

Parallel Simulation of Discrete Events

with special attention to simulation of the immune system

J.P. Daigle

Department of Computer Science
Georgia State University

Parallel Algorithms, Spring, 2006

Outline

Parallel Simulation

Simulating the Immune System

The Immune System
Simulation
Challenges

IMMSIM

the IMMSIM model
Improvements in ParIMM
Problems for Each Processor
The Parallel Process

Conclusion

overall conclusions
For Further Reading

1 Simulating the Immune System

- The Immune System
- Simulation Challenges

2 IMMSIM

- the IMMSIM model
- Improvements in ParIMM
- Problems for Each Processor
- The Parallel Process

3 Conclusion

- overall conclusions

Our motives for wanting to simulate the immune system are:

- 1 To cure cancer: immune system specific challenges
- 2 To make power grids more stable: applying the lessons to other physical systems.
- 3 Complex adaptive systems are just an interesting challenge

Understanding the Immune System

Parallel
Simulation

Simulating the
Immune
System

The Immune System
Simulation
Challenges

IMMSIM

the IMMSIM model
Improvements in
PartIMM

Problems for Each
Processor

The Parallel Process

Conclusion

overall conclusions
For Further Reading

The immune system of vertebrates consists of two parts
Innate Immune System The innate immune system is the first line of defense against common infections.

Includes:

- skin, saliva, mucos, outer barriers
- macrophages, phagocytes, non-discriminate inner barriers

Adaptive Immune System The adaptive immune system acquires new information over an organisms life.

Includes:

- bone marrow
- the lymphatic system (lymph, lymphatics, lymph nodes, spleen, thymus, etc.)
- specialized cells (T-Cells, B-Cells, mast cells, macrophages, dendritic cells, more)

The Adaptive Immune System

Parallel Simulation

Simulating the Immune System

- The Immune System
- Simulation
- Challenges

IMMSIM

- the IMMSIM model
- Improvements in PartIMM
- Problems for Each Processor
- The Parallel Process

Conclusion

- overall conclusions
- For Further Reading

The adaptive immune system is what we are interested in simulating. In particular, we are interested in acquired immunity, such as for chicken pox. This is a complicated process that involves many different components, all with a lifecycle and a variety of different states.

The Immune Response

Parallel Simulation

Simulating the Immune System

The Immune System Simulation Challenges

IMMSIM

the IMMSIM model Improvements in PartIMM

Problems for Each Processor

The Parallel Process

Conclusion

overall conclusions For Further Reading

A simple view of an immune response:

- 1 A novel *pathogen* enters the body, through a breach in the skin or other outer defense.
- 2 The innate immune system begins responding to the invader. Dendritic cells engulf the pathogen and carry antigen to the lymph nodes.
- 3 The dendritic cells present antigen to T-Cells in the lymph nodes.
- 4 The T-Cells activate B-Cells, which in turn produce antibodies.
- 5 The antibodies return to the site of the infection and destroy the pathogens

Adaptation

Parallel Simulation

Simulating the Immune System

- The Immune System
- Simulation
- Challenges

IMMSIM

- the IMMSIM model
- Improvements in PartIMM
- Problems for Each Processor
- The Parallel Process

Conclusion

- overall conclusions
- For Further Reading

What is missing from the previous slide is the knowledge that the immune system will respond much more quickly to a known pathogen than a novel pathogen. This is because some B-cells become *memory cells*, which retain the ability to quickly make particular antibodies. For new antibodies, populations of B-Cells have to undergo mutation in order to “find” the correct antibody.

Simulation of Complex Adaptive Systems

Parallel Simulation

Simulating the Immune System

The Immune System
Simulation Challenges

IMMSIM

the IMMSIM model
Improvements in ParIMM
Problems for Each Processor
The Parallel Process

Conclusion

overall conclusions
For Further Reading

- A typical immune response involves about 10^{12} cells of different kinds.
- Large populations are traditionally simulated through the use of ordinary differential equations, which average the behavior of large populations.
- But the history and lineage of each cell in the simulation is important to determine when and where antibodies will be produced.
- Another approach is needed.

Simulation of Complex Adaptive Systems

Parallel Simulation

Simulating the Immune System

The Immune System
Simulation Challenges

IMMSIM

the IMMSIM model
Improvements in ParIMM
Problems for Each Processor
The Parallel Process

Conclusion

overall conclusions
For Further Reading

- A typical immune response involves about 10^{12} cells of different kinds.
- Large populations are traditionally simulated through the use of ordinary differential equations, which average the behavior of large populations.
- But the history and lineage of each cell in the simulation is important to determine when and where antibodies will be produced.
- Another approach is needed.

Simulation of Complex Adaptive Systems

Parallel Simulation

Simulating the Immune System

The Immune System
Simulation
Challenges

IMMSIM

the IMMSIM model
Improvements in
ParIMM
Problems for Each
Processor
The Parallel Process

Conclusion

overall conclusions
For Further Reading

- A typical immune response involves about 10^{12} cells of different kinds.
- Large populations are traditionally simulated through the use of ordinary differential equations, which average the behavior of large populations.
- But the history and lineage of each cell in the simulation is important to determine when and where antibodies will be produced.
- Another approach is needed.

Simulation of Complex Adaptive Systems

Parallel Simulation

Simulating the Immune System

The Immune System
Simulation
Challenges

IMMSIM

the IMMSIM model
Improvements in
PartIMM
Problems for Each
Processor
The Parrallel Process

Conclusion

overall conclusions
For Further Reading

- A typical immune response involves about 10^{12} cells of different kinds.
- Large populations are traditionally simulated through the use of ordinary differential equations, which average the behavior of large populations.
- But the history and lineage of each cell in the simulation is important to determine when and where antibodies will be produced.
- Another approach is needed.

Simulation of Complex Adaptive Systems

Parallel Simulation

Simulating the Immune System

The Immune System
Simulation
Challenges

IMMSIM

the IMMSIM model
Improvements in
PartIMM
Problems for Each
Processor
The Parrallel Process

Conclusion

overall conclusions
For Further Reading

- A typical immune response involves about 10^{12} cells of different kinds.
- Large populations are traditionally simulated through the use of ordinary differential equations, which average the behavior of large populations.
- But the history and lineage of each cell in the simulation is important to determine when and where antibodies will be produced.
- Another approach is needed.

Discrete Event Simulation

Parallel Simulation

Simulating the Immune System

The Immune System Simulation Challenges

IMMSIM

the IMMSIM model Improvements in PartIMM

Problems for Each Processor

The Parallel Process

Conclusion

overall conclusions

For Further Reading

- Models each event at a particular level of granularity separately
- Draws on fields such as cellular automata and Object Oriented Programming
- Drawbacks include loss of efficiency
- Advantages include increase in accuracy

Simulating the Immune System with IMMSIM

Parallel Simulation

Simulating the Immune System

The Immune System Simulation Challenges

IMMSIM

the IMMSIM model

Improvements in PartIMM

Problems for Each Processor

The Parallel Process

Conclusion

overall conclusions

For Further Reading

- The first DES for the IS: 1992 (Kleinstein and Seiden)
- Based on a modified cellular automaton
- Focuses on a particular behavior, the original scope is limited to portions of individual lymph nodes

IMMSIM

IMMSIM is represented as a triangular lattice, each node of which may contain 0 or more entities such as T-Cells, B-Cells, antibodies, or antigens.

Simulating the Immune System with IMMSIM

Parallel Simulation

Simulating the Immune System

The Immune System Simulation Challenges

IMMSIM

the IMMSIM model

Improvements in PartIMM

Problems for Each Processor

The Parallel Process

Conclusion

overall conclusions

For Further Reading

- The first DES for the IS: 1992 (Kleinstein and Seiden)
- Based on a modified cellular automaton
- Focuses on a particular behavior, the original scope is limited to portions of individual lymph nodes

IMMSIM

IMMSIM is represented as a triangular lattice, each node of which may contain 0 or more entities such as T-Cells, B-Cells, antibodies, or antigens.

Cellular Automata

Parallel Simulation

Simulating the Immune System

The Immune System Simulation Challenges

IMMSIM

the IMMSIM model

Improvements in PartIMM

Problems for Each Processor

The Parallel Process

Conclusion

overall conclusions

For Further Reading

Cellular Automata

- consist of a discrete lattice of sites
- evolve in discrete time steps
- each site takes on a finite set of possible values
- each site evolves according to the same *deterministic rules*
- The rules for evolution depend on a states in the *local neighborhood*

IMMSIM

- consists of a triangular lattice of sites
- evolves in discrete time steps
- each site takes on a finite set of possible values
- each site evolves according to the same *probabalistic rules*
- The rules for evolution depend on the state of the site

Other parts of the simulation

Parallel Simulation

Simulating the Immune System

The Immune System
Simulation
Challenges

IMMSIM

the IMMSIM model

Improvements in PartIMM

Problems for Each Processor

The Parallel Process

Conclusion

overall conclusions
For Further Reading

Entities

IMMSIM defines 6 cellular and 5 molecular entities. Cellular entities can be viewed as *stochastic finite state machines*. Simply, the state of an entity is determined by a probability function on its current state and a molecular input.

Repertoire

- binding cells have a receptor, a bit string of length l
- all entities have a number of distinct classes, equal to $2^{N_e l}$ where N_e is the number of receptors for the cell type.
- There are potentially 2^{21} different classes of B and T cells in the simulation. This is less than the natural lower limit of 2.5×10^7 , but its a start!

Other parts of the simulation

Parallel Simulation

Simulating the Immune System

The Immune System Simulation Challenges

IMMSIM

the IMMSIM model

Improvements in PartIMM

Problems for Each Processor

The Parallel Process

Conclusion

overall conclusions

For Further Reading

interactions

Interactions are defined by the following syntax:

```
INTERACTION:<involved entities>  
SPECIFIC:YES|NO  
MATCH:<involved molecule>  
CONDITION:<allowed state for <involved entities>  
ACTION : <final state for <involved entities>
```

mutation

Parallel Simulation

Simulating the Immune System

The Immune System
Simulation
Challenges

IMMSIM

the IMMSIM model

Improvements in PartIMM

Problems for Each Processor

The Parallel Process

Conclusion

overall conclusions

For Further Reading

- Each B-Cell produces antibodies, which must match a given antigen. This is accomplished by measuring the Hamming Distance between two strings, but...
- In the real IS, cells mutate to match antigens
- This must be simulated as well

the process of the simulation

Parallel Simulation

Simulating the Immune System

The Immune System
Simulation
Challenges

IMMSIM

the IMMSIM model

Improvements in
PartIMM

Problems for Each
Processor

The Parallel Process

Conclusion

overall conclusions

For Further Reading

- 1 population & antigen inputs chosen
- 2 all possible interactions in a site are examined
- 3 a subset of of interactions are simulated
- 4 new cells are born, divide, or die
- 5 diffusion of entities into neighboring sites
- 6 repeat from step two

Why make a parallel version?

- 1 To simulate the same space faster
- 2 To simulate much larger spaces (accuracy)

ParIMM, a high performance adaptation of IMMSIM, accomplishes both of these goals, but is not without its limitations.

- 1 Memory Allocation
- 2 Data extraction

Broad View of the Parallel Simulator

Parallel Simulation

Simulating the Immune System

The Immune System Simulation Challenges

IMMSIM

the IMMSIM model

Improvements in ParIMM

Problems for Each Processor

The Parallel Process

Conclusion

overall conclusions

For Further Reading

- The grid is divided among the processors
- Each processor keeps a list of its local entities
- Each processor calculates its local reactions and results
- Entities may move between processors as they travel through the grid
- Data is “dumped” at the end to a main processor

Data Storage

Parallel Simulation

Simulating the Immune System

The Immune System Simulation Challenges

IMMSIM

the IMMSIM model Improvements in PartIMM

Problems for Each Processor

The Parallel Process

Conclusion

overall conclusions For Further Reading

- Data Structure
 - Originally, Arrays were considered for cellular entities
 - Would require 2^{Cl} space, where l is the size of an entity in bits and C is a constant.
 - alternate choice was to use linked lists, one per entity type
- Single Linked Lists are accessed in only one order
 - This would introduce bias, some cells would react more than others
 - The lists are permuted before each round.
- These changes allow storage to be handled in $O(N)$ time and $O(N)$ space.

Afinity Potential

Parallel Simulation

Simulating the Immune System

The Immune System Simulation Challenges

IMMSIM

the IMMSIM model Improvements in PartIMM

Problems for Each Processor

The Parrallel Process

Conclusion

overall conclusions For Further Reading

Antibodies and Antigens are well matched if their strings oppose. This affinity potential is a function of the Hamming Distance between two strings.

```
int dist=0, k, xyxor
xyxor= x ^ y
for k=0; k<l; k++ do
    dist += ((xyxor » k) & 0x1)
```

This prefix sum of the bitwise difference approach is potentially very slow.

Solution:

- Build a hash table that contains the number of 1's for each integer 0 to $2^l - 1$
- A single bitwise XOR between the two strings produces the key
- The Hamming Distance is the entry at that key.

ParImm Communication Pattern

Parallel Simulation

Simulating the Immune System

The Immune System Simulation Challenges

IMMSIM

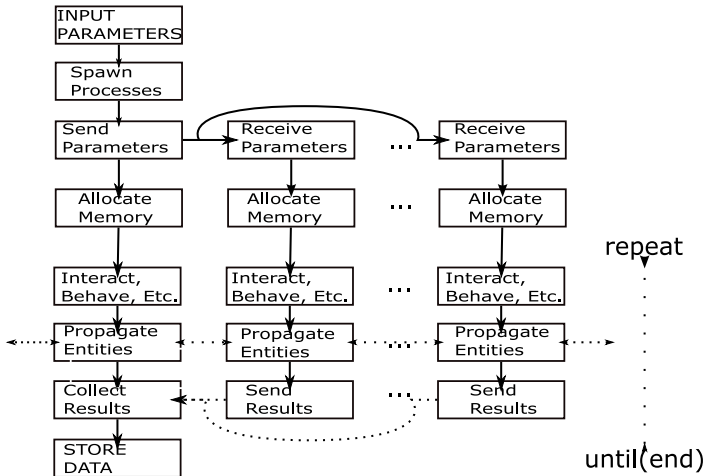
the IMMSIM model Improvements in ParImm

Problems for Each Processor

The Parallel Process

Conclusion

overall conclusions For Further Reading



Problems with this Approach

Parallel Simulation

Simulating the Immune System

The Immune System
Simulation
Challenges

IMMSIM

the IMMSIM model
Improvements in
ParImm

Problems for Each
Processor

The Parallel Process

Conclusion

overall conclusions
For Further Reading

Essentially, PARImm takes linear code and combines it with a message passing interface (PVM) to expand the number of sites that can be simulated.

- 1 The simulation is conservative
- 2 The simulation is time-stepped
- 3 The simulation is fixed

The simulation is conservative

Parallel Simulation

Simulating the Immune System

The Immune System Simulation Challenges

IMMSIM

the IMMSIM model

Improvements in PartIMM

Problems for Each Processor

The Parallel Process

Conclusion

overall conclusions

For Further Reading

- Each site must update and pass messages before the next move
- This can slow the simulation down, especially if most of the interactions site s in round $k + 1$ involve entities that were already in s during round k .
- This may be an appropriate place to use an *Optimistic* simulation approach.
- Recent work in this area suggests that this could be very beneficial.[4, 5]

The simulation is time-based

Parallel Simulation

Simulating the Immune System

The Immune System Simulation Challenges

IMMSIM

the IMMSIM model

Improvements in PartIMM

Problems for Each Processor

The Parallel Process

Conclusion

overall conclusions

For Further Reading

- Each entity must update its state in each pass.
- This may not be necessary. A particular cell may not engage in any interactions for many turns, and remain in a single site.
- This may be an appropriate problem to use a DEVS architecture, to make the simulation interaction based, rather than time stepped.[6]
- Some work has been done in this direction with a serial Java Implementation, CAFFISS.[1]

The simulation is fixed

Parallel Simulation

Simulating the Immune System

The Immune System
Simulation
Challenges

IMMSIM

the IMMSIM model
Improvements in
PartIMM

Problems for Each
Processor

The Parallel Process

Conclusion

overall conclusions
For Further Reading

- Each processor handles a fixed number of nodes, regardless of activity
- Processor allocation is done by card dealing
- There may be a way to distribute nodes dynamically.
- The advantage to this would be that those nodes with more activity would get more processors dedicated to them.
- Because activity *migrates* from node-to-node, it seems likely that this approach would work very well, especially if the architecture mimicked the lattice.

- This is an enormously complex problem. PARIMM only simulates a portion of the immune system, and that to a large degree of approximation (2^{21} vs 2.5×10^7)
- While good results have been obtained from this approach, better parallelization technique could increase both the size and accuracy of the simulation.

For Further Reading I

Parallel Simulation

Simulating the Immune System

The Immune System Simulation Challenges

IMMSIM

the IMMSIM model
Improvements in PartIMM
Problems for Each Processor
The Parallel Process

Conclusion

overall conclusions
For Further Reading



J. C. Tay and A. Jhavar, “CAFISS: a complex adaptive framework for immune system simulation,” in *SAC '05: Proceedings of the 2005 ACM symposium on Applied computing*, (New York, NY, USA), pp. 158–164, ACM Press, 2005.



M. Bernaschi and F. Castiglione, “Design and implementation of an immune system simulator,” *Computers in Biology and Medicine*, vol. 31, no. 5, pp. 303–331, 2001.



S. H. Kleinstein and P. E. Seiden, “Simulating the immune system,” *IEEE MultiMedia*, vol. 2, no. 4, pp. 69–77, 2000.

For Further Reading II

Parallel Simulation

Simulating the Immune System

The Immune System
Simulation
Challenges

IMMSIM

the IMMSIM model
Improvements in PartIMM
Problems for Each Processor
The Parallel Process

Conclusion

overall conclusions
For Further Reading



S. K. Prasad and Z. Cao, “Parallel distributed simulation and modeling methods: Syncsim: a synchronous simple optimistic simulation technique based on a global parallel heap event queue,” in *WSC '03: Proceedings of the 35th conference on Winter simulation*, pp. 872–880, Winter Simulation Conference, 2003.



R. M. Fujimoto, “Parallel simulation: distributed simulation systems,” in *WSC '03: Proceedings of the 35th conference on Winter simulation*, pp. 124–134, Winter Simulation Conference, 2003.

For Further Reading III

Parallel Simulation

Simulating the Immune System

The Immune System Simulation Challenges

IMMSIM

the IMMSIM model
Improvements in PartIMM
Problems for Each Processor
The Parallel Process

Conclusion

overall conclusions

For Further Reading



B. P. Zeigler and H. S. Sarjoughian, “Distributed simulation: creating distributed simulation using DEVS m&s environments,” in *WSC '00: Proceedings of the 32nd conference on Winter simulation*, (San Diego, CA, USA), pp. 158–160, Society for Computer Simulation International, 2000.