Scalable Parallel Simulations of Wireless Networks with WiPPET: Modeling of Radio Propagation, Mobility and Protocols *

O. E. Kelly, J. Lai, N. B. Mandayam, A. T. Ogielski, J. Panchal, R. D. Yates

Wireless Information Network Laboratory (WINLAB)
Rutgers University, 73 Brett Road
Piscataway, NJ 08854–8060, U.S.A.

We review the design, selected applications and performance of WiPPET (Wireless Propagation and Protocol Evaluation Testbed), a general parallel simulation testbed for various types of wireless networks. WiPPET has been written in TeD/C++, an object-oriented modeling framework that isolates network modeling from the underlying parallel discrete event simulator. We describe the techniques for modeling radio propagation (long and short-scale fading and interference) and protocols that promote scalability of parallel simulations at session and packet time-scales. We outline two selected applications of WiPPET: integrated radio resource management in a mobile wireless voice network; and packet losses due to mobility and short-scale fading over a radio link. We illustrate the efficiency of the simulator under two partitioning schemes with parallel performance data obtained using the Georgia Time Warp optimistic simulator.

Keywords: Wireless Networks; Parallel simulation; Radio resource management

1. Introduction

Modeling of radio networks is very distinct from modeling of wired networks in that the physical channel properties, i.e., radio propagation and interference, cannot be separated from the higher network protocols, because strong interactions at all levels drive engineering design decisions. From the simulation point of view, recurring interference calculations, due to changes in each transmit-

* This work is partially supported by DARPA Contract N66001-96-C-8530 and by NSF Grant NCR-9527163.

ter's power or position, lead to grand demand for computational resources. The parallel simulator WiPPET[1] extends the functionality of WINLAB's previous sequential simulator MADRAS [1], the Mobility And Dynamic Resource Allocation Simulator. MADRAS has figured importantly in research (e.g., [2–5]) and has demonstrated the necessity of an integrated approach to radio resource allocation. It also showed the desirability of the reusable testbed approach. Prior results regarding parallel simulations with WiPPET can be found in [6] which casts the wireless network into reusable and substitutable parts, and [7] which evaluates the efficiency of parallel simulations. Other work on parallel simulators for wireless systems includes [8–10].

WiPPET has been implemented in the TeD/C++ framework [11–14], and is used with the Georgia Time Warp (GTW) optimistic parallel simulator [15]. WiPPET consists of a collection of base TeD objects modeling *mobility*, *traffic*, *radio propagation* and *protocols* that together with a configuration file (specifying interconnections and system parameters) can be instantiated as a virtual testbed [6]. The objectives of the testbed are to: (i) Shift the work to *design of models* rather than low–level programming of custom simulators, and (ii) Enable reuse of a growing database of objects for creation of new simulation models. Therefore, the design of base entities requires the efficient implementation of the *fundamental* behaviors, while expressing the *unifying* features of the different design options to enable entity substitutions.

Current and near-term research involving simulations with WiPPET focuses on radio resource management and media access protocols, effects of packet losses due to interference and fading on the performance of Internet protocols, and studies of interference among nearby autonomous systems employing distinct modulation and radio resource management algorithms.

In order to be useful in multiple research projects, a wireless system simulator must have the ability to model radio phenomena at different temporal and spatial scales, as well as the ability to model multiple protocol layers, as sketched in Figure 1. The levels of granularity are dictated both by the hierarchy of transmission units in the system (e.g., frames, packets, bits or chips), and by characteristic time-scales for the physical processes with which transmissions interact (e.g., short–scale Rayleigh fading [16], or long–scale shadow fading [17]).

[1] Wireless Propagation and Protocol Evaluation Testbed. Also a breed of fast racing dogs (whippet).

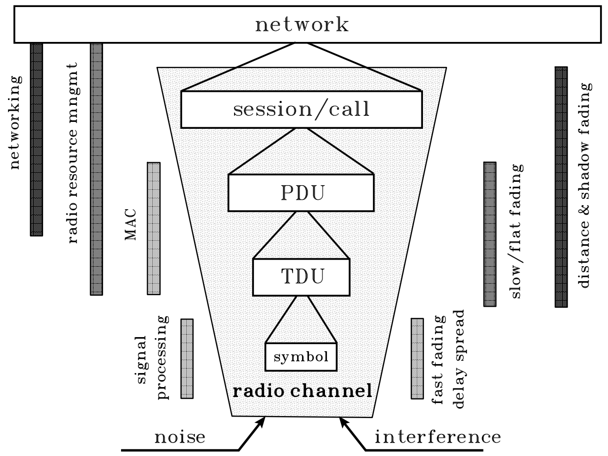

Figure 1. WiPPET provides support and "hooks" for model development at a variety of time scales and protocol layers.

Moreover, a versatile wireless simulator should support research studies *across* levels and time-scales. Such a system facilitates the incorporation of interactions between layers in the design and optimization process.

2. The TeD Simulation Framework

We begin by briefly reviewing the object oriented modeling approach of the TeD framework [18]. TeD allows us to design new system elements or protocol models as well as to reuse "black box" designs created by others. TeD carefully separates the core modeling constructs from the executable constructs. These modeling constructs, defined in terms of a meta language, enclose executable code fragments written in any general-purpose programming language. The current TeD compiler accepts executable content expressed in C++ , with a macro interface to mediate between the executable language and the runtime simulator services. In the following, we describe the principal TeD language constructs. More complete, detailed documentation and examples can be found in the language manuals.

Entity The *entity* declaration is used to define a black-box view of both real-life objects, such as base stations, as well as as abstract objects, such as protocols. Entities are connected to each other by typed *channels*. Channels transmit the *events* that describe dynamic interactions between entities. An output channel

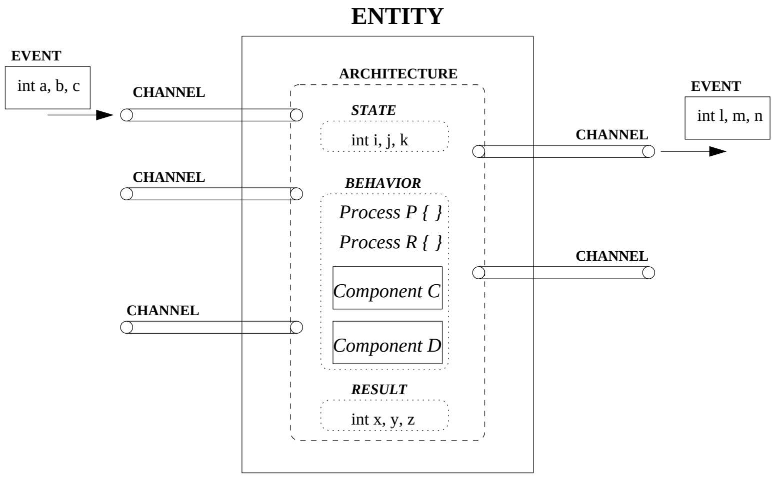

Figure 2. Basic framework of an Entity

of an entity can be mapped to an input channel of another entity. The behavior of an entity is defined solely in terms of its dynamic reactions to events on its input channels, and the production of events on its output channels.

Architecture The *architecture* of an entity defines the actions of the entity upon event arrivals on its input channels, and the production of events on its output channels. An architecture is composed of a small number of code blocks:

- *Deferred Constants*: Variables whose values are held constant during run-time, but can be initialized to different values for different architecture instances.
- *State*: A set of variables that together form a part of the abstract state of the modeled entity.
- *Processes*: Threads of computation that act on the events arriving on the input channels of the entity, or that act upon time advances. Processes may generate events on output channels as part of their computation.
- *Components*: Sets of entities that *logically* form sub-entities of an entity's behavior.

Structural and behavioral descriptions of an entity can be inherited, thus allowing object-oriented hierarchical design and development. A derived architecture inherits all the items of the parent architecture (state, processes, etc.), which can also be redefined (overridden).

TeD models are designed without reference to any specific parallel simulation engine. TeD models are thus translated by a special TeD compiler that generates the C++ code which can directly use the services of the parallel simulation kernel. The generated C++ code is then compiled with a standard C++ compiler and linked with the parallel simulator. Once the models TeD objects have been compiled, a separate *configuration* specification is used to instantiate a desired network model configuration, thus avoiding recompilation when a modeler wishes to change model parameters, or even a modeled network topology. The network thus instantiated can be simulated on a multiprocessor machine. While the discussion in this section has been in general terms about the simulation framework, we next review some of the modeling details that are used in WiPPET .

3. Mobility, Propagation, and Protocols in WiPPET

The design of primary TeD entities for a wireless network simulation has to be convenient for modelers. The design of WiPPET takes into account selection of appropriate entities that model the earlier mentioned aspects of *mobility*, *traffic*, *radio propagation* and *protocols* that are necessary to execute a realistic simulation of a wireless system. Natural candidates for entities are mobile stations, base stations and switches. Thus entities `MS`, `BS`, and `switch` each represent their physical counterpart. These entities contain TeD processes simulating session level radio resource management protocols for call admission, transmitter power control and handoff. The entity modeling traffic, `callgen` generates calls (session arrivals). Entities called `zones` encapsulate geography, spectral occupancy, propagation, and SIR calculations. While the creation and design of entities follows more or less a logical analogy to physical systems, one of the key issues is to capture the details of *mobility*, *traffic*, *radio propagation* and *protocols* according to the realistic communication environment that is being simulated. In the following, we discuss details of appropriate models for the above.

3.1. Geography and Mobility

A mobile wireless simulator must include a description of the terrain in which a mobile station (MS) can travel. Moreover, the geometry of the terrain determines gross radio propagation characteristics. In WiPPET, the continuous terrain is approximated by an adjacency graph of accessible MS locations.

Mobility describes the process by which an MS moves from one point to another within the modeled geography. WiPPET allows the modeler to import an approximation to a real map, or a synthetic geography, and to model either trajectories that evolve at runtime, or precomputed trajectories that simulate a particular traffic scenario. The current version of WiPPET uses common model of mobility in microcellular systems that is a biased random walk in a synthetic Manhattan–style urban environment with a 5–10 meter geographic resolution.

3.2. Propagation

The performance of a radio receiver depends on receiving sufficient signal energy from a desired transmitter while not receiving too much interference from other sources. Further, a received signal may be smeared by the delay spread in multipath propagation. Radio reception quality exhibits very large fluctuations in space and time (fading) due to environmental motion and activity of other sources.

There is a vast body of literature on the measurements and modeling of propagation in various environments at various frequencies, for a review see [19, 20]. The classical models are derived as follows: The ith terminal, located at position $\mathbf{x}_i$, transmits signal $s_i(t)$. The received signal at location $\mathbf{b}$ over a general time varying radio channel is

$$R_i(t, \mathbf{b}) = \int_{-\infty}^{t} h(t, \tau, \mathbf{x}_i, \mathbf{b}) s_i(\tau)\, d\tau \tag{1}$$

where $h(t, \tau, \mathbf{x}, \mathbf{b})$ is the time varying and spatially varying channel impulse response.

The modeling of the channel response $h(t, \tau, \mathbf{x}, \mathbf{b})$ depends on performance assessment needs: it may be site–specific, that is based on actual measurements or on ray-tracing approximation for a particular town or building; or it may employ stochastic processes with values and correlations mimicking those measured in a class of of similar environments (e.g., typical American suburbs, or generic multi-story office buildings). While the design of WiPPET allows site–specific propagation models, in our research we focus on statistical characterization of radio networks, and thus we limit the discussion to stochastic propagation models.

Signal attenuation in a radio channel is commonly represented as the product of long-scale and short-scale fading: $h(t, \tau, \mathbf{x}, \mathbf{b}) = g(\mathbf{x}, \mathbf{b}) \cdot h(t, \tau)$. The long-scale fading, g, represents the attenuation of the signal averaged in a region larger than 10 or more carrier wavelengths around the receiver, thus due to distance and shadowing by large objects. The short-scale factor $h(t, \tau)$ represents time-varying effects of interference and delay dispersion of carrier waves arriving over distinct paths.

Long-scale Fading

Long-scale fading is constant in time and may be precomputed. For cellular networks WiPPET employs the path gain matrix $g(i,j) = \log g(\mathbf{x}_i, \mathbf{b}_j)$ for all $\mathbf{x}_i \in \mathcal{M}$ and $\mathbf{b}_j \in \mathcal{B}$, where $\mathcal{M}$ is the set of mobile locations, and $\mathcal{B}$ is the set of base station locations. In WiPPET the $|\mathcal{M}| \times |\mathcal{B}|$ gain matrix g is a large state constant that may be read from a file, precomputed in a site-specific or stochastic propagation model. The classic stochastic model represents the geographic path losses as $g(\mathbf{x}, \mathbf{b}) = g_d(|\mathbf{x} - \mathbf{b}|) \cdot g_s(\mathbf{x}, \mathbf{b})$ where g_d is deterministic function of distance and g_s is a log-normally distributed spatial random process. Measured $g_d(|\mathbf{x} - \mathbf{b}|)$ decays as an inverse power of distance, $g_d \sim |\mathbf{x} - \mathbf{b}|^{-\eta}$. In free space, $\eta = 2$; however, in complex environments with propagation dominated by phenomena ranging from waveguide effects to extensive shadowing and scattering, the values $1 < \eta < 4$ are found [20]. A standard stochastic description of the shadow fading fluctuations due to large obstacles is as a single realization of a two dimensional log-normal random process, $\log g_s \sim \mathcal{N}(0, \sigma^2)$, with exponentially decaying spatial correlations, $E\left[\log g_s(\mathbf{x}, \mathbf{b}) \log g_s(\mathbf{x}', \mathbf{b})\right] = \sigma^2 \exp(-|\mathbf{x} - \mathbf{x}'|/d_0)$ [17].

Short–scale fading

The term short-scale fading encompasses the effects of the interference of the components of multipath propagation of radio waves, and their varying travel time from the source to the receiver. Sklar [16] provides a useful taxonomy: First, the relation between the maximum excess delay T_m in impulse response (due to multipath propagation) and the symbol duration time T_s determines two signal degradation categories: i) **Frequency–selective fading** for $T_m > T_s$, producing channel-induced intersymbol interference (ISI); ii) **Flat fading** for $T_m < T_s$, when multipath components are not resolvable. Second, the relation between the channel coherence time T_0 (characterizing temporal variability due to changing multipath components) and the symbol duration time T_s determines two other signal degradation categories: i) **Fast fading** for $T_0 < T_s$, producing symbol distortion, e.g. due to high Doppler; ii) **Slow fading** for $T_0 > T_s$, giving rise to

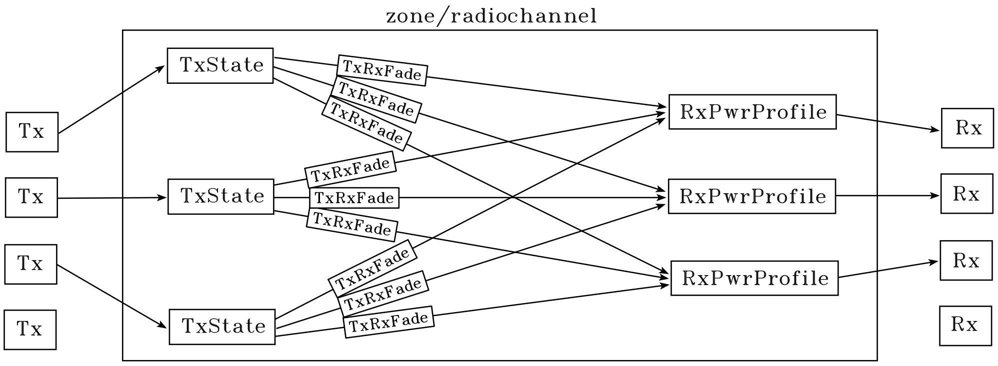

Figure 3. Interference calculations must combine the contributions of *all* actively transmitting Tx to the received power profile for *every* Rx. Each time the power, or position, or the value of short–scale fading loss of a Tx changes, *all* active Rx power profiles in that radio channel must be updated.

SIR fluctuations.

Modeling of frequency-selective fading, e.g., for a RAKE receiver, employs temporal fading (like Rayleigh) separately on each resolved path, together with a stochastic process of arrivals and departures of multipath components, with appropriate delay and power distribution. Such models have been derived in literature both for urban and indoors environments. There are also various models of temporal fading: Rayleigh, Rice (LOS) are typical. With these, temporal variation appears through relative motion of the Tx and Rx. However, omnipresent environmental motion introduces temporal fading even if both Tx and Rx are stationary, but that case is difficult to characterize because it is highly environment dependent. The case in which relative motion is the larger effect is characterized by Rayleigh and Ricean fading.

In contrast to long–scale fading, modeling of short–scale fading need not be geographically indexed due to short–range correlations. In WiPPET the stochastically generated short–scale temporal fading is modeled independently for each Tx–Rx pair in the shared radio channel and combined at the Rx.

3.3. Radio Resource Management Protocols

Radio resource management includes any actions that are necessary in a wireless system to provide the required Quality of Service (QoS) over the radio link. In wireless voice networks the term has generally included protocols that control call admission, channel allocation, power control, and hand–off. In a wireless data network, there are also transport protocols acting above the resource management activities.

Almost all aspects of radio resource management depend on accurate measurements of the quality of the wireless link. Moreover, for a large class of systems and protocols, link quality is measured in terms of the *signal to interference ratio* (SIR). From equation (1), the received power of mobile i at position $\mathbf{b}$ is

$$P_{\mathrm{RX}}^{(i)}(t,\mathbf{b}) = E\left[R_i^2(t,\mathbf{b})\right] \tag{2}$$

At position $\mathbf{b}$, the SIR of terminal i is the ratio

$$\Gamma_i(t,\mathbf{b}) = \frac{P_{\mathrm{RX}}^{(i)}(t,\mathbf{b})}{\sum_{j\neq i} P_{\mathrm{RX}}^{(j)}(t,\mathbf{b}) + \eta^2} \tag{3}$$

where η^2 is the additive noise power.

In a typical mobile radio propagation situation, the received signal will experience both lognormal shadow fading and Rayleigh fading. They can be considered to be superimposed because of the independence assumption between the two phenomena. A direct analysis (or simulation) of performance in this superimposed situation is possible by considering the Rayleigh analysis with the local mean of the lognormal random variable. However, in these two fading situations, the active time-scales in performance considerations are different [21]. For instance, if we consider an application (for e.g., call dropping in cellular voice communications) where this time-scale is a relatively large duration (i.e., on the order of seconds), we only need to consider the effect of lognormal shadow fading since the longest fade caused by Rayleigh fading will be much shorter for a given mobile velocity. Similarly, for an application with a very small active time-scale (i.e., on the order of milliseconds), such as packet transmission, we need to consider primarily time-dependent performance analysis for Rayleigh fading. Even though the mean power level is still affected by the lognormal shadow fading, it changes much more slowly than the Rayleigh fading fluctuations. Other variables that could affect the relative importance of the type of fading are the use of interleavers and coders.

It is also important to note that the details of the fading model and the interference measurements should be matched to the signal processing capabilities of the modeled receiver to obtain the meaningful statistics of data loss. For example, when the receiver can track the radio channel, the expectation in equation (2) may be conditioned by knowledge of the channel impulse response h. Further, the evaluation of a radio resource management protocol dictates SIR measurements on a time scale consistent with the operation of the protocol. For example, simulations of intercell handoff protocols employ SIR measurements averaged over seconds[22] whereas packet-level simulations need multiple SIR values on the short time-scale of a packet transmission to determine if the packet's data have been corrupted.

The interaction of transmitted signals due to users sharing the same bandwidth results in a cross–connection of transmitters to receivers as illustrated in Figure 3. However, a transmitter and receiver may act independently if they do not share the same spectrum, or if they are sufficiently separated in distance that $g(\mathbf{x}, \mathbf{b}) \approx 0$. That is, depending on the structure of the system being simulated, the cross–connection may be essentially complete, or relatively sparse. Nonetheless, we will see that the interference calculations dominate the computational effort of the simulation.

4. Applying WiPPET to Radio Resource Management

In this section, we describe the use of WiPPET on problems of *radio resource management*. As mentioned earlier, the two versions of WiPPET detailed below namely $\mathrm{WiPPET}_{\mathrm{session}}$ and $\mathrm{WiPPET}_{\mathrm{packet}}$ differ from each other primarily in that the performance evaluation in each case is on time-scales that are orders of magnitude apart. Accordingly, radio resource management algorithms also differ in their structure in that the signal quality used as an indicator for management, is assessed over disparate time-scales (see [23] for details).

In each of the above two versions, a mobile continuously moves under the control of the WiPPET *Mobility* module. When a mobile gets to a new geographic point, all the cross–connections (link gains) shown in Figure 3 have to be updated accordingly. This includes the pathloss/shadow (long scale) fading value, the short scale fading value (due to Rayleigh or Ricean fading) as well as the interference powers at the current time index. The objective of radio resource management algorithms is then to update radio resources (such as power, frequency and connectivity) according to the SIR measured at that time. $\mathrm{WiPPET}_{\mathrm{session}}$ and $\mathrm{WiPPET}_{\mathrm{packet}}$ differ in evaluating this SIR that constitutes the fundamental difference between session and packet level communications. In $\mathrm{WiPPET}_{\mathrm{session}}$, the active time-scales are much larger than the scale on which short scale variations of signal occur, and thus signal measurements made may assumed to be averaged sufficiently to remove short scale effects. This precludes the need for short scale fading in session level algorithms and countermeasures. However, in WiPPET$_{\text{packet}}$, the active time-scales in performance evaluation are at the level of packet durations and and hence there is a need to account for both long scale and short scale fading. Further, there is also a need to account for error-control techniques such as error-correction coding [21].

4.1. WiPPET*: Session level simulation*

WiPPET$_{\text{session}}$, a version of WiPPET that models session or call level traffic, has radio propagation models including distance loss and shadow fading, mobility and resource management algorithms – call admission, uplink power control, and six handoff algorithms. Among the six are three standards based handoff algorithms (AMPS, IS-136, GSM) and three other handoff algorithms (SIR based handoff, received-signal-strength (RSSI) based handoff, and a combined SIR/RSSI based handoff) derived from [3]. Related work also includes modeling and simulation of the recent TDMA based North American standard IS-136 [24]

Figure 4 shows an interaction diagram for AMPS handoff that exemplifies the level of implementation detail. The power control algorithm is an asynchronous distributed uplink power control in which each user attempts to minimize its transmit power while maintaining a required SIR [25]. Long–scale fading is precomputed for the specified geography by overlaying correlated log–normal shadow fading [17] on power–law distance loss [20]. For call admission there are two variants of base station and channel assignment: use a channel from the base with the strongest pilot tone, or search a dynamic list of bases in order of decreasing pilot strength to find an acceptable channel. From [1] we have incorporated a family of dynamic channel allocation methods that employ interference measurements. Finally, WiPPET$_{\text{session}}$ provides system performance statistics in terms of call blocking, call dropping, average number of handoffs per call, etc., which are used to evaluate the performance of various radio resource management protocols.

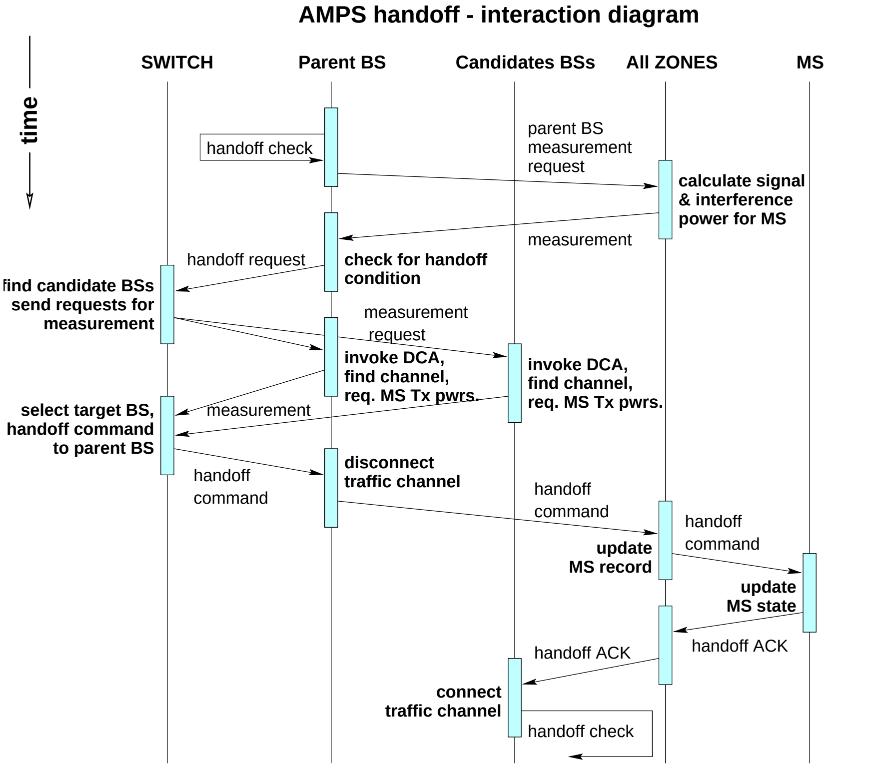

Figure 4. Interaction diagram for AMPS handoff.

4.2. WiPPET*: Packet level simulation*

In WiPPET$_{\text{packet}}$, a version of WiPPET for packet transmission, we simulate packet level traffic with distance loss, short scale (Rayleigh) and long scale (shadow) fading, mobility, and power control. In this work, we characterize packet loss performance for an elementary version of the packet level simulation with limited functionality and protocols. Related work also includes studying the interaction of TCP with radio link protocols [26,27]

Since the packet loss depends upon the bit or code symbol losses inside the packet, the SIR profile used for determining the packet loss must be finely sampled. In this case, we need to include short scale fading as well as long scale fading for the *TxRxFade* module shown in Figure 3. We have implemented models of short scale Rayleigh and Ricean fading [28] with i) wave-superposition method [29], ii) spectrum-shaping method [30] ; and also iii) a quantized channel state Markovian model [31]. WiPPET$_{\text{packet}}$ generates Rayleigh fading using the spectrum shaping method based on Clarke's model. The Rayleigh fading for

Tx-Rx pairs are modeled as mutually independent.

Once we obtain the detailed SIR samples at the receiver, a key issue is how to use this profile to determine whether the packet is corrupted. This is the function of the *receiver discriminator*. The determination of packet loss depends upon the modulation and channel coding scheme used in the packet transmission. Currently $\text{WIPPET}_{\text{packet}}$ employs a simple discriminator in the RX module of each receiver. This discriminator collects and compares each uplink sampled SIR with the threshold SIR Γ_{DROP}. If the sampled SIR is greater than Γ_{DROP}, a counter is incremented. If the counter exceeds a predefined number (*DROP_COUNT*), the packet is judged to be in error and the discriminator updates PER (packet error rate) statistics.

The verification of the receiver discriminator was done by comparing simulation results of $\text{WIPPET}_{\text{packet}}$ with the analytical results shown in [32] in terms of the PER vs. the fade margin for the CDPD (Cellular Digital Packet Data) uplink channel. CDPD is the data standard based on AMPS. In CDPD, 20 msec long packets are transmitted from the `MS` to the parent `BS` at 19.2 kbps. Packets are channel coded with the 8-error-correcting (63,47) Reed-Solomon code over $GF(2^6)$. The packet consists of 63 code symbols and each code symbol consists of 6 bits. The SIR value is assumed to be constant during a code symbol duration, a reasonable assumption for short scale fading under normal mobile speeds in the 800 MHz AMPS band. Thus the discriminator uses 63 samples of SIR corresponding to each code symbol to determine the packet loss. If more than 8 samples are below the threshold SIR then the packet is considered as in error. A detailed explanation appears in [32].

We performed a single cell experiment with only background noise and zero interference on the uplink radio channel. A single `MS` transmitted packets to its parent `BS`. The discriminator at the parent `BS` evaluates the 63 SIR samples per packet to decide whether the packet is in error. The integrity of $\text{WIPPET}_{\text{packet}}$ was verified by comparing the results of the simulations with previously published analytical results [32]. Figure 5 shows PER versus fade margin, which is comparable to the results shown in [32]. The reader may also refer to [32] for the details of the analysis and interpretation of the results.

The packet level simulation in WIPPET are an initial step toward multilayer protocol simulation. Our current work is on using this packet level simulation at the physical layer, to examine MAC layer retransmission schemes and their interaction with transport layer protocols such as TCP or UDP [26].

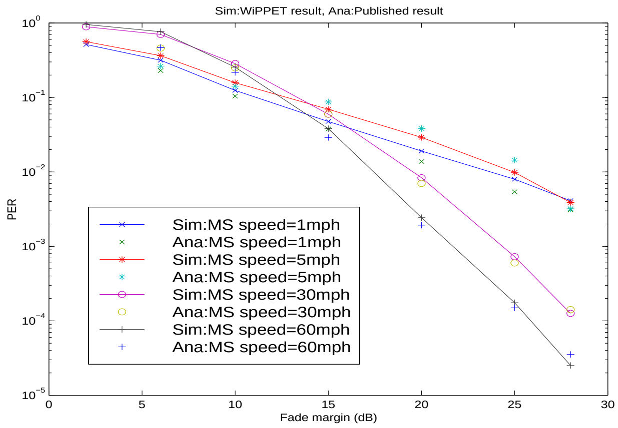

Figure 5. PER versus fade margin as evaluated by $\text{WiPPET}_{\text{packet}}$ and by analysis.

5. Simulator Performance

In this section, we describe how simulator design choices affect the speed of the simulation. For both $\text{WiPPET}_{\text{session}}$ and $\text{WiPPET}_{\text{packet}}$, frequent calculations of SIR throughout the network are a necessary input to the radio resource management algorithms and these consume most of the computational effort of the simulation. Parallelizing the interference calculations is the main consideration in parallelizing the entire simulation. There are two straightforward approaches to this parallelization.

Geographic Parallelization In this approach, the geography is divided into multiple zones. Each zone entity contains a fraction of the geography points together with the corresponding path loss sub-matrix. Each zone entity with sub-geography points and path loss sub-matrix is mapped on one processor.

Radio Channel Parallelization In radio channel parallelization, instead of dividing geography into multiple zones, the radio channels are divided into multiple zones. Each zone is then mapped on to a separate processor for execution. Each radio channel zone contains all the geographic locations and the path loss matrix, but only a fraction of the radio channels.

We now evaluate the parallel performance of geographic and channel zone partitioning and the sensitivity of that performance to simulation parameters. although our experiments are with $\text{WiPPET}_{\text{session}}$, similar conclusions can be drawn for $\text{WiPPET}_{\text{packet}}$.

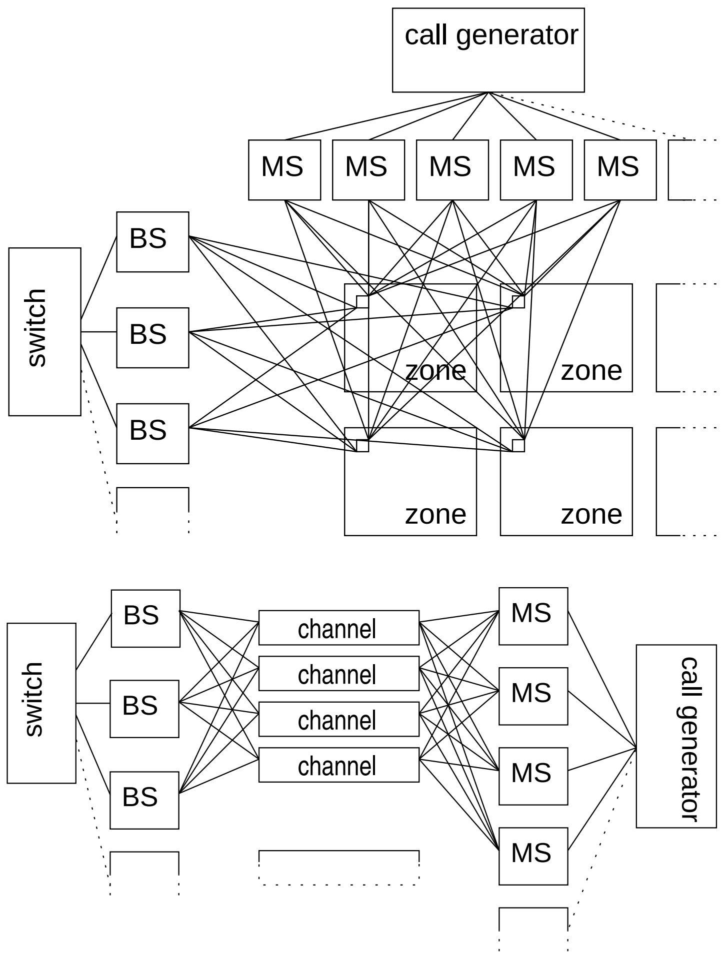

Figure 6. Two approaches to parallel simulation of mobile, multi-channel wireless networks: either by partitioning the geographical terrain among `zone` entities (top, WiPPET 0.3), or by partitioning the spectrum into orthogonal radio channels among `zone` entities (bottom, WiPPET 0.4). Hybrid approaches are also possible. Lines between entities denote the principal event flows.

In our experiments, the geography is a synthetic Manhattan-style urban environment of size 12×12 km (24×24 city blocks). Mobile trajectories are modeled by a biased random walk. Radio propagation is modeled with a spatial stochastic process of correlated log-normal fluctuations superimposed on distance dependent path loss factors characteristic for this category of urban environments. The study area contains 48 base stations sharing 40 radio channels with reuse factor one. Call admission and handoff are based on RSSI; power control is based

parameter [units]	**scenario** A	B	C	D	E	F	G
traffic [calls/s]	2	6	2	2	6	2	2
resolution [m]	10	10	2	10	10	10	10
PWR period [s]	–	–	–	0.5	0.5	0.1	0.5
speed [m/s]	25	25	25	25	25	25	2.2

Table 1
Seven experimental scenarios to evaluate geographic and channel zone partitioning.

on SIR. Call holding time is exponentially distributed with mean 120 seconds. Total simulated time in each case has been 10,000 seconds. For speed comparisons we used the same TED/GTW simulator both for the one- and multi-processor simulations. Table 1 specifies the Poisson call arrival rate, the geographic resolution, the signal checking period, and the travel speed for seven scenarios that were tested with geographic and channel partitioning on up to 8 processors of a multiprocessor Sun Enterprise 4000 server. Signal quality checks for power adjustment and potentially to initiate handoff are executed asynchronously and independently for each `MS` with the PWR period specified in Table 1. Scenarios A, B, and C simulate only mobility and call admission. A is the nominal case among the mobility scenarios, B has 3× offered traffic, and C has 5× finer geographic resolution. Scenarios D, E, F, and G simulate mobility, call admission, power control and handoff. D is the nominal case among the interference scenarios, E has 3× offered traffic, F has 5× more frequent power adjustments, and G has 11× slower moving subscribers. In the nominal interference scenario (D), an average of 240 calls are active at any moment.

Model partitioning depends on the number of processors used. The mapping strategy has the number of processors equal to the number of zones, with one zone mapped to each processor. An equal number of mobiles are mapped to each processor, and similarly for bases. `Callgen` and `switch`, of which there are one each, are mapped to different processors.

The performance of the simulator under the specified scenarios is given in the two by two panel layout of Figure 7. In Figure 7, the left column refers to geographic zone partitioning, the right column refers to channel zone partitioning, the top row refers to mobility scenarios (A, B, C), and the bottom row refers to interference scenarios (D, E, F, G).

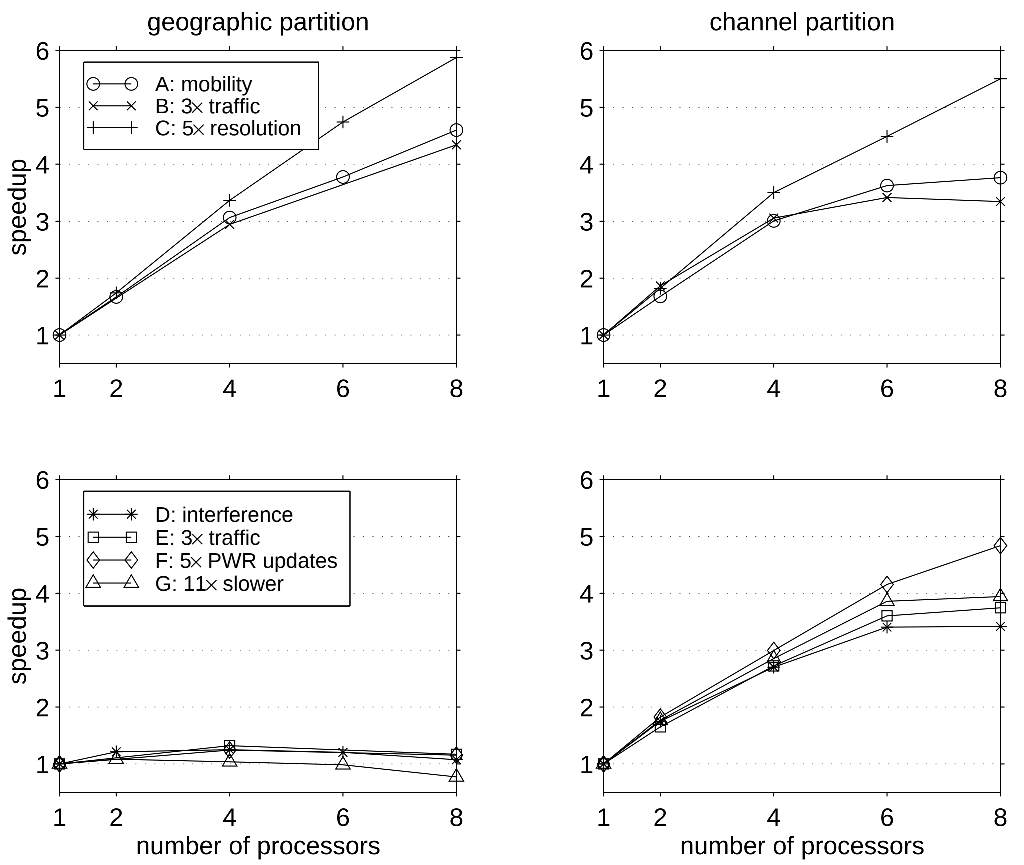

Figure 7. The speedup for N processors is defined as the ratio of computing time on one processor to the computing time on N processors.

The results in Figure 7 show a strong speedup in some but not all of the scenarios. It is interesting to analyze how this happens. First, there is a considerable concurrency in the system itself — a prerequisite for any speedup at all. From the simulation standpoint, however, that is not enough; we must also consider the costs of partitioning the problem into communicating sub-models, and the mechanisms that force synchronization between sub-models. It is helpful to think of each scenario as a blend of pure mobility activities and pure interference calculation activities. In the following, we examine partitioning costs and synchronization mechanisms for the four combinations of activity and partitioning method.

- **Mobility/geographic** Mobility is an autonomous activity of each mobile that is managed by its `zone`. Partitioning introduces zone transfer events.

Synchronization is forced when an MS transfers from one zone to another by moving; however, zone transfers are relatively rare events.

- **Mobility/channel** In channel partitioning, zone transfers are also introduced and synchronization is forced when an `MS` transfers from one zone to another by changing channel. In both channel and geographic partitioning, zone transfers happens at most a handful of times per call. The low cost and rarity of synchronization required for pure mobility is supported by the strong speedup curves in the top panels of Figures 7. In those panels, note especially scenario C which has the largest fraction of mobility effort among the seven scenarios.
- **Interference/channel** For orthogonal radio channels, interference in one channel is completely independent of interference on other channels. Zone transfer is the only event introduced by partitioning. Here again, synchronization is forced only when the MS transfers from one zone to another by changing channel. The rare synchronization required with orthogonal radio channels is supported by the strong speedup curves in the lower right panel of Figures 7.
- **Interference/geographic** Within a radio channel, the cost of geographic partitioning is that for each of the hundreds of interference requests that are processed, a wave of request/response events carrying portions of the sum in the denominator of the SIR measurement of equation (3) ripples between the current BS and the entire array of zones. The benefit is that the processing time for a portion of the sum is less than the processing time for the whole sum. However, the number of events needed to propagate interference calculations from each base station to all zones is simply too large; see [7] for details.

Preliminary experiments in reducing event counts, and in artificially increasing the work of interference calculation show that both mechanisms serve to improve the speedup seen in geographical parallelization. There is also scope to reduce the parallelization costs by simple programming techniques that reduce event usage, such as event aggregation or subscriber notifications. When a system is partitioned along the spectral direction into zones containing a subgroup of radio channels each, asynchronous computations of mobility and interference both contribute towards high exploitation of inherent concurrency that promotes fast parallel execution. However, channel parallelization does not accommodate systems such as CDMA that lack an orthogonal decomposition. In certain cases it may be best to have only one instance of a zone entity. In contrast, signal quality calculations in geographically partitioned zones force substantial communication between sub-models that can worsen parallel performance considerably.

6. Future Directions

In this paper we have summarized the parallel simulation of wireless networks that is ongoing at WINLAB. Future work on simulation of third generation (3G) wireless systems includes simulating interference cancelation in wideband code-division-multiple-access (W-CDMA). This is being accomplished using the next generation modeling framework SSF (Scalable Simulation Framework) [33] that was designed on the basis of experiences with TeD. The specific simulator that is currently under development, $\text{WiPPET}_{\text{signal}}$, is a generic tool capable of multicell simulations at the level of bits, chips and waveforms.

Acknowledgements

We acknowledge Kalyan Perumalla for his invaluable assistance in TeD implementations of WiPPET.

Author Biographies

OWEN E. KELLY is a Postdoctoral Fellow at WINLAB, Rutgers University. He received M.S. and Ph.D. degrees (1994,1996) from Rice University, Houston, TX in electrical engineering, and the B.A.Sc. degree (1989) from the University of Waterloo, Canada. His research emphasizes statistical modeling and decision problems in the field of communication and currently includes fitting discrete stochastic channel models, and simulation of large wireless networks. Since December 1998, he has been with Algorex Inc. in Iselin, NJ.

JIE LAI is a doctoral candidate at the Wireless Information Network Laboratory (WINLAB), Rutgers University. His research interests include wireless system performance analysis, fading channel modeling, channel coding, signal processing, and simulation of wireless networks.

NARAYAN MANDAYAM received the B.Tech (Hons.) degree in 1989 from the Indian Institute of Technology, Kharagpur, and the M.S. and Ph.D. degrees in 1991 and 1994 from Rice University, Houston, TX, all in electrical engineering. From 1994 to 1996, he was a Research Associate at the Wireless Information Network Laboratory (WINLAB), Dept. of Electrical & Computer Engineering, Rutgers University. In September 1996, he joined the ECE department at Rutgers as an Assistant Professor. His research interests are in communication theory, spread spectrum, wireless system modeling and performance, multi-access protocols, multiuser detection and multimedia wireless communications.

Dr. Mandayam is a recipient of the Institute Silver Medal from the Indian Institute of Technology, Kharagpur. He also received the National Science Foundation CAREER Award in 1998.

ANDREW T. OGIELSKI is a Research Professor with joint appointment at DIMACS and Electrical and Computer Engineering/WINLAB, Rutgers University since 1997. He held academic positions in Europe and in USA from 1976 to 1981, was a Member of Technical Staff in Physics Research and in Mathematics Research at Bell Laboratories from 1982 to 1989, then joined Bellcore where he was a Director of Parallel Computing and Algorithms Research, and later of Internet Communications Research until 1997. He received the M.Sc. in Physics in 1972 and the Ph.D. in Theoretical Physics in 1976, both from The University of Wroclaw, Poland. His current research focuses on modeling techniques and parallel simulations of large telecommunication networks, both mobile wireless and wireline, parallel simulation software design, broadband network traffic monitoring and analysis, and data management and simulations in neurophysiology.

JIGNESH PANCHAL received the B.E. degree in 1996 from the L.D. College of Engineering, Gujarat University (gold medal). He is an M.S. student at WINLAB. His current research interest is simulation and evaluation of wireless network resource management algorithms. Since July 1998, he has been with Lucent Technologies in Whippany, NJ.

ROY YATES received the B.S.E. degree in 1983 from Princeton University, and the S.M. and Ph.D. degrees in 1986 and 1990 from M.I.T., all in Electrical Engineering. Since 1990, he has been with the Wireless Information Networks Laboratory (WINLAB) and the ECE department at Rutgers University where he is currently an associate professor. His research interests include power control, interference suppression, and media access protocols for wireless communications systems. He is a co-author (with David Goodman) of the text *Probability and Stochastic Processes: A Friendly Introduction for Electrical and Computer Engineers*, published by John Wiley and Sons.

References

[1] K.G. Chen. Integrated dynamic radio resource management of wireless communication systems. Masters thesis, Rutgers, The State University of New Jersey, New Brunswick, NJ, May 1996.

[2] C.N. Chuah and R. Yates. Evaluation of a minimum power handoff algorithm. In *Sixth IEEE International Symposium on Personal, Indoor and Mobile Radio Communications PIMRC'95*, pages 623–627, 1995.

[3] C.N. Chuah, R. Yates, and D. Goodman. Integrated dynamic radio resource managemant. In *Proc. IEEE Vehicular Technology Conference, VTC-95*, pages 584–588, 1995.

[4] R. Yates, S. Gupta, C. Rose, and S. Sohn. Soft dropping power control. In *Proceedings of The IEEE Vehicular Technology Conference VTC 97*, May 1997.

[5] R. Yates and C.Y. Huang. Integrated power control and base station assignment. *IEEE Transactions on Vehicular Technology*, 44(3):638–644, August 1995.

[6] J. Panchal, O. Kelly, J. Lai, N. Mandayam, A. T. Ogielski, and R. Yates. Parallel simulations of wireless networks with TED: Radio propagation, mobility and protocols. *Performance Evaluation Review*, 25(4):30–39, Mar 1998.

[7] J. Panchal, O. Kelly, J. Lai, N. Mandayam, A. T. Ogielski, and R. Yates. WiPPET, a virtual testbed for parallel simulations of wireless networks. In *Proceedings of the 1998 Parallel and Discrete Simulations Conference PADS 98*, pages 162–169, Banff, Canada, May 1998.

[8] M. Liljenstam and R. Ayani. A model for parallel simulation of mobile telecommunication systems. In *Proceedings of the International Workshop on Modelling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS)*, 1996.

[9] W. Liu, C-C. Chiang, H-K. Wu, V. Jha, M. Gerla, and R. Bagrodia. Parallel simulation environment for mobile wireless networks. In *Proceedings of the 1996 Winter Simulation Conference, WSC'96*, pages 605–612, Coronado, CA, 1996.

[10] C. D. Carothers, R. M. Fujimoto, Y-B. Lin, and P. England. Distributed simulation of large-scale pcs networks. In *Proceedings of the 1994 MASCOTS Conference*, pages 2–6, January 1994.

[11] D. M. Nicol, editor. *Special Issue on the Telecommunications Description Language*, volume 25 of *Performance Evaluation Review.* ACM SIGMETRICS, March 1998.

[12] K. Perumalla, A. Ogielski, and R. Fujimoto. MetaTeD: A meta language for modeling telecommunication networks. TR GIT-CC-96-32, College of Computing, Georgia Institute of Technology, 1996.

[13] K. Perumalla and R. Fujimoto. A C++ instance of TeD. TR GIT-CC-96-33, College of Computing, Georgia Institute of Technology, 1996.

[14] K. S. Perumalla, M. Andrews, and S. Bhatt. A virtual PNNI network testbed. In *Winter Simulation Conference Proceedings*, Dec 1997.

[15] S. Das, R. M. Fujimoto, K. Panesar, D. Allison, and M. Hybinette. GTW: A time warp system for shared memory multiprocessors. In *Winter Simulation Conference Proceedings*, pages 1332–1339, 1994.

[16] B. Sklar. Rayleigh fading channels in mobile digital communication systems. *IEEE Communications Magazine*, 35(9):136–146, 1997.

[17] M. Gudmundson. Correlation model for shadow fading in mobile radio systems. *Electron. Lett.*, 27(23):2145–2146, 1991.

[18] S. Bhatt, R. Fujimoto, A. Ogielski, and K. Perumalla. Parallel Simulation Techniques for Large-Scale Networks. *IEEE Communications Magazine*, 36(8):42–47, August 1998.

[19] W. C. Jakes, editor. *Microwave Mobile Communications*. Wiley, 1974.

[20] Theodore S. Rappaport. *Wireless Communications*. Prentice Hall, 1996.

[21] J. Lai and N.B. Mandayam. Fade margins for minimum duration outages in lognormal shadow fading and rayleigh fading. In *Proceedings of Asilomar'97*, pages 609–613, Monterrey, CA, November 1997.

[22] R. Vijayan and J. M. Holtzman. A Model for Analyzing Handoff Algorithms. *IEEE Transactions on Vehicular Technology*, 42(3):351–356, August 1993.

[23] J. Panchal. Parallel simulator for wireless networks. Masters thesis, Rutgers, The State University of New Jersey, New Brunswick, NJ, July 1998. See http://www.winlab.rutgers.edu/∼ryates/thesis/panchal-ms.ps.

[24] D. Pandian. Channel allocation and power control in IS-136. Masters thesis, Rutgers, The State University of New Jersey, New Brunswick, NJ, January 1999. See http://www.winlab.rutgers.edu/∼ryates/thesis/pandian.ps.

[25] R. Yates. A framework for uplink power control in cellular radio systems. *IEEE J. Sel. Areas Commun.*, 13(7):1341–1346, Sep 1995.

[26] Y. Bai, A. Ogielski, and G. Wu. TCP - RLP coordination and interprotocol signaling for wireless internet. In *Proceedings of VTC'99*, To Appear, May 1999.

[27] Y. Bai, A. Ogielski, and G. Wu. TCP over IS-707. In *PIMRC'99*, Submitted, May 1999.

[28] R. H. Clarke. A statistical theory of mobile-radio reception. *Bell System Technical Journal*, 47:957–1000, 1968.

[29] W. C. Jakes, Jr. *Microwave Mobile Communications*. New York: John Wiley and Sons, 1974.

[30] J. I. Smith. A computer generated multipath fading simulation for mobile radio. *IEEE Trans. on Veh. Technol.*, VT-24(3):39–40, August 1975.

[31] F. Babich, O.E. Kelly, and G. Lombardi. A variable-order discrete model for the mobile propagation channel. In *Proc. 9th Tyrrhenian International Workshop on Digital Communications*, Lerici, Italy, Sept. 1997.

[32] J. Lai and N.B. Mandayam. Packet error rate for burst-error-correcting codes in rayleigh fading channels. In *Proceedings of VTC'98*, pages 1568–1572, Ottawa, Canada, May 1998.

[33] J. H. Cowie, D. M. Nicol, and A. Ogielski. Modeling the Global Internet. *Computer Science and Engineering*, pages 42–50, January/February 1999.