

Behavioral Simulations in MapReduce

Guozhang Wang, Marcos Vaz Salles, Benjamin Sowell, Xun Wang, Tuan Cao,
Alan Demers, Johannes Gehrke, Walker White

Cornell University
Ithaca, NY 14853, USA

{guoz, vmarcos, sowell, tuancao, ademers, johannes, wmwhite}@cs.cornell.edu
{xw239}@cornell.edu

ABSTRACT

In many scientific domains, researchers are turning to large-scale behavioral simulations to better understand real-world phenomena. While there has been a great deal of work on simulation tools from the high-performance computing community, behavioral simulations remain challenging to program and automatically scale in parallel environments. In this paper we present BRACE (Big Red Agent-based Computation Engine), which extends the MapReduce framework to process these simulations efficiently across a cluster. We can leverage spatial locality to treat behavioral simulations as iterated spatial joins and greatly reduce the communication between nodes. In our experiments we achieve nearly linear scale-up on several realistic simulations.

Though processing behavioral simulations in parallel as iterated spatial joins can be very efficient, it can be much simpler for the domain scientists to program the behavior of a single agent. Furthermore, many simulations include a considerable amount of complex computation and message passing between agents, which makes it important to optimize the performance of a single node and the communication across nodes. To address both of these challenges, BRACE includes a high-level language called BRASIL (the Big Red Agent SIMulation Language). BRASIL has object-oriented features for programming simulations, but can be compiled to a dataflow representation for automatic parallelization and optimization. We show that by using various optimization techniques, we can achieve both scalability and single-node performance similar to that of a hand-coded simulation.

1. INTRODUCTION

Behavioral simulations, also called agent-based simulations, are instrumental in tackling the ecological and infrastructure challenges of our society. These simulations allow scientists to understand large complex systems such as transportation networks, insect swarms, or fish schools by modeling the behavior of millions of individual agents inside the system [6, 11, 13].

For example, transportation simulations are being used to address traffic congestion by evaluating proposed traffic management

systems before implementing them [11]. This is a very important problem as traffic congestion cost \$87.2 billion and required 2.8 billion gallons of extra fuel and 4.2 billion hours of extra time in the U.S. in 2007 alone [44]. Scientists also use behavioral simulations to model collective animal motion, such as that of locust swarms or fish schools [6, 13]. Understanding these phenomena is crucial, as they directly affect human food security [24].

Despite their huge importance, it remains difficult to develop large-scale behavioral simulations. Current systems either offer high-level programming abstractions, but are not scalable [22, 33, 36], or achieve scalability by hand-coding particular simulation models using low-level parallel frameworks, such as MPI [50].

This paper proposes to close this gap by bringing database-style programmability and scalability to agent-based simulations. Our core insight is that behavioral simulations may be regarded as computations driven by large *iterated spatial joins*. We introduce a new simulation engine, called BRACE (Big Red Agent-based Computation Engine), that extends the popular MapReduce dataflow programming model to these iterated computations. BRACE embodies a high-level programming language called BRASIL, which is compiled into an optimized shared-nothing, in-memory MapReduce runtime. The design of BRACE is motivated by the requirements of behavioral simulations, explained below.

1.1 Requirements for Simulation Platforms

(1) Support for Complex Agent Interaction. Behavioral simulations include frequent local interactions between individual entities in the simulation system, or agents. In particular, agents may affect the behavior decisions of other agents, and multiple agents may issue concurrent writes to the same agent. A simulation framework should support a high degree of agent interaction without excessive synchronization or rollbacks. This precludes discrete event simulation engines or other approaches based on task parallelism and asynchronous message exchange.

(2) Automatic Scalability. Scientists need to scale their simulations to millions or billions of agents to accurately model phenomena such as city-wide traffic or swarms of insects [6, 13, 21]. These scales make it essential to use data parallelism to distribute agents across many nodes. This is complicated by the interaction between agents, which may require communication between several nodes. Rather than requiring scientists to write complex and error-prone parallel code, the platform should automatically distribute agents to achieve scalability.

(3) High Performance. Behavioral simulations are often extremely complex, involving sophisticated numerical computations and elaborate decision procedures. Much existing work on behavioral simulations is from the high-performance computing commu-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Articles from this volume were presented at The 36th International Conference on Very Large Data Bases, September 13-17, 2010, Singapore.

Proceedings of the VLDB Endowment, Vol. 3, No. 1
Copyright 2010 VLDB Endowment 2150-8097/10/09... \$ 10.00.

nity, and they frequently resort to hand-coding specific simulations in a low-level language to achieve acceptable performance [21, 38]. A general purpose framework must be competitive with these hand-coded applications in order to gain acceptance.

(4) Commodity Hardware. Historically, many scientists have used large shared-memory supercomputer systems for their simulations. Such machines are tremendously expensive, and cannot scale beyond their original capacity. We believe that the next generation of simulation platforms will target shared-nothing systems and will be deployed on local clusters or in the cloud on services such as Amazon’s EC2 [2].

(5) Simple Programming Model. Domain scientists have shown their willingness to try simulation platforms that provide simple, high-level programming abstractions, even at some cost in performance and scalability [22, 33, 36]. Nevertheless, a behavioral simulation framework should provide an expressive and high-level programming model without sacrificing performance.

1.2 Contributions

The main contributions we make with this work are summarized in the following:

- We show how behavioral simulations can be abstracted in the state-effect pattern, a programming pattern we developed for scaling the number of non-player characters in computer games [51, 52]. This pattern allows for a high degree of concurrency among strongly interacting agents (Section 2).
- We show how MapReduce can be used to scale behavioral simulations expressed in the state-effect pattern across clusters. We abstract these simulations as iterated spatial joins and introduce a new main memory MapReduce runtime that incorporates optimizations motivated by the spatial properties of simulations (Section 3).
- We present a new scripting language for simulations that compiles into our MapReduce framework and allows for algebraic optimizations in mappers and reducers. This language hides all the complexities of modeling computations in MapReduce and parallel programming from domain scientists (Section 4).
- We perform an experimental evaluation with two real-world behavioral simulations that shows our system has nearly linear scale-up and single-node performance that is comparable to a hand-coded simulation. (Section 5).

We further discuss related work to our approach in Section 6.

2. BEHAVIORAL SIMULATIONS IN THE STATE-EFFECT PATTERN

Behavioral simulations model large numbers of individual agents that interact in a complex environment. Unlike scientific simulations that can be modeled as systems of equations, agents in a behavioral simulation can execute complex programs that include non-trivial control flow. In order to illustrate the functioning of behavioral simulations, we use a traffic simulation [54] and a fish school simulation [13] as running examples. Details on these simulations can be found in Appendix C.

Concurrency in Behavioral Simulations. Logically, all agents in a behavioral simulation execute actions concurrently. These actions usually result from interactions with other agents. In a fish school simulation, for example, a fish continuously updates its direction of

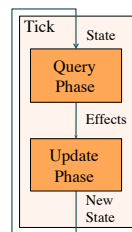
movement based on the orientations of other fish within its visibility range. This implies that it is necessary to ensure consistency on multiple concurrent reads and writes among several fish.

Traditional discrete-event simulation platforms handle concurrency by either preempting or avoiding conflicts [9, 39, 15, 35, 57]. These systems implement variants of pessimistic or optimistic concurrency control schemes. The frequency of local interactions among agents in a behavioral simulation, however, introduces many conflicts. This leads to poor scalability due to either excessive synchronization or frequent rollbacks, depending on the technique employed. Thus, these previous solutions, while programmable, do not allow us to scale behavioral simulations to large scenarios.

The State-Effect Pattern. In order to deal with concurrency in behavioral simulations, we observe that most behavioral simulations can be modeled under a similar structure, which we introduce below. Behavioral simulations use a time-stepped model in which time is discretized into “ticks” that represent the smallest time period of interest. Events that occur during the same tick are treated as *simultaneous* and can be reordered or parallelized. This means that an agent’s decisions cannot be based on previous actions made during the same tick. An agent can only read the state of the world as of the previous tick. For example, in the traffic simulation, each car inspects the positions and velocities of other cars as of the beginning of the tick in order to make lane changing decisions.

In previous work on scaling computer games, we proposed a model for this kind of time-stepped behavior called the *state-effect pattern* [51, 52]. The basic idea is to separate read and write operations in order to limit the synchronization necessary between agents. In the state-effect pattern, the attributes of an agent are separated into *states* and *effects*, where states are public attributes that are updated only at tick boundaries, and effects are used for intermediate computations as agents interact. Therefore state attributes remain fixed during a tick, and only need to be synchronized at the end of each tick. Furthermore, each effect attribute has an associated decomposable and order-independent *combinator function* for combining multiple assignments, i.e., updates to this attribute, during a tick. This allows us to compute effects in parallel and combine the results without worrying about concurrent writes. For example, in the fish school simulation every fish agent has an orientation vector as an effect attribute. This effect attribute uses vector addition as its combinator function, which is invoked on assignments of the orientations of nearby fish. Since vector addition is commutative, we can process these assignments in any order.

In the state-effect pattern, each tick is divided into two phases: the *query phase* and the *update phase*, as shown in the figure on the right. In the query phase, each agent queries the state of the world and assigns effect values, which are combined using the appropriate combinator function. To ensure the property that the actions during a tick are conceptually simultaneous, state variables are read-only during the query phase and effect variables are write-only.



In the update phase, each agent can read its state attributes and the effect attributes computed from the query phase; it uses these values to compute the new state attributes for the next tick. In the fish school simulation, the orientation effects computed during the query phase are read during the update phase to compute a fish’s new velocity vector, represented as a state attribute. In order to ensure that updates do not conflict, each agent can only read and write its own attributes during the update phase. Hence, the only way that agents can communicate is through effect assignments in

effects	map ₁ ^t	reduce ₁ ^t	map ₂ ^t	reduce ₂ ^t	map ₁ ^{t+1}
local	update ^{t-1} distribute ^t	query ^t	—	—	update ^t distribute ^{t+1}
non-local	update ^{t-1} distribute ^t	non-local effect ^t	—	effect aggregation ^t	update ^t distribute ^{t+1}

Table 1: The state-effect pattern in MapReduce

the query phase. We classify effect assignments into *local* and *non-local* assignments. In a local assignment, an agent updates one of its own effect attributes; in a non-local assignment, an agent writes to an effect attribute of a different agent.

The Neighborhood Property. The state-effect pattern addresses concurrency in a behavioral simulation by limiting the synchronization necessary during a tick. However, the synchronization at tick boundaries may still be very expensive, as it is possible that every agent needs to query every other agent in the simulated world to compute its effects. We observe that this rarely occurs in practice. Most behavioral simulations are spatial, and simulated agents can only interact with other agents that are close according to a distance metric [30]. For example, a fish can only observe other fish within a limited distance ρ [13]. We will take advantage of this property of a large class of behavioral simulations to optimize communication within and across ticks.

3. MAPREDUCE FOR SIMULATIONS

In this section, we abstract behavioral simulations modeled according to the state-effect pattern and neighborhood property as computations driven by iterated spatial joins (Section 3.1). We proceed by showing how these joins can be expressed in the MapReduce framework (Section 3.2). We then propose a system called BRACE to process these joins efficiently (Section 3.3).

3.1 Simulations as Iterated Spatial Joins

In Section 2, we observed that behavioral simulations can be modeled by two important properties: the state-effect pattern and the neighborhood property. The state-effect pattern essentially characterizes behavioral simulations as iterated computations with two phases: a query phase in which agents inspect their environment to compute effects, and an update phase in which agents update their own state.

The neighborhood property introduces two important restrictions on these phases, called visibility and reachability. We say that the *visible region* of an agent is the region of space containing agents that this agent can read from or assign effects to. An agent needs access to all the agents in its visible region to compute its query phase. Thus a simulation in which agents have small visible regions requires less communication than one with very large or unbounded visible regions. Similarly, we can define an agent’s *reachable region* as the region that the agent can move to after the update phase. This is essentially a measure of how much the spatial distribution of agents can change between ticks. When agents have small reachable regions, a spatial partitioning of the agents is likely to remain balanced for several ticks. Frequently an agent’s reachable region will be a subset of its visible region (an agent cannot move farther than it can see), but this is not required.

We observe that since agents only query other agents within their visible regions, processing a tick is similar to a *spatial self-join* [34]. We join each agent with the set of agents in its visible region and perform the query phase using only these agents. During the update phase, agents move to new positions within their reachable regions and we perform a new iteration of the join during the next tick. We will use this observation next to parallelize behavioral simulations efficiently in the MapReduce framework. As

we will see in Section 4, the observation also enables us to apply efficient database indexing techniques, e.g., [18, 37].

3.2 Iterated Spatial Joins in MapReduce

In this section, we show how to model spatial joins in MapReduce. A formal version of this model appears in Appendix A. MapReduce has been criticized for being inefficient at processing joins [53] and also inadequate for iterative computations without modification [19]. However, the spatial properties of simulations allow us to process them effectively without excessive communication. Our basic strategy uses the technique of Zhang et al. to compute a spatial join in MapReduce [56]. Each map task is responsible for spatially partitioning agents into a number of disjoint regions, and the reduce tasks join the agents using their visible regions.

The set of agents assigned to a particular partition is called that partition’s *owned set*. Note that we cannot process partitions completely independently because each agent needs access to its entire visible region, which may intersect several partitions. To address this, we define the *visible region* of a partition as the union of the visible regions of all points in the partition. The map task *replicates* each agent a to every partition that contains a in its visible region.

Table 1 shows how the phases of the state-effect pattern are assigned to map and reduce tasks. For simulations with only local effect assignments, a tick t begins when the first map task, map_1^t , assigns each agent to a partition (distribute^t). Each reducer is assigned a partition and receives every agent in its owned set as well as replicas of agents within its visible region. These are exactly the agents necessary to process the query phase of the owned set (query^t). The reducer, reduce_1^t , outputs a copy of each agent it owns after executing the query phase and updating the agent’s effects. The tick ends when the next map task, map_1^{t+1} , executes the update phase (update^t).

This two-step approach works for simulations that have only local effects, but does not handle non-local effect assignments. Recall that in a non-local effect assignment some agent a updates an effect in some other agent b within a ’s visible region. For example, if the fish simulation included predators, we could model a shark attack as a non-local effect assignment by a shark to a nearby fish. Non-local effects require communication during the query phase. We implement this communication using two MapReduce passes, as illustrated in Table 1. The first map task, map_1^t , is the same as before. The first reduce task, reduce_1^t , performs non-local effect assignments to its replicas (non-local effect^t). These partially aggregated effect values are then distributed to the partitions that own them, where they are combined by the second reduce task, reduce_2^t . This computes the final value for each aggregate (effect aggregation^t). As before, the update phase is processed in the next map task, map_1^{t+1} . Note that the second map task, map_2^t , is only necessary for distribution, but does not perform any computation and can be eliminated in an implementation. We call this model map-reduce-reduce.

Our map-reduce-reduce model relies heavily on the neighborhood property. The number of replicas that each map task must create depends on the size of the agent’s visible regions, and the frequency that partitions change their owned set depends on the size of their reachable regions. So, as shown in the next section, we rely on the fact that long-distance interactions are uncommon in behavioral simulations to optimize the communication pattern when processing the simulation.

3.3 The BRACE MapReduce Runtime

In this section we describe a MapReduce implementation that takes advantage of the state-effect pattern and the neighborhood

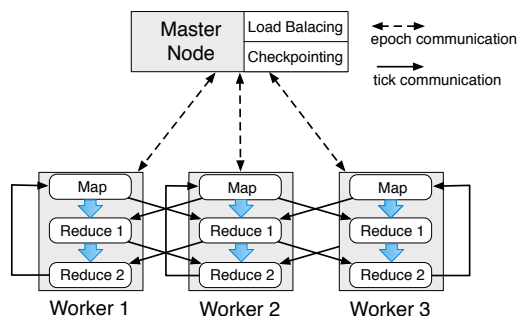


Figure 1: BRACE Architecture Overview

property. We introduce BRACE, the Big Red Agent Computation Engine, our platform for scalable behavioral simulations. BRACE includes a MapReduce runtime especially optimized for the iterated spatial joins discussed in Section 3.1. We have developed a new system rather than using an existing MapReduce implementation such as Hadoop [25] because behavioral simulations have considerably different characteristics than traditional MapReduce applications such as search log analysis. The goal of BRACE is to process a very large number of ticks efficiently, and to avoid I/O or communication overhead while providing features such as fault tolerance. As we show below, our design allows us to apply a number of techniques developed in the HPC community [14, 20, 26, 46] and bridge them to a map-reduce model.

Shared-Nothing, Main-Memory Architecture. In behavioral simulations, we expect data volumes to be modest, so BRACE executes map and reduce tasks entirely in main memory. For example, a simulation with one million agents whose state and effect fields occupy 1 KB on average requires roughly 1 GB of main memory. Even larger simulations with orders of magnitude more agents will still fit in the aggregate main memory of a cluster. Since the query phase is computationally expensive, partition sizes are limited by CPU cycles rather than main memory size.

Figure 1 shows the architecture of BRACE. As in typical MapReduce implementations, a master node is responsible for cluster coordination. However, BRACE’s master node only interacts with worker nodes every *epoch*, which corresponds to a fixed number of ticks. This amortizes overhead related to fault tolerance and load balancing. In addition, we allocate tasks of map-reduce-reduce iterations to workers so as to minimize communication overheads within and across iterations.

Fault Tolerance. Traditional MapReduce runtimes provide fault tolerance by storing output to a replicated file system and automatically restarting failed tasks. Since we expect ticks to be quite short and they are processed in main memory, it would be prohibitively expensive to write output to stable storage between every tick. Furthermore, since individual ticks are short, the benefit from restarting a task is likely to be small.

We employ epoch synchronization with the master to trigger *coordinated checkpoints* [20] of the main memory of the workers. As the master determines a pre-defined tick boundary for checkpointing, the workers can write their checkpoints independently without global synchronization. As we expect iterations to be short, failures are handled by re-execution of all iterations since the last checkpoint, a common technique in scientific simulations. In fact, we can leverage previous literature to tune the checkpointing interval to minimize the total expected runtime of the whole computation [14].

Partitioning and Load Balancing. As we have observed in Section 3.2, bounded reachability implies that a given spatial partitioning will remain effective for a number of map-reduce-reduce iterations. Our runtime uses that observation to keep data partitioning

stable over time and re-evaluates it at epoch boundaries.

At the beginning of the simulation, the master computes a partitioning function based on the visible regions of the agents and then broadcasts this partitioning to the worker nodes. Each worker becomes responsible for one region of the partitioning. Agents change partitions slowly, but over time the overall spatial distribution may change quite dramatically. For example, the distribution of traffic on a road network is likely to be very different at morning rush hour than at evening rush hour. This would cause certain nodes to become overloaded if we used the same partitioning in both cases. To address this, the master periodically receives statistics from the workers about the number of agents in the owned region and the communication and processing costs. The master then decides on repartitioning by balancing the cost of redistribution with its expected benefit [26]. If the master decides to modify the partitioning, it broadcasts the new partitioning to all workers. The workers then switch to the new partitioning at a specified epoch boundary.

Collocation of Tasks. Since simulations run for many iterations, it is important to minimize communication between tasks. We accomplish this by collocating the map and reduce tasks for a partition on the same node so that agents that do not switch partitions can be exchanged using shared memory rather than the network. Since agents have limited reachable regions, the owned set of each partition is likely to remain relatively stable across ticks, and so will remain on the same node. Agents still need to be replicated, but their primary copies do not have to be redistributed. This idea was previously explored by the Phoenix project for SMP systems [55] and the Map-Reduce-Merge project for individual joins [53], but it is particularly important for long-running behavioral simulations.

Figure 1 shows how collocation works when we allow non-local effect assignments. Solid arrows (both thick and thin) indicate the flow of agents during a tick. Each node processes a map task and two reduce tasks as described in Section 3.1. The map task replicates agents as necessary and sends them to the appropriate reduce tasks. The first-level reducers compute local effects and send non-local effects to the second-level reducers, which aggregate all (local and non-local) effects and send them to the map tasks for the next tick. Because of the neighborhood property, only agents that are near partition boundaries need to be replicated, and agents change partitions infrequently. Most messages are between tasks on the same node (as indicated by the thick blue arrows), enabling communication by shared memory rather than the network. This optimization exploits dependencies on the data between tasks, much in line with schemes for relaxing communication in bulk-synchronous computations, e.g., [46].

When we put together all of the optimizations above, the runtime of BRACE resembles an optimized shared-nothing message-passing environment. However, BRACE still exposes an intuitive map-reduce programming abstraction as an interface to this runtime. As we show in the next section, this interface enables us to compile a high-level language for domain scientists into efficient map-reduce-reduce pipelines over BRACE.

4. PROGRAMMING AGENT BEHAVIOR

In this section, we show how to offer a simple programming model for a domain scientist, targeting the last requirement of Section 1.1. MapReduce is set-based; a program describes how to process all of the elements in a collection. Simulation developers prefer to describe the behavior of their agents individually, and use message-passing techniques to communicate between agents. This type of programming is closer to the scientific models that describe agent behavior.

```

class Fish {
  // The fish location
  public state float x : (x+vx); #range[-1,1];
  public state float y : (y+vy); #range[-1,1];

  // The latest fish velocity
  public state float vx : vx + rand() + avoidx / count * vx;
  public state float vy : vy + rand() + avoidy / count * vy;

  // Used to update our velocity
  private effect float avoidx : sum;
  private effect float avoidy : sum;
  private effect int count : sum;

  /** The query-phase for this fish. */
  public void run() {
    // Use "forces" to repel fish too close
    foreach(Fish p : Extent<Fish>) {
      p.avoidx <- 1 / abs(x - p.x);
      p.avoidy <- 1 / abs(y - p.y);
      p.count <- 1;
    }
  }
}

```

Figure 2: Class for Simple Fish Behavior

We introduce a new programming language – BRASIL, the Big Red Agent SIMulation Language. BRASIL embodies agent centric programming with explicit support for the state-effect pattern, and performs further algebraic optimizations. It bridges the mental model of simulation developers and our MapReduce processing techniques for behavioral simulations. We provide an overview of the main features of BRASIL (Section 4.1) and describe algebraic optimization techniques that can be applied to our scripts (Section 4.2). Formal semantics for our language as well as the proofs of theorems in this section are provided in Appendix B.1.

4.1 Overview of BRASIL

BRASIL is an object-oriented language in which each object corresponds to an agent in the simulation. Agents in BRASIL are defined in a class file that looks superficially like Java. The programmer can specify fields, methods, and constructors, which can be either public or private. Unlike in Java, each field in a BRASIL class must be tagged as either *state* or *effect*. The BRASIL compiler enforces the read-write restrictions of the state-effect pattern over those fields as described in Section 2. Figure 2 is an example of a simple two-dimensional fish simulation, in which fish swim about randomly, avoiding each other with imaginary repulsion “forces”.

Recall that the state-effect pattern divides computation into query and update phases. In BRASIL, the query phase for an agent class is expressed by its `run()` method. State fields are read-only and effect assignments are aggregated using the functions specified at the effect field declarations. This is similar to aggregator variables in Sawzall [42]. In our fish simulation example, each fish repels nearby fish with a “force” inversely proportional to distance. The update phase is specified by *update rules* attached to the state field declarations. These rules can only read values of other fields in this agent. In our example, fish velocity vectors are updated based on the avoidance effects plus a random perturbation.

BRASIL has some important restrictions. First, it only supports iteration over a finite set or list via a `foreach`-loop. This eliminates unbounded looping, which is not available in algebraic database languages. Second, there is an interplay between `foreach`-loops and effects: effect variables can only be read outside of a `foreach`-loop, and all assignments within a `foreach`-loop are aggregated. This powerful restriction allows us to treat the entire program, not just the communication across map and reduce operations, as a dataflow query plan.

BRASIL also has a special programming construct to enforce the

neighborhood property outlined in Section 2. Every state field that encodes spatial location is tagged with a visibility and reachability constraint. For example, the constraint `range[-1,1]` attached to the `x` field in Figure 2 means that each fish can inspect others whose `x` coordinate is at most $[-1, 1]$ different to its own in the query phase, and can change its `x` coordinate by at most $[-1, 1]$ in the update phase. The language runtime enforces these constraints, as described in Appendix B.

4.2 Optimization

We compile BRASIL into a well-understood dataflow language. In our previous work on computer games, we used the relational algebra to represent our data flow [51]. However, for distributed simulations, we have found the monad algebra [7, 31, 41, 48] – the theoretical foundation for XQuery [31] – to be a much more appropriate fit. In particular, the monad algebra has a `MAP` primitive for descending into the components of its nested data model; this makes it a much more natural companion to MapReduce than the relational algebra.

We present the formal translation to the monad algebra in Appendix B, together with several theorems regarding its usage in optimization. Most of these optimizations are the same as those that would be present in a relational algebra query plan: algebraic rewrites and automatic indexing. In fact, any monad algebra expression on flat tables can be converted to an equivalent relational algebra expression and *vice versa* [41]; rewrites and indexing on the relational form carry back into the monad algebra form. In particular, many of the techniques used by Pathfinder [49] to process XQuery with relational query plans apply to the monad algebra.

Effect Inversion. An important optimization that is unique to our framework involves eliminating non-local effects. When non-local effect assignments can be eliminated, we are able to process each tick with a single reduce pass instead of two (Section 3). Consider again the program of Figure 2. We may rewrite its `foreach`-loop as

```

foreach(Fish p : Extent<Fish>) {
  avoidx <- 1 / abs(p.x - x);
  avoidy <- 1 / abs(p.y - y);
  count <- 1;
}

```

This rewritten expression does not change the results of the simulation, but only assigns effects locally. We call this transformation *effect inversion*, and it is always possible at some cost in visibility range. Indeed, in Appendix B.2, we formally prove the following:

THEOREM 3 If the visibility constraint on a script is distance d , there is an equivalent script with a visibility constraint at most $2d$ that uses only local effect assignments.

Increasing the visibility bound increases the number of replicas required at each node. Hence this optimization eliminates one communication round at the cost of sending more information in the remaining round.

5. EXPERIMENTS

In this section, we present experimental results using two distinct real-world behavioral simulation models we have coded using BRACE. We focus on the following: (i) We validate the effectiveness of the BRASIL optimizations introduced in Section 4.2. In fact, these optimizations allow us to approach the efficiency of hand-optimized simulation code (Section 5.2); (ii) We evaluate BRACE’s MapReduce runtime implementation over a cluster of machines. We measure simulation scale-up via spatial data partitioning as well as load balancing (Section 5.3).

5.1 Setup

Implementation. The prototype BRACE MapReduce runtime is implemented in C++ and uses MPI for inter-node communication. Our BRASIL compiler is written in Java and directly generates C++ code that can be compiled with the runtime. Our prototype includes a generic KD-tree based spatial index capability [4]. We use a simple rectilinear grid partitioning scheme, which assigns each grid cell to a separate slave node. A one-dimensional load balancer periodically receives statistics from the slave nodes, including computational load and number of owned agents; from these it heuristically computes a new partition trying to balance improved performance against estimated migration cost. Checkpointing is not yet integrated into BRACE’s implementation at the time of writing.

We plan to integrate more sophisticated algorithms for all these components in future work. But our current prototype already demonstrates good performance and automatic scaling of realistic behavioral simulations written in BRASIL.

Simulation Workloads. We have implemented realistic traffic and fish school simulations in BRASIL. The traffic simulation includes the lane-changing and acceleration models of the state-of-the-art, open-source MITSIM traffic simulator [54]. MITSIM is a single-node program, so we compare its performance against our BRASIL reimplementations of its model also running on a single node. We simulate a linear segment of highway, and scale-up the size of the problem by extending the length of the segment.

The fish simulation implements a recent model of information flow in animal groups [13]. In this model the “ocean” is unbounded, and the spatial distribution of fish changes dramatically as “informed individuals” guide the movements of others in the school.

Neither of these simulations uses non-local effect assignments; therefore we need only a single reducer per node. To evaluate our effect inversion optimization, we modified the fish simulation to create a predator simulation that uses non-local assignments. It is similar in spirit to artificial society simulations [30]. Appendix C describes these simulation models in more detail. We measure total simulation time in our single-node experiments and tick throughput (agent ticks per second) when scaling up over multiple nodes. In all measurements we eliminate start-up transients by discarding initial ticks until a stable tick rate is achieved.

Hardware Setup. We ran all of our experiments in the Cornell Web Lab cluster [3]. The cluster contains 60 nodes interconnected by a pair of 1 gigabit/sec Port Summit X450a Ethernet Switches. Each node has two Quad Core Intel Xeon, 2.66GHz, processors with 4MB cache each and 16 GB of main memory.

5.2 BRASIL Optimizations

We first compare the single-node performance of our traffic simulation to the single-node program MITSIM. The main optimization in this case is spatial indexing. For a meaningful comparison, we validate the aggregate traffic statistics produced by our BRASIL reimplementations against those produced by MITSIM. Details of our validation procedure appear in Appendix C.

Figure 3 compares the performance of MITSIM against BRACE using our BRASIL reimplementations of its model. Without spatial indexing, BRACE’s performance degrades quadratically with increasing segment length. This is expected: In this simulation, the number of agents grows linearly with segment length; and without indexing every vehicle enumerates and tests every other vehicle during each tick. With spatial indexing enabled, BRACE converts this behavior to an orthogonal range query, resulting in log-linear growth, as confirmed by Figure 3. BRACE’s spatial indexing achieves performance that is comparable, but inferior to MITSIM’s

hand-coded nearest-neighbor implementation. Our optimization techniques generalize to nearest-neighbor indexing, and adding this to BRACE is planned future work. With this enhancement, we expect to achieve performance parity with MITSIM.

We observed similar log-linear versus quadratic performance when scaling up the number of agents in the fish simulation in a single node. We thus omit these results. When we increase the visibility range, however, the performance of the KD-tree indexing decreases, since more results are produced for each index probe (Figure 4). Still, indexing yields from two to three times improvement over a range of visibility values.

In addition to indexing, we also measure the performance gain of eliminating non-local effect assignments through effect inversion. Only the predator simulation has non-local effect assignments, so we report results exclusively on this model. We run two versions of the predator simulation, one with non-local assignments and the other with non-local assignments eliminated by effect inversion. We run both scripts with and without KD-tree indexing enabled on 16 slave nodes, and with BRACE configured to have two reduce passes in the first case and only a single reduce pass in the second case. Our results are displayed in Figure 5. Effect inversion increases agent tick throughput from 3.59 million (Idx-Only) to 4.36 million (Idx+Inv) with KD-tree indexing enabled, and from 2.95 million (No-Opt) to 3.63 million (Inv-Only) with KD-tree indexing disabled. This represents an improvement of more than 20% in each case, demonstrating the importance of this optimization.

5.3 Scalability of the BRACE Runtime

We now explore the parallel performance of BRACE’s MapReduce runtime on the traffic and fish school simulations as we scale the number of slave nodes from 1 to 36. The size of both simulations is scaled linearly with the number of slaves, so we measure scale-up rather than speed-up.

The traffic simulation represents a linear road segment with constant up-stream traffic. As a result, the distribution on the segment is nearly uniform, and load is always balanced among the nodes. Therefore, throughput grows linearly with the number of nodes even if load balancing is disabled (Figure 6). The sudden drop around 20 nodes is an artifact of IP routing in the multi-switch configuration of the WebLab cluster on which we ran our experiments.

In the fish simulation, fish move in schools led by informed individuals [13]. In our experiment, there are two classes of informed individuals, trying to move in two different fixed directions. The spatial distribution of fish, and consequently the load on each slave node, changes over time. Figure 7 shows the scalability of this simulation with and without load balancing. Without load balancing, two fish schools eventually form in nodes at the extremes of simulated space, while the load at all other nodes falls to zero. With load balancing, partition grids are adjusted periodically to assign roughly the same number of fish to each node, so throughput increases linearly with the number of nodes.

Figure 8 confirms this, where still 16 slave nodes are used and we sample the time per epoch every 100 epochs, for readability, but also display all outliers. With load balancing enabled, the time per simulation epoch is essentially flat; with load balancing disabled, the epoch time gradually increases to a value that reflects all agents being simulated by only two nodes.

6. RELATED WORK

Much of the existing work on behavioral simulations has focused on task-parallel discrete event simulation systems, where computations are executed according in response to events instead of in time steps [9, 39, 15, 35, 57]. In contrast to our data-parallel so-

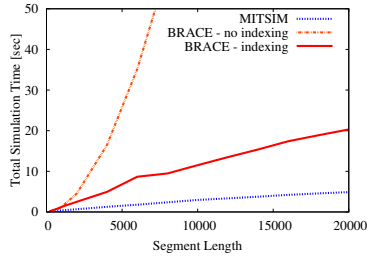


Figure 3: Traffic: Indexing vs. Seg. Length

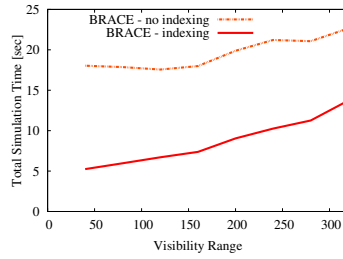


Figure 4: Fish: Indexing vs. Visibility

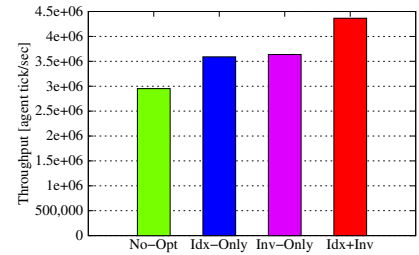


Figure 5: Predator: Effect Inversion

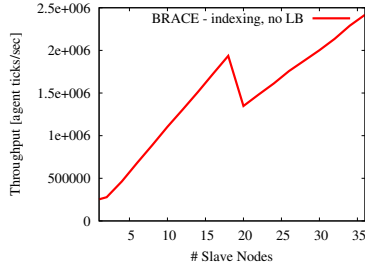


Figure 6: Traffic: Scalability

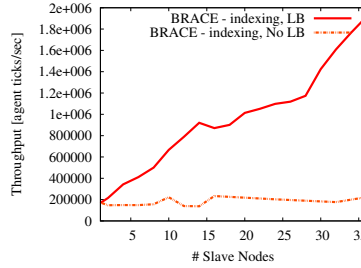


Figure 7: Fish: Scalability

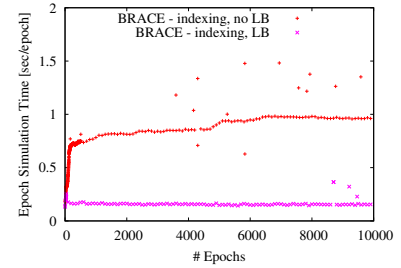


Figure 8: Fish: Load Balancing

lution for behavioral simulations based on the state-effect pattern, these systems implement either conservative or optimistic protocols to detect conflicts and preempt or rollback simulation tasks. The strength of local interactions and the time-stepped model used in behavioral simulations lead to unsatisfactory performance, as shown in attempts to adapt discrete event simulators to agent-based simulations [29, 28]. Other task parallel approaches, such as using one concurrent transaction per agent [40], also fail to scale due to similar reasons.

Platforms specifically targeted at agent-based models have been developed, such as Swarm [36], Mason [33], and Player/Stage [22]. These platforms offer tools to facilitate simulation programming, but most rely on message-passing abstractions with implementations inspired by discrete event simulators, so they suffer in terms of performance and scalability. A few recent systems attempt to distribute agent-based simulations over multiple nodes without exploiting application properties such as visibility and time-stepping [27, 43]. This leads either to poor scale-up or to unrealistic restrictions on agent interactions.

Regarding join processing with MapReduce, Zhang et al. [56] compute spatial joins by an approach similar to ours when only local effect assignments are allowed. Their mapper partitions are derived using spatial index techniques rather than by reasoning about the application program, and they do not discuss *iterated* joins, an important consideration for our work. Locality optimizations have been studied for MapReduce on SMPs [55] and for MapReduce-Merge [53]; in this paper we consider the problem in a distributed main memory MapReduce runtime.

Data-driven parallelization techniques have also been studied in parallel databases [17, 23] and data parallel programming languages [5, 45]. However, it is unnatural and inefficient to use either SQL or set-operations exclusively to express flexible computation over individuals as required for behavioral simulations.

Given this situation, behavioral simulation developers have resorted to hand-coding parallel implementations of specific simulation models [21, 38], or trading model accuracy for scalability and ease of implementation [8, 50]. Our work, in contrast, aligns in spirit with recent efforts to bring dataflow programming techniques to complex domains, such as distributed systems [1] and networking [32], with huge benefits in performance and program-

ing productivity. To the best of our knowledge, our approach is the first to bring both programmability and scalability through data parallelism to behavioral simulations.

7. CONCLUSIONS

In this paper we show how MapReduce can be used to scale behavioral simulations across clusters by abstracting these simulations as iterated spatial joins. To efficiently distribute these joins we leverage several properties of behavioral simulations to get a shared-nothing, in-memory MapReduce framework called BRACE, which exploits collocation of mappers and reducers to bound communication overhead. In addition, we present a new scripting language for our framework called BRASIL, which hides all the complexities of modeling computations in MapReduce and parallel programming from domain scientists. BRASIL scripts can be compiled into our MapReduce framework and allow for algebraic optimizations in mappers and reducers. We perform an experimental evaluation with two real-world behavioral simulations to show that BRACE has nearly linear scalability as well as single-node performance comparable to a hand-coded simulator.

Acknowledgments. This material is based upon work supported by the New York State Foundation for Science, Technology, and Innovation under Agreement C050061, by the National Science Foundation under Grants 0725260 and 0534404, by the iAd Project funded by the Research Council of Norway, by the AFOSR under Award FA9550-10-1-0202, and by Microsoft. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding agencies.

8. REFERENCES

- [1] P. Alvaro, T. Condie, N. Conway, K. Elmeleegy, J. M. Hellerstein, and R. Sears. Boom analytics: exploring data-centric, declarative programming for the cloud. In *EuroSys*, pages 223–236, 2010.
- [2] Amazon Elastic Compute Cloud (Amazon EC2). <http://aws.amazon.com/ec2/>.
- [3] W. Arms, S. Aya, P. Dmitriev, B. Kot, R. Mitchell, and L. Walle. Building a research library for the history of the web. In *Proc. Joint Conference on Digital Libraries*, 2006.
- [4] J. Bentley. K-d trees for semidynamic point sets. In *Proc. SGC*, 1990.

- [5] G. Blelloch. Programming parallel algorithms. *CACM*, 39(3):85–97, 1996.
- [6] J. Buhl, D. Sumpter, I. Couzin, J. Hale, E. Despland, E. Miller, and S. Simpson. From disorder to order in marching locusts. *Science*, 312(5778):1402–1406, 2006.
- [7] P. Buneman, S. Naqvi, V. Tannen, and L. Wong. Principles of programming with complex objects and collection types. *Theor. Comput. Sci.*, 149(1):3–48, 1995.
- [8] N. Cetin, A. Burri, and K. Nagel. A large-scale agent-based traffic microsimulation based on queue model. In *Proc. Swiss Transport Research Conference (STRC)*, 2003.
- [9] K. Chandy and J. Misra. Distributed simulation: A case study in design and verification of distributed programs. *IEEE Transactions on Software Engineering*, 5(5):440–452, 1978.
- [10] C. Choudhury, M. Ben-Akiva, T. Toledo, A. Rao, and G. Lee. NGSIM cooperative lane changing and forced merging model. Technical report, Federal Highway Administration, 2006. FHWA-HOP-07-096.
- [11] C. Choudhury, T. Toledo, and M. Ben-Akiva. NGSIM freeway lane selection model. Technical report, Federal Highway Administration, 2004. FHWA-HOP-06-103.
- [12] Cornell Database Group. The brasil language specification. <http://www.cs.cornell.edu/bigreddata/games/simulations.php>, 2010.
- [13] I. Couzin, J. Krause, N. Franks, and S. Levin. Effective leadership and decision-making in animal groups on the move. *Nature*, 433(7025):513–516, 2005.
- [14] J. Daly. A higher order estimate of the optimum checkpoint interval for restart dumps. *Future Generation Computer Systems*, 22(3):303–312, 2006.
- [15] S. Das, R. Fujimoto, K. Panesar, D. Allison, and M. Hybinette. GTW: A time warp system for shared memory multiprocessors. In *Proc. Winter Simulation Conference (WSC)*, 1994.
- [16] J. Dean and S. Ghemawat. Mapreduce: Simplified data processing on large clusters. In *Proc. OSDI*, pages 137–150, 2004.
- [17] D. Dewitt, S. Ghandeharizadeh, D. Schneider, A. Bricker, H. Hsiao, and R. Rasmussen. The gamma database machine project. *IEEE Trans. on Knowledge and Data Engineering*, 2(1):44–62, 1990.
- [18] J. Dittrich, L. Blunski, and M. A. V. Salles. Indexing moving objects using short-lived throwaway indexes. In *Proc. SSTD*, 2009.
- [19] J. Ekanayake, S. Pallickara, and G. Fox. Mapreduce for data intensive scientific analyses. In *Proc. eScience*, 2008.
- [20] M. Elnozahy, L. Alvisi, Y.-M. Wang, and D. B. Johnson. A Survey of Rollback-Recovery Protocols in Message-Passing Systems. *ACM Computing Surveys*, 34(3):375–408, 2002.
- [21] U. Erra, B. Frola, V. Scaranò, and I. Couzin. An efficient gpu implementation for large-scale individual based simulation of collective behavior. In Press.
- [22] B. Gerkey, R. Vaughan, and A. Howard. The player/stage project: Tools for multi-robot and distributed sensor systems. In *Proc. Int. Conf. on Advanced Robotics (ICAR)*, 2003.
- [23] G. Graefe. Encapsulation of parallelism in the volcano query processing system. In *Proc. SIGMOD*, 1990.
- [24] D. Grünbaum. Behavior: Align in the sand. *Science*, 312(5778):1320–1322, 2006.
- [25] Hadoop. <http://hadoop.apache.org/>.
- [26] B. Hendrickson and K. Devine. Dynamic load balancing in computational mechanics. *Computer Methods in Applied Mechanics and Engineering*, 184(2-4):485–500, 2000.
- [27] D. Ho, T. Bui, and N. Do. Dividing agents on the grid for large scale simulation. *Intelligent Agents and Multi-Agent Systems*, 5357:222–230, 2008.
- [28] B. Horling, R. Mailler, and V. Lesser. Farm: A scalable environment for multi-agent development and evaluation. *Software Engineering for Multi-Agent Systems II*, 2940:364–367, 2004.
- [29] M. Hybinette, E. Kraemer, Y. Xiong, G. Matthews, and J. Ahmed. SASSY: A design for a scalable agent-based simulation system using a distributed discrete event infrastructure. In *Proc. Winter Simulation Conference (WSC)*, 2006.
- [30] Joshua Epstein and Robert Axtell. *Growing Artificial Societies: Social Science from the Bottom Up*. The MIT Press, 1996.
- [31] C. Koch. On the complexity of nonrecursive xquery and functional query languages on complex values. In *Proc. PODS*, 2005.
- [32] B. Loo, T. Condie, M. Garofalakis, D. Gay, J. Hellerstein, P. Maniatis, R. Ramakrishnan, T. Roscoe, and I. Stoica. Declarative networking: Language, execution and optimization. In *Proc. SIGMOD*, 2006.
- [33] S. Luke, C. Cioffi-Revilla, L. Panait, K. Sullivan, and G. Balan. MASON: A multiagent simulation environment. *Simulation*, 81, 2005.
- [34] G. Luo, J. Naughton, and C. Ellmann. A non-blocking parallel spatial join algorithm. In *Proc. ICDE*, 2002.
- [35] F. Mattern. Efficient algorithms for distributed snapshots and global virtual time approximation. *Journal of Parallel and Distributed Computing*, 18(4):423–434, 1993.
- [36] N. Minar, R. Burkhart, C. Langton, and M. Askenazi. The swarm simulation system: A toolkit for building multi-agent simulations. Working Papers 96-06-042, Santa Fe Institute, 1996.
- [37] K. Mouratidis, D. Papadias, and M. Hadjieleftheriou. Conceptual partitioning: an efficient method for continuous nearest neighbor monitoring. In *Proc. SIGMOD*, 2005.
- [38] K. Nagel and M. Rickert. Parallel implementation of the transims micro-simulation. *Parallel Computing*, 27(12):1611–1639, 2001.
- [39] D. Nicol. The cost of conservative synchronization in parallel discrete event simulations. *Journal of the ACM*, 40(2):304–333, 1993.
- [40] T. Ohmori and M. Hoshi. Gaming-simulations of multi-agent information systems using large databases: The concept and database algorithms. In *Proc. DASFAA*, 1995.
- [41] J. Paredaens and D. van Gucht. Possibilities and limitations of using flat operators in nested algebra expressions. In *Proc. PODS*, 1998.
- [42] R. Pike, S. Doward, R. Griesemer, and S. Quinlan. Interpreting the data: Parallel analysis with Sawzall. *Scientific Programming Journal*, 13(4):227–204, 2005.
- [43] M. Scheutz, P. Schermerhorn, R. Connaughton, and A. Dingler. SWAGES—an extendable parallel grid experimentation system for large-scale agent-based alife simulations. In *Proc. Artificial Life X*, 2006.
- [44] D. Schrank and T. Lomax. The 2009 urban mobility report. Technical report, Texas Transportation Institute, 2009.
- [45] J. T. Schwartz, R. B. Dewar, E. Schonberg, and E. Dubinsky. *Programming with sets; an introduction to SETL*. Springer-Verlag New York, Inc., New York, NY, USA, 1986.
- [46] J. soo Kim, S. Ha, and C. S. Jhon. Relaxed barrier synchronization for the bsp model of computation on message-passing architectures. *Information Processing Letters*, 66, 1998.
- [47] D. Suciú. *Parallel programming languages for collections*. PhD thesis, University of Pennsylvania, 1995.
- [48] V. Tannen, P. Buneman, and L. Wong. Naturally embedded query languages. In *Proc. ICDT*, 1992.
- [49] J. Teubner. *Pathfinder: XQuery Compilation Techniques for Relational Database Targets*. PhD thesis, Technische Universität München, 2006.
- [50] Y. Wen. *Scalability of Dynamic Traffic Assignment*. PhD thesis, Massachusetts Institute of Technology, 2008.
- [51] W. White, A. Demers, C. Koch, J. Gehrke, and R. Rajagopalan. Scaling games to epic proportions. In *Proc. SIGMOD*, 2007.
- [52] W. White, B. Sowell, J. Gehrke, and A. Demers. Declarative processing for computer games. In *Proc. SIGGRAPH Sandbox Symposium*, 2008.
- [53] H.-C. Yang, A. Dasdan, R.-L. Hsiao, and D. S. Parker. Map-reduce-merge: simplified relational data processing on large clusters. In *Proc. SIGMOD*, 2007.
- [54] Q. Yang, H. Koutsopoulos, and M. Ben-Akiva. A simulation laboratory for evaluating dynamic traffic management systems. In *Annual Meeting of Transportation Research Board*, 1999. TRB Paper No. 00-1688.
- [55] R. M. Yoo, A. Romano, and C. Kozyrakis. Phoenix rebirth: Scalable mapreduce on a large-scale shared-memory system. In *IISWC*, 2009.
- [56] S. Zhang, J. Han, Z. Liu, K. Wang, and Z. Xu. Sjm: Parallelizing spatial join with mapreduce on clusters. In *CLUSTER*, 2009.
- [57] G. Zheng, G. Kakulapati, and L. Kale. Bigsim: A parallel simulator for performance prediction of extremely large parallel machines. In *Proc. IPDPS*, 2004.

APPENDIX

A. SPATIAL JOINS IN MAPREDUCE

In this appendix, we first briefly review the MapReduce model (Appendix A.1). We then formalize the spatial joins run at each tick of a behavioral simulation in MapReduce (Appendix A.2).

A.1 MapReduce

Since its introduction in 2004, MapReduce has become one of the most successful models for processing long running computations in distributed shared-nothing environments [16]. While it was originally designed for very large batch computations, MapReduce is ideal for behavioral simulations because it provides automatic scalability, which is one of the key requirements for next-generation platforms. By varying the degree of data partitioning and the corresponding number of map and reduce tasks, the same MapReduce program can be run on one machine or one thousand.

The MapReduce programming model is based on two functional programming primitives that operate on key-value pairs. The map function takes a key-value pair and produces a set of intermediate key-value pairs, $\text{map} : (k_1, v_1) \rightarrow [(k_2, v_2)]$, while the reduce function collects all of the intermediate pairs with the same key and produces a value, $\text{reduce} : (k_2, [v_2]) \rightarrow [v_3]$. Since simulations operate in ticks, we use an iterative MapReduce model in which the output of the reduce step is fed into the next map step. Formally, this means that we change the output of the reduce step to be $[(k_3, v_3)]$.

A.2 Formalizing Simulations in MapReduce

In the following, we formally develop the map and reduce functions for processing a single tick of a behavioral simulation.

Formalizing Agents and Spatial Partitioning. We first introduce our notation for agents and their state and effect attributes. We denote an agent a as $a = \langle \text{oid}, \mathbf{s}, \mathbf{e} \rangle$, where \mathbf{s} is a vector of the agent’s state attributes and \mathbf{e} is a vector of its effects. To refer to an agent or its attributes at a tick t , we will write a^t , \mathbf{s}^t , or \mathbf{e}^t . Since effect attributes are aggregated using combinator functions, they need to be reset at the end of every tick. We will use θ to refer to the vector of idempotent values for each effect. Finally, we use \oplus to denote the aggregate operator that combines effect variables according to the appropriate combinator.

The neighborhood property implies that some subset of each agent’s state attributes are spatial attributes that determine an agent’s position. For an agent $a = \langle \text{oid}, \mathbf{s}, \mathbf{e} \rangle$, we denote this spatial location $\ell(\mathbf{s}) \in \mathcal{L}$, where \mathcal{L} is the spatial domain. Given an agent a at location l , the visible region of a is $VR(l) \subseteq \mathcal{L}$.

Both the map and reduce tasks in our framework will have access to a spatial partitioning function $P : \mathcal{L} \rightarrow \mathbf{P}$, where \mathbf{P} is a set of partition ids. This partitioning function can be implemented in multiple ways, such as a regular grid or a quadtree. We define the *owned* set of a partition p as the inverse image of p under P , i.e., the set of all locations assigned to p . Since each location has an associated visible region, we can also define the visible region of a partition as $VR(p) = \bigcup_{l \in \mathcal{L}, P(l)=p} VR(l)$. This is the set of all locations that might be visible by some agent in p .

Simulations with Local Effects Only. Since the query phase of an agent can only depend on the agents inside its visible region, the visible region of a partition contains all of the data necessary to execute the query phase for its entire owned region. We will take advantage of this by *replicating* all of the agents in this region at p so that the query phase can be executed without communication.

Figure 9 shows the map and reduce functions for processing tick t when there are only local effect assignments. At tick t , the map function performs the update phase from the previous tick, and the

$$\begin{aligned} \text{map}^t(\cdot, a^{t-1}) &= [(p, \langle \text{oid}, \mathbf{s}^t, \theta \rangle) \text{ where } \ell(\mathbf{s}^t) \in VR(p)] \\ \text{reduce}^t(p, [\langle \text{oid}_i, \mathbf{s}_i^t, \theta \rangle]) &= [(p, \langle \text{oid}_i, \mathbf{s}_i^t, \mathbf{e}_i^t \rangle), \forall i \text{ s.t. } P(\ell(\mathbf{s}_i^t)) = p] \end{aligned}$$

Figure 9: Map and reduce functions with local effects only

$$\begin{aligned} \text{map}_1^t(\cdot, a^{t-1}) &= [(p, \langle \text{oid}, \mathbf{s}^t, \theta \rangle) \text{ where } \ell(\mathbf{s}^t) \in VR(p)] \\ \text{reduce}_1^t(p, [\langle \text{oid}_i, \mathbf{s}_i^t, \theta \rangle]) &= [(P(\ell(\mathbf{s}^t)), \langle \text{oid}_i, \mathbf{s}_i^t, \mathbf{f}_i^t \rangle)] \\ \text{map}_2^t(k, a) &= (k, a) \\ \text{reduce}_2^t(p, [\langle \text{oid}_i, \mathbf{s}_i^t, \mathbf{f}_i^t \rangle]) &= [(p, \langle \text{oid}_i, \mathbf{s}_i^t, \oplus_j \mathbf{f}_j^t \rangle), \forall j \text{ s.t. } \text{oid}_i = \text{oid}_j] \end{aligned}$$

Figure 10: Map and reduce functions with non-local effects

reduce function performs the query phase. The map function takes as input an agent with state and effect variables from the previous tick (a^{t-1}), and updates the state variables to \mathbf{s}^t and the effect attributes to θ . During the very first tick of the simulation, \mathbf{e}^{t-1} is undefined, so \mathbf{s}^t will be set to a value reflecting the initial simulation state. The map function emits a copy of the updated agent keyed by partition for each partition containing the agent in its visible set ($\ell(\mathbf{s}^t) \in VR(p)$). This has the effect of replicating the agent a to every partition that might need it for query phase processing. The amount of replication depends on the partitioning function and on the size of each agent’s visible region.

The reduce function receives as input all agents that are sent to a particular partition p . This includes the agents in p ’s owned region, as well as replicas of all the agents that fall in p ’s visible region. The reducer will execute the query phase and compute effect variables for all of the agents in its owned region (agent i s.t. $P(\ell(\mathbf{s}_i^t)) = p$). This requires no communication, since the query phase of an agent in p ’s owned region can only depend on the agents in p ’s visible region, all of which are replicated at the same reducer. The reducer outputs agents with updated effect attributes to be processed in the next tick.

Simulations with Non-Local Effects. The method above only works when all effect assignments are local. If an agent a makes an effect assignment to some agent b in its visible region, then it must communicate that effect to the reducer responsible for processing b . Figure 10 shows the map and reduce functions to handle simulations with non-local effect assignments. The first map function task is same as in the local effect case. Each agent is partitioned and replicated as necessary. As before, the first reduce function computes the query phase for the agents in p ’s owned set and computes effect values. In this case, however, it can only compute intermediate effect values \mathbf{f}^t , since it does not have the effects assigned at other nodes. This reducer outputs one pair for every agent, including replicas. These agents are keyed with the partition that owns them, so that all replicas of the same agent go to the same node.

The second map function is the identity, and the second reduce function performs the aggregation necessary to complete the query phase. It receives all of the updated replicas of all of the agents in its owned region and applies the \oplus operation to compute the final effect values and complete the query phase. Each reducer will output an updated copy of each agent in its owned set.

B. FORMAL SEMANTICS OF BRASIL

In this section, we provide a more formal presentation of the semantics of BRASIL than the one presented in Section 4. In particular, we show how to convert BRASIL expressions into monad algebra expressions for analysis and optimization. We also prove several results regarding effect inversion, introduced in Section 4.2, and illustrate the resulting trade-offs between computation and

communication.

For the most part, our work will be in the traditional monad algebra. We refer the reader to the original work on this algebra [7, 31, 41, 48] for its basic operators and nested data model. We also use standard definitions for the derived operations like cartesian product and nesting. For example, we define cartesian product as

$$f \times g := \langle 1 : f, 2 : g \rangle \circ \text{PAIRWITH}_1 \circ \text{FLATMAP}(\text{PAIRWITH}_2) \quad (1)$$

For the purpose of readability, composition in (1) and the rest of our presentation, is read left-to-right; that is, $(f \circ g)(x) = g(f(x))$.

We assume that the underlying domain is the real numbers, and that we have associated arithmetic operators. We also add traditional aggregate functions like COUNT and SUM to the algebra; these functions take a set of elements (of the appropriate type) and return a value.

In order to simplify our presentation, we do make several small changes that relax the typing constraints in the classic monad algebra. In particular, we want to allow union to combine sets of tuples with different types. For this end, we introduce a special NIL value. This value is the result of any query that is undefined on the input data, such as projection on a nonexistent attribute. This value has a form of “null-semantics” in that values combined with NIL are NIL, and NIL elements in a set are ignored by aggregates. In addition, we introduce a special aggregate function GET. When given a set, this function returns its contents if it is a singleton, and returns NIL otherwise. Neither this function, nor the presence of NIL significantly affects the expressive power of the monad algebra [47].

B.1 Monad Algebra Translation

For the purpose of illustration, we assume that our simulation has only one class of agents, all of which are running the same simulation script. It is relatively easy to generalize our approach to multiple agent classes or multiple scripts. Given this assumption, our simulation data is simply a set of tuples $\{t_0, \dots, t_n\}$ where each tuple t_i represents the data inside of an agent. Every agent tuple has a special attribute KEY which is used to uniquely identify the agent; variables which reference another agent make use of this key. The state-effect pattern requires that all data types other than agents be processed by value, so they can safely be stored inside each agent.

We let τ represent the type/schema of an agent. In addition to the key attribute, τ has an attribute for each state and effect field. The value of a state attribute is the value of the field. The value of an effect attribute is a pair $\langle 1 : n, 2 : \text{AGG} \rangle$ where n is a unique identifier for the field and AGG is the aggregate for this effect.

During the query phase, we represent effects as a tuple $\langle k : \mathbb{N}, e : \mathbb{N}, v : \sigma \rangle$, where k is the key of the object being effected, e is the effect field identifier, and v is the value of the effect. As a shorthand, let ρ be this type. Even though effects may have different types, because of our relaxed typing, this will not harm our formalism.

The syntax of BRASIL forces the programmer to clearly separate the code into a query script (i.e. `run()`) and an update script (the update rules). A query script compiles to an expression whose input and output are the tuple $\langle 1 : \tau', 2 : \{\tau\}, 3 : \{\rho\} \rangle$. The first element represents the active agent for this script; τ' “extends” type τ in that it is guaranteed to have an attribute for the key and each state field, but it may have more attributes. The second element is the set of all other agents with which this agent must communicate. The last element is the set of effects generated by this script.

Let Q be the monad expression for the query script. Then the effect generation stage is the expression

$$\mathfrak{Q}(Q) = (\text{ID} \times \text{ID}) \circ \text{NEST}_2 \circ \text{MAP}(\widehat{Q}) \quad (2)$$

where \widehat{Q} is defined as

$$\widehat{Q} = \langle 1 : \pi_1, 2 : \pi_2, 3 : \{\} \rangle \circ Q \circ \langle 1 : \pi_1, 2 : \pi_3 \rangle \quad (3)$$

$$\begin{aligned} \llbracket \text{const } \tau \ x = E \rrbracket_V &= \langle 1 : \chi_x(\llbracket E \rrbracket_V), 2 : \pi_2, 3 : \pi_3 \rangle \\ \llbracket \text{effect } \tau \ x : f \rrbracket_V &= \langle 1 : \chi_x(\langle 1 : \rho(x), 2 : f \rangle), 2 : \pi_2, 3 : \pi_3 \rangle \\ \llbracket x <- E \rrbracket_V &= \langle 1 : \pi_1, 2 : \pi_2, 3 : \pi_3 \oplus \\ &\quad \langle \langle 1 : \pi_1 \circ \pi_{\text{KEY}}, 2 : \rho(x), 3 : \llbracket E \rrbracket_V \rangle \circ \text{SNG} \rangle \rangle \\ \llbracket R.x <- E \rrbracket &= \langle 1 : \pi_1, 2 : \pi_2, 3 : \pi_3 \oplus \\ &\quad \langle \langle 1 : \llbracket R \rrbracket_V, 2 : \rho(x), 3 : \llbracket E \rrbracket_V \rangle \circ \text{SNG} \rangle \rangle \\ \llbracket \text{if } (E) \ \{B_1\} \ \text{else } \{B_2\} \rrbracket_V &= \\ &\quad \langle 1 : \pi_1, 2 : \pi_2, 3 : \text{SNG} \circ \sigma_{\llbracket E \rrbracket_V} \circ \text{GET} \circ \llbracket B_1 \rrbracket_V \oplus \\ &\quad \text{SNG} \circ \sigma_{\neg \llbracket E \rrbracket_V} \circ \text{GET} \circ \llbracket B_2 \rrbracket_V \rangle \\ \llbracket \text{foreach } (\tau \ x : E) \ \{B\} \rrbracket_V &= \\ &\quad \langle 1 : \pi_1, 2 : \pi_2, \\ &\quad 3 : \langle 1 : \pi_1 \circ \chi_x(\llbracket E \rrbracket_V) \circ \text{PAIRWITH}_x, 2 : \pi_2, 3 : \pi_3 \rangle \\ &\quad \circ \text{FLATMAP}(\llbracket B \rrbracket_V \circ \pi_3) \rangle \end{aligned}$$

Figure 11: Translation for Common Commands

This produces a set of agents and the effects that they have generated (which may or may not be local). In general, we will aggregate aggressively, so each agent will only have one effect for each $\langle k, e \rangle$ pair. For the effect aggregation stage, we must aggregate the effects for each agent and inline them into the agent tuple. If we only have local effect assignments, then this expression is $\mathfrak{Q}(Q) \circ \mathfrak{C}$ where

$$\mathfrak{C} = \text{MAP}(\langle \langle \text{KEY} : \pi_{\text{KEY}}, s_i : \pi_{s_i}, e_j : \pi_2 \circ \sigma_{\pi_e = \pi_{e_j} \circ \pi_1} \circ (\pi_{e_j} \circ \pi_2) \rangle \rangle_{i,j} \rangle) \quad (4)$$

where the s_i are the state fields and the e_j are the value of effect fields. However, in the case where we have non-local effects, we must first redistribute them via the expression

$$\mathfrak{R} = (\pi_1 \times \pi_2) \circ \text{MAP}(\langle 1 : \pi_1, 2 : \text{FLATTEN} \circ \sigma_{\pi_k = \pi_1 \circ \pi_{\text{KEY}}} \rangle) \quad (5)$$

So the entire query phase is $\mathfrak{Q}(Q) \circ \mathfrak{R} \circ \mathfrak{C}$. Finally, for the update phase, each state s_i has an update rule which corresponds to an expression U_{s_i} . These scripts read the output of the expression \mathfrak{C} . Hence the query for our entire simulation is the expression

$$\mathfrak{Q}(Q) \circ \mathfrak{R} \circ \mathfrak{C} \circ \mathfrak{U}(U_{s_0}, \dots, U_{s_n}) \quad (6)$$

where the update phase is defined as

$$\mathfrak{U}(U_{s_0}, \dots, U_{s_n}) = \text{MAP}(\langle \langle \text{KEY} : \pi_{\text{KEY}}, s_0 : U_{s_0}, \dots, s_n : U_{s_n} \rangle \rangle) \quad (7)$$

The only remaining detail in our formal semantics is to define semantics for the query scripts and update scripts. Update scripts are just simple calculations on a tuple, and are straightforward. The nontrivial part concerns the query scripts. A script is just a sequence of statements $S_0; \dots; S_n$ where each statement is a variable declaration, assignment, or control structure (e.g. conditional, foreach-loop). See the BRASIL Language manual for more information on the complete grammar [12]. It suffices to define, for each statement S , a monad algebra expression $\llbracket S \rrbracket$ whose input and output are the triple $\langle 1 : \tau', 2 : \{\tau\}, 3 : \{\rho\} \rangle$; we handle sequences of commands by composing these expressions.

Our query script semantics in BRASIL also depends upon the visibility constraints in the script. In particular, it is possible that a reference to another agent is fine initially, but violates the visibility constraint as that agent moves relative to the one holding the reference. For that reason, BRASIL employs *weak reference semantics* for agent references, similar to weak references in Java. If another agent moves outside of the visible region, then all references to it will resolve to NIL. To formally support the notion of weak references, we represent visibility as a predicate $V(x, y)$ which compares two agents; for any statement S , we let $\llbracket S \rrbracket$ be its interpretation with this constraint and $\llbracket S \rrbracket_V$ be the semantics without.

Before translating statements, we must translate expressions that may appear inside of them. The only nontrivial expressions are ref-

erences; arithmetic expressions or other complex expressions translate to the monad algebra in the natural way. References return either the variable value, or the key for the agent referenced. Ignoring visibility constraints, for any identifier x , we define

$$\llbracket x \rrbracket = \begin{cases} \text{PAIRWITH}_3 \circ \sigma_{\pi_1 \circ \pi_x \circ \pi_1 = \pi_3 \circ \pi_x} & E \text{ is effect} \\ \circ \sigma_{\pi_1 \circ \pi_{\text{KEY}} = \pi_3 \circ \pi_k} \circ \text{GET} & \\ \pi_1 \circ \pi_x & \text{otherwise} \end{cases} \quad (8)$$

In general, for any reference $E.x$, we define

$$\llbracket E.x \rrbracket = \langle 1 : \pi_2 \circ \sigma_{\pi_{\text{KEY}} = \llbracket E \rrbracket} \circ \text{GET}, 2 : \pi_2, 3 : \pi_3 \rangle \circ \llbracket x \rrbracket \quad (9)$$

If we include visibility constraints, $\llbracket E \rrbracket_V$ is defined in much the same way as $\llbracket E \rrbracket$ except when E is an agent reference. In that case,

$$\llbracket x \rrbracket_V = \langle 1 : \text{ID}, 2 : \pi_2 \circ \sigma_{\pi_{\text{KEY}} = \llbracket E \rrbracket} \circ \text{GET} \rangle \circ \langle 1 : V, 2 : \pi_2 \rangle \circ \text{SNG} \circ \sigma_{\pi_1} \circ \text{GET} \circ \pi_2 \circ \pi_k \quad (10)$$

This expression temporarily retrieves the object, tests if it is visible, and returns NIL if not.

To complete our semantics, we introduce the following notation.

- $\chi_a(f)$ is an operation that takes a tuple and extends it with an attribute a having value f . It is definable in the monad algebra, but its exact definition depends on its usage context.
- \oplus is an operation that takes two sets of effects and aggregates those with the same key and effect identifier. It is definable on in the monad algebra, but its exact definition depends on the effect fields in the BRASIL script.
- $\rho(x)$ is the effect identifier for a variable x . In practice, this is the position of the declaration of x in the BRASIL script.

Given these three expressions, Figure 11 illustrates the translation of some of the more popular statements in the monad algebra. In general, variable declarations modify the first element of the input triple (i.e. the active agent), while assignments and control structures modify the last element (i.e. the effects).

As we discussed in Section 4.2, this formalism allows us to apply standard algebraic rewrites from the monad algebra for optimization. For example, many of the operators in Figure 11 – particularly the tuple constructions – are often unnecessary. They are there to preserve the input and output format, in order to facilitate composition. There are rewrite rules that function like dead-code elimination, in that they remove tuples that are not being used. One of the consequences is that many `foreach`-loops simplify to the form

$$F(E, B) = \langle 1 : \text{ID}, 2 : E \rangle \circ \text{PAIRWITH}_2 \circ \text{FLATMAP}(B) \quad (11)$$

Note that this form is “half” of the cartesian product in (1); it joins a single value with a set of values. Thus when we simplify the `foreach`-loop to this form, we can often apply join optimization techniques to the result.

Another advantage of this formalism is that it allows us to prove correctness results. Note that the usage of weak references in BRASIL gives a different semantics for visibility than the one present in Section 3; BRASIL uses visibility to determine how agent references are resolved, while the BRACE runtime uses visibility to determine agent replication and communication. The BRASIL semantics are preferable for a developer, because they are easy to understand and hide MapReduce details. Fortunately, our formalism allows us to establish that these two are equivalent.

THEOREM 1 Let Q be a BRASIL query script whose references are restricted by visibility predicate V . Then

$$\text{NEST}_2 \circ \text{MAP}(\llbracket Q \rrbracket_V) = \sigma_V \circ \text{NEST}_2 \circ \text{MAP}(\widehat{\llbracket Q \rrbracket}) \quad (12)$$

Furthermore, let

$$O(F) = F \circ (\pi_2 \times \pi_3) \circ \sigma_{\pi_1 \circ \pi_{\text{KEY}} = \pi_2 \circ \pi_k} \circ \text{MAP}(\pi_1) \quad (13)$$

be the set of objects affected by an expression F . Then

$$\text{MAP}(\langle 1 : \pi_1, 2 : O(\llbracket Q \rrbracket_V) \rangle) \circ \sigma_V = \text{MAP}(\langle 1 : \pi_1, 2 : O(\widehat{\llbracket Q \rrbracket}) \rangle) \quad (14)$$

The significance of (12) is that, instead of implementing the overhead of checking for weak references, we can filter out the agents that are not visible and eliminate any further visibility checking. The significance of (14) is that weak references insure one agent can only affect other agents falling in its visible region.

B.2 Effect Inversion

As we saw in Section 4.2, there is an advantage to writing a BRASIL script so that all effects assignments are local. It may not always be natural to do so, as the underlying scientific models may be expressed in terms of non-local effects. However, in certain cases, we may be able to automatically rewrite a BRASIL program to only use local effects. In particular, if there are no visibility constraints, then we can always invert effect assignments to make them local-only.

THEOREM 2 Let Q be a query script with no visibility constraints. There is a script Q' with only local effects such that $\llbracket Q \rrbracket = \llbracket Q' \rrbracket$.

PROOF SKETCH. Our proof takes advantage of the fact that effect fields (as opposed to effect variables) may not be read during the query phase, and that effects are aggregated independent of order. We start with Q and create a copy script Q_1 . Within this copy, we remove all syntactically non-local effect assignments (e.g. $E.x \leftarrow v$). Some of these may actually be local in the semantic sense, but this does not effect our proof.

We construct another copy Q_2 . For this copy, we pick a variable a that does not appear in Q . We replace every local state reference x in Q with $a.x$. We also remove all local effect assignments. Finally, we replace each syntactically non-local assignment $E.x \leftarrow v$ with the conditional assignment `if (E == this) {x ← v}`. We then let Q_3 be the script

```
foreach (Agent a : Extent<Agent>) { Q2(a) }
```

That is, Q_3 is the act of an agent running the script for each other agent, searching for effects to itself, and then assigning them locally. The script $Q_1; Q_3$ is our desired script. \square

Note that this conversion comes at the cost of an additional `foreach`-loop, as each agent simulates the actions of all other agents. Thus, this conversion is much more computationally expensive than the original script. However, we can often simplify this to remove the extra loop. As mentioned previously, a `foreach`-loop can often be simplified to the form in (11). In the case of two nested loops over the same set E , the merging of these two loops is a type of self-join. That is,

$$\begin{aligned} F(E, F(E, B)) &= \langle 1 : \text{ID}, 2 : E \rangle \circ \text{PAIRWITH}_2 \circ \\ &\quad \text{FLATMAP}(\langle 1 : \text{ID}, 2 : E \rangle \circ \text{PAIRWITH}_2 \circ \text{FLATMAP}(B)) \\ &= \langle 1 : \text{ID}, 2 : E, 3 : E \rangle \circ \text{PAIRWITH}_2 \circ \\ &\quad \text{FLATMAP}(\text{PAIRWITH}_3 \circ \text{FLATMAP}(B')) \\ &= \langle 1 : \text{ID}, 2 : (E \times E) \rangle \circ \text{PAIRWITH}_2 \circ \text{FLATMAP}(B'') \end{aligned}$$

where B' and B'' are B rewritten to account for the change in tuple positions. As part of this rewrite, one may discover that self-join is redundant in the expression B'' and eliminate it; this is how we get simple effect inversions like the one illustrated in Section 4.2.

In the case of visibility constraints, the situation becomes a little more complex. In order to do the inversion that we did the proof of Theorem 2, we must require that any agent a_1 that assigns effects to another agent a_2 must restrict its visibility to agents visible to a_2 ; that way a_2 can get the same results when it reproduces the actions

of a_1 . This is fairly restrictive, as it suggests that every agent needs to be visible to every other agent.

We can do better by introducing an information flow analysis. We only require that, for each non-local effect assigned to agent, that effect is computed using only information from agents visible to the one being assigned. However, this property depends on the values of the agents, and cannot (in generally) be inferred statically from the script. Thus it is infeasible to exploit this property in general.

However, there is another way to invert scripts in the phase of visibility constraints. Suppose the visibility constraint for a script Q is a distance bound, such as $d(x, y) < R$. If we relax the visibility constraint for the script in the proof of Theorem 2 to $d(x, y) < 2R$, then the proof carries through again. We state this modified result as follows:

THEOREM 3 Let Q be a query script with visibility constraint V . Let V' be such that $V'(x, y)$ if and only if $\exists z V(x, z) \wedge V(z, y)$. Then there is a script Q' with only local effects such that $\llbracket Q \rrbracket_V = \llbracket Q' \rrbracket_{V'}$.

PROOF SKETCH. The proof is similar to that of Theorem 2. The only difference is that we have to ensure that the increased visibility for Q' does not cause the weak references in a script to resolve to agents that would have otherwise evaluated to NIL. In the construction of Q_2 , we use local constants to normalize the expressions so that any agent reference in the original script becomes a local constant. For example, suppose each agent has a field `friend` that is a reference to another agent. If we have a conditional of the form

```
if (friend.x - x < BOUND) { ... }
```

then we normalize this expression as

```
const agent temp = friend;
if (temp.x - x < BOUND) { ... }
```

We then wrap these introduced constants with conditionals that test for visibility with respect to the old constraints. For example, the code above would become

```
const agent temp = (visible(this, friend) ?
                    friend : null);
if (temp.x - x < BOUND) { ... }
```

where `visible` is a method evaluating the visibility constraint and evaluates to NIL in the monad algebra. Given the semantics of NIL, this translation has the desired result. \square

C. DETAILS OF SIMULATION MODELS

This section describes the simulation models we have implemented for BRACE single-node performance and scalability experiments.

Traffic Simulation. Traffic simulation is required to provide accurate, high-resolution, and realistic transportation activity for the design and management of transportation systems. MITSIM, a state-of-the-art single-node behavioral traffic simulator has several different models covering different aspects of driver behavior [54]. For example, during each time step, a lane selection model will make the driver inspect the lead and rear vehicles as well as the average velocity of the vehicles in her current, left, and right lanes (within lookahead distance parameter ρ) to compute the utility function for each lane. A probabilistic decision of lane selection is then made according to the lane utility. If the driver decides to change her lane, she needs to inspect the gaps from herself to the lead and rear vehicles in the target lane to decide if it is safe to change to the target lane in the next time step. Otherwise, the vehicle following model is used to adapt her velocity based on the lead vehicle. The newly computed velocity will replace the old velocity in the next time step. Note that if the driver cannot find a lead

Lane	Change Frequency	Avg. Density	Avg. Velocity
1	8.93%	7.42%	0.007%
2	5.57%	10.38%	0.007%
3	7.67%	9.38%	0.007%
4	21.37%	19.72%	0.007%

Table 2: RMSPE for Traffic Simulation (LookAhead = 200)

or rear vehicle within ρ , she will just assume the distance to the lead or rear vehicle is infinite, and adjust the velocity according to a free-flow submodel.

One note is that since the MITSIM implementation hand-coded nearest neighbor indexing for accessing the lead and rear vehicles for performance reasons, its lookahead distance actually varies for each vehicle. In our reimplementation we fix the lookahead distance to 200 and apply single-node spatial indexing. In order to make sure this implementation difference does not generate drastically different aggregate driving behavior, we validate consistency of the MITSIM model encoded in BRASIL in terms of the simulated traffic conditions. We compare lane changing frequencies, average lane velocity and average lane density with the segment length 20,000 on both simulators. The statistical difference is measured by RMSPE (Relative Mean Square Percentage Error), which is often used as a goodness-of-fit measure in the traffic simulation literature [10]. The results for all these three statistics are shown in Table 2. We can see that except for Lane 4’s average density and changing frequency, all the other statistics demonstrate strong agreement between the two simulators. This exception is due to the fact that in the MITSIM lane changing model drivers have a reluctance factor to change to the right most lane (i.e., Lane 4). As a result there are only a few vehicles on that lane (56.33 vehicles on average compared to 351.42 on other lanes), and small lane changing record deviations due to the fixed lookahead distance approximation can contribute significantly to the error measurement.

Fish School Simulation. Couzin et al. have built a behavioral fish school simulation model to study information transfer in groups of fish when group members cannot recognize which companions are informed individuals who know about a food source [13]. This computational model proceeds in time steps, i.e., at each time period each fish inspects its environment to decide on the direction which it will take during the next time period. Two basic behaviors of a single fish are avoidance and attraction. Avoidance has the higher priority: Whenever a fish is too close to others (i.e., distance less than a parameter α), it tries to turn away from them. If there is no other fish within distance α , then the fish will be attracted to other fish within distance $\rho > \alpha$. The influence will be summed and normalized with respect to the current fish. Therefore, any other individuals out of the visibility range ρ of the current individual will not influence its movement decision. In addition, informed individuals have a preferred direction, e.g., the direction to the food source or the direction of migration. These individuals will balance the strength of their social interactions (attraction and avoidance) with their preferred direction according to a weight parameter ω .

Predator Simulation. Since both the traffic and the fish school simulations only use local effect assignments, we designed a new predator simulation, inspired by simulations of artificial societies [30]. In this simulation, a fish can “spawn” new fish and “bite” other fish, possibly killing them, so density naturally approaches an equilibrium value at which births and deaths are balanced. Since effect inversion is not yet implemented in the BRASIL Compiler, we program biting behavior either as a non-local effect assignment (fish assign “hurt” effects to others) or as a local one (fish collect “hurt” effects from others) in otherwise identical BRASIL scripts.