

# A Scalable Massive MIMO Array Architecture Based on Common Modules

Antonio Puglielli\*, Nathan Narevsky\*, Pengpeng Lu\*, Thomas Courtade\*, Gregory Wright†,  
Borivoje Nikolic\*, Elad Alon\*

\*University of California, Berkeley, CA 94704 USA

†Alcatel-Lucent Bell Labs, Holmdel, NJ 07733 USA

**Abstract**—Massive MIMO is envisioned as one of the key enabling technologies for 5G wireless and beyond. While utilizing the spatial dimension to reduce interference and increase capacity in multi-user scenarios, massive MIMO base stations present several unique implementation challenges due to their large physical size and the high datarate generated by all the elements. To be cost-effective and energy efficient, practical designs must leverage the particular characteristics of massive MIMO to ensure scalability. Here, we propose an array architecture based on a common module which serves a small number of antennas with RF transceivers, data converters, and several support functions. Multiple chips are tiled into a grid and interconnected through a digital nearest-neighbor mesh network, avoiding the severe problems associated with analog signal distribution. Scalability across a wide range of array sizes is achieved by using distributed beamforming algorithms. It is demonstrated that by using this approach, the maximum backhaul datarate scales as the number of users rather than the number of antennas. Finally, we present a detailed accounting of the power consumption of the array and use the resulting optimization problem to show that per-element overhead limits the minimum achievable power consumption.

## I. INTRODUCTION

There has been significant recent interest in massive MIMO, which refers to the use of base-station antenna arrays with a very large number of elements to communicate with a much smaller number of spatially dispersed users. This attention is largely due to the fact that, as the base-station array size increases, simple linear beamforming asymptotically achieves the capacity of the multi-user channel [1], [2], [3]. In essence, massive MIMO base-stations exploit a very high degree of spatial resolution to distinguish the users' unique spatial signatures and apply beamforming to cancel inter-user interference. Under the relatively modest (and reasonable, according to measurements [4], [5]) condition that the users' channels become asymptotically orthogonal, it is possible to near-perfectly cancel inter-user interference as the base-station array size grows large [1], [6], [7], [8].

Accompanying the theoretical interest in massive MIMO, there have been several recent demonstrations of test systems aimed at verifying theoretical results and understanding practical design considerations. The Argos [9] and ArgosV2 [10] projects demonstrate systems with 64 and 96 antennas, respectively, operating in the 2.4GHz ISM band. The system is designed in a hierarchical manner, with the central controller serving several hubs, each of which connects to a number of radio modules and provides both a backhaul connection as well

as digital signal processing capability through a local FPGA. The 64-element Argos system has achieved a capacity up to 80 bits/s/Hz serving 15 users [9], [11].

The LuMaMi system built at Lund University [12] consists of a 100-element array communicating with 10 users over a 20MHz channel at 2.6GHz. The system uses 50 FPGAs to implement the baseband processing after data aggregation. The baseband samples are communicated to the FPGAs and the central processor over a hierarchical backhaul network which achieves an aggregate throughput of 384Gbps using a series of interconnected PCI Express interfaces and serializers. Very preliminary results indicate that this system is capable of uplink spatial multiplexing [12].

The Ngara system built by CSIRO in Australia [13] implements an array of up to 32 antennas. This system is structured as a cascade of modules, with a bank of FPGAs performing all the baseband processing and connecting to 32 data converters. The RF/analog circuits are divided into two modules with analog signals routed between them. Operating at VHF frequencies, Ngara achieves a spectral efficiency of up to 67 bits/s/Hz, supporting 50Mb/s uplink and downlink rates to all users over 28MHz bandwidth.

Finally, the recently proposed USC SDR [14] system is assembled hierarchically using servers, FPGAs, and custom-designed RFICs. One or more servers control a series of FPGAs, each of which is connected to up to four radios. The backplane is designed using high-speed PCIe interfaces to perform fully centralized processing on the servers.

The demonstrator systems discussed above have focused on using off-the-shelf hardware to verify theoretical predictions. However, since deployments will vary widely in the required channel capacity and consequently in the array size, it is critical to implement massive MIMO base-stations in a scalable manner to ensure that cost and power consumption are weak functions of the number of antennas. To meet this objective, it is necessary to rethink the array architecture to design solutions specifically tailored to massive MIMO. On this note, most of the testbeds described above have assumed that fully centralized processing yields maximum system flexibility. With this approach, the backhaul network has emerged as a key implementation bottleneck, requiring complex, expensive, and power-hungry components to handle the enormous bandwidth of the array.

In this paper, we discuss the implementation challenges that

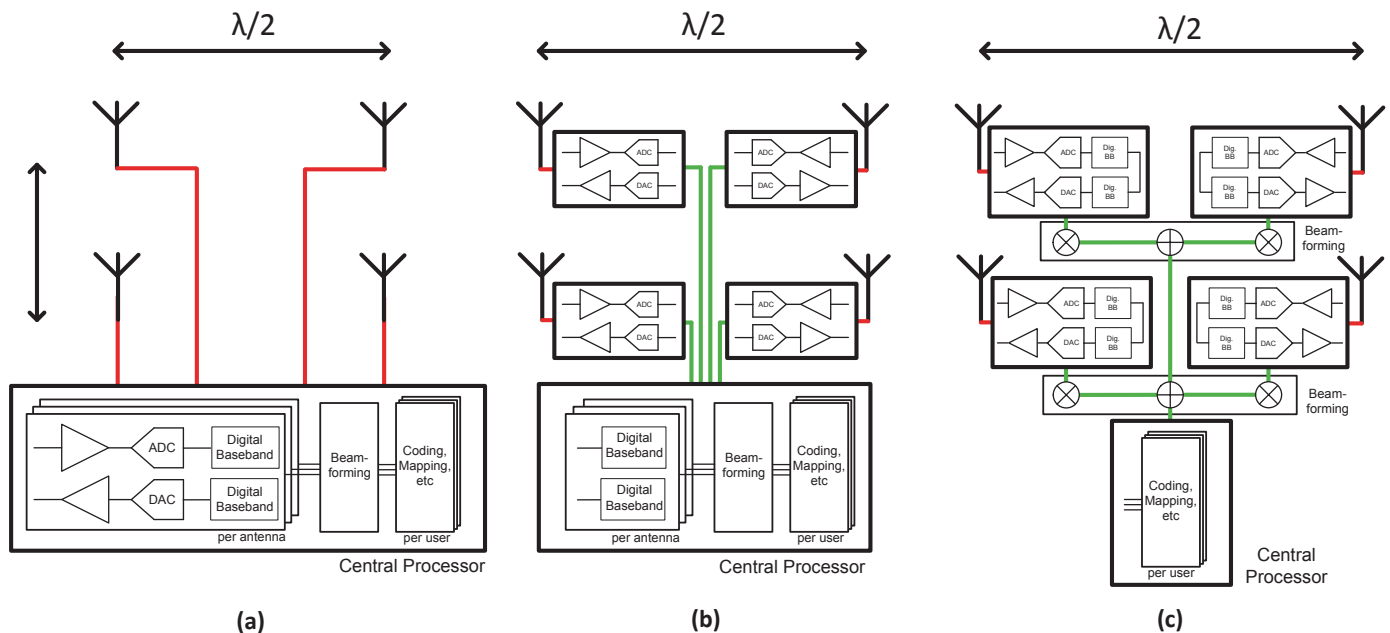


Fig. 1. Three possible array architectures. Analog traces are shown in red and digital links in green. (a) Analog-connected array: analog signals are routed from the antenna to the central processor. (b) Digitally-connected array: the RF front end and data converters are collocated with each antenna and ADC/DAC samples are exchanged with the central processor. (c) Digitally-connected array with distributed beamforming: per-antenna digital baseband processing is performed at the antenna and the beamforming computation is distributed through the routing network.

arise specifically in massive MIMO arrays and propose an architecture that addresses them. The key feature is that the baseband processing is distributed across a number of common modules, where each module is connected to a small number of antennas. The resulting data throughput grows only slowly with the number of modules, enabling the use of low-cost and energy-efficient links between them. Additionally, the common module-based architecture significantly simplifies the signal distribution and the physical layout of the array.

## II. ARRAY ARCHITECTURE

Unlike traditional MIMO systems for wireless communication, massive MIMO demands uniquely large base-station arrays. This presents a challenge in terms of merely transporting signals between the data source/sink and the antenna elements. At RF frequencies, a 100-element half-wavelength-spaced array can easily measure almost a meter on a side. Though this problem is alleviated at millimeter-wave frequencies, it is by no means solved, particularly for arrays with a large number of elements. This implementation issue poses two closely related questions. First, how should the antennas be connected to the central processor? Second, is the required backhaul bandwidth achievable in practice?

Any array architecture requires partitioning the transceiver functions into those implemented close to the respective antennas and those implemented at the main system controller. Since a transceiver can be subdivided into three major blocks — (i) RF/analog front end; (ii) analog-to-digital converter (ADC) and digital-to-analog converter (DAC); and (iii) digital baseband processor — three qualitatively different array architectures arise from distributing different portions of the transceiver chain (Fig. 1):

- 1) **Analog-connected array.** In an analog-connected array, the ADC, DAC, and all the digital baseband hardware are at the central processor while analog signals are routed to and from the antenna elements.
- 2) **Digitally-connected array.** In a digitally-connected array, each antenna is equipped with a complete and collocated analog front end, ADC, and DAC. Each antenna element directly exchanges digitized samples of its analog waveform with the central processor.
- 3) **Digitally-connected array with distributed beamforming.** In a digitally-connected array, any per-antenna baseband processing can be performed locally at the antenna. Consequently, beamforming can be distributed into computations performed throughout the array. With this scheme, the signals exchanged with the central processor correspond to samples of the users' data signals rather than samples of each antenna's RF waveform.

There are several issues with an analog-connected array that limit its use to only a small number of elements. First, analog routing introduces loss that depends exponentially on distance, degrading the signal-to-noise ratio at the receiver and increasing the power consumption in the transmitter. Second, analog-connected arrays are very susceptible to crosstalk and external interference which may limit the performance of beamforming and spatial filtering. Third, analog routing scales very poorly since arrays with more elements require both more and longer transmission lines. Fourth, the routing loss increases with frequency, presenting an obstacle to working at high carrier frequencies. Finally, even moderate quality transmission lines are expensive and bulky.

For all of these reasons, one can qualitatively conclude that, to minimize the length of analog routing, each antenna should




Fig. 2. Block diagram of example array and common module. Analog wires are shown in red and digital links in green. For clarity, this example only shows two antennas per common module. Each antenna is served by a separate RF chain and data converters.

be as close as possible to its data converters. This naturally points to an array composed of a grid of identical modules, each consisting of a single RF transceiver, ADC, and DAC. However, since each module also requires support circuits such as frequency generation and a backhaul, equipping each transceiver with independent copies of these blocks introduces a large amount of redundancy. Instead, multiple transceivers can be fused into one common module equipped with several RF chains, ADCs, and DACs and a single set of support hardware. Interestingly, this suggests that the optimal array is a hybrid of analog and digital architectures, trading off between sharing support hardware and adding analog routing loss. This is discussed further in Section IV. To form the overall array, multiple common modules are tiled together with digital interconnections to their nearest neighbors. Additionally, each common module is equipped with the digital hardware to perform distributed beamforming, which substantially improves the scalability of the backhaul network. Fig. 2 shows a block diagram of a module and, as an example, an array created from nine identical modules, each driving two antennas.

### A. System Model

The system considered here is a massive MIMO array with  $M$  antennas communicating with  $K$  spatially dispersed users. The base-station array is divided into a grid of  $N$  modules, with  $N$  being an integer divisor of  $M$  such that each module serves  $M/N$  antennas. Communication to all users is conducted over a total bandwidth  $B$ , such that the sampling rate of the data converters must be at least  $2B$ . If each ADC and DAC uses  $N_{bit}$  bits and is oversampled by a factor of  $N_{os}$ , then the datarate generated or consumed by one antenna is  $R = 2BN_{os}N_{bit}$  and the datarate of one module is  $(M/N)R$ .

## III. BACKHAUL TOPOLOGIES AND DISTRIBUTED BEAMFORMING

In a digitally interconnected array, the required capacity of the backhaul network emerges as the main limitation on the array size. Indeed, the Lund University testbed needs a total backhaul capacity of 384Gb/s to support 100 elements using 20MHz of channel bandwidth [12].

With a digital interconnection fabric, two main types of backhaul networks are possible. At one extreme, each common module can have its own dedicated physical link to the central unit. At the other, all modules can be daisy-chained on a single link over which all the antennas' signals are transmitted. These topologies correspond to extremes that favor either low per-link datarate on one hand or physical resource sharing on the other. Other types of backhaul networks can be constructed by combining these two extremes, such as having multiple parallel chains. Furthermore, a mesh network can be considered an extension of the daisy-chain concept where each node can communicate with all its nearest neighbors.

The fully parallel backhaul (Fig. 3(a)) requires the lowest per-link datarate but has limited scalability. To serve all the elements, the interconnect length must grow with the size of the array. This requires progressively higher-performance and more costly and power-hungry links and substrates to support reliable bit transmission. In addition, routing complexity and crosstalk between links increases with the number of modules.

These challenges are addressed by implementing the backhaul as a nearest-neighbor mesh network, which requires connections only at the scale of the inter-module distance regardless of the total array size. Since connections are only local, the challenge of globally routing  $N$  links while maintaining acceptable crosstalk levels is avoided entirely. The mesh also presents a level of redundancy that allows reconfiguration of the routing path to circumvent failures.

Nevertheless, a mesh backhaul by itself does nothing to reduce the total bandwidth required at the central processor, as shown in Fig. 3(b). When performing centralized beamforming, regardless of the backhaul topology there is a fundamental requirement to exchange  $M$  unique waveforms with the central processor, for a maximum datarate of  $MR$ . Furthermore, there is an additional penalty in aggregate datarate due to the multi-hop communication. Suppose the  $N$  modules are connected to the central processor with  $N_{chain}$  parallel daisy-chains where  $N_{chain}$  is an integer divisor of  $N$ . At any point along the chain, the datarate is proportional to the number of preceding elements. Therefore, the aggregate datarate through the entire array is

$$R_{tot} = N_{chain} \sum_{i=1}^{N_{chain}} i \frac{M}{N} R = \frac{1}{2} M \left( \frac{N}{N_{chain}} + 1 \right) R \quad (1)$$

The total power consumed by the backhaul network increases as the product of the number of antennas and the number of modules, corresponding to the penalty incurred by sending data through multiple hops. Similar effects occur in the fully parallel backhaul since some links must communicate over a large distance, requiring increased power consumption.

A simple example shows the limitation of centralized processing. Consider an array with  $M = 100$ ,  $B = 20\text{MHz}$ , no oversampling ( $N_{os} = 1$ ) and  $N_{bit} = 10$ ; even under these very modest conditions and ignoring overhead such as addresses and time-stamps, the datarate entering and exiting the central unit is 40Gb/s. A state-of-the-art link operating at this rate

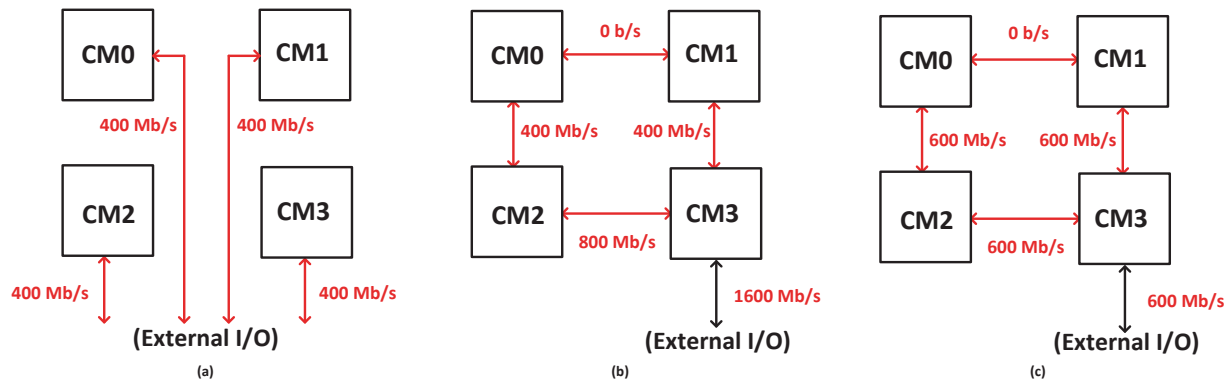


Fig. 3. Three possible routing schemes. Datarates are shown for an illustrative example where each chip generates a datarate of 400Mb/s and the total users' sample rate is 600Mb/s. (a) Fully parallel backhaul: each chip has an independent connection to the central processor. (b) Mesh network backhaul with centralized beamforming: chips are chained together into a mesh network and each antenna exchanges its unique waveform with the central processor. (c) Mesh network with distributed beamforming: each chip computes its estimate of the users' signals and these are summed throughout the routing network to generate the overall beamformed samples.

over a few centimeters could achieve an energy efficiency of 1 pJ/bit [15]. If the array is designed with  $N = 25$  (4 antennas per module) and  $N_{chain} = 5$ , merely transporting bits into and out of the central processor consumes 120mW. For comparison, this is approximately one-sixth the power consumption of an entire 2x2 MIMO 802.11n receiver, including RF, PHY, and MAC [16]. Even exploiting the greatest possible link parallelism at the central processor, it would be difficult to achieve array size-channel bandwidth product much greater than 10GHz-antenna and the required circuits could easily be the most expensive and power-hungry components in the array.

The solution to this problem is to perform distributed beamforming at each module. This idea was originally suggested by [9]; here we extend the discussion, quantitatively compare the routing capacity, and discuss the impact on the common module design. The key insight is that the  $M$  waveforms at the array antennas are not linearly independent but instead lie in a  $K$ -dimensional subspace generated by the  $K$  distinct users. By exploiting this redundancy, it is possible to exchange  $K$  rather than  $M$  unique signals with the central processor by performing distributed beamforming. In the massive MIMO regime, where  $M \gg K$ , the required backhaul capacity is substantially reduced.

Since linear beamforming is simply a matrix multiplication, this computation can be easily distributed. In the uplink, each element multiplies its received signal by a beamforming weight vector containing one entry per user. These vectors are then summed across the array to generate the per-user spatially filtered signals. This task can be embedded in the digital link to be very low-latency and low-energy. The process is reversed in the downlink: all the users' data streams are broadcast to all the modules and each element combines them with the appropriate beamforming weights to generate its DAC samples.

In an OFDM-based communication system where beamforming is performed independently on each subcarrier, each common module requires a timing and frequency recovery block, downsampling and upsampling filters, an FFT unit, and a small number of complex multipliers and adders. At the




Fig. 4. Maximum and aggregate datarates using either centralized or distributed beamforming. In this example,  $B = 20\text{MHz}$ ,  $N_{chain} = 10$ ,  $N_{bit} = 10$  and  $N_{bit,bf} = 15$ . The number of users grows with the size of the base station with fixed ratio  $M/K = 32$ .

beginning of each frame, each user transmits a synchronization and training preamble which is used to estimate the timing parameters for each user at each module. This training preamble can also include channel estimation fields for computation of beamforming weights. Subsequently, during data packet transmission, distributed linear beamforming is performed independently for each subcarrier as described above. Note that in both uplink and downlink the processing is exactly the same as in the fully centralized case. The only additional hardware required is a configurable-depth buffer on each module to match the latency of the backhaul network.

The key result is that, by routing the users' signals around the array rather than the antennas', the maximum required datarate is proportional to the number of users rather than the number of antennas (Fig. 3(c)). Consider a case where each user's modulated data stream is represented by real and complex samples of  $N_{bit,bf}$  bits each. Then the maximum datarate at the central processor is given by

$$R_{max} = 2KBN_{bit,bf} \quad (2)$$

and the aggregate datarate to deliver all  $K$  signals to all  $N$  modules is

$$R_{tot} = 2NKBN_{bit,bf} \quad (3)$$

From these equations, it is evident that the capacity of the central processor's link is only proportional to the number of users served, a substantial improvement over fully centralized processing. Fig. 4 illustrates these benefits for an example case with  $B = 20\text{MHz}$ ,  $N_{bit} = 10$ ,  $N_{bit,bf} = 15$ , and  $N_{chain} = 10$ . Maintaining a constant ratio of base-station antennas to users ( $M/K = 32$ ), both maximum and aggregate datarates are reduced substantially when performing distributed beamforming.

#### IV. ENERGY OPTIMIZATION OF ARRAY AND COMMON MODULE

Large arrays offer potentially large improvements in radiated power efficiency, exploiting high directivity to reduce the actual power consumption required to deliver a certain equivalent isotropic radiated power (EIRP). However, radiated power is only one part of the total power consumption, since there are several sources of overhead that contribute a fixed power cost for every transceiver. This overhead limits the minimum achievable power consumption.

To formulate a power optimization problem, we follow a similar procedure as in [17]. In the downlink, the total power consumption of the array can be split into three contributions, parametrized by the array parameters  $M$  and  $N$ .

$$P_{tot} = \frac{EIRP}{M\eta} e^{(\frac{k\lambda}{2}\sqrt{\frac{M}{N}})} + MP_{tx} + NP_{chip} \quad (4)$$

The first term expresses the power required to achieve a desired EIRP with an array of  $M$  transmitters, each with system-level efficiency  $\eta$  and a routing loss. The routing loss is modeled assuming that each module's  $M/N$  antennas are arranged in a square with half-wavelength separation and, for simplicity, that all wires are the same length. The loss per unit length,  $k$ , depends on both the substrate and carrier frequency. The second term accounts for the fixed power consumption of each transmitter which is independent of the radiated power. The final contribution consists of the overhead power per common module, arising from shared blocks such as voltage regulation, frequency generation, and backhaul. This contribution increases when more modules are added to the array but is amortized across all the antennas connected to a single common module.

As can be seen from (4), there exist optimum values for both  $N$  and  $M$ . The optimum value for  $N$  represents the tradeoff between analog loss and sharing of functions across a common module, while the optimum value for  $M$  represents the tradeoff between radiated power and all sources of overhead. The per-element overhead is therefore the limiting factor to the achievable transmit-mode power consumption.

Fig. 5 shows the array power consumption as a function of the number of elements. Each curve corresponds to a different per-transmitter power overhead  $P_{tx}$ . As expected, reducing  $P_{tx}$  reduces the power consumption of the array, increases the




Fig. 5. Total array power consumption as a function of number of elements, for various per-transmitter overhead powers. Per-chip overhead is held constant at  $1\text{mW}$  and the number of elements per chip is held constant at 4.




Fig. 6. Total array power consumption as a function of number of elements per module for various per-chip overhead powers. Per-transmitter overhead is held constant at  $1\text{mW}$  and the total array size is fixed at 256 elements.

optimum number of elements, and flattens out the optimum. Note that, in this figure, the spectral efficiency depends only on the number of elements and converges to a fixed value asymptotically with the array size. Fig. 6 shows the array power consumption as a function of the number of elements per module, for constant total array size and spectral efficiency. Increasing the overhead power consumption of the module increases the optimal number of elements per module.

In the receive direction, the analysis is somewhat different. If the noise at each antenna is dominated by thermal noise in the receiver itself rather than noise in the environment, it will be uncorrelated between antenna elements. In this regime, beamforming will average out the noise contribution of each element, providing an SNR gain of  $M$ . Consider a reference receiver which achieves the desired output SNR  $\rho$  with power consumption  $P_{rx}$ . Constructing an array from  $M$  such receivers, the system will achieve an output SNR of  $M\rho$  with a power consumption of  $MP_{rx}$ . Since noise figure is, to first order, inversely proportional to power consumption, each receiver's power consumption can be reduced by a factor

of  $M$  to achieve an output SNR of  $\rho$  with an overall power consumption of  $P_{rx}$ . In essence the array breaks a single high performance receiver into many lower performance ones that, in aggregate, recover the original performance. The bottom line is that to achieve a target SNR, the array consumes the same amount of power as a single-element receiver. In reality, the power consumption is increased by the amount of per-element overhead, requiring careful design to minimize the performance loss compared to a single element receiver.

Note that the above analysis only considers operations that must be performed on a per-element basis. Most of the digital baseband processing in both transmit and receive direction is performed on a per-user basis (e.g. coding, modulation, interleaving, etc) and therefore does not present a power cost that depends on the size of the array.

## V. CONCLUSION

In this paper we have analyzed the implementation tradeoffs of massive MIMO arrays. To address challenges associated with the large number of elements, we have proposed designing the array as a grid of identical common modules, each of which serves a (small) number of antennas. Each common module consists of a radio and data converters for each antenna it serves along with several shared functions such as frequency generation and a backhaul connection to the central processor. These modules are digitally interconnected to their nearest neighbors in a mesh network which provides connectivity to the central processor. Furthermore, each module possesses a small amount of digital baseband hardware to perform distributed beamforming.

There are several advantages to this approach:

- 1) Because identical chips are connected to their nearest neighbors, the array size can be increased or decreased simply by adding or removing elements at the periphery and providing short connections to their neighbors.
- 2) This architecture is applicable to many deployments since the flat hierarchy provides flexibility to accommodate a variety of array sizes, aspect ratios, and substrates.
- 3) With distributed beamforming, the maximum data rate at the central processor is proportional only to the number of users rather than the number of base-station elements. This significantly improves the scalability of massive MIMO arrays since this is frequently limited by the cost and power burdens of a very high capacity backhaul.

In addition, we presented a framework that shows how to pick the array parameters to minimize total power consumption. These optima trade off the radiated power and the overhead incurred by adding more elements to the array. Based on the insights gained from this system-level analysis, it can be envisioned to implement a module as a single CMOS integrated circuit and the entire array assembled on a flexible or conforming substrate.

## ACKNOWLEDGMENT

This work was supported in part by the DARPA Arrays on Commercial Timescales program, contract HR0011-14-1-

0055. The authors would also like to acknowledge the students, faculty, and sponsors of the Berkeley Wireless Research Center, in particular Lingkai Kong, Greg LaCaille, Kosta Trotskovsky, Amy Whitcombe, Vladimir Milovanovic, Simon Scott, and Stephen Twigg.

## REFERENCES

- [1] T. L. Marzetta, "Noncooperative Cellular Wireless with Unlimited Numbers of Base Station Antennas," *IEEE Trans. Wireless Communications*, vol. 9, no. 11, pp. 1436–1449, Nov. 2010.
- [2] F. Rusek, D. Persson, B. K. Lau, E. G. Larsson, T. L. Marzetta, O. Edfors, and F. Tufvesson, "Scaling up MIMO: Opportunities and Challenges with Very Large Arrays," *IEEE Sig. Process. Mag.*, vol. 30, no. 1, pp. 40–60, Jan. 2013.
- [3] E. G. Larsson, F. Tufvesson, O. Edfors, and T. L. Marzetta, "Massive MIMO for Next Generation Wireless Systems," *IEEE Commun. Mag.*, vol. 52, no. 2, pp. 186–195, Feb. 2014.
- [4] X. Gao, O. Edfors, F. Rusek, and F. Tufvesson, "Linear pre-coding performance in measured very large MIMO channels," in *Proc. IEEE Veh. Tech. Conf (VTC)*, San Francisco, CA, USA, Sep. 2011.
- [5] J. Hoydis, C. Hoek, T. Wild, and S. ten Brink, "Channel measurements for large antenna arrays," in *Proc. Int. Symp. Wireless Commun. Syst. (ISWCS)*, Paris, France, Aug. 2012.
- [6] H. Huh, G. Caire, H. C. Papadopoulos, and S. A. Ramprasad, "Achieving "Massive MIMO" Spectral Efficiency with a Not-so-Large Number of Antennas," *IEEE Trans. Wireless Communications*, vol. 11, no. 9, pp. 3226–3239, Sept. 2012.
- [7] J. Hoydis, S. ten Brink, and M. Debbah, "Massive MIMO in the UL/DL of Cellular Networks: How Many Antennas Do We Need?" *IEEE J. Sel. Areas Commun.*, vol. 31, no. 2, pp. 160–171, Feb. 2013.
- [8] H. Yang and T. L. Marzetta, "Performance of Conjugate and Zero-Forcing Beamforming in Large-Scale Antenna Systems," *IEEE J. Sel. Areas Commun.*, vol. 31, no. 2, pp. 172–179, Feb. 2013.
- [9] C. Shepard, H. Yu, N. Anand, E. Li, T. Marzetta, R. Yang, and L. Zhong, "Argos: practical many-antenna base stations," in *Proc. ACM Int. Conf. Mobile Computing and Networking (Mobicom)*. Istanbul, Turkey: ACM, 2012.
- [10] C. Shepard, H. Yu, and L. Zhong, "ArgosV2: A Flexible Many-Antenna Research Platform," in *Proc. ACM Int. Conf. Mobile Computing and Networking (Mobicom)*. ACM, 2013.
- [11] C. Shepard, N. Anand, and L. Zhong, "Practical performance of MU-MIMO precoding in many-antenna base stations," in *Proc. ACM workshop on Cellular Networks: operations, challenges, and future design (Cellnet)*. ACM, 2013.
- [12] J. Vieira *et al.*, "A flexible 100-element testbed for Massive MIMO," in *IEEE Globecom 2014 Workshop - Massive MIMO: From Theory to Practice*. Austin, TX, USA: IEEE, 2014.
- [13] H. Suzuki, R. Kendall, K. Anderson, A. Grancea, D. Humphrey, J. Pathikulangara, K. Bengston, J. Matthews, and C. Russel, "Highly spectrally efficient Ngaru Rural Wireless Broadband Access Demonstrator," in *Communications and Informations Technologies (ISCIT), 2012 International Symposium on*. Austin, TX, USA: IEEE, 2012.
- [14] H. V. Balan, M. Segura, S. Deora, A. Michaloliakos, R. Rogalin, K. Psounis, and G. Caire, "USC SDR, an easy-to-program, high data rate, real time software radio platform," in *SRIF '13, Proceedings of the second workshop on Software radio implementation forum*. ACM, 2013.
- [15] B. Raghavan *et al.*, "A Sub-2 W 39.844.6 Gb/s Transmitter and Receiver Chipset With SFI-5.2 Interface in 40 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 48, no. 12, pp. 3219–3228, Dec. 2013.
- [16] M. Zargari *et al.*, "A Dual-Band CMOS MIMO Radio SoC for IEEE 802.11n Wireless LAN," *IEEE J. Solid-State Circuits*, vol. 43, no. 12, pp. 2882–2895, Dec. 2008.
- [17] L. Kong, "Energy-efficient 60ghz phased-array design for multi-gb/s communication systems," Ph.D. dissertation, EECS Department, University of California, Berkeley, Dec 2014. [Online]. Available: <http://www.eecs.berkeley.edu/Pubs/TechRpts/2014/EECS-2014-191.html>