Reconfigurable Intelligent Surfaces:
Three Myths and Two Critical Questions

Emil Björnson, Özgecan Özdogan, Erik G. Larsson

Abstract—The search for physical-layer technologies that can play a key role in beyond-5G systems has started. One option is reconfigurable intelligent surfaces (RIS), which can collect wireless signals from a transmitter and passively beamform them towards the receiver. The technology has exciting prospects and is quickly gaining traction in the communication community, but in the current hype we have witnessed how several myths and overstatements are spreading in the literature.

In this article, we take a neutral look at the RIS technology. We first review the fundamentals and then explain specific features that can be easily misinterpreted. In particular, we debunk three myths: 1) Current network technology can only control the transmitter and receiver, not the environment in between; 2) A better asymptotic array gain is achieved than with conventional beamforming; 3) The pathloss is the same as with anomalous mirrors.

To inspire further research, we conclude by identifying two critical questions that must be answered for RIS to become a successful technology: 1) What is a convincing use case for RIS?; 2) How can we estimate channels and control an RIS in real time?

INTRODUCTION

The electromagnetic waves that carry information in wireless communications interact with objects and surfaces on their way from the transmitter to the receiver. Although the superposition of many propagation paths gives rise to random-like fading phenomena, every propagation path has a constant behavior. However, there exist engineered materials whose interactions with electromagnetic waves are not constant but reconfigurable. These materials are not naturally occurring but can be manufactured and deployed to shape the propagation environment. The prospects of including such *reconfigurable intelligent surfaces (RIS)* as a part of beyond-5G network architectures are attracting much attention [1], [2]. RIS have also been called *software-controlled metasurfaces* [3] and *intelligent reflecting surfaces* [4].

A basic use case of RIS is illustrated in Fig. 1, where a rooftop-mounted base station (BS) is transmitting to an indoor user. Suppose the building materials are such that the direct path through the wall experiences massive penetration losses, while the path through the window only experiences minor losses. Inside the window, an RIS is deployed to capture signal energy proportional to its area and re-radiate it in the shape of a beam towards the receiver. To ensure the beam is focused

towards the user device, wherever it is in the room, the RIS must be reconfigurable. By using an RIS in this setup, the signal-to-noise ratio (SNR) can be improved.

An RIS is a thin surface composed of N elements, each being a reconfigurable scatterer: a small antenna that receives and re-radiates without amplification, but with a configurable time-delay [3]. For narrowband signals, this delay corresponds to a phase-shift. Assuming the phase-shifts are properly adjusted, the N scattered waves will add constructively at the receiver. This principle resembles traditional beamforming: each element has a fixed radiation pattern but the collection of phase-shifts determines where constructive interference among the scattered waves occurs. The color pattern at the RIS in Fig. 1 represents the phase-shifts necessary to steer a beam towards the receiver. Each element is substantially smaller than the wavelength (e.g., a fifth of the wavelength in each direction [5]) so it scatters signals almost uniformly, giving the surface the ability to form equally strong beams in any direction [6].

The propagation analysis of an RIS essentially entails (i) finding the Green's function of the signal source (a sum of spherical waves if close, or a plane wave if far away), (ii) computing the impinging field at each RIS element, (iii) integrating this field over the surface of each element to find the current density, (iv) computing the radiated field from each element, and (v) applying the superposition principle to find the field at the receiver. Since the elements are small, one can approximate the re-radiated field by pretending each element is a point source and then the received signal is a superposition of phase-shifted, amplitude-scaled source signals [6].

There are many prospective use cases for RIS-aided wireless communications, in addition to improving the SNR as in Fig. 1. The RIS can also mitigate interference between users that are spatially multiplexed or limit the signal-leakage outside the intended coverage area, to mitigate eavesdropping [2]–[4]. Support for wireless power transfer, backscattering, and spatial modulation is also conceivable; most things that can be implemented using antenna arrays can also be carried out by an RIS [5].

The definition of an RIS is a surface with real-time reconfigurable scattering properties (e.g., amplitude, delay, and polarization) that is controlled to improve the communication performance. The concept is often connected with metasurfaces, which are two-dimensional surfaces consisting of arrays of reconfigurable elements of metamaterial [5]. However, there are other potential ways of implementing RIS [2]. One example is using small patch antennas terminated with an adjustable impedance. In any case, the reconfigurability will likely be limited to a finite set of states per element (with

Fig. 1. A typical use case of an RIS, where it receives a signal from the transmitter and re-radiates it focused towards the receiver. To focus the beam in the right direction, the RIS must be configured properly.

given delays and amplitudes) and mutual coupling between adjacent elements is another limitation. We refer to [5] for a recent survey on the implementation aspects for metasurfaces. There are decades of research on reflectarrays and array lenses [7], which are architectures for building transmitters consisting of a feed antenna that sends the signal via a reconfigurable surface capable of electronically tunable beamforming. The key difference is that an RIS is not co-located with the transmitter or receiver, but deployed in between to aid the communication.

BASIC FEATURES AND RELATED MYTHS

We will now describe three fundamental features that the RIS technology possess. Along the way, we will also debunk three myths that are flourishing in the literature.

Feature 1: Creating Controllable Radio Environments

A key feature of RIS is the ability to alter how wireless signals propagate between the transmitter and receiver. It is a technology for creating controllable/smart/programmable radio environments, which are defined as environments that can customize how signals propagate from the transmitter to the receiver [2]. This feature enables joint optimization of the transmitter/receiver and the controllable entities in the environment, using channel state information (CSI). When motivating the novelty of this feature, the following claim has been made repeatedly [3], [4], [8].

Myth 1: Current network technology can only control the transmitter and receiver, not the environment in between.

Many wireless systems indeed consist of a transmitter that communicates with a receiver without the involvement of other entities. The radio environment is then uncontrollable; the transmitter and receiver must conform to it by adaptive modulation/coding, beamforming, and power control. However, this is a choice made by the network provider because the technology for controlling the signal propagation between the end points has existed all along. The wireless repeater was invented in 1899 and advanced relaying technology, capable of adaptively improving the channel between the transmitter and receiver, has been supported by cellular standards since 3G [9]. Hence, the statement above is a myth.

We will now put the RIS technology into a historical context. The term *cooperative communications* is broadly used to refer to network architectures containing entities between the transmitter and receiver that enhance the physical channel, by exploiting diversity, beamforming, and/or multiplexing gains [9]. These entities are co-optimized with the transmitter and receiver, thus satisfying the definition of controllable radio environments. Two main categories are *transparent relaying* and *regenerative relaying*. In the former category, each relay is an entity that receives a signal from the transmitter and processes it in analog (or digitally) before re-radiating it towards the receiver. Amplify-and-forward is a classic protocol for creating additional signal paths by re-radiating an amplified signal in a way that can be transparent to the receiver. No
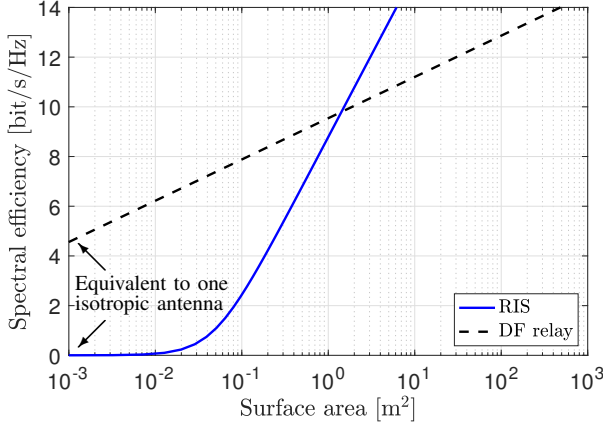
Fig. 2. This figure revisits the setup in Fig. 1 and compares the use of an RIS with the use of a multi-antenna DF relay deployed at the same place. The direct path is assumed non-existing, while there are line-of-sight paths via the RIS/relay. The transmitter is 300 m from the RIS/relay, while the user is 10 m from it. The transmit power is equivalent to 10 W per 20 MHz at the BS and 0.1 W at the relay. The antenna gain is 10 dBi at the BS and 0 dBi at the relay and user. The glass penetration loss is −20 dB and the noise figure is 10 dB. The figure shows the SE achieved by an RIS and a DF relay for different surface areas. For practical SEs (below 8 bit/s/Hz), the DF relay is smaller, while the advantage of the RIS is the lack of power amplifiers and full-duplex mode.

baseband processing is required, only amplification. In regenerative relaying, each relay is decoding the received signal and processes it in the digital baseband, before retransmitting it in an optimized manner towards the receiver. Decode-and-forward (DF) is a common example. Classical relays operate in half-duplex, where reception and retransmission are separated in time, but regenerative full-duplex relays capable of receiving and transmitting simultaneously are emerging [10].

The RIS technology is unique in that it fills an empty slot in the relaying taxonomy: it is a transparent relay with a full-duplex protocol [1], thus it affects the propagation in real-time. The potential advantage over traditional relays is that large surfaces can be implemented with reduced energy consumption and cost since the use of printed metamaterial requires no amplifiers, but only power dissipation in the hardware controlling the reconfigurability. The drawback is the reduced signal range due to the lack of amplification.

Fig. 2 illustrates this in a setup that resembles the one illustrated in Fig. 1, but where the RIS is possibly replaced by a multi-antenna half-duplex repetition-coded DF relay (a simple but suboptimal relaying scheme). Perfect CSI is assumed and each RIS element scatters all the incoming signal energy with a perfectly controlled phase. The figure shows the array's surface area required to achieve a particular spectral efficiency (SE) when using either an RIS or a DF relay. The results are frequency-independent but the number of elements that fits into the surface area grows quadratically with the wavelength. We observe that the DF relay can have a much smaller form factor than the RIS, except if very high SE is required. The reason is that the DF relay achieves a much higher SNR but it also needs a higher SNR to achieve the same SE since it operates in half-duplex, whereas the RIS operates in full-duplex.

In summary, the RIS technology can control/optimize the propagation environment between the transmitter and receiver, just as previous relaying technologies. The unique feature of RIS is that it reduces the hardware complexity at the price of requiring a larger surface.

Feature 2: Passive Beamforming

Beamforming appears when delayed copies of the same signal are emitted from multiple antennas. This gives rise to constructive interference at spatial locations where the copies are received synchronously and destructive interference elsewhere. If the time-delays at N transmit antennas are tuned to achieve constructive interference at the receiver, it will receive N times more power than if the same total power was transmitted from a single antenna. This is the conventional *array gain* that shows how the beamformed signal becomes more spatially focused as the array size grows.

An RIS is capable of passive beamforming. It receives signal power from the transmitter proportional to its surface area, which is proportional to the number of elements, N. When the RIS re-radiates the signal, with time-delays selected to beamform at the receiver, an array gain of N is obtained just as with conventional beamforming. The combination of these two effects, both being proportional to N, leads to an SNR at the receiver proportional to N^2. This is called the "square law" [4].

Suppose we compare the setup in Fig. 1 with the case when the RIS is replaced with a transmitter array having the same size. The RIS setup will achieve an SNR that grows as N^2, while the SNR in the latter setup only grows as N. It has been claimed that these are asymptotic scaling laws [4], which means that the SNR with the RIS grows unboundedly with N at the order of N^2 and eventually becomes larger than with the transmitter array. This is incorrect.

Myth 2: A better asymptotic array gain is achieved than with conventional beamforming.

The first issue with this statement is that array gains of the type described above only appear when the surface area (of the RIS or transmitter array) is small compared to the propagation distances. The transmitter/receiver must be in the geometric far-field of the surface so that the pathloss is approximately the same to all parts of the surface. Since the surface area grows with N, the far-field approximation eventually breaks down as N increases. Neither linear nor quadratic asymptotic power scaling laws can exist since the law of conservation of energy dictates that we cannot receive more energy than was transmitted. Nevertheless, the SNR achieved with an RIS actually grows quadratically with the number of elements for many practically-sized surfaces [11]. Hence, it might seem possible that a better SNR can be achieved in the RIS setup when considering large, equal-sized arrays. The second issue with the statement is the premise that the quadratic power scaling is advantageous. The pathloss from the transmitter to each RIS element is huge in the far-field, thus it is more accurate to say that the power loss between the transmitter and RIS reduces as $1/N$ [11]. The SNR achieved with an RIS cannot surpass the SNR achieved when replacing it by an

Fig. 3. This figure continues the example from Fig. 2 by comparing the end-to-end SNR achieved by the RIS and the DF relay. The RIS benefits from the "square law" by achieving a steeper slope for practical surface areas (below 100 m^2). Nevertheless, the DF relay is consistently providing a better SNR and both curves converge to a finite number since there are no asymptotic scaling laws.

equal-sized antenna array transmitting with the same power as in the RIS case, but the difference reduces as $1/N$.

The reason that the DF relay outperformed the RIS for most SE values in Fig. 2 is that the RIS suffers from the power loss inherent in the "square law". To demonstrate this, Fig. 3 revisits the example by showing the end-to-end SNR achieved with the RIS and DF relay for different surface areas. Since we use logarithmic scales, the quadratic array gain is observed from the steeper slope of the RIS curve. However, this curve begins at a much smaller value and when it approaches the DF relay curve, the steeper slope has tapered off. Both curves will eventually converge to a finite number [11]. The reason that the RIS became preferable for very high SEs in Fig. 2 is that the SNR gap eventually becomes so small that the half-duplex operation of the DF relay becomes the bottleneck.

In summary, an RIS is capable of passively beamforming a signal towards the receiver. Due to the faster-than-linear SNR-scaling, physically large surfaces are highly preferable.

Feature 3: Synthesizing a Different Surface Shape

The RIS can not only form a beam, it can synthesize the scattering behavior of an arbitrarily-shaped surface of the same size. For example, it can create a superposition of multiple beams or act as a diffuse scatterer [5].

A common example is to synthesize an *anomalous mirror/reflector*. A mirror is a surface that reflects an impinging plane wave as an outgoing plane wave, also known as specular reflection. A conventional mirror satisfies the law of reflection: the angles of the impinging and reflected waves to the surface normal are the same but on opposite sides, as illustrated by the blue ribbons in Fig. 4. An anomalous mirror reflects impinging plane waves as outgoing plane waves with a different "unnatural" angle to the surface normal [2]. A conventional mirror is an infinitely large homogeneous surface and approximations thereof appear naturally (e.g., a metal plate or water surface). In contrast, an anomalous mirror is not naturally appearing

but can be synthesized by an engineered inhomogeneous surface. A property of mirrors is that the receiver observes the transmitting source as being behind the mirror. One can analyze the wave propagation as if the transmitter is moved to the location of the mirror image, as illustrated in Fig. 4.

It has been stated that an RIS can generally be viewed as an anomalous mirror if it has a width and length larger than ten wavelengths [8]. If that is the case, the pathloss in Fig. 4 can be computed based on the sum of the distance from the transmitter to the RIS and from the RIS to the receiver [8], which is the distance from the mirror image to the receiver. These are myths that are summarized as follows.

Myth 3: The pathloss is the same as with anomalous mirrors.

An ideal mirror reflects a signal with zero beamwidth. If a plane wave is impinging on a finite-sized RIS that is configured to focus the signal towards a receiver located in the far-field, then the radiated field will be strongest in the angular direction of the receiver but it will not be a plane wave. Far-field focusing is called beamforming and the beamwidth is the same as for beamforming from an equal-sized transmitter array. Hence, the half-power beamwidth of the reflected signal is inversely proportional to the size of the RIS (measured in wavelengths) and becomes 6° for a surface that is ten wavelengths in each dimension [6].

Mirrors and plane waves are theoretical idealizations that only appear approximately in practice. They can be fairly accurate approximations in visible light and are, thus, used in geometrical optics to analyze imaging. The situation is different in the radio spectrum used for communications. A surface that our eyes perceive as a mirror might be far from mirror-like for radio signals. Since the wavelength is roughly 100000 times larger in radio spectrum than in visible light (e.g., comparing green light at 600 THz with a radio signal at 6 GHz), a surface must be 100000 times larger in each dimension to identically reflect signals. The transmitter must be 100000 times further away if its emitted spherical waves should be approximated as planar, and the receiver must be 100000 times further away to perceive the reflected signals as plane waves. Since mirrors only exist in asymptotic limits, there are no finite-sized surfaces that always can be approximated as mirrors. If the RIS is viewed from far enough away, its radiated field will have a beamwidth that is inversely proportional to its size.

Even if we limit the scope to setups where conventional mirrors approximately exist, the statement above remains a myth since the pathloss achieved by an equal-sized RIS is widely different. An RIS can both affect the direction and shape of the reflected signal [11], as illustrated by the red ribbons in Fig. 4 where the signal is focused at the receiver. For this reason, the SNR achieved by the RIS is proportional to N^2 and is inversely proportional to the product of the squared distances to the RIS [6], rather than inversely proportional to the squared sum of the distances as with a mirror.

To explain the fundamental differences, Fig. 5 continues the example from Figs. 2 and 3 by showing how the end-to-end pathloss depends on how far the receiver is from the RIS (the distance between the transmitter and RIS is as before).

Fig. 4. A mirror reflects an impinging plane wave as a plane wave in an angular direction determined by the law of reflection, so the receiver perceives the transmitter as being located at the mirror image location. An RIS can both configure the angle of the reflected beam and its shape, thus it should not be interpreted as an anomalous mirror. The figure illustrates how the RIS focuses the signal at the receiver to maximize the SNR.
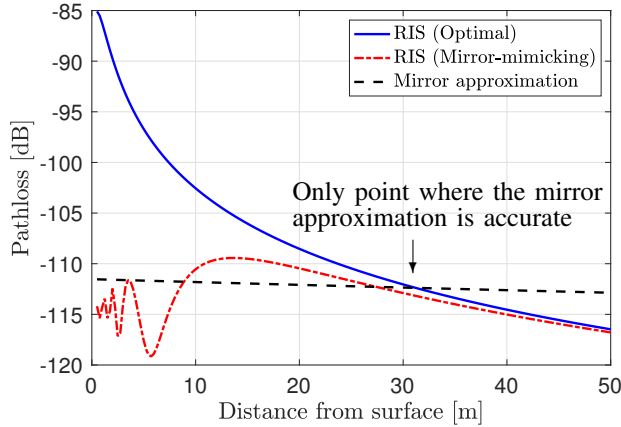

Fig. 5. This figure revisits the setup from the previous figures and considers an RIS that is 2×2 m, which represents 20×20 wavelengths at a 3 GHz frequency and 200×200 wavelengths at a 30 GHz frequency. The figure shows the end-to-end pathloss as a function of the distance between the RIS and the receiver. An optimally configured RIS is compared with an RIS that is configured to mimic a mirror and the pathloss obtained if it was an ideal mirror. It is clear that an RIS can generally not been interpreted as a mirror.

The solid curve is for an RIS that is optimized to achieve the highest SNR, while the dashed curve represents an anomalous mirror. We notice that a mirror is a poor approximation of an RIS at most distances. When the receiver is far from the RIS, the pathloss is worse than with a mirror since the RIS is too small to emit approximately plane waves. When the receiver is close to the RIS, the pathloss is instead much better than with a mirror. This is like when you look into a large mirror and your reflection only appears in a small part of it; the rest of the mirror is not needed. A well-configured RIS makes use of the entire surface by focusing the signal at the receiver as illustrated in Fig. 4. The dash-dotted curve in Fig. 5 represents a mirror-mimicking RIS that is configured to delay the signals as a cutout from an infinitely large anomalous mirror would do. This curve is close to the optimized RIS when the receiver is far from the surface, while it begins to oscillate in the vicinity of the mirror approximation at shorter distances. This indicates that the mirror analogy can be used for identifying suitable time-delays when the transmitter and receiver are far away. This is the reason why (approximate) anomalous reflection is a canonical use case in the electromagnetic literature [5], where pathloss modeling is not considered.

In summary, the pathloss achieved with an RIS does not coincide with that of an anomalous mirror. When the receiver is far from the surface, it is too small to behave like a mirror. When the receiver is near the surface, the RIS can approximate the mirror behavior but it would be suboptimal; a mirror beamforms to points infinitely far away, while a RIS can focus on the actual receiver location. One way to describe the capabilities of an RIS is as a parabolic reflector with curvature and direction that can be electronically steered, but that is also a simplification since an RIS is capable of mimicking the scattering of arbitrarily-shaped objects having the same size.

CRITICAL QUESTION 1: WHAT IS A CONVINCING USE CASE FOR RIS?

An immense amount of time and resources are required to bring a new technology concept, such as RIS, from theory to practice. Very convincing benefits compared to existing technologies must be established to motivate such an investment; we essentially need to demonstrate 10 times improvements with respect to a practically important performance metric, not just 20% gains that might disappear in an imperfect implementation. Massive MIMO (multiple-input multiple-output) and mmWave communications passed this test in the 5G development since the former can increase the number

simultaneously served users by ten times while the latter can increase the data rate per user by ten times using much wider bandwidths. Several other "5G-branded" technologies failed the test because the gains were too limited.

RIS technology has many technical features beyond current mainstream technology [2], [5]. However, to motivate the practical development of RIS technology, the critical question is: *what is a convincing use case?* The question is open; RIS is a hammer looking for a nail. There is no shortage of visions on what an RIS can be used for (some ideas were listed in the introduction) but will it excel at anything? Coverage extension is one option but Fig. 2 showed that conventional half-duplex relaying is a competitive solution, and full-duplex regenerative relays are emerging [10]. Since each RIS element must be identically configured over the entire frequency band, the RIS technology has a further competitive disadvantage over wideband channels. Improved spatial multiplexing and interference mitigation is another potential use case, but then it needs to beat Cell-free Massive MIMO, which is the emerging deployment of distributed jointly-operating antennas. Perhaps it is in terahertz bands, where the implementation of coherent transceivers is truly challenging and the sparse channels make additional propagation paths useful even if they are weak, that the RIS technology will be most beneficial. These are just speculations since there is no hard evidence yet.

CRITICAL QUESTION 2: HOW CAN WE ESTIMATE CHANNELS AND CONTROL AN RIS IN REAL TIME?

The envisioned use cases of RIS critically depend on a proper configuration of the elements based on CSI. There are two reasons why channel acquisition is particularly challenging with RIS. Firstly, unlike conventional transceiver architectures, an RIS is not inherently equipped with transceiver chains. It lacks sensing capabilities but simply "reflects" the impinging signals. Therefore, conventional channel estimation methods cannot be utilized. Secondly, introducing an RIS into an existing setup will increase the number of channel coefficients proportionally to the number of elements, N. As shown earlier, a large N is needed for RIS to be competitive, thus the estimation overhead might be huge. A key question is: *can an RIS be real-time reconfigured to manage user mobility?*

The literature contains initial approaches to tackle the problem. One approach is to transmit a pilot sequence repeatedly and measure the received signal when using different RIS configurations. For example, the elements can be turned on/off according to a pattern or the array geometry can be used to sweep through changes of the main reflection angle. At least N reconfigurations must be tested in different time slots to excite all the channel dimensions. Only a concatenation of the channels to/from the RIS are observed and mutual coupling between RIS elements complicates the estimation. This approach is illustrated in Fig. 6 and requires a wireless control loop between the receiver and the RIS controller circuit with a capacity proportional to N. Even when CSI has been acquired, it is computationally complex to select appropriate time-delays, particularly in wideband channels [12]. To reduce complexity, adjacent RIS elements can be grouped to have the same configuration [12], at the cost of a performance loss.
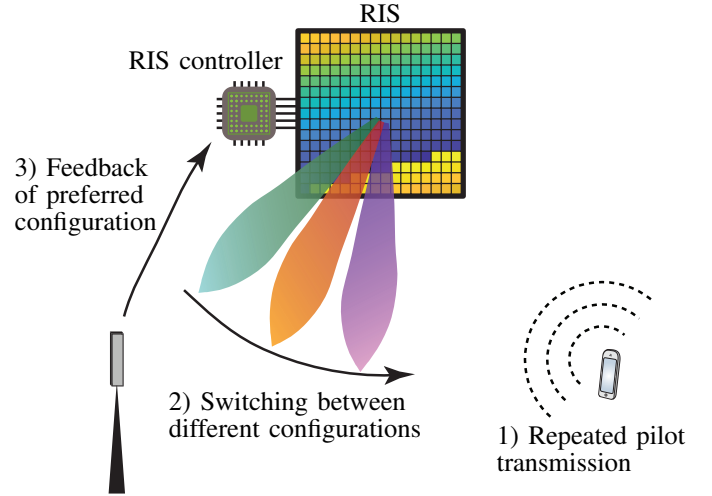

Fig. 6. One approach to configure the RIS is to transmit pilots that the RIS scatters using different configurations. The receiver feeds back a preferred configuration to the RIS.

Another approach is to alter the passive nature of the RIS by having a few elements with receiver chains [13], which enables sensing and channel estimation directly at the RIS. The ability to extrapolate a few measurements to estimate the entire wideband channel requires spatially sparse channels with a known parametrization. This might be reasonable in mmWave or terahertz bands but further work on channel and hardware modeling is required. The sparseness can also make the channels flat over rather wide bandwidths. Learning-based and sparsity-based estimation algorithms were considered in [13], [14]. Even if the RIS has sensing capabilities, a control loop is needed to jointly select the RIS configuration and the beamforming at the transmitter/receiver.

Estimation algorithms can leverage special channel characteristics to reduce the pilot overhead. For instance, the channel between the BS and RIS is semi-static and common for all users, which makes the end-to-end channels correlated between users. An estimation algorithm exploiting this correlation was proposed in [15]. The BS-to-RIS channel can contain many coefficients if the BS has many antennas but since this channel is semi-static, it can be estimated less frequently than the RIS-to-user channel, which typically contains fewer coefficients since users have fewer antennas.

There is no doubt that RIS can be used for fixed communication links, but mobile operation requires real-time channel estimation and reconfiguration, even in indoor use cases. A few millimeters of movement will change the channels in mmWave bands and above. It remains to be demonstrated if any estimation protocol can enable real-time reconfigurability and under what mobility conditions. Since the array is passive, the RIS technology is potentially more energy-efficient than alternative technologies [8] but this remains to be demonstrated quantitatively. The RIS will require a power source for reconfigurability and wireless control channels. It is likely that the control interface will consume most of the power at the RIS, so one cannot predict the total power consumption

until the channel estimation and reconfigurability have been solved and validated.

SUMMARY

An RIS is a full-duplex transparent relay that synthesizes the scattering behavior of an arbitrarily shaped object. Since the RIS is not amplifying the signal, a larger surface area is required to achieve a given SNR than using conventional relays or multi-antenna transceivers. RIS-aided communication is an emerging research topic where the main open problems are to identify convincing use cases and designing practical protocols for reconfigurability.

REFERENCES

[1] C. Huang, A. Zappone, G. C. Alexandropoulos, M. Debbah, and C. Yuen, "Reconfigurable intelligent surfaces for energy efficiency in wireless communication," *IEEE Trans. Wireless Commun.*, vol. 18, no. 8, pp. 4157–4170, 2019.

[2] M. D. Renzo, A. Zappone, M. Debbah, M.-S. Alouini, C. Yuen, J. de Rosny, and S. Tretyakov, "Smart radio environments empowered by reconfigurable intelligent surfaces: How it works, state of research, and road ahead," *arXiv, abs/2004.09352v1*, 2020.

[3] C. Liaskos, S. Nie, A. Tsioliaridou, A. Pitsillides, S. Ioannidis, and I. Akyildiz, "A new wireless communication paradigm through software-controlled metasurfaces," *IEEE Commun. Mag.*, vol. 56, no. 9, pp. 162–169, 2018.

[4] Q. Wu and R. Zhang, "Towards smart and reconfigurable environment: Intelligent reflecting surface aided wireless network," *IEEE Commun. Mag.*, vol. 58, no. 1, pp. 106–112, 2020.

[5] O. Tsilipakos *et al.*, "Toward intelligent metasurfaces: The progress from globally tunable metasurfaces to software-defined metasurfaces with an embedded network of controllers," *Advanced Optical Materials*, no. 2000783, 2020.

[6] Ö. Özdogan, E. Björnson, and E. G. Larsson, "Intelligent reflecting surfaces: Physics, propagation, and pathloss modeling," *IEEE Wireless Commun. Lett.*, vol. 9, no. 5, pp. 581–585, 2020.

[7] S. V. Hum and J. Perruisseau-Carrier, "Reconfigurable reflectarrays and array lenses for dynamic antenna beam control: A review," *IEEE Trans. Antennas Propag.*, vol. 62, no. 1, pp. 183–198, 2014.

[8] E. Basar, M. Di Renzo, J. De Rosny, M. Debbah, M. Alouini, and R. Zhang, "Wireless communications through reconfigurable intelligent surfaces," *IEEE Access*, vol. 7, pp. 116753–116773, 2019.

[9] M. Dohler and Y. Li, *Cooperative Communications: Hardware, Channel and PHY*. John Wiley & Sons, 2010.

[10] M. Heino, D. Korpi, T. Huusari, E. Antonio-Rodriguez, S. Venkatasubramanian, T. Riihonen, L. Anttila, C. Icheln, K. Haneda, R. Wichman, and M. Valkama, "Recent advances in antenna design and interference cancellation algorithms for in-band full duplex relays," *IEEE Commun. Mag.*, vol. 53, no. 5, pp. 91–101, 2015.

[11] E. Björnson and L. Sanguinetti, "Power scaling laws and near-field behaviors of massive MIMO and intelligent reflecting surfaces," *IEEE Open J. Commun. Soc.*, 2020.

[12] B. Zheng and R. Zhang, "Intelligent reflecting surface-enhanced OFDM: Channel estimation and reflection optimization," *IEEE Wireless Commun. Lett.*, vol. 9, no. 4, pp. 518–522, 2020.

[13] A. Taha, M. Alrabeiah, and A. Alkhateeb, "Deep learning for large intelligent surfaces in millimeter wave and massive MIMO systems," in *IEEE GLOBECOM*, 2019, pp. 1–6.

[14] G. C. Alexandropoulos and E. Vlachos, "A hardware architecture for reconfigurable intelligent surfaces with minimal active elements for explicit channel estimation," in *IEEE ICASSP*, 2020, pp. 9175–9179.

[15] Z. Wang, L. Liu, and S. Cui, "Channel estimation for intelligent reflecting surface assisted multiuser communications: Framework, algorithms, and analysis," *IEEE Trans. Wireless Commun.*, 2020, to appear.

BIOGRAPHIES

Emil Björnson is Associate Professor at Linköping University, Sweden. He has co-authored the textbooks *Optimal Resource Allocation in Coordinated Multi-Cell Systems* (2013) and *Massive MIMO Networks: Spectral, Energy, and Hardware Efficiency* (2017). He received the 2018 IEEE Marconi Prize Paper Award in Wireless Communications and the 2019 IEEE ComSoc Fred W. Ellersick Prize.

Özgecan Özdogan is a Ph.D. student at Linköping University, Sweden. She received her M.Sc. degree from Izmir Institute of Technology, Turkey, in 2017.

Erik G. Larsson is Professor at Linköping University, Sweden, and Fellow of the IEEE. He co-authored *Fundamentals of Massive MIMO* (Cambridge University Press, 2016). He received, among others, the 2015 IEEE ComSoc Stephen O. Rice Prize in Communications Theory, the 2017 IEEE ComSoc Leonard G. Abraham Prize, the 2018 IEEE ComSoc Best Tutorial Paper Award, and the 2019 IEEE ComSoc Fred W. Ellersick Prize.