

# A Java-based Simulator for Ad Hoc Mobile Distributed Systems

Fernando Pereira<sup>1</sup>, Marco Tulio Valente<sup>2</sup>, Roberto S. Bigonha<sup>1</sup>, and Mariza A. S. Bigonha<sup>1</sup>

<sup>1</sup> Department of Computer Science, Federal University of Minas Gerais, Brazil

<sup>2</sup> Department of Computer Science, Catholic University of Minas Gerais, Brazil  
{fernandm, bigonha, mariza}@dcc.ufmg.br, mtov@pucminas.br

**Abstract.** This paper presents a Java-based simulator for the PeerSpaces coordination model. PeerSpaces is a shared space coordination model designed for ad hoc mobile networks. The model is based on the concept of tuple spaces firstly proposed by Linda. In order to address the dynamic environment that characterizes ad hoc networks, PeerSpaces does not assume the presence of any centralized structure for communication and coordination. This paper also describes an algorithm for termination detection in ad hoc networks built on PeerSpaces.

## 1 Introduction

Ad hoc networks are wireless networks that do not rely on any based station infrastructure for communication [13, 6]. Instead, hosts depend on each other to send and receive messages. In ad hoc networks, two hosts can exchange messages whenever they happen to be at communication range. Bluetooth is an example of network technology that makes this form of wireless communication possible [15].

Designing applications for such networks presents many interesting problems [18, 17]. Particularly, coordination is a challenging task. Since a user can find himself in a different network at any moment, the services available to him change along the time. Thus, computation should not rely on any predefined and well known context. Also, coordination should not assume the existence of any centralized node because permanent availability of this node cannot be granted. Communication should also be uncoupled in time and space, meaning that two communicating entities do not need to establish a direct connection to exchange data nor must know the identity of each other.

This paper describes a Java based simulator for the PeerSpaces model and a distributed algorithm built on this simulator. PeerSpaces is a coordination model for ad hoc networks based on Linda [5]. In Linda, several process communicate through a central data repository called *tuple space*. Processes communicate by inserting (**out**), reading (**rd**) and removing (**in**) ordered sequences of data from this space. Tuple retrieving is associative because it is based on a pattern against which a matching tuple is non-deterministically chosen from the space. If a matching tuple is not found, the caller process is suspended until such tuple is posted in its tuple space.

Communication in Linda presents many characteristics that are desirable in mobile environment. Particularly, communication is asynchronous and uncoupled in time and space. Communicating processes do not need to create a socket-like connection to exchange data. The associative mechanism allows communication based on the contents of the messages rather than on their addresses or other identifiers. Moreover, the blocking semantics used to retrieve tuples automatically provides synchronization among processes. All these features are important in mobile systems, since they are characterized by dynamic and short-lived patterns of communication

In traditional Java-based Linda systems, like TSpaces [21] and JavaSpaces [4], the tuple space is a centralized and global data structure that runs in a predefined service provider. In the base station scenario this server can easily be located in the fixed network. However, if operations in ad hoc mode are a requirement, this is not possible because in this case the fixed infrastructure simply does not exist. This suggests that standard client/server implementations of Linda are not suitable for ad hoc scenarios, since they assume a tight coupling between client and servers and the permanent availability of the latter.

Aiming to answer the new requirements posed by ad hoc mobile computing systems, PeerSpaces departs from the traditional client/server architecture used in Linda and push towards a completely decentralized one. In this model, each node (or peer) has the same capabilities, acting as client, shared space provider and as router of messages. In order to provide support to operations in ad hoc mode, service lookup is distributed along the network and does not require any previous knowledge about the network topology.

This paper is organized as follows. In Section 2, the PeerSpaces model is informally presented, including its main design goals, concepts and primitives. Section 3 provides an overview of a stochastic simulator for PeerSpaces. Section 4 presents the implementation and simulation in PeerSpaces of a termination detection algorithm. Detecting the termination of diffusing computation is an important problem in distributed systems and is used in this paper as an example of algorithm that can be built on PeerSpaces. Finally, Section 6 concludes the paper.

## 2 PeerSpaces: A Coordination Model for Ad Hoc Mobile Systems

PeerSpaces assumes an ad hoc network of mobile devices. Thus, there is no infrastructure network and hosts may connect or disconnect at any moment. As usual in ad hoc settings, two hosts can communicate when their wireless interfaces are in the same vicinity. The model does not assume any centralized structure and does not promise to provide any kind of shared memory abstraction encompassing connected hosts. Instead, it fosters a peer to peer model of computation, where any connected node has the same capabilities. Furthermore, hosts can discover each other using a decentralized lookup service and then communicate using remote primitives.

The main concepts used in PeerSpaces are the following:

**Hosts** The model assumes that hosts are mobile devices. Each host has its own local tuple space and a set of running threads. The host-level tuple space has three main purposes. Firstly, it is used for local coordination among threads running in the host. Secondly, it is used for remote communication, since there are primitives in the model to retrieve and output messages in the space of remote hosts. Thirdly, it is used to publish *services* and to retrieve the results of *lookup queries*. A service is any entity available in the host that can be useful to other hosts. Services in PeerSpaces are defined by tuples, whose fields describe the attributes of the service. Finally, a lookup query is a query performed in the network to discover services.

**Network** Mobile hosts in the model are connected by a wireless and ad hoc network. As usual in such networks, connectivity is transient and determined by the distance among hosts. Consequently, the topology of the network is continuously changing. Moreover, any host in the network can act as router, relaying messages between nodes that are not directly connected.

## 2.1 PeerSpaces Primitives

PeerSpaces defines a set of primitives to assemble applications using the previous defined concepts. The set of primitives of PeerSpaces is a superset of the primitives originally proposed by Linda.

As in Linda, the **out**  $v$  primitive inserts tuple  $v$  into the tuple space of the local node. The remote **out**  $h, v$  primitive inserts tuple  $v$  in the tuple space of host  $h$ . Remote **out** is asynchronous meaning that if host  $h$  is not reachable the operation immediately returns after leaving tuple  $v$  stored in space of the issuing host  $h'$ . When a communication path is established between  $h'$  and  $h$  the tuple is automatically transmitted to its final destination.

The **in**  $p, x$  primitive removes a tuple that matches pattern  $p$  from the local space and binds it to  $x$ . If there are several matching tuples, one of them is chosen non-deterministically. If there is no matching tuple, the calling thread remains blocked until the operation can be completed. The remote operation specifies a remote host name as a parameter as in **in**  $h, p, x$ .

The non-blocking version of **in** is called **inp** (*probe in*). If it is not possible for **inp**  $p, x$  to match a tuple, a null reference is bound to name  $x$ . There is also a remote version for the **inp** operation. The last primitive of the **in** family is called **ing**. The operation **ing**  $p, a[]$  retrieves from the local tuple space all tuples which match pattern  $p$  and stores them into the array  $a$ . The operation does not block the calling thread if no tuple is found.

In order to retrieve information from tuple spaces without removing data PeerSpaces provides a reading operation. The **rd**  $p, x$  operation is similar to a local **in** but it does not change the state of the tuple space. The remote version of the operation is **rd**  $h, p, x$  and the non blocking version is named **rdp**. Like

**rd**, the non-blocking operation also has a remote version. Finally, **rdg**  $p, a[]$  fills the array  $a$  with all tuples in the local space matching pattern  $p$ .

PeerSpaces extends Linda with a lookup primitive used to discover services in the network. Since the model is designed for use in ad hoc networks, this primitive does not make use of any central authority, as a directory service. Instead, the execution of the lookup primitive, called **find**, is distributed along the federation of connected devices. The **find**  $p$  operation queries hosts in the network for tuples matching pattern  $p$ . All matching tuples found are copied asynchronously to the local space of the host that has called the operation.

For example, a PDA-based auction system that wants to find other PDA's users selling 21-inch TVs despite their brand, price and seller can issue the following operation: **find**  $\langle \text{mall}, \text{sellors} \rangle, \langle \text{tv}, 21, ?, ?, ? \rangle$ . The operation will trigger a query for tuples that match the pattern  $\langle \text{tv}, 21, ?, ?, ? \rangle$  in the hosts of subgroup **sellors** of the root group **mall**. Matching tuples, like  $\langle \text{tv}, 21, \text{foo}, 325, \text{h} \rangle$ , where **foo**, **325** and **h** are respectively the TV's brand, price and the name of the mobile host offering the TV, will be outputted in the local space of the PDA sometime after the operation was issued. The PDA system can then retrieve these answers using the local **in** primitive and place an offer  $v$  using the remote operation **out**  $h, v$ .

The **find** primitive originates a query that is propagated to all nodes of the network. Basically, the host that originated the query transmits it to its neighbors, that retransmit it to their neighbors and so on, until the network graph is covered. This protocol is similar to the one used by distributed file sharing systems in the Internet, like Gnutella [7]. Not surprisingly, the logical network created by Gnutella over the fixed Internet presents many characteristics that are typical of wireless and ad hoc networks.

Often it is useful to query connected hosts for a service and keep the query effective until such service is available. In this way, a client does not need to periodically send lookup queries to detect new services that may become available since the last query was issued. In PeerSpaces, lookup queries that remain active after their first execution are called continuous queries. Continuous lookup queries are issued adding the lifetime  $t$  to the **find** primitive: **find**  $p, t$ . This primitive will search hosts of the network for available services matching pattern  $p$  and for services that may become available in  $t$  time units.

### 3 A Java-based Simulator for PeerSpaces

In order to design and test distributed ad hoc systems in PeerSpaces, a simulator was implemented for the model. The simulator is written in Java and is available as a Java package. For tuple space manipulation, the simulator uses a Java-based tuple space system named LighTS [8].

The simulator supports a stochastic model of simulation. The behavior of network elements is based on probabilistic parameters. The simulation is driven by a discrete time and random events are generated following the parameters defined by the user. The user specifies the number of nodes in the network and

the size of the grid where the nodes are located. For each node, the user defines its pattern of movement and the range of its wireless interface.

In order to implement and test a distributed algorithm, the user has to extend some predefined classes and interfaces. Each algorithm is a set of classes that implements the `Command` interface. This interface has two methods: `eval` (to evaluate the command guard) and `exec` to execute the command action. Several algorithms can be composed together in order to create a more complex one.

The simulator was implemented in two layers. In the first layer the user has the illusion that communication between two entities is carried out directly. In this layer, the user makes use of the primitives of the model. The second layer implements the routing protocol used by the `PeerSpaces` primitives. In this layer, messages are transmitted only between adjacent nodes.

The routing algorithm implemented in the second layer of the simulator is based on flooding [19]. According to this technique, after receiving a message whose final address is not its own, a node send the message to all its neighbors. Although the implementation of this algorithm is straightforward, the exchange of information between two nodes that are not neighbors can generate a large number of messages in the network. On the other hand, flooding is fast and reliable in the sense that if there is a path connecting sender and receiver, then a message will be delivered in the shortest possible time. The problem of loops in the transmission of messages is solved by assigning identifiers to each message.

## 4 Example: A Termination Detection Algorithm

In order to introduce the design of distributed algorithms within `PeerSpaces`, an algorithm for termination detection of diffusing computation is described next. Termination detection is an important problem in distributed systems. Basically the problem consists in a node getting the information that a computation previously spread through the network has been completed by every element in the system. For instance, in some public key cryptographic systems each node must change its public key after a certain amount of time for the sake of security. If such a system is used in a distributed environment, when a node changes its key it should be assured that each other element of the network knows the new key before sending secure messages.

The solution described next was first presented in [16] and is based on a well-known solution proposed by Dijkstra and Scholten [3]. The approach adopted is divided in three phases. In the first phase, a partial ordering is built on the network, starting with a root node in such a way that in the end of the process, any nodes  $a$ ,  $b$  and  $c$  will have different identifiers. Those identifiers must follow the transitive property on the *less than* operation, meaning that if  $a < b$  and  $b < c$  then  $a < c$ <sup>1</sup>. The second phase of the algorithm involves the activation of idle nodes through the propagation of the job. Upon activation, a node becomes the child of the activating node and enters the active state. In the third and

---

<sup>1</sup> In the algorithm,  $a < b$  means that the rank of node  $a$  is smaller than the rank of  $b$ .

last phase, the termination of the distributed computation is actually detected. Each node after terminating the requested job, passes this information to a node of higher rank and eventually all termination reports will reach the root node, which has the highest rank.

Figure 1 describes the main variables and auxiliary functions used by the guarded commands of the algorithm.

$n$	Host identifier
$root$	Root node identifier
$\pi$	Name of the local tuple space of the node
$count(t)$	Number of occurrences of tuples matching pattern $t$
$isBetter(r_1, r_2)$	True if $r_1$ is a better rank than $r_2$

**Fig. 1.** Global variables for node  $n$  and auxiliary functions

The termination detection algorithm is implemented in three separated modules. The first one creates a partial ordering in the network. Since this module implements a well-known algorithm, it will be not described in this paper. The result of this algorithm is the insertion of a tuple in the format  $\langle \text{"rank"}, root, rank \rangle$  into the tuple space of each connected node. In this tuple,  $root$  is the name of the node which started the ranking process and  $rank$  is the position of the node in the ordered sequence. The complete implementation of the partial ordering algorithm can be found at [12].

#### 4.1 Job Diffusion

Job diffusion begins with the root node and follows a flooding pattern of transmission. Each node, after receiving a job, *activates* all of its neighbors by re-transmitting the job. The job executed in each node is represented by a tuple that matches the pattern  $\langle \text{"job"}, n_o, ttl \rangle$ , where  $n_o$  is the node that sent the tuple and  $ttl$  is an integer used to simulate the time needed to finish the execution of the job.

The activation of a node by another one is recorded by both participants, but in different ways. The node  $n_a$  that has received the job posts the tuple  $\langle \text{"child"}, n_a \rangle$  in the tuple space of the node that sent the job. Node  $n_a$  then changes its states from idle to active, which is indicated by the insertion of the tuple  $\langle \text{"activeJob"}, ?, ? \rangle$  in the tuple space of the just activated node. Figure 1 describes the commands responsible for the job diffusion. The value *job duration* in the command **StartJobDiffusion** simulates the time that each node will need to finish the task.

#### 4.2 Detecting Termination

The tuple  $\langle \text{"activeJob"}, ?, ttl \text{ as Integer} \rangle$  simulates the job assumed by a node.  $ttl$  is an integer representing the job complexity. At each iteration of the command

**StartJobDiffusion**

Guard:

 $\langle \text{"job"}, n, \text{Integer} \rangle \notin \pi \wedge n = \text{root}$ 

Action:

**out**  $\langle \text{"job"}, n, \text{job duration} \rangle$ **JobPropagation**

Guard:

 $\langle \text{"activeJob"}, \text{String}, \text{Integer} \rangle \notin \pi \wedge \mathbf{count}(\langle \text{"job"}, \text{String}, \text{String} \rangle) \geq 1$ 

Action:

**rdp**  $\langle \text{"job"}, s_1 \text{ as String}, ttl \text{ as Integer} \rangle, \tau$ **if**  $(s_1 \neq n)$ **out**  $s_1, \langle \text{"child"}, n \rangle$  $\forall$  neighbor  $v$  of  $n$  do**out**  $v, \langle \text{"job"}, n, ttl \rangle$ **out**  $\langle \text{"activeJob"}, s_1, ttl \rangle$ **Table 1.** Job diffusion

*nodeExecuteTask* the *ttl* value is decremented by one until it reaches zero, meaning the job is completed. When a node finishes the task it was given, it generates an *idle report* by putting a tuple matching  $\langle \text{"idleReport"}, n, \text{children} \rangle$  in its local tuple space. In the next step the idle report is propagated to a node with a better rank. A node performs this action by sending to every neighbor a tuple containing information about the idle reports it currently holds. This tuple, called *node info*, contains the rank, activated children and the identifier of the node. Upon receiving such tuple, the receiving node compares its rank with the rank of the remote node. If the local rank wins, the receiving node attempts to retrieve the idle report tuple from the remote node. A non-blocking **inp** operation is used in this case because as the node info may be sent to several nodes more than one **in** operation could be performed over the same tuple.

The propagation of idle reports does not continue forever because there is not a node with a better rank than that of the root node. When receiving the idle reports the root attempts to build and prune a tree of activated nodes. If it is possible for the root to recursively remove all the leaves from the tree the termination algorithm is finished.

## 5 Related Work

Many characteristics of PeerSpaces have been inspired in file sharing applications popular in the Internet, like Napster [10], Freenet [2] and Gnutella [7]. Particularly, the peer to peer network created by Gnutella over the fixed Internet presents many properties that are interesting in mobile settings, like absence of centralized control, self-organization and adaptation to failures. PeerSpaces is an effort to transport and adapt such characteristics to mobile computing systems.

Jini [1] is a distributed object infrastructure that adds support to dynamic service registration and lookup to Java RMI [20]. However, the system assumes the existence of a central server to run the lookup service, which restricts its use

**NodeExecuteTask**

Guard:

 $\langle \text{"activeJob"}, s_1 \text{ as String}, ttl \text{ as Integer} \rangle \in \pi \wedge ttl > 0$ 

Action:

```

inp  $\langle \text{"activeJob"}, s_1 \text{ as String}, ttl \text{ as Integer} \rangle, \tau$ 
out  $\langle \text{"activeJob"}, s_1, ttl - 1 \rangle$ 
if  $(ttl - 1 = 0)$ 
  Tuple [] children;
  ing  $\langle \text{"child"}, \text{String} \rangle, \text{children}$ 
   $\forall \langle \text{"child"}, s_i \text{ as String} \rangle \in \text{children}$  do
    out  $\langle \text{"idleReport"}, n, s_0 | s_1 | \dots | s_k \rangle$ 

```

**PropagateIdleReports**

Guard:

 $\langle \text{"nodeInfo"}, \text{String}, \text{String}, \text{String}, \text{String} \rangle \notin \pi \wedge$   
 $\langle \text{"idleReport"}, \text{String}, \text{String} \rangle \in \pi \wedge n \neq \text{root}$ 

Action:

```

rdp  $\langle \text{"rank"}, \text{root as String}, \text{rank as String} \rangle, \tau$ 
rdp  $\langle \text{"idleReport"}, \text{father as String}, \text{children as String} \rangle, \sigma$ 
 $\forall$  neighbor  $v$  of  $n$  do
  out  $v, \langle \text{"nodeInfo"}, \text{father}, \text{children}, \text{rank}, n \rangle$ 

```

**AcceptIdleReports**

Guard:

 $\langle \text{"nodeInfo"}, \text{String}, \text{String}, \text{String}, \text{String} \rangle \in \pi$ 

Action:

```

inp  $\langle \text{"nodeInfo"}, f \text{ as String}, ch \text{ as String}, r \text{ as String}, s \text{ as String} \rangle, \sigma$ 
rdp  $\langle \text{"rank"}, \text{root as String}, \text{localRank as String} \rangle, \tau$ 
if  $(\text{isBetter}(\text{localRank}, r))$ 
  out  $\langle \text{"idleReport"}, f, ch \rangle$ 
  inp  $s, \langle \text{"idleReport"}, f, ch \rangle, x$ 

```

**PruneTree**

Guard:

 $\text{count}(\langle \text{"idleReport"}, \text{String}, \text{String} \rangle) > 1 \wedge \langle \text{"done"} \rangle \notin \pi$ 

Action:

```

inp  $\langle \text{"idleReport"}, \text{father as String}, \text{children as String} \rangle, \sigma$ 
inp  $\langle \text{"idleReport"}, \text{child as String}, \emptyset \rangle, \tau$ 
if  $(\tau = \langle \text{"idleReport"}, \text{root}, \emptyset \rangle)$ 
  out  $\langle \text{"done"} \rangle$ 
else
  out  $\langle \text{"idleReport"}, \text{father}, \text{children} - \text{child} \rangle$ 

```

**Table 2.** Termination detection

to networks with base station support. The Jini framework also includes a Linda-like shared data space implementation, called JavaSpaces [4]. Once more, the system assumes that the data space resides in a central server, which precludes its utilization when operating in ad hoc mode. The same problem is shared by other client/server implementations of Linda, like TSpaces [21].

Lime [14, 9] introduces the notion of transiently shared data space to Linda. In the model, each mobile host has its own tuple space. The contents of the local spaces of connected hosts are transparently merged by the middleware creating the illusion of a global and virtual data space. Applications in Lime perceive the effects of mobility by atomic changes in the contents of this virtual space. However, even when used in a small federation of hosts, the main problems of transiently shared spaces are efficiency and scalability. The reason is the amount of global synchronization required to assure the consistency of the virtual space.

Particularly, query operations must run as a distributed transaction to retrieve matching tuples. Moreover, the model allows users to define the destination tuple space of an outputted tuple. This leads to the notion of misplaced tuples, i.e., tuples that are temporally in a wrong tuple space waiting for the connection to its target host. Thus, the host engagement protocol also requires a distributed transaction to deliver misplaced tuples. Finally, disengagements in Lime should be announced in order to remove event handlers placed at remote hosts.

There are several simulators for mobile distributed systems. Probably, the most well known are `ns` [11] and `GloMoSim` [22]. They do not support, however, the implementation of distributed algorithms using high level abstractions, like tuple spaces.

## 6 Conclusions and Future Work

In this paper we have presented a simulator for PeerSpaces, a coordination model for mobile computing systems. PeerSpaces was designed to overcome the main shortcoming of shared space coordination models when used in ad hoc wireless networks – the strict reliance on the traditional client/server architecture – while preserving the main strengths of such models – the asynchronous and uncoupled style of communication. The design of the model has privileged observance to ad hoc networks principles. As usual in such models, transparency is sacrificed in name of scalability and soundness. In order to illustrate the programming approach proposed by PeerSpaces, a termination detection algorithm was presented.

There are many research opportunities in the PeerSpaces model. For instance, PeerSpaces lacks mechanisms to control the access to the shared space. Also, we intend to investigate the use of more efficient routing algorithms.

The PeerSpaces simulator, including its source code, can be downloaded from the URL: <http://www.dcc.ufmg.br/~fernandm/~peerspaces>.

## References

1. K. Arnold. *The Jini Specifications*. Addison-Wesley, 2nd edition, 2000.
2. I. Clarke, O. Sandberg, B. Wiley, and T. Hong. Freenet: A distributed anonymous information storage and retrieval system. In *ICSI Workshop on Design Issues in Anonymity and Unobservability*, International Computer Science Institute, 2000.
3. E. W. Dijkstra and C. S. Scholten. Termination detection for diffusing computation. *Information Processing Letter*, i 4(11), Aug 1980.
4. E. Freeman, S. Hupfer, and K. Arnold. *JavaSpaces principles, patterns, and practice*. Addison-Wesley, Reading, MA, USA, 1999.
5. D. Gelernter. Generative communication in Linda. *ACM Transactions on Programming Languages and Systems*, 7(1):80–112, Jan. 1985.
6. S. Giordano. *Mobile Ad-Hoc Networks*, chapter of Handbook of Wireless Networks and Mobile Computing. John Wiley & Sons, 2002.
7. Gnutella Home Page. <http://gnutella.wego.com>.
8. LighTS Home Page. <http://lights.sourceforge.net>.

9. A. L. Murphy, G. P. Picco, and G.-C. Roman. Lime: A middleware for physical and logical mobility. In *Proceedings of the 21<sup>st</sup> International Conference on Distributed Computing Systems*, May 2001.
10. Napster Home Page. <http://www.napster.com>.
11. NS Home Page. <http://www.isi.edu/nsnam/ns/>.
12. F. M. Q. Pereira. Implementation of a termination detection algorithm in the peerspaces model. Technical report, Universidade Federal de Minas Gerais, Brazil, Sept. 2002.
13. C. Perkins. *Ad Hoc Networking*. Addison-Wesley, 2000.
14. G. P. Picco, A. L. Murphy, and G.-C. Roman. Lime: Linda meets mobility. In D. Garlan, editor, *Proceedings of the 21<sup>st</sup> International Conference on Software Engineering*, pages 368–377. ACM Press, May 1999.
15. R. Mettala. Bluetooth Protocol Architecture. <http://www.bluetooth.com>.
16. C. Roman and J. Payton. A termination detection protocol for use in mobile ad hoc environment. Technical Report WUCS-01-29, Washington University, Department of Computer Science, St. Louis, Missouri, 2001.
17. G.-C. Roman, G. P. Picco, and A. L. Murphy. Software Engineering for Mobility: A Roadmap. In A. Finkelstein, editor, *The Future of Software Engineering*, pages 241–258. ACM Press, 2000.
18. M. Satyanarayanan. Fundamental challenges in mobile computing. In *ACM Symposium on Principles of Distributed Computing*, May 1996.
19. A. S. Tanenbaum. *Computer Networks*. Prentice Hall, 3rd edition, 1996.
20. A. Wollrath, R. Riggs, and J. Waldo. A distributed object model for the Java system. In *2nd Conference on Object-Oriented Technologies & Systems*, pages 219–232. USENIX Association, 1996.
21. P. Wycko, S. W. McLaughry, T. J. Lehman, and D. A. Ford. TSpaces. *IBM Systems Journal*, 37(3):454–474, Aug. 1998.
22. X. Zeng, R. Bagrodia, and M. Gerla. Glomosim: A library for parallel simulation of large-scale wireless networks. In *Workshop on Parallel and Distributed Simulation*, pages 154–161, 1998.