

Five Disruptive Technology Directions for 5G

Federico Boccardi, Vodafone

Robert W. Heath Jr., University of Texas at Austin

Angel Lozano, Universitat Pompeu Fabra

Thomas L. Marzetta, Bell Labs, Alcatel-Lucent

Petar Popovski, Aalborg University

ABSTRACT

New research directions will lead to fundamental changes in the design of future fifth generation (5G) cellular networks. This article describes five technologies that could lead to both architectural and component disruptive design changes: device-centric architectures, millimeter wave, massive MIMO, smarter devices, and native support for machine-to-machine communications. The key ideas for each technology are described, along with their potential impact on 5G and the research challenges that remain.

INTRODUCTION

The fifth generation (5G) cellular network is coming. What technologies will define it? Will 5G be just an evolution of 4G, or will emerging technologies cause a disruption requiring a wholesale rethinking of entrenched cellular principles? This article focuses on potential disruptive technologies and their implications for 5G.

We classify the impact of new technologies, leveraging the Henderson-Clark model [1], as follows:

- Minor changes at both the node and architectural levels (e.g., the introduction of codebooks and signaling support for a higher number of antennas). We refer to these as *evolutions* in the design.
- Disruptive changes in the design of a class of network nodes (e.g., the introduction of a new waveform). We refer to these as *component changes*.
- Disruptive changes in the system architecture (e.g., the introduction of new types of nodes or new functions in existing ones). We refer to these as *architectural changes*.
- Disruptive changes that have an impact at both the node and architecture levels. We refer to these as *radical changes*. We focus on disruptive (component, architec-

tural, or radical) technologies, driven by our belief that the extremely higher aggregate data rates and the much lower latencies required by 5G cannot be achieved with a mere evolution of the status quo. We believe that the following five potentially disruptive technologies could lead to both architectural and component design changes, as classified in Fig. 1.

1) **Device-centric architectures:** The base-station-centric architecture of cellular systems may change in 5G. It may be time to reconsider the concepts of uplink and downlink, as well as control and data channels, to better route information flows with different priorities and purposes toward different sets of nodes within the network. We present device-centric architectures.

2) **Millimeter wave (mmWave):** While spectrum has become scarce at microwave frequencies, it is plentiful in the mmWave realm. Such a spectrum “el Dorado” has led to an mmWave “gold rush” in which researchers with diverse backgrounds are studying different aspects of mmWave transmission. Although far from being fully understood, mmWave technologies have already been standardized for short-range services (IEEE 802.11ad) and deployed for niche applications such as small-cell backhaul. We discuss the potential of mmWave for broader application in 5G.

3) **Massive MIMO:** Massive multiple-input multiple-output (MIMO)¹ proposes utilizing a very high number of antennas to multiplex messages for several devices on each time-frequency resource, focusing the radiated energy toward the intended directions while minimizing intra- and intercell interference. Massive MIMO may require major architectural changes, particularly in the design of macro base stations, and it may also lead to new types of deployments. We discuss massive MIMO.

4) **Smarter devices:** 2G-3G-4G cellular networks were built under the design premise of having complete control at the infrastructure

This work was carried out when Federico Boccardi was with Bell Labs, Alcatel-Lucent.

¹ Most work in massive MIMO has assumed operation at frequencies of 5 GHz or less. While the same principles may prove useful at millimeter wavelengths, a successful marriage of massive MIMO and millimeter waves may take on a considerably different form.

side. We argue that 5G systems should drop this design assumption and exploit intelligence at the device side within different layers of the protocol stack, for example, by allowing device-to-device (D2D) connectivity or exploiting smart caching at the mobile side. While this design philosophy mainly requires a change at the node level (component change), it also has implications at the architectural level. We argue for smarter devices.

5) **Native support for machine-to-machine (M2M) communication:** A native² inclusion of M2M communication in 5G involves satisfying three fundamentally different requirements associated with different classes of low-data-rate services: support of a massive number of low-rate devices, sustaining a minimal data rate in virtually all circumstances, and very-low-latency data transfer. Addressing these requirements in 5G requires new methods and ideas at both the component and architectural levels, and such is the focus of a later section.

DEVICE-CENTRIC ARCHITECTURES

Cellular designs have historically relied on the axiomatic role of “cells” as fundamental units within the radio access network. Under such a design postulate, a device obtains service by establishing a downlink and an uplink connection, carrying both control and data traffic, with the base station commanding the cell where the device is located. Over the last few years, different trends have been pointing to a disruption of this cell-centric structure:

- The base station density is increasing rapidly, driven by the rise of heterogeneous networks. While heterogeneous networks were already standardized in 4G, the architecture was *not* natively designed to support them. Network densification could require some major changes in 5G. The deployment of base stations with vastly different transmit powers and coverage areas, for instance, calls for a decoupling of downlink and uplink in a way that allows the corresponding information to flow through different sets of nodes [5].

- The need for additional spectrum will inevitably lead to the coexistence of frequency bands with radically different propagation characteristics within the same system. In this context, [6] proposes the concept of a *phantom cell* where the data and control planes are separated: the control information is sent by high-power nodes at microwave frequencies, whereas the payload data is conveyed by low-power nodes at mmWave frequencies.

- A new concept called *centralized baseband* related to the concept of cloud radio access networks is emerging ([7]), where virtualization leads to a decoupling between a node and the hardware allocated to handle the processing associated with this node. Hardware resources in a pool, for instance, could be dynamically allocated to different nodes depending on metrics defined by the network operator.

- Emerging service classes, described later in this article, could require a complete redefinition of the architecture. Current works are looking at architectural designs ranging from centralization

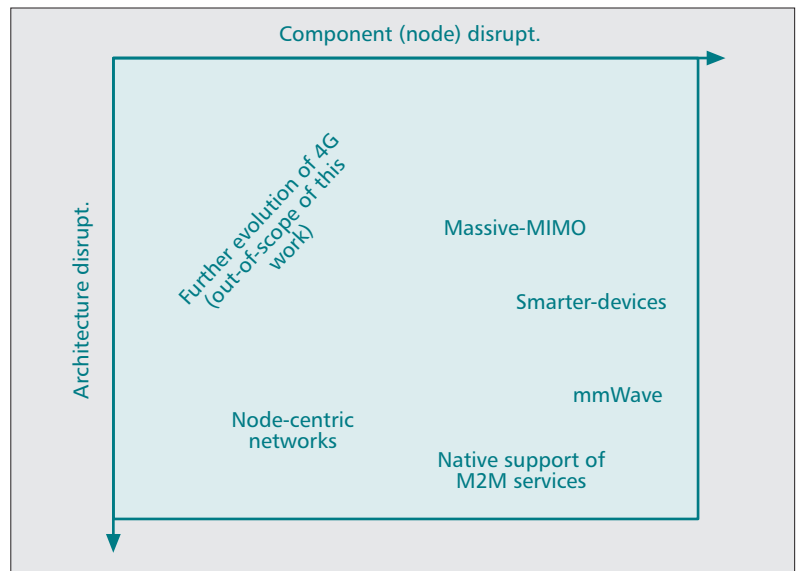


Figure 1. The five disruptive directions for 5G considered in this article, classified according to the Henderson-Clark model.

or partial centralization (e.g., via aggregators) to full distribution (e.g., via compressed sensing and/or multihop).

- Cooperative communications paradigms such as cooperative multipoint (CoMP) or relaying, which despite falling short of their initial hype are nonetheless beneficial [8], could require a redefinition of the functions of the different nodes. In the context of relaying, for instance, recent developments in wireless network coding [9] suggest transmission principles that would allow recovering some of the losses associated with half-duplex relays. Moreover, recent research points to the plausibility of full-duplex nodes for short-range communication in the not-so-distant future.

- The use of smarter devices could impact the radio access network. In particular, both D2D and smart caching call for an architectural redefinition where the center of gravity moves from the network core to the periphery (devices, local wireless proxies, relays).

Based on these trends, our vision is that the cell-centric architecture should evolve into a device-centric one: a given device (human or machine) should be able to communicate by exchanging multiple information flows through several possible sets of heterogeneous nodes. In other words, the set of network nodes providing connectivity to a given device and the functions of these nodes in a particular communication session should be tailored to that specific device and session. Under this vision, the concepts of uplink/downlink and control/data channel should be rethought (Fig. 2).

While the need for a disruptive change in architectural design appears clear, major research efforts are still needed to transform the resulting vision into a coherent and realistic proposition. Since the history of innovations [1] indicates that architectural changes are often the drivers of major technological discontinuities, we believe that the trends above might have a major influence on the development of 5G.

² As was learned with MIMO, first introduced in 3G as an add-on and then natively included in 4G, major improvements come from native support (i.e., from a design that is optimized from its inception rather than amended a posteriori).

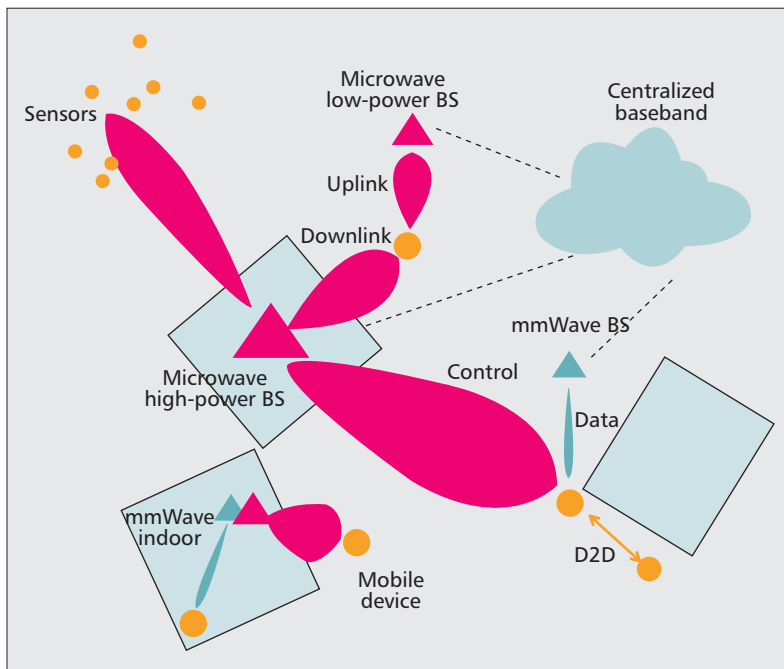


Figure 2. Example of device-centric architecture.

MILLIMETER WAVE COMMUNICATION

Microwave cellular systems have precious little spectrum: around 600 MHz are currently in use, divided among operators [10]. There are two ways to gain access to more microwave spectrum:

- To repurpose or *refarm* spectrum. This has occurred worldwide with the repurposing of terrestrial TV spectrum for applications such as rural broadband access. Unfortunately, repurposing has not freed up that much spectrum, only about 80 MHz, and at a high cost associated with moving the incumbents.
- To share spectrum utilizing, for instance, cognitive radio techniques. The high hopes initially placed on cognitive radio have been dampened by the fact that an incumbent not fully willing to cooperate is a major obstacle to spectrum efficiency for secondary users.

Altogether, it appears that doubling the current cellular bandwidth is the best case scenario at microwave frequencies. Alternatively, there is an enormous amount of spectrum at mmWave frequencies ranging from 3 to 300 GHz. Many bands therein seem promising, including most immediately the local multipoint distribution service at 28–30 GHz, the license-free band at 60 GHz, and the E-band at 71–76 GHz, 81–86 GHz, and 92–95 GHz. Foreseeably, several tens of gigahertz could become available for 5G, offering well over an order of magnitude increase over what is available at present. Needless to say, work needs to be done on spectrum policy to render these bands available for mobile cellular.

Propagation is not an insurmountable challenge. Recent measurements indicate similar general characteristics as at microwave frequencies, including distance-dependent path loss and the possibility of non-line-of-sight communication. A main difference between microwave and

mmWave frequencies is the sensitivity to blockages: the results in [11], for instance, indicate a path loss exponent of 2 for line-of-sight propagation but 4 (plus additional power loss) for non-line-of-sight. MmWave cellular research will need to incorporate sensitivity to blockages and more complex channel models in the analysis, and also study the effects of enablers such as higher density infrastructure and relays. Another enabler is the separation between control and data planes, already mentioned.

Antenna arrays are a key feature in mmWave systems. Large arrays can be used to keep the antenna aperture constant, eliminating the frequency dependence of path loss relative to omnidirectional antennas (when utilized at one side of the link) and providing a net array gain to counter the larger thermal noise bandwidth (when utilized at both sides of the link). Adaptive arrays with narrow beams also reduce the impact of interference, meaning that mmWave systems could more often operate in noise-limited rather than interference-limited conditions. Since meaningful communication might only happen under sufficient array gain, new random access protocols are needed that work when transmitters can only emit in certain directions and receivers can only receive from certain directions. Adaptive array processing algorithms are required that can adapt quickly when beams are blocked by people or some device antennas become obscured by the user's own body.

MmWave systems also have distinct hardware constraints. A major one comes from the high power consumption of mixed signal components, chiefly the analog-to-digital converters (ADCs) and digital-to-analog converters (DACs). Thus, the conventional microwave architecture where every antenna is connected to a high-rate ADC/DAC is unlikely to be applicable to mmWave without a huge leap forward in semiconductor technology. One alternative is a hybrid architecture where beamforming is performed in analog at RF, and multiple sets of beamformers are connected to a small number of ADCs or DACs; in this alternative, signal processing algorithms are needed to steer the analog beamforming weights. Another alternative is to connect each RF chain to a 1-bit ADC/DAC, with very low power requirements; in this case, the beamforming would be performed digitally but on very noisy data. There are abundant research challenges in optimizing different transceiver strategies, analyzing their capacity, incorporating multiuser capabilities, and leveraging channel features such as sparsity.

A data rate comparison between technologies is provided in Fig. 3, for certain simulation settings, in terms of mean and 5 percent outage rates. MmWave operation is seen to provide very high rates compared to two different microwave systems. The gains exceed the 10 \times spectrum increase because of the enhanced signal power and reduced interference thanks to directional beamforming at both the transmitter and receiver.

From the discussion above, and referring again to the Henderson-Clark model, we conclude that mmWave requires radical changes in the system, as it has a strong impact in both the component and architecture designs. Conse-

quently, we view mmWave as a potentially disruptive technology for 5G, which, provided the above discussed challenges can be tackled, could lead to unrivaled data rates and a completely different user experience.

MASSIVE MIMO

Massive MIMO (also referred to as “Large-Scale MIMO” or “Large-Scale Antenna Systems”) is a form of multiuser MIMO in which the number of antennas at the base station is much larger than the number of devices per signaling resource [14]. Having many more base station antennas than devices renders the channels to the different devices quasi-orthogonal and very simple spatial multiplexing/de-multiplexing procedures quasi-optimal. The favorable action of the law of large numbers smooths out frequency dependencies in the channel and, altogether, huge gains in spectral efficiency can be attained (Fig. 4).

In the context of the Henderson-Clark framework, we argue that massive MIMO has a disruptive potential for 5G:

- At a node level, it is a scalable technology. This is in contrast with 4G, which, in many respects, is not scalable: further sectorization therein is not feasible because:
 - There is limited space for bulky azimuthally directive antennas.
 - There is an inevitable angle spread of the propagation; in turn, single-user MIMO is constrained by the limited number of antennas that can fit in certain mobile devices.

In contrast, there is almost no limit on the number of base station antennas in massive MIMO provided that time-division duplexing is employed to enable channel estimation through uplink pilots.

- It enables new deployments and architectures. While one can envision direct replacement of macro base stations with arrays of low-gain resonant antennas, other deployments are possible, such as conformal arrays on the facades of skyscrapers or arrays on the faces of water tanks in rural locations. Moreover, the same massive MIMO principles that govern the use of collocated arrays of antennas also apply to distributed deployments in which a college campus or an entire city could be covered with a multitude of distributed antennas that collectively serve many users (in this framework, the centralized baseband concept presented earlier is an important architectural enabler).

While very promising, massive MIMO still presents a number of research challenges. Channel estimation is critical and currently represents the main source of limitations. User motion imposes a finite coherence interval during which channel knowledge must be acquired and utilized, and consequently there is a finite number of orthogonal pilot sequences that can be assigned to the devices. Reuse of pilot sequences causes pilot contamination and coherent interference, which grows with the number of antennas as fast as the desired signals. The mitigation of pilot contamination is an active research topic. Also, there is still much to be learned about massive MIMO propagation, although

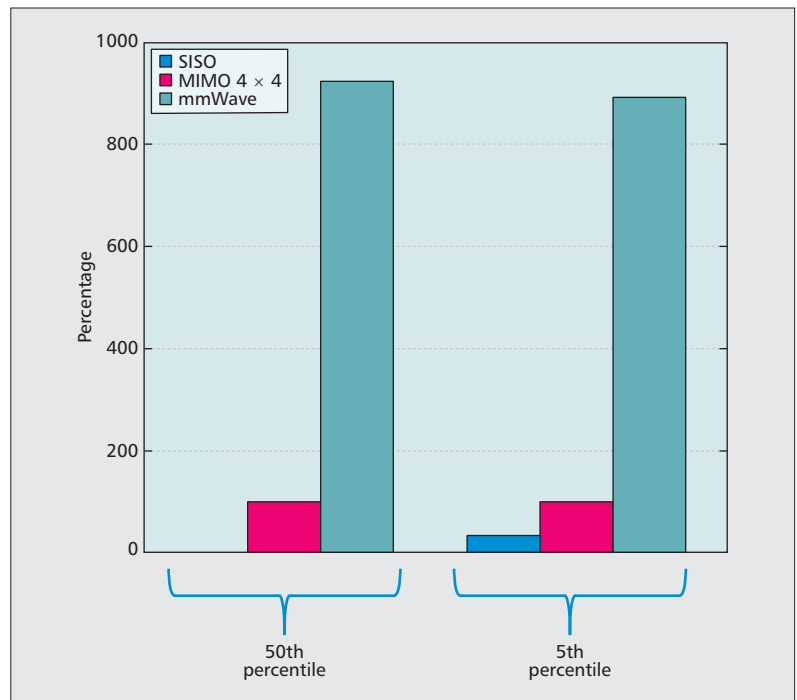


Figure 3. Cell data rate comparison between microwave systems using 50 MHz of bandwidth (single-user single-antenna and single-user MIMO) and a mmWave system with 500 MHz of bandwidth and a single user. Results are given in terms of gain (%) w.r.t. the MIMO 4 × 4 baseline.³ More details about the comparison setup are provided in [12].

experiments thus far support the hypothesis of channel quasi-orthogonality. From an implementation perspective, massive MIMO can potentially be realized with modular low-cost low-power hardware with each antenna functioning semi-autonomously, but a considerable development effort is still required to demonstrate the cost effectiveness of this solution. Note that at the microwave frequencies considered in this section, the cost and energy consumption of ADCs/DACs are sensibly lower than at mmWave frequencies.

From the discussion above, we conclude that the adoption of massive MIMO for 5G could represent a major leap with respect to today’s state of the art in system and component design. To justify these major changes, massive MIMO proponents should further work on solving the challenges emphasized above and showing realistic performance improvements by means of theoretical studies, simulation campaigns, and testbed experiments.

SMARTER DEVICES

Earlier generations of cellular systems were built on the design premise of having complete control at the infrastructure side. In this section, we discuss some of the possibilities that can be unleashed by allowing the devices to play a more active role and then how 5G’s design should account for an increase in device smartness. We focus on three different examples of technologies that could be incorporated into smarter devices: D2D, local caching, and advanced interference rejection.

³ The results given in this article have been obtained by considering the current state-of-the-art understanding of the technologies considered. However, we emphasize that at this point it is not yet possible to provide a fully realistic assessment and a comparison with deployed 4G systems. Undeniably, some research efforts are still in the so-called hype phase, and much work is still required before a steady understanding of the performance and the required enablers can be reached.

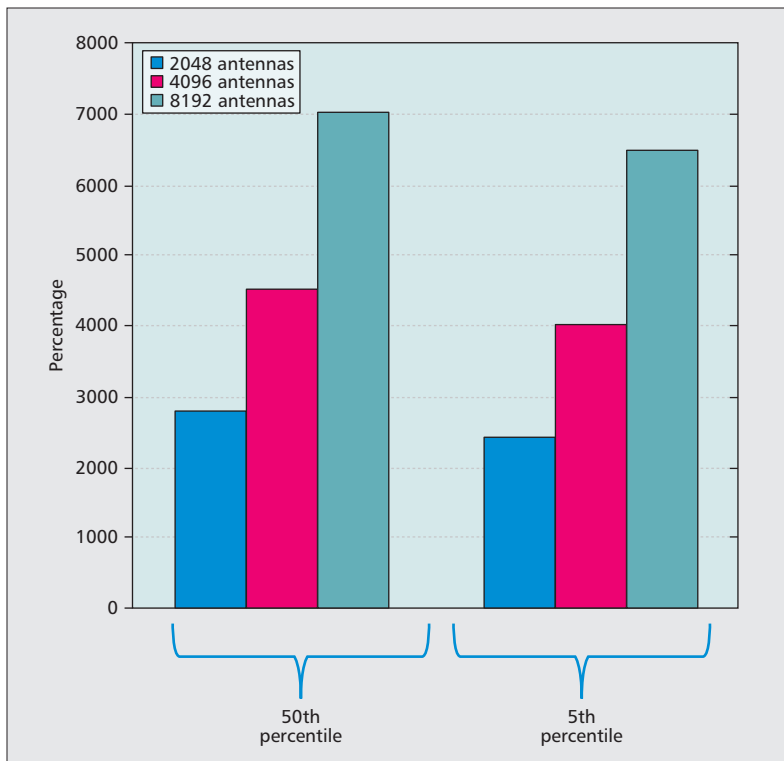


Figure 4. Cell data rate comparison for a fixed access application of massive MIMO. An array of 2048, 4096, or 8192 antennas, utilizing 50 MHz and radiating a total of 120 W, serves 1000 users randomly located in a cell of radius 6 km. Results are given in terms of gain (%) w.r.t. the MIMO 4×4 baseline.

D2D

In voice-centric systems it was implicitly accepted that two parties willing to establish a call would not be in close proximity. In the age of data, this premise might no longer hold, and it could be common to have situations where several co-located devices would like to wirelessly share content (e.g., digital pictures) or interact (e.g., video gaming or social networking). Handling these communication scenarios via simply connecting through the network involves gross inefficiencies at various levels:

- Multiple wireless hops are utilized to achieve what requires, fundamentally, a single hop. This entails a multifold waste of signaling resources as well as higher latency.
- Transmit powers of a fraction of a Watt (in the uplink) and several Watts (in the downlink) are consumed to achieve what requires, fundamentally, a few milliWatts. This, in turn, entails unnecessary levels of battery drain and interference to all other devices occupying the same signaling resources elsewhere.
- Given that the path losses to possibly distant base stations are much stronger than direct link ones, the corresponding spectral efficiencies are also lower.

While it is clear that D2D has the potential to handle local communication more efficiently, local high-data-rate exchanges could also be handled by other radio access technologies such as Bluetooth or WiFi direct. Use cases requiring a mixture of local and nonlocal content or a mixture of low-latency and high-data-rate con-

straints (e.g., interaction between users via augmented reality) could represent more compelling reasons for the use of D2D. In particular, we envision D2D as an important enabler for applications requiring low latency,⁴ especially in future network deployments utilizing baseband centralization and radio virtualization.

From a research perspective, D2D communication presents relevant challenges:

- Quantification of the real opportunities for D2D. How often does local communication occur? What is the main use case for D2D: fast local exchanges, low-latency applications, or energy saving?
- Integration of a D2D mode with the uplink/downlink duplexing structure.
- Design of D2D-enabled devices, from both a hardware and a protocol perspective, by providing the needed flexibility at both the PHY and medium access control (MAC) layers.
- Assessing the true net gains associated with having a D2D mode, accounting for possible extra overheads for control and channel estimation.

Finally, note that while D2D is already being studied in the Third Generation Partnership Project (3GPP) as a 4G add-on (refer to footnote 2), the main focus of current studies is proximity detection for public safety [15]. What we are discussing here is having a D2D dimension *natively* supported in 5G.

LOCAL CACHING

The current paradigm of cloud computing is the result of a progressive shift in the balance between data storage and data transfer: information is stored and processed wherever it is most convenient and inexpensive because the marginal cost of transferring it has become negligible, at least on wireline networks [2]. For wireless devices, though, this cost is not always negligible. The understanding that mobile users are subject to sporadic abundance of connectivity amidst stretches of deprivation is hardly new, and the natural idea of opportunistically leveraging the former to alleviate the latter has been entertained since the 1990s [3]. However, this idea of caching massive amounts of data at the edge of the wireline network right before the wireless hop only applies to delay-tolerant traffic, and thus made little sense in voice-centric systems. Caching might finally make sense now in data-centric systems [4].

Thinking ahead, it is easy to envision mobile devices with truly vast amounts of memory. Under this assumption, and given that a substantial share of the data that circulates wirelessly corresponds to the most popular audio/video/social content that is in vogue at a given time, it is clearly inefficient to transmit such content via unicast, but it is frustratingly impossible to resort to multicast because the demand is asynchronous. We hence see local caching as an important alternative, at both the radio access network edge (e.g., at small cells) and mobile devices, also thanks to enablers such as mmWave and D2D.

ADVANCED INTERFERENCE REJECTION

In addition to D2D capabilities and massive volumes of memory, future mobile devices may also have varying form factors. In some instances, the

⁴ Low-latency local communications are discussed also in the next section. However, while the focus in this section is on use cases requiring local human interaction (e.g., video gaming or augmented reality), the focus in the next section is on use cases requiring local interaction between objects (e.g., vehicles).

devices might accommodate several antennas with the consequent opportunity for active interference rejection therein, along with beamforming and spatial multiplexing. A joint design of transmitter and receiver processing, and proper control and pilot signals, are critical to allow advanced interference rejection. As an example, in Fig. 5 we show the gains obtained by incorporating the effects of nonlinear, and intra- and inter-cluster interference awareness into devices with one, two, and four antennas.

While this section has mainly focused on analyzing the implications of smarter devices at a component level, in the previous section we discussed the impact at the radio access network architecture level. We regard smarter devices as having all the characteristic of a disruptive technology for 5G, and therefore we encourage researchers to further explore this direction.

NATIVE SUPPORT FOR M2M COMMUNICATION

Wireless communication is becoming a commodity, just like electricity or water [13]. This commoditization, in turn, is giving rise to a large class of emerging services with new types of requirements. We point to a few representative such requirements, each exemplified by a typical service.

- **A massive number of connected devices:** Whereas current systems typically operate with, at most, a few hundred devices per base station, some M2M services might require over 104 connected devices. Examples include metering, sensors, smart grid components, and other enablers of services targeting wide area coverage.

- **Very high link reliability:** Systems geared at critical control, safety, or production have been dominated by wireline connectivity largely because wireless links did not offer the same degree of confidence. As these systems transition from wireline to wireless, it becomes necessary for the wireless link to be reliably operational virtually all the time.

- **Low latency and real-time operation:** This can be an even more stringent requirement than the ones above, as it demands that data be transferred reliably within a given time interval. A typical example is vehicle-to-X connectivity, whereby traffic safety can be improved through the timely delivery of critical messages (e.g., alert and control).

Figure 6 provides a perspective on the M2M requirements by plotting the data rate vs. the device population size. The figure illustrates where systems currently stand and how the research efforts are expanding them. Area R1 reflects the operating range of today's systems, outlining the fact that the device data rate decreases as its population increases. In turn, R2 is the region that reflects current research aimed at improving the spectral efficiency. Finally, R5 indicates the region where operation is not feasible due to fundamental physical and information-theoretical limits.

Regions R3 and R4 correspond to the emerging services discussed in this section:

- R3 refers to massive M2M communication where each connected machine or sensor transmits small data blocks sporadically.

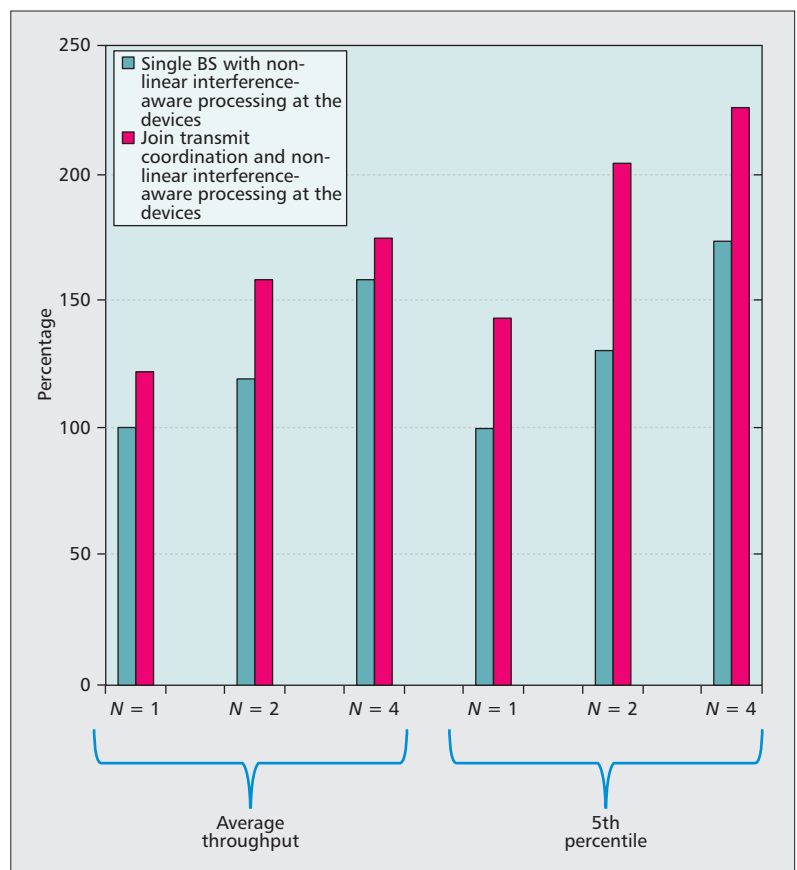


Figure 5. Throughput gains obtained by incorporating the effects of nonlinear, and intra- and inter-cluster interference awareness into devices, with $N = 1, 2,$ and 4 antennas. Results are given in terms of gain (%) w.r.t. the single-base single-antenna baseline. More details can be found in [16].

Current systems are not designed to simultaneously serve the aggregated traffic accrued from a large number of such devices. For instance, a current system could easily serve 5 devices at 2 Mb/s each, but not 10,000 devices each requiring 1 kb/s.

- R4 demarks the operation of systems that require high reliability and/or low latency, but with a relatively low average rate per device. The complete description of this region requires additional dimensions related to reliability and latency.

There are services that simultaneously pose more than one of the above requirements, but the common point is that *the data size of each individual transmission is small*, going down to several bytes. This profoundly changes the communication paradigm for the following reasons:

- Existing coding methods that rely on long codewords are not applicable to very short data blocks.
- Short data blocks also exacerbate the inefficiencies associated with control and channel estimation overheads. Currently, the control plane is robust but suboptimal as it represents only a modest fraction of the payload data; the most sophisticated signal processing is reserved for payload data transmission. An optimized design should aim at much tighter coupling between the data and control planes.

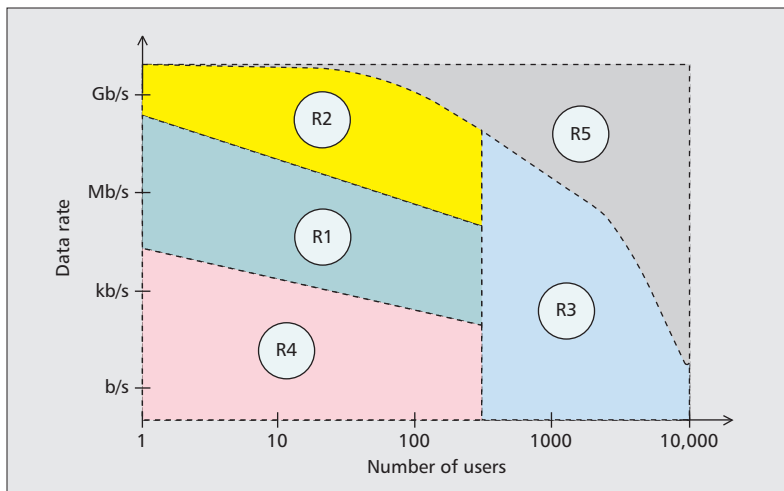


Figure 6. Operating regions in terms of data rate vs. size of the population.

- As mentioned earlier, the architecture needs a major redesign, looking at new types of nodes. At a system level, the frame-based approaches that are at the heart of 4G need rethinking in order to meet the requirements for latency and flexible allocation of resources to a massive number of devices.

From the discussion above, and from the related architectural consideration earlier, and referring one last time to the Henderson-Clark model, we conclude that native support of M2M in 5G requires radical changes at both the node and architecture levels. Major research work remains to be done to come up with concrete and interworking solutions enabling M2M inside 5G systems.

CONCLUSION

This article has discussed five disruptive research directions that could lead to fundamental changes in the design of cellular networks. We have focused on technologies that could lead to both architectural and component design changes: device-centric architectures, mmWave, massive MIMO, smarter devices, and native support of M2M. It is likely that a suite of these solutions will form the basis of 5G.

REFERENCES

- [1] A. Afuah, *Innovation Management: Strategies, Implementation and Profits*, Oxford Univ. Press, 2003.
- [2] J. Zander and P. Mähönen, "Riding the Data Tsunami in the Cloud: Myths and Challenges in Future Wireless access," *IEEE Commun. Mag.*, vol. 51, no. 3, Mar. 2013, pp. 145–51.
- [3] D. Goodman et al., "Infostations: A New System Model for Data and Messaging Services," *Proc. IEEE VTC*, vol. 2, Rome, Italy, May 1997, pp. 969–73.
- [4] N. Golrezaei et al., "FemtoCaching and Device-to-Device Collaboration: A new architecture for Wireless Video Distribution," *IEEE Commun. Mag.*, vol. 51, no. 1, Apr. 2013, pp. 142–49.
- [5] J. Andrews, "The Seven Ways HetNets Are a Paradigm Shift," *IEEE Commun. Mag.*, vol. 51, no. 3, Mar. 2013, pp. 136–44.
- [6] Y. Kishiyama et al., "Future Steps of LTE-A: Evolution towards Integration of Local Area and Wide Area Systems," *IEEE Wireless Commun.*, vol. 20, no. 1, Feb. 2013, pp. 12–18.
- [7] "C-RAN: The Road Towards Green RAN," China Mobile Research Inst., Beijing, China, white paper, v. 2.5, Oct. 2011.

- [8] A. Lozano, R. W. Heath Jr., and J. G. Andrews, "Fundamental Limits of Cooperation," *IEEE Trans. Info. Theory*, vol. 59, no. 9, Sept. 2013, pp. 5213–26.
- [9] C. D. T. Thai et al., "Multi-Flow Scheduling for Coordinated Direct and Relayed Users in Cellular Systems," *IEEE Trans. Commun.*, vol. 61, no. 2, Feb. 2013, pp. 669–78.
- [10] Z. Pi and F. Khan, "An Introduction to Millimeter-Wave Mobile Broadband Systems," *IEEE Commun. Mag.*, vol. 49, no. 6, June 2011, pp. 101–07.
- [11] T. Rappaport et al., "Millimeter Wave Mobile Communications for 5G Cellular: It Will Work!," *IEEE Access*, vol. 1, 2013, pp. 335–49.
- [12] R. W. Heath Jr., "What Is the Role of MIMO in Future Cellular Networks: Massive? Coordinated? mmWave?," *ICC Wksp. Plenary: Beyond LTE-A*, Budapest, Hungary; slides at http://users.ece.utexas.edu/~rheath/presentations/2013/Future_of_MIMO_Plenary_Heath.pdf.
- [13] M. Weiser, "The Computer for the 21st Century," *Sci. Amer.*, Sept. 1991.
- [14] T. L. Marzetta, "Noncooperative Cellular Wireless with Unlimited Numbers of Base Station Antennas," *IEEE Trans. Wireless Commun.*, vol. 9, no. 11, Nov. 2010, pp. 3590–3600.
- [15] 3GPP TR 23.703 v.0.3.0, "Study on Architecture Enhancements to Support Proximity Services (ProSe)," Apr. 2013.
- [16] P. Baracca, F. Boccardi, and N. Benvenuto, "A Dynamic Clustering and Resource Allocation Algorithm for Downlink CoMP Systems with Multiple Antenna UEs," <http://arxiv.org/abs/1311.5114>, 2013.

BIOGRAPHIES

FEDERICO BOCCARDI (federico.boccardi@vodafone.com) is a principal engineer in Vodafone. He received his M.Sc and Ph.D. degrees in telecommunication engineering from the University of Padova, Italy, in 2002 and 2007 respectively. He was with Bell Labs UK from 2006 to 2010 and with Bell Labs Germany from 2010 to 2013. He has participated in different EU collaborative projects and in the 3GPP standardisation activity for LTE and LTE-Advanced. Besides his research activity in wireless communication, his interests fall in the intersection between technology innovation and strategy.

ROBERT W. HEATH JR. [F'11] (rheath@utexas.edu) is a Cullen Trust Endowed Professor in the Electrical and Communications Engineering Department at the University of Austin, Texas, and director of the Wireless Networking and Communications Group. He received B.S. and M.S. degrees in electrical engineering from the University of Virginia and his Ph.D. in electrical engineering from Stanford University. He is the president and CEO of MIMO Wireless Inc. and chief innovation officer at Kuma Signals LLC. He is a registered Professional Engineer in Texas.

ANGEL LOZANO [F'14] (angel.lozano@upf.edu) is a professor and the vice-rector for research at Universitat Pompeu Fabra, Barcelona. He received his Ph.D. from Stanford University in 1998 and was with Bell Labs between 1999 and 2008. He was also, between 2005 and 2008, an adjunct associate professor at Columbia University. He is Chair of the IEEE Communication Theory Technical Committee and an elected member of the Board of Governors of the IEEE Communications Society. He serves as Associate Editor for *IEEE Transactions on Information Theory* and has participated in standardization activities for 3GPP, 3GPP2, and IEEE 802.20.

THOMAS L. MARZETTA [F'03] (tim@research.bell-labs.com) received his Ph.D. in electrical engineering from MIT in 1978. He joined Bell Laboratories in 1995. Within the former Mathematical Sciences Research Center he was director of the Communications and Statistical Sciences Department. He is the originator of Massive MIMO. He was the recipient of the 1981 ASSP Paper Award and received the 2013 IEEE Guglielmo Marconi Best Paper Award.

PETAR POPOVSKI (petarp@es.aau.dk) received his Dipl.-Ing. (1997) and Magister Ing. (2000) in communication engineering from Sts. Cyril and Methodius University, Skopje, Macedonia, and his Ph.D. from Aalborg University, Denmark (2004). He is currently a professor at Aalborg University. He is an Editor in *IEEE Communications Letters*, the *IEEE JSAC Cognitive Radio Series*, and *IEEE Transactions on Communications*. He is a Steering Committee member for the *IEEE Internet of Things Journal*. His research interests are in wireless communications and networking.