

Optimistic Parallel Discrete Event Simulations of Physical Systems using Reverse Computation

Yarong Tang, Kalyan S. Perumalla, Richard M. Fujimoto
College of Computing, Georgia Institute of Technology
Atlanta, Georgia 30332-0280
{yarongt, kalyan, fujimoto}@cc.gatech.edu

and

Homa Karimabadi, Jonathan Driscoll, Yuri Omelchenko
SciberNet Inc.
Solana Beach, CA 92075
{homak, driscoll, yurio}@scibernet.com

Abstract

Efficient computer simulation of complex physical phenomena has long been challenging due to their multi-physics and multi-scale nature. In contrast to traditional time-stepped execution methods, we describe an approach using optimistic parallel discrete event simulation (PDES) and reverse computation techniques. We show that reverse computation-based optimistic parallel execution can significantly reduce the execution time of a plasma simulation without requiring a significant amount of additional memory compared to conservative execution techniques. We describe an application-level reverse computation technique that is efficient and suitable for complex scientific simulations involving floating point operations.

1. Introduction

Parallel Discrete Event Simulation (PDES) has been an active research area in the high performance computing community for many years. Synchronization techniques for PDES systems are usually classified into two principal categories: conservative approaches that avoid violating the local causality constraint, and optimistic approaches that allow violations to occur, but provide a mechanism to recover. The operation of recovering a previous state in an optimistic parallel simulation is known as a *rollback*, and involves undoing incorrect computations.

A widely used technique for implementing rollback is

state-saving that saves the values of state variables prior to an event computation and restores them by referring to these saved values upon rollback. *Copy state saving* creates an entire copy of a process's state; *incremental state saving* keeps a log of changes to individual state variables; *infrequent state saving* periodically saves the entire copy of the state. A relatively new technique for rollbacks, *reverse computation* [1], realizes rollbacks by performing the inverses of the individual operations executed in the event computation. These techniques have been exploited in small- and large-scale parallel simulations.

However, advances in PDES research to date have yet to be explored in space physical science, where multi-physics and multi-scale physical systems are modeled by partial differential equations and particles. Traditionally, simplified models of such physical systems have been simulated using time-driven or time-stepped approaches [2]. The inherent limitations of the time-stepped approach prevent the simulation of more complex physical systems that are important in plasma physics and other areas of science. Even the latest techniques developed in time-stepped research such as adaptive mesh refinement (AMR) [3] are constrained by excessive computational requirements and the necessity to update all cells within a given mesh patch based on the Courant-Friedrichs-Levy (CFL) condition in that patch. Discrete event simulation was recently used to model spatially discretized physical systems [4] where it was shown that this approach not only alleviates the constraint of the CFL condition but also provides a significant performance advantage over the time-stepped approach.

In our work, we extend this work and apply optimistic

parallel discrete event simulation techniques to this problem. Because memory constraints are often a severe limitation in the size of the computations that can be performed, reverse computation offers greater promise than traditional state saving techniques. We explore the use of reverse execution for plasma simulations to gain new insights for such challenging, complex physical systems. The combination of DES methodology and reverse computing techniques offers the potential to dramatically reduce the amount of time required to perform plasma simulations without incurring a large penalty in additional memory requirements.

The main contributions in this work can be summarized as follows. To our knowledge, this is the first work to apply reverse computation techniques to parallel physical system simulations and to show performance advantages using this approach. In addition, we provide a simple model and guidelines for creating reverse simulation codes at the application level that may help physicists or astrophysicists to develop simulation prototypes without comprehensive knowledge of PDES mechanisms.

The remainder of this paper is organized as follows. The next section discusses related research. Section 3 provides an overview of the physical system we simulate and the reverse computation approach. Section 4 gives an in-depth discussion of the reverse computation implementation and discusses its challenges. Section 5 presents experimental results from a preliminary performance evaluation study. We conclude by reporting current and future work in this area and provide guidelines for reversing parallel physical simulation codes.

2. Related Work

A limited amount of research has examined physical system simulation using parallel discrete event simulation techniques. Perhaps the earliest was the “colliding pucks” application developed for the Time Warp Operating System (TWOS) [5]. This work, modeling a set of pucks traveling over a frictionless plane, was used to benchmark an early implementation of the Time Warp protocol. Lubachevsky discusses the use of conservative simulation protocols to create cellular automata models of Ising spin [6]. Other work describes challenges in using discrete event simulation techniques for a few other physical system problems [7]. A formal approach to both discrete event and continuous simulation modeling based on DEVS (Discrete Event System Specification), was proposed by Zeigler et al. [8] and some numerical solutions have been examined based on the DEVS formalism [9].

Seminal work in optimistic parallel discrete event simulation was completed by Jefferson [10]. State saving has historically been the dominant approach to enabling the rollback of computations. Much work has been done to

reduce the cost of state saving both using hardware [11] and software support. The less costly software approaches reported in the literature include copy state saving, incremental state saving [12, 13] and infrequent state saving [14, 15, 16, 17]. A different approach avoiding the cost of state saving was first described in [1], where reverse execution is used to roll back computations. Their reverse procedures were automatically generated by a compiler. More recent work using reverse execution for parallel network simulations, using manually generated code, was reported in [18]. Our work is different in that it applies reverse execution techniques to the simulation of physical systems that involve complex floating point operations and generates reverse code based on application semantics.

Traditionally, complex physical systems described by partial differential equations and particles are modeled by time-stepped simulations. The authors in [4] recently demonstrated both the feasibility and efficiency of applying discrete event simulation (DES) methodology to model such complex systems. Their study shows that a performance improvement of up to two orders of magnitudes is achievable by switching from a time-stepped approach to an event-driven approach. Our study is based on their work and focuses on *parallel* execution techniques for their DES models. In particular, we examine the feasibility of optimizing the parallel synchronization mechanism by applying reverse computation techniques.

3. Overview

In this section we describe the computational plasma simulation model that is used as a case study. A more detailed discussion on the DES modeling methodology and the physics behind the models can be found in [4]. Since our focus is on the use of reverse execution techniques in optimistic simulation, a high level description of the model suffices for our purpose.

3.1. Computational Model: PIC Simulation

One of the great challenges in space physics is to understand how the solar wind interacts with the earth’s magnetosphere. The work presented in [4] was the first to use a DES approach to simulating such complex physical systems. A feasibility study based on the well-known particle-in-cell (PIC) model [19] was described to provide the foundation for development of such multi-physics and multi-scale simulations. We choose this same model for our feasibility study.

The PIC modeling approach serves as a good starting point in theoretical research in plasma dynamics. This model is conceptually simple, yet captures the characteristics of physical phenomena involved and can be extended to more complex models. However, the PIC

simulation is nontrivial in that it simulates movements of millions of particles and their interactions with the fields. There has been considerable interest in improving its performance in the high performance computing community. Although we apply reverse computation method only to this particular model in our study, the programming model and methodology used here are representative of grid-based physical systems. Therefore, the approach presented here may be applicable to a wider range of physical systems.

We limit our simulations to a one-dimensional electrostatic model using spherical coordinates. Figure 1 illustrates a plasma PIC simulation of charging a spacecraft immersed in neutral plasma by injecting a charged beam from its surface [19]. The spacecraft is initially charge-neutral and immersed in the charge-neutral plasma of the solar wind. A charged beam is periodically injected from the spacecraft surface. The effect of beam injection is two-fold: the surface charge of the spacecraft incurs a change equal in magnitude but with an opposite sign to the injected beam particles; the field in the plasma space incurs a change equal to the beam both in magnitude and sign. The change in plasma field would affect the plasma particles' movement which in turn affects the plasma field.

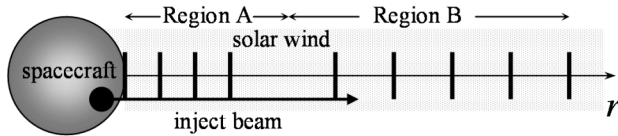


Figure 1. Schematic of the PIC model

3.2. Parallelization

In PDES of a PIC model, the simulation domain is divided into “cells” with each cell mapped onto one logical process (LP). In our spacecraft model, each cell is called a “shell”, and is modeled as a Shell class (we will use the terms cell and shell interchangeably). Two distinct regions are shown in figure 1, based on different grid spacing. The state of each LP includes the cell field variables and the states of all particles within the cell boundaries. The dynamic behavior of the system is driven by particle movements that are modeled by events. Whenever a particle moves across a cell boundary, the electric field in the affected cells is updated by keeping track of the charge that crosses the boundary [4]. Note that field updates are limited to active cells in which events are allowed to be scheduled.

There are three types of events associated with particle movements: *ParticleArrivalEvent*, *ParticleDepartureEvent* and *ParticleInjectEvent*. The corresponding event handlers are outlined in figure 2. When the simulation starts, all

active cells are initialized, including electric fields and particle states (velocities, positions, etc.). In particular, each particle's movement is determined by its *MoveTime* or cell exit time (a time in the future when the particle will exit the hosting cell). Whenever a particle is created or inserted in a cell, its exit time must be calculated and a pair of departure and arrival events scheduled at the exit time. In addition, a particle's exit time needs to be recalculated whenever the hosting cell “wakes up”. A wakeup happens when a cell's field changes beyond a threshold value, resulting in the need for recalculation of exit times of all its particles. The physical activities associated with each component shown in figure 2 will be explained in detail in section 4.

```

Shell::arrival( ParticleArrivalEvent *e ) {
    if ( this cell is active ) {
        update cell state;
        insert particle in cell;
    } else if ( e is a beam particle ) {
        activate cell;
    }
}
Shell::departure( ParticleDepartureEvent *e ) {
    if ( particle bounced from right neighbor ) {
        bounce particle back; // no cell state change
    } else {
        update cell state;
    }
}
Shell::inject( ParticleInjectEvent *e ) {
    update cell state;
    insert beam particles;
}

```

Figure 2. A simplified PIC model

3.3. Reverse Computation Approach

The characteristics of this plasma simulation present three major challenges concerning the synchronization of parallel computations. The rationale for the parallelization approach used here is based on the following considerations.

- **Lookahead.** The simulation is highly dynamic. The amount of parallelism can vary dramatically as the simulation progresses. Dependencies among events are governed by each particle's exit time, but this time can be arbitrarily close into the near future. This low level of “predictability” results in a low, dynamically changing value of lookahead that makes efficient execution using conservative synchronization techniques difficult. This suggests that optimistic

synchronization [10] may be a more natural choice for this simulation.

- **Memory.** Realistic plasma simulations involve a large number of events, on the order of billions. Complex data structures are often needed. This makes traditional approaches to optimistic execution using state saving problematic: the amount of memory required can be prohibitively large. Further, the amount of model-level computation performed for each event tends to be relatively small, on the order of a few microseconds on a contemporary CPU. This suggests that the time overhead for state saving may be significant, and hence could significantly degrade performance, even utilizing techniques such as incremental state saving. While various techniques to reduce the cost of state-saving-based rollbacks have been developed [20, 21], our experiences indicate much memory is still needed to achieve efficient parallel execution [22]. For this reason, reverse computation was selected for the parallelization of this plasma simulation code. However, comparison of the reverse execution approach with other advanced check-pointing approaches remains an area of future investigation.
- **Floating point.** The reverse computation approach proposed in [1] uses an automated approach to creating the reverse execution code for each line of forward execution code. For example, a decrement statement is generated to undo an increment statement in the forward execution code. This approach becomes problematic when floating point arithmetic is used because the computation may not be easily reversed due to effects such as round-off error. Here, we explore a different approach where the program is viewed at a higher level of abstraction, and suitable reverse computation code is developed manually.

These factors motivate the approach that was adopted for optimistic synchronization using manually derived reverse computation code. We believe this can be used to build a foundation for future work in developing scalable parallel simulators for complex physical systems.

4. Parallel Simulation Code

Here we use a one-dimensional model of the spacecraft electrostatic particle code as an illustrative example to discuss some of the challenges in generating the reverse code for this physical system simulation. There are two types of distinct physical entities in this simulation: particles and cells¹. Particles move across cells and each

cell keeps track of the particles residing within its own domain. The communication between adjacent cells occurs via particle movement events that contain information of particle physical states. The complex data structures housing the particles and physical processes being captured require a careful modeling of the system. The object-oriented design used here allows one to encapsulate physical properties via classes..

The code in Figure 2 includes three event handlers, one for each type of event. Much of the complexity of the event computation is encapsulated within the **insert** and **update** operations. An **insert** operation includes an “insert” queue operation (data structures representing particles in the cell are organized in a priority queue) and computation of the exit time of that particle from the cell based on an equation of motion. An **update** operation recomputes the cell’s field value based on the arrival particle’s charge. The change in the field value may trigger an expensive wakeup computation that again scans the particle list and updates each particle’s exit time.

It may be noted that the reverse computation techniques introduced in [1] would generate the reverse code for each instruction without taking into account the semantics of the higher level operations that are being performed. This will clearly lead to inefficiencies for the queue management operations used in this code, and as mentioned earlier, leads to difficulties concerning the reversibility of floating point operations. The model-specific approach taken here involves generating the reverse code for the application by exploiting knowledge of the higher level semantics of the operations being performed. We call this approach *application-level* reverse computation. Here, each cell must manage a large number of particles within its domain. The choice of container class directly affects the efficiency of the simulation. We use the **list** class because of its efficient insertion and deletion operations. However, both insertion and deletion are destructive due to pointer assignments and thus irreversible. However, it is clear from a higher level examination of the operations being performed that insertion and deletion are perfect inverse operations of each other. Indeed, reverse computation in such queue operations are more memory-efficient than state-saving.

Careful readers may notice that the particle departure event handler depicted in figure 2 does not specify any deletion operation that matches the insertion operations in the particle arrival event handler. In fact, deletions are performed aggregately on an as-needed basis. This is done based on performance considerations. Instead of deleting each particle at its exiting time, a near-periodic deletion operation is used to amortize the cost of deletions in

treated as a special cell that does not keep the physical states of particles. We use “cells” thereafter to refer to the regular cells unless otherwise specified.

¹ The spacecraft situated at one end of the spatial coordinate can be

queues. It is observed that wakeup events happen almost periodically with a frequency determined by physical conditions within the simulation; furthermore, each wakeup requires a scan of the cell's particle queue. We find that restricting deletion operations only at cell wakeup times reduces the total overhead of particle deletions without introducing excessive memory usage. Since the aggregate deletions are used to clean up the obsolete states for particles that have already exited, no rollbacks are needed to recover these obsolete states in the event of undoing a wakeup.

Figure 3 shows the reverse code of the simulation. Notice that after decomposing the forward code into components based on simulating physical processes, the reverse code is relatively easy to construct based on the operations shown in figure 3. The next step for generating the complete reverse code is just a matter of reversing each physical process. Examples of reversing some of the more difficult processes are shown in figure 3.

```

Shell::undo_arrival( ParticleArrivalEvent *e ) {
    if ( cell was activated ) {
        undo_activate cell;
    } else if ( cell already active ) {
        delete particle in cell;
        undo_update cell state;
    }
}
Shell::undo_departure( ParticleDepartureEvent *e ) {
    if ( particle was bounced from right neighbor ) {
        undo_bounce particle;
    } else {
        undo_update cell state;
    }
}
Shell::undo_inject( ParticleInjectEvent *e ) {
    delete beam particles;
    undo_update cell state;
}

```

Figure 3. A simplified reverse code of the PIC model

The **insert** operation appears in both particle arrival and injection event handlers. Its effect includes assigning memory for the new particle states in the queue and scheduling arrival/departure event pairs at each particle's future exiting time. Conversely, the **delete** operation in the reverse code should perform corresponding inverses of these processes. Particles in each cell are organized in a FILO (first-in last-out) queue, so the **delete** operation always removes the particle at the head of the queue that is exactly the same particle that was inserted in the forward computation. As for "undoing" event scheduling, it is assumed the underlying simulation engine provides the

application with a primitive for explicitly retracting scheduled events; this is very useful to implement the **delete** operation.

The effect of the **activate** process is to load an inactive cell where previously no particle movements are allowed with particles of uniform distribution. Its basic operation is in fact multiple particle insertions. Its reverse code can utilize the **insert-delete** pair operations described above.

The cell state **update** process is in fact the most complex process and its reverse code is not trivial. Each **update** performs two major computations: computes the cell's new field values and then updates the cell's particle queue. Each cell computes its field locally and keeps track of field values at its left and right boundaries by summing the charges passing through that boundary. During a cell's field update, an addition of charge at its boundary can be simply reversed as a subtraction of the charge at the boundary upon rollback. The update on a particle queue, however, is not as easy. It requires a check of the wakeup condition (i.e., the field change exceeding a threshold) and triggers a wakeup event if necessary. Otherwise it makes no changes. As previously described, a cell wakeup is the single most costly operation in the plasma simulation. In the event of a wakeup, all particles' states are recomputed based on the new cell field values and obsolete particle states are erased. The destructive nature of the re-computation is one of the principal challenges in generating the reverse code for this simulation.

One example of an irreversible operation is the calculation of a particle's exit time *MoveTime*. It is calculated by finding the roots *dt* of the quadratic equations [4]:

$$\frac{1}{2} * Acc * dt^2 + Vel * dt + Pos - CellWidth = 0 \quad (1)$$

$$\frac{1}{2} * Acc * dt^2 + Vel * dt + Pos = 0 \quad (2)$$

where *Acc*, *Vel*, *Pos* are particle acceleration, velocity and position in the cell, respectively. Equations (1) and (2) represent the right and left exit conditions, respectively. *dt* is the time difference between the current simulation time and the particle's last movement time. The smallest real value of *dt* is used for *MoveTime*. An initial inspection of the quadratic equations seems to suggest the impossibility of applying reverse computation to this process. However, if we apply the reverse computation approach at the application level, we find that the movement of the particle is highly reversible. Based on the physical laws of particle motion, the recovery of *dt* does not require the direct inverse of the quadratic equations. Indeed, as illustrated in figure 4, the particle's acceleration, velocity and position states can be simply rolled back by reverse computation and then the critical state *dt* can be reconstructed by carrying the forward computation using the recovered particle states. Note that the parameters *cell_field* and *dt* in the reverse code refer to the rolled-

back cell_field value and the time difference between now and the time when the particle moved right before the wakeup. Finally, we note that random number generation is essential to this parallel simulation. Therefore, an efficient random number generator (RNG) that is reversible and has a long period is required. For this purpose, we use the reversible RNG that is described in [1].

```

Particle::update_position( double dt ) {
    Pos += Vel * dt + 0.5 * Acc * dt * dt;
}
Particle::update_velocity( double dt ) {
    Vel += Acc * dt;
}
Particle::update_acceleration( double cell_field ) {
    Acc = cell_field * particle_charge / particle_mass;
}
Particle::reverse_position( double dt ) {
    Pos = Pos - Vel * dt - 0.5 * Acc * dt * dt;
}
Particle::update_state( double cell_field, double dt ) {
    update_position(dt);
    update_velocity(dt);
    update_acceleration(cell_field);
    dt = update_dt(); // solve eq. (1) and (2)
    MoveTime = now + dt;
}
Particle::reverse_state( double cell_field, double dt ) {
    update_acceleration(cell_field);
    update_velocity(-dt);
    reverse_position(dt);
    dt = update_dt();
    MoveTime = now + dt;
}

```

Figure 4. The reverse code example of the particle states

5. Performance Evaluation

The application-level reverse computation approach is best implemented in a system that decouples implementation details of the simulation engine from the application in order to allow one to focus one's efforts on application semantics. As a result, a simulation engine that can support reverse computation at an application level and provide efficient management of large numbers of events with minimal storage requirement is needed. In addition, the simulation engine should have the flexibility and extensibility to support future refinement of the parallel simulation.

The parallel simulation code using reverse execution described in the previous section was implemented using *μsik*, a general-purpose parallel/distributed simulation

engine based on a micro-kernel architecture [23]. *μsik* provides primitives supporting multiple synchronization approaches, including optimistic and conservative synchronization, as well as means to relax event ordering rules and mixing different approaches to synchronization within a single parallel execution. It therefore provides the capabilities needed for the parallel physical system simulations described here.

In a *μsik* simulation, logical (simulation) processes (LP) are fully autonomous entities that communicate via events. In our simulation model, each cell is implemented as an LP and can choose to run conservatively or optimistically. A conservative implementation of the simulation described here performed very poorly, due to poor lookahead, and is not discussed further. We focus on optimistic execution using our reverse handlers to support rollback.

5.1. Experiment Configuration

To demonstrate the feasibility and efficiency of reverse computation in the electrostatic plasma simulation, we carried out all experiments on a Symmetric Multi-Processor (SMP) machine running Red Hat Linux 7.3 with a customized 2.4.18-10smp kernel. The SMP machine is equipped with eight Pentium III 550MHz Xeon processors that share 4GB of memory.

We use normalized units throughout our simulation, where length, time and velocity are normalized to electron Debye length, electron plasma frequency and electron thermal velocity, respectively. The spacecraft is assumed to have a 500 unit radius and each cell has a width of 0.24 units. The solar wind plasma is initially loaded with uniformly distributed electrons and protons. We choose the initial values of 30 electrons and 30 protons per cell. The injected positron beam has energy of 10 keV with an injection period of 0.004. Upon initialization, there are up to 7000 cells of which the first 70 close to the spacecraft are "active"; as the simulation progresses up to time 60, the beam travels further away from spacecraft surface and thus more cells are activated.

5.2. Parallel Performance

Figure 5 shows a snapshot visualization of phase space structures for the solar wind electrons, protons and beam particles from a time-stepped simulation and the optimistic PDES simulation. Both simulations are run up to 60 time units, and with the same simulation parameters except using different random number generators (RNG). The PDES used a specialized reversible RNG in contrast to the generic single-stream RNG used in the time-stepped simulation. Despite this difference, the two phase space structures at the end of the simulations are rather close in form. The result from the PDES execution with reverse computation is verified to accurately capture the main features of movement for all three species of particles at

the end of the simulation. In particular, it can be seen that the beam front in both simulations has propagated to the same distance and beam particles display a similar shape in phase space. It is also evident that the electron phase space has a finer resolution in the PDES case for up to 4000 cells. This is the result of its fine time-scale based on individual particles. A minor difference to note is that, in the PDES case, the phase space does not extend all the way to the right wall, whereas in the time-stepped model, it does. This is because we model an expanding box in PDES but not the time-stepped model. Overall, the results helped serve as validation of our optimistic simulation model against the original sequential simulation model.

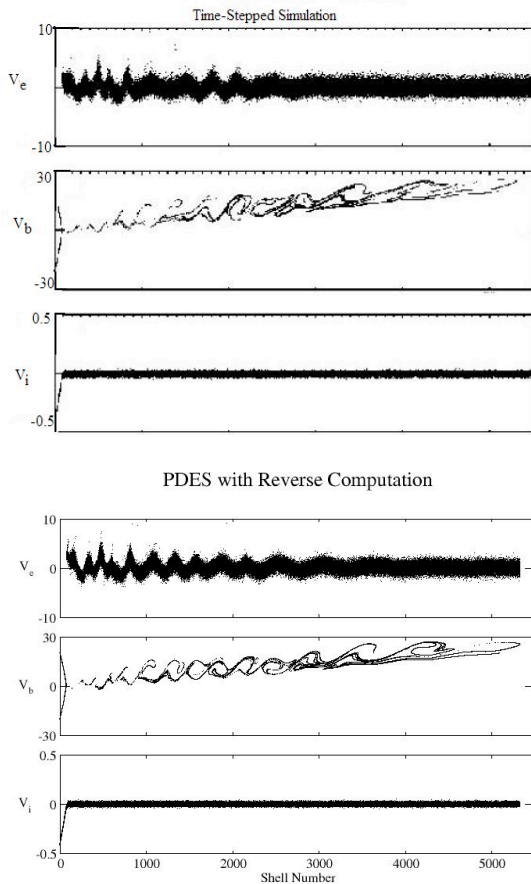


Figure 5. Validation by phase space comparison of time-stepped simulation and PDES with reverse computation.

The speedup of DES over TDS has been discussed in great detail in [4]. In our work, we focus on further improving the DES performance by realizing parallelization in the simulation and utilizing optimistic synchronization. All the parallel experiments discussed in the following section were run with the same physical

parameters and resulted in the same number of committed events as the sequential runs. Figure 6 shows the parallel speedup in terms of execution time for up to 8 processors. The sequential data is measured by running the parallel code on a single processor. It should be noted that the single processor execution incurs neither rollbacks nor state saving overhead. Because μsik was designed for both efficient sequential and parallel execution, we believe these measurements reflect the performance one could expect to see using a reasonably efficient sequential simulation engine.

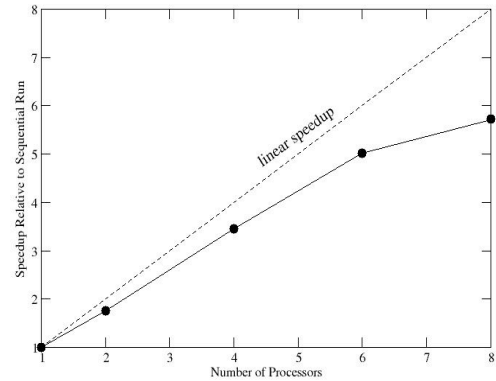


Figure 6. PDES vs. sequential DES

We observe that the optimistic parallel execution achieves a nearly linear speedup up to 4 processors, but the performance improvement is somewhat less in going from 4 to 8 processors. This phenomenon is largely due to the fact that there is relatively little computation per particle event. As the computation is distributed over more and more processors, the amount of computation between event communications decreases, resulting in reduced speedup. We expect that this problem will not persist if a larger, more complex physical model such as a three dimensional plasma code were used. An initial test with an increased simulation time of 2 units did show better speedup performance due to the fact that the longer the simulation runs, the more cells are activated, resulting in more balanced computation for each processor.

A second factor that results in less than optimal performance concerns the distribution of the workload. Figure 7 shows the amount of computation assigned to each processor in each of the runs. Here, the load is distributed by first dividing the physical area encompassed by the simulation into two regions (as illustrated in figure 1) with the initially active cells closer to the spacecraft in the “heavy activity” region, and other cells forming the “less active” region. Cells in the active regions are evenly grouped and distributed among the available processors, while other cells are grouped into sub-regions or “blocks”

and distributed among processors in a round-robin fashion. All simulations shown in figure 7 have a fixed “block” size. We observe that during the lifetime of each simulation, the processor load shows significant variation as more and more cells become active. Upon simulation termination, the simulation with the largest number of processors tends to be the least balanced. The imbalance is inherent of such simulations due to their highly dynamic nature and the static load-balancing scheme. Further investigation of characteristics of electrostatic plasma simulations is needed to aid in the development of a more efficient load-balancing algorithm for this application that can lead to better parallel performance for large numbers of processors.

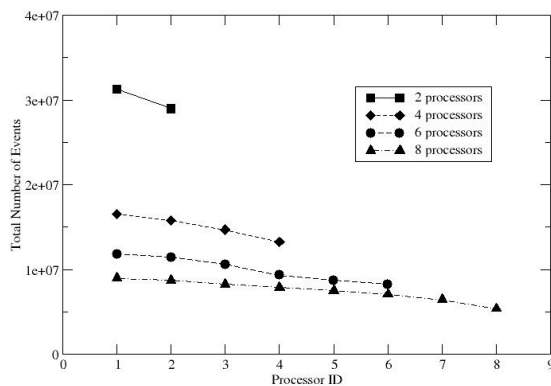


Figure 7. Event rate distribution

5.3. Efficiency

Intuitively, grid-based physical systems such as the electrostatic plasma simulation studied here have the desirable features of locally solved field values and queuing/dequeuing operations that are time-reversible, but the evolution of the system itself (beam injections, cell wakeups in our case) is not time-reversible. However, with the application-level reverse computation illustrated here, we have shown that numerical operations in the electrostatic plasma simulation chosen for this study are truly reversible, despite round-off errors and irreversible evolution processes. The most important discovery from our study is that application-level reverse computation may be quite efficient for these scientific simulations.

The efficiency mainly comes from two contributing factors: the smaller amount of additional memory required for optimistic execution, particularly, queue operations where no additional state is required to perform rollbacks; note the fact that the simulation is not constrained by arbitrarily small look-ahead values. However, there are still important practical issues related to reverse computation.

Ideally, one would like to apply reverse computation to all reversible operations. But reverse computation also comes at a cost: if the number of destructive operations is sufficiently large and no efficient application-level reverse computation can be found, employing reverse computation can result in worse performance than state-saving. One such case as pointed out in [1] is when a rollback spans several processed events. Merely switching pointers to restore a state based on the earliest rolled back event incurs a small cost in copy state-saving; while reverse computation must roll back one event at a time and thus excessive rollbacks can cause performance to degrade considerably. The effect of this is particularly severe in our simulation when a rollback spans multiple wakeup events.

Our solution to reducing the rollbacks of costly wakeup events is to limit the “optimism” of the parallel execution. The idea of controlling optimism is not new. One such approach is first described in the Moving Time Window (MTW) protocol [24], where LPs are not allowed to advance beyond a time window above the GVT. Although other approaches such as probabilistic rollbacks [25], local rollbacks [26], Breathing Time Buckets [27] and Wolf Call mechanism [28], have been proposed to counter the effect of over-optimism, we simply choose to a time window-based approach because of its simplicity and low overhead. μ sik supplies simulation applications with a convenient facility for our purpose. A “run-ahead” parameter, which is in fact the moving window size, can be set by the model upon simulation initialization. This limits how far in simulation time each LP can run ahead of other LPs during optimistic execution. By tuning the run-ahead parameter based on cell wakeup frequency, we are able to reduce or eliminate consecutive rollbacks of wakeup events. The significant performance gain in our experiment indicates the extra operation associated with the window maintenance is a small cost to pay. Further improvement in parallel performance is possible by fine-tuning the run-ahead value. This is left for future investigation.

In addition to the basic reverse computation techniques discussed here, advanced reverse techniques can be applied to the plasma simulation. For example, compiler-supported reverse computation can be used to further optimize the parallel performance at run-time. This approach is beyond the scope of our discussion and will be studied in the future.

6. Conclusions

In this work, we have applied reverse execution to perform parallel discrete event simulations of a physical system. We demonstrated that application-level reverse computation can be used to manually generate efficient reverse code. These results suggest that reverse

computation merits further investigation as an approach for parallel/distributed simulation of physical systems modeled using a discrete event simulation paradigm.

As previously mentioned, the PIC simulation considered in this paper is only a simplified example of reverse execution in simulating physical systems. The examples given in section 4 are representative and certainly do not encompass the diversity and complexity of all physical system simulations. However, the underlying reverse techniques can be used in other grid-based models without extensive modifications. Here we provide some guidelines for the development of parallel physical discrete event simulations using reverse computation. Since our exploration of reverse computation is an on-going research effort, the guidelines provided here should be used as references rather than strict rules for applying reverse computation in scientific simulations.

- Reverse computation is well-suited for fine-grained applications such as 1D electrostatic grid-based plasma models. It is especially useful where efficient queue management is needed. But other optimization techniques should also be considered in order to fully optimize parallel performance.
- Good knowledge of the application semantics, especially the underlying physics, can be beneficial in producing reverse code for physical systems. Model-specific optimization can be quite efficient but requires knowledge of application-level operations. The simple example of reversing the quadratic equation would not have been efficient, if at all possible, without knowledge of the physics involved (particle's motion in this case).
- The modeling process largely determines how successfully reverse computation will improve parallel performance. Initial analysis in [1] shows that complex use of jump instructions such as `goto`, `break` and `continue` are difficult to optimize in terms of memory usage.
- In modeling physical systems, one should attempt to avoid monolithic code for event handlers and use functions calls that are associated with each physical process. If an event handler only consists of a long sequence of simple instructions, it is difficult to extract application semantics and therefore reverse computation will degenerate to instruction-by-instruction reverse execution. Using many small function calls that reflect physical processes helps to develop reverse codes based on physical properties of the system. Another advantage is easier debugging and testing for the reverse code.

The work presented here is only an initial step based on a simplified physical system. Yet, the results show promise. Our goal is to build a scalable parallel simulator

for complex physical systems by exploitation of more advanced reverse computation techniques.

Acknowledgements

This work was supported in part by the National Science Foundation grants ATM-0326431 and 0325046. Additionally the authors would like to thank the anonymous reviewers for their helpful comments on improving the paper.

References

- [1] Carothers, C. D., K. Perumalla and R. M. Fujimoto. Efficient Optimistic Parallel Simulation Using Reverse Computation. *ACM Transactions on Modeling and Computer Simulation*, 9(3): 224-253, 1999.
- [2] Birdsall, C. K. and A. B. Langdon. Plasma Physics via Computer Simulation. *McGraw-Hill Book Company*, 1985.
- [3] Dawson, C. and R. Kirby. High Resolution Schemes for Conservation Laws with Locally Varying Time Steps. *SIAM Journal of Scientific Computing*, 22(6): p. 2256, 2001.
- [4] Karimabadi, H., J. Driscoll, Y. A. Omelchenko, and N. Omidi. A New Asynchronous Methodology for Modeling of Physical Systems: Breaking the Curse of Courant Condition. *J. Computational Physics*, 205(2): 755 - 775, 2005.
- [5] Hontalas, P., B. Beckman, M. DiLorento, L. Blume, P. Reiher, K. Sturdevant, L. V. Warren, J. Wedel, F. Wieland and D. Jefferson. Performance of the Colliding Pucks Simulation on the Time Warp Operating System. *Distributed Simulation, Society for Computer Simulation International*, 1989.
- [6] Lubachevsky, B. D. Efficient Distributed Event-Driven Simulations of Multiple-Loop Networks. *Communications of the ACM*, 32(1): 111-123, 1989.
- [7] Lubachevsky, B. D. Several Unsolved Problems in Large-Scale Discrete Event Simulations, in *Proceedings of the 7th Workshop on Parallel and Distributed Simulation*, 1993. p. 60-67.
- [8] Zeigler, B. P., H. Praehofer and T. G. Kim. Theory of Modeling and Simulation: Second Edition. *Academic Press*, 2000.
- [9] Nutaro, J. J. Parallel Discrete Event Simulation with Application to Continuous Systems. *Ph. D. Dissertation*. 2003. Electrical and Computer Engineering Dept, University of Arizona.
- [10] Jefferson, D. Virtual Time. *ACM Transactions on Programming Languages and Systems*, 7(3): 404-425, 1985.
- [11] Fujimoto R. M., J. J. Tsai and G. Gopalakrishnan. The Roll Back Chip: Hardware Support For Distributed Simulation Using Time Warp, in *Proceedings of SCS Distributed Simulation Conference*, 19(3): 81-86, 1988.
- [12] Bauer H. and C. Sporrer. Reducing Rollback Overhead in Time Warp Based Distributed Simulation with Optimized Incremental State Saving, in *Proceedings of the 26th Annual Simulation Symposium*, 1993.
- [13] Steinman J. Incremental State Saving in SPEEDES Using C++, in *Proceedings of the 1993 Winter Simulation*

- Conference*, 1993.
- [14] Lin. Y. B. and E.D.Lazowska. The Optimal Checkpoint Interval in Time Warp Parallel Simulation, *Technical Report 89-09-04*, Dept. of Computer Science and Engineering, University of Washington, 1989.
 - [15] Bellenot S. State Skipping Performance with the Time Warp Operating System, in *Proceedings of the 6th Workshop on Parallel and Distributed Simulation*, 1992. p. 53-61.
 - [16] Franks S., F. Gomes, B. Unger, and J. Cleary. State Saving for Interactive Optimistic Simulation, in *Proceedings of the 11th Workshop on Parallel and Distributed Simulation*, 1997. p. 72-79.
 - [17] Fleischmann J. and P. Wilsey. Comparative Analysis of Periodic State Saving Techniques in Time Warp Simulators, in *Proceedings of the 9th Workshop on Parallel and Distributed Simulation*, 1995. p. 50-58.
 - [18] Yuan, G., C. D. Carothers and S. Kalyanaraman. Large-Scale TCP Models Using Optimistic Parallel Simulation, in *Proceedings of the 17th Workshop on Parallel and Distributed Simulation*, 2003. p. 153.
 - [19] Pritchett, P. L. and R. M. Winglee. The Plasma Environment during particle beam injection into space plasmas: 1. Electron-beams, *J. of Geophysical Res. – Space Physics*, 1987. **92** (A7): 7673-7688.
 - [20] Jefferson, D.R., Virtual Time II: Storage Management in distributed Simulation, in *Proceedings of the Ninth Annual ACM Symposium on Principles of Distributed Computing*, 1990. p. 75-89.
 - [21] Preiss, B.R. and W.M. Loucks. Memory Management Techniques for Time Warp on a Distributed Memory Machine, in *Proceedings of the 9th Workshop on Parallel and Distributed Simulation*, 1995. p. 30-39.
 - [22] Das, S.R. and R.M. Fujimoto. An Empirical Evaluation of Performance-Memory Tradeoffs in Time Warp. *IEEE Transactions on Parallel and Distributed Systems*, 1997. **8**(2): p. 210-224.
 - [23] Perumalla, K. μ sik -- A Micro-kernel for Parallel/Distributed Simulation Systems, in *Proceedings of the 19th Workshop on Principles of Advanced and Distributed Simulation*, 2005.
 - [24] Sokol, L. M. and B. K. Stucky. MTW: Experimental results for a constrained optimistic scheduling paradigm, in *Proceedings of the SCS Multiconference on Distributed Simulations*, 1990. **22**: 169-173.
 - [25] Madiseti V., D. Hardaker and R. Fujimoto. The MIMDIX Operating System for Parallel Simulation. *Journal on Parallel and Distributed Computing*, 1993. **18**(4) : 473-483.
 - [26] Dickens, P.M. and P.F. Reynolds. SRADS with Local Rollback, in *Proceedings of the SCS Multiconference on Distributed Simulation*, 1990. **22**: 161-164.
 - [27] Steinman, J. S. Breathing Time Warp, in *Proceedings of the 7th Workshop on Parallel and Distributed Simulation*, 1993. p. 109-118.
 - [28] Madiseti V., J. Walrand and D. Messerschmitt, WOLF: a rollback algorithm for optimistic distributed simulation systems, in *Proceedings of the 20th Conference on Winter Simulation*, 1988. p. 296-305.