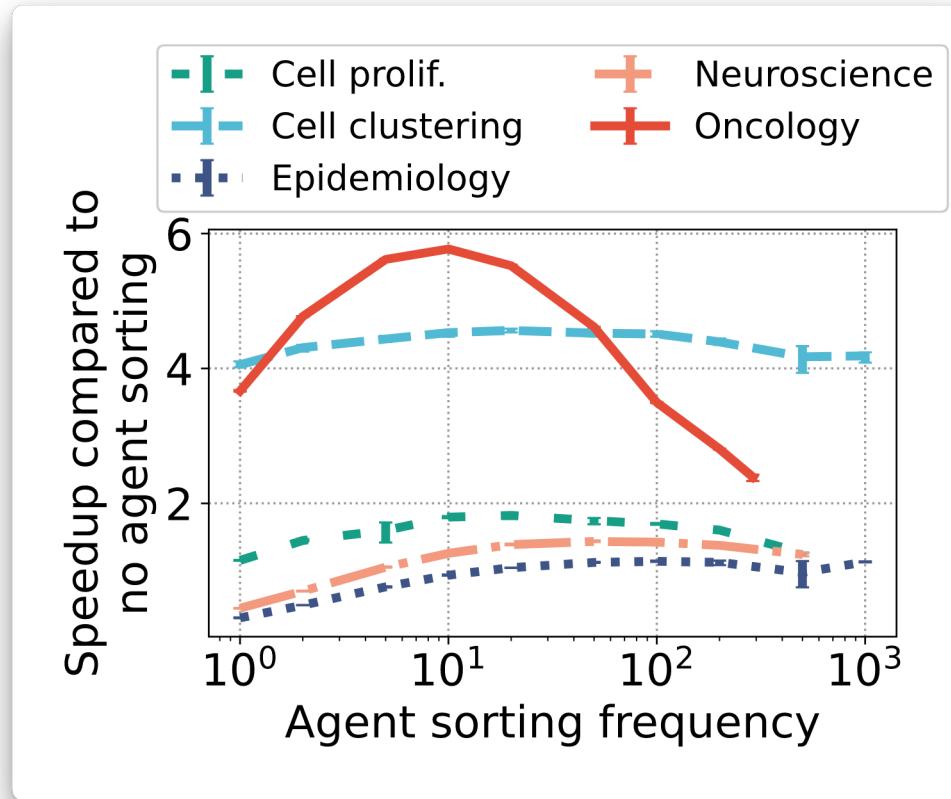
Agent sorting speeds up simulations up to 6×



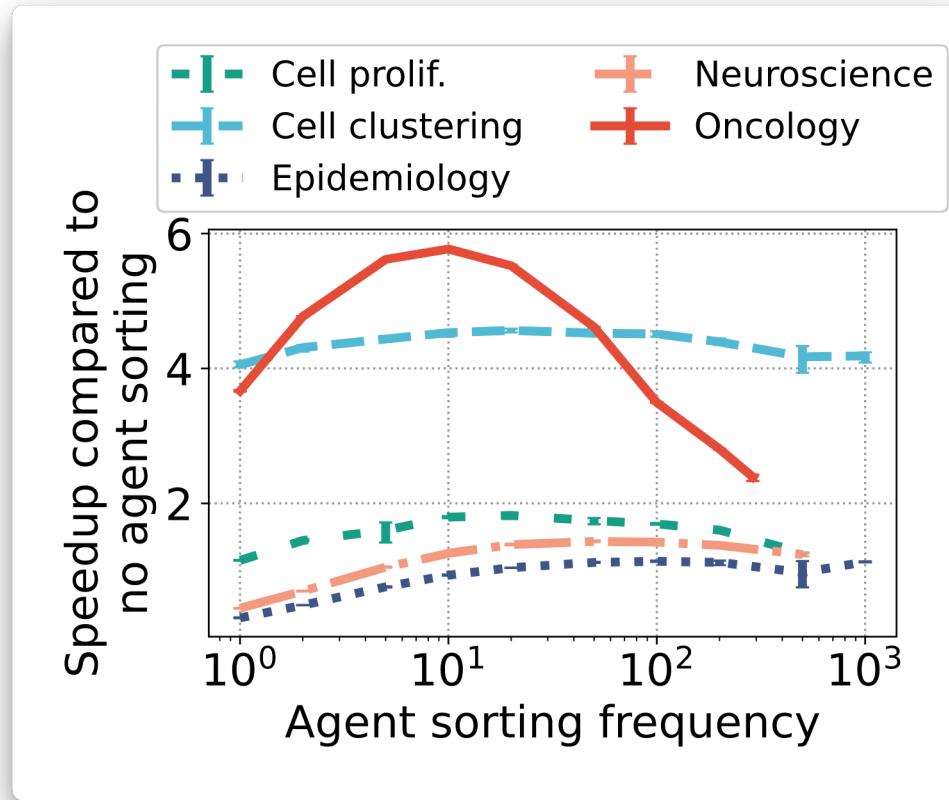
Agent sorting speeds up simulations up to 6×



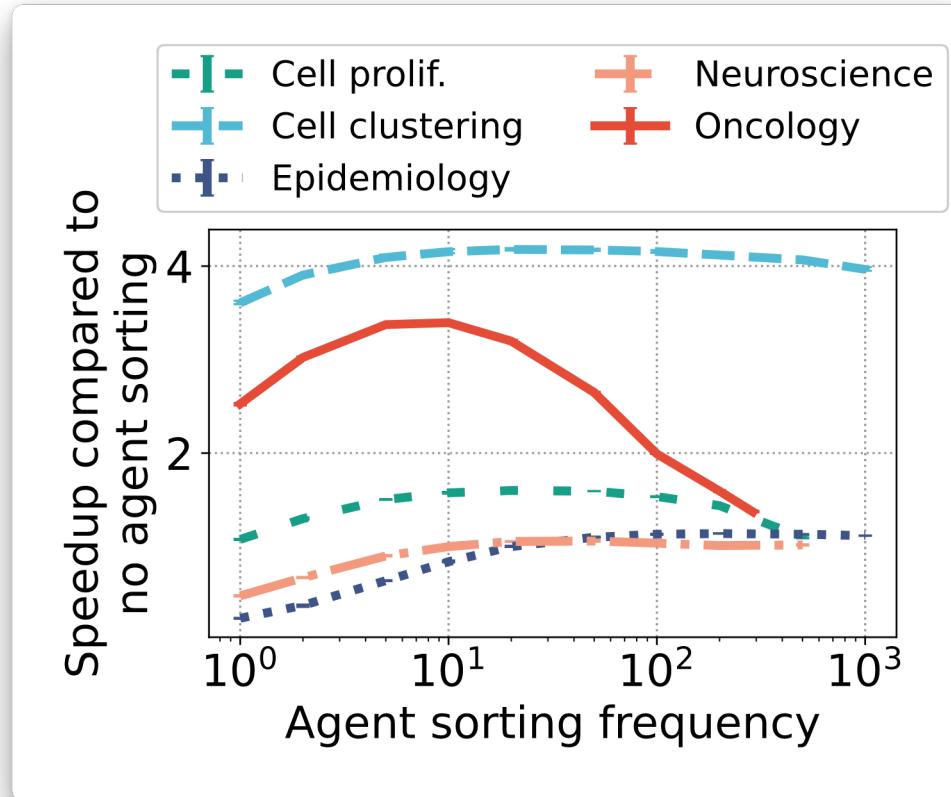
Agent sorting speeds up simulations up to 6×



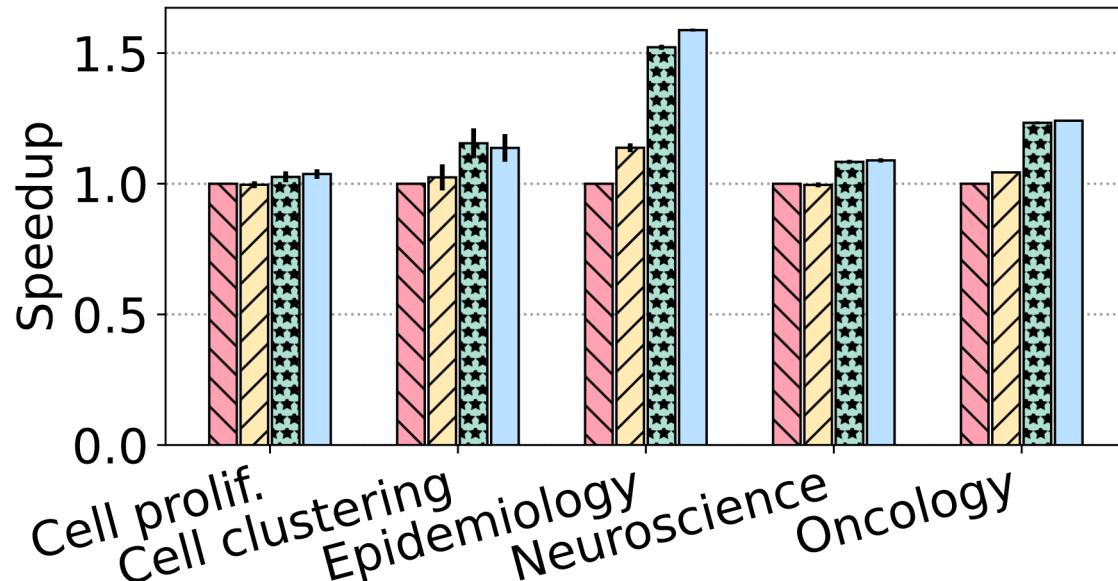
Agent sorting speeds up simulations up to 6×



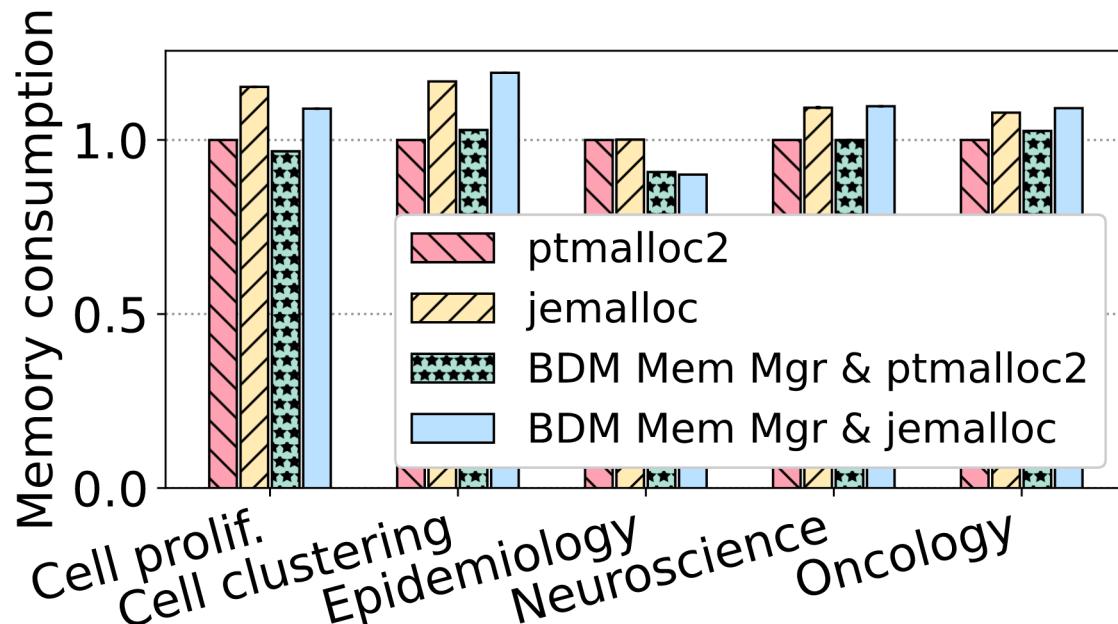
Agent sorting speedup for one NUMA domain



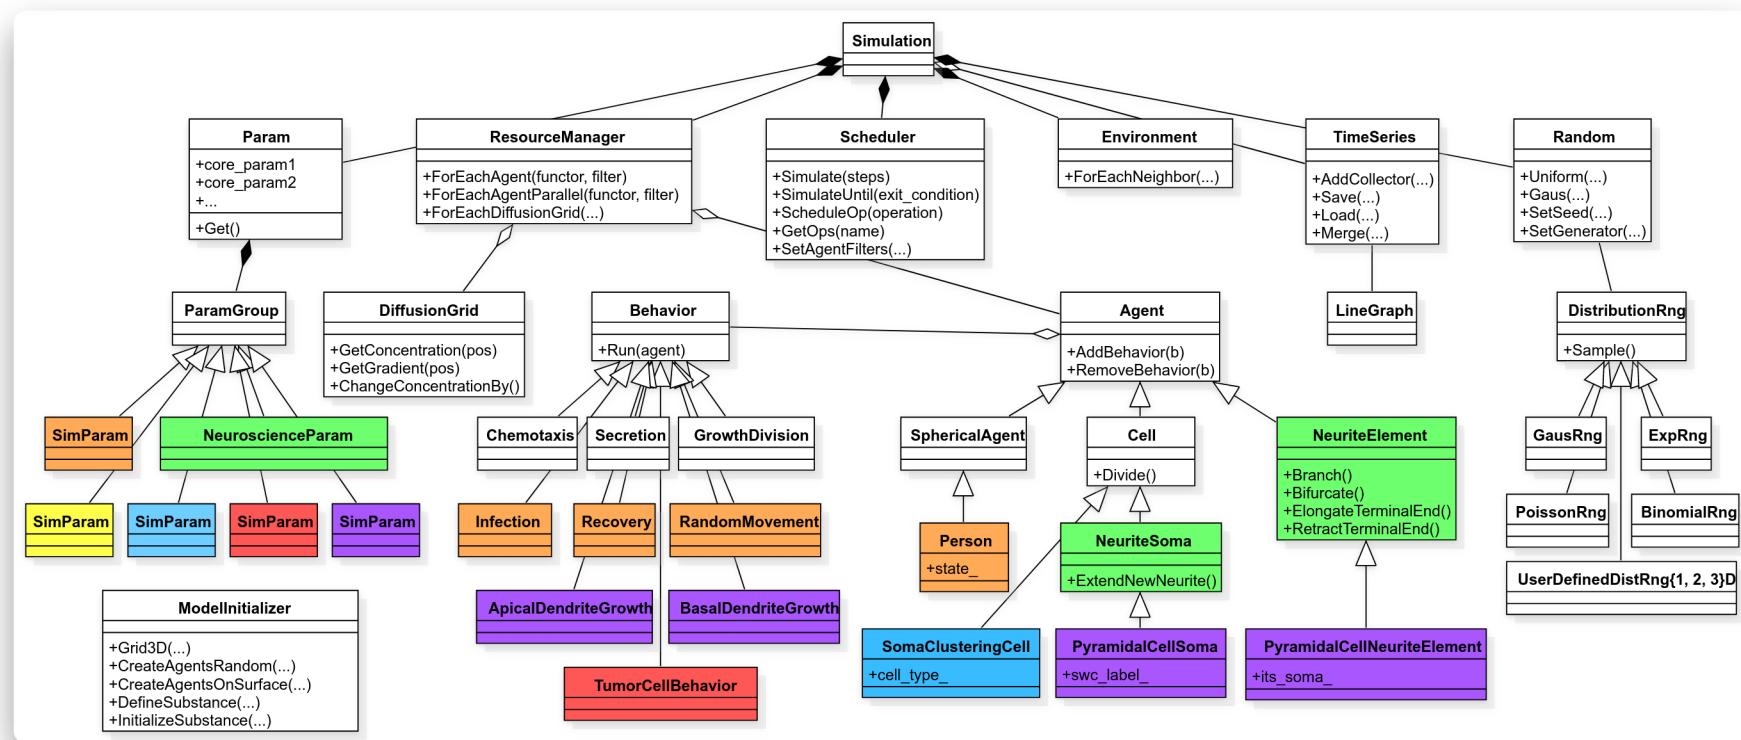
Memory allocator comparison



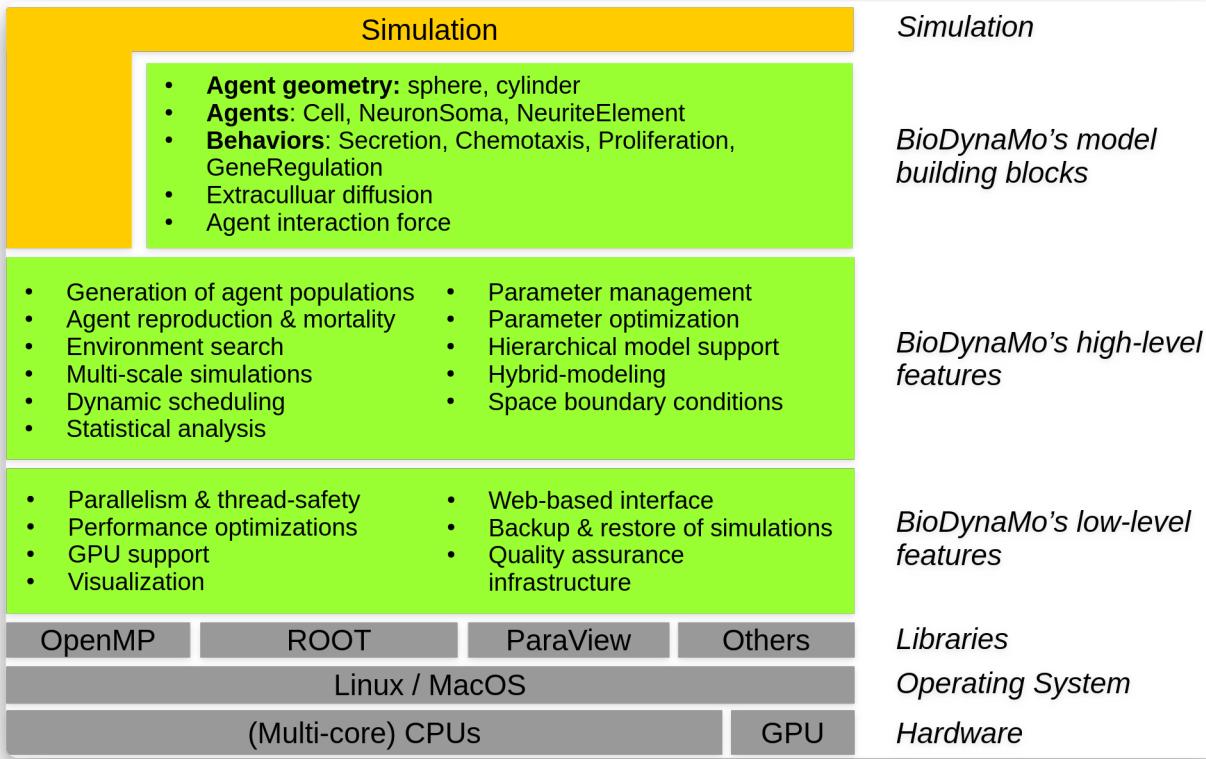
Memory allocator comparison



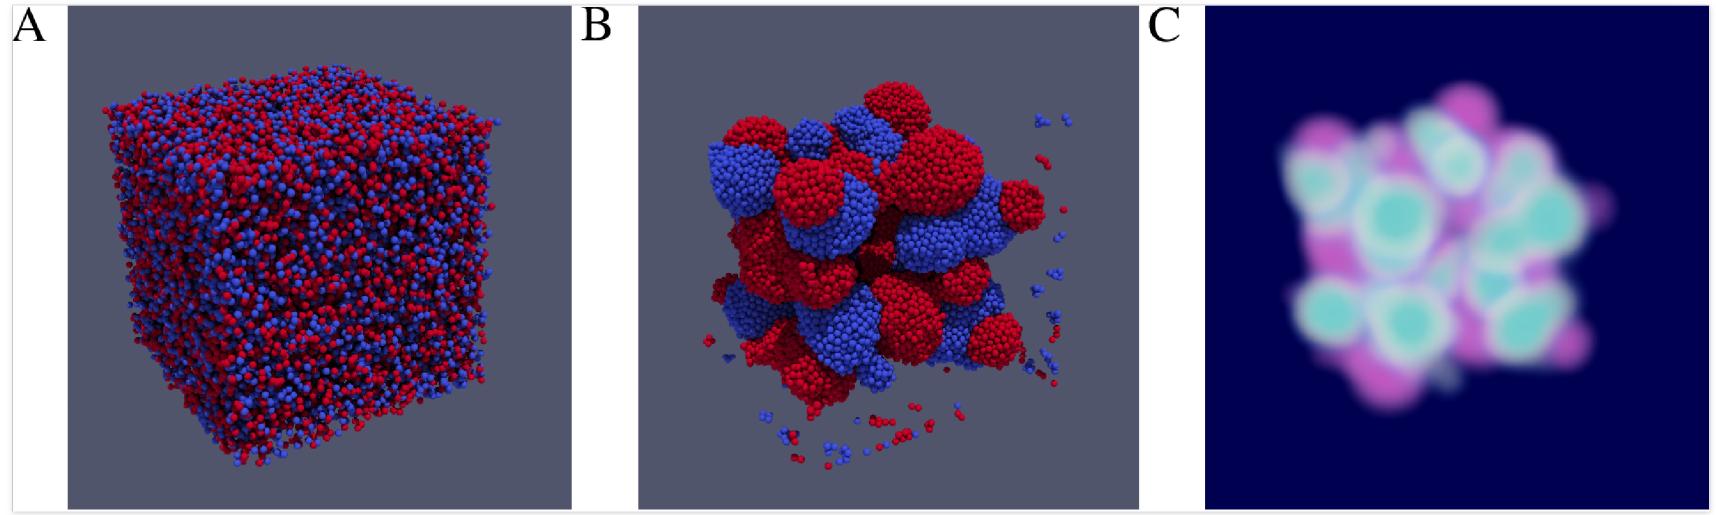
Modularity



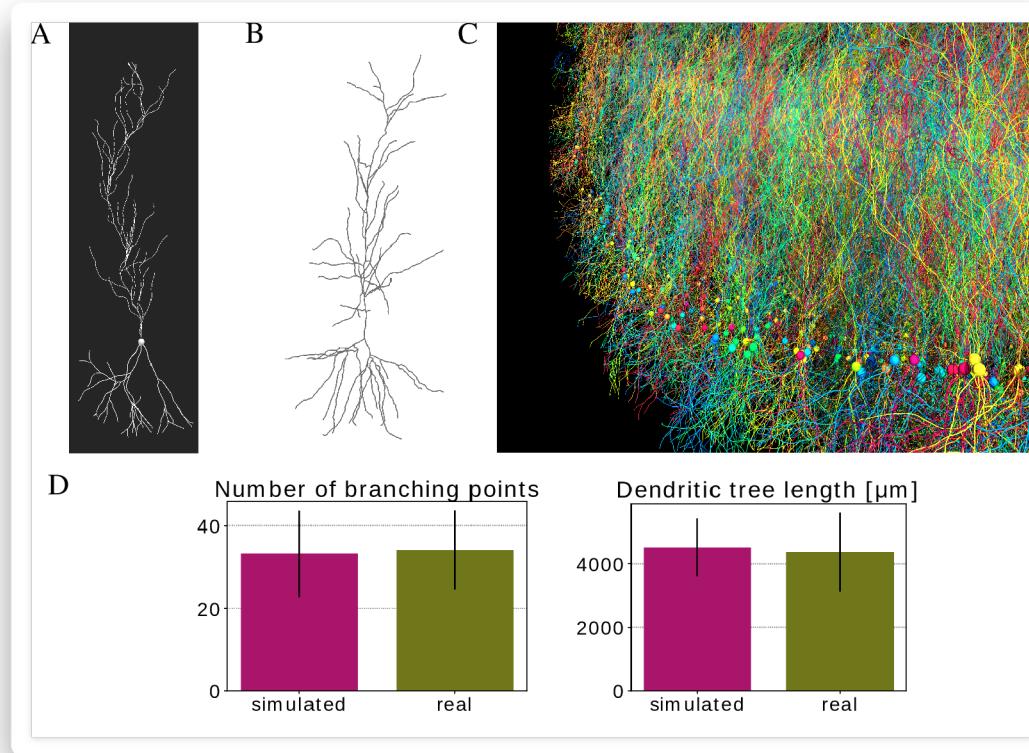
Abstraction layers and modeling features



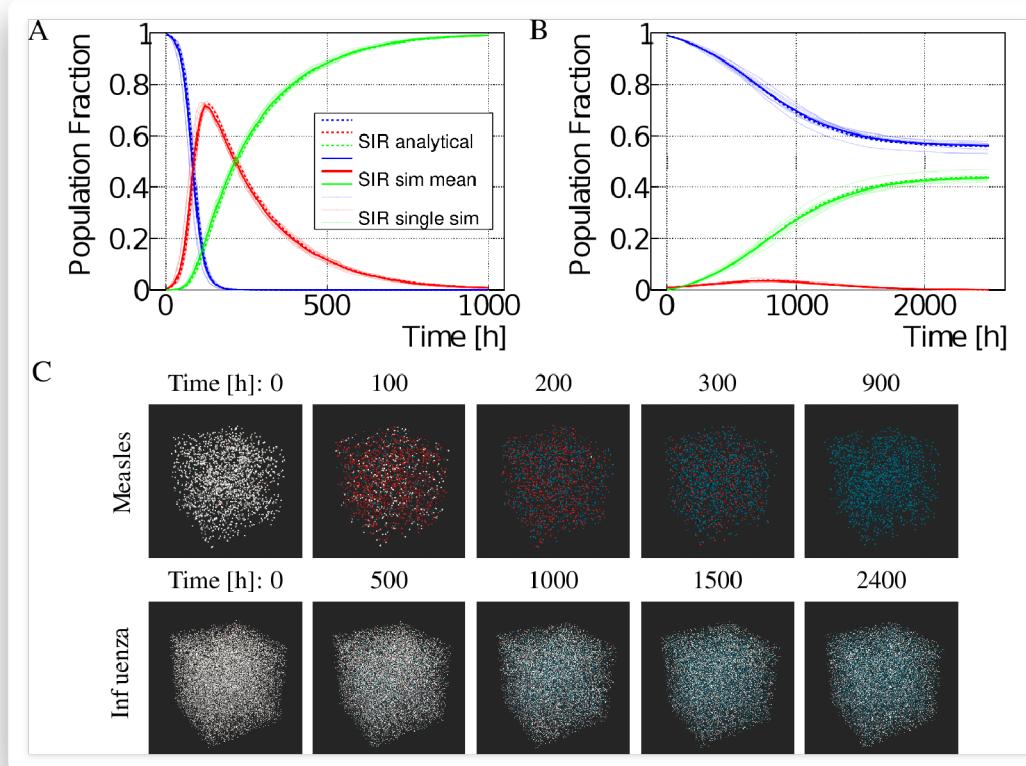
Cell sorting



Neuroscience use case

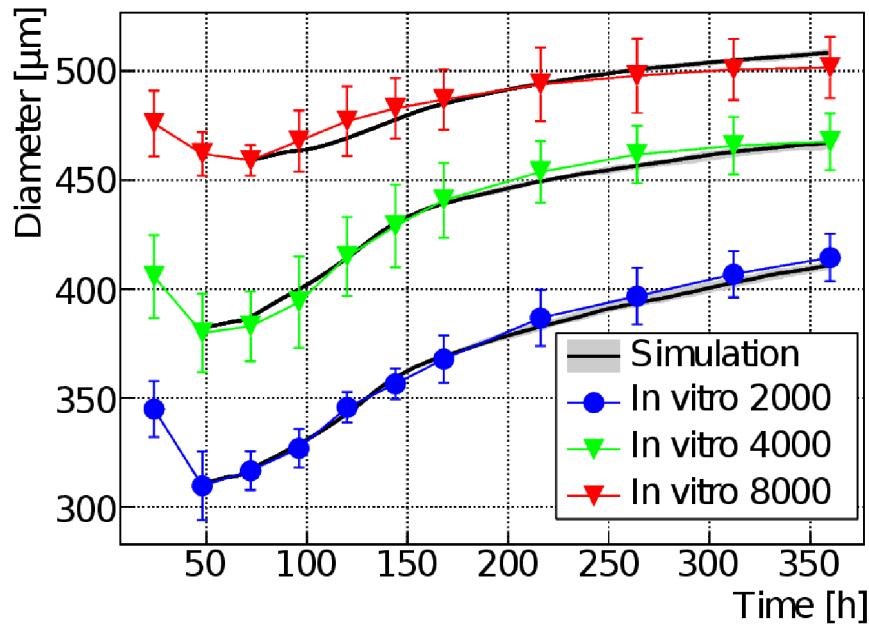


Epidemiology use case

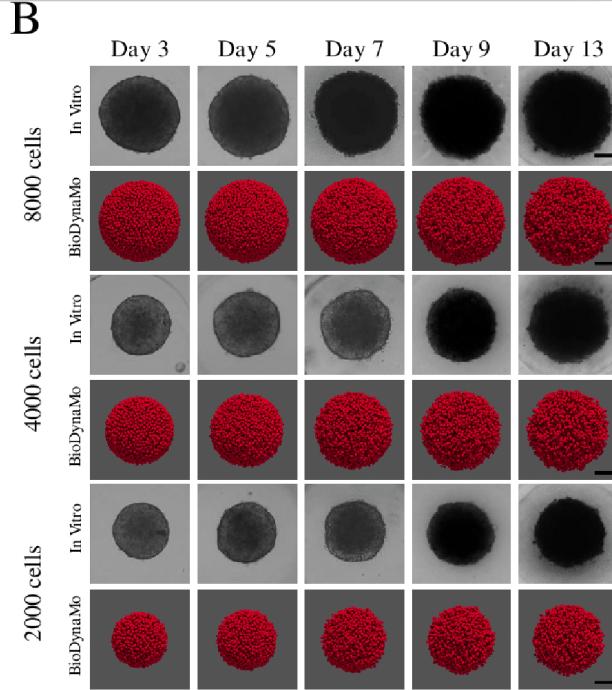


Oncology use case

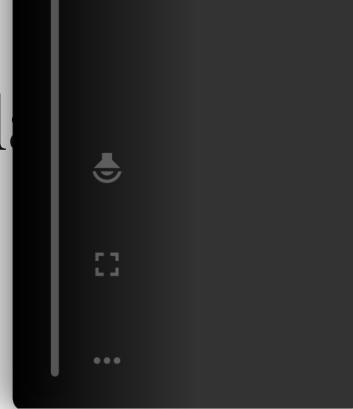
A



B



Pyramid with video



Performance data table

Table 6. Performance data. The values in column “Agents” and “Diffusion volumes” are taken from the end of the simulation. Runtime measures the wall-clock time to simulate the number of iterations. It excludes the time for simulation setup and visualization.

Simulation	Agents	Diffusion volumes	Iterations	System (Table 5)	Physical CPUs	Runtime	Memory
Neuroscience use case							
Single (Figure 4A in the main manuscript)	1 494	250	500	A	1	0.16 s	382 MB
				D	1	0.12 s	479 MB
Large-scale (Figure 4C in the main manuscript)	9 036 986	65 536	500	A	72	35 s	6.47 GB
				D	2	11 min 28 s	5.37 GB
Very-large-scale	1 018 644 154	5 606 442	500	B	72	1 h 24 min	438 GB
Oncology use case (Figure 5 in the main manuscript)							
2000 initial cells	4 177	0	312	A	1	1.05 s	382 MB
				D	1	0.832 s	480 MB
4000 initial cells	5 341	0	312	A	1	1.76 s	382 MB
				D	1	1.34 s	480 MB
8000 initial cells	7 861	0	288	A	1	3.27 s	384 MB
				D	1	2.60 s	482 MB
Large-scale	1 000 3925	0	288	A	72	1 min 42 s	7.42 GB
				D	2	43 min 56 s	5.84 GB
Very-large-scale	986 054 868	0	288	B	72	6 h 21 min	604 GB
Epidemiology use case (Figure 6C in the main manuscript)							
Measles	2 010	0	1000	A	1	0.53 s	381 MB
				D	1	0.42 s	479 MB
Seasonal Influenza	20 200	0	2500	A	1	16.41 s	383 MB
				D	1	16.40 s	479 GB
Medium-scale (measles)	100 500	0	1000	A	72	1.36 s	1 GB
Large-scale (measles)	10 050 000	0	1000	A	72	59.19 s	5.87 GB
				D	2	19 min 18 s	5.41 GB
Very-large-scale (measles)	1 005 000 000	0	1000	B	72	2 h 0 min	495 GB
Soma clustering (Figure 2)	32 000	1 240 000	6 000	A	72	12.91 s	1.02 GB
				D	2	2 min 7 s	522 MB

BioDynaMo's GPU offloading capabilities

Bioinformatics, 38(2), 2022, 453–460
doi: 10.1093/bioinformatics/btab649
Advance Access Publication Date: 16 September 2021
Original Paper

OXFORD

Systems biology
BioDynaMo: a modular platform for high-performance agent-based simulation

Lukas Breitwieser  ^{1,2,*}, Ahmad Hesam ^{1,3,*}, Jean de Montigny¹, Vasileios Vavourakis^{4,5}, Alexandros Iosif⁴, Jack Jennings⁶, Marcus Kaiser^{6,7,8}, Marco Manca  ⁹, Alberto Di Meglio¹, Zaid Al-Ars³, Fons Rademakers¹, Onur Mutlu^{2,10,*} and Roman Bauer^{11,*}

¹CERN openlab, IT Department, CERN, Geneva 1211, Switzerland, ²Department of Computer Science, ETH Zurich, Zurich 8092, Switzerland, ³Department of Quantum & Computer Engineering, Delft University of Technology, Delft 2628CD, The Netherlands, ⁴Department of Mechanical & Manufacturing Engineering, University of Cyprus, Nicosia 2109, Cyprus, ⁵Department of Medical Physics & Biomedical Engineering, University College London, London WC1E 6BT, UK, ⁶School of Computing, Newcastle University, Newcastle upon Tyne NE4 5TG, UK, ⁷Department of Functional Neurosurgery, Ruijin Hospital, Shanghai Jiao Tong University School of Medicine, Shanghai 200025, China, ⁸Precision Imaging Beacon, School of Medicine, University of Nottingham, Nottingham NG7 2UH, UK, ⁹ScimPulse Foundation, Geleen 6162 BC, The Netherlands, ¹⁰Department of Information Technology and Electrical Engineering, ETH Zurich, Zurich 8092, Switzerland and ¹¹Department of Computer Science, University of Surrey, Guildford GU2 7XH, UK

*To whom correspondence should be addressed.
Associate Editor: Jonathan Wren

Received on February 4, 2021; revised on September 2, 2021; editorial decision on September 3, 2021; accepted on September 13, 2021

GPU Acceleration of 3D Agent-Based Biological Simulations

Ahmad Hesam
ABS group
Delft University of Technology
Delft, Netherlands
a.hesam@tudelft.nl

Lukas Breitwieser
CERN openlab
CERN
Geneva, Switzerland
lukas.breitwieser@cern.ch

Fons Rademakers
CERN openlab
CERN
Geneva, Switzerland
fons.rademakers@cern.ch

Zaid Al-Ars
ABS group
Delft University of Technology
Delft, Netherlands
z.al-ars@tudelft.nl

Abstract—Researchers in biology are faced with the tough challenge of developing high-performance computer simulations of their increasingly complex agent-based models. BioDynaMo is an open-source agent-based simulation platform that aims to alleviate researchers from the intricacies that go into the development of high-performance computing. Through a high-level interface, researchers can implement their models on top of BioDynaMo's multi-threaded core execution engine to rapidly develop simulations that effectively utilize parallel computing hardware. In biological agent-based modeling, the type of operations that are typically the most compute-intensive are those that involve agents interacting with their local neighborhood. In this work, we investigate the currently implemented method of handling neighborhood interactions of cellular agents in BioDynaMo, and ways to improve the performance to enable large-scale and complex simulations. We propose to replace the kd-tree implementation to find and iterate over the neighborhood of each agent with a uniform grid method that allows us to take advantage of the massively parallel architecture of graphics processing units (GPUs). We implement the uniform grid method in both CUDA and OpenCL to address GPUs from all major vendors and evaluate several techniques to further improve the performance. Furthermore, we analyze the performance of our implementations for models with a varying density of neighboring agents. As a result, the performance of the mechanical interactions method improved by up to two orders of magnitude in comparison to the multithreaded baseline version. The implementations are open-source and publicly available on GitHub.

Index Terms—agent-based modeling, simulation, GPU, co-processing, biological models, acceleration

becoming increasingly more parallelized as a result of Dennard scaling [4] and the stagnation of Moore's law [5], as pointed out in [6]. Moreover, general-purpose computing on graphics processing units (GPUs) is an attractive solution to improve the computational efficiency of ABS applications in particular [7], [8], and parallel applications in general [9], [10]. By porting applications to, either fully or partially, run on GPUs it is possible to observe speedups of several orders of magnitude in comparison to the CPU-only execution [11]. Although several ABS frameworks exist that achieve significant speedups using GPUs in the field of ABS, there is still significant room for improvement, which we wish to address in this article.

BioDynaMo [6] is an open-source software platform for life scientists for simulating biological agent-based models. Each agent in BioDynamO is programmed to follow a specified set of rules, imposed by the modeler, that can trigger specified actions affecting itself or other agents. Agents in biological systems often interact with their local environment, and their behavior can be influenced by other agents that reside within a certain range. An example is the mechanical interactions a cellular agent undergoes when it physically collides with another agent. Local interactions are an extremely important concept in biological systems since it is the driving force behind key biological processes, such as tissue development [12].

BioDynaMo is fully parallelized using OpenMP and its

- Lukas Breitwieser et al., 2022, DOI: 10.1093/bioinformatics/btab649
- Ahmad Hesam et al., 2021, DOI: 10.1109/IPDPSW52791.2021.00040