MultiGigabit Millimeter Wave Communication: System Concepts and Challenges

Upamanyu Madhow
Department of Electrical and Computer Engineering
University of California
Santa Barbara, CA 93106, USA
Email: madhow@ece.ucsb.edu

***Abstract*— The millimeter wave band from 60-95 GHz offers large swathes of unlicensed and semi-unlicensed spectrum, which may well form the basis for the next revolution in wireless communication, in which wireless catches up with wires. With the rapid scaling of silicon processes, low-cost implementations for radio frequency front-ends are on the horizon. A key challenge now is to parlay these breakthroughs into innovative system concepts. We review three such concepts here.**

***Millimeter wave MIMO:* The small carrier wavelength enables spatial multiplexing in line-of-sight environments, potentially resulting in point-to-point outdoor wireless links at optical speeds (40 Gbps) using bandwidths of the order of 5 GHz.**

***Directional multihop networking:* Indoor Gigabit wireless links based on 60 GHz unlicensed spectrum are subject to disruption due to line-of-sight blockage by obstacles such as furniture and humans. We show that a multihop architecture with a small number of relays assures full network connectivity.**

***All-digital multiGigabit baseband:* Since high-speed analog-to-digital conversion (ADC) is costly and power-hungry, in order to design all-digital baseband processing that can be implemented inexpensively by riding Moore's law, we must be able to perform signal processing with sloppy ADC. We discuss Shannon-theoretic limits and signal processing challenges in this context.**

I. Introduction

Millimeter wave communication has the potential of sparking the next revolution in wireless communication in particular, and telecommunications in general. While cellular and wireless local area networks at lower frequencies (1-5 GHz) constantly struggle with the scarcity of spectrum, the mm wave band has huge swathes of spectrum available at no cost. In the US, there is 7 GHz of unlicensed spectrum in the 60 GHz ”oxygen absorption” band that is well suited to short-range indoor and outdoor links. Further, the Federal Communications Commission (FCC) has made available 13 GHz of spectrum in the 70-95 GHz (away from the oxygen absorption band, in order to facilitate longer range communication) for semi-unlicensed use for directional point-to-point ”last mile” links. Effective use of this spectrum may well enable wireless to finally ”catch up” with wires, leading to systems such as ”wireless USB,” ”Gigabit wireless Ethernet,” and ”wireless fiber.”

Mm wave systems have been deployed by the military for decades. However, it is only now, with the prospect of using inexpensive silicon processes (as CMOS and SiGe processes provide faster and faster transistors) [1] and packaging to realize mm wave RF integrated circuits (ICs), that we can anticipate obtaining the economies of scale necessary for proliferation of mm wave systems in commercial applications. The time is ripe, therefore, to develop and validate system concepts that address some of the fundamental challenges associated with high-speed, mm wave communication. First, millimeter wave links are fundamentally a line of sight (LOS) technology, being highly vulnerable to blockage. This is because diffraction around obstacles is severely curtailed (relative to that at lower carrier frequencies) because of the smaller wavelengths: as a rule of thumb, an obstacle that is a few wavelengths across effectively blocks the wave. Second, according to the Friis formula for free space propagation, the propagation loss scales as λ^2, where λ is the carrier wavelength, so that, for given transmit and receive antenna directivities, mm wave links incur much higher losses than, say, existing cellular and WLAN links. Fortunately, it is far easier to synthesize highly directive antennas at smaller wavelengths, and for a given antenna size, the directivity scales as $\frac{1}{\lambda^2}$. Thus, if we fix the area of the antennas used at each end, the overall propagation loss scales as $\frac{1}{\lambda^2}$, so that we actually win by using lower wavelengths. Of course, this approach has significant associated challenges, including the issue of how to direct the transmit and receive antennas at each other (electronic beamsteering becomes crucial for enabling flexible and robust deployments), and how to design network protocols that work with highly directional beams. Third, at speeds of multiGigabits/sec, the modern paradigm of digital communication transceiver design, which is heavily centered around digital signal processing (DSP) becomes difficult to realize. This is because high-speed, high-precision analog-to-digital conversion (ADC) is too costly and power-hungry. The solution in the short term is to come up with creative hybrid analog/digital baseband designs. However, an all-digital, DSP-centric approach is attractive in the longer term, because it enables us to ride Moore's law to obtain the economies of scale associated with low-cost IC implementations. It is of interest, therefore, to explore all-digital baseband transceiver design with imperfect, or ”sloppy,” ADC. Finally, based on the spectral chunks available, and the feasible bandwidths for RF and baseband processing, channel bandwidths for mm wave systems are expected to be of the order of 1-5 GHz. Attaining optical speeds of 40 Gigabits/sec over such a channel is a challenging task. The use of very large constellations is prohibited not only by the complexity of processing them, but also by the excessive power requirements for operating in a bandwidth-efficient regime. One approach is to employ power-efficient signaling, but to increase spectral efficiency by employing spatial multiplexing. However, unlike lower frequency systems, where spatial multiplexing arises from rich scattering, we must now obtain it in the LOS regime. We shall see that this is made possible by the small wavelengths at mm wave frequencies.

We now discuss three system concepts which respond to the preceding challenges associated with mm wave communication. As mentioned in the acknowledgements, this work is in collaboration with a number of graduate students and colleagues at UCSB. Our goal is to provide an overview of these activities. We therefore provide a brief review of work that either has been or will be reported in detail elsewhere.

II. Millimeter wave MIMO

Consider an outdoor point-to-point mm wave LOS link with a directional transmit antenna aimed at a directional receive antenna. Such beamsteering can be done either manually or electronically, and there has been significant progress recently on devising innovative IC realizations [3] for the latter. Let us call each such directive antenna a *subarray,* and form an array of such subarrays, pointing at each other, as shown in Figure 1. As shown in the figure, each subarray can be electronically steerable using a beamsteering IC. A wave sent from a given subarray has a particular complex-valued response at the receive array of subarrays. The key observation now is that, for mm wave carrier frequencies, even at ranges of the order of kilometers, the receive array response due to different transmit subarrays is different enough to permit spatial multiplexing, as long as the transmit and receive subarrays spacing is of the order of a meter. More specifically, for the linear arrays shown in Figure 1, the receive array response corresponding to two transmit elements is actually orthogonal if the spacing between subarrays satisfies

$$d = \sqrt{\frac{R\lambda}{N_R}}$$

This is equivalent to the well-known Rayleigh criterion providing the diffraction-limited resolution of optical systems.

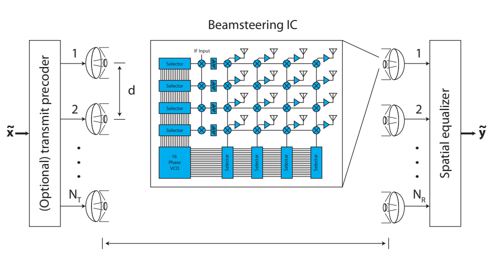

Fig. 1. Spatial multiplexing in a highly directive, LOS environment

In this case, if we sent independent data streams from different transmit subarrays, they can be separated at the receiver simply by spatial matched filtering. Of course, the Rayleigh criterion may not be satisfied: we may wish to space the subarrays closer in order to attain a more compact form factor, or the range we operate at may not be exactly the same as the nominal range for which the subarrays satisfy the Rayleigh criterion. In this case, a mix of transmit precoding and spatial equalization is required to support spatial multiplexing.

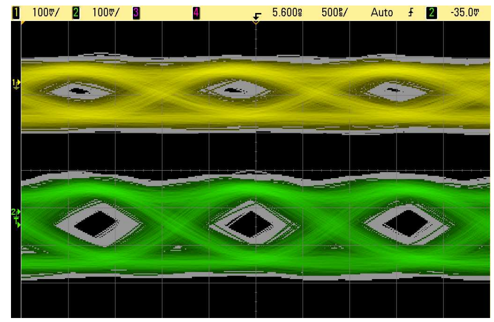

Fig. 2. Eye diagrams for the two spatially multiplexed channels, each running at 600 Mbps, after receiver spatial matched filtering.

The mm wave MIMO concept has been verified experimentally at UCSB, using two transmit and two receive subarrays at Rayleigh spacing, pointed manually at each other over a range of 6 meters. The spatial matched filtering at the receiver is performed by tuning carrier phases at intermediate frequency, and the channel separation thus obtained is 17 dB, which essentially corresponds to complete removal of the co-channel interference. Each transmit subarray sends at 600 Mbps, for an aggregate data rate of 1.2 Gbps, and the noise margin is such that the measured bit error rate is less than 10^{-6}. The eye diagram for the two channels after the spatial matched filtering are shown in Figure 2. Detailed reports on the results of these and other experiments in progress will be reported in forthcoming publications.

III. Directional multihop networking

We now turn to indoor mm wave networking using the unlicensed 60 GHz band. Oxygen absorption is not a factor for the short ranges (of the order of 10 meters or less) for such networks. The key bottleneck here is the blockage due to walls, furniture, and moving obstacles such as humans. Given that transmit power is expensive at mm wave frequencies, it is a bad idea to try to blast through obstacles (e.g., the power loss due to blockage by a human may be of the order of 20 dB). Rather, this is a scenario tailormade for multihop communication, wherein we try to maintain network connectivity using power-efficient LOS links whenever possible by routing around obstacles.

Figure 3 shows a typical office, with a number of wireless terminals communicating at 60 GHz. It is clear that, despite the small confines of the scenario shown, the network is not fully connected: some links are blocked by stationary obstacles, while others which are nominally available can be blocked by the movement of humans within the office. We can apply Huygen's principle to evaluate the propagation loss due to

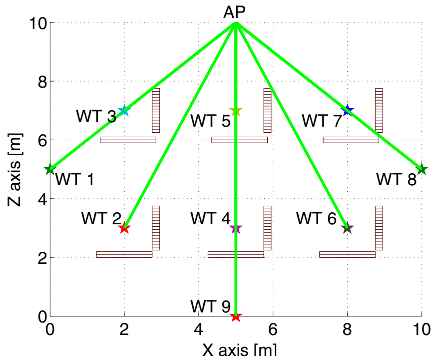

Fig. 3. An office with cubicles, with blockage of LOS by furniture and humans moving around (not shown).

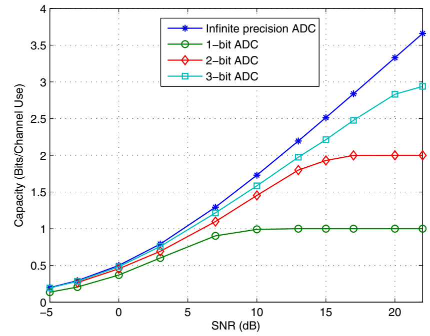

Fig. 5. Capacity with 1-, 2-, and 3-bit quantization.

obstacles, using simple cylindrical models for humans, for example. This, together with a desired link margin for reliable operation, provides a model for link connectivity as a function of time. This is used to demonstrate in [4] that, for a small number of relays placed in strategic locations, a multihop architecture enables the network to maintain full connectivity. For a network controlled by a single access point, it is possible to react effectively to time-varying link connectivities using relatively simple medium access control strategies. As shown by Figure 4, most terminals in the network of Figure 3 maintain connectivity to the access point through a multihop route.

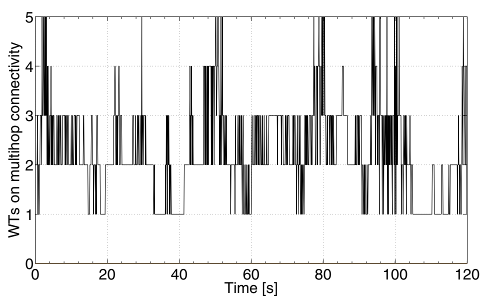

Fig. 4. Multihop is crucial to maintain connectivity in a mm wave network.

IV. Baseband design with sloppy ADC

Finally, we address the issue of baseband design at multiGigabit rates. We typically operate using power-efficient signaling with small constellations, which implies that multiGigabit data rates correspond to multiGigahertz sampling rates. As we have discussed, high-precision ADC at such high rates is either not available, or too costly and power-hungry. On the other hand, DSP-centric transceiver architectures are attractive in terms of cost-effective implementations, hence we wish to examine whether we can relax the requirements on the ADC while maintaining an "all-digital" baseband.

We begin with Shannon theory for the classical AWGN channel with ideal synchronization and Nyquist sampling: how is channel capacity affected if we severely reduce ADC precision, from typical values of 8-12 bits to 1-4 bits? Preliminary results were reported in [5], and further results are being submitted separately. The basic insights from our analysis and computations are that the optimal input distribution is discrete, with the number of levels matched to the number of quantizer levels, that standard uniform PAM constellations are close to optimal, and that there is only a moderate loss of spectral efficiency relative to unquantized observations even at moderately high SNR (e.g., at 10 dB SNR, the capacity with 2-bit quantization is 85% of that with unquantized observations).

Our Shannon-theoretic analysis shows that, under ideal conditions, the performance loss due to severe quantization is often quite acceptable. However, even for an ideal LOS link, carrier and timing synchronization is required in order to reduce the continuous-time channel to the discrete-time channel considered in the preceding analysis. Algorithms for efficiently achieving this with severely quantized samples are therefore an important topic of ongoing research. A more difficult problem is how to handle dispersive channels when the receiver employs low-precision ADC. Even when small constellations and singlecarrier modulation are employed, the increase in dynamic range created by a dispersive channel may be excessive for an all-digital receiver employing heavily quantized samples. One potential approach in such situations that we are exploring is the use of transmit precoding based on feedback from the receiver. An alternative approach for handling large dynamic range is to use higher-precision, lower-speed ADCs in parallel; the performance of such a time-interleaved ADC is limited by mismatch in gain, timing, and frequency response between the component parallel ADCs. For generic ADC applications, such mismatch can be corrected in the digital domain, as shown in our recent work [7], [8]. Taking this a step further for the communication applications considered here, it is possible to adapt the receiver to address ADC mismatch and channel imperfections jointly, leveraging the training available in communication signals. For a indoor wireless channel which is a realization of the CM 1 statistical model (corresponding to a channel with a strong LOS component) developed for UWB system design [6], we see from Figure 6 that mismatch correction is effective in approaching the performance without mismatch for an OFDM system.

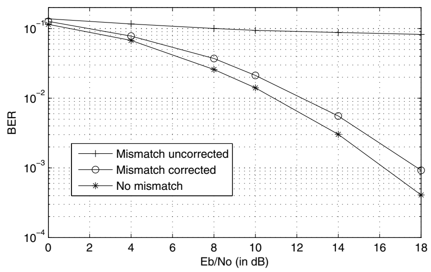

Fig. 6. Bit error rates for a 128 sub-carrier OFDM system with 16-QAM modulation over a typical indoor channel with a dominant LOS component, using 8 time-interleaved ADCs in parallel. The performance without ADC mismatch correction is poor, but joint mismatch correction and channel compensation approaches the performance without mismatch.

Further details for both singlecarrier and multicarrier systems will be reported in subsequent publications.

V. Conclusions

As illustrated by the three examples discussed here, there is plenty of room for innovative system design in the context of emerging, ultra high-speed mm wave communication networks. Ongoing research focuses on signal processing/hardware co-design for mm wave MIMO, protocols for networks with steerable, highly directional links, and signal processing algorithms for all-digital architectures using imperfect ADC.

Acknowledgements

This work was supported by the National Science Foundation under grants CCF-0431205, ECS-0636621 and CCF-0729222. The work surveyed here is in collaboration with a number of graduate students and colleagues. The work on mm wave MIMO involves Colin Sheldon, Eric Torkildson, Prof. Mark Rodwell, and Prof. Patrick Yue. The work on directional wireless networking involves Sumit Singh, Prof. Elizabeth Belding, and Prof. Mark Rodwell. The work on all-digital baseband involves Jaspreet Singh, Sandeep Ponnuru, Dr. Munkyo Seo, Prof. Onkar Dabeer, and Prof. Mark Rodwell.

References

[1] B. Floyd, *et al,* "SiGe vs. CMOS for 60-100 GHz: technology, circuits, packages, and systems," Govt. Microcircuit Applications Conf. Dig. Papers, Mar. 2007.

[2] E. Torkildson, B. Ananthasubramaniam, U. Madhow and M. Rodwell, "Millimeter-wave MIMO: wireless links at optical speeds," *Proc. 44th Allerton Conference on Communication, Control and Computing,* University of Illinois at Urbana-Champaign, September 2006.

[3] C. Carta, M. Seo, M. Rodwell, "A mixed-signal row-column architecture for very large monolithic mm-wave phased arrays," *2006 IEEE Lester Eastman Conference on High Performance Devices,* August 2006.

[4] S. Singh, F. Ziliotto, U. Madhow, E. M. Belding and M. J. W. Rodwell, "Millimeter wave WPAN: cross-layer modeling and multihop architecture," *Proc. IEEE Infocom 2007 Minisymposium,* Anchorage, Alaska, USA, May 2007.

[5] J. Singh, O. Dabeer and U. Madhow, "Communication limits with low-precision analog-to-digital conversion at the receiver," *Proc. IEEE International Conference on Communications (ICC'07),* Glasgow, Scotland, June 2007.

[6] J. R. Foerster, M. Pendergrass, A. Molisch, "A channel model for Ultrawideband indoor communication," Mitsubishi Electric Research Laboratory, Inc., TR-2003-73, November 2003.

[7] M. Seo, M. J. W. Rodwell, U. Madhow, "Comprehensive digital correction of mismatch errors for a 400 Msamples/s 80-dB SFDR time-interleaved analog-to-digital converter," *IEEE Trans. Microwave Theory and Techniques,* vol. 53, no. 3, pp. 1072-1082, March 2005.

[8] M. Seo, M. Rodwell and U. Madhow, "Generalized blind mismatch correction for two-channel time-interleaved A-to-D converters," *Proc. IEEE International Conf. on Acoustics, Speech and Signal Processing (ICASSP '07),* Hawaii, USA, April 2007.