



LUND UNIVERSITY

Real-Time Implementation Aspects of Large Intelligent Surfaces

Tataria, Harsh; Tufvesson, Fredrik; Edfors, Ove

Published in:

IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) 2020

DOI:

[10.1109/ICASSP40776.2020.9053897](https://doi.org/10.1109/ICASSP40776.2020.9053897)

2020

Document Version:

Peer reviewed version (aka post-print)

[Link to publication](#)

Citation for published version (APA):

Tataria, H., Tufvesson, F., & Edfors, O. (2020). Real-Time Implementation Aspects of Large Intelligent Surfaces. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) 2020* (pp. 9170-9174). IEEE - Institute of Electrical and Electronics Engineers Inc.. <https://doi.org/10.1109/ICASSP40776.2020.9053897>

Total number of authors:

3

Creative Commons License:

Unspecified

General rights

Unless other specific re-use rights are stated the following general rights apply:

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Read more about Creative commons licenses: <https://creativecommons.org/licenses/>

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

LUND UNIVERSITY

PO Box 117
221 00 Lund
+46 46-222 00 00

REAL-TIME IMPLEMENTATION ASPECTS OF LARGE INTELLIGENT SURFACES

Harsh Tataria, Fredrik Tufvesson, and Ove Edfors

Department of Electrical and Information Technology, Lund University, Lund, Sweden
e-mail: {harsh.tataria, fredrik.tufvesson, ove.edfors}@eit.lth.se

ABSTRACT

With the potential to provide a clean break from massive multiple-input multiple-output, *large intelligent surfaces (LISs)* have recently received a thrust of research interest. Various proposals have been made in the literature to define the exact functionality of LISs, ranging from fully *active* to largely *passive* solutions. Nevertheless, almost all studies in the literature investigate the fundamental spectral efficiency performance of these architectures. In stark contrast, this paper investigates the *implementation aspects* of LISs. Using the *fully active* LIS as the basis of our exposition, we first present a rigorous discussion on the relative merits and disadvantages of possible implementation architectures from a *radio-frequency circuits* and *real-time processing* viewpoints. We then show that a *distributed* architecture based on a common module interfacing a smaller number of antennas can be scalable. To avoid severe losses with analog signal distribution, multiple common modules can be interconnected via a digital nearest-neighbor network. Furthermore, we show that with such a design, the maximum backplane throughput scales with the number of served *user terminals*, instead of the number of antennas across the surface. The discussions in the paper can serve as a guideline toward the real-time design and development of LISs.

Index Terms— Backplane complexity, common modules, distributed architectures, LIS, real-time implementation.

1. INTRODUCTION

Massive multiple-input multiple-output (MIMO) systems are now starting to appear in commercial fifth-generation wireless networks around the world [1]. With tens-to-hundreds of service antennas at a cellular base station (BS), massive MIMO can achieve significant spatial multiplexing gains [2–6]. In addition, linear digital beamforming techniques to serve multiple user terminals within the same time-frequency resource have resulted in close to optimal performance over real propagation channels [7]. Relative to fourth-generation systems, this has resulted in an order-of-magnitude increase in the system spectral efficiency. Since the inception of massive MIMO in 2010 [5], phenomenal progress has been made in order to understand its theoretical and implementation aspects (see e.g., [2, 3, 5–10] for a taxonomy). To this end, our understanding of massive MIMO systems has greatly matured over almost a decade.

As a result, researchers in academia and industry have now begun to investigate what lies *beyond massive MIMO*. The authors in [2] highlight *five* promising research directions for antenna array-based technologies. One possible direction is the use of a *large intelligent surface (LIS)* [11–13]. As stated in [11], the original intention of LISs is that *man-made structures* in the environment (such as building walls) which have substantially large physical apertures can be made *electromagnetically active*. This implies that each part of the surface can transmit and receive electromagnetic radiation,

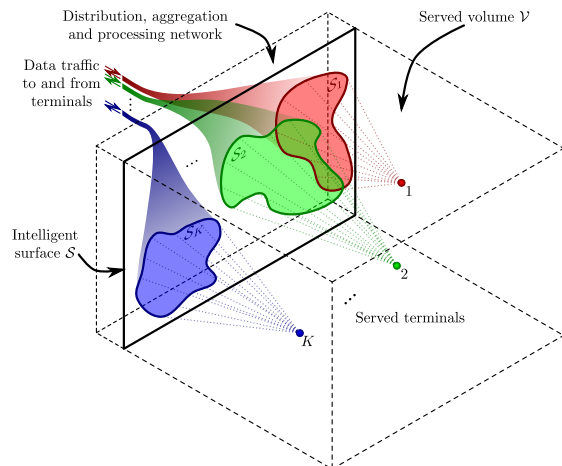


Fig. 1. A fully active LIS, \mathcal{S} , serving K randomly located terminals in volume \mathcal{V} with the necessary backplane processing.

acting as a BS with fully digital functionality. As a result of this, each antenna element on the surface needs to be able to transmit and receive radio-frequency (RF) signals. In this context, the term *intelligent* implies that the surface is able to operate *autonomously*, and is able to *adapt* to changes in the propagation channel. Prior to the seminal work of [11], the authors of [14] have made an attempt to study the impact of LISs on the overall communication system via University of California Berkley’s *e-Wallpaper* project. Since then, numerous studies have analyzed and evaluated the fundamental information-transfer capabilities of LISs, with and without hardware impairments, see e.g., [11, 15, 16]. Collectively, it has been shown that an *infinite continuous* LIS can be replaced by its *discrete* version *without* loss in performance. Nonetheless, due to its large physical size and high backplane data throughput (generated from the many elements across the surface), such architectures may not be realized with *centralized* processing. This is *unlike* the case for massive MIMO systems where centralized processing, designed in a careful manner, does yield in practically implementable solutions [9, 17]. Henceforth, from a theoretical performance viewpoint, research on *distributed* architectures for LIS is under way [18], where a smaller number of distributed units, each having a separate signal processing chain are connected to a centralized processor. Figure 1 depicts the concept of an active LIS, \mathcal{S} , with the required backplane processing to serve K terminals within a given three-dimensional volume \mathcal{V} . It is important to note that the surface as a *whole* does not need to serve a given terminal, since some terminals may be in closer proximity to a given part of the surface relative to others. Instead, for cost effective operation, smaller *regions* (subsets) of the overall surface can be *activated*, i.e., $\mathcal{S}_k \subseteq \mathcal{S}$ with $k = \{1, 2, \dots, K\}$. Multiple activated areas do not need to be electrically connected, and different subsets may overlap, where the surface transmits and receives to/from more than one physical location at a given time.

This work is partially supported by ELLIIT – an Excellence Center at Linköping-Lund University in Information Technology, and Ericsson AB.

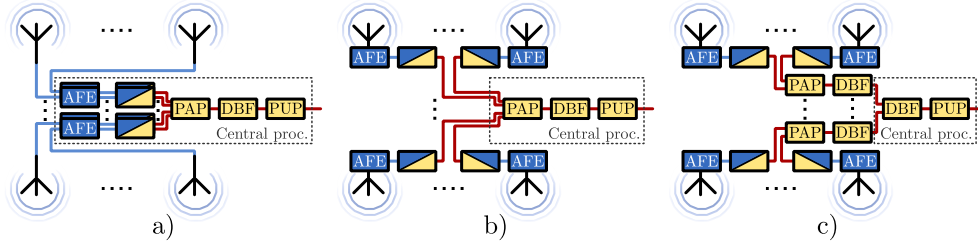


Fig. 2. Three possible types of implementation architectures, where the *analog* links are shown in blue and the *digital* links are shown in red color. (a) Analog connections, (b) digital connections, and (c) digital connections with distributed beamforming functionality. The acronyms AFE, PAP, DBF and PUP denote analog front-end, (digital) per-antenna processing, (digital) beamforming, and (digital) per-user processing. The dual colored boxes (blue-yellow) denote AD and DA connections, respectively.

This helps to *reduce* the per-link backplane data throughput, as well as to reduce the processing energy consumption - things which are extremely important when the surface aperture grows with a fixed number of served terminals.

The literature has proposed and analyzed the use of more *passive intelligent reflecting surfaces (IRSs)* [19–26]. IRS is also often referred to as *reconfigurable intelligent surfaces*, *software controllable metasurfaces*, and *reconfigurable reflect arrays*, respectively. The idea is to deploy an array of phase shifters (in hardware or software) across the IRS, which is located between a conventional BS and terminals for passive beamforming. At first glance, the passive operation seems appealing, since one may be able to manipulate the links between the BS and terminals without the use of RF up/down-conversion. However, *synchronously* operating a large number of phase shifters with a low *insertion loss* is a non-trivial task [27]. Standard varactor-based or capacitive phase shifters are known to be rather lossy, and this loss scales at least linearly with the number of phase shifters [28]. More sophisticated phase shifters with adaptive tuning capability require *active* control circuitry to inject a biasing voltage *per-phase shifter*, greatly increasing the overall design complexity and energy consumption [29]. Even more critically, in order to carry out passive beamforming from the IRS, one would have to compute the necessary beamforming weights based on some knowledge of channel state information. However, without active RF circuitry, there is no mechanism to sense the RF signals at the IRS. This implies that channel estimation can not be done passively. Furthermore, the fundamental performance benefits of IRS relative to conventional relays and massive MIMO systems are questionable (see e.g., [26, 30]). It has recently been proven that an IRS-based system *can not* outperform a conventional massive MIMO system in terms of higher achievable signal-to-noise ratio (SNR) at a given terminal [30]. *Due to these uncertainties in the performance aspects, we confine the discussions of the paper to the active LISs.*

Despite the rapid progress in assessing the performance limits of LISs, a study which discusses its implementation aspects and challenges is missing. This paper aims to close this gap. We first present a detailed discussion on the different contending implementation architectures for LISs, and evaluate their relative trade-offs. In terms of scalability, a distributed architecture across a number of common modules, where each module is connected to a smaller number of antennas may be the most suitable. We show that the decentralized nature of implementation can greatly simplify the analog signal distribution, and the physical surface layout. Multiple common modules can be connected with digital links, which are used to show that the resulting backplane throughput grows slowly as a function of the terminals, and *not* the number of antennas on the LIS.

2. REAL-TIME IMPLEMENTATION ARCHITECTURES

In steep contrast to the massive MIMO, LISs require an unconventionally large physical area/size. For canonical frequencies up to

the C-band (around 3.5 GHz), a 1024 element uni-polarized, half-wavelength spaced LIS with a 128×8 panel array can easily occupy several meters of physical area in the horizontal dimension. This presents the first implementation challenge in *efficiently routing* RF signals between the *data source/sink* and the *antenna elements* on the surface. This further raises *two* closely related implementation questions: (1) How should the antenna elements across the LIS be connected to the central processing unit? (2) Is the required backplane bandwidth achievable in practice to facilitate the necessary signalling? Fundamentally, any transceiver architecture requires partitioning its functionality into blocks which need to be *electrically* close to the radiating elements, and blocks which need to be near the centralized processor. Since the functionality of an RF transceiver can be split into the analog front-end, analog-to-digital (AD)/digital-to-analog (DA) converters, and baseband processing, there are *three* qualitatively different architectures from distributing different portions of the transceiver chain. Following the same order as the arguments presented in [31], these can be summarized as:

- Analog connections.** Here the AD, DA converters and all baseband hardware is at the central processor, while analog signals are routed from-and-to the radiating elements.
- Digital connections.** Each antenna element has a complete analog front-end, AD, DA converters and directly communicates *digitized* samples of the received analog waveform(s) with the central processing unit.
- Digital connections with distributed beamforming.** Here any *per-antenna* baseband processing can be performed *locally* at the element. Consequently, beamforming can be *distributed* into computations performed across the LIS. With such an architecture, the signals exchanged with the central processor correspond to samples of the *terminal* data, rather than each element's waveform.

The aforementioned architectures are presented in Fig. 2, where a varying degree of centralized processing is observed as we transition from analog to the digital distributed processing. The analog architecture poses a couple of critical challenges that limit its use only to a small number of elements on the surface. Firstly, the routing of analog signals induces a loss that *exponentially increases* with distance. Analog routing scales rather poorly with growing array size (and hence surface aperture), since a large number of elements require a greater number of transmission lines covering longer electrical distances. Rather interestingly, this issue is not alleviated by using higher quality transmission lines, due to their unusual form factor and cost (see [32]). The overall impact of the above is the reduction of received SNR at a terminal and at the surface. Secondly, analog architectures are highly sensitive to *distortion* and *external interference*, which limits the *beamforming* performance to/from the surface. To this end, in order to minimize the length of analog tracks, each antenna element needs to be in close proximity to its associated

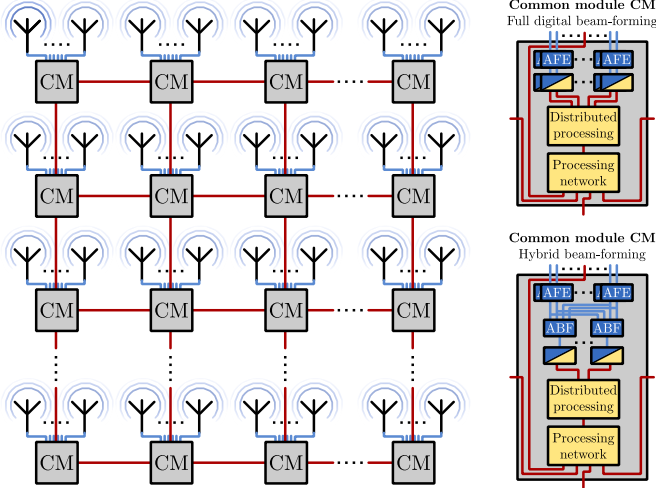


Fig. 3. An example LIS architecture with a network of common modules. Analog links are depicted in *blue* while the digital links are highlighted in *red*. The acronyms are equivalent to Fig. 2.

data converters. Naturally, doing this lends itself to an architecture which composes of a grid of identical modules, each comprising of a single RF transceiver, as well as AD/DA converters. An alternative could be distribution of coherent RF signals over fiber [33]. Despite this, the signal distribution constitutes as a major challenge. Instead, *multiple* transceivers can be fused into a *common module* equipped with *several* RF chains, AD/DA converters and a *single* set of support hardware. This concept has been initially explored for massive MIMO systems [31], and here we extend it to the case of LIS. Rather interestingly, this implies that the optimal architecture contains a mixture of analog and digital processing, striking the right balance between sharing support hardware and analog routing loss. As such, the overall architecture could be formed with a *multiplicity* of common modules tiled together with high speed digital interconnects to the nearest neighbors. Additionally, each common module is equipped with the digital hardware to perform *distributed beamforming*, which substantially improves the scalability of the surface backplane network. Figure 3 depicts the block diagram of a LIS architecture consisting of common modules, as well as the two functional options for each common module in fully digital beamforming and hybrid (analog-digital) beamforming.

In view of the above, the data throughputs generated or consumed by *one antenna element* and *one common module* on the LIS can be computed via the methodology in [31]. To do this, we assume that the LIS contains an array of M antennas serving K terminals. The LIS can be divided into a discrete grid of N modules, with N being an integer divisor of M , such that each module drives M/N radiating elements within a certain surface area. A system bandwidth B is assumed such that the sampling rate of the AD/DA converters must be $\geq 2B$. If each AD/DA converter has a resolution of N_{res} bits, and is over-sampled by $N_{\text{over-samp}}$, then the data throughput generated/consumed by a given element on the surface is $R = 2BN_{\text{res}}N_{\text{over-samp}}$, and the throughput of one module is $(M/N)R$. The trade-off between the number of antennas and connected modules is an important one. This is since with a fixed R and N , increasing M yields an increase in the data throughput per module. To this end, an increased surface aperture with a higher number of antennas should be complimented with a relatively higher number of modules for efficient backplane processing.

3. BACKPLANE INTERCONNECT TOPOLOGIES AND IMPACT OF DISTRIBUTED BEAMFORMING

In a digitally interconnected architecture, the required backplane throughput serves the biggest hurdle limiting an increase in the number of antenna elements. As seen from the Lund University massive MIMO testbed, a total backplane throughput of 384 Gb/s was required to facilitate 100 elements over 20 MHz channel bandwidth [9]. From this one can imagine the required throughput for a LIS containing hundreds-thousands of elements. We foresee *two* types of interconnect topologies: First is when each common module can have a *dedicated physical link* to the central control unit. Second is when all modules can be *daisy-chained on a single link* over which all antenna signals are transmitted. These configurations seem to lean towards low per-link data throughput or low physical resource sharing. Naturally, other topologies can be designed by combining the two approaches, such as having multiple parallel chains. Moreover, a *mesh network* can be considered as an extension of the daisy-chain concept, where each node can communicate with its nearest neighbors. The fully parallel backplane, as demonstrated in Fig. 4 a) requires the lowest per-link data throughput. However, it has limitations in terms of scalability. To serve all elements of the LIS, the interconnect length must grow with the size of the surface/array. This requires high performing, costly and high energy consuming links to support reliable transmission. In addition, routing complexity and the cross-interference between links also increases with the number of modules. The aforementioned challenges can be addressed by implementing the surface backplane as a nearest-neighbor mesh network, which require interconnects only at the scale of the *inter-module distance*, regardless of the number of elements deployed on the surface. Since the connections are only local, the challenge of globally routing N links while maintaining acceptable cross-interference levels is avoided. The mesh network also presents a level of redundancy that allows for reconfiguration of the routing path to circumvent failures.

Nonetheless, an interconnect topology consisting of a mesh network *by itself* does not help to reduce the total required bandwidth at the surface control/processing level, as shown in Fig. 4 b). When performing centralized beamforming, irrespective of the backplane topology, there is a fundamental requirement to exchange M waveforms with the central processor, for a maximum data throughput of MR . Moreover, there is an additional penalty in the aggregate throughput due to the multi-hop nature of links. Assuming that the LIS contains N modules which are connected to the central processor with N_{ch} parallel daisy-chains, where N_{ch} is an integer divisor of N . At any point along the chain, the data throughput is proportional to the number of *preceding* elements. Thus, the *aggregate* data throughput through the surface is given by [31]

$$R_{\text{tot}} = N_{\text{ch}} \sum_{\ell=1}^{\Delta} \ell \left(\frac{MR}{N} \right) = \frac{M}{2} (\Delta + 1) R, \quad (1)$$

where $\Delta = N/N_{\text{ch}}$. In addition, the total power consumed by the backplane network *increases* as the product of the number of antennas and the number of modules, corresponding to the penalty incurred by sending data through multiple hops. Similar effects take place in the fully parallel backplane since some links must communicate over a large distance, requiring increased energy consumption. To demonstrate the limitation of centralized routing, we consider a LIS with $M = 1024$ elements with $B = 20$ MHz, no oversampling, i.e., $N_{\text{over-samp}} = 1$, and $N_{\text{res}} = 10$. Under such modest conditions, neglecting additional system-related overheads, the throughput coming in-and-out of the central unit is 381.4 Gb/s. As stated in [34], a

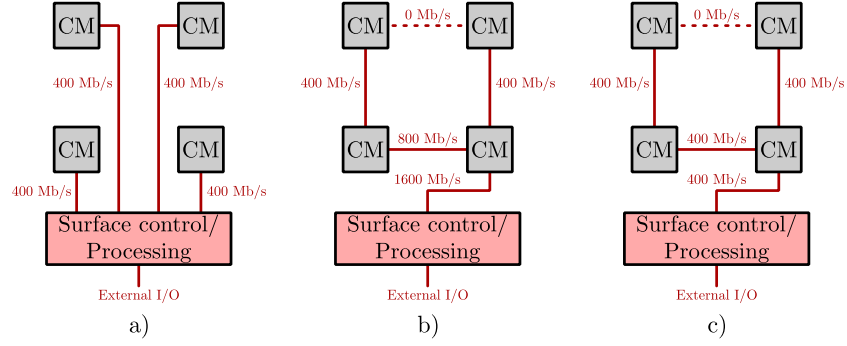


Fig. 4. Possible routing schemes inspired by the arguments in [31]. Throughputs are shown for an example where each block generates 400Mb/s. a) Fully parallel backplane interconnect. b) Mesh network with *centralized* beamforming. c) Mesh network with *distributed beamforming*: each block computes its estimate of the terminals’ signals and these are summed throughout the routing network to generate the overall beamformed samples.

link operating at this rate over a few centimeters could achieve an energy efficiency of 1 pJ/bit. To this end, just transporting bits in and out of the central processor will consume hundreds of milliwatts of power. Even exploiting the greatest possible link parallelism at the central processor, it would be difficult to achieve LIS size-bandwidth product, especially for scalable architectures.

The above problem can be resolved by distributing the processing at each module. This was originally suggested by [35], and was evaluated for distributed massive MIMO systems by [31]. Below we extend these discussions for LIS, where we qualitatively compare the routing throughput, as well as discuss the impact of the common modules. The key observation is that the M waveforms at the radiating elements of the LIS are not linearly independent but instead lie in a K -dimensional subspace generated by the K distinctly located terminals [31]. By exploiting this redundancy, it is possible to exchange only K rather than M unique signals with the central processor, performing distributed beamforming. Just as in massive MIMO systems, we foresee the regime of $M \gg K$, from which the required backplane throughput is substantially reduced. Since linear digital beamforming is performing matrix manipulations, such computations can be distributed. In the uplink, each element would multiply its received signal by a weight vector containing *one entry per-terminal*. These vectors are then summed across the array to generate the per-terminal spatially filtered signals. This task can be embedded in the digital link to be low-latency and low-energy [31]. The process is reversed in the downlink, where the terminal data streams are sent to all the modules and each element combines them with the appropriate beamforming weights to generate DA samples.

In an orthogonal frequency-division multiplexing-based system where beamforming is performed independently on each subcarrier, each common module would require a timing and frequency recovery block, downsampling and upsampling filters, a fast Fourier transform unit, as well as a small number of complex multipliers and adders. At the beginning of each frame, a given terminal transmits a synchronization and training preamble which can be used to estimate the timing parameters for each terminal at each module. This training preamble can also include channel estimation for computation of beamforming weights. Subsequently, during data transmission, distributed linear beamforming is performed independently for each subcarrier. Note that in both uplink and downlink, the processing remains exactly the same. The only additional hardware required is a *configurable-depth buffer* on each common module to match the latency of the backplane network [31]. By routing the respective terminal signals around the array rather than the antennas, *the maximum required throughput is proportional to the number of terminals rather than the number of antennas on the LIS*. This is

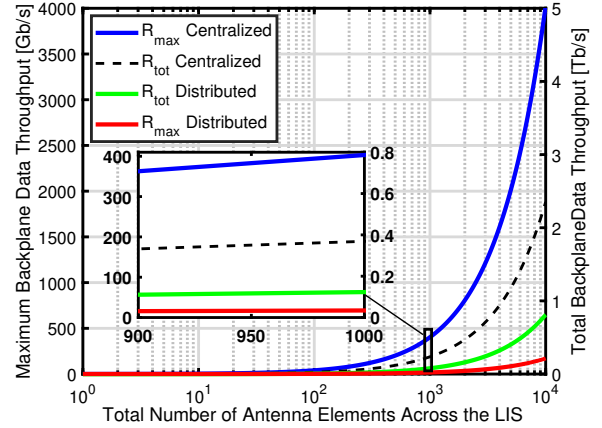


Fig. 5. Maximum and aggregate backplane throughput with centralized and distributed beamforming for constant $M/K = 32$.

seen in Fig. 4 c). We note that the above discussions are inspired by [31]. We assume that each terminal’s modulated data stream is represented by in-phase and quadrature samples of $N_{\text{bit,beamf}}$ bits each, then the maximum throughput at the central processor is given by $R_{\text{max}} = 2KB N_{\text{bit,beamf}}$, and the aggregate throughput to deliver all K signals to all N modules is given by $R_{\text{tot}} = 2NKB N_{\text{bit,beamf}}$ [31]. The substantial improvement over fully centralized processing is evident since the throughput of the central processor’s link is now proportional to the number of terminals. Figure 5 illustrates these benefits for an example with $B = 20$ MHz, $N_{\text{res}} = 10$, $N_{\text{bit,beamf}} = 15$ and $N_{\text{ch}} = 10$. With a constant ratio of $M/K = 32$, both the maximum and aggregate throughputs are seen to reduce tremendously with distributed computations, highlighting its importance in such systems.

4. CONCLUSION

The paper presents a first look at the implementation aspects of active LIS. An investigation into the relative merits and disadvantages of different architectures was carried out. To address the challenges associated with the large surface aperture and a large number of antennas, a distributed approach with a grid of common modules seems to serve as the most scalable solution. Each module is intended to drive a small number of elements and consists of a RF up/down-converter, data converters and a backplane connection to a central processor. The modules were digitally linked to the nearest neighbor via a mesh network, providing connectivity to the central processor. The maximum backplane communication throughput of a LIS was shown to slowly increase with the number of served terminals, providing an enormous advantage against centralized architectures.

5. REFERENCES

- [1] E. Dahlman, *et al.*, *5G NR: Next Generation Wireless Access Technology*, Elsevier Press, 2018.
- [2] E. Björnson, *et al.*, “Massive MIMO is a reality - what is next?: Five promising research directions for antenna arrays,” *Digital Signal Process.*, 2019, to appear.
- [3] M. Shafi, *et al.*, “5G: A tutorial overview of standards, trials, challenges, deployment, and practice,” *IEEE J. Sel. Areas Commun.*, vol. 35, no. 6, pp. 1201-1221, Jun. 2017.
- [4] M. Shafi, *et al.*, “Microwave vs. millimeter-wave propagation channels: Key differences and impact on 5G cellular systems,” *IEEE Commun. Mag.*, vol. 56, no. 12, pp. 14-20, Dec. 2018.
- [5] T. L. Marzetta, “Noncooperative cellular wireless with unlimited numbers of base station antennas,” *IEEE Trans. Wireless Commun.*, vol. 9, no. 11, pp. 1436-1449, Nov. 2010.
- [6] E. G. Larsson, *et al.*, “Massive MIMO for next generation wireless systems,” *IEEE Commun. Mag.*, vol. 52, no. 2, pp. 186-195, Feb. 2014.
- [7] X. Gao, *et al.*, “Massive MIMO performance evaluation based on measured propagation data,” *IEEE Trans. Wireless Commun.*, vol. 14, no. 7, pp. 3899-3911, Jun. 2015.
- [8] H. Q. Ngo, *et al.*, “Energy and spectral efficiency of very large multiuser MIMO systems,” *IEEE Trans. Commun.*, vol. 61, no. 4, pp. 1436-1449, Apr. 2013.
- [9] J. Viera, *et al.*, “A flexible 100-antenna testbed for massive MIMO,” in *Proc. IEEE Global Conf. Commun. (GLOBECOM) Workshops*, Dec. 2014, pp. 287-293.
- [10] E. Björnson, *et al.*, “Massive MIMO in sub-6 GHz and mmWave: Physical, practical, and use-case differences,” *IEEE Wireless Commun.*, vol. 26, no. 2, pp. 100-108, Apr. 2019.
- [11] S. Hu, *et al.*, “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.
- [12] C. Liaskos, *et al.*, “A new wireless communication paradigm through software-controlled metasurfaces,” *IEEE Commun. Mag.*, vol. 56, no. 9, pp. 162-169, Sep. 2018.
- [13] S. Shi, *et al.*, “Ultrawideband optically fed tightly coupled phased array,” *J. Lightwave Technol.*, vol. 33, no. 23, pp. 4781-4790, Dec. 2015.
- [14] P. Lu, *et al.*, “RF TX Design Considerations for eWallpaper,” May 2014. [Online]. Accessible from <https://bwrc.eecs.berkeley.edu/projects/5605/ewallpaper>.
- [15] S. Hu, *et al.*, “Capacity degradation with modeling hardware impairment in large intelligent surface,” in *Proc. IEEE Global Conf. Commun. (GLOBECOM)*, Dec. 2018.
- [16] L. Subrt, *et al.*, “Controlling propagation environments using intelligent walls,” in *Proc. European Conf. Antennas and Propag. (EuCAP)*, Mar. 2012.
- [17] C. Shepard, *et al.*, “Practical performance of MU-MIMO precoding in many-antenna base stations,” in *Proc. ACM Workshop on Cellular Networks: Operations, Challenges and Future Design (CELLNET)*, Jun. 2013.
- [18] S. Hu, *et al.*, “User assignment with distributed Large Intelligent Surface (LIS) systems,” in *Proc. IEEE Personal Indoor Mobile Radio Commun. (PIMRC)*, Sep. 2018.
- [19] Q. Wu and R. Zhang, “Intelligent reflecting surface enhanced wireless network via joint active and passive beamforming,” *IEEE Trans. Wireless Commun.*, 2019, to appear.
- [20] Q. Wu and R. Zhang, “Beamforming optimization for intelligent reflecting surface with discrete phase shifts,” in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Process. (ICASSP)*, May 2019, pp. 7830-7833.
- [21] Q. Wu and R. Zhang, “Intelligent reflecting surface enhanced wireless network: Joint active and passive beamforming design,” in *Proc. IEEE Global Conf. Commun. (GLOBECOM)*, Dec. 2018.
- [22] C. Huang, *et al.*, “Reconfigurable intelligent surfaces for energy efficiency in wireless communication,” *IEEE Trans. Wireless Commun.*, vol. 18, no. 8, pp. 4157-4170, Aug. 2019.
- [23] C. Huang, *et al.*, “Energy efficient multi-user MISO communication using low resolution large intelligent surfaces,” in *Proc. IEEE Global Conf. Commun. (GLOBECOM)*, Dec. 2018.
- [24] M. Di Renzo, *et al.*, “Smart radio environments empowered by reconfigurable AI meta-surfaces: An idea whose time has come,” *EURASIP J. Wireless Commun. and Networking*, vol. 2019, no. 1, pp. 1-20, May 2019.
- [25] D. Mishra and H. Johansson, “Channel estimation and low-complexity beamforming design for passive intelligent surface assisted MISO wireless energy transfer,” *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Process. (ICASSP)*, May 2019, pp. 4659-4663.
- [26] E. Björnson, *et al.*, “Intelligent reflecting surface vs. decode-and-forward: How large surfaces are needed to beat relaying?,” *CoRR*, vol. abs/1906.03949, 2019.
- [27] D. M. Pozar, *Microwave Engineering*, Wiley Press, 1990.
- [28] D. Uttamchandani, *Handbook of MEMS for Wireless and Mobile Applications*, Elsevier Press, 2013.
- [29] B. An, *et al.*, “Wideband tunable phase shifter with low in-band phase deviation using coupled line,” *IEEE Microwave and Wireless Components Lett.*, vol. 28, no. 8, pp. 678-680, Aug. 2008.
- [30] E. Björnson and L. Sanguinetti, “Demystifying the power scaling law of intelligent reflecting surfaces and metasurfaces,” in *Proc. IEEE Int. Workshop on Computational Adv. Multisensor Adaptive Process. (CAMSAP)*, Dec. 2019, to appear.
- [31] A. Puglielli, *et al.*, “A scalable massive MIMO array architecture based on common modules,” in *Proc. IEEE Int. Conf. Commun. (ICC) Workshops*, May 2015, pp. 1310-1315.
- [32] A. Rocha, *et al.*, “A demonstrator for impedance matching systems in transmission lines with time and frequency simulation,” *IEEE Antennas and Propag. Mag.*, vol. 61, no. 3, pp. 92-103, Jun. 2019.
- [33] W. Ahmad, *et al.*, “CMOS integrated remote antenna unit for fiber-fed distributed MIMO systems,” *IEEE Trans. Microwave Theory and Techn.*, vol. 65, no. 1, pp. 173-186, Jan. 2017.
- [34] B. Ragahavan, *et al.*, “A sub-2W 39.8446 Gb/s transmitter and receiver chipset with SFI-5.2 interface in 40 nm CMOS,” *IEEE J. Solid-State Circuits*, vol. 48, no. 12, pp. 3219-3228, Dec. 2013.
- [35] C. Shepard, *et al.*, “Argos: Practical many-antenna base stations,” in *Proc. ACM Int. Conf. Mobile Computing and Networking (MOBICOM)*, Aug. 2012, pp. 53-64.