

MULTITHREADED AGENT-BASED SIMULATION

Michael E. Goldsby
Carmen M. Pancerella

7011 East Avenue
Sandia National Laboratories
Livermore, CA 94550, USA

ABSTRACT

Multithreading can significantly increase the performance of large agent-based simulations on multicore systems, but agent-based software packages do not commonly offer adequate support for multithreading. This report describes alterations and additions made to the MASON agent-based simulation package that allow the application programmer to make use of multiple threads easily and without radical change to conventional agent-based programming style. The report confirms performance gains with the results of test runs.

1 INTRODUCTION

The work reported here grew out of a project that needed the results of runs made with sizable populations (Armstrong, Schoeniger, Sale, Goldsby, and Mayo 2011). A multicore server was available, but we could find no agent-based simulation package with adequate support for multithreading. Some packages would allow the programmer to use multiple threads, but the data structures were not thread-safe, and they left all issues of division of work and avoidance of conflicting accesses up to the application programmer.

This report describes alterations and additions to the MASON (Luke, Cioffi-Revilla, Panait, Sullivan, and Balan 2005) agent-based simulation package that allow the application programmer to use threads with ease. We realized at the beginning that widespread use of locks would have been complicated and an almost sure-fire performance killer. As it turned out, we devised a scheme requiring no locks at all. The altered system shows good performance gains in the test runs.

The next section briefly describes agent-based simulation. Section 3 describes MASON, the agent-based simulation package chosen for this work. Section 4 describes the basic approach. Section 5 gives details of the implementation. Section 6 comments on related work. Section 7 reports the results of performance tests with the altered system. Finally, section 8 contains conclusions and suggestions for future work.

2 AGENT-BASED SIMULATION

In agent-based modeling (ABM), a system is modeled as a collection of autonomous decision-making entities (agents) that encapsulate the behavior of individuals in a system. Each agent individually assesses its state, makes decisions based on rules, and interacts with other agents. Collectively the agents emulate a real-world system. Multi-agent systems (MAS) or agent-based simulation (ABS) is the corresponding computational technique, where a software agent executes on a processor and communicates with other agents in the system, in order to model a complex system; see (Epstein and Axtell 1996), (Jennings 2000) and (Uhrmacher and Weyns 2009) for more detail. Each agent maintains its own state, and there is no centralized control.

Agent-based simulation has become an important method of modeling not only how individuals behave but also how the behavior and interactions of individuals affects the larger system. Thus, ABS is used to model socio-economic systems in economics, finance, supply chain, traffic modeling, social networking, and others.

There are a number of ABM software platforms that can be used to develop agent-based simulations (Railsback, Lytinen, and Jackson 2006). Two such agent-based platforms are Repast (North, Collier, and Vos 2006) and MASON, both implemented in Java. Repast is arguably the most popular agent-based simulation package in use and has active support. MASON is a much smaller, less mature, and leaner package, which nevertheless implements the core constructs needed for agent-based simulation. MASON was designed with execution speed a priority.

Such agent-based systems are typically single-threaded, with support for parallelism rudimentary or lacking. Both Repast and MASON have a minimal level of support for multithreading, but in neither are the data structures thread-safe, nor is there support for partitioning data among the threads. An exception is the high-performance computing version of Repast, written in C++, which uses MPI (MPI Forum 2013) as the underlying communication layer and runs a single user thread in each MPI process ??.

3 MASON

The agent-based system chosen for this work is MASON. MASON is a full-featured simulation package with a relatively small and well-commented code base. We chose MASON over Repast, which is more widely used than MASON and has features that are roughly a superset of MASON's, because MASON's codebase is smaller and easier to adapt.

MASON provides several forms of shared data structures or *fields* for the environment: 2- and 3-dimensional discrete grids, which have both dense and sparse forms, 2- and 3-dimensional continuous spaces, and networks. There are three types of discrete dense grids distinguished by their values, integer, floating point (double) and single arbitrary object. Sparse grids can contain an arbitrary collection of objects at any grid point, and continuous spaces, which are backed by sparse grids in implementation, can associate an arbitrary collection of objects with any location within a rectangular area or volume.

MASON also includes network fields, which define nodes connected by edges and allow the construction of arbitrary graphs. We have not included networks in the present implementation, but we believe it is possible to extend our scheme to them.

MASON has no distinct agent type. Any object implementing the `Steppable` interface, which defines the `step` method, may play the role of an agent. Although agents typically interact through the environment, one agent may call the methods of another directly if it has a reference to it; however, we will place certain restrictions on this form of interaction, as explained below.

MASON already contains low-level support for parallelism. It provides a `Steppable` class named `ParallelSequence` that contains an array of threads and fans a single `step()` call out into `step()` calls by each of the threads on the respective members of an array of user-provided `Steppable` objects. If a program uses the single-program multiple-data (SPMD) paradigm, each `Steppable` in the array is the same except for parameterization. The threads synchronize at each step boundary by means of a semaphore class defined in MASON. MASON leaves all other issues involving synchronization, concurrent access to data and division of work in the hands of the user, however.

4 APPROACH

The approach is intuitive and leaves almost all of the existing package intact. We decompose the area or volume of each field into as nearly as possible equal non-overlapping parts, which we call partitions, and assign a thread to each partition. The thread is responsible for processing the parts of each field that belong to its partition. The intention is to use one thread per available hardware processor. Figure 1 shows a 200 by 200 field of size divided into four partitions.




Figure 1: A 200x200 field with four partitions.

To resolve the issues of thread safety and conflicting data accesses, we use a combination of barrier synchronization and coding convention. We set the convention that a partition's thread may write only to locations (in a grid or space) belonging to the partition. However, it is often necessary for a partition's thread to read values from a neighboring partition. In order to avoid reading a value while the owning partition's thread is writing it, we divide each simulation step into two parts separated by a barrier synchronization. In the first part, a partition's thread may not write to any field (more precisely, it may not write any value that may be read by another partition's thread), but it may read from anywhere in a field, including parts of the field belonging to other partitions. During the second part of the step, a thread may read and write any value in its own partition.

It is also often necessary for an agent to move from a part of a field belonging to one partition to a part belonging to another. We accomplish such movement transparently to the application programmer by means of interprocessor message queues. See section 5 for details.

Our approach allows already-existing applications written in object-oriented agent-based style to operate in parallel with relatively small changes (see example code in the appendices).

5 IMPLEMENTATION

This section describes the changes and additions to the MASON codebase. The implementation shares each MASON field among all the partitions, so there is a single data structure for each field on which all the partitions threads operate.

MASON has a graphical user interface that provides a visual representation of the actions of the agents in the fields, but we made no attempt to preserve it in the multithreaded version, which we see as being most suitable for large batch runs.

Both MASON and Repast have discrete-event scheduling (for a single host), but most agent-based simulation uses simple time-stepping (as indeed the name of the `step` method suggests.) Hence we gear our implementation toward time-stepped simulations.

Many applications allow agents to move in their environment (for instance, the two test programs described in 7). In the course of such movement, an agent may cross from one partition into another. The usual field of play for agents is the sparse field (sparse grid or continuous space), since it contains the machinery to allow multiple agents to share a location. The sparse field method `setObjectLocation` allows the program to move an agent from one position in the field to another. We added logic to the `setObjectLocation` method that sends the agent to a neighboring partition wrapped in an interpartition message if the destination location lies in a different partition than the agent already occupies. Thus transition from one partition to another is handled automatically by the run-time system.

The application may need to make its own adjustments when an agent passes from one partition to another. For instance, if each partition's logic keeps track of the agents in that partition, the source partition

may need to remove the agent from its collection when it leaves, and the destination partition may need to add the agent to its collection when it arrives. For that reason, we define the `Transportable` interface, with methods `transporting` and `transported`. Any agent that may travel from one partition to another must implement this interface. When the `setObjectLocation` method learns that the object is leaving a partition, it calls the `transporting` method on the object, and when the run-time system delivers the message, it calls the object's `transported` method.

To support interpartition messaging, we define an interpartition message queue for each partition (using class `java.util.concurrent.LinkedQueue`) and provide a statically available method for sending interpartition messages. The run-time system delivers interpartition messages immediately after the start of a simulation step.

In addition to the changes and additions to MASON, we require certain coding conventions. We make the convention that a partition's thread may write only to locations (in the grid or space) belonging to the partition. It is frequently necessary for a partition's thread to read values from a neighboring partition (for instance, if the value at a grid point on the border of the partition depends on the values at the surrounding grid points). In order to avoid reading a value while the owning partition's thread is writing it, we set the convention that each simulation step be divided into two parts. During the first part, a partition's thread may not write to a field (or more precisely, it may not write any value that may be read by another partition's thread), but it may read from anywhere in a field, even from parts of the field belonging to other partitions. During the second part, a partition may read and write any value in its own partition. If an agent directly calls a method of another agent, it must do so in accordance with these conventions; that is, the call must not cause a write to another partition or a write to the same partition during the read part of the step.

Note that it is the user's responsibility to abide by the convention that the simulation step be divided into a read part and a write part; there is no enforcement by the run-time system or the compiler. It is necessary to observe this convention to ensure correctness and avoid exceptions (in particular `java.util.concurrent.ModificationException`).

Separating the reading and writing parts of a simulation step also lays the foundation for ensuring that the runs are repeatable in the sense that two runs done with the same random number generator seeds and the same number of partitions must produce exactly the same results.

It is necessary for the partitions to synchronize after the delivery of the interpartition messages, since the delivery may alter a field value later read by a neighboring partition. For that purpose, we include a statically available `synchronize` method that synchronizes all partitions at a barrier (using `java.util.concurrent.CyclicBarrier`.) The synchronization at the beginning of the step ensures that all partitions have sent all interpartition messages for the previous step; the synchronization after the delivery of the interpartition messages ensures that all writing done as the result of their delivery is finished, so that the read part of the step can begin.

The `synchronize` method is available also to the application to enable it to abide by the convention mentioned above that divides each step into a read part and a write part. The application must call `synchronize` at the end of the read part of the step, to guarantee that no cross-partition reads are in progress when the application writes a field value.

Figure 2 represents a single simulation step, initiated by synchronization at a semaphore. The run-time does a barrier synchronization after delivery of the interpartition messages, and the user (by convention) does a barrier synchronization before the write part of the step.

The sparse fields in MASON maintain common data structures for the entire field, so two partition threads that are trying to write to their respective parts of the field could nonetheless suffer a conflict. In order to avoid locks and bottlenecks, we multiplexed these data structures over the partitions, replacing them with an array of structures, giving each partition access to one element of each array. Doing so adds the overhead of an access to a map (`java.util.concurrent.ConcurrentHashMap`) to select the correct element on every access to the sparse field. This was the most invasive of all the changes, although it was mechanically a simple one.




Figure 2: A single timestep.

The dense grids provide a number of global actions (sum, min, max, mean, etc.). The partitioned system uses reductions to implement these actions. Support for the reductions includes a reduction message queue for each partition, implemented using `java.util.concurrent.LinkedBlockingQueue`. The reductions are invisible to the application programmer.

We give each partition its own random number generator. Users may force the partitions to use any seed values they want, but by default the system derives a different seed for each partition from the current time.

In all, the implementation required changes to 13 MASON modules (of 104 total MASON modules, exclusive of the display and demo code.) The changes were not rewrites but additions to already existing functionality. To that we added 16 modules of our own. Many of these are short and simple (an interpartition message class, nine classes associated with reductions, the `Transportable` interface mentioned above). There are four modules to compute the extents of the 2D or 3D grids or spaces owned by each partition.

The new `ThreadState` module uses a `java ThreadLocal` variable to contain information pertaining to each individual partition. The `ThreadState` class contains the partition identifier (an integer in the range 0 to $np - 1$, where $np =$ the number of partitions) and a reference to the partition's random number generator. The programmer may extend the `ThreadState` object to contain all information that is local to a partition for a particular application. In the original MASON system, the `SimState` object contained the entire state of the simulation; in our partitioned version, the simulation state is spread over the `SimState` object and the `ThreadState` object of each partition.

The new `Harness` module makes it easier to write a partitioned application. The main class of a MASON application must extend the `SimState` class. `Harness` extends `SimState`, and the main class of a partitioned application extends `Harness`.

The `Harness` constructor takes the run-time class of the application's `ThreadState` object as an argument and passes it to `ThreadState`'s initialization method to initialize the partition's thread-local state. It also wrangles the application's `ParallelSequence` object, scheduling the members of the array of `Steppable` objects the application passes to it for execution at each time step. It contains the inner class `Initializer`, which the application may optionally override to schedule a single initial step in which the application may do individual initialization of each partition.

In writing the `Steppable` objects, we found it far more efficient to schedule a single call to the `step` method per timestep and then to iterate through the agents or over the field locations than to schedule `step` calls for every possible event.

Appendices A and B show the code for the `step` method of the original MASON Flockers demo program and the `step` method used in our partitioned version, respectively. Note that they contain the

identical statements and differ only in that the altered version segregates the statements into different methods according to whether they execute in the read part or the write part of a timestep.

6 RELATED WORK

Voss, You, Yen, Chen, Lin, Turner, and Lin (2010) describe an approach for the Repast system similar to that followed here for MASON. As in the work described here, it uses partitioned data with multiple threads. It does not explicitly describe cross-partition access or travel, but mentions that synchronized blocks are used for thread safety when necessary. The work presented here avoids the overhead of synchronized blocks by using barrier synchronization to rule out conflicting accesses and transparently uses lock-free message queues for interpartition travel by agents. They note, as we do, that iterating over the agents sustains far less overhead than scheduling each agent action.

Armstrong, Schoeniger, Sale, Goldsby, and Mayo (2011) present an earlier version of the work described here. That version gave each partition separate data structures for each field and maintained shadow copies in each partition of the nearby elements of neighboring partitions, which were the only elements outside its partition that a partition's thread could read. It required the overhead of a translation from a global field location to a location in a local data structure for each field access. The changes and additions required by the earlier version were more extensive and more complicated than those described in this report.

7 PERFORMANCE

The platform used for performance runs was a shared memory machine with 32 Intel Xeon 2.27 GHz processors, cache size 24576 kB and a 64GB memory space. All processors remained enabled during all tests.

The performance tests used two demo programs distributed with the MASON package, *Flockers* and *HeatBugs*. *Flockers* is similar to Craig Reynolds' *Boids* (Reynolds 1987) and simulates the flocking of birds in flight. The agents fly in a 2-dimensional continuous space. The agents in *HeatBugs* exchange heat with their environment and seek an ideal temperature. If an agent (bug) is above (below) that temperature, it moves to the coldest (hottest) location in its immediate (Moore) neighborhood. *HeatBugs* operates in two alternating phases: in the first phase the bugs move, and in the second the environment (represented by a dense floating point grid) diffuses its heat by allowing conduction from a cell's Moore neighborhood to or from the cell depending on the temperature difference and a specified conduction rate. The nature of both programs is such that their computational complexity increases linearly with problem size.

Both test programs are canonical agent-based simulations representing swarm models, swarming being an emergent behavior from the local rules followed by individual agent. These test programs have significance to the types of problems that are being modeled in multi-agent simulations. For example, the self-organizing behavior of flocking in an agent-based simulation has been used to study how leaders emerge in a group as a consequence of members interactions (Quera, Beltran, and Dolado 2010), and a heat bugs model was adapted to model pollution in a metropolitan city (Ahat, Amor, Bui, Lamure, and Courel 2009).

In addition, the patterns of computation followed by the two programs can apply to a much larger class of models. The *Flockers* program iterates through the agents and updates each according to information gained from the other agents in its immediate neighborhood. The *HeatBugs* program iterates through the agents and updates each according to information gained from the environment in its immediate vicinity, then iterates through the locations of the environment and alters the information at each location according to information gained from the surrounding locations. These rather general patterns suggest that the performance results may be indicative of performance over a larger class of applications.

We use three different problem sizes each for *HeatBugs* and *Flockers*. The density of agents per field area is the same for all three sizes for *HeatBugs* and similarly for *Flockers*. The problem size remains the same in each case as the number of processors increases (strong scaling). The test programs never used more than about 16GB of memory.




Figure 3: HeatBugs: 1600000 agents on 12649x12649 field.




Figure 4: HeatBugs: 100000 agents on 3162x3162 field.




Figure 5: HeatBugs: 6250 agents on 791x791 field.




Figure 6: Flockers: 1280000 agents on 12649x12649 field.




Figure 7: Flockers: 80000 agents on 3162x3162 field.




Figure 8: Flockers: 5000 agents on 791x791 field.

The graphs and tables below show the results of the performance runs. They show both speedup and efficiency versus the number of partitions p , where speedup is the time of one simulation step performed with one partition divided by the time of one simulation step performed with p partitions, and efficiency is speedup divided by p . We show the average of five runs for each configuration except the two largest (Flockers with 1280000 agents and HeatBugs with 1600000 agents), for which we used the average of three runs.

Scaling varies with the problem and the problem size, as Figures 3 through 8 show. HeatBugs with 1600000 agents scales extremely well up to about 10 processors (speedup of 8.42 with 10 processors), scales moderately well up to 20 processors (speedup of 14.3 with 20 processors) and does not flatten out entirely until it reaches 28 processors. HeatBugs with 100000 agents scales very well up to three processors (speedup of 2.6 with with three processors) and moderately well up to about 10 processors (speedup of 6.8 with 10 processors.) The scaling of HeatBugs with 6250 agents is not impressive and reaches its maximum speedup of 4.5 at 11 processors. These results show that for good scaling, it is important to have a large enough problem over which to amortize the overhead of parallelization.

The largest Flockers problem size scales about half as well as the largest HeatBugs size. The mid-size Flockers runs scale nearly as well as HeatBugs up to five or six processors and then falls off somewhat more rapidly. The smallest Flockers scales about as well as the smallest HeatBugs up to five processors but tapers off rapidly after that. We suspect that the tendency of the flocks to combine into larger flocks may have caused unbalanced loads on the processors, but we had no opportunity to verify it.

8 CONCLUSIONS

The work presented here shows that it is not overly difficult to add multithreading to an agent-based simulation system with a clearly written codebase, and that multithreading on multicore hardware can significantly decrease the run time of large agent-based simulation runs. Moreover, multithreading can be accomplished in such a way that the code written for the multithreaded system closely resembles code for a conventional single-threaded system.

Our performance results on two different agent-based simulation problems indicate that other large agent-based simulations will benefit from multithreading.

It would be desirable to do more extensive performance studies. One possible enhancement would be to do load balancing. It would not be difficult to keep track of the amount of time each partition's thread spends waiting for synchronization and periodically adjust the size of the partitions to distribute the waiting more evenly over them. As yet we have no data on how load balancing would affect performance; obviously it would have significant effect only on applications in which the load tends to become unbalanced because of agent movement or other factors.

A logical next step would be to extend the scheme to multiple network-connected hosts by distributing each grid or space over all the hosts, each implementing a section of it. Interhost communication latency would be a critical factor for performance. One means of diminishing the effect of latency is to trade frequency of communication for duplication of data and processing. By restricting the depth to which a read access can penetrate into a neighboring partition and the depth to which an agent can travel into a neighboring partition in a single simulation step, and by duplicating the contents and the processing of the border regions of a partition in the neighboring partitions, we could decrease the frequency of interhost communication and amortize the communication latency over several simulation steps (Aaby, Perumalla, and Seal 2010) .

We are exploring ways to make the source code available as open source.

ACKNOWLEDGMENTS

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U. S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

A ORIGINAL MASON CODE

```
public void step(SimState state)
{
    final Flockers flock = (Flockers)state;
    loc = flock.flockers.getObjectLocation(this);

    if (dead) return;

    Bag b = getNeighbors();

    Double2D avoid = avoidance(b, flock.flockers);
    Double2D cohe = cohesion(b, flock.flockers);
    Double2D rand = randomness(flock.random);
    Double2D cons = consistency(b, flock.flockers);
    Double2D mome = momentum();

    double dx = flock.cohesion * cohe.x
                + flock.avoidance * avoid.x
                + flock.consistency* cons.x
                + flock.randomness * rand.x
                + flock.momentum * mome.x;
    double dy = flock.cohesion * cohe.y
                + flock.avoidance * avoid.y
                + flock.consistency* cons.y
```

Goldsby and Pancerella

```
        + flock.randomness * rand.y
        + flock.momentum * mome.y;

// renormalize to the given step size
double dis = Math.sqrt(dx*dx+dy*dy);
if (dis>0)
    {
        dx = dx / dis * flock.jump;
        dy = dy / dis * flock.jump;
    }

lastd = new Double2D(dx,dy);
loc = new Double2D(flock.flockers.stx(loc.x + dx),
                  flock.flockers.sty(loc.y + dy));
flock.flockers.setObjectLocation(this, loc);
}
```

B PARTITIONED CODE

```
/** Called during the read part of a timestep. */
public void step_r(Flockers flockState, Flockers.LocalState local)
    {
        // get space that flocker inhabits, flocker's current location
        final Continuous2D flockSpace = flockState.flockSpace;
        Double2D currLoc = loc;

        if (dead) return;

        Bag nbrs = getNeighbors();

        Double2D avoid = avoidance(nbrs, flockSpace);
        Double2D cohe = cohesion(nbrs, flockSpace);
        Double2D rand = randomness(local.random());
        Double2D cons = consistency(nbrs, flockSpace);
        Double2D mome = momentum();

        double dx = flockState.cohesion * cohe.x
            + flockState.avoidance * avoid.x
            + flockState.consistency* cons.x
            + flockState.randomness * rand.x
            + flockState.momentum * mome.x;
        double dy = flockState.cohesion * cohe.y
            + flockState.avoidance * avoid.y
            + flockState.consistency* cons.y
            + flockState.randomness * rand.y
            + flockState.momentum * mome.y;

        // renormalize to the given step size
        double dis = Math.sqrt(dx*dx+dy*dy);
        if (dis>0)
```

```
{
    dx = dx / dis * flockState.jump;
    dy = dy / dis * flockState.jump;
}

currD = new Double2D(dx, dy);
nextLoc = new Double2D(flockSpace.stx(currLoc.x + dx),
                      flockSpace.sty(currLoc.y + dy));
}

/** Called during the write part of a timestep. */
public void step_w(Flockers flockState, Flockers.LocalState local)
{
    Continuous2D flockSpace = flockState.flockSpace;
    lastD = currD;
    loc = nextLoc;
    flockSpace.setObjectLocation(this, loc);
}
```

REFERENCES

- Aaby, B. G., K. S. Perumalla, and S. K. Seal. 2010, March. "Efficient simulation of agent-based models on multi-GPU and multi-core clusters". In *3rd International ICST Conference on Simulation Tools and Techniques*, edited by L. F. Perrone, G. Stea, J. Liu, A. Uhrmacher, and M. Villén-Altamirano. Institute for Computer Sciences, Social Informatics and Telecommunications Engineering: ICST/ACM.
- Ahat, M., S. B. Amor, M. Bui, M. Lamure, and M.-F. Courel. 2009. "Pollution modeling and simulation with multi-agent and pretopology". In *Complex Sciences: First International Conference, Complex 2009*, edited by J. Zhou, 225–231. Institute for Computer Sciences, Social Informatics and Telecommunications Engineering: Springer.
- Armstrong, R. C., J. S. Schoeniger, K. L. Sale, M. Goldsby, and J. Mayo. 2011, December. "Community-Based Resistance to Intrusion in Information Technology Systems". Sandia report SAND2011-8839, Sandia National Laboratories.
- Epstein, J. M., and R. Axtell. 1996. *Growing artificial societies*. Cambridge, MA, USA: Bradford.
- Jennings, N. R. 2000. "On agent-based software engineering". *Artificial Intelligence* 117 (2): 277–296.
- Luke, S., C. Cioffi-Revilla, L. Panait, K. Sullivan, and G. C. Balan. 2005. "MASON: a multiagent simulation environment". *Simulation* 81 (7): 517–527.
- MPI Forum 2013. "The Message Passing Interface (MPI) Standard". Accessed Mar. 3, 2013. <http://www.mcs.anl.gov/research/projects/mpi/>.
- North, M. J., N. T. Collier, and J. R. Vos. 2006, January. "Experiences creating three implementations of the repast agent modeling toolkit". *ACM Transactions on Modeling and Computer Simulation* 16 (1): 1–25.
- Quera, S., F. S. Beltran, and R. Dolado. 2010. "Agent-based simulation platforms: review and development recommendations". *Journal of Artificial Societies and Social Simulation* 13 (2): 8.
- Railsback, S. F., S. L. Lytinen, and S. K. Jackson. 2006. "Agent-based simulation platforms: review and development recommendations". *SIMULATION* 82:609–623.
- Reynolds, C. 1987. "Flocks, herds and schools: A distributed behavioral model". In *Proceedings of the 14th annual conference on computer graphics and interactive techniques*, edited by M. C. Stone, 25–34. Association for Computing Machinery: ACM.
- Uhrmacher, A. M., and D. Weyns. 2009. *Multi-agent systems: simulation and applications*. Boca Raton, FL, USA: CRC Press.

Voss, A., J.-Y. You, E. Yen, H.-Y. Chen, S. Lin, A. Turner, and J.-P. Lin. 2010, December. "Scalable social simulation: investigating population-scale phenomena using commodity computing". In *Proceedings of the IEEE 6th Conference on e-Science*. Brisbane, Queensland, Australia: Institute of Electrical and Electronics Engineers, Inc.

AUTHOR BIOGRAPHIES

MICHAEL E. GOLDSBY is a contractor to Sandia National Laboratories, California, with special interests in concurrent and distributed systems and simulation. He is a member of ACM. His email address is megolds@sandia.gov and his web page is <http://michaelegoldsby.com>.

CARMEN M. PANCERELLA is a Principal Member of the Technical Staff at Sandia National Laboratories in Livermore, California. She holds a M.S. and Ph.D. In Computer Science from the University of Virginia. Over the years, her modeling and simulation research interests have included parallel discrete event simulation, agent-based modeling, and the use of models in support of emergency preparedness. Her email address is carmen@sandia.gov.