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

Jignesh Panchal, Owen Kelly, Jie Lai, Narayan Mandayam, Andrew T. Ogielski, Roy Yates

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

Abstract

We review the design, selected applications and parallel performance of WiPPET, a general parallel simulation testbed for various types of wireless networks. WiPPET has been written in TeD/C++, a new 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 levels. We illustrate their efficiency under two partitioning schemes with parallel performance data obtained using the Georgia Time Warp optimistic simulator. Finally, 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.

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 transmitter'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, 3, 4, 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.

WIPPET is a versatile simulator for wireless networks, which has been implemented in the TED/C++ framework [8, 9, 10, 11], and is used with the Georgia Time Warp (GTW) optimistic parallel simulator [12]. WIPPET is 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, especially in the 5 GHz U-NII bands.

In order to be useful in multiple research projects, a versatile 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, long–scale shadow fading). A versatile wireless simulator should support research studies at various levels and time-scales, and more importantly *across* levels and time-scales. Such a system facilitates the incorporation of

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

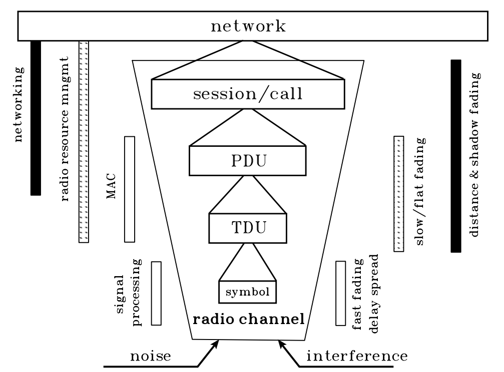

Figure 1: WIPPET provides support and "hooks" for model development at a variety of time scales and protocol layers.

interactions between layers in the design and optimization process.

2 Mobility, Propagation, and Protocols in WIPPET

2.1 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. A common model of mobility in microcellular systems is a biased random walk in a synthetic Manhattan–style urban environment with a 5–10 meter geographic resolution.

2.2 Propagation and Interference

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 [13, 14]. 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 gernal 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 is the channel impulse response. The corresponding received power at $\mathbf{b}$ is $P_{\mathrm{RX}}^{(i)}(t,\mathbf{b}) = E\left[R_i^2(t,\mathbf{b})\right]$. A measure of the received signal quality at position $\mathbf{b}$ for terminal i is *signal to interference ratio* (SIR), which is defined by the ratio

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

where η^2 is the additive noise power.

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 multistory office buildings). While the design of WIPPET allows to use 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.

2.2.1 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 [14]. 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, $E[\log g_s(\mathbf{x}, \mathbf{b}) \log g_s(\mathbf{x}', \mathbf{b})] = \sigma^2 \exp(-|\mathbf{x} - \mathbf{x}'|/d_0)$ [15].

2.2.2 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 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. It is important to note that the details of the fading model should be matched to the signal processing capabilities of the modeled receiver to obtain the meaningful statistics of data loss; i.e., further abstractions of fading may be made according to the level

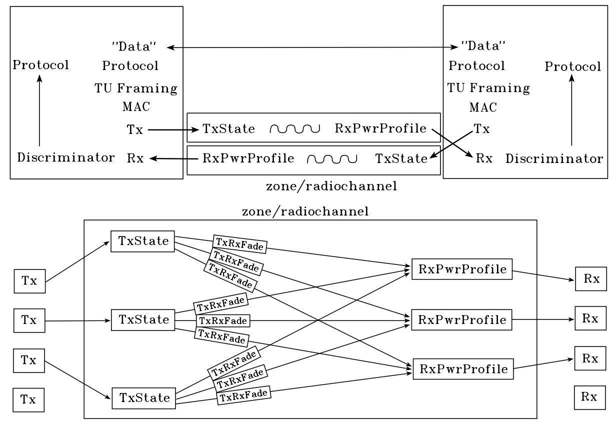

Figure 2: Top: Abstraction of the generic communicating packet radio entities. Bottom: The main role of a zone entity is to combine the contributions of *all* actively transmitting Tx to the received power profile for *every* Rx. Each time the power, or position, or the value of short–scale fading loss of a Tx changes, *all* active Rx power profiles in that radio channel must be updated.

in Figure 2 (top) that is of interest. For example, studies of radio resource management typically require averaged fading and interference, whereas for packet-level simulations one needs SIR values on a fine time scale during a packet 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 2 (bottom). 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 system being simulated, the cross–connection may be essentially complete, or relatively sparse.

2.3 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). 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. WIPPET's primary purpose is evaluation of protocols, and evaluation means performance comparison. Within TED, comparison is facilitated by the ability to substitute architectures of entities. A canonical comparative experiment requires running two nearly identical models, the only difference being that certain protocol processes are different. Unlike propagation and interference, which require considerable thought to cast into the event-passing paradigm of TED, protocols fall easily into place by defining an event for each message type of the protocol, and a TED process for each protocol, to evaluate these events.

2.4 Entity Design 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 radio resource management protocols at the connection level (session admission, transmitter power control and handoff). The entity modelling traffic, `callgen` generates calls (session arrivals), and the entity `zone` encapsulates geography and propagation. Figures 3 and 4 show the details of approaches to parallelization and the inheritance heirarchy in WIPPET. We refer the reader to [7] for details of entity design.

3 WIPPET: Session level simulation

WIPPET$_{\text{session}}$, a version of WIPPET that models session or call level traffic, has radio propagation model 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].

Figure 5 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 [17]. Long–scale fading is precomputed for the specified geography by overlaying correlated log–normal shadow fading [15] on power–law distance loss [14]. 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 provides system

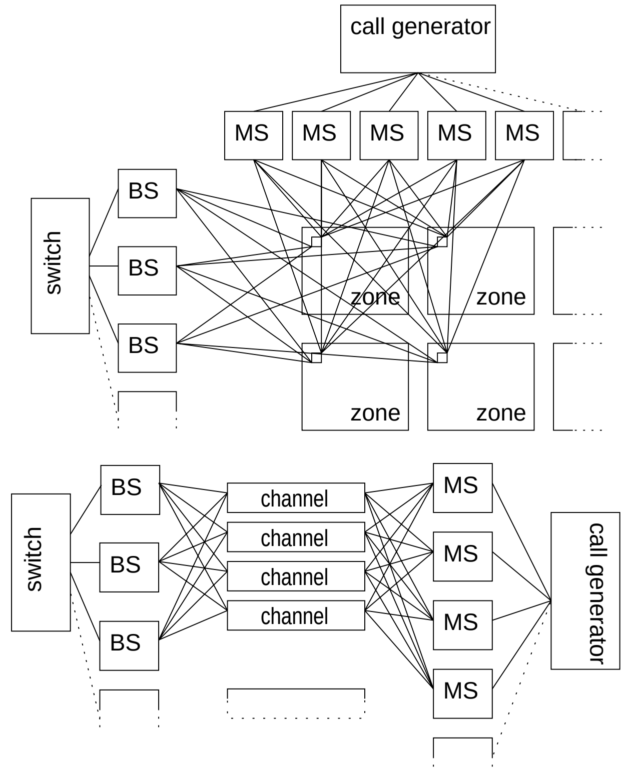

Figure 3: Two approaches to parallel simulation of mobile, multi-channel wireless networks: either by partitioning the geographical terrain among `zone` entities (top, WIPPET 0.3), or by partitioning the spectrum into orthogonal radio channels among `zone` entities (bottom, WIPPET 0.4). Hybrid approaches are also possible. Lines between entities denote the principal event flows.

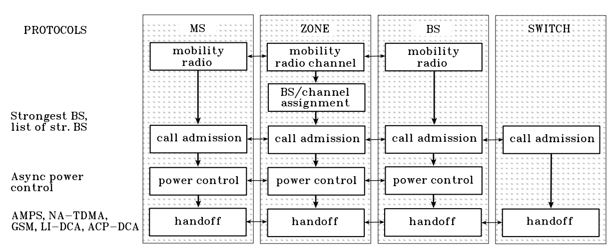

Figure 4: Entity inheritance hierarchy (vertical) and their compatibility for protocol substitutions (horizontal). We note that in the current version of TED it is computationally expensive to embed component entities in parent entities when they need to send or receive events through the parent's channels. Therefore, our design builds up each entity type by inheritance of architectures, with each level of inheritance adding appropriate processes and state fragments.

performance statistics in terms of call blocking, call dropping, average number of handoffs per call, etc., which are used to analyze the performance of various radio resource

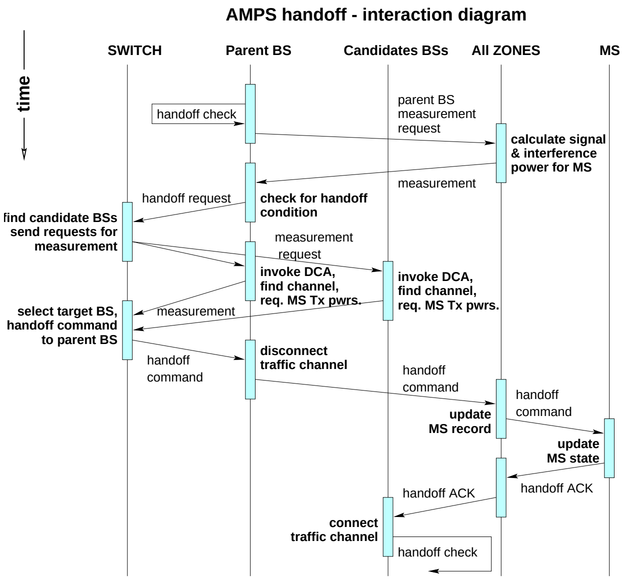

Figure 5: Interaction diagram for AMPS handoff.

management protocols.

Frequent calculations of SIR throughout the network are a necessry 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. We now describe two approaches to this parallelization.

3.1 Geographic Parallelization

In this approach, the geography is divided into multiple zones. Each zone entity contains a fraction of the geography points together with corresponding path loss sub-matrix. Each zone entity with sub-geography points and path loss sub-matrix is mapped on one processor. Equal sharing of the simulation load among processors for uniprocessor and 4–processor cases are shown in Figure 6. One can extend the same mapping approach to the N–processor simulation.

3.2 Radio Channel Parallelization

In radio channel parallelization, instead of dividing geography into multiple zones, the radio channels are divided into multiple zones and map the radio channel zones on different processors. Each radio channel zone contains all geography and path loss matrix, but only a fraction of the radio channels. The sharing of the simulation load among processors for uniprocessor and 4–processor cases are shown in Figure 7. This mapping extends easily to an N–processor simulation.

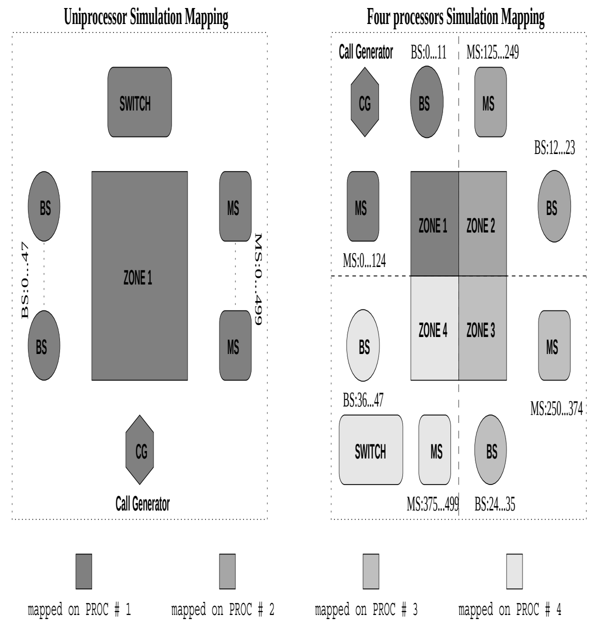

Figure 6: Load balancing in geography parallelization: (a)[left] uniprocessor simulation; (b)[right] 4 processor simulation.

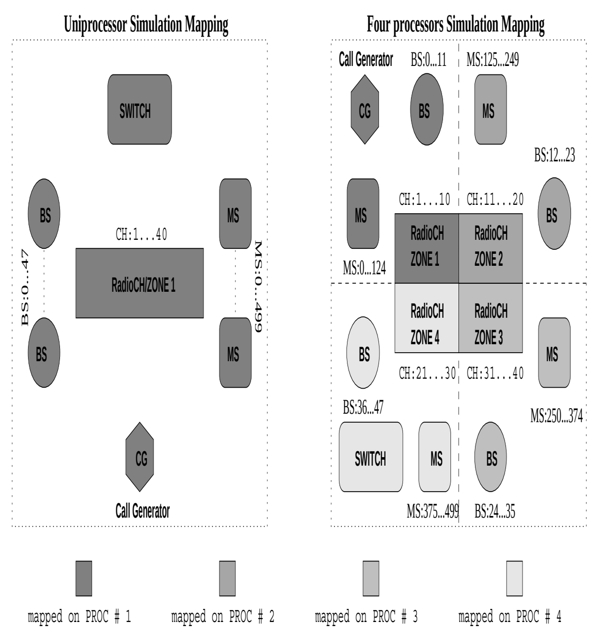

Figure 7: Load balancing in radio channel parallelization: (a)[left] uniprocessor simulation; (b)[right] 4 processor simulation.

3.3 Parallel Simulation Performance

First, we must define terminology used in the discussion of parallel simulation performance.

- The *speed up* of an N processor simulation is the ratio of the uniprocessor execution time to the execution time on N processors.
- The *net* events of N processor simulation is defined as total events required, excluding rollback events, to complete the simulation.
- The *excess events* of an N processor simulation is defined as difference between net events of the N processor simulation and net events of the uniprocessor simulation. Thus, as a function of the number of processors N,

$$\%\text{ excess events}(N) = 100 \times \frac{\text{excess events }(N)}{\text{net events(uniproc.)}} \quad (2)$$

We have performed experiments in a multi-processor environment to evaluate geography and radio channel parallelization.

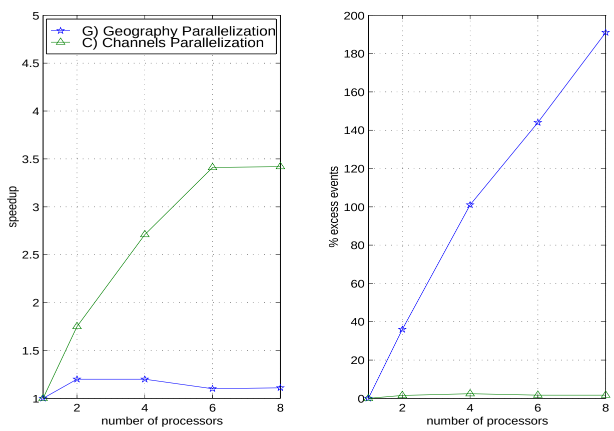

Figure 8: Experimental results: (a) [left] speed up; (b) [right] % excess events

The results in Figure 8(a) shows a significant speed up for channel parallelization but not for geographic parallelization. This conclusion is consistent with results in [18]. It is interesting to analyze how this happens. Within a radio channel, interference is a simultaneous phenomenon across the geographical extent of the network. Therefore, within one radio channel, geographic partitioning creates excess events necessary to carry partial results to a place where the interference in the channel is summed. For orthogonal radio channels, interference in one channel is completely independent of interference on other channels (local activity). A single zone does all the calculations of the interference on a single channel. This dramatically reduces the number of excess events associated with interference calculations. Partitioning also creates excess events that we call zone transfer events. These occur when mobiles cross boundaries, in the case of geographic parallelization, or when mobiles change channels, in the case of channel parallelization. Here again, synchronization is forced only when the MS transfers from one zone to another. Under either parallelization, these zone transfer events occur perhaps a few times during the lifetime of a call, unlike interfernce calculations which can occur hundreds of times per second. Consequently, in Figure 8(b) we see that the excess events due to zone transfer are rare. Hence, orthogonal radio channels result in the strong speedup curves shown in Figure 8(a).

parameter [units]	2D Macrocell
Number of BSs	48
Number of Channels	40
Geography Style	Manhattan
Geographical Area	$12km \times 12km$
Cell Radius [m]	1000
Antenna Height [m]	50
Distance Resolution [m]	5
Shadow Fading Standard Deviation [dB]	6
Shadow Fading Correlation Distance [m]	50
Noise Power [dBW]	-150
Maximum Power [dBW]	0
Mobile Speed [m/sec]	25
Probability of Going Straight	0.6
Time Between Handoff Checks [sec]	1.0
Time Between Power Updates [sec]	1.0
BS Assignment Algorithm	Strongest Pilot BS
Call Admission Channel-Allocation Algorithm	LI-DCA
Handoff Algorithm	RSSIBHO
Handoff Channel-Allocation Algorithm	LI-DCA
Target SIR [dB]	20
Drop SIR [dB]	14
Admission SIR [dB]	23

Table 1: WiPPET simulation parameters

3.4 WIPPET$_{\text{session}}$: verification

WIPPET$_{\text{session}}$ was validated by verifying that a TeD/C++ implementation of MADRAS [1] was in close agreement with previously published MADRAS results [1]. Specifically, Figures 9 and 10 show the agreement in call blocking probability and call dropping probability vs. the offered traffic respectively with reference curves from [1]. Table 1 shows the WIPPET$_{\text{session}}$ configuration parameters and algorithms used for comparison of the simulation results.

4 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.

During packet transmission, a mobile continuously moves under the control of the WIPPET$_{\text{packet}}$ *Mobility*

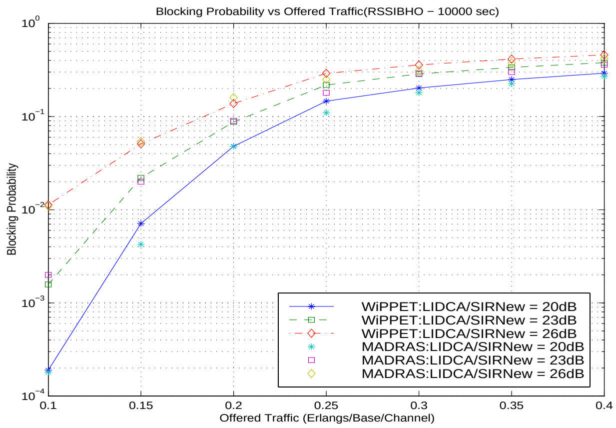

Figure 9: Call blocking rate as a function of offered traffic at different SIR admission thresholds using least interference dynamic channel allocation (LIDCA).

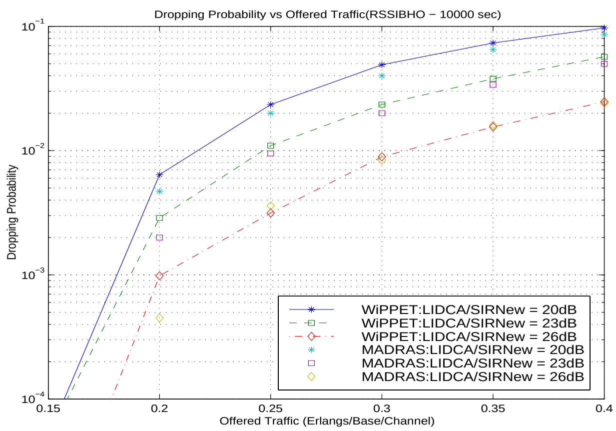

Figure 10: Call dropping rate as a function of offered traffic at different SIR admission thresholds using least interference dynamic channel allocation (LIDCA).

module. When a mobile gets to a new geographic point, the pathloss/shadow (long scale) fading value is updated. The uplink power is also updated according to the target SIR value, long scale fading value, and interference power. The objective of the power control module is to adjust the the required MS transmitter power to compensate the signal power loss due to distance/shadow fading and the average interference power from the co-channel interfering mobiles at the parent BS. The required transmitter power for the ith MS, denoted by a_i, is calculated at the parent BS as

$$a_i = \gamma_i \frac{\sum_{j \neq i} a_j g(\mathbf{x}_j, \mathbf{b}) + \eta^2}{g(\mathbf{x}_i, \mathbf{b})} \qquad (3)$$

where γ_i is the target SIR of the ith user.

Since the packet loss depends upon the bit or code symbol losses inside the packet, the SIR profile used for determining the packet loss has 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 2. We have implemented models of short scale Rayleigh and Ricean fading [19] with i) wave-superposition method [20], ii) spectrum-shaping method [21] ; and also iii) a quantized channel state Markovian model [22]. WIPPET$_{\text{packet}}$ has used Rayleigh fading using 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 *receiver discriminator*. The determination of packet loss depends upon the modulation and channel coding scheme used in the packet transmission. Currently 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 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 the WIPPET$_{\text{packet}}$ with the analytical results shown in [23] 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 [23].

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 the simulator testbed (WiPPET-packet level) was verified by comparing the results of the simulations with previously published analytical results [23]. Figure 11 shows PER versus fade margin, which is comparable to the results shown in [23]. The reader

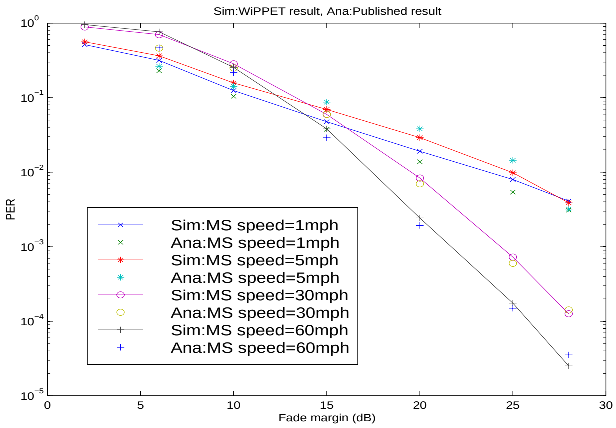

Figure 11: PER versus fade margin as evaluated by WIPPET$_{\text{session}}$ and by analysis.

may also refer to [23] for the details of the analysis and interpretation of the results.

The packet level simulation of WIPPET are the initial step toward multi-layer protocal simulation. We will use this packet level simulation at the physical layer, to examine MAC layer retransmission schemes and transport layer protocols such as TCP or UDP.

Acknowledgments

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

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*, Banff, Canada, May 1998.

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

[9] 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.

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

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

[12] 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.

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

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

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

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

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

[18] 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'96)*, February 1996.

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

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

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

[22] 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.

[23] J. Lai and N.B. Mandayam. Packet error rate for burst-error-correcting codes in rayleigh fading channels. In *Proceeding VTC'98*, May 1998.