Efficient DSP and Circuit Architectures for Massive MIMO: State of the Art and Future Directions

Liesbet Van der Perre , Liang Liu , and Erik G. Larsson , *Fellow, IEEE*

Overview Article

Abstract—Massive MIMO is a compelling wireless access concept that relies on the use of an excess number of base-station antennas, relative to the number of active terminals. This technology is a main component of 5G New Radio and addresses all important requirements of future wireless standards: a great capacity increase, the support of many simultaneous users, and improvement in energy efficiency. Massive MIMO requires the simultaneous processing of signals from many antenna chains, and computational operations on large matrices. The complexity of the digital processing has been viewed as a fundamental obstacle to the feasibility of Massive MIMO in the past. Recent advances on system-algorithm-hardware co-design have led to extremely energy-efficient implementations. These exploit opportunities in deeply-scaled silicon technologies and perform partly distributed processing to cope with the bottlenecks encountered in the interconnection of many signals. For example, prototype ASIC implementations have demonstrated zero-forcing precoding in real time at a 55 mW power consumption (20 MHz bandwidth, 128 antennas, and multiplexing of 8 terminals). Coarse and even error-prone digital processing in the antenna paths permits a reduction of consumption with a factor of 2 to 5. This article summarizes the fundamental technical contributions to efficient digital signal processing for Massive MIMO. The opportunities and constraints on operating on low-complexity RF and analog hardware chains are clarified. It illustrates how terminals can benefit from improved energy efficiency. The status of technology and real-life prototypes discussed. Open challenges and directions for future research are suggested.

Index Terms—Antenna arrays, circuits, MIMO communication.

I. INTRODUCTION

MASSIVE MIMO is an efficient sub-6 GHz physical-layer technology for wireless access, and a key component of the 5G New Radio (NR) interface [1]. The main concept is to use large antenna arrays at base stations to simultaneously

Fig. 1. Massive MIMO exploits large antenna arrays at the base stations, to spatially multiplex many terminals.

serve many autonomous terminals, as illustrated in Figure 1 [2], [3]. Smart processing at the array exploits differences among the propagation signatures of the terminals to perform spatial multiplexing. Massive MIMO offers two main benefits:

- 1) Excellent spectral efficiency, achieved by spatial multiplexing of many terminals in the same time-frequency resource [4], [5]. Efficient multiplexing requires channels to different terminals to be sufficiently distinct. Theory as well as experiments have demonstrated that this can be achieved both in line-of-sight and in rich scattering.
- 2) Superior energy efficiency, by virtue of the array gain, that permits a reduction of radiated power. Moreover, the ability to achieve excellent performance while operating with low-accuracy signals and linear signal processing further enables considerable savings in the power required for signal processing.

This overview paper focuses on sub-6 GHz Massive MIMO systems implemented with fully digital per-antenna signal processing. Massive MIMO at mmWave frequencies is also possible, and can benefit from the large bandwidth available at these frequencies. Propagation and hardware implementation aspects are different at mmWaves; for example, hybrid analog-digital beamforming approaches are typically considered [6]. However, this is not discussed further here.

The complexity of the signal processing has been considered a potential obstacle to actual deployment of Massive MIMO technology. An obvious concern is how operations on large matrices and the interconnection of the many antenna signals can be efficiently performed in real-time. Moreover, real-life experiments have shown that the channel responses to different terminals can be highly correlated in some propagation

Manuscript received January 11, 2018; revised May 31, 2018 and July 16, 2018; accepted July 17, 2018. Date of publication July 20, 2018; date of current version August 10, 2018. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Marios Kountouris. The work of E. G. Larsson was supported by the Swedish Research Council (VR) and ELLIIT. (Corresponding author: Liesbet Van der Perre.)

L. Van der Perre is with the KU Leuven, Leuven 3000, Belgium and also with Lund University, Lund 22100, Sweden (e-mail: liesbet.vanderperre@kuleuven.be).

L. Liu is with the Lund University, Lund 22100, Sweden (e-mail: liang.liu@eit.lth.se).

E. G. Larsson is with the Department of Electrical Engineering, Linköping University, Linköping 581 83, Sweden (e-mail: erik.g.larsson@liu.se).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TSP.2018.2858190

Fig. 2. Massive MIMO opens up new hardware-software co-design opportunities for low-complexity circuitry.

environments. Appropriate digital signal processing hence needs to feature interference suppression capabilities, which further increases complexity.

This paper discusses the digital signal processing required to realize the Massive MIMO system concept, and examines in detail the co-design of algorithms, hardware architecture, and circuits (Figure 2). Unconventional, low-complexity digital circuitry implementations in deeply scaled silicon are possible, despite (and thanks to) the excess number of antenna signals. A careful choice of algorithmic and circuit parameters permits considerable reduction of the average energy consumption. Terminals in turn can be implemented at low complexity while benefiting from the channel hardening effect, that offers increased reliability.

Proof of concept implementations and demonstrations have revealed constraints that turned out more harsh than anticipated in initial theoretical assessments. This concerns the interconnection of the signals from all antennas, which poses a bottleneck that partly necessitates distributed processing. Also, relaxing the specifications of the analog and RF chains can result in higher distortion both in-band and out-of-band than initially anticipated, as hardware imperfections can in general not be considered uncorrelated.

The rest of the paper is organized as follows. First, basic concepts and notation are introduced. Next, we provide a complexity analysis considering computation as well as data transfer. The following section zooms in on the RF and front-end, highlighting the opportunities and constraints of relaxing their specifications in the large-number-of-antennas regime. Subsequently, the central detector and precoder blocks are detailed and major complexity reductions facilitated by algorithm-hardware co-design are demonstrated. Signal processing leveraging on error-resilient circuits in the per-antenna functionality is discussed next, and consequent energy savings are illustrated. Further we introduce introduces the increased reliability that can be delivered on complexity terminals. Finally, in the conclusions we discuss validation performed in real-life test beds, summarize opportunities and constraints in efficient

processing for Massive MIMO systems, and suggest future research directions.

II. MASSIVE MIMO SYSTEM MODEL

This section introduces the notation for MIMO transmission that is used in the paper. Further details can be found in, for example, [3]. We consider the block-fading model where the time-frequency domain is partitioned into coherence intervals within which the channel is static. The number of samples in each coherence interval is equal to the coherence time in seconds multiplied by the coherence bandwidth in Hertz. For the signal processing algorithms discussed in this paper, it does not matter whether there is coding across coherence intervals or not.

In every coherence interval, a flat fading complex baseband channel model applies. Let M be the number of antennas at the base station, and K the number of terminals served simultaneously. Also, denote by $\mathbf{g}_k$ the M -vector of channel responses between the k th terminal and the array. Then on uplink, for every sample in the coherence interval,

$$\mathbf{y} = \sum_{k=1}^K \mathbf{g}_k x_k + \mathbf{w}, \quad (1)$$

where $\mathbf{y}$ is an M -vector comprising samples received at the base station array, $\{x_k\}$ are symbols sent by the k th terminal, and $\mathbf{w}$ is noise. On downlink, assuming linear precoding,

$$y_k = \mathbf{g}_k^T \sum_{k'=1}^K \mathbf{a}_{k'} x_{k'} + w_k. \quad (2)$$

where y_k is the sample received by the k th terminal, $\mathbf{a}_k$ is a precoding vector associated with the k th terminal, x_k is the symbol destined to the k th terminal, and w_k is $CN(0, 1)$ receiver noise.

The base station forms a channel estimate, $\hat{\mathbf{g}}_k$, of $\mathbf{g}_k$ for each terminal k by measurements on uplink pilots. Channel estimation is discussed extensively in for example [3] (for independent Rayleigh fading) and [7] (for correlated fading models).

On uplink, the data streams from the terminals are detected through linear processing. This entails multiplication of $\mathbf{y}$ with a vector, $\mathbf{a}_k$ for each terminal, yielding the scalar $\mathbf{a}_k^H \mathbf{y}$. Common choices of the detection vector $\mathbf{a}_k$ include

$$\begin{cases} \text{max.-ratio:} & \mathbf{a}_k = \alpha_k \hat{\mathbf{g}}_k \\ \text{zero-forcing:} & \mathbf{a}_k = \alpha_k \left[\hat{\mathbf{G}} (\hat{\mathbf{G}}^H \hat{\mathbf{G}})^{-1} \right]_{(:,k)} \\ \text{MMSE:} & \mathbf{a}_k = \alpha_k \left[\hat{\mathbf{G}} (\hat{\mathbf{G}}^H \hat{\mathbf{G}} + \mathbf{I})^{-1} \right]_{(:,k)} \end{cases} \quad (3)$$

where α_k is a normalizing constant (different for the three methods), and $\hat{\mathbf{G}} = [\hat{\mathbf{g}}_1, \dots, \hat{\mathbf{g}}_K]$. The result of this linear processing will comprise the desired signal, embedded in additive interference and noise.

On downlink, channel reciprocity is leveraged. Low-complexity front-ends typically introduce non-reciprocity and this non-reciprocity needs to be compensated for; see Section IV. The base station forms the transmitted vector

 Fig. 3. Signal processing in an OFDM-based Massive MIMO system for M BS antennas and K UEs.

$\sum_k \mathbf{a}_k x_k$ in (2) where the precoding vector $\mathbf{a}_k$ is given by:

$$\begin{cases} \text{max.-ratio:} & \mathbf{a}_k = \alpha_k \hat{\mathbf{g}}_k^* \\ \text{zero-forcing:} & \mathbf{a}_k = \alpha_k \left[\hat{\mathbf{G}}^* (\hat{\mathbf{G}}^T \hat{\mathbf{G}}^*)^{-1} \right]_{(:,k)} \\ \text{regularized zero-forcing:} & \mathbf{a}_k = \alpha_k \left[\hat{\mathbf{G}}^* (\hat{\mathbf{G}}^T \hat{\mathbf{G}}^* + \lambda \mathbf{I})^{-1} \right]_{(:,k)} \end{cases} \quad (4)$$

where, again, $\{\alpha_k\}$ are normalizing constants and λ is a regularization parameter. The signal received at the terminal will contain the symbol of interest, plus additive interference and noise.

Many variations are possible and detection and precoding that take multi-cell interference into account are also possible [7], [8].

III. SIGNAL PROCESSING AND DATA TRANSFER COMPLEXITY ASSESSMENT

Both Massive MIMO base stations and terminals can be implemented with significantly better energy efficiency compared to in conventional systems. This is possible owing to a combination of effects. First, the array gain permits a reduction of the radiated power. Second, the large number of constituent signals promotes excellent performance while operating relatively simple algorithms on coarse signals.

In this section we focus on the processing at the base station side. The opportunity to reduce terminal-side complexity is discussed in Section VII. First, a high-level assessment of the signal processing requirements, in terms of number of computations, is presented. The data transfer and interconnection of signals poses a distinct bottleneck. Hence, next a distributed processing approach is presented to balance performance and complexity.

A. Computational Complexity

We first analyze the computational complexity of a Massive MIMO base station. Figure 3 shows a high-level block diagram of the signal processing for an OFDM-based Massive MIMO

TABLE I

ESTIMATED NUMBER OF DSP OPERATIONS IN GOPS, FOR $M = 100$ AND $K = 10$, 20 MHz BANDWIDTH, AND 3 BPS/Hz (16-QAM, CODE RATE 3/4)

Subcomponent	Downlink data (DL) [GOPS]	Uplink data (UL) [GOPS]	Training [GOPS]
Inner modem	175	520	290
Outer modem	7	40	0
Per-antenna DSP	920	920	920

system. Other modulation options can be used, and single-carrier schemes may be preferred. The overall partition of the processing presented here will still hold.

The processing in Massive MIMO systems is logically grouped into three categories:

- 1) The outer modem performing symbol (de)mapping, (de)interleaving and channel (de)coding. This processing performed on the transmit/receive bits applies to each User Equipment (UE) individually.
- 2) The inner modem comprising channel estimation, and detection and precoding of the uplink and downlink data, respectively. This central processing aggregates/distributes data from/to all the antenna chains.
- 3) The per-antenna processing which primarily consists of the analog and digital front-end (mainly re-sampling and filtering) and OFDM processing.

We identify inherent parallelism and observe that the processing complexity scales with the number of BS antennas, M , the number of UEs, K , or both [9]:

- Per-antenna processing: Scales with M as each antenna requires OFDM (de)modulation and a digital/analog front-end.
- Central processing: Scales with M and K .
- Per-user processing: Scales with K .

The number of digital signal processing operations performed in the sub-systems provides a high-level estimate of complexity. Table I gives numbers for a sample system with $M = 100$ antennas at the base-station and $K = 10$ simultaneous terminals. It is acknowledged that these estimates represent an

over-simplification, as the nature and precision of the operations will be an important determining factor in the eventual hardware complexity and power consumption.

Table I demonstrates that the collective per-antenna digital processing is demanding, and requires a minimal-complexity implementation. Interestingly, the per-antenna processing does not need to be performed with high precision to offer very good performance. An in-depth analysis and efficient implementation options are presented in Section VI.

For the inner modem processing in Massive MIMO, a high degree of reconfigurability is desired in order to adapt to changing operating conditions, such as the number of connected UEs, and their SNRs/path losses. Section V discusses efficient algorithm-hardware co-design solutions for the Massive MIMO precoding and detection.

Furthermore, reciprocity calibration needs to be performed occasionally. Elegant solutions have been proposed and demonstrated, see Section IV.

Channel coding clearly is an essential component of the wireless transmission, yet it is not Massive MIMO-specific and therefore not further treated in this paper.

B. Signal Interconnection and Data Transfer Complexity

The transfer of data between processing components creates a significant challenge, as the amount of signals and data to be aggregated/distributed from/to all the antennas is very high. The required data shuffling rate between the per-antenna processing and the central processing is [9]

$$R_{\text{antennas2central}} = M \times R_{\text{OFDM}} \times W, \quad (5)$$

where R_{OFDM} is the sampling rate after OFDM processing and W is the word-length of one data sample. For a 100-antenna 20 MHz bandwidth system, the sampling rate R_{samp} at each antenna is 30.72 MS/s and thus

$$R_{\text{OFDM}} = R_{\text{samp}} \times \frac{N_{\text{data}}}{N_{\text{sub}} + N_{\text{CP}}} = 16.8 \text{ MS/s}, \quad (6)$$

where N_{data} , N_{sub} , and N_{CP} are the number of data subcarriers, the total number of subcarriers, respectively the number of cyclic prefix samples. Assuming that 24 bits are used for one complex sample, $R_{\text{antennas2central}}$ equals 40.32 Gb/s. This requirement is an order of magnitude higher than in a conventional system.

Additionally, the data transfer network must re-organize data among different dimensions. Figure 4 illustrates the uplink data shuffling between the per-antenna and the central processing. First, ① in the figure, the data shuffling network aggregates data samples of all subcarriers from all antenna chains. Next, ② in the figure, it divides the entire data into bandwidth chunks depending on the number of central processing units in the system, and distributes the data to the corresponding processing unit.

This high data transfer requirements has motivated the development of decentralized processing architectures, which are introduced next.

Fig. 4. Illustration of the data shuffling between the per-antenna and central processing.

Fig. 5. Decentralized processing architecture, performing group-based operations between the per-antenna processing (PAP) and the central unit.

C. Decentralized Processing

Depending on the selected MIMO processing algorithms, both the processing performed in the per-antenna and in the central units, and the communication between these two, will influence the resulting system performance and overall complexity. For instance, the maximum-ratio precoding operation $\sum_k \alpha_k \hat{\mathbf{g}}_k^* x_k$ can be performed in each antenna path in a distributed manner, whereas the zero-forcing algorithm requires centralized processing, specifically for the inversion of the Gram matrix $(\hat{\mathbf{G}}^H \hat{\mathbf{G}})^{-1}$.

Decentralized processing enables parallel computing and offers a balanced trade-off between system performance and data transfer requirements [10]–[13]. The authors of [12] propose a decentralized architecture for both uplink and downlink, illustrated in Figure 5. Instead of aggregating the full channel state information and transmit/received data vectors at the centralized processing node, M antenna nodes are grouped into B equally sized groups, each serving C antenna nodes. A middle-level processing node, labeled group processor, is introduced between the per-antenna and central processor to handle the corresponding data dedicated to the group of C antenna nodes. As a result, a limited amount of data is then aggregated/distributed to/from the central processor, relaxing the requirements on the

data transfer network. For instance, the Gram matrix calculation $\mathbf{Z} = \hat{\mathbf{G}}^H \hat{\mathbf{G}}$ can be rewritten as

$$\mathbf{Z} = \sum_{b=1}^B \hat{\mathbf{G}}_b^H \hat{\mathbf{G}}_b, \quad (7)$$

where $\hat{\mathbf{G}}_b \in \mathbb{C}^{C \times K}$ is the local channel estimate for each group of C antennas. The decentralized processing is performed such that the terms $\hat{\mathbf{G}}_b^H \hat{\mathbf{G}}_b$ are computed at each group processor locally, and the results are aggregated at the central processor for the final summation. The tree-like distributed processing architecture is further elaborated in [14], with special focus on modularity and scalability. Especially, the trade-off between data processing, storage, and shuffling is investigated for maximum-ratio transmission, zero-forcing, and MMSE algorithms.

IV. ANALOG AND RF PROCESSING: RELAX WITH CAUTION!

In traditional base stations, the RF electronics and analog front-ends, and the power amplifiers specifically, consume most of the power [15]. In Massive MIMO, thanks to the array gain provided by the closed-loop beamforming, much less radiated power is needed for the data transmission. This facilitates a significant reduction of the RF complexity and power consumption compared to conventional systems.

The hardware in any wireless transceiver will introduce distortion, and the most important source of distortion is nonlinearities in power amplifiers and quantization noise in A/D-converters. A commonly used model in the literature has been that this distortion is additive and uncorrelated among the antennas [7]. If this were the case, then the effects of hardware imperfections would average out as the number of antennas is increased, in a similar way as the effects of thermal noise average out. In more detail, consider the linear processing in the uplink, $\mathbf{a}_k^H \mathbf{y}$; see Section II. The essence of the argument is that if the received signal at the array, $\mathbf{y}$, is affected by uncorrelated additive distortion noise $\mathbf{d}$, then the effective power of the useful signal after beamforming processing, $\mathbf{a}_k^H \mathbf{g}_k$, would grow as M whereas the power of the distortion, $\mathbf{a}_k^H \mathbf{d}$, would be constant with respect to M (see [7] for more precise analyses). But unfortunately, this model does not accurately describe the true nature of the hardware distortion.

To understand why, fundamentally, the distortion is correlated among the antennas, consider the downlink in the special case of a single terminal in line-of-sight. Then the signal radiated by the m th antenna is simply a phase-shifted version of the signal radiated by the first antenna ($m = 1$). The distortion arising from an amplifier nonlinearity at the m th antenna is phase-shifted by the same amount as the signal. Hence, if all amplifiers have identical characteristics (a weak assumption in practice), the distortion is beamformed into the same direction as the signal of interest, and receives the same array gain as that signal of interest. That is, the effects of the distortion do scale proportionally to M rather than disappearing as M is increased. In this case, the covariance matrix of the distortion, when viewed as an M -vector $\mathbf{d}$, has rank one. A similar effect exists on the uplink, when the nonlinearities in low-noise amplifiers are considered [16].

In the remainder of this section, we discuss the specifics of distortion arising from amplifier nonlinearities and finite-resolution A/D-converters in more detail. We furthermore discuss the impact and calibration of RF front-end non-reciprocity.

A. Power Amplifiers Benefit from the Large Array

The required output power of a Massive MIMO base station can be reduced inversely proportionally to the square root of number of BS antennas, or even linearly in operating regimes with good channel estimation quality, thanks to the coherent combination of all antenna signals. This results in significantly reduced output specifications of the Power Amplifiers (PAs). The power amplification stage typically accounts for $> 70\%$ of the power consumption of base stations in wireless broadband macro-cells [15]. Moreover they necessitate cooling, causing a $\sim 10\%$ overhead. The reduced output power in Massive MIMO hence can reduce the total power by a factor of 3 in an exemplary 100-antenna base station, assuming that all other contributions remain equal.

The PA mostly operates at a low efficiency as a consequence of a considerable back-off, required to avoid entering the saturation region. For OFDM-based systems such as 3GPP-LTE, the PA typically operates with a back off of 8–12 dB. Best-in-class solutions need complex techniques that achieve an efficiency of $\sim 30\%$ [17]. Entering the saturation region introduces non-linear distortion, which comes with two detrimental effects: distortion of the intended signal within the band of interest, and out-of-band (OOB) emissions that result in adjacent channel leakage.

We consider a polynomial memoryless model [18] for the non-linear behaviour of the PA. The impact on the signal at RF can be expressed as:

$$y(t) = \sum_p \alpha_p x_{\text{RF}}^p(t), \quad (8)$$

where $x_{\text{RF}}(t)$ is the input signal to the PA, $y(t)$ is the output signal, and α_p is the non-linear distortion coefficient of the PA for the p th harmonic component. The third-order harmonic will have the largest impact both in terms of in-band distortion and adjacent channel leakage. Furthermore, the amplitude will be limited to the saturation amplitude $a_{\text{out,sat}}$ for input values exceeding the input saturation amplitude $a_{\text{in,sat}}$:

$$|y(t)| = a_{\text{out,sat}}; \quad |x_{\text{RF}}(t)| > a_{\text{in,sat}}. \quad (9)$$

The non-linear distortion resulting from the PAs in the many antenna paths is hence signal dependent. The input signals to the PAs can be correlated, depending on the specific communication scenario in terms of users, channel responses, and power (im)balance among the users. In [19] we analyzed how the distortion terms can combine by means of a basic dual-tone modulation scheme. The following effects can occur:

- 1) The distortions may add up coherently in the channel and generate considerable out-of-band emissions. This will be the case for example in a single-user situation with one strongly dominating propagation direction.

Fig. 6. Increasing number of base station antennas improves the EVM with PAs operating in saturation.

- 2) In most multi-user scenarios the precoder will provide significant different compositions of signals to the antenna paths and hence power amplifiers. In general, this will randomize the harmonic distortion terms.

The constellation diagrams in Figure 6 illustrate the impact of increasing the number of antennas at the base station on the Error Vector Magnitude (EVM), for a case with equal-strength signals for the different users and i.i.d. Rayleigh fading channels. The results were simulated based on a cubic polynomial model for the PA, which operates in saturation (0 dB with respect to the 1 dB compression point). With $M = 30$ antennas at the base station, the constellation points are seriously dispersed and an EVM of -10 dB is measured. When increasing the number of antennas, in steps of 10 in the graph, the clarity of the constellation diagram greatly improves and for $M = 100$ an EVM of -22 dB is observed.

In conclusion, the power amplifiers benefit from the large array owing to the drastically reduced total output power requirement. Moreover in many typical conditions, Massive MIMO systems will not transmit predominantly to one user and in one direction. One could then operate the PAs efficiently in their non-linear region. Hence, a considerable further improvement of the power consumption could be achieved. However, the inconvenient truth is that in general, directive emissions of OOB radiation can arise under some conditions. More detailed mathematical models and results can be found in [20].

B. Coarse and Lean Convertors

The impact of low-resolution data converters on system performance has been investigated. We give an overview of these theoretical results and discuss them in perspective of actual design constraints and merits of state-of-the-art data converters. These reveal that minimizing the resolution strictly (e.g., below 6 bits) does not result in a significant power reduction in a conventional base station. One should hence question any penalty in system performance and/or additional DSP complexity when considering very low resolution data converters.

A specific type of hardware distortion arises if low-precision A/D converters are used at the base station. Such converters are highly desirable owing to their low cost and power consumption. In principle for each bit reduction in resolution, the A/D

converter power is halved. Doubling the sampling frequency will double the power. This is reflected in the common figure-of-merit (F.o.M.) in terms of energy consumption per conversion step (cs) [21] used to assess the design merit of A/D converters implementing different architectural principles and resolution/bandwidth specifications:

$$\text{F.o.M.}_{\text{A/D}} = \frac{\text{Power}_{\text{A/D}}}{2^{\text{ENOB}} \cdot f_s} \quad (10)$$

where ENOB is the Effective Number of Bits resolution as measured and f_s is the sampling frequency.

The resulting quantization noise of A/D conversion is fairly easy to model accurately, and rigorous information-theoretic analyses of its effect are available. In some cases, line-of-sight with a single terminal, the quantization noise may combine constructively. However, in frequency-selective, Rayleigh fading channels with large delay-spreads and multi-user beamforming, the distortion averages out over the antennas to a significant extent. Specifically, with 1-bit quantization, the quantization noise has a power equal to $(\pi/2 - 1)P$ where P is the received signal power [22], and the aggregate effect of the quantization is approximately a loss in effective SINR of 4 dB. The 1-bit A/D converter case is of particular interest as it allows operation without automatic gain control (AGC), which simplifies hardware complexity. With N -bit quantization, $N > 1$, corresponding results can be found in [23], and when N grows eventually the capacity formulas for the un-quantized case [3, Ch. 3] are rediscovered. Other authors have derived similar results subsequently [24] – and earlier, using heuristic arguments, [25], [26]. Importantly, these analyses take into account the fact that both the received pilots and the payload data will be affected by quantization noise.

The loss in effective SINR due to quantization needs to be considered relatively to the extra power consumption resulting from adding bits resolution in the A/D converters. Circuit innovation in data converters has brought great improvements in power efficiency. State-of-the-art designs for A/D converter cores achieve figures-of-merit following 10 in the order of 10 fJ/cs [27], [28]. A 6-bit ADC with a speed of several 100 Mbit/s consumes < 1 mW.

Massive MIMO systems operating with low-resolution Digital-to-Analog (D/A) converters at the base station in the downlink transmission have also been studied. There is some evidence that they are sufficient to attain a good performance in terms of achievable link rate [29], [30]. Also, while these analyses are independent of the actual modulation and coding used in the system, numerical end-to-end link simulations have independently arrived at essentially the same conclusion that the degradation of BER performance due to low-precision (< 6 bits) D/A converters is negligible [31]. It is however a misconception that the number of bits resolution affects the D/A converter power consumption in a similar way as it does for A/D converters. The constraint on OOB emission in combination with the swing to be delivered to the analog output signal are the dominant factors in the power and complexity of a D/A converter [21]. A relevant standard figure-of-merit (F.o.M.) for

current-steering D/A converters is given by

$$\text{F.o.M.}_{\text{D/A}} = \frac{V_{\text{pp}} \cdot f_{\text{out}} \cdot 10^{\text{SFDR}/20}}{\text{Power}_{\text{D/A}}} \quad (11)$$

where SFDR is the spurious free dynamic range, being the distance between the signal and the largest single unwanted component – the spurious signal, and V_{pp} is the peak-to-peak signal swing which accounts for the power (and design problems) needed for generating the analog signal in a digital-to-analog converter. D/A converters with a resolution < 10 bits are conveniently implemented by current injection or resistive architectures whose power consumption is typically not directly impacted by their resolution. In contrast, the complexity of the reconstruction filter in the D/A converter is mostly determined by the SFDR specification, which will eventually determine the out-of-band (OOB) harmonic distortion. Digital predistortion and analog filtering to reduce OOB emissions have been proposed for coarsely quantized precoding in Massive MIMO [32]. The extra processing complexity in deeply scaled technology will be very reasonable, yet a degradation of the in-band signal-to-interference-noise-and-distortion ratio (SINDR) on the link is introduced. This presents the same trade-off between in-band transmission versus out-of-band rejection encountered in D/A converter design.

The trend in broadband wireless systems to increase spectral efficiency through a combination of higher order modulation constellations and conventional multi-layer MIMO has raised the resolution requirement for data converters $> \sim 12$ bits. Massive MIMO can operate without noticeable implementation loss with only 4–6-bit A/D and D/A converter resolution. This reduces the power consumption of an individual A/D converter specifically with a factor $> \sim 50$, which more than compensates for the fact that 10–30 times more converters are needed. It is however neither necessary nor overall beneficial to reduce the resolution of A/D and D/A converters below 6 bits:

- On uplink, reducing the A/D resolution further will save less than 100 mW in a 100 antenna basestation.
- On downlink, a potential implementation loss of 0.5 dB or more due to a D/A converters with a lower resolution may require 10% more power in the PA stage. More importantly, the constraints on OOB emission will not be met. Dedicated processing will hence be needed to avoid or filter out unacceptable leakage in adjacent bands.

C. Reciprocity Calibration in RF Front-Ends

Channel estimates are obtained from uplink pilots; see Section II. In practice, the response observed by the digital baseband processing for each user includes both the propagation channel and the transceiver transfer functions. The full responses for uplink and downlink can be expressed as:

$$\begin{aligned} \mathbf{g}_{k,UL} &= \mathbf{R}_B \tilde{\mathbf{g}}_k t_k \\ \mathbf{g}_{k,DL} &= r_k \tilde{\mathbf{g}}_k^T \mathbf{T}_B, \end{aligned} \quad (12)$$

where $\mathbf{R}_B$ and $\mathbf{T}_B$ are complex diagonal matrices containing the base station receiver and transmitter responses, and t_k and r_k are the responses of the transmitter and receiver of user

terminal k . While the responses of the propagation channel $\tilde{\mathbf{g}}_k$ are reciprocal, the responses of the front-ends will typically cause non-reciprocity in the full response. In the precoded Massive MIMO downlink reception the following holds:

$$\begin{aligned} \mathbf{R}_B &\neq \mathbf{T}_B \\ r_k &\neq t_k \\ \Rightarrow \mathbf{g}_{k,DL} &\neq \mathbf{g}_{k,UL}. \end{aligned} \quad (13)$$

When the corresponding estimates $\hat{\mathbf{g}}_k$ of $\mathbf{g}_{UL}$ are used to calculate the precoding coefficients, they will introduce Multi-User Interference (MUI) and potentially an SNR loss, depending on the precoding vectors $\mathbf{a}_k$. We include the derivation for the zero-forcing precoder, and refer to [33] for a comprehensive treatment. Under the assumption of negligible channel estimation errors and considering normalized responses to simplify notation,¹ the received signals at the terminals $\mathbf{y} = [y_1, \dots, y_K]^T$ are given by

$$\mathbf{y} = \mathbf{G}_{DL}^T \mathbf{G}_{UL}^* (\mathbf{G}_{UL}^T \mathbf{G}_{UL}^*)^{-1} \mathbf{x} + \mathbf{w} \quad (14)$$

where $\mathbf{x}$ and $\mathbf{w}$ are the K -vectors of transmitted symbols and received noise samples, respectively. Writing out the front-end responses gives the following expression:

$$\mathbf{y} = (\mathbf{R}_U \tilde{\mathbf{G}}^T \mathbf{T}_B) (\mathbf{R}_B^* \tilde{\mathbf{G}}^* \mathbf{T}_U^*) (\mathbf{G}_{UL}^T \mathbf{G}_{UL}^*)^{-1} \mathbf{x} + \mathbf{w}, \quad (15)$$

where $\mathbf{R}_U$ and $\mathbf{T}_U$ are diagonal matrices containing the transmitter and receiver responses of terminals t_k and r_k . Equation (15) shows that in general the combined precoding, channel, and transceiver responses will not result in a diagonal matrix. As a result, MUI will occur. Structurally it is the multiplication of the base station's front-end responses $\mathbf{T}_B \mathbf{R}_B^*$ that is responsible for the MUI. The terminal responses appear as scalar multiplications on the received symbols and will be contained in the equalization processing in the terminal. A suitable calibration procedure operating locally at the base station can restore the reciprocity. Calibration data needs to be obtained through measurements of the transceiver front-end responses, for which several approaches have been proposed and validated:

- Utilization of an auxiliary front-end, which sequentially measures the RF transceiver front-ends. The method works well in conventional MU-MIMO systems [34]. However, it does not scale well to large numbers of antennas.
- Exploitation of the coupling, essentially radio propagation, between antennas in the array to derive the relative differences among the transceiver responses. This solution has been implemented in real-life testbeds and performs well [35].

Analysis has shown that non-reciprocity requirements are not as severe for Massive MIMO as in conventional systems [33] and depend on the system load and precoding algorithms. The RF transceiver responses may vary in time mainly due to temperature differences. The calibration procedure hence needs to be repeated on a regular basis. In typical conditions the required updating frequency is in the order of hours. It thus introduces only very limited overhead.

¹Power control does not impact reciprocity, and it will show up as a scalar multiplication on the individual terminal signals.

V. ALGORITHM-HARDWARE CO-DESIGN FOR PRECODING AND DETECTION

The central detector and precoder perform the crucial operations to achieve spatial multiplexing. This section zooms in on the hardware implementation of the precoding and detection algorithms.

A. Implementation Challenges and Design Considerations

Linear processing provides good precoding and detection performance under favorable propagation conditions. However, linear processing in Massive MIMO does not necessarily result in low computational complexity given that the operations need to be performed on large matrices. For instance, the complexity of computing $(\mathbf{G}^H \mathbf{G})^{-1}$ for an $M \times K$ matrix $\mathbf{G}$ is

$$MK^2 + K^3. \quad (16)$$

This number is in the order of 10^4 for an $M = 128$, $K = 16$ system. In TDD Massive MIMO systems, processing latency is a crucial design consideration, especially for high-mobility scenarios. The analysis in [9] shows that the time budget for operating the precoding is $150 \mu\text{s}$ to support a moderate mobility of 70 km/h . The high computational complexity and processing speed need to be handled with reasonable hardware cost and power consumption. These implementation challenges necessitate meticulously optimized solutions following a systematic algorithm-hardware co-design methodology.

A central property of Massive MIMO is that the column vectors of the channel matrix are asymptotically orthogonal under favorable propagation conditions. As a result, the Gram matrix, $\mathbf{Z} = \mathbf{G}^H \mathbf{G}$, becomes diagonally dominant, i.e.,

$$|z_{i,i}| \gg |z_{i,j}|, \text{ for } i \neq j \text{ and } M \gg K, \quad (17)$$

and for i.i.d. channels,

$$\frac{1}{M} \mathbf{Z} \rightarrow \mathbf{I}, \text{ for } M \rightarrow \infty \text{ and for fixed } K. \quad (18)$$

The extent of the diagonal dominance varies with the characteristics of the antenna array, the propagation environment, and the number of users served. Exploiting this dominance, approximate matrix inversion can be performed to reduce the computational complexity. Matrix inversion approaches can be categorized into three types: explicit computation, implicit computation, and hybrid methods. We next assess the complexity and suitability of these methods.

B. Explicit Matrix Inversion

Explicit matrix inversion can be performed using approaches such as Gauss-elimination, Neumann series expansion [36], and truncated polynomial expansion [37]. Recently, the Neumann series approximation has been identified as one of the most hardware-friendly algorithms for Massive MIMO systems [38], [39]. If a $K \times K$ matrix $\mathbf{Z}$ satisfies

$$\lim_{n \rightarrow \infty} (\mathbf{I} - \mathbf{X}^{-1} \mathbf{Z})^n \simeq \mathbf{0}_K, \quad (19)$$

its inverse can be approximated by a Neumann series with L terms as:

$$\mathbf{Z}^{-1} \approx \sum_{n=0}^L (\mathbf{I} - \mathbf{X}^{-1} \mathbf{Z})^n \mathbf{X}^{-1}, \quad (20)$$

where $\mathbf{X}$ is a pre-conditioning matrix. The number of terms, L , can be used as a tuning parameter to trade off between complexity and accuracy. It is shown in [39] that using the main diagonal of the Gram matrix,

$$\mathbf{Z}_d = \text{diag}[\mathbf{Z}_{1,1}, \dots, \mathbf{Z}_{K,K}], \quad (21)$$

as the pre-conditioning matrix, the Neumann series approximation can provide close-to-exact-inversion performance with $L = 3$ when $K \ll M$. However, a significant performance loss is demonstrated when $M/K < 8$. To improve the accuracy, the following weighted Neumann series approximation was introduced in [40], [41]:

$$\mathbf{Z}^{-1} \approx \sum_{n=0}^L \alpha_n (\mathbf{I} - \mathbf{X}^{-1} \mathbf{Z})^n \mathbf{X}^{-1}. \quad (22)$$

In [40], the coefficients α_n are selected by solving the equation

$$\sum_{n=0}^{\infty} \mathbf{B}^n \approx \sum_{n=0}^L \alpha_n \mathbf{B}^n, \quad (23)$$

where

$$\mathbf{B} = -\mathbf{Z}_d^{-1/2} (\mathbf{Z} - \mathbf{Z}_d) \mathbf{Z}_d^{-1/2}. \quad (24)$$

At the price of extra computational complexity, the method in (22) improves the performance significantly, especially in cases with a high user load.

C. Implicit Matrix Inversion

Implicit matrix inversion uses linear-solvers such as conjugate-gradient [42], coordinate-descent [43], and Gauss-Seidel [44] to perform linear precoding and detection, without explicitly calculating the Gram matrix inverse. In [43], the coordinate-descent method is adopted to realize an MMSE detector. The regularized squared Euclidean distance,

$$f(\mathbf{x}) = \|\mathbf{y} - \mathbf{G}\mathbf{x}\|_2^2 + N_0 \|\mathbf{x}\|_2^2, \quad (25)$$

is minimized sequentially for each variable in $\mathbf{x}$ in a round-robin fashion. In (25), N_0 is the variance of each complex entry in the noise vector $\mathbf{w}$. In each iteration, the solution for the i th element in $\mathbf{x}$ is

$$\hat{x}_i = \frac{1}{\|\mathbf{g}_i\|_2^2 + N_0} \mathbf{g}_i^H \left(\mathbf{y} - \sum_{j \neq i} \mathbf{g}_j x_j \right). \quad (26)$$

This procedure is then repeated for L iterations.

D. Hybrid Method

Matrix decomposition algorithms factorize the Gram matrix into intermediate matrices, which are generally triangular. Forward or backward substitution is then performed to accomplish

the corresponding precoding and detection operation. The solution in [45] utilizes QR-decomposition. The Gram matrix $\mathbf{Z}$ is decomposed as

$$\mathbf{Z} = \mathbf{Q}\mathbf{R}, \quad (27)$$

where $\mathbf{Q}$ is unitary and $\mathbf{R}$ is upper triangular. The linear equation $\hat{\mathbf{s}} = \mathbf{Z}^{-1}\mathbf{s}$ is then rewritten as

$$\mathbf{R}\hat{\mathbf{s}} = \mathbf{Q}^H \mathbf{s}, \quad (28)$$

which can be solved using backward substitution. This method avoids the explicit computation of matrix inverses, relaxing (to some extent) the requirements on data representation accuracy. By exploiting the diagonally dominant property of the Gram matrix, modified QR-decomposition can be performed [45]. For instance, the original solutions

$$\begin{aligned} c &= a/r \\ s &= b^*/r \\ r &= \sqrt{|a|^2 + |b|^2} \end{aligned} \quad (29)$$

to the Givens rotation operation

$$\begin{bmatrix} c & s \\ -s^* & c \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} r \\ 0 \end{bmatrix}, \quad (30)$$

are approximated by

$$\begin{aligned} c &= c_{\text{const}} \\ s &= b^*/a. \end{aligned} \quad (31)$$

Equation (31) makes use of the fact that $|a| \gg |b|$ and results in 50% complexity savings by introducing the constant c_{const} .

Cholesky-decomposition ($\mathbf{Z} = \mathbf{L}\mathbf{L}^*$) has also been studied for Massive MIMO precoding and detection implementation [46], [47]. It has lower computational complexity than the Neumann series expansion method (with $L \geq 4$) [39] and provides accurate processing independent of M and K . More importantly, the Cholesky decomposition imposes lower memory requirements, since only the lower triangular matrix $\mathbf{L}$ needs to be stored.

E. Complexity versus Accuracy Trade-Off

To select appropriate processing algorithms for Massive MIMO is non-trivial, and an analysis of the trade-off between computational complexity and processing performance is necessary. Reference [48] presents such an analysis for different MMSE detection techniques.

To evaluate the processing accuracy, we simulate the performance of different detection techniques including Neumann series approximation (NSA), Cholesky decomposition (ChD), modified QRD (MQRD), and coordinate descent (CD). The effects of fixed-point arithmetics is also taken into consideration to examine the required data precision. In the simulations, $M = 128$, K sweeps from 8 to 32, and an i.i.d. block Rayleigh fading channel with perfect channel estimation and synchronization was considered. A rate-1/2 convolutional code with generator polynomial [171, 133] and a constraint length of 7 was used. Figure 7 shows the performance at 10^{-4} BER relative to floating-point ZF detection. The number of iterations L for the

Fig. 7. Simulated performance of different detection methods. The subscripts in the legend indicate the fixed-point resolution of the fractional part. Markers at -4 dB performance loss mean that the corresponding detection scheme has a performance loss greater than 4 dB or shows an error floor before reaching a BER of 10^{-4} .

NSA and CD was set to 3. Implicit and hybrid methods are more robust to lower resolutions, while NSA requires a larger number of bits to calculate the matrix inverse explicitly. When M/K is small the Gram matrix becomes less diagonally dominant and approximate matrix inversion methods suffer from a larger performance loss. CD offers better interference cancellation when the user load is relatively high.

Table II lists the corresponding computational complexity in terms of number of real multiplications. The computation is divided into two parts depending on how frequently it needs to be executed, i.e., per channel realization and per channel use (instance of the detection problem). The Gram matrix calculation, matrix decomposition, and matrix inversion are performed when the channel changes, while matched-filtering and backward/forward substitution are performed for each received vector. Thereby, the computational complexity depends on the channel dynamics, i.e., the number of samples (P) during which the channel is constant. Figure 8 depicts the results. Different system setups and channel conditions are analyzed. While changing M , K , and P in the three sub-figures, the other two are fixed to $M = 128$, $K = 16$, and $P = 5$, respectively. Several observations can be made. The detection complexity grows linearly with M , enabling large savings in transmit power by deploying large numbers of antennas, with a mild increase in the processing power. Moreover, the processing complexity (for explicit and hybrid matrix inversion algorithms) can be dramatically reduced in static environments, in which case the channel matrix-dependent operations are performed very rarely.

In addition to the processing accuracy and computational complexity, parallelism is an important aspect to be considered, and it highly impacts the processing latency. Iterative algorithms such as Neumann series approximation and coordinate descent can suffer from a long processing latency for MUI-dominant channels. On the other hand, matrix decomposition can be performed in a more parallel fashion and was thus selected for the first Massive MIMO precoder-detector chip introduced in the next section. Moreover, the intermediate results $\mathbf{Z}^{-1}$, $\mathbf{L}$, and $\mathbf{Q}\mathbf{R}$ can be shared between the uplink and downlink processing, further simplifying the hardware.

TABLE II
COMPUTATIONAL COMPLEXITY (NUMBER OF REAL MULTIPLICATIONS) OF DIFFERENT DETECTION TECHNIQUES

Algorithms	Per channel realization	Per channel use
Neumann Series	$2MK(K+1) + 8K^2 + 4(L-1)K^3$	$4K^2 + 4KM$
Cholesky Decomp.	$2MK(K+1) + 4K(K+1)(K+2)/6$	$4K^2 + 4K + 4KM$
Modified QR Decomp.	$2MK(K+1) + 4K^3/3 + 3K^2/2 - 31K/3$	$6K^2 - 2K + 4KM$
Coordinate Descent	-	$4M(L-1) + 4KML$

Fig. 8. Computational complexity (per instance of the detection problem) of different implementations of ZF detection, for different numbers of base station antennas, numbers of users, and channel coherence duration.

F. 128×8 Massive MIMO Precoder-Detector Chip Achieving 300 Mb/s at 60 pJ/b

Integrated hardware implementations will ultimately define both the performance and power consumption of Massive MIMO systems. Hence, algorithms should be selected such that the corresponding operations can be mapped into simple,

configurable, and scalable hardware architectures to enable high throughput, low latency, and flexible implementation. The re-configurability and scalability are essential to enable efficient operation in a wide range of conditions. In this section we present a design [45] demonstrating such an algorithm and hardware architecture co-design, where the QR-decomposition based ZF precoding is mapped onto a systolic array architecture; see Figure 9. The systolic array consists of a homogeneous network of elementary processing nodes, where each node performs the same pre-defined tasks. Due to the homogeneity, the architecture is scalable to support different M and K . The data flow in a systolic array is straightforward and parallel, leading to a simple and high-speed hardware implementation.

The QR-decomposition based precoder, together with a Cholesky decomposition based detector, was fabricated using 28 nm FD-SOI (Fully Depleted Silicon On Insulator) technology. Figure 10(a) shows a photograph of the chip. It occupies only a 1.1 mm^2 silicon area and consumes $\sim 50 \text{ mW}$ power for precoding and detection for a 128×8 Massive MIMO system with a 300 Mb/s throughput. The fabricated chip and the measurement results prove that the Massive MIMO concept works in practice and that system-algorithm-hardware co-optimization enables record energy-efficient signal processing. The cross-level design approach also applies advanced circuits techniques leveraging on the flexible FD-SOI body bias feature [50]. Using forward body bias or reverse body bias allows systems to dynamically adjust processing speed and power consumption of the chip towards the most efficient operating point.

The algorithm-hardware co-design method is further exploited in [49] to map an iterative expectation-propagation detection (EPD) onto a condensed systolic array for higher hardware resource utilization. This detector chip (Figure 10(b)) is fabricated using 28 nm FD-SOI technology and provides 1.8 Gb/s throughput with 127 mW power consumption. It offers 3 dB processing gain comparing to [45], equivalent to a $2\times$ boost in link margin that can be utilized to lower the TX power and relax the front-end requirements.

VI. PER-ANTENNA CHAIN PROCESSING AT THE SEMICONDUCTOR EDGE

An obvious concern is how the large number of antennas and the associated signal processing will affect the cost and energy consumption of the base station. The individual antenna signals may have low precision, but regardless of that, the coherent combination yields excellent SNR eventually. We demonstrate below that the resolution of digital signals and operators, such

Fig. 9. Simple, configurable, and scalable architecture for QRD-based Massive MIMO precoder (From [45]).

Fig. 10. Microphotographs of Massive MIMO precoder and detector chips: (a) From [45] (b) From [49].

as filtering coefficients, can be scaled back sharply. Furthermore, we advocate processing of the per-antenna functionality without the conventional circuit design margins that are used to cope with uncertainties in the semiconductor technology. This approach has been called “at the semiconductor’s edge” to indicate an operation point where the performance-energy benefit of the technology is maximally exploited at the expense of reliability [51]. Specifically, voltage over-scaling offers significant energy reductions in deeply scaled CMOS, up to more than 50%, at the risk of occasional processing errors. Massive MIMO systems can be designed to meet required performance levels when operating with error-prone digital signal processing circuits. Circuits remain functional even for the worst-case scenario in which the DSP circuitry in some antenna paths fails completely, for example by a broken power supply. We will call the situation where the signal in an antenna branch is fully lost “antenna outage”.

A. Per-Antenna Functions: Coarse Processing Provides Excellent Performance

Massive MIMO can operate well with low-resolution signals. A profiling of the per-antenna functionality in terms of generic operations per second shows that for an LTE-like setup, about 80% of the complexity is in the filtering and the remaining

20% is in the (I)FFT operation. The filtering functionality is the most demanding because of the need to over-sample and hence process at high speed. Significant savings in complexity are therefore possible by minimizing the resolution of this processing. An exploration of the word lengths of the data signals, $n_s f$, and of the filtering coefficients, n_f , is reported on in [52]. The circuit area complexity C_{Filt} of the T -tap FIR filtering of I- and Q-signals as a function of the word lengths is calculated using basic formulas for the complexity of adders and multipliers, which are dependent on the word-lengths n and m of the operands as follows:

$$C_{\text{add}} = n \cdot \log_2 n$$

$$C_{\text{mult}} = n \cdot m \quad (32)$$

$$C_{\text{Filt}} = 2 \cdot T \cdot (m + n) \cdot \log_2(m + n) + 2 \cdot T \cdot m \cdot n.$$

If a smaller number of bits is used to represent the signals and the filter coefficients, the hardware complexity as given in (32) is reduced. However, decreasing the word length will increase the quantization noise. For a desired transmission quality the just-sufficient precision can be determined. Considering that the quantization noise will be independent among the antennas, its combined impact will be smaller for larger numbers of antennas. This effect is illustrated in Figure 11 for the rather demanding 64-QAM case, and an uncoded Bit-Error Rate (BER) of 10^{-3} . The curves were generated based on individual BER vs. SNR simulations for different coefficient and signal resolutions, from which the equal performance points were extracted. Dotted lines show equal-complexity (in terms of area) solutions. For a 128×4 Massive MIMO system, 4 and 5 bits are sufficient for the signals and the coefficients, respectively, for the targeted performance. This brings a 62% complexity reduction for the filters compared to the 8×4 case. The outer right points on the curves are clearly always suboptimal and demonstrate that high-precision filter coefficients do not improve performance, while they can cause a significant complexity penalty. A similar observation holds for the upper left points. For higher system loads, more bits are needed. At the system level one could

Fig. 11. Representation of the relative circuit complexity (area) as function of the signal and filter coefficient word lengths. The markers show possible operating points with a BER of 10^{-3} . The dashed lines with numbers show operating points with equal complexity. The graphs demonstrate that low-resolution processing is feasible with large antenna arrays. From [52].

trade-off system load for constellation order to satisfy throughput requirements.

This analysis provides evidence that low-complexity, coarse processing in the digital filters of the individual antenna signals can offer the required performance in Massive MIMO. In the downlink the signals will next be passed to the D/A converters. The latter could be low resolution as well. The more demanding design challenge for D/A converters however is to meet out-of-band emission specifications, as introduced in Section IV.

The (I)FFT operations required in Massive MIMO systems with multicarrier modulation can also be designed for Massive MIMO operation specifically and benefit from the complexity reduction brought by the coarse quantization. A thorough optimization is however quite complex and should consider varying quantization at the different butterfly stages.

B. Processing at the Semiconductor's Edge

Applications have benefited over the last decades from Moore's law, providing ever higher performance at lower power consumption. Integrated Circuits (ICs) have been able to operate at lower dynamic power thanks to the scaling of the supply voltage V_{dd} . For digital circuits, the average dynamic power consumption is

$$P_{\text{dyn, av}} = (\alpha C) \cdot V_{\text{dd}}^2 \cdot f_s, \quad (33)$$

where αC is the effective switching capacitance of the module and f_s is the switching frequency. Clearly, $P_{\text{dyn, av}}$ scales quadratically with the supply voltage V_{dd} .

However with scaling towards deep sub-micron CMOS technologies (65 nm and smaller), designers are facing ever-increasing variability challenges. The process, voltage and temperature (PVT) variabilities are considered to be the three main contributors to circuit variability. Conventionally, to cope with this challenge, ICs are designed at the worst PVT corners, to ensure that they always operate correctly. Figure 12 illustrates the different operating regions for ICs suffering from manufacturing variability.

Fig. 12. Different approaches to scaling of the supply voltage V_{dd} to cope with speed variability. Operation at the worst-case corner misses out on potential energy savings. Adaptive Voltage Scaling (AVS) provides the just-needed V_{dd} for the circuit to function error-free. A further reduction of V_{dd} by Voltage Over-Scaling (VOS) would save more power, yet would introduce processing errors.

The conventional design approach for worst-case conditions introduces considerable margins, leading to reduced peak performance and wasted power consumption. The worst-case synthesis assumes that all devices in the circuit operate in the slow-process corner and experience the least favorable voltage and temperature conditions. Temperature variations can yield up to 20% speed differences for a single D flip-flop. For instance, [53] shows that for 28-nm technology, the performance (speed) difference for a representative circuit is as large as 2.2 times between the typical case and the worst case. Adaptive scaling techniques manage power dissipation and temperature by using a variable supply voltage V_{dd} .

Scaling down the supply voltage is regarded as an error-free power saving method as long as the signal timing constraints are met. However, the critical (minimum) V_{dd} that guarantees timing closure cannot be determined at design time due to PVT variabilities and aging effects.

A third design approach has recently gained interest, namely to scale the V_{dd} below the critical supply voltage, which is called Voltage Over-Scaling (VOS). In the VOS approach, the designer accepts that sporadic errors might occur: for logic components, the signal from the longest propagation paths can be mis-captured; for memory components, it may lead to incorrect write/read data/address or data loss. This methodology of approximate computing enables very energy-efficient processing [54]. Wireless communication systems are designed to cope with distortions and errors occurring on the channel. They are hence inherently good candidates for error-resilient processing solutions. In Massive MIMO, the large number of antennas implies redundancy in the system. It is promising to apply VOS specifically in the per-antenna processing, reaching beyond the reliability margins of the circuits, but still operating at a point where the computations are more often correct than wrong.

C. Massive MIMO Resilience to Circuit Errors

Massive MIMO inherently is resilient to some circuits errors in the per-antenna processing. Hardware errors in a number of antenna paths can be absorbed by the system thanks to the averaging induced by the large number of antennas – reminiscent of

Fig. 13. BER performance versus channel SNR. Randomly affected “victim antennas” from significant digital hardware errors for uncoded and coded (3/4 soft LDPC) QPSK, and uncoded and coded (3/4 soft LDPC) 16-QAM. From [51]. The legend denotes: i) error-free (star markers), ii) 3% victim antennas (circle markers), and iii) 10% victim antennas (triangle markers).

how the effects of small-scale fading average out in the coherent multi-user MIMO processing [3]. Semiconductor process variability was at first experienced globally, between wafers or circuits separated in space on a silicon wafer, hence die-to-die. Designers have thus realistically assumed transistor parameters to be correlated for nearby circuits on a specific die and chip. However, in deeply scaled technologies, device variability is mostly caused by the inaccuracy of lithography and etch technology. Intra-die (local) variations have consequently become significant, and are even reported dominant over global variations [55]. This apparent design challenge comes with a new opportunity to shave margins in the implementation of Massive MIMO. Indeed, different from the distortion resulting from non-linearities, the digital distortion is independent of the signal and hence uncorrelated over the antennas. The Massive MIMO system will continue functioning even when, sporadically, one or a few individual antenna signals is subject to full failure. As mentioned in Section VI-B, this opens the door to operation of circuits with much lower design margins compared to traditional specifications, and most interestingly at lower supply voltages and hence power consumption.

The digital hardware errors in (I)FFT and filters introduced by silicon unreliability and by ambitious design methodologies result in incorrect bits during signal processing. This can be regarded as digital circuit distortion. We characterize the impact on the purity of the signal in terms of the signal-to-digital distortion ratio (SDDR):

$$\text{SDDR} = 10 \cdot \log \frac{\sigma_s^2}{\sigma_d^2} \quad (34)$$

where σ_s^2 and σ_d^2 are the powers of the error-free digital antenna signal output, and the noise power of the digital distortion due to circuit unreliability, respectively. First, we consider VOS errors which are temporary and local in nature. The BER-performance is shown in Figure 13 for a severe SDDR distortion,

Fig. 14. Impact of antenna outage on Massive MIMO system performance depends on the system load, for the pessimistic case where the errors are not detected. Disabling antennas will limit the impact of antenna outage on the Massive MIMO system performance. From [51].

where signals get stuck at a fixed value. Results for different modulation orders and both uncoded and coded performance (rate 3/4 soft decoded LDPC) are shown. The resulting SNR degradation remains limited to < 1 dB for 3% of the antennas being a “victim” of circuit errors in the coded 16-QAM, and even up to 10% of the antennas in QPSK case.

When operating deeply scaled circuits without margins, occasionally a full circuit failure may occur. The impact of this effect on the Massive MIMO system performance is called “antenna outage”. The digital output are then permanently stuck at a fixed value, which is assumed to be its maximum possible value. The SDDR of the outage antenna is $-\infty$, as the signals from the victim antennas are completely lost. This model is regarded a worst-case hardware failure. Note that the $-\infty$ SDDR does not imply infinite noise to the whole system, as only the victim antennas are affected and their PA power is normalized among all antennas. Therefore, a single antenna outage will not cause the system to fail entirely. The impact on the system performance is shown in Figure 14 for different antenna outage and system loads, for the pessimistic case where the errors are not detected.

As demonstrated, Massive MIMO can operate well with rather severe circuits errors, and thus allows significant VOS. The impact increases with higher system load and modulation constellations. The V_{DD} may be adapted according to the system parameters to always offer just sufficient performance. In-situ monitoring based on test signals can be applied to perform adequate V_{DD} scaling [51].

In order to further improve the system robustness towards hardware errors, techniques to first detect hardware errors, and next either neglect, or if needed disable, defective hardware can be applied. Importantly, the distortion originating from digital circuit errors fundamentally differs from pure random noise. While process variations may feature continuous random distributions, their effect typically results in discrete error events. Dedicated monitoring circuitry can be established [56] for the

Fig. 15. BER performance is only slightly degraded for up to $\sim 10\%$ of antennas failing. Systematic failure of circuits is detected and corresponding antenna signals are discarded. From [51].

functional components such as (I)FFTs and filters, that will detect these errors. If the Massive MIMO system is operated whereby it receives information from the hardware level on failing circuits, it can adapt its signal processing accordingly. One option is to disable systematically failing antenna paths and no longer consider them in the central processing. The BER results are given in Figure 15 for a case with moderate system load (10×100 in this simulation). It shows that excluding defective circuits limits the degradation level to < 0.5 dB on uncoded QPSK for up to $\sim 10\%$ of the antenna paths failing. This approach is equivalent to operating the Massive MIMO system with a reduced number of BS antennas M . For a representative case of QPSK transmission in a 100-antenna, 10-user scenario and with 28 nm standard CMOS technology, up to 40% power savings can be achieved with negligible performance degradation [57].

In conclusion, lean per-antenna processing can be performed in Massive MIMO systems. The very large number of operations, due to the large number of antenna paths, can be performed with low precision and with a profoundly scaled supply voltage. Combined, these techniques can reduce the power consumption due to the digital processing on each antenna path by an order of magnitude. For an exemplary system with 100 antennas at the base station, the total is comparable to a conventional MIMO system with an order of magnitude less antennas.

VII. TERMINALS: INCREASED RELIABILITY WITH LOW-COMPLEXITY SIGNAL PROCESSING

A. Increased Service Levels on Low Complexity Terminals

It has been shown that the Massive MIMO system concept does not require any additional specific functionality at the UE side. Massive MIMO terminals that have a single antenna, or apply simple diversity reception, will only be able to receive a single spatial stream. However, large numbers of terminals can be multiplexed in the same time-frequency slot, and every

Fig. 16. Envisioned use cases for future international mobile telecommunication. (source: Recommendation ITU-R M.2083-0 "Framework and overall objectives of the future development of IMT for 2020 and beyond" [58])

Fig. 17. The array gain and channel hardening effect demonstrated experimentally, for a $M = 128$, $K = 8$ setup. With permission and ©Ove Edfors, Lund University.

terminal can be allocated the full bandwidth of the system. This results in a throughput per terminal comparable with that of conventional UEs that receive multiple spatial streams in parallel.

5G terminals are expected to come in large numbers and support a diverse set of service requirements. Next to the continued traffic growth towards terminals allocated to human users, a variety of devices will require Machine Type Communication (MTC). Figure 16 illustrates three main use cases envisioned by industry alliances and the International Telecommunication Union (ITU) [58].

Figure 16 demonstrates that 5G technologies not only need to enhance mobile broadband links. New solutions are needed to connect a very large number of (ultra-) low-power devices and machines requiring very reliable and low-latency services. Massive MIMO can simultaneously support many broadband terminals in sub-6 GHz bands in indoor, outdoor, and mobile environments. The technology can also be tailored to optimally serve new MTC-based applications. Especially for narrowband MTC, the high array gain and the high degree of spatial diversity offered by Massive MIMO will help. The spatial diversity specifically gives rise to channel hardening.

The effects of array gain and channel hardening are illustrated for a 128-antenna setup in Figure 17. Consistently boosted signal levels over all terminal positions, thanks to the array gain, are

Fig. 18. A conventional wideband receiver for multiple spatial layers requires complex MIMO detection.

observed. Terminals can potentially transmit data at several tens of dB lower output powers. The latter however requires high-quality CSI to be available, and the power allocated to pilots will limit the savings in practice. The channel hardening effect enhances the reliability of the links and improves the quality of service; most notably:

- 1) Increased performance at the cell edges, where terminals may experience limited or worst case no connectivity in current networks. Massive MIMO addresses this challenge, provided good uplink pilot-based CSI acquisition is ensured.
- 2) Power savings and hence longer autonomy for battery-powered devices.
- 3) Improved reliability. Fewer packet retransmissions can also reduce the end-to-end latency. The specifications put forward for Ultra Reliable Low Latency Communication (URLLC) in 5G is to support a 99.9999% reliability, and an end-to-end latency better than 1ms.
- 4) Sustained good service levels in conditions with many simultaneously active users.

In the next paragraphs, first a typical broadband user equipment is zoomed in on. It is indicated how low-power operation can be achieved while keeping backward compatibility with 4G air interfaces. Next, we discuss how tailored Massive-MIMO systems have great potential to address the challenging requirements of MTC terminals.

B. Energy Efficient Broadband Terminals

No advanced processing is required at the UE in Massive MIMO systems. In contrast, 4G systems deliver broadband services to UEs through multiplexing of several spatial layers. We compare a typical Massive MIMO terminal with the reference case of a 4×4 MIMO link. The latter requires MIMO detection at the terminal side in the downlink. Figure 18 shows a functional block scheme of a conventional broadband, multiple-antenna terminal receiver.

The complexity breakdown of a typical MIMO-OFDM baseband chain identifies channel estimation and MIMO detection as the main bottlenecks. We take as a reference 4×4 MIMO-OFDM case where the multiple-antenna processing can be conveniently performed per subcarrier resulting in relatively low-complexity implementations [59]. We consider a basic linear MIMO detector, and non-linear detectors implementing (ordered) Successive Interference Cancellation (SIC). The latter are required to achieve acceptable system performance especially in the low-SNR regime and in high-mobility scenarios. The power consumption of the inner modem receiver of the terminal in

TABLE III
RELATIVE POWER CONSUMPTION ESTIMATES FOR UE INNER MODEM RECEIVERS

4×4 linear detector	4×4 non-linear detector	Massive MIMO single antenna	Massive MIMO 2-antenna diversity
P_{ref}	$1.5 - 5P_{ref}$	$\sim 10\%P_{ref}$	$\sim 20\%P_{ref}$

TABLE IV
POWER CONSUMPTION IN DIFFERENT MODES MEASURED
ON A LPWAN IoT NODE

Operation mode	Power consumption (mW)
Transmit	$\geq 140^3$
Receive	40
Sense	13
Sleep	0.1

a Massive MIMO system is estimated relatively to published VLSI implementations for conventional MIMO receivers [60], [61]. A range of algorithms and implementations for MIMO detectors have been reported on, differing substantially in complexity. Our analysis is based on typical data for the specific components, and our own design know-how. Table III summarizes the assessment for both single-antenna and dual-antenna diversity-reception terminals, demonstrating an expected reduction in power consumption of a factor 5 to 50.² The instantaneous throughput will be higher for conventional MIMO terminals receiving several spatial layers. To compare the energy efficiency (in Joule/bit), the same average throughput needs to be considered.

C. Tailored Solutions Fit for Low-Power Connected Devices

MTC for sensors and actuators opens the door for a variety of new IoT applications. Low energy consumption is essential to enable long autonomy of devices powered by batteries or even relying on harvested energy. The physics of radio propagation dictates a strong attenuation on the link with distance, d :

$$P_{Rx} \propto G_{Tx} G_{Rx} d^{-n} P_{Tx}, \quad n = 2 \text{ in free space, } n > 2 \text{ typically,} \quad (35)$$

where P_{Rx} and P_{Tx} are the received and transmitted powers, respectively, and G_{Rx} and G_{Tx} are directivity gains at the receiving and transmitting end of the link. The above is especially unfortunate for mostly uplink-dominated MTC. Low Power Wide Area Network (LPWAN) technologies are dedicated to connect IoT nodes at long ranges. We performed measurements with an IoT node communicating via a LORA gateway [62]. Inspection of the power consumption of this illustrative node in Table IV provides valuable insights. The transmit power is relatively high since the power amplifier needs to provide sufficient power to cope with large-scale fading losses. The energy consumption, which will ultimately determine the autonomy of the node, is shown in Figure 19.

² A similar reduction in hardware complexity could be achieved for UE radios custom-designed to operate in Massive MIMO networks specifically. Backward compatibility with previous broadband systems may require the presence of MIMO detection hardware in broadband UEs in practice.

Fig. 19. The transmit energy will dominate the battery time on a LPWAN IoT node.

This pinpoints the fierce challenge of connecting sensor nodes and other autonomous devices at a longer range. Their traffic is mostly dominated by uplink, hence putting the node in the most energy-consuming transmitting mode. Equation (35) reveals that fundamentally only few parameters can be influenced to improve the link budget. Antennas at IoT nodes, due to size and cost constraints, can hardly offer any gain and on the contrary not seldom introduce losses. Massive MIMO systems exhibit a large-antenna array gain and apply an adaptive channel-matched beamforming approach. They offer the opportunity to reduce the transmit power in constrained MTC nodes proportionally to the square root of number of BS antennas M or even proportionally to M if accurate CSI is acquired. This enables the simultaneous service of a large number of devices. This asset is important to keep up with the predicted evolution towards Massive MTC. A Massive MIMO-based LPWAN could also offer extended coverage and increased reliability, provided that a power-efficient solution for the pilot-based CSI-acquisition is implemented. This challenge, to develop Massive MIMO technology for MTC services is further discussed in Section VIII.

VIII. DEMONSTRATIONS, CONCLUSION AND FUTURE DIRECTIONS

A. Signal Processing at Work in Massive MIMO Demonstrations

Demonstrations that have proven the superior spectral efficiency of Massive MIMO and the adequacy of DSP solutions in real-life testbeds are illustrated here below. Furthermore we summarize the conclusions of this paper and outline future research directions.

To prove a new wireless technology, it is very important to build up testbeds to conduct verification and evaluate performance in real-life environments with over-the-air transmission. For Massive MIMO it is especially crucial, since performance is dependent on propagation characteristics, and measurement-based channel models themselves are still under development. Thanks to recent advances in Software-Defined Radio (SDR) technology, several Massive MIMO prototype systems have been built by both industry and academia, including the Argos testbed with 96 antennas [10], Eurecom's 64-antenna testbed [64], Facebook's ARIES project [65], the 100-antenna LuMaMi testbed from Lund University (Figure 20a) [63], SEU's

Fig. 20. Two different Massive MIMO testbeds: (a) the LuMaMi testbed at Lund University a with collocated antenna array (from [63]) and (b) the KU Leuven testbed with separated antenna arrays.

128-antenna testbed [5], and testbeds exploring distributed arrays from the KU Leuven (Figure 20b) [66] and University of Bristol [67].

1) *World-Record in Spectral Efficiency and Massive MIMO in Mobility*: The signal processing techniques discussed in this paper, especially the cross-level optimization methodology, have been exploited in the development of Massive MIMO testbeds to enable real-time processing of wide-band signals for large numbers of antennas. For instance, the LuMaMi testbed adopts the processing distribution scheme in Figure 3, where 50 SDRs with Field-Programmable Gate-Arrays (FPGAs) are used to perform per-antenna processing in a parallel fashion. Four centralized FPGAs are responsible for per-subcarrier processing, and the Peripheral Component Interconnect Express (PCIe) with direct memory access (DMA) channels handles the data shuffling. QR-decomposition based ZF processing has been implemented to fully leverage the available parallel processing resources in the FPGAs.

Diverse field trials, both indoors and outdoors with static and mobile users, have been conducted using the Massive MIMO testbeds. In a 2016 experiment, a 128-antenna Massive MIMO base station served 22 users, each transmitting with 256-QAM modulation, on the same time-frequency resource [67]. The spectral efficiency benefits from the spatial multiplexing as well as from the high constellation order, enabled by the array gain. In practice, protocol overhead and FEC redundancy will determine the actual net spectral efficiency. In the actual demonstration a spectral efficiency of 145.6 bits/s/Hz was achieved on a 20 MHz radio channel, representing a ~ 20 times increase with respect to the current 4G air interface. The performance was achieved in an environment without mobility and multi-cell interference, which would constitute the limiting factors performance in a practical deployment.

The same research group also demonstrated Massive MIMO operation in an outdoor scenario with moderate mobility [4]. Figure 21 shows the measurement scenario where the 100-antenna LuMaMi testbed is placed on the rooftop of a building facing a parking lot ~ 75 m away. Ten single-antenna users are served in real time at 3.7 GHz, including six users moving at pedestrian speed and four terminals on vehicles moving at a speed up to around 50 km/h. The spatial multiplexing was fully

Fig. 21. Overview of the testbed demonstration of Massive MIMO in a mobility scenario, at the campus of Lund University, Sweden.

achieved and the communication quality was on average well maintained for all terminals [68]. Sporadic interruptions could be traced back to temporary loss of synchronization. It should be noted that both the speed of the cars and the number of terminals could be larger in a real deployment. In the proof of concept they were limited by the available test space and equipment. In fact, at 3.7 GHz carrier frequency and with a slot length of 0.5 ms, the maximum permitted mobility (assuming a two-ray model with Nyquist sampling, and a factor-of-two design margin, as in [3]) is over 140 km/h [69].

2) *Further Investigation Needed for Synchronization:* A critical challenge requiring further investigation is the initial synchronization between the base station and the user terminals. This initial synchronization has to start without any knowledge of the channels, and therefore cannot benefit from an array gain. How to efficiently perform initial time and frequency synchronization acquisition without the Massive array gain and how to explore the (partial) array gain to provide faster and more robust synchronization are still open questions. Two methods were studied during the LuMaMi testbed experiments. One method is to reserve a dedicated RF chain for the synchronization signal, which is transmitted using an omni-directional antenna. In this case, a higher-power PA (which is not available in LuMaMi) is needed to provide coverage. Another method is to use beam-sweeping for the synchronization signal [70], but this is inefficient, as it is essentially equivalent to repetition coding, and also there is risk of synchronization loss when the users are not hit by a beam. Improved techniques, based on space-time block codes, have been investigated [71]–[73]. Iterative search and tracking methods [74] may have potential, especially for mobile users.

B. Concluding on the Signal Processing

Appropriate co-design of algorithms, hardware architectures, and circuits in Massive MIMO implementations brings significant benefits:

- Energy efficient implementations of “theoretically optimal” Massive MIMO DSP architectures are nontrivial but possible. We have detailed some of the most important innovations required, and explained their analysis. The

power consumption of conventional macro base stations is dominated by the PA stage. They benefit in Massive MIMO from the ability to operate on an order of magnitude less transmit power.

- The sufficiency of low-precision quantization and processing, predicted by information-theoretic studies, has now also been validated through real signal processing experiments. A reduction in word-length up to 6 times compared to conventional systems translates into corresponding savings in complexity, power consumption and memory.
- Dedicated and scalable hardware architectures implementing tailored algorithms for large matrix processing facilitate zero-forcing precoding at the base station in real time, at 30 mW power consumption in relevant scenarios for a 128×8 system.
- Voltage over-scaling, a speculative concept just 5 years ago, has found appropriate application in the Massive MIMO per-antenna processing.
- Smart control of algorithmic modes and scalable devices, including body bias adaptation, can guarantee suitable performance-power trade-offs over a wide range of communication scenarios and channel propagation conditions.
- Lean terminals could operate in typical broadband cellular Massive MIMO networks at about 10%–20% of the power consumption of equivalent conventional terminals, both in data transmission and reception.
- The efficiency of Massive MIMO base stations can be further improved by relaxing the requirements of the RF and analogue hardware. However, caution is needed as (non-linear) distortion may under specific conditions combine coherently.

C. Future Directions

1) *Progress Massive MIMO Deployment in Actual Networks:* Integration of all components into deployment in actual networks represents a vast design and development effort, that will include:

- Overcoming challenges related to connection of the many antenna paths to the central processing units. This involves implementing high-speed interconnects and coping with potential coupling effects in the front-end modules.
- Devising efficient schedulers for large numbers of users. Achieving the high spatial multiplexing gains offered by Massive MIMO fundamentally requires that many terminals are scheduled for service simultaneously. Tuning or re-design of higher-layer protocols could be beneficial to shape the traffic patterns, such that aggressive spatial multiplexing can be performed.
- Designing antenna arrays. Massive MIMO arrays do not have to be linear, rectangular or cylindrical. Small antenna elements could be naturally integrated into the environment, onto the surface of existing structures, or faces of buildings, for example, in an aesthetically pleasing manner. Insights from electromagnetics may guide the design of new types of arrays. Specifically, for a given volume V , consider the corresponding smallest possible sphere that

contains V . If one covers the surface of this sphere with antennas at a density of $\sim 1/\lambda^2$ elements per square meter, then there is no point in installing any additional elements inside of the interior of V [75]. Sampling the surface on a $\lambda \times \lambda$ -grid captures all information in the radiated field. In conclusion, what goes into the interior of V is unimportant, only the surface matters.

Industrial recognition of the value of Massive MIMO technology is evidenced by the large number of contributions on the topic in the 3GPP-LTE standardization of New Radio (NR) for 5G systems. Leading operators have already started to perform commercial field trials of the technology [69].

2) *Enhanced Functionalities*: Large antenna arrays can also be used to perform accurate positioning and localization. This feature can offer improved context-awareness to services. Also the Massive MIMO communication system itself could exploit this information to perform smart pilot allocation, for example.

3) *Scale Up Capacity and Efficiency*: The call for more and higher-quality wireless services is expected to increase for many years, and the quest for wireless systems offering higher spectral and energy efficiency will continue. Higher peak-rates can be offered in Massive MIMO by performing spatial multiplexing of several streams to one terminal. Actual gains may be limited due to insufficient rank of the channel, yet for two streams this will mostly be achievable with co-located antennas exploiting cross-polarization.

Wider bandwidth channels can be allocated especially in mmWave bands. Radio propagation and in particular absorption is considerably different at these frequencies. Arrays with a large number of antennas can be small in size, yet their effective gain may suffer from high losses on the interconnect. Consequently, Massive MIMO systems in these bands call for other architectures and their deployment will best suit particular use cases, for example hotspots.

With larger antenna arrays, both better spatial multiplexing and array gains can be achieved. The new concepts of cell-free Massive MIMO [76] and intelligent surfaces [77] accelerate this trend to a next level. With cell-free Massive MIMO, coherently cooperating antennas are spread out over a larger geographical area, providing improved macro-diversity and improved channel rank for multiple-antenna terminals. The intelligent surface concept envisages distributed nodes that form electromagnetically active walls, floors, and planar objects. New research is urgently needed to bring these new concepts to their full potential.

ACKNOWLEDGMENT

The authors would like to thank their colleagues and collaborators in the European FP7-MAMMOET project for the nice cooperation that truly progressed Massive MIMO Technology.

REFERENCES

- [1] "First 5G NR specs approved," Dec. 2017. [Online]. http://www.3gpp.org/news-events/3gpp-news/1929-nsa_nr_5g
- [2] E. G. Larsson, O. Edfors, F. Tufvesson, and T. L. Marzetta, "Massive MIMO for next generation wireless systems," *IEEE Comm. Mag.*, vol. 52, no. 2, pp. 186–195, Feb. 2014.
- [3] T. L. Marzetta, E. G. Larsson, H. Yang, and H. Q. Ngo, *Fundamentals of Massive MIMO*. Cambridge, U.K.: Cambridge Univ. Press, 2016.
- [4] P. Harris *et al.*, "Performance characterization of a real-time Massive MIMO system with LOS mobile channels," *IEEE J. Sel. Areas Commun.*, vol. 35, no. 6, pp. 1244–1253, Jun. 2017.
- [5] X. Yang *et al.*, "Design and implementation of a TDD-based 128-antenna Massive MIMO prototype system," *China Commun.*, vol. 14, no. 12, pp. 162–187, Dec. 2017.
- [6] A. F. Molisch *et al.*, "Hybrid beamforming for Massive MIMO: A survey," *IEEE Commun. Mag.*, vol. 55, no. 9, pp. 134–141, Sep. 2017.
- [7] E. Björnson *et al.*, "Massive MIMO networks: Spectral, energy, and hardware efficiency," *Found. Trends® Signal Process.*, vol. 11, pp. 154–655, 2017.
- [8] X. Li, E. Björnson, E. G. Larsson, S. Zhou, and J. Wang, "Massive MIMO with multi-cell MMSE processing: Exploiting all pilots for interference suppression," *EURASIP J. Wireless Commun. Netw.*, vol. 117, no. 1, 2017.
- [9] S. Malkowsky *et al.*, "Implementation of low-latency signal processing and data shuffling for TDD Massive MIMO systems," in *Proc. IEEE Int. Workshop Signal Process. Syst.*, Dec. 2016, pp. 260–265.
- [10] C. Shepard *et al.*, "Argos: Practical many-antenna base stations," in *Proc. ACM Int. Conf. Mobile Comput. Netw.*, Istanbul, Turkey, Aug. 2012, pp. 53–64.
- [11] A. Puglielli *et al.*, "Design of energy- and cost-efficient Massive MIMO arrays," *Proc. IEEE*, vol. 104, no. 3, pp. 586–606, Mar. 2016.
- [12] K. Li, R. R. Sharan, Y. Chen, T. Goldstein, J. R. Cavallaro, and C. Studer, "Decentralized baseband processing for Massive MU-MIMO systems," *IEEE J. Emerg. Sel. Topics Circuits Syst.*, vol. 7, no. 4, pp. 491–507, Dec. 2017.
- [13] Q. Yang *et al.*, "Bigstation: Enabling scalable real-time signal processing in large MU-MIMO systems," in *Proc. ACM SIGCOMM*, Aug. 2013, pp. 399–410.
- [14] E. Bertilsson, O. Gustafsson, and E. G. Larsson, "A scalable architecture for Massive MIMO base stations using distributed processing," in *Proc. Asilomar Conf. Signals, Syst. Comput.*, Nov. 2016, pp. 864–868.
- [15] G. Auer *et al.*, "How much energy is needed to run a wireless network?" *IEEE Wireless Commun.*, vol. 18, no. 5, pp. 40–49, Oct. 2011.
- [16] C. Mollén, U. Gustavsson, T. Eriksson, and E. G. Larsson, "Impact of spatial filtering on distortion from low-noise amplifiers in Massive MIMO base stations," *IEEE Trans. Commun.*, to be published.
- [17] Y. Li, J. Lopez, P. H. Wu, W. Hu, R. Wu, and D. Y. C. Lie, "A SiGe envelope-tracking power amplifier with an integrated CMOS envelope modulator for mobile WiMAX/3GPP LTE transmitters," *IEEE Trans. Microw. Theory Techn.*, vol. 59, no. 10, pp. 2525–2536, Oct. 2011.
- [18] P. Horlin and A. Bourdoux, *Digital Compensation for Analog Front-Ends: A New Approach to Wireless Transceiver Design*. Hoboken, NJ, USA: Wiley, 2008.
- [19] E. G. Larsson and L. Van der Perre, "Out-of-band radiation from antenna arrays clarified," *IEEE Wireless Commun. Lett.*, to be published.
- [20] C. Mollén, U. Gustavsson, T. Eriksson, and E. G. Larsson, "Spatial characteristics of distortion radiated from antenna arrays with transceiver nonlinearities," arXiv:1711.02439, 2017. [Online]. Available: <http://arxiv.org/abs/1711.02439>
- [21] M. Pelgrom, *Analog-to-Digital Conversion*. Berlin, Germany: Springer-Verlag, 2017.
- [22] C. Mollén, J. Choi, E. G. Larsson, and J. R. W. Heath, "Uplink performance of wideband Massive MIMO with one-bit ADCs," *IEEE Trans. Wireless Commun.*, vol. 16, no. 1, pp. 87–100, Jan. 2017.
- [23] C. Mollén, J. Choi, E. G. Larsson, and R. W. Heath, Jr., "Achievable uplink rates for Massive MIMO with coarse quantization," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, Mar. 2017, pp. 6488–6492.
- [24] Y. Li, C. Tao, G. Seco-Granados, A. Mezghani, A. L. Swindlehurst, and L. Liu, "Channel estimation and performance analysis of one-bit Massive MIMO systems," *IEEE Trans. Signal Process.*, vol. 65, no. 15, pp. 4075–4089, Aug. 2017.
- [25] L. Fan, S. Jin, C.-K. Wen, and H. Zhang, "Uplink achievable rate for Massive MIMO with low-resolution ADC," *IEEE Commun. Lett.*, vol. 19, no. 12, pp. 2186–2189, Dec. 2015.
- [26] J. Zhang, L. Dai, S. Sun, and Z. Wang, "On the spectral efficiency of Massive MIMO systems with low-resolution ADCs," *IEEE Commun. Lett.*, vol. 20, no. 5, pp. 842–845, May 2016.
- [27] G. Van der Plas and B. Verbruggen, "A 150 MS/s 133 μ W 7 bit ADC in 90nm digital CMOS," *IEEE J. Solid-State Circuits*, vol. 43, no. 12, pp. 2631–2640, Dec. 2008.
- [28] K. Choo, J. Bell, and M. Flynn, "Area-efficient 1GS/s 6b SAR ADC with charge-injection-cell-based DAC," in *Proc. IEEE Int. Solid-State Circuits Conf.*, Feb. 2016, pp. 460–461.

- [29] Y. Li, C. Tao, A. L. Swindlehurst, A. Mezghani, and L. Liu, "Downlink achievable rate analysis in Massive MIMO systems with one-bit DACs," *IEEE Commun. Lett.*, vol. 21, no. 7, pp. 1669–1672, Jul. 2017.
- [30] S. Jacobsson, G. Durisi, M. Coldrey, T. Goldstein, and C. Studer, "Quantized precoding for Massive MU-MIMO," *IEEE Trans. Commun.*, vol. 65, no. 11, pp. 4670–4684, Nov. 2017.
- [31] C. Desset and L. Van der Perre, "Validation of low-accuracy quantization in Massive MIMO and constellation EVM analysis," in *Proc. Eur. Conf. Netw. Commun.*, Jun. 2015, pp. 21–25.
- [32] S. Jacobsson, G. Durisi, M. Coldrey, and C. Studer, "On out-of-band emissions of quantized precoding in Massive MU-MIMO-OFDM," in *Proc. Asilomar Conf. Signals, Syst., Comput.*, Oct. 2017, pp. 21–26.
- [33] MAMMOET project deliverable D2.4. 2015. [Online]. Available: <https://mammoet-project.eu/downloads/publications/deliverables/MAMMOET-D2.4-Analysis-non-reciprocity-impact-PU-M20.pdf>
- [34] A. Bourdoux, B. Come, and N. Khaled, "Non-reciprocal transceivers in OFDM/SDMA systems: Impact and mitigation," in *Proc. Radio Wireless Conf.*, Aug. 2003, pp. 183–186.
- [35] J. Vieira, F. Rusek, O. Edfors, S. Malkowsky, L. Liu, and F. Tufvesson, "Reciprocity calibration for Massive MIMO: Proposal, modeling and validation," *IEEE Trans. Wireless Commun.*, vol. 16, no. 5, pp. 3042–3056, Mar. 2017.
- [36] D. Zhu, B. Li, and P. Liang, "On the matrix inversion approximation based on Neumann series in Massive MIMO systems," in *Proc. IEEE Int. Conf. Commun.*, Jun. 2015, pp. 1763–1769.
- [37] A. Mueller, A. Kammoun, E. Björnson, and M. Debbah, "Linear precoding based on polynomial expansion: Reducing complexity in Massive MIMO," *EURASIP J. Wireless Commun. Netw.*, vol. 63, Dec. 2016.
- [38] H. Prabhu, O. Edfors, J. Rodrigues, L. Liu, and F. Rusek, "Hardware efficient approximative matrix inversion for linear pre-coding in Massive MIMO," in *Proc. IEEE Int. Symp. Circuits Syst.*, Jun. 2014, pp. 260–265.
- [39] M. Wu, B. Yin, G. Wang, C. Dick, J. Cavallaro, and C. Studer, "Large-scale MIMO detection for 3GPP LTE: Algorithms and FPGA implementations," *IEEE J. Sel. Topics Signal Process.*, vol. 8, no. 5, pp. 916–929, Oct. 2014.
- [40] K. Lee and C. Chen, "An eigen-based approach for enhancing matrix inversion approximation in Massive MIMO systems," *IEEE Trans. Veh. Technol.*, vol. 66, no. 6, pp. 5483–5487, Jun. 2017.
- [41] B. Nagy, M. Elsabrouty, and S. Elramly, "Fast converging weighted Neumann series precoding for Massive MIMO systems," *IEEE Wireless Commun. Lett.*, vol. 7, no. 2, pp. 154–157, Oct. 2017.
- [42] B. Yin, M. Wu, J. R. Cavallaro, and C. Studer, "Conjugate gradient-based soft-output detection and precoding in Massive MIMO systems," in *Proc. IEEE Global Commun. Conf.*, Dec. 2014, pp. 3696–3701.
- [43] M. Wu, C. Dick, J. Cavallaro, and C. Studer, "High-throughput data detection for Massive MU-MIMO-OFDM using coordinate descent," *IEEE Trans. Circuit Syst. I, Reg. Papers*, vol. 63, no. 12, pp. 2357–2367, Dec. 2016.
- [44] X. Gao, L. Dai, J. Zhang, S. Han, and I. Chih-Lin, "Capacity-approaching linear precoding with low-complexity for large-scale MIMO systems," in *Proc. IEEE Int. Conf. Commun.*, Jun. 2015, pp. 1577–1582.
- [45] H. Prabhu, J. Rodrigues, L. Liu, and O. Edfors, "A 60 pJ/b 300 Mb/s 128×8 Massive MIMO precoder-detector in 28 nm FD-SOI," in *Proc. IEEE Int. Solid-State Circuits Conf.*, Feb. 2017, pp. 60–61.
- [46] H. A. J. Alshamary and W. Xu, "Efficient optimal joint channel estimation and data detection for Massive MIMO systems," in *Proc. IEEE Int. Symp. Inf. Theory*, Jul. 2016, pp. 875–879.
- [47] R. Gangarajiah, H. Prabhu, O. Edfors, and L. Liu, "A Cholesky decomposition based Massive MIMO uplink detector with adaptive interpolation," in *Proc. IEEE Int. Symp. Circuits Syst.*, May 2017, pp. 1–4.
- [48] M. Cirkic and E. Larsson, "On the complexity of very large multi-user MIMO detection," in *Proc. IEEE Int. Workshop Signal Process. Adv. Wireless Commun.*, Jun. 2014, pp. 55–59.
- [49] T. Wei, H. Prabhu, L. Liu, O. Viktor, and Z. Zhengya, "A 1.8 Gb/s 70.6 pJ/b 12816 link-adaptive near-optimal Massive MIMO detector in 28nm UTBB-FDSOI," in *Proc. IEEE Int. Solid-State Circuits Conf.*, Feb. 2018, pp. 224–226.
- [50] N. Planes *et al.*, "28 nm FDSOI technology platform for high-speed low-voltage digital applications," in *Proc. IEEE Symp. VLSI Technol.*, Jun. 2012, pp. 133–134.
- [51] Y. Huang, C. Desset, A. Bourdoux, W. Dehaene, and L. Van der Perre, "Massive MIMO processing at the semiconductor edge: Exploiting the system and circuit margins for power savings," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, Mar. 2017, pp. 3474–3478.
- [52] S. Gunnarsson, M. Bortas, Y. Huang, C.-M. Chen, L. Van der Perre, and O. Edfors, "Lousy processing increases energy efficiency in Massive MIMO systems," in *Proc. Eur. Conf. Netw. Commun.*, Jun. 2017, pp. 1–5.
- [53] Y. Huang, M. Li, C. Li, P. Debacker, and L. Van der Perre, "Computation-skip error mitigation scheme for power supply voltage scaling in recursive applications," *J. Signal Process. Syst.*, vol. 84, no. 3, pp. 413–424, Sep. 2016.
- [54] J. Han and M. Orshansky, "Approximate computing: An emerging paradigm for energy-efficient design," in *Proc. Eur. Test Symp.*, May 2013, pp. 1–6.
- [55] S. K. Saha, "Modeling process variability in scaled CMOS technology," *IEEE Des. Test Comput.*, vol. 27, no. 2, pp. 8–16, Mar. 2010.
- [56] M. Fojtik *et al.*, "Bubble Razor: An architecture-independent approach to timing-error detection and correction," in *Proc. IEEE Int. Solid-State Circuits Conf.*, Feb. 2012, pp. 488–490.
- [57] Y. Liu, T. Zhang, and K. K. Parhi, "Computation error analysis in digital signal processing systems with overscaled supply voltage," *IEEE Trans. Very Large Scale Integr. Syst.*, vol. 18, no. 4, pp. 517–526, Apr. 2010.
- [58] ITU vision on 5G usage scenarios. 2015. [Online]. Available: https://www.itu.int/dms_pubrec/itu-r/rec/m/R-REC-M.2083-0-201509-I!!PDF-E.pdf
- [59] P. Vandenameele, L. Van der Perre, M. G. E. Engels, B. Gyselinckx, and H. J. D. Man, "A combined OFDM/SDMA approach," *IEEE J. Sel. Areas Commun.*, vol. 18, no. 11, pp. 2312–2321, Nov. 2000.
- [60] Y. Shingo and Y. Miyana, "VLSI implementation of a 4×4 MIMO-OFDM transceiver with an 80-MHz channel bandwidth," in *Proc. IEEE Int. Symp. Circuits Syst.*, May 2009, pp. 1743–1746.
- [61] J. Ketonen, M. Juntti, and J. R. Cavallaro, "Performance-complexity comparison of receivers for a LTE MIMO-OFDM system," *IEEE Trans. Signal Process.*, vol. 58, no. 6, pp. 3360–3372, Jun. 2010.
- [62] Connecting sensors with low power wireless technologies. 2018. [Online]. Available: <https://dramco.be/tutorials/low-power-lora/>
- [63] S. Malkowsky *et al.*, "The world's first real-time testbed for Massive MIMO: Design, implementation, and validation," *IEEE Access*, vol. 5, pp. 9073–9088, May 2017.
- [64] "Openairinterface Massive MIMO testbed: A 5G innovation platform," 2018. [Online]. <http://www.openairinterface.org/>
- [65] "Introducing Facebook's new terrestrial connectivity systems—terragraph and project ARIES," 2016. [Online]. <https://code.facebook.com/posts/1072680049445290/introducing-facebook-s-new-terrestrial-connectivity-systems-terragraph-and-project-aries/>
- [66] C. M. Chen, V. Volskiy, A. Chiumento, L. Van der Perre, G. A. E. Vandenbosch, and S. Pollin, "Exploration of user separation capabilities by distributed large antenna arrays," in *Proc. IEEE Global Commun. Conf. Workshops*, Dec. 2016, pp. 1–6.
- [67] P. Harris *et al.*, "Serving 22 users in real-time with a 128-antenna Massive MIMO testbed," in *Proc. IEEE Int. Workshop Signal Process. Syst.*, Oct. 2016, pp. 266–272.
- [68] MAMMOET project deliverable D4.2. 2016. [Online]. Available: <https://mammoet-project.eu/downloads/publications/deliverables/MAMMOET-D4.2-Testbed-assessment-PU-M33.pdf>
- [69] "Massive MIMO in mobile environments," in the *Massive MIMO blog*, 2016. [Online]. <http://Massive-mimo.net>
- [70] C. Barati *et al.*, "Directional cell discovery in millimeter wave cellular networks," *IEEE Trans. Wireless Commun.*, vol. 14, no. 2, pp. 6664–6678, Dec. 2015.
- [71] M. Karlsson, E. Björnson, and E. G. Larsson, "Performance of in-band transmission of system information in Massive MIMO systems," *IEEE Trans. Wireless Commun.*, vol. 17, no. 3, pp. 1700–1712, Mar. 2018.
- [72] X. G. Xia and X. Gao, "A space-time code design for omnidirectional transmission in Massive MIMO systems," *IEEE Wireless Commun. Lett.*, vol. 5, no. 5, pp. 512–515, Oct. 2016.
- [73] X. Meng, X. Gao, and X. G. Xia, "Omnidirectional precoding based transmission in Massive MIMO systems," *IEEE Trans. Commun.*, vol. 64, no. 1, pp. 174–186, Jan. 2016.
- [74] G. Marco, M. Mezzavilla, and M. Zorzi, "Initial access in 5G mmwave cellular networks," *IEEE Commun. Mag.*, vol. 54, no. 11, pp. 40–47, Nov. 2016.
- [75] M. Franceschetti, "On Landau's eigenvalue theorem and information cut-sets," *IEEE Trans. Inf. Theory*, vol. 61, no. 9, pp. 5042–5051, Sep. 2015.
- [76] H. Q. Ngo, A. Ashikhmin, H. Yang, E. G. Larsson, and T. L. Marzetta, "Cell-free Massive MIMO versus small cells," *IEEE Trans. Wireless Commun.*, vol. 16, no. 3, pp. 1834–1850, Mar. 2017.
- [77] S. Hu, F. Rusek, and O. Edfors, "Beyond Massive MIMO: The potential of data transmission with large intelligent surfaces," *IEEE Trans. Signal Process.*, vol. 66, no. 10, pp. 2746–2758, May 2018.

Liesbet Van der Perre received the Ph.D degree from KU Leuven, Leuven, Belgium, in 1997. She was with the nano-electronics research institute IMEC in Belgium from 1997 to 2015, where she took up responsibilities as a Senior Researcher, System Architect, a Project Leader, and a Program Director. She was appointed Honorary Doctor at Lund University, Lund Sweden, in 2015. She was a part-time Professor with the University of Antwerp, Belgium, from 1998 to 2002. She is currently a Professor with the Department of Electrical Engineering, KU Leuven, Leuven,

Belgium, and a Guest Professor with the Electrical and Information Technology Department, Lund University, Lund, Sweden. She has authored and co-authored over 300 scientific publications and the paper winning the DAC/ISSCC 2006 design contest. Her research interests include wireless communication, with a focus on physical layer and energy efficiency in transmission, implementation, and operation. She was the Scientific Leader of FP7-MAMMOET, Europe's prime project on Massive MIMO technology. She has been serving as a Scientific and Technological Advisor, Reviewer and Jury Member for companies, institutes, and funding agencies. She is a Member of the Board of Directors of the company Zenitel since 2015. She was a system architect for the OFDM ASICs listed in the IEEE International Solid State Circuit Conference Best of 50 Years papers in 2003.

Liang Liu received the Ph.D degree from Fudan University, Shanghai, China, in 2010. In 2010, he was a Visiting Researcher with Electrical, Computer and Systems Engineering Department, Rensselaer Polytechnic Institute, USA. He joined Lund University as a Postdoctoral Researcher in 2010. He is an Associate Professor with Electrical and Information Technology Department, Lund University, Lund, Sweden. His research interest includes signal processing for wireless communication and digital integrated circuits design. He is active in several EU and Swedish

national projects, including FP7 MAMMOET, VINNOVA SoS, and SSF HiPEC, DARE. He is a Board Member of the IEEE Swedish Solid-State Circuits/Circuits and Systems chapter. He is also a Member of the Technical Committees of VLSI systems and applications and CAS for communications of the IEEE Circuit and Systems Society.

Erik G. Larsson (F'16) received the Ph.D. degree from Uppsala University, Uppsala, Sweden, in 2002. He was with the KTH Royal Institute of Technology, Stockholm, Sweden, the George Washington University, USA, the University of Florida, USA, and Ericsson Research, Sweden. In 2015, he was a Visiting Fellow with Princeton University, USA, for four months. He is currently a Professor of Communication Systems with Linköping University, Linköping, Sweden. He has co-authored more than 150 journal papers on these topics, the two Cambridge University

Press textbooks *Space-Time Block Coding for Wireless Communications* (2003) and *Fundamentals of Massive MIMO* (2016). He is co-inventor on 18 issued and many pending patents on wireless technology. His research interests include the areas of wireless communications and signal processing. He is a member of the IEEE Signal Processing Society Awards Board during 2017–2019. He is an editorial board member of the IEEE SIGNAL PROCESSING MAGAZINE during 2018–2020. From 2015 to 2016, he was a Chair of the IEEE Signal Processing Society SPCOM Technical Committee. From 2014 to 2015, he was a Chair of the steering committee for the IEEE WIRELESS COMMUNICATIONS LETTERS. He was the General Chair of the Asilomar Conference on Signals, Systems and Computers in 2015, and its Technical Chair in 2012. He was an Associate Editor for, among others, the IEEE TRANSACTIONS ON COMMUNICATIONS (2010–2014) and the IEEE TRANSACTIONS ON SIGNAL PROCESSING (2006–2010).

He was the recipient of the IEEE Signal Processing Magazine Best Column Award twice, in 2012 and 2014, the IEEE ComSoc Stephen O. Rice Prize in Communications Theory in 2015, the IEEE ComSoc Leonard G. Abraham Prize in 2017, and the IEEE ComSoc Best Tutorial Paper Award in 2018.