

# Time-parallel simulation of wireless ad hoc networks

Guoqiang Wang · Damla Turgut · Ladislau Bölöni ·  
Dan C. Marinescu

© Springer Science+Business Media, LLC 2007

**Abstract** In this paper, we study time-parallel simulation of wireless networks based upon the concept of the perturbation induced by a networking event and present a layer-by-layer analysis of the impact of perturbations on the wireless network. This analysis allows us to propose several methods to improve the accuracy of time-parallel simulation. We describe an implementation based on the widely used ns-2 simulator and on the iterative extension of the warmup period. We introduce a method for initial state approximation which can improve the accuracy of the simulation for table-driven ad hoc routing protocols. A series of experiments show that on typical scenarios time-parallel simulation leads to a significant speedup while maintaining a high level of accuracy.

## 1 Introduction

The development of communication systems requires rigorous performance studies. Analytical performance studies can only be carried for simple models of complex systems; such studies tend to provide a basic understanding of the

system behavior and qualitative results. Simulation, on the other hand, is often used for quantitative performance study of more realistic models of complex systems. Wireless ad hoc networks are rarely amenable to analytical performance analysis due to the complexity of the models involved. Simulation has emerged as an important tool for the study of wireless ad hoc networks, but it is commonly used to study networks with a few hundred nodes for a relatively short period of time.

Nowadays systems consisting of a few thousand nodes are rather common. Indeed, let us consider a simple scenario: there are 120 students in a classroom; each student has a cell phone (GSM source), a PDA, and laptop (two 802.11b WiFi sources). There are five Bluetooth sources: PDA, laptop, cell phone, headset, and mouse. Some of the students might have WiFi enabled cameras, Bluetooth enabled audio players, and matching head phones. All in all, it does not seem out of the ordinary to have three WiFi and seven Bluetooth sources per person. Thus, even without considering that many of the WiFi nodes have a transmission range long enough to cover neighboring classrooms as well, in order to study the networking environment of one classroom we have to simulate a system with 1,200 wireless sources operating in the same frequency band.

A wireless network with 1,200 nodes pushes the limits of serial simulators such as ns-2. Even when feasible, such simulations require a significant amount of computing power and can take a very long time. An alternative is to resort to space-parallel or time-parallel simulation. In many instances, e.g., in the classroom environment, every transmission can be received by, or at least interfere with the transmissions of many if not all the other nodes; the fully connected dependency defeats the purpose of spatial partitioning.

---

G. Wang · D. Turgut (✉) · L. Bölöni · D. C. Marinescu  
School of Electrical Engineering and Computer Science,  
University of Central Florida, P.O. Box. 162450, Orlando,  
FL 32816-2450, USA  
e-mail: turgut@eecs.ucf.edu

G. Wang  
e-mail: gwang@eecs.ucf.edu

L. Bölöni  
e-mail: lboloni@eecs.ucf.edu

D. C. Marinescu  
e-mail: dcm@eecs.ucf.edu

In this paper, we present an approach for time-parallel simulation of wireless ad hoc networks. Time-parallel simulation has been used for some time; its applicability is strongly dependent on the simulated phenomena and can yield exact results only for systems whose behavior is modeled by regenerative stochastic processes. However, as we will show, good approximations can be obtained quickly, and the quality of the approximation can be improved through iteration.

Our time-parallel simulation of wireless ad hoc networks uses a traditional simulation package coupled with new techniques for the practical implementation of the simulation. We develop a methodology to analyze the accuracy of the simulation and introduce the notion of perturbation of measurements; then we present a layer-by-layer analysis of the impact of the perturbations on the performance of the network. This approach allows us to predict the accuracy of the simulation results through an analysis of the protocols involved at each layer.

This paper is organized as follows. We survey related work in Sect. 2, then we present the basic concepts for time-parallel simulation of wireless networks in Sect. 3. We also introduce a model of the propagation of perturbations in wireless networks and the impact of the protocols at the various layers of the networking stack. We apply these considerations to propose several methods for improving the speed and accuracy of time-parallel simulation. Section 4 presents the implementation of the theoretical models developed in the previous section with the widely used ns-2 simulator. We find that some of the theoretical concepts can be immediately incorporated in simulation studies conducted with existing simulators such as ns-2 while others are very difficult due to the particularities of the implementation. A series of experiments investigating the speedup and precision of the proposed method for typical wireless network simulation scenarios are presented in Sect. 5. We summarize our results in Sect. 6.

## 2 Related work

Parallel discrete event simulation (PDES) reduces the overall execution time by parallel execution of the simulation on multiple processors. There are two main avenues for parallel simulation: space-parallel simulation (distributed simulation), and time-parallel simulation. In the space-parallel simulation approach [17, 27], the simulation model is decomposed into a number of components on a spatial basis. Each component is modeled by a *logical processor*. Logical processors establish a communication mechanism among each other to avoid or fix possible causality errors. There are two general mechanisms to avoid/correct causality errors: optimistic mechanisms and

conservative mechanisms. With optimistic mechanisms [22], a processor can execute an event  $e$  without the knowledge of its prior events, and state recovery methods [34, 35, 37] are required to restore the state of the simulation once causality errors are detected. Instances of optimistic space-parallel simulations include [9, 11, 12, 16, 35, 43]. In conservative simulations, a processor does not execute an event  $e$  until all events that may affect event  $e$  are executed. Instances of conservative space-parallel simulations include [19, 25, 28, 45].

Load balancing in PDES refers to distributing the workload over different processors evenly. An efficient load balancing scheme can greatly improve the speedup of PDES, due to the fact that the overall progress of the simulation is decided by the progress of the slowest processor. Various load balancing approaches for space-parallel simulations can be found in [2, 4, 8, 13–15, 18, 36, 38, 44].

The parallel/distributed network simulator (PDNS) [45] project uses a space-parallel simulation approach based on the ns-2 network simulator [39]. However, the applicability of PDNS is limited to *wired* networks, and the traffic simulated at different spatial partitions cannot affect each other.

The SWiMNet parallel simulator [5–7], is used for the simulation of personal communication services (PCS) networks with fixed channel assignment by specifying fine grained mobility, variable call process, and arbitrary coverage area. It is based on a combination of optimistic and conservative paradigms and makes use of the event pre-computation by the model independence within the PCS model.

WiPPET [31], an optimistic parallel simulator for evaluating the performance of wireless protocols, exploits the parallelism of multi-channel radio networks either by geographic locations of the resources or by different radio channels.

Table 1 presents a comparison of PDNS, SWiMNet, WiPPET and the time parallel simulation approach proposed in this paper.

In the time-parallel simulation approach [1, 20, 21, 26, 40, 41], the long simulation interval is partitioned into smaller adjacent simulation intervals, and each simulation interval is assigned to a processor with a *guessed* initial state. The simulation terminates when the final state of each interval matches the initial state of its successive interval. Thus, *state matching* is one of the key problems of time-parallel simulation. In [26], the authors propose a time-parallel simulation algorithm based on state matching. A simulation is defined as *partial regenerative* if there exists a subset of the system state variables such that the subsystem represented by the subset can repeat its state infinitely many times. The system is then partitioned at the

**Table 1** A comparison of PDNS, SWiMNet, WiPPET and the time-parallel simulation approach (TPS) proposed in this paper

	PDNS	SWiMNet	WiPPET	TPS
Application domain	Wired networks	PCS networks	Multi-channel radio networks	Wireless ad hoc networks
Parallelism	Space-parallel	Space-parallel	Space-parallel	Time-parallel
Error control scheme	Conservative	Conservative/optimistic	Optimistic	State matching
Main issues	The interference between different logic processors must be limited. The optimistic approach requires a large amount of memory			Efficient estimation of the initial state of the simulation segments

*regeneration points* which mark a *regenerative sub-state*. In some cases the *regeneration points* of a regenerative simulation can be found without performing a detailed simulation; the state matching problem can be solved by performing a pre-computation [29, 30]. Wang and Abrams propose a *pre-simulation* to identify regenerative points based upon Markovian modeling [41].

Kiesling [24] mentioned that the widespread use of time parallel simulation is restricted by the state-match problem, due to the difficulty in identifying regeneration points, especially for models with complex states. Unfortunately, time-parallel simulation of wireless networks belongs to this category. Kiesling pointed out that the use of approximate solutions can facilitate the temporal decomposition of simulation models. Since errors are introduced, an error control method must be provided for approximate time-parallel simulation. The simulation starts with guessed incorrect initial states, and fix-up computations are re-executed to reduce the error. Although time-parallel simulation rarely allows us to obtain accurate results, approximate results [21, 23, 24, 40] can be produced efficiently. The initial version of the time-parallel simulation of wireless ad hoc networks are presented in [3].

### 3 Time-parallel simulation of ad hoc networks

#### 3.1 The one-step time-parallel model

A simulation  $S$  of an ad hoc network is specified by a quadruple  $(\mathcal{E}, A, \tau, \mathcal{I})$  where  $\mathcal{E} = \{e_1, e_2, \dots\}$  is a set of *planned events*,  $A$  is the geographic area,  $\tau = [\tau^s, \tau^e]$  is a time interval and  $\mathcal{I}$  is the initial state of the network at time  $\tau^s$ . The output of the simulation is the simulation trace  $\mathcal{T} = \{t_1, t_2, \dots\}$  and the final state  $\mathcal{F}$  at time  $\tau^e$ . The simulation trace is a series of events which includes the planned events ( $\mathcal{E} \subset \mathcal{T}$ ) as well as events which are consequences of the planned events.

The idea behind time-parallel simulation is to replace the simulation  $S$  with a number of smaller *simulation segments*, which operate on the full area  $A$ , but on subsets of the time interval  $\tau$ . First, we describe the most straightforward way for implementing time-parallel

simulation through disjoint time intervals of equal length. In the following sections, we consider modifications to this basic model in order to achieve higher accuracy and/or speedup. In this basic model, we partition the time interval into  $m$  disjoint intervals of length  $\tau_d = \frac{\tau^e - \tau^s}{m}$ :

$$\tau_i = [\tau_i^s, \tau_i^e] = [(i-1) \cdot \tau_d, i \cdot \tau_d], i \in \{1, \dots, m\} \quad (1)$$

This also implies the partitioning of the planned event set:

$$\mathcal{E}_i = \{e_i | e_i \in \mathcal{E}, \tau_i^s \leq \text{time}(e_i) < \tau_i^e\} \quad (2)$$

The segments of the time-parallel simulation are specified by the quadruples  $(\mathcal{E}_i, A, \tau_i, \mathcal{I}_i)$ . As the segments are  $m$  times shorter than in the original simulation, their execution will take roughly  $m$  times shorter time, and they can be executed in parallel on  $m$  independent processors. The simulation trace of the full simulation will be the union of the simulation traces  $\mathcal{T} = \mathcal{T}_1 \cup \dots \cup \mathcal{T}_m$ , while the final state of the network will be  $\mathcal{F} = \mathcal{F}_m$ .

The main difficulty pertains the calculation of the initial state  $\mathcal{I}_i$ . For the first segment, we have  $\mathcal{I}_1 = \mathcal{I}$ . For the other segments, the correct initial state is the final state of the previous segment  $\mathcal{I}_i = \mathcal{F}_{i-1}$ . However, using the final state of the previous simulation segment would require the in-order simulation of the segments, preventing any parallelism.

One crude approximation would be to assume that every simulation segment starts with the original state of the network  $\mathcal{I}$ . In this *one-step time-parallel simulation* the simulation segments are specified by  $(\mathcal{E}_i, A, \tau_i, \mathcal{I})$ .

Experiments show that the accuracy of the one-step time-parallel simulation is not sufficient for the practical purposes of wireless ad hoc network simulation. In the following, we analyze the source of errors in the time-parallel simulation of wireless ad hoc networks, estimate the magnitude of these errors, and propose methods to improve the accuracy of the simulation.

#### 3.2 Measurements on the simulation trace

A simulation trace includes *planned events*, ( $\mathcal{E} \subset \mathcal{T}$ ), as well as *new events* triggered by the planned events and

calculated by the simulation process. For example the event corresponding to sending a packet may trigger a new event indication a collision with another transmission. Every event  $e_i$  has a set of *properties*, e.g., the time of the event time( $e_i$ ), the type of the event, the packet size, etc. Such information allows us to determine important measures of performance such as packet delay, packet loss ratio, throughput, and so on.

The purpose of running a simulation is to collect the information associated with the events in the trace. However, we are rarely interested in individual events, we are concerned with observations, or measurements, spanning longer time intervals, or the entire trace.

Measurements can be:

- (a) *instantaneous*  $M^i(t_c)$ , reflecting the situation at the current time  $t_c$ .
- (b) *cumulative*  $M^c(t_c)$ , reflecting the evolution of the measurement from the beginning of the simulation to the current time  $t_c$ . For every instantaneous measurement  $M^i$ , there is an associated cumulative measurement  $M^c$ :

$$M^c(t_c) = \int_0^{t_c} M^i(t) dt \quad (3)$$

- (c) *average*  $M^a(t_c)$ . For every instantaneous measurement  $M^i$ , there is an associated average measurement  $M^a$  defined by:

$$M^a(t_c) = \frac{\int_0^{t_c} M^i(t) dt}{t_c} \quad (4)$$

An example of instantaneous measurement is the *current transmission rate*, with the associated cumulative measurement *total transmitted data*, and the associated average measurement *average transmission rate*.

Many average measurements are expressed as *ratios* of measured quantities. For some of these measurements, we find it useful to define the complementary measure  $C(M^a) = 1 - M^a$ . For instance, the *average packet loss ratio* has as complement the *average packet delivery ratio*.

### 3.3 Perturbations on a measurement

In the following, we investigate how an event *perturbs* the measurements of a simulation of wireless ad hoc networks. Our model accurately captures the underlying phenomena and provides a framework in which observations can be analyzed and predictions can be made.

We call an *approximate simulation for the measurement*  $M$ , related to simulation  $S$ , a simulation  $S'$  which generates a trace  $T'$  such that  $M(T') \approx M(T)$ .

The *absolute error of simulation* for measurement  $M$  is

$$\delta_M = |M(T) - M(T')| \quad (5)$$

The *relative error of simulation* for measurement  $M$  is

$$\epsilon_M = \frac{|M(T) - M(T')|}{|M(T)|} \quad (6)$$

We define the relative accuracy for measurement  $M$  as  $1 - \epsilon_M$ . In most cases, the relative accuracy, expressed in percentages, is the most intuitive measure of the quality of an approximate simulation. For instance, a statement such as “the simulation has an accuracy of 98%” can conveniently express the fact that the relative error is smaller than 2%. The relative accuracy is a non-dimensional quantity, allowing us to compare the accuracy of different measurements.

However, the relative error suffers from the phenomena of error amplification for measurements which represent the average ratios for rare events. Let us consider a simulation scenario where out of 1,000 packets sent two are lost; then we say that the packet loss ratio is  $M_{\text{PLR}} = 0.002$ . If we have an approximate simulation which finds only one lost packet, the relative error will be:

$$\epsilon_{\text{PLR}} = \frac{|M'_{\text{PLR}} - M_{\text{PLR}}|}{|M_{\text{PLR}}|} = \frac{|0.002 - 0.001|}{|0.002|} = 50\% \quad (7)$$

However, repeating the same calculation for the complementary measure packet delivery ratio  $M_{\text{PDR}} = C(M_{\text{PLR}}) = 1 - M_{\text{PLR}}$ :

$$\epsilon_{\text{PDR}} = \frac{|M'_{\text{PDR}} - M_{\text{PDR}}|}{|M_{\text{PDR}}|} = \frac{|0.998 - 0.999|}{|0.998|} \approx 0.1\% \quad (8)$$

Thus, the simulation appears very inaccurate for one measurement and very accurate for the complementary measurement. The physical phenomena behind this phenomenon is that the appearance of a rare event, for instance of a packet loss in a lightly loaded network is dependent on the coincidence of a number of factors. An approximate simulation might approximate well the probability of occurrence of those factors, but not the exact location and time where the event will take place. Therefore, the absolute error might be a better indicator of the accuracy of simulation for the measurements on rare events. The absolute error has the additional advantage that the absolute error of the measurement and the complementary measurement is identical

$$\begin{aligned} \delta_M &= |M(T) - M(T')| = |(1 - M(T)) - (1 - M(T'))| \\ &= \delta_{C(M)} \end{aligned} \quad (9)$$

Now we introduce the notion of a *perturbing event*  $e_p$ . We consider a simulation with the planned events  $\mathcal{E} = \{e_1 \dots e_n\}$  and a perturbed set  $\mathcal{E}' = \mathcal{E} \cup e_p$ . The perturbed set of planned events will lead to a simulation trace  $\mathcal{T}'$  and, obviously, to perturbed measurements  $M'(t)$ . We are interested in the size and the temporal extent of the perturbations  $\Delta M(t) = M'(t) - M(t)$ . We say that

- The perturbation has *no effect* on measurement  $M$  if  $\Delta M(t) = 0, \forall t$ .
- The event  $e_p$  creates a *time limited perturbation* in the measurement  $M$  which lasts until the *extinction time*  $t_e$  if  $\Delta M(t) = 0, \forall t > t_e$ .
- The event  $e_p$  has a *shift effect perturbation* on the measurement  $M$  if  $\lim_{t \rightarrow \infty} (\Delta M(t)) = c$ , with  $c$  being the *shift constant*. If there is a time point  $t_s$  with the property that  $\Delta M(t) = c, \forall t > t_s$ , we call  $t_s$  the *stabilization point*.

**Property 1** If an event  $e_p$  causes a time-limited perturbation on the instantaneous measurement  $M^i$  with extinguishing time  $t_e$ , it creates a shift effect perturbation on the corresponding cumulative measurement  $M^c$ , with the stabilization point  $t_s = t_e$ . The shift constant can be expressed as:

$$c = \int_{t_p}^{t_e} \Delta M^i(t) dt \tag{10}$$

- We say that an event causes a *destabilizing perturbation* of measurement  $M$  if  $\exists t_s > t_p$  such that  $\Delta M = c, \forall t > t_s$ .

### 3.4 Propagation of perturbations in wireless ad hoc networks

As a first approximation, the events in a wireless ad hoc networks are caused by individual transmissions of packets. We consider the perturbing event to be the insertion of a new packet in the network. The removal of the packet is not considered separately, because it is equivalent to simply reversing the perturbed and the original system. As we are concerned about the absolute values of the difference between the measurements of these systems, the reversed system will yield the same conclusions. An additional type of perturbation we consider is the perturbation of the initial state, i.e., the simulation starts with a different initial state than expected. We discuss this type of perturbation separately.

For events representing the addition or removal of a single packet, the immediate affect of the perturbations is

usually minimal. However, the networking protocols deployed at various levels of the networking stack can amplify or (in some cases) reduce the impact of the perturbations. In the following, we investigate the impact of the protocols deployed at the various layers of the networking stack on the perturbation produced by the insertion of a single packet.

#### 3.4.1 Physical layer

An inserted packet will perturb the physical layer measurements only for the duration of the packet transmission. The time it takes to transmit a packet can be calculated as follows. Assume that we send a packet of 1,536 bit, the maximum length supported by the 802.11b protocol. The transmission rate of 802.11b is 1.375 Mbps. The time required to send this packet is composed of:

- Distributed inter-frame space (DIFS), set to 50  $\mu s$
- Data packet transmission: 192  $\mu s$  for the preamble + 1,536/1.375 Mbps = 192  $\mu s$  + 1,118  $\mu s$
- Small inter-frame space (SIFS), set to 10  $\mu s$
- 802.11 ACK packet: 192  $\mu s$  + 14/1.375 Mbps = 203  $\mu s$

Thus, the total time becomes 2084  $\mu s$ , or approximately, 2 ms. This might change slightly for different protocols, but the order of magnitude remains the same. We conclude that the perturbation in the physical layer due to a new packet is time-limited, with a very short (2 ms) extinction time.

#### 3.4.2 MAC layer

The perturbation caused by an inserted packet in the MAC layer depends on the influence of the load, and the type of the packet. The influence of a control packet such as ACK, RTS, or CTS, its influence extends beyond the time frame covered. For instance, receiving an RTS packet in a carrier sense multiple access with collision detection (CSMA-CA) network forces the node to refrain from transmitting for the duration specified in the packet. The size of this interval can be as long as the maximum packet transmission time, on the order of magnitude of 2  $\mu s$ .

Delaying the sending of a packet can in its turn delay the sending of the response or follow-up packets, creating a ripple effect. Normally, this will appear as a time-limited perturbation. For a time division multiple access (TDMA) type protocol, for a channel with the capacity of  $n$  bps, with a load of  $v \leq 1$ , and a packet size of  $l$ , the extinction time can be estimated as:



$$t_e = \frac{l}{n \cdot (1 - v)}$$

Note that for a full channel,  $v = 1$ , the perturbation becomes a shift effect perturbation.

For CSMA-CD protocols, the perturbation will extend to the length of the contention window. 802.11b uses an exponential back-off algorithm where the contention window can vary between CVmin and CVmax. Typical values of CVmin are between 7–15 and of CVmax are between 7–255 with the numbers being multiples of a slot time which is 20  $\mu$ s in 802.11b. Thus, the influence of initial collision can extend to 5,000  $\mu$ s. However, if the channel is very busy, the initial collision can lead to further collisions down the line. Furthermore, any collision is extending the collision window through exponential back-off, which will be only gradually reduced.

In conclusion, on highly loaded networks, the insertion of a packet creating a collision can cause perturbations of up to several seconds in the worst case.

### 3.4.3 Routing layer

Perturbations triggered by packet insertion at the routing layer are heavily dependent on the nature of the protocol and the inserted packet. The critical question is whether the packet will change the routing of future packets or not. If the packet is part of an established flow of application layer data, it will most likely not affect the routing of the other packets. If the packet is the first packet of a new flow, and the routing protocol is a reactive one, the packet will establish a new route. This frequently requires broadcast of routing information, which in its turn triggers the transmission of additional packets. Although this appears to be a destabilizing perturbation, the extent of this flooding is carefully constrained by the routing protocol and it will last at most several seconds.

Proactive routing algorithms deploy a routing table periodically updated by routing packets. The insertion of a routing packet triggers a perturbation in the network by changing the routing table of a node, and this will affect the routing of future packets. This is a major perturbation, affecting a large number of packets and a time interval on the order of minutes. A routing table perturbation is extinct when the routing tables of the original and the perturbed system re-converge. This situation occurs when:

- The modified routes are superseded by new, independently discovered routes, in both the original and the perturbed system (for instance, as a result of node mobility).

- The original system acquires the same routes as the perturbed one.
- The modified routes expire through a timeout and the routing table returns to the unperturbed version.
- An external command or a predetermined timeout flushes partially or completely the routing tables, forcing the recomputing of all routes.

Although it is technically possible to imagine a routing protocol where the loss or addition of a single routing packet would change the routes indefinitely, virtually all protocol designers, in their quest to make the protocols more reliable, have adapted features which make the routing tables converge; this has the indirect effect of limiting the perturbations and improving the accuracy of time-parallel simulations.

### 3.4.4 Transport layer and application layer

Finally, we investigate how a perturbing event affects the transport layer and the application layer protocols. The major difference here is between reliable and non-reliable protocols. Let us consider the case of the most frequently used reliable transport protocol, TCP. The loss of a single TCP packet can significantly perturb the subsequent TCP flow: the packet will be retransmitted, the transmission window reset to its minimal value, which will then extend through the slow start algorithm. Thus, the loss of a single packet can exert an influence over the network for several tens of seconds.

In the case of the UDP protocol, which does not implement reliable transmission, the perturbation is minimal or nonexistent at the transport layer. However, applications which deploy UDP at the transport layer frequently use application layer protocols to control the flow of data. For instance, the multimedia streaming application RealPlayer is using an application level byte stream protocol (real-time transport protocol (RTP), specified in RFC 1889) for the transfer of multimedia information, with UDP being the transport protocol. Packet losses are handled by a complex logic and actions involving the real time streaming protocol (RTSP), real time control protocol (RTCP), session description protocol (SDP) and, of course RTP.

We conclude that perturbations have effects up to the application layer. However, at the application layer, the behavior of the system becomes very complex, and in some cases, small perturbations can have major effects on the state of the system. This problem is not specific to time-parallel simulation. Most simulation studies of wireless ad hoc networks consider scenarios with constant or variable bit rate generic UDP sources, without application level protocols deployed in the simulator. When application

level flow is simulated, it is done by an artificially generated stream which replicates the properties of a packet stream generated by the application level protocol, without actually deploying the application in the simulator [42]. In these types of models perturbations propagate only over very short time frames, thus creating an environment favorable for time-parallel simulation.

Until now, we have considered only perturbations caused by a single event. In the following, we discuss the perturbations caused by the interaction among two or more independent events:

*Dual events.* Dual events cancel each other; their successive occurrence does not change the state of the system. For instance, a reboot sequence is represented by two events: a node failure, followed after a short time by the node recovery. Individually, each event causes a perturbation, however their effects cancel each other. The state of the system will be the same as if none of the events occurred. This favorable case rarely occurs in practice.

*Idempotent events.* Idempotent events lead to the same state regardless if only one, a subset, or all of them occur. For instance, two consecutive missed routing table updates have the same effect, an incorrect routing table entry, which requires the need to run a path discovery process. Idempotent events appear whenever some component of the system has a special “correct” state, while the other, “incorrect” states are functionally identical. Such system components are relatively frequent; examples are routing table entries or established end-to-end connections.

*Independent events.* Independent events lead to the same state of the system regardless of the order in which they occur; the overall perturbation of state caused by the events is additive. Their compound effect can be studied considering the perturbations caused by the individual events independently.

*Mutually amplifying events.* Such events lead to drastic alteration of state; the perturbation caused by a sequence of such events is significantly larger than the perturbations caused by the individual events occurring independently. For example, an event causes a routing table to be overwritten with incorrect values and the second event propagates the incorrect routing table to all the nodes in the network, causing the network failure.

Naturally, mutually amplifying events are the worst case scenario from the point of view of the accuracy of the time parallel simulation. We need to carefully consider whether they can occur in the considered scenarios. Note that in our hypothetical example, the mutually amplifying events caused a total interruption of the communication in the simulation. Practically deployed networking stacks contain build-in guards against such types of unstable behavior. Mutually amplifying events are clearly possible during the experimental testing of new routing or MAC protocols. The

observation of such an unstable behavior is an important feedback to the developers; at the same time, they need to be aware that the errors in the time-parallel simulation of such unstable systems are significantly higher.

We conclude that inter-event interactions yield perturbations which are identical or smaller than the perturbations considered individually, for most practically deployed, stable protocols. Mutually amplifying events, however, are possible if the deployed protocol stack shows unstable behavior (most frequently, as a result of experimental code). In these cases the accuracy of the simulation will be much lower.

### 3.5 Techniques for improving the accuracy of time-parallel simulations

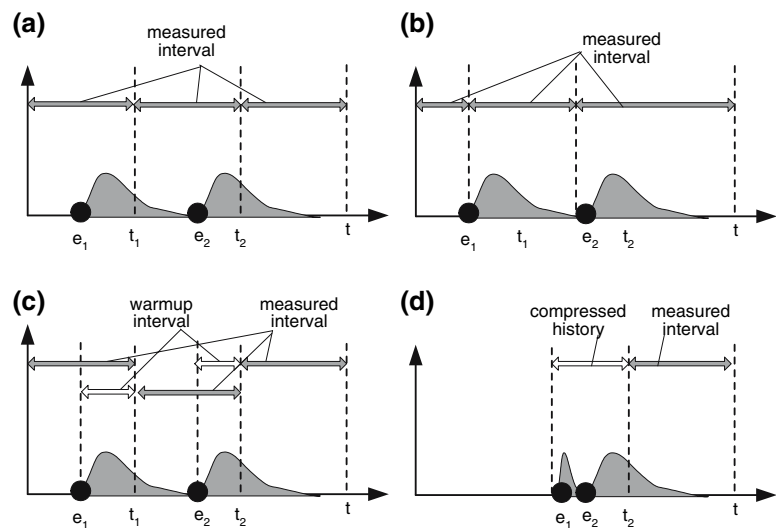
Let us now return to the one-step time-parallel simulation described in Sect. 3.1, and consider the sources of errors in the approximation. For instance, in Fig. 1(a) we see a partitioning of the simulation interval  $[0, t]$  into three equal time intervals:  $[0, t_1 = t/3]$ ,  $[t_1 = t/3, t_2 = 2t/3]$  and  $[t_2 = 2t/3, t]$ . The event  $e_1$ , in the first segment is creating a perturbation which is not extinguished until the middle of the second segment. As the two segments are simulated independently, the processor simulating the second segment does not know about the perturbation which leads to an inaccurate simulation. Let us now review several techniques through which this inaccuracy can be reduced or eliminated.

#### 3.5.1 Time interval shift

One way to increase the accuracy of the simulation is to select the boundaries of the simulation intervals such that the perturbations are completely contained in the segments, as shown in Fig. 1(b). These points are the equivalent of the regenerative states of stochastic processes. If we can not find states where all the perturbations are extinguished, we can search for points where there are a comparatively smaller number of ongoing perturbations (partial regenerative states).

There are several problems with this approach. First, there might not be any regenerative states in the simulation, or their distribution might be such that it does not lead to segments of size appropriate for parallelization. The most difficult problem, however, is finding regenerative or partially regenerative states in a simulation without first running the simulation itself. There are certain circumstances when the identification of such points is possible. If a routing protocol periodically flushes the complete set of routing information, that instance corresponds to a

**Fig. 1** An illustration of several techniques for improving the accuracy of time-parallel simulations. (a) One-step time-parallel simulation with segments of equal length. (b) Time interval shift. (c) Warmup intervals. (d) Warmup with history compression



regenerative state. Similarly, long periods of silence can be used as regenerative states, and they can be identified with an initial analysis of the scenario.

### 3.5.2 Warmup interval

This technique relies on separating the timespan of the simulation segment into the *warmup interval* followed by the *measurement interval*. During the warmup interval, we perform the simulation but do not record any measurements. Thus, the simulation performed during the measurement interval will be more accurate, because it will consider not only the events occurring during the measurement interval, but also the perturbations caused by events which occurred during the warmup.

The ideal size of the warmup interval is the shortest period which contains all the perturbing events which generate perturbations extending into the measurement interval (see Fig. 1(c)). The size of the required warmup can vary between various segments, the first segment does not require a warmup. Unfortunately, for practical cases we can not accurately compute the ideal length of the warmup interval without first running the simulation.

### 3.5.3 Warmup with compressed history

In a practical run of time-parallel simulation with warmup, the size of the warmup period can be significantly longer than the measured interval; thus most of the computation time is spent into the simulation of the warmup interval, which does not contribute to the measurements. Traditionally, the warmup interval is the exact copy of the simulation scenario for a period before the measured

interval. We can replace the warmup interval with a shorter and/or simpler simulation interval which, however, would yield the same results. For instance, we can remove all the events from the warmup period which do not produce perturbations at the measured point. For the events which produce perturbations, we might be able to replace them with events easier and faster to simulate. We call this modified warmup interval a *compressed history* (see Fig. 1(d)).

### 3.5.4 Initial state approximation (ISA)

In this technique, we are computing an approximation of the initial state of the segments without previous execution of the simulation. We had seen that for a typical simulation of a wireless ad hoc network the perturbation with the largest impact on the measurements relates to changes in the routing table. As most simulations start with an empty routing table which will be filled in during the initial phase of the simulation, this is a major source of errors. Approximating the routing table at the beginning of each simulated segment can thus significantly increase the accuracy of the simulation.

### 3.5.5 Composite methods

The methods of improving the precision of the time-parallel simulation can be used in conjunction with each other. The warmup can be composed of two periods: one of compressed history (for instance, covering the complete timeframe from the start of the simulation), followed by a warmup period operating on the unmodified planned event list. The initial state approximation method can be used to



approximate the state at the beginning of the warmup period. Time period shift can be deployed in conjecture with any of the methods.

#### 4 Implementation of time-parallel simulation for ad hoc networks

In the following, we present our implementation of the time-parallel simulation using the ns-2 simulator [39]. The first method, based on the iteratively extended warmup of the simulation segments, can be deployed without changes to the ns-2 implementation of the protocols. An alternative method combines iteratively extended warmup and initial state approximation and produces better results; this method requires some knowledge of the implementation of the routing protocol and additional code to allow the initialization of the routing tables.

Let us now highlight the relevant features of the ns-2 simulator. We defined a simulation as the transformation the quadruplet  $(\mathcal{E}, A, \tau, \mathcal{I})$  into the output consisting of the final state  $\mathcal{F}$  and the simulation trace  $\mathcal{T}$ . While ns-2 is compatible to this definition, there are differences in our ability to extract, interpret and modify the different components.

The ns-2 trace files,  $\mathcal{E}$  and  $\mathcal{F}$  are very easy to create, parse and filter, they consist of lines of text, one line per event. Once the simulation passes a given timeline, the output is immediately available, thus data can be extracted while the simulation is in progress.

In contrast, it is relatively difficult to generate and extract the initial and final state of the simulation. The data associated with the state is distributed across the protocol implementations; there is no standard format and no standard API through which this information can be accessed. While in principle there is no obstacle to access this information, this task needs to be handled from case to case, and it requires extensive knowledge of the implementation details of the protocols involved.

##### 4.1 Time-parallel simulation with iteratively extended warmup

Our first method for implementing time-parallel simulation relies on iterative extension of the warmup interval. The approach allows us to obtain simulation results with increasingly higher accuracy *during the runtime of the simulation*. The simulation can be stopped when the desired accuracy is reached. This approach does not require the manipulation of the initial or the final state of the simulation, thus it is independent of the simulated protocols. Nevertheless, different protocols might converge at a different speed to the correct solution.


To implement the approach, we first decide on the number of simulation threads,  $m$  and choose  $m$  equally spaced starting points  $(i - 1)\tau/m$ , where  $i \in \{1, \dots, m\}$ . The  $m$  threads of the simulation are started independently and they are defined by the quadruplets  $(\mathcal{E}_i, A, [i\tau/m, \tau), \mathcal{I})$ , respectively. The planned event set  $\mathcal{E}_i$  of simulation thread  $i$  is defined as:

$$\mathcal{E}_i = \{e | e \in \mathcal{E}, \text{time}(e) \geq \frac{i\tau}{m}\} \tag{11}$$

Note that each simulation thread starts with the same, normally empty, initial state  $\mathcal{I}$ . The first thread, started at  $\tau_1=0$  is performing a traditional, linear simulation. An example of this process, for the case of  $m = 10, \tau=200$  s is illustrated in Fig. 2.

We can obtain the first approximate solution starting from the moment when all the simulation threads have progressed at least  $\tau/m$  in their simulation process. As the simulation progresses further, most time intervals were simulated by more than one thread, but as the threads started at different time points, the length of the warmup period with which the simulations were performed varies. The best approximation is obtained by stitching together the segments from the various simulation runs, *taking every segment from the simulation run where it was executed with the largest amount of warmup*.

For instance, the shaded rectangles in Fig. 2 show the results obtained after each thread progressed over at least  $3\tau/m$  simulation time. The time interval  $[60, 80)$  is simulated by threads 2, 3, and 4, the best approximation is provided by thread 2 with the longest warmup period.



**Fig. 2** Time-parallel simulation with iteratively extended warmup interval. The simulation progressed  $3 \times 200/10 = 60$  s in each thread. The shaded rectangles show the current best approximations for the corresponding time segments

## 4.2 Time-parallel algorithm with initial state approximation

In the previous section, we have seen how our implementation of iteratively extended warmup involves the threads of simulation specified by the quadruplets  $(\mathcal{E}_i, A, [i\tau/m, \tau], \mathcal{I})$ , that is, the simulation threads start with the same, usually empty, initial state  $\mathcal{I}$ . This is one of the major source of errors. To obtain an accurate simulation, we would need to start the simulation with an initial state  $\mathcal{I}_i$ , which describes the state of the simulation as would appear at the corresponding moment in a linear simulation run.

The initial state approximation method relies on running the simulation threads specified by  $(\mathcal{E}_i, A, [i\tau/m, \tau], \mathcal{I}'_i)$  where  $\mathcal{I}'_i \approx \mathcal{I}_i$  is an approximation of the initial state. The whole point of the method is that  $\mathcal{I}'_i$  is computable without the need to run a serial simulation to the corresponding timepoint.

The analysis of Sect. 3.4 shows that the perturbations with the largest effect on the behavior of the network are the changes in the routing tables. Therefore, in our implementation of the initial state approximation, we will concentrate on the estimation of the state of the routing table at the simulated time point. Naturally, this method can be applied only to the routing protocols which use a routing table. The method can not be applied to purely reactive routing protocols such as AODV. On the other hand, reactive routing protocols have less state information, and approximating their initial state with an empty state yields less error. This conjecture is supported by our experiments in Sect. 5.2.

The approximation of the routing table needs to be implemented separately for every routing protocol, as both the routing table and its representation in the simulation code can vary between the different protocols. In the following, we describe our implementation for the DSDV routing protocol.

The implementation requires the understanding of the procedure used by the protocol to build its routing table. In DSDV, every node maintains a routing table which contains all available destinations, the next node to reach a given destination and the number of nodes necessary to reach that destination. In essence, this is a shortest path problem, complicated by the fact that (a) no global view of the network is available and (b) the mobility of the nodes can change the network connectivity and make the routing table obsolete. DSDV solves the first problem by deploying a distributed version of the Bellman-Ford algorithm. This algorithm relies on the iterative improvement of the routing tables through periodic or event-triggered broadcast of the routes to the immediate neighbors of the node. Even in a static network, it needs several send/receive cycles until the

correct distance information is distributed all over the network. In a mobile network, the routing tables contain only an approximation of the correct shortest path information.

Let us now consider how we approximate the initial state of the routing table at time  $t$ . When running a simulation, the researcher is in the position of an omniscient external observer. Knowing the mobility models deployed in the simulation, one can obtain the correct location information of the nodes at any time point  $t_i$ . Using this information, an “ideal” routing table can be obtained by simply running Dijkstra’s shortest path algorithm for each node. Note, that having a global view of the network, there is no point in running a distributed algorithm. The resulting routing table is the one the DSDV algorithm would converge to *if the nodes would be immobile* for a sufficiently long time; this table will be used as the approximate initial state  $\mathcal{I}_i$  in the time-parallel simulation. In practice, the actual routing tables will be somewhat out of date due to node mobility.

This approach can be readily adapted to routing protocols which use a distance vector based routing table; routing protocols which use tables constructed based on different principles (e.g., geographic routing algorithms, directed diffusion, and so on) need different approximation methods.

## 5 Simulation study

We present the results of an investigation of the performance, accuracy, and benefits of our methodology for time-parallel simulation. We first performed the simulation using the serial ns-2 simulator, obtaining a baseline result. Next, we repeated the simulation using the iteratively extended warmup model of time-parallel simulation as described in Sect. 4.1, with data concerning the throughput and packet loss ratio collected after every iteration. We repeated the simulations for the widely used ad-hoc on-demand distance vector routing (AODV) and destination sequenced distance vector routing (DSDV) protocols; for the table-driven DSDV protocol we run the simulation both with and without initial state approximation.

We are concerned only the *relative error* of the results; the absolute values of the measured quantities are not relevant for our study, they are dependent on the scenario and the deployed protocols, rather than the parallelization model. However, the considerations in Sect. 3.4 show that certain parameters of the scenario, such as the mobility of the nodes and the network load, affect the number and nature of perturbations. To quantify this influence we investigate the accuracy of the time-parallel simulation function of these scenario parameters.

In all our experiments, we found that the first iterations lead to a major decrease in the relative approximation error, followed by a smaller improvement in subsequent iterations. Thus, we find it useful to use a logarithmic scale for the presentation of the results.

### 5.1 Simulation setup

To run our simulations in a realistic setting, we chose a setting representative for practical scenarios. We consider a set of mobile nodes moving in a rectangular area, with a set of pre-determined communication patterns. The transmission range of the nodes is significantly smaller than the simulation area. The node-to-node communication is facilitated by a wireless ad hoc network routing protocol. One of the challenges of the scenario is that the mobility of the nodes changes the network topology and the routing of the packets.

Our scenarios use two well known protocols, representative of the major classes of wireless ad hoc routing protocols. DSDV [32] is a pro-active routing protocol in which the nodes maintain routing tables and perform actions to keep them up-to-date. AODV [33] is a reactive, on-demand routing protocol where the routes are established only as a result of explicit demand. These two classes of protocols exhibit different behavior in relation to the time-parallel simulation.

We use the “random waypoint” model [10] to simulate the node movement. Traffic patterns are generated by *constant bit rate* (CBR) sources sending 512-byte UDP packets at a rate of 1 packet per second. The simulation area is  $500 \times 500$  and the default number of nodes is 80. All the nodes have a transmission range of 100 m. The scenario extends over a time interval of 600 s. Table 2 shows the default settings and the range of the parameters for our experiments.

**Table 2** The default values and the range of the parameters for the simulation scenario

Parameter	Default	Range
simulation area	$500 \times 500$ (m <sup>2</sup> )	
number of nodes	80	
transmission range	100 (m)	
speed	1 (m/s)	1–21 (m/s)
pause time	15 (s)	
simulation time	600 (s)	
segment duration	30 (s)	10–60 (s)
number of CBR sources	20	4–40
CBR packet size	512 (bytes)	
CBR sending rate	4 (kbps)	

### 5.2 Achievable speedup

In our implementation  $m$  threads start at the simulation time  $\tau_i = (i - 1)\tau/m$  and run concurrently. The first approximate solution is obtained when all the threads have progressed at least to time  $\tau_d = \tau/m$ , thus the first approximation can be obtained  $m$  times faster compared to a serial simulation, ignoring the overhead required to extract and assemble the approximate trace file.

It would appear therefore, that we want to increase  $m$  up to the number of available processors. However, the accuracy of the results decreases when  $m$  increases. Our approach is to run the simulation for several more iterations  $k$ , in each iteration the simulation progressing another  $\tau/m$  on all threads.

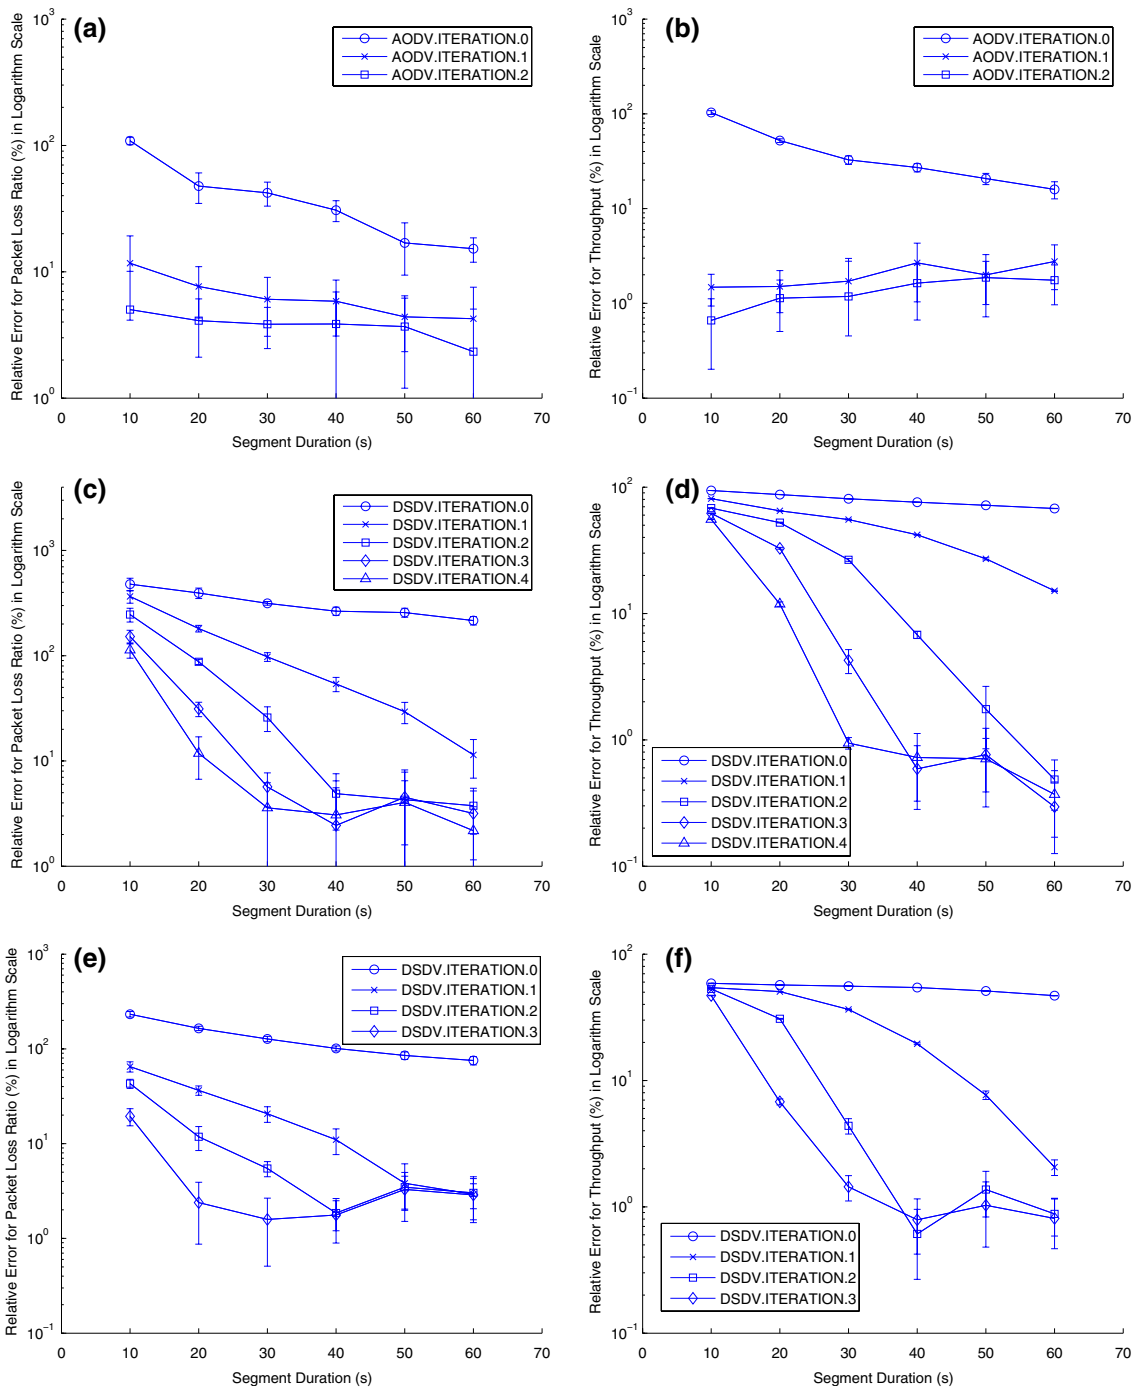
The iterations will stop when the desired accuracy is reached. For this paper, we assume that the desired accuracy is 95%. Let us assume that the simulation requires  $k_{95}$  iterations to reach this level of accuracy. The speedup of the time-parallel simulation will be  $\eta = m/k_{95}$ . The value of  $k_{95}$  depends on many factors: the choice of  $m$ , the length of the simulation, the deployed protocols and the simulation scenario. The use of initial state approximation can reduce the required number of iterations.

In a series of simulation experiments, we have measured the accuracy for various segment sizes  $\tau_d = 10, 20, 30, 40, 50,$  and  $60$  s. In our experiments the simulated time being  $\tau = 600$  s, this corresponds to values of  $m = 60, 30, 20, 15, 12,$  and  $10$ , respectively.

We show the relative error for the packet loss ratio, Fig. 3(a), (c), (e), and the relative error for the throughput, Fig. 3(b), (d), (f). The lines marked *protocolname.ITERATION.i* show the relative error after the  $i$ th iteration. The results show the average and the 95% confidence interval of 10 simulation runs.

The first observation is that for all simulations the accuracy in general is increasing with the number of iterations (although some accidental reversals are possible). As expected, for a given iteration, the accuracy is higher for longer segment sizes (i.e., smaller  $m$  values). The accuracy of the throughput is in general higher than the packet loss ratio, thus we conjecture that the throughput is relatively less sensitive to perturbations than the packet loss ratio.

There is a significant difference between the accuracy obtained at a given iteration for AODV (diagrams a and b) and DSDV (diagrams c and d). The proactive, table-driven DSDV protocol has a much lower accuracy for a given iteration, than the reactive AODV protocol. This is especially noticeable for the packet loss ratio. This validates our inference from Sect. 3.4. We note a significant increase in the accuracy for every iteration.



**Fig. 3** The relative measurement error function of the segment duration at various iterations of the time-parallel simulations. The left side of the figure diagrams (a, c and e) show the measured packet loss ratio, the right side of the figure diagrams (b, d and f) the measured

throughput. Diagrams (a and b) show the results for the AODV routing protocol, diagrams (c and d) for DSDV while diagrams (e and f) for DSDV with initial state approximation used in the time-parallel simulation

Table 3 shows a different view of this data, which might be more important for a researcher. With our assumption that the desired accuracy is 95%, the table shows for each  $m$  and corresponding  $\tau_d$  value the number of iterations necessary to achieve that accuracy  $k_{95}$ , the actual accuracy obtained  $(1 - \varepsilon)$  and the overall speedup of the simulation  $\eta$ .

A conclusion to be drawn from Table 3, is that time-parallel simulation leads to a significant speedup. The speedup is 20 times for AODV, 4.28 times for DSDV without ISA, and 8.6 times for DSDV with ISA. In general, the speedup is greater for reactive routing protocols. Initial state approximation with the model described in Sect. 4.2

**Table 3** The number of iterations  $k_{95}$  needed, the accuracy  $(1 - \varepsilon)$  and the speedup  $\eta$ , as a function of segment duration, for the simulation of AODV, DSDV without ISA and DSDV with ISA

$m$ -no. of threads	$\tau_d$ -segment duration (s)	AODV			DSDV w/o ISA			DSDV with ISA		
		$k_{95}$	$1 - \varepsilon$ (%)	$\eta$	$k_{95}$	$1 - \varepsilon$ (%)	$\eta$	$k_{95}$	$1 - \varepsilon$ (%)	$\eta$
60	10	3	95.6	20.0	14	97.3	4.28	7	98.4	8.6
30	20	2	95.4	15.0	8	96.9	3.75	4	97.6	7.5
20	30	2	95.4	10.0	5	96.7	4	4	98.4	5.0
15	40	2	95.8	7.5	4	96.5	3.75	3	98.1	5.0
12	50	2	95.8	6.0	4	97.7	3	2	96.3	6.0
10	60	2	95.7	5.0	3	97.3	3.3	2	97.2	5.0

can significantly increase the speedup for proactive, table-driven protocols as well, but it comes with the disadvantage that it requires additional code besides the standard ns-2 protocol implementation.

### 5.3 The influence of the scenario parameters on the accuracy

Our results show that a significant speedup can be achieved, while the accuracy is dependent on the choice of the routing protocols. In a series of experiments, we investigated the influence of other parameters of the scenario on the accuracy of the time-parallel simulation. In these experiments, we used a segment duration  $\tau_d = 30$  s corresponding to  $m = 20$  simulation threads. For all the measurements, we measured the relative errors of packet loss ratio and throughput for AODV, DSDV without ISA, and DSDV with ISA.

Figure 4 shows the relative errors at various iterations function of the network load. The network load is simulated by varying the number of CBR sources (from 5 to 45). Figure 4(a) shows the relative error of the packet loss ratio and Fig. 4(b) shows the relative error for the throughput. We see a very slight tendency for the relative error on packet loss ratio to increase with the throughput for the early iterations. In general, however, the relative error remains almost constant across the range of the CBR sources considered. At the third iteration, the relative error was below 5% for all CBR values.

Figure 4(c), (d) shows the relative error for packet loss ratio and throughput, respectively, for the simulation of the DSDV protocol. Interestingly, the relative error for packet loss ratio shows a decreasing trend with the number of CBR sources for the early iterations and a slightly increasing trend for the later iterations. For the throughput, the relative error is almost constant for a given iteration, independently of the network load. These slight trends notwithstanding (some of which might be an artifact of the experimental setup), in general the relative error consistently decreases with the number of iterations for both

packet loss ratio and throughput, for all the tested values of CBR sources.

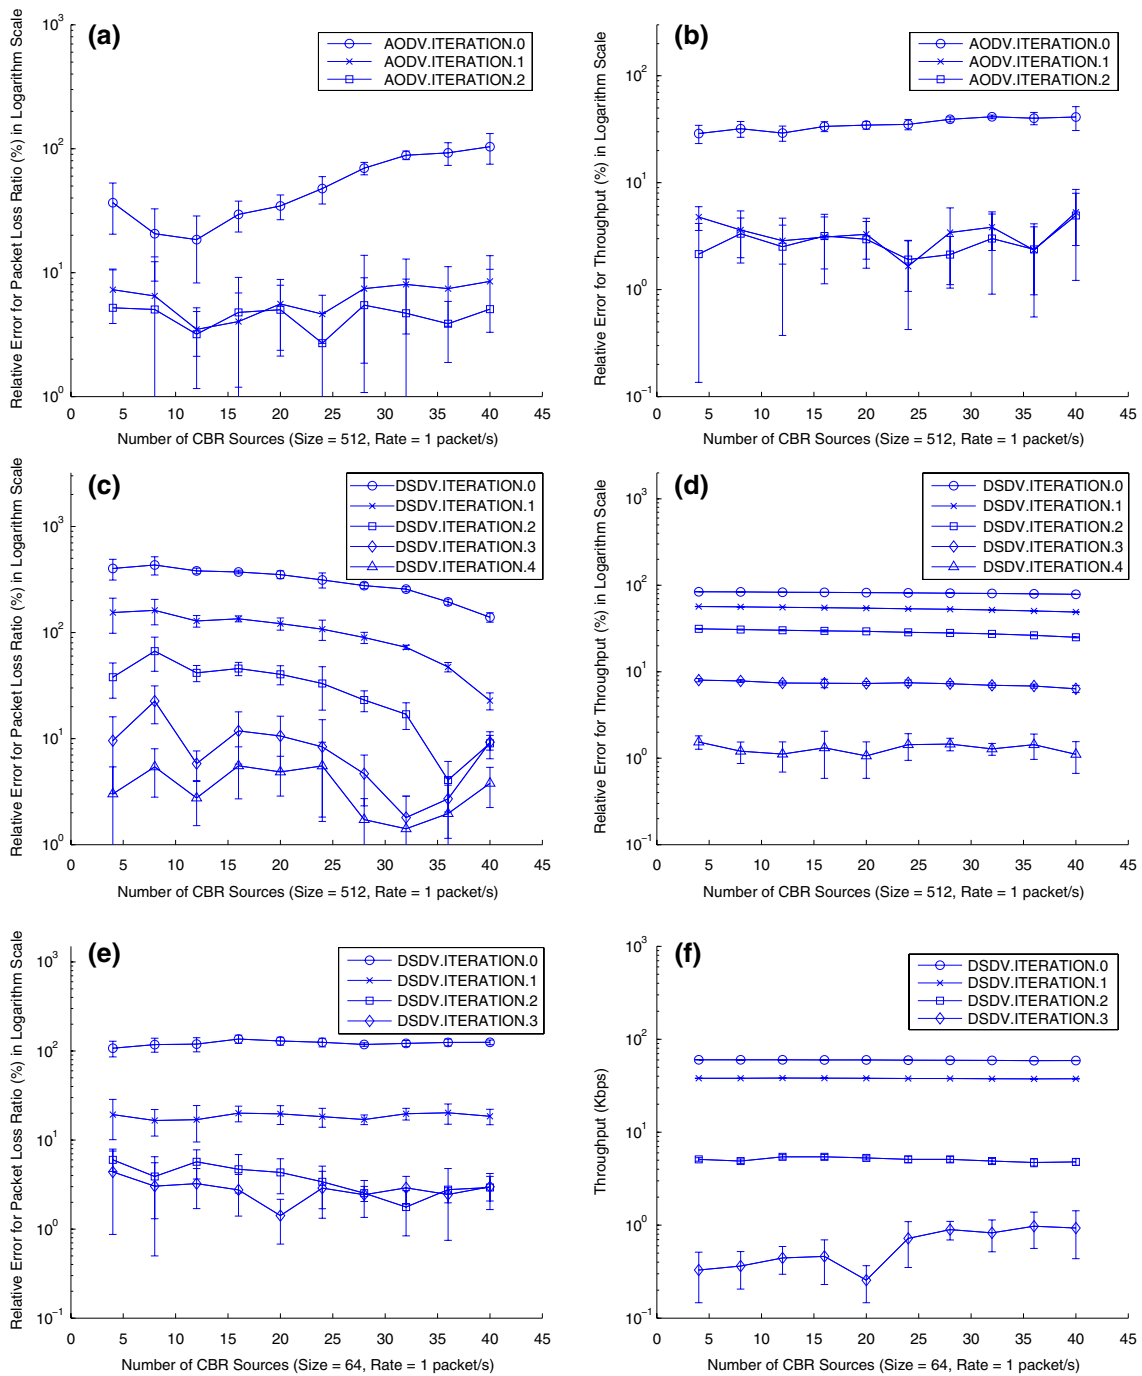
Figure 4(e), (f) shows the values for DSDV with the initial state approximation approach being used in the time-parallel simulation. The trends are essentially the same as in the case without initial state approximation. However, for the equivalent iterations the relative error is significantly lower. For higher network loads, we see that the relative error for packet loss ratio is actually very close for iterations 2 and 3, and in one case, for 32 CBR sources, the iteration 3 shows a slight increase in the relative error over iteration 2 (although both errors are less than 5%). This reversal indicates the limits of the improvement obtainable with the initial state approximation approach.

In conclusion, we find that the relative error of the packet loss ratio and throughput shows only very slight dependence on the network load. The general trend is that the relative error consistently decreases with the number of iterations.

Figure 5 shows the same set of measurements test function of the average mobility of the network nodes, ranging from 1 to 21 m/s. For most experiments, the results show that the relative error has a slight tendency to decrease when node mobility increases. A plausible explanation is that in a highly mobile network the routing tables and the cached flow entries are recomputed more frequently, thus limiting the influence of perturbations. Other than this, all the previously observed tendencies remain valid for all the possible values of node mobility. The relative error decreases with the number of iterations for both the packet loss ratio and the throughput. In general, AODV converges to a relative error of less than 5% in three iterations. DSDV converges much slower, but the convergence can be sped up using initial state approximation. Overall, the influence of the mobility on the accuracy is minor and predictable.

We conclude that the only significant parameter of the scenario is the choice of the routing protocol and, for proactive routing protocols, whether initial state approximation was deployed or not. This is a favorable result, because it limits the number of variables a researcher needs to control for a time-parallel simulation.





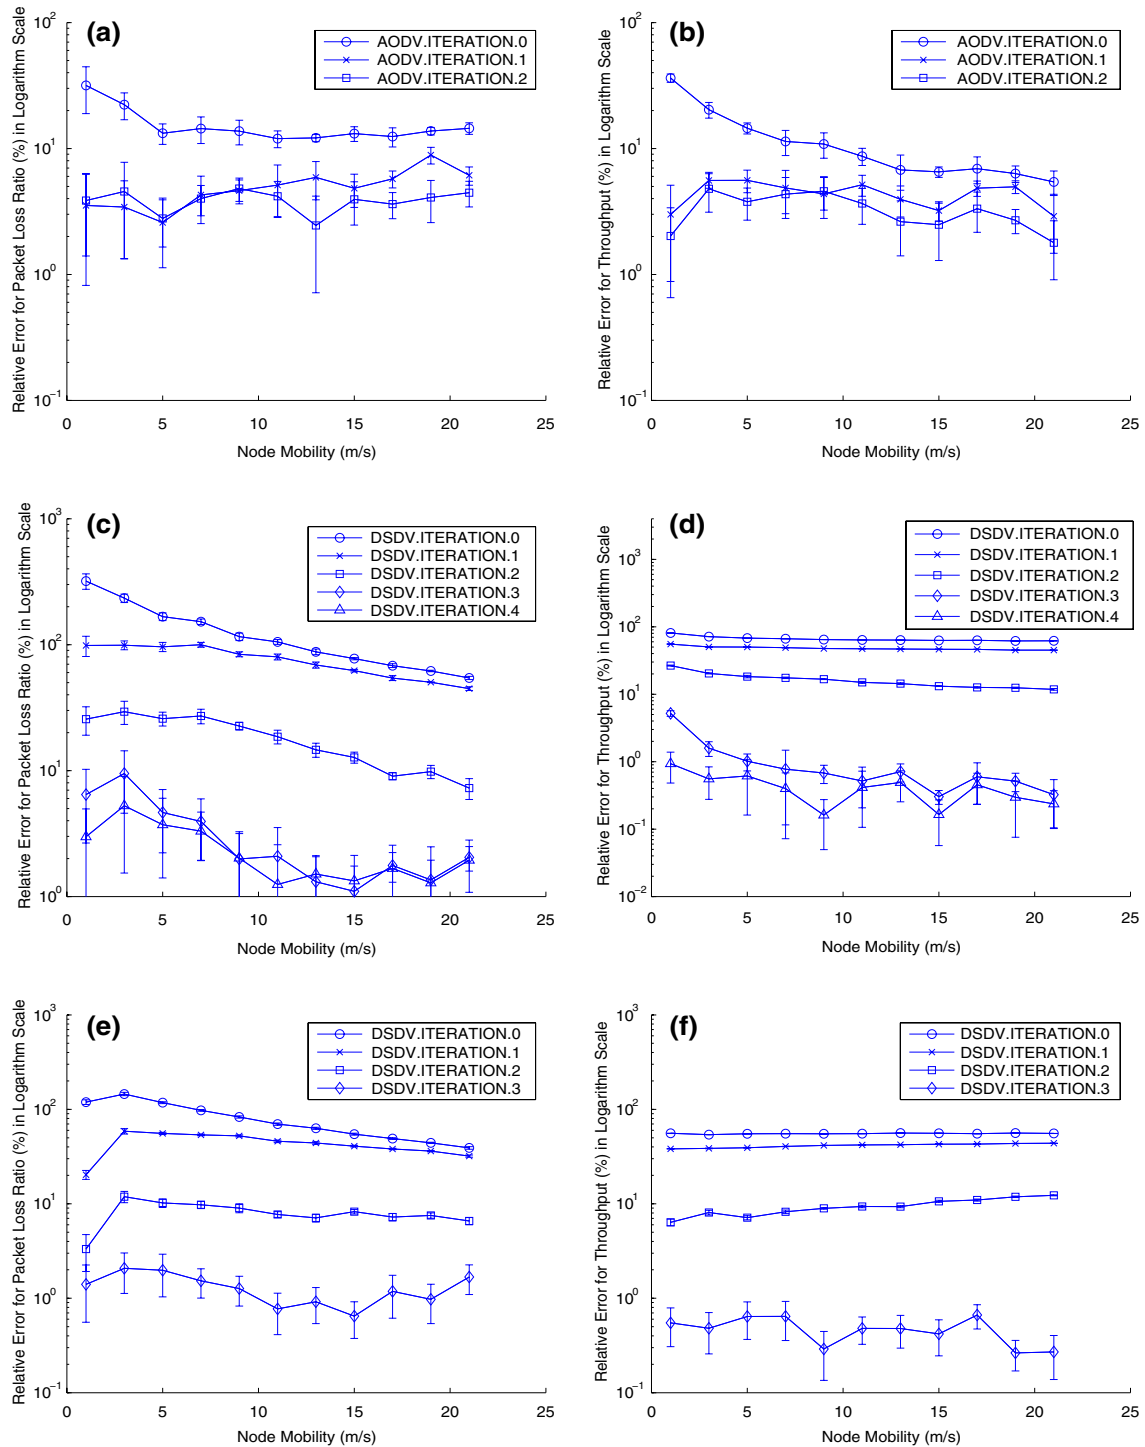
**Fig. 4** The relative measurement error function of the network load at various iterations of the time-parallel simulations. The left side of the figure diagrams (a, c, and e) show the measured packet loss ratio, the right side of the figure diagrams (b, d, and f) the measured

throughput. Diagrams (a and b) show the results for the AODV routing protocol, diagrams (c and d) for DSDV while diagrams (e and f) for DSDV with initial state approximation used in the time-parallel simulation

## 6 Summary and future work

In this paper, we described a methodology for time-parallel simulation of wireless ad hoc networks. We presented a quasi-formal analysis of perturbations, which gives us

some understanding of the source of errors in one-shot time-parallel simulations. Based on a layer-by-layer analysis of the propagation of perturbations in the wireless networking stack, we proposed several avenues for improving the accuracy of the time-parallel simulation.



**Fig. 5** The relative measurement error function of the node mobility at various iterations of the time-parallel simulations. The left side of the figure diagrams (a, c, and e) show the measured packet loss ratio, the right side of the figure diagrams (b, d and f) the measured

throughput. Diagrams (a and b) show the results for the AODV routing protocol, diagrams (c and d) for DSDV while diagrams (e and f) for DSDV with initial state approximation used in the time-parallel simulation

Building on these considerations, we described an implementation of time-parallel simulation based on the ns-2 simulator. The techniques deployed are the iterative

extension of the warmup period, and initial state approximation for the proactive, table-driven routing protocols. A series of experiments showed that a speedup between 5 and

20 times can be obtained for an accuracy of 95%, depending on the choice of the routing protocol and whether initial state approximation was deployed or not. In general, the simulation of reactive routing protocols shows a higher accuracy and/or speedup than the one for proactive, table-driven protocols. We show, however, that the accuracy is relatively independent on other parameters of the scenario, such as network load or node mobility.

Time-parallel simulation can be used to study measures of performance such as packet loss ratio and throughput, but there are other measures, such as end-to-end delay, that require a different approach.

Future work include improved initial state approximation techniques that would bring the speedup of the table driven protocols closer to the one for reactive protocols. We are also considering more sophisticated models to predict the accuracy of the simulation for arbitrary scenarios and combinations of networking protocols. Finally, we plan to develop software components for time-parallel simulation of wireless ad hoc networks on cluster computers, without the requirement of in-depth knowledge of the protocol implementation.

**Acknowledgments** The research reported in this paper was partially supported by National Science Foundation grants ACI0296035 and EIA0296179. One of the authors (dcm) acknowledges support from the Science Foundation of Ireland (SFI) through a 2007 Ernst T.S. Walton award.

## References

- Andradóttir, S., & Ott, T. J. (1995). Time segmentation parallel simulation of networks of queues with loss or communication blocking. *ACM Transactions on Modeling and Computer Simulation*, 5(4), 269–305.
- Avril, H., & Tropper, C. (1996). The dynamic load balancing of clustered time warp for logic simulation. In *Proceedings of the 10th Workshop on Parallel and Distributed Simulation (PADS)* (pp. 20–27).
- Boloni, L., Turgut, D., Wang, G., & Marinescu, D. (2006). Challenges and benefits of time-parallel simulation of wireless ad hoc networks. In *Proceedings of the 1st International Conference on Performance Evaluation Methodologies and Tools (Valuetools)*, October.
- Boukerche, A., & Das, S. K. (1997). Dynamic load balancing strategies for conservative parallel simulations. In *Proceedings of the 11th Workshop on Parallel and Distributed Simulation (PADS)* (pp. 20–28).
- Boukerche, A., Das, S. K., & Fabbri, A. (2001). SWimNet: A scalable parallel simulation testbed for wireless and mobile networks. *Wireless Networks* 7(5), 467–486.
- Boukerche, A., Das, S. K., Fabbri, A., & Yildiz, O. (1999). Exploiting model independence for PCS parallel simulation. In *Proceedings of 13th ACM/IEEE Workshop on Parallel and Distributed Simulation* (pp. 166–173).
- Boukerche, A., & Fabbri, A. (2000). Partitioning parallel simulation of wireless networks. In *Proceedings of the 2000 Winter Simulation Conference* (pp. 1449–1457).
- Boukerche, A., & Tropper, C. (1994). A static partitioning and mapping algorithm for conservative parallel simulations. *ACM SIGSIM Simulation Digest* 24(1), 164–172.
- Briner, J. (1991). Fast parallel simulation of digital systems. In *Proceedings of Multiconference on Advances in Parallel and Distributed Simulation* (pp. 71–77).
- Broch, J., Maltz, D. A., Johnson, D. B., Hu, Y. C., & Jetcheva, J. (1998). A performance comparison of multi-hop wireless ad hoc network routing protocols. In *Proceedings of the 4th International Conference on Mobile Computing and Networking (MobiCom)* (pp. 85–97).
- Carothers, C. D., Fujimoto, R. M., England, P., & Lin, Y. B. (1994). Distributed simulation of large-scale PCS networks. In *Proceedings of the 2nd IEEE International Workshop on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems* (pp. 2–6).
- Carothers, C. D., Fujimoto, R. M., & Lin, Y. B. (1995). A case study in simulating PCS networks using Time Warp. In *Proceedings of the 9th Workshop on Parallel and Distributed Simulation (PADS)* (pp. 87–94).
- Choe, M., & Tropper, C. (1999). On learning algorithms and balancing loads in time warp. In *Proceedings of the 13th Workshop on Parallel and Distributed Simulation (PADS)* (pp. 101–108).
- Deelman, E., & Szymanski, B. K. (1998). Dynamic load balancing in parallel discrete event simulation for spatially explicit problems. In *Proceedings of the 12th Workshop on Parallel and Distributed Simulation (PADS)* (pp. 46–53).
- Devine, K., Hendrickson, B., Boman, E., John, M. S., & Vaughan, C. (2000). Design of dynamic load-balancing tools for parallel applications. In *Proceedings of the 14th International Conference on Supercomputing (ICS)* (pp. 110–118).
- Fujimoto, R. M. (1989). Time warp on a shared memory multi-processor. *Transactions of the Society for Computer Simulation* 6(3), 211–239.
- Fujimoto, R. M. (1990). Parallel discrete event simulation. *Communications of the ACM* 33(10), 30–53.
- Glazer, D. W., & Tropper, C. (1993). On process migration and load balancing in time warp. *IEEE Transactions on Parallel and Distributed Systems* 4(3), 318–327.
- Heidelberger, P., & Nicol, D. (1993). Conservative parallel simulation of continuous time Markov chains using uniformization. *IEEE Transactions on Parallel and Distributed Systems* 4(8), 906–921.
- Hoseyni-Nasab, M., & Andradóttir, S. (1996). Parallel simulation by time segmentation: Methodology and applications. In *Proceedings of the 1996 Winter Simulation Conference* (pp. 376–381).
- Hoseyni-Nasab, M., & Andradóttir, S. (1998). Time segmentation parallel simulation of tandem queues with manufacturing blocking. In *Proceedings of the 1998 Winter Simulation Conference* (pp. 1487–1492).
- Jefferson, D. R. (1985). Virtual time. *ACM Transactions on Programming Languages and Systems* 7(3), 404–325.
- Kiesling, T. (2004). Approximate time-parallel cache simulation. In *Proceedings of the 2004 Winter Simulation Conference* (pp. 345–354).
- Kiesling, T. (2005). Using approximation with time-parallel simulation. *Simulation* 81(4), 255–266.
- Lin, Y. B., & Lazowska, E. D. (1990). Exploiting lookahead in parallel simulation. *IEEE Transactions on Parallel and Distributed Systems* 1(4), 457–469.
- Lin, Y.-B., & Lazowska, E. D. (1991). A time-division algorithm for parallel simulation. *ACM Transactions on Modeling and Computer Simulation* 1(1), 73–83.

27. Madiseti, V. K., Walrand, J. C., & Messerschmitt, D. G. (1991). Asynchronous algorithms for the parallel simulation of event-driven dynamical systems. *ACM Transactions on Modeling and Computer Simulation* 1(3), 244–274.
28. Nicol, D. (1992). Conservative parallel simulation of priority class queuing networks. *IEEE Transactions on Parallel and Distributed Systems*, 3(3), 294–303.
29. Nikolaidis, I., Fujimoto, R., & Cooper, C. A. (1993). Parallel simulation of high-speed network multiplexers. *IEEE Conference on Decision and Control*, 3(1), 2224–2229.
30. Nikolaidis, I., Fujimoto, R., & Cooper, C. A. (1994). Time-parallel simulation of cascaded statistical multiplexers. In *Proceedings of the ACM SIGMETRICS Conference on Measurement and Modeling of Computer Systems* (pp. 231–240).
31. Panchal, J., Kelly, O., Lai, J., Mandayam, N., Ogielski, A. T., & Yates, R. (1998). WiPPET, a virtual testbed for parallel simulations of wireless networks. *ACM SIGSIM Simulation Digest*, 28(1), 162–169.
32. Perkins, C., & Bhagwat, P. (1994). Highly dynamic destination-sequenced distance-vector routing (DSDV) for mobile computers. In *Proceedings of ACM SIGCOMM* (pp. 234–244).
33. Perkins, C., & Royer, E. (1999). Ad hoc on-demand distance vector routing. In *Proceedings of the 2nd IEEE Workshop on Mobile Computing Systems and Applications* (pp. 99–100).
34. Quaglia, F. (2001). A cost model for selecting checkpoint positions in time warp parallel simulation. *IEEE Transactions on Parallel and Distributed Systems* 12(4), 346–362.
35. Quaglia, F., & Santoro, A. (2003). Nonblocking checkpointing for optimistic parallel simulation: Description and an implementation. *IEEE Transactions on Parallel and Distributed Systems*, 14(6), 593–610.
36. Reiher, P. L., & Jefferson, D. (1990) Dynamic load management in the time warp operating system. *Transactions of the Society for Computer Simulation*, 7(2), 91–120.
37. Soliman, H. M., & Elmaghraby, A. S. (1998). An analytical model for hybrid checkpointing in time warp distributed simulation. *IEEE Transactions on Parallel and Distributed Systems*, 9(10), 947–951.
38. Som, T. K., & Sargent, R. G. (2000). Model structure and load balancing in optimistic parallel discrete event simulation. In *Proceedings of the 14th Workshop on Parallel and Distributed Simulation (PADS)* (pp. 147–154).
39. VINT. The UCB/LBNL/VINT network simulator-ns (version 2). URL <http://www.isi.edu/nsnam/ns>
40. Wang, J. J., & Abrams, M. (1992). Approximate time-parallel simulation of queueing systems with losses. In *Proceedings of the 1992 Winter Simulation Conference* (pp. 700–708).
41. Wang, J. J., & Abrams, M. (1993). Determining initial states for time-parallel simulations. In *Proceedings of the 7th Workshop on Parallel and Distributed Simulation (PADS)* (pp. 19–26).
42. Weigle, M. C., Adurthi, P., Hernandez-Campos, F., Jeffay, K., & Smith, F. D. (2006). Tmix: A tool for generating realistic TCP application workloads in ns-2. *ACM SIGCOMM Computer Communication Review*, 36(3), 65–76.
43. Wieland, F. (2001). Practical parallel simulation applied to aviation control. In *Proceedings of the 15th Workshop on Parallel and Distributed Simulation* (pp. 109–116). ACM/IEEE Computer Society.
44. Wilson, L. F., & Nicol, D. M. (1996). Experiments in automated load balancing. In *Proceedings of the 10th Workshop on Parallel and Distributed Simulation (PADS)* (pp. 4–11).
45. PDNS – Parallel/Distributed NS. URL <http://www.cc.gatech.edu/computing/compass/>, 2004.

## Author Biographies



**Guoqiang Wang** received a B.S. degree from Southeast University, Nanjing, China, in 2001. He is currently a Ph.D. Candidate in the School of Electrical Engineering and Computer Science at the University of Central Florida, Orlando. His research areas include ad hoc routing, Grid computing, and parallel simulation.



**Damla Turgut** is an Assistant Professor with the School of Electrical Engineering and Computer Science at University of Central Florida. She received her B.S., M.S., and Ph.D. degrees from the Computer Science and Engineering Department of University of Texas at Arlington in 1994, 1996, and 2002, respectively. She has been included in the WHO's WHO among students in American Universities and Colleges in 2002. She has been awarded outstanding research award

and has been recipient of the Texas Telecommunication Engineering Consortium (TxTEC) fellowship. She is a member of IEEE, member of the ACM, and the Upsilon Pi Epsilon honorary society. Her research interests include wireless networking and mobile computing, distributed systems, and embodied agents.



**Ladislau Bölöni** is an Assistant Professor with the School of Electrical Engineering and Computer Science at University of Central Florida. He received a Ph.D. degree from the Computer Sciences Department of Purdue University in May 2000. He received a Master of Science degree from the Computer Sciences Department of Purdue University in 1999 and Diploma Engineer degree in Computer Engineering with Honors from the Technical University of Cluj-Napoca, Romania in 1993.

He received a fellowship from the Computer and Automation Research Institute of the Hungarian Academy of Sciences for the 1994-95 academic year. He is a senior member of IEEE, member of the ACM, AAAI and the Upsilon Pi Epsilon honorary society. His research interests include autonomous agents, grid computing and wireless networking.



**Dan C. Marinescu** joined the Computer Science Department at University of Central Florida in August 2001 as Professor of Computer Science. He has been an Associate and then Full Professor of Computer Science at Purdue University, in West Lafayette, Indiana, since 1984. He is the Scientific Director of the I2.Lab, an organization supporting interdisciplinary research in computer and information sciences (<http://i2lab.ucf.edu>) at UCF. He

is conducting research in parallel and distributed systems, computational biology, quantum computing, parallel and distributed

computing, and has published more than 180 papers in journals and referred conference proceedings in these areas. He is the author of “Internet-Based Workflow Management: Towards a Semantic Web,” published by Wiley in 2002 and has co-edited “Process Coordination and Ubiquitous Computing.” The book “Approaching Quantum Computing,” co-authored with Gabriela M. Marinescu was published in September 2004 by Prentice Hall.