

Network-Scale Emulation of General Wireless Channels

Xiaohui Wang, Kevin Borries, Eric Anderson, and Peter Steenkiste
Carnegie Mellon University
Pittsburgh, PA

Abstract—This paper presents a framework for signal-level emulation of propagation effects over generalized fading channels at the scale of entire networks. Network emulation enables research into network-scale systems – which would otherwise be limited to low-fidelity network simulators and one-off field experiments – to use real radio hardware and realistic channel models. Our hardware and software architecture goes beyond previous work in that it supports real-time emulation of a very general and parametric class of channels, which includes vehicular (broadband mobile-to-mobile) and indoor channels in addition to classical stationary-to-mobile and stationary-to-stationary channels.

I. INTRODUCTION

Science and engineering rely heavily on the ability to perform controlled, repeatable experiments. Channel simulation and emulation tools are meant to bring some of this ability to wireless research. Traditionally, wireless *communication* researchers have used high-fidelity small-scale models, while wireless *networking* researchers have considered larger-scale systems, but with much less detailed channel models. Network emulation bridges that gap by providing real-time emulated channels between a reasonable number of interacting nodes. Prior work simulates Rayleigh and Rician fading at stationary-to-pedestrian speeds, but cannot support more general wireless channels [1], [2]. This paper presents a software and hardware architecture for emulating a broader class of channel conditions, including indoor channels and mobile-to-mobile channels at vehicular speeds, in real time.

Channel response can be modeled in the time domain using a tapped delay line model, illustrated in Figure 1. Each tap effectively represents a resolvable propagation path, and the evolution of tap weights provides a statistical approximation of Doppler spreading and non-resolvable multipath effects. Time-varying tap weights must be generated to reflect large scale attenuation and small scale fading characteristics. The delay between taps and their relative magnitude determine the *frequency selectivity* of the channel and the degree of *inter-symbol interference* experienced. The rate at which tap weights change and the patterns of those changes determine *coherence time*, the (simulated) *Doppler spreading*, and higher-order statistics such as the *average fade duration*. This paper discusses the design and implementation of a flexible fading

simulator that is capable of producing a wide variety of weight sequences at each tap in real time.




Fig. 1. Tapped delay line model of an emulated signal path.

Two example channel definitions from GSM 3GPP [3] are shown in Figure 2. The rural area channel (RAx) is defined by four taps with a maximum excess delay of 0.5 μ s; the first has a Jakes Doppler spectrum with a line-of-sight (LOS) component and the remaining are classical (Jakes). The hilly terrain (HTx) channel is defined by twelve taps with a maximum excess delay of 20 μ s; all taps have a classical Doppler spectrum. Reference channels for wideband mobile-to-mobile communication are not similarly standardized; however theoretical models and recent measurements suggest an Akki spectrum like those shown in Figure 5 and terrain-dependent delay spreads ranging from 0.3 to 5 μ s [4]–[8].

The remainder of this paper presents our architecture for generating appropriate tap weights and a discussion of the resources required to effectively implement various channel models.




Fig. 2. Tapped delay line models for reference channels from 3GPP TS 05.05. 4-tap RAx in solid blue, 12-tap HTx in dotted black.

This research was funded in part by NSF under award number CNS-0434824, and in part by the Air Force Research Laboratory under award number FA89750.10.1.0232.

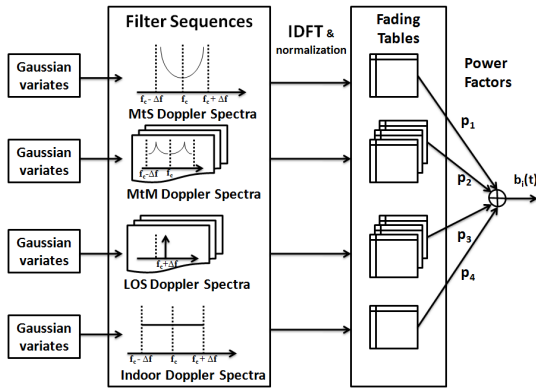


Fig. 3. Tap weight generation process.

II. TAP FADING SPECTRA SIMULATION

Viewed in the frequency domain, the important properties of a tap weight sequence are captured by the Doppler power spectral density (Doppler spectrum) it produces. We therefore describe the tap weight generation process in terms of the *fading spectrum* at each tap. The architecture described below can produce the dramatically varying tap fading spectra associated with a wide range of wireless channels.

A. Principle of Operation

For each tap, weights with the desired characteristics are computed as illustrated in Figure 3. The desired spectrum is defined as a weighted sum of simpler spectra and is produced at run time by combining samples drawn from its components. An attractive feature of this design is that a modest number of component spectra suffices to generate a broad range of channels (including all the models discussed in this paper), and the computational work of generating samples for those spectra can be done off-line. This leaves only light-weight sampling and summing operations to be done in real time. This process is described in more detail below:

For each component Doppler spectrum, a frequency-domain channel response is generated as a sequence of zero-mean i.i.d. Gaussian random numbers. Each sequence is filtered with the appropriate Doppler spectrum [9]. Each spectrum represents one type of Doppler shifting or spreading of the propagated signal, for example the classic mobile-to-stationary (*MtS*) channel corresponds to the U-shaped Jakes Doppler spectrum. After filtering, each sequence is then converted to a time domain sequence, normalized, and stored in a lookup table which we refer to as a *fading table*. Each fading table represents not just a single Doppler spectrum but a *family* of spectra, because different levels of Doppler shifting can be produced by iterating through the same table in larger or smaller steps. In addition, many uncorrelated sequences can be drawn from the same table using different starting indices.¹ Therefore, a single fading table can represent *all* Jakes spectra,

¹Multiple tables may optionally be used if the possibility of a spaced-time correlation must be avoided.

while the Akki spectra require a separate table for each value of a speed ratio a , but not for each f_{m_1} in Eq. (3).

For any given tap i , at each time t , the tap weight $b_i(t)$ is produced by taking the “next” value from every fading table, and combining these values using the (possibly 0) weights $p_j(t)$. Therefore, each tap’s fading spectrum can be any of the component spectra, or a combined spectrum representing multiple fading effects. For example, a Rayleigh-distributed classical Doppler spectrum combined with a line-of-sight factor produces a Rician distribution. Several important component Doppler spectra are described in the following sections.

B. The Mobile-to-X Channel Models

The mobile-to-stationary (*MtS*) scattering channel, as modeled by [10], assumes a ring of scatters around the receiver producing a uniformly distributed angle of arrival. The Doppler spectrum $S(f)$ of *MtS* fading channels can be described as Eq. (1):

$$S(f) = \begin{cases} \frac{1}{4\pi f_m} \frac{1}{\sqrt{1 - \frac{(f-f_c)^2}{f_m^2}}}, & |f - f_c| \leq f_m \\ 0, & \text{otherwise} \end{cases} \quad (1)$$

where f_c is the center frequency and f_m is the maximum Doppler shift, defined as:

$$f_m = \frac{V}{c} f_c \quad (2)$$

In Eq. (2), V is the speed of the mobile and c is the speed of light. A filter sequence generated for f_{max} can be used for any $f_m \in (0, f_{max}]$ by stepping through the sequence at a lower speed. f_m and f_{max} correspond directly to V and some V_{max} , which is the greatest (physical) speed for which the sequence can be used.

The mobile-to-mobile (*MtM*) scattering channel, as modeled by [4], considers independent double-ring scattering (that is, at both transmitter and receiver). Let the speeds of the first and second mobile be V_1 and V_2 (without loss of generality $V_1 \geq V_2$). The corresponding maximum Doppler shifts are f_{m_1} and f_{m_2} respectively. The Doppler spectrum $S(f)$ of *MtM* fading channel is described by Eq. (3):

$$S(f) = \begin{cases} \frac{1}{\pi^2 f_{m_1} \sqrt{a}} \text{Re} \left[K \left[\frac{(1+a)}{2\sqrt{a}} \sqrt{1 - \left(\frac{f-f_c}{(1+a)f_{m_1}} \right)^2} \right] \right], & \text{when } |f - f_c| \leq (1+a)f_{m_1} \\ 0, & \text{otherwise} \end{cases} \quad (3)$$

where $K(\cdot)$ is the complete elliptic integral of the first kind, and a is the speed ratio, defined as: $a = V_2/V_1$, ($0 < a \leq 1$).

For any given a , a single fading table can be generated to produce tap sequences for any V_1 up to a predefined maximum, as with the *MtS* channel. At any time t , $V_1(t)$ determines the rate at which new values are read from the table.

C. The Indoor Channel

Another Doppler spectrum of interest is an indoor scattering spectrum. Indoor channels are generally distinguished by relatively low velocities and a rich three-dimensional scattering environment. Both measurement studies and theoretical models suggest that this combination leads to a “flat” Doppler spectrum [11]–[13]. The flat spectrum is not a special case of




Fig. 4. System Architecture.

the mobile-to-mobile channel described above, but can still be simulated using the general procedure from § II-A.

D. General Wireless Channel Model

As described in § II-A, multiple fading models can be combined to represent an aggregate fading spectrum. This allows for combinations beyond the specific models shown in this paper; for example, [14] improves the *MtM* scattering model by combining multiple Doppler spectra with an LOS component. Near-arbitrary user-defined component spectra can be included, as long as they are known in advance.

III. FADING TAP IMPLEMENTATION AND VALIDATION

This section discusses the implementation and validation of the tap fading spectrum generation process described in the preceding section. Our prototype is implemented on the CMU Wireless Network Emulator [2] for the 2.4 GHz ISM band. The architecture is shown in Figure 4: A control computer generates tap configuration values in real time using the process described here, and the weight and delay values are sent to the signal processing FPGA. Both the mobile-to-x and indoor fading models are implemented on this platform. Discussion will be focused on the mobile-to-x case.

A. Generating Fading Tables

Tap weights are generated using the process in Figure 3. A long sequence of frequency sampling points is created to represent one Doppler spectrum. We choose a maximum speed $V_{max} = 70$ m/s (≈ 156 mile/h). Relative to a center frequency f_c of 2.437 GHz, the maximum Doppler shift is $f_{max} = \frac{V_{max}}{c} f_c \approx 600$ Hz. This requires a sampling frequency $f_s \geq 1.2$ kHz to smoothly simulate spectra for Doppler shifts up to f_{max} . We select f_s to be 1.2 kHz. As a result, the time between data points in the time-domain fading tables is $\frac{1}{f_s} \approx 0.42$ μ s. The more sampling points used for a given frequency range, the longer the resulting time-domain fading table can be. We choose a sequence length of 65536, which implies a time-domain fading table covering ≈ 30 seconds at V_{max} . This means that the small-scale fading process is periodic, but we consider the period to be acceptably long.

The prototype generates sequences for the following component Doppler spectra:




Fig. 5. Akki (*MtM*) filter spectra for speed ratios $a = 0.5$ and $a = 1$.




Fig. 6. *MtM* channel time domain samples.

- A normalized Jakes' spectrum (Eq. (1)) to represent the *MtS* Doppler spectrum.
- A set of Akki spectra (Eq. (3)) for the *MtM* Doppler spectra for different (binned) speed ratios. The $K(\cdot)$ values in Eq. (3) are static, and thus generated off-line.
- A normalized "flat" indoor Doppler spectrum.
- Two LOS spectra with a spike at $f_{LOS} = \pm \frac{V}{c} f_c$.

Figure 5 shows two examples of the Akki spectrum filter generated for the *MtM* Doppler spectrum, with speed ratios $a = 0.5$ and 1, with no LOS component. Figure 6 shows a corresponding sequence of time domain fading samples generated for $a = 1$.

B. Run Time Simulation

At run time, the control software selects fading tables and assigns power factors to each fading table. For example, when both nodes are moving, an *MtM* fading table is selected: the current speed ratio $a = \frac{V_2}{V_1}$ determines which *MtM* Doppler spectrum is applied. The power factors can be either configured by the user, or set to default values for typical channels. For example, when there is a direct LOS between sender and receiver, the LOS components will be included with higher power factors. The run-time software infers channel requirements from higher-layer environment models, switching fading tables and changing power factors when necessary.

C. Verification

This section presents a qualitative statistical evaluation of the samples produced by our prototype implementation. Figure 7 shows an averaged frequency domain view of a




Fig. 7. Generated MtM channel spectrum: $a = 0.5$, no LOS.




Fig. 8. Q-Q plot: MtM channel tap weights vs. Gaussian.

large number of (time domain) fading sequences for the MtM channel with $a = 0.5$. The result shows random variation around the filter spectrum shape.

We also evaluate the distribution of the filtered samples. The result of a linear transformation of Gaussian random vectors should be Gaussian random vectors. Because the filtering process and IDFT in Figure 3 are both linear, the resulting sequences are expected to be Gaussian variables. Figure 8 shows a Q-Q plot of normalized fading table values against a zero-mean Gaussian distribution with standard deviation of 1.0. The straight-line fit suggests that the Gaussian distribution is maintained.

IV. SIMULATION RESOURCE COSTS

An n -node (or n -antenna) network involves $O(n^2)$ channels, and both signal propagation through each channel and the time-varying properties of the channels themselves must be modeled in real time. Resource efficiency is therefore a fundamental concern for network-scale emulation.

A. System Model

The architecture considered here is based on the CMU Wireless Network Emulator [2], and is shown in Figure 4. Signal processing, which must be fast and synchronous, is performed in a central FPGA. Channel and higher-level modeling, which are more complicated and less timing-sensitive, are performed on a general-purpose computer. The two systems are linked by a serial channel. In this architecture, FPGA resources limit operations on the signal path, host CPU resources limit channel modeling, and the bandwidth of the host-FPGA link

limits updates to the signal path. Timing precision on the host is a cross-cutting issue. The techniques discussed in this paper are applicable to other architectures, but we will use this as a reference point in considering the resource costs of various simulation options.

Design space variables having a significant impact on resource use are shown in Table I. Any given resource limit will allow trade-offs over these variables; specifics are discussed below.

Variable	Design Points
Bandwidth	5 MHz (min. IMT-Advanced channel) 22 MHz (one 802.11 channel) 84 MHz (2.4 GHz ISM or upper UHF TV)
Channels	2 (single bi-directional link) 18 (3x3 MIMO, bi-directional) 380 (20 node SISO network) 3420 (20 node 3x3 MIMO)
Taps per channel	1 (minimum) 2 (TS 05.05 indoor model) 6 (TS 05.05 rural and reduced models) 12 (All TS 05.05 models)
Maximum excess delay	$\approx 0.4 \mu\text{s}$ (802.11 card equalizer) $0.4\text{-}0.6 \mu\text{s}$ (rural, indoor models) $5.0 \mu\text{s}$ (urban model) $20.0 \mu\text{s}$ (hilly terrain model)
Minimum coherence time	$23 \mu\text{s}$ (v2v worst-case 300 km/h, 5.9 GHz) $50 \mu\text{s}$ (v2v, worst-case 300 km/h, 2.4 GHz) $260 \mu\text{s}$ (v2v worst measured, 5.9 GHz [8])

TABLE I
WIRELESS NETWORK EMULATOR DESIGN SPACE VARIABLES

B. Resource Limits and Feasible Ranges

For the sake of simplicity, we consider only homogeneous designs in which every node, channel, and path gets the same level of resources as every other. A well-chosen inhomogeneous configuration might be more useful, however the homogeneous case serves to illustrate general limits.

Some preliminary definitions:

n	number of nodes	(4)
A	number of antennas per node	(5)
p	number of taps per channel	(6)
τ	maximum excess delay	(7)
$N = A^2 n(n-1)p$	total number of taps (paths)	(8)
$f_m = \frac{v}{\lambda}$	maximum Doppler spread	(9)
$T_C \approx \frac{.25}{f_m}$	minimum coherence time	(10)
B	bandwidth	(11)
$f_{samp} \geq 2B$	sampling rate	(12)
d	dynamic range (dB)	(13)
$s \geq \frac{d}{6.02}$	sample size (bits)	(14)

1) *FPGA Multipliers Blocks*: The main limiting operation in the tapped delay line channel model is the signal multiplication for each tap. These are implemented using dedicated multiplier components in the FPGA fabric. For reference consider a recent FPGA, the Virtex-6 XC6VS475T, contains

$\approx 2,000$ multipliers with a maximum rate of 600 MHz [15]. This gives a limit of $f_{mult} \lesssim 1.2 * 10^{12}$ multiplications per second. Each tap requires one multiplication for every sample processed. The number of distinct paths is therefore bounded by Eq. (15):²

$$N \leq \frac{f_{mult}}{f_{samp}} \leq \frac{f_{mult}}{2B} \quad (15)$$

If $B = 100$ MHz (\approx the 2.4 GHz ISM band or upper (ch. 38 – 51) UHF TV band), this gives $N \leq 6000$. This allows, for example, a 22-node SISO network with 12 taps per channel, or a 15-node 3x3 MIMO network with 3 taps. Referring to Eq. (8), the multiplier cost is $O(n)$ in bandwidth and taps per channel, but $O(n^2)$ in nodes or antennas per node.

2) *Block RAM*: Delay on the simulated signal paths is implemented using on-chip block RAM. The memory required depends on the sample size (which depends on the required dynamic range), the sample rate, and the duration for which a sample must remain in RAM, which is the channel delay spread. Define the duration of data that can be stored as:

$$S_{max} \leq \frac{\text{RAM}}{sf_{samp}} \quad (16)$$

Our reference FPGA has a $3.6 * 10^7$ bits of block RAM. Considering $B = 100$ MHz, and an 18-bit sample size ($d = 108$ dB), $S_{max} \leq 1$ S. The delay spread is bound by Eq. (17). The number of taps is not a factor, as each sample needs to be held until the final tap, regardless of the total number.

$$\tau \leq \frac{S_{max}}{A^2 n(n-1)} \quad (17)$$

For any of the channels discussed in Table I, multipliers will be the limiting factor on our FPGA: Considering the limiting 15-node 3x3 MIMO network from above, $\tau \leq 530$ μ s.

3) *Host Computation*: Tap weights are generated on the host, as shown in Figure 4. The processing time limits the update frequency of tap weights. Our current implementation requires 3.2 μ s per tap using one 3.80 GHz Xeon core, which corresponds to 15 nodes with 3-tap channels at an update rate of ≈ 2 ms. This implementation was not designed for vehicular speeds; we believe that it can be accelerated dramatically.

4) *Host-FPGA Communication*: If tap weights are computed on the host and evolve according to a random process, they must be sent to the FPGA at an interval of k times the tap's coherence time. The k will depend on the on-FPGA interpolation, but the value cannot be forecast beyond the duration of significant correlation. Let s_t be the size of a tap weight, in bits. The data rate r will be bound by Eq. (18):

$$\frac{s_t N}{k T_c} \leq r \leq \text{host-fpga bandwidth} \quad (18)$$

Consider $k = 1$, $s_t = 24$ bits, and 1 Gb/s of bandwidth. This gives $N \lesssim 4 * 10^7 T_c$. This is limiting for vehicular coherence time estimates: For our worst-case coherence time of $T_c = 23$ μ s, this allows about 960 taps, enough for a 9-node

²Using complex samples and taps, each multiplication requires 4 real multiplications, but the required sample rate is halved, giving $N \leq \frac{f_{mult}}{4B}$.

(SISO) network with 12-tap channels, or 15 nodes with 4-tap channels. However, the worst *observed* value $T_c = 260$ μ s allows 30 nodes with 12-tap channels. The number of possible taps is $O(n)$ in the coherence time and the available bandwidth. The number of taps required scales as in § IV-B1.

V. CONCLUSION

Wireless network emulation enables controlled repeatable experimentation with reasonable-sized wireless networks. The combination of real radio hardware and signal-level emulation provides a level of realism significantly greater than that of network simulators, while the scale allows network-wide issues such as routing, forwarding, and resource management to be studied. This work presents a single architecture which can implement channel models for situations as varied as an indoor Bluetooth scatternet, an outdoor macrocell, or a high-speed vehicular DSRC network.

Given the ability to implement a wide range of channel models, the next challenge is determining *which* models to use for each channel – and for inter-channel interactions – given arbitrary and dynamic user-specified environments.

REFERENCES

- [1] G. Judd and P. Steenkiste, "A software architecture for physical layer wireless network emulation," in *WiNTECH*, Los Angeles, 2006.
- [2] K. Borries, G. Judd, D. Stancil, and P. Steenkiste, "FPGA-based channel simulator for a wireless network emulator," in *67th Vehicular Technology Conference (VTC2009-Spring)*, 2009.
- [3] *Digital cellular telecommunications system (Phase 2+); Radio Transmission and Reception*, GSM 3GPP Technical Specification ETSI TS 100 910, Rev. V8.20.0 (2005-11), 3GPP TS 05.05.
- [4] A. S. Akki and F. Haber, "A statistical model of mobile-to-mobile land communication channel," *IEEE Transactions on Vehicular Technology*, vol. 35, pp. 2 – 7, February 1986.
- [5] A. G. Zajić and G. L. Stüber, "Space-time correlated mobile-to-mobile channels: Modelling and simulation," *IEEE Transactions on Vehicular Technology*, vol. 57, pp. 715 – 726, 2008.
- [6] A. Paier, J. Karedal, N. Czink, H. Hofstetter, C. Dumard, T. Zemen, F. Tufvesson, A. F. Molisch, and C. F. Mecklenbräuker, "Car-to-car radio channel measurements and 5ghz: Pathloss power-delay profile, and delay-doppler spectrum," in *IEEE ISWCS 2007*, 2007, pp. 224 – 228.
- [7] A. Paier, J. Karedal, N. Czink, C. Dumard, T. Zemen, F. Tufvesson, C. F. Mecklenbräuker, and A. F. Molisch, "Comparison of lund'07 vehicular channel measurements with the IEEE 802.11p channel model," EURO-COST, Technical Report COST 2100 TD(08) 436, 2008.
- [8] I. Tan, W. Tang, K. Laberteaux, and A. Bahai, "Measurement and analysis of wireless channel impairments in DSRC vehicular communications," in *ICC 2008*, 2008, pp. 4882 – 4888.
- [9] J. Smith, "A computer generated multipath fading simulation for mobile radio," *Vehicular Technology, IEEE Transactions on*, vol. 24, no. 3, pp. 39 – 40, Aug. 1975.
- [10] R. H. Clarke, "A statistical theory of mobile-radio reception," *The Bell System Technical Journal*, vol. 47, no. 6, pp. 957 – 1000, 1968.
- [11] R. Clarke and W. L. Khoo, "3-d mobile radio channel statistics," *Vehicular Technology, IEEE Transactions on*, vol. 46, pp. 798 – 799, Aug. 1997.
- [12] F. Vatalaro and A. Forcella, "Doppler spectrum in mobile-to-mobile communications in the presence of three-dimensional multipath scattering," *IEEE Trans. on Vehicular Technology*, vol. 46, 1997.
- [13] X. Zhao, J. Kivinen, P. Vainikainen, and K. Skog, "Characterization of doppler spectra for mobile communications at 5.3 ghz," *IEEE Transactions on Vehicular Technology*, vol. 52, pp. 14 – 23, 2003.
- [14] T. Tank and J.-P. M. G. Linnartz, "Statistical characterization of rician multipath effects in a mobile-to-mobile communication channel," *Int. Jnl. of Wireless Information Networks*, vol. 2, pp. 17–26, 1995.
- [15] "Virtex-6 family overview," http://www.xilinx.com/support/documentation/data_sheets/ds150.pdf, January 2010, Doc. DS150 (v2.2).