Networks and Devices for the 5G Era

Boyd Bangerter, Shilpa Talwar, Reza Arefi, and Ken Stewart, Intel

ABSTRACT

Mobile services based on 4G LTE services are steadily expanding across global markets, providing subscribers with the type of responsive Internet browsing experience that previously was only possible on wired broadband connections. With more than 200 commercial LTE networks in operation as of August 2013 [1], LTE subscriptions are expected to exceed 1.3 billion by the end of 2018 [2]. LTE's rapid uptake, based on exponential growth in network data traffic, has opened the industry's eyes to an important reality: the mobile industry must deliver an economically sustainable capacity and performance growth strategy; one that offers increasingly better coverage and a superior user experience at lower cost than existing wireless systems, including LTE. This strategy will be based on a combination of network topology innovations and new terminal capabilities. Simple network economics also require that the industry's strategy enable new services, new applications, and ultimately new opportunities to monetize the user experience. To address these pressing requirements, many expert prognosticators are turning their attention to future mobile broadband technologies and standards (i.e., 5G) as well as evolutions of the 3GPP's existing LTE standard and IEEE 802.11 standards.

INTRODUCTION

This article will examine the network technologies, device technologies, and spectrum considerations on the road to fifth generation (5G). The mobile industry's transition from 4G to 5G, which could take a decade or longer, will see network operators, infrastructure vendors, and device manufacturers progressively implement next-generation technologies even as one or more standardized 5G technologies are being defined. This transition period, which we call the 5G Era, has effectively already begun. Several technologies expected to be key ingredients of 5G, such as small cell base stations, are available today. Many others will be developed and deployed over the coming years.

There are a few notable trends driving the 5G Era transition.

Macrocellular network capacity cannot continue to increase infinitely: Increasing capacity to keep pace with usage and future demands will require networks that are flatter and more distributed. Next-generation networks will also require advanced source coding (e.g., H.265 and beyond), advanced radio access networks (RANs) such as heterogeneous networks (HetNets), and advanced radio access technologies (RATs) such as new wireless wide area network (WWAN) and wireless local area network (WLAN) technologies. At the same time, transport technologies at cell sites (fronthaul and backhaul) will have to be significantly improved in terms of both speed and deployment flexibility.

The industry's concept of mobile performance is evolving, and so are performance metrics: Today, 3G/4G network performance is evaluated on "hard" metrics, including peak data rates, coverage, and spectral efficiency. The 5G Era will see expanded performance metrics centered on the user's quality of experience (QoE), including factors such as ease of connectivity with nearby devices and improved energy efficiency. 5G networks will offer a more user-centric and context-aware experience, delivering personalized content and assistance services. 5G network elements will need to cooperate in new ways to deliver this level of personalization.

The variety of RATs and wireless devices is growing: In the 5G Era, many devices will use multiple RATs and modes ranging from device-to-device (D2D) communications based on WiFi Direct or Long Term Evolution (LTE) Direct, to short-range millimeter-wave (mm-wave) technologies such as WiGig, and even new body area networks oriented toward wearable devices. What is more, the Internet of Things concept contemplates a world where everything from alarm clocks to washing machines to automobiles is connected through mobile-based machine-to-machine (M2M) communications. Maintaining an optimum user experience in such a complex network environment will require closely coordinated RAT selection and management at both the network and device levels.

In short, the 5G Era will achieve these needed improvements in network and service performance and efficiency by providing a technology framework where networks, devices, and applications are co-optimized. We begin by examining the network technologies leading to 5G, followed by an overview of 5G Era device and spectrum issues.

0163-6804/14/$25.00 © 2014 IEEE

NETWORK INNOVATIONS ON THE ROAD TO 5G

Today's 3G and 4G networks are designed primarily with a focus on peak rate and spectral efficiency improvements. In the 5G Era, we will see a shift towards network efficiency with 5G systems based on dense HetNet architectures. HetNets are among the most promising low-cost approaches to meet the industry's capacity growth needs and deliver a uniform connectivity experience.

As is well known, a HetNet comprises a group of small cells that supports aggressive spectrum spatial reuse. However, HetNets will be architected to incorporate an increasingly diverse set of frequency bands within a range of network topologies, including macrocells in licensed bands (e.g., LTE) and small cells in licensed or unlicensed bands (e.g., WiFi). New higher frequency spectrum (e.g., mm-wave) may also be deployed in small cells to enable ultra-high-data-rate services.

In addition to small cells, client devices will become an integral part of the 5G Era network. Together, small cells and D2D communication will form a new underlay tier of low-cost infrastructure that complements the coverage and capacity of conventional cellular networks.

Cost and flexibility of deployment will also be important factors in 5G networks, requiring a shift toward software-based implementations and virtualization technologies. In particular, 5G systems will be able to create multiple virtual core networks tailored to the specialized requirements of particular applications. For example, the system could feature a virtual core network to support M2M, a separate virtual core network to support over-the-top Internet content, and another virtual core network to support operator-differentiated media services, all of which can be configured by dynamically utilizing the network resources from the same or different networks.

Furthermore, the 5G Era network will need flexible and powerful nodes at the edge to offload the traffic from the core network, to manage data flows efficiently by dynamically adjusting network resources to ensure high QoE for each application flow, and to process the raw information coming from the multitude of sensors/Internet of Things devices. More content will be cached at the edge of the network to reduce core network traffic during busy hours and reduce latency when content is being retrieved. Pre-caching of user generated content and Internet content based on estimated popularity, social trends, and user presence and preferences will allow network operators to better utilize their network pipelines based on context information.

The next several paragraphs describe the 5G HetNet architecture and enabling technologies. Figure 1 shows some key elements in HetNet evolution, including network densification through small cells and D2D communications, multi-cell cooperation through anchor-booster and cloud-based architecture, and multi-radio interworking at both the network and device levels.

THE EVOLUTION OF HETNET ARCHITECTURE

SMALL CELLS

In 5G HetNets, macro and small cells may be connected to each other via ideal or non-ideal backhaul, resulting in different levels of coordination across the network for mobility and interference management. Increasing degrees of network cooperation, from loose network node coordination to completely centralized control, will provide increasing levels of network capacity. When access to ideal backhaul is not available, anchor-booster architecture may be used to coordinate between macro and small cells. In this architecture, the macro cell operates as an *anchor* base station, and is primarily responsible for control and mobility, while the small cell operates as a *booster* base station and is mainly responsible for offloading data traffic [3]. The separation of data and control plane in anchor-booster architecture eases the integration of other RATs, such as WiFi or future mm-wave RATs, as booster cells within the LTE framework.

CLOUD RAN

In scenarios where small cells can be connected to macrocells with low-latency high-rate ("ideal") backhaul, the baseband signals from several hundred cells can be received and processed at a centralized server platform. This architecture, known as Cloud RAN (C-RAN), creates a super base station with distributed antennas supporting multiple RAN protocols and dynamically adapting its signal processing resources based on the varying traffic load within its geographical coverage [4]. The techniques rely on real-time low-latency virtualization, which provides a pool of resources that can be dynamically allocated for baseband processing. C-RAN architecture saves on operational cost by locating all the processing of multiple base stations in one unit, and simplifies implementation of LTE-Advanced features such as coordinated multipoint (CoMP) and enhanced intercell interference coordination (eICIC) by centralizing baseband processing. The evolution of C-RAN will include even more advanced techniques such as joint processing and demodulation of multiple users' signals, and joint resource allocation across multiple RATs to further increase 5G capacity. The Cloud RAN architecture is typically favored by operators with access to optical fiber and low-cost wireless fronthaul, or in extremely high-density scenarios such as sports stadiums.

DEVICE-TO-DEVICE COMMUNICATIONS

D2D communication enables the exchange of data traffic directly between user equipment without the use of base stations or the core network other than for assistance in setting up direct connections. D2D communication supports new usage models based on the proximity of users, including social networking applications, peer-to-peer content sharing, and public safety communications in the absence of network coverage. Additionally, D2D communication serves as another "cell tier" in the 5G

Today's 3G and 4G networks are designed primarily with a focus on peak rate and spectral efficiency improvements. In the 5G Era, we will see a shift towards network efficiency with 5G systems based on dense HetNet architectures.

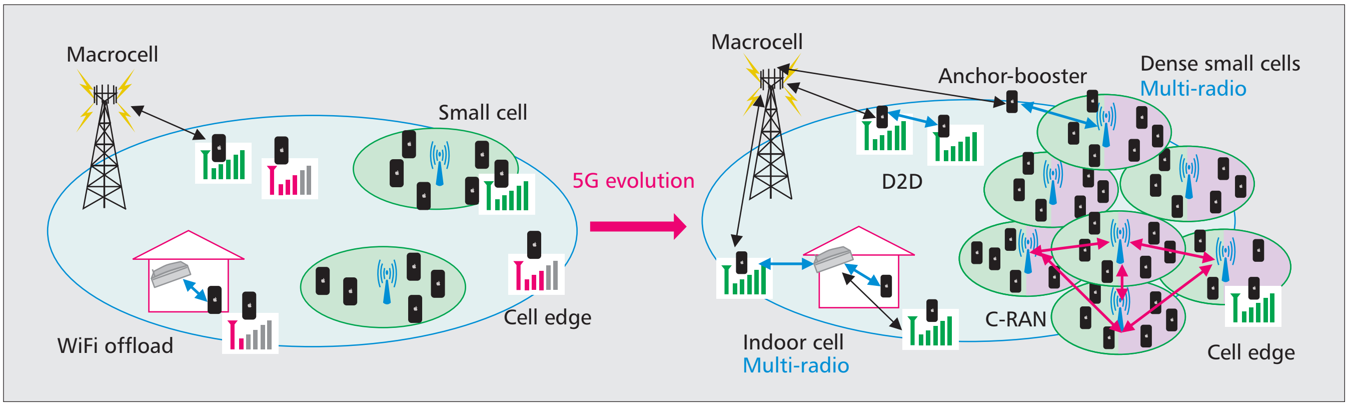

Figure 1. *The evolution of HetNet architecture.*

HetNet, where clusters of devices cooperate with each other to dramatically increase network capacity, either reusing the same spectrum as the macrocell or using unlicensed spectrum. D2D communication also confers additional benefits beyond increased area spectral efficiency, including improved cellular coverage, reduced end-to-end latency, and reduced power consumption.

MULTI-RAT VIRTUAL RADIO ACCESS NETWORKS

As previously noted, 5G networks are expected to support multiple RATs with overlapping coverage deployed as part of a single multi-radio HetNet, creating rich opportunities for intelligently combining and aggregating capacity across the different RATs. To accomplish this, 5G systems will need to support end-to-end network architectures and protocols that seamlessly combine multiple RATs and technologies (operating over licensed, unlicensed, and higher-frequency bands) together into a single virtual RAN, and do so in a manner transparent to the end users. An integrated virtual radio network will enable joint management and simultaneous use of radio resources across different radio technologies to significantly improve radio capacity, and enhance coverage and wireless link reliability. It will also ensure seamless application sessions across the virtual radio network by enabling simultaneous transport and dynamic switching of application flows and radio bearers over multiple physical radio networks.

ENABLING TECHNOLOGIES FOR IMPROVING NETWORK EFFICIENCY

NETWORK COOPERATION AND ADVANCED INTERFERENCE MITIGATION

The aggressive spectral reuse envisioned with dense HetNet architectures will not be realizable without advanced interference management to control the resulting network interference. 5G systems will need to manage such interference through cooperation across densely deployed small cells and end-user devices to provide an "edgeless" experience to mobile users.

In practice, both availability of channel state information and backhaul capacity/latency may limit the achievable gains in multi-cell cooperation technologies. Several practical strategies have recently been proposed to approach the rates promised by theory [5–7]. Interference coordination techniques promise linear scaling of capacity with the number of network nodes [8, 9]. Interference alignment tries to align multiple interferers in a sub-space with smaller dimension than the number of interferers [10]. 5G systems will leverage fundamental ideas from the last decade of developments in this area, but they will need to address practical limitations of channel feedback and backhaul imperfections.

MASSIVE MIMO

Advanced multiple-input multiple-output (MIMO) techniques are at the heart of achieving higher capacity for cellular systems. Multiuser MIMO (MU-MIMO) offers increased multiplexing gains [11], and even though it has been included in the Third Generation Partnership Project (3GPP) LTE-Advanced standard, its full potential has yet to be realized. Drastically higher capacity can be obtained by very large MIMO (VLM) arrays [10] employed at the base station. Increasing transmit array size has desirable implications for coverage, intersymbol and intracell interference control, and transmit power budget optimization [12, 13]. Fortunately, most of the gains can be realized even at manageable antenna dimensions [14]. It is expected that VLM will be a core technology to create significantly higher capacity either in the form of distributed radio heads with centralized processing [15] or in deployment of hundreds of antenna elements in higher frequency bands such as mm-wave, where antenna dimensions become more practical.

SIMULTANEOUS TRANSMISSION AND RECEPTION

As physical layer technologies approach Shannon capacity, simultaneous transmission and reception (STR) at the same frequency and time can provide higher spectral efficiency. Doubling of spectral efficiency can be immediately realized in point-to-point communication such as wireless backhaul [16]. In addition, STR can improve network efficiency in contention networks such as WiFi by mitigating the hidden node problem. When a node receives a packet designated to it and meanwhile has a packet to transmit, STR capability enables it to transmit the packet while receiving the designated packet. This not only doubles the throughput, but also enables hidden nodes to better detect active nodes in their neighborhood. On the other hand, when the node has no packet to send, it can transmit a dummy signal so that any hidden node can detect the activity in its vicinity and realize that the channel is in use. Likewise, STR makes device discovery easier in D2D communication systems where a device can discover neighboring devices easily by monitoring uplink signals from proximate user equipment without suspending transmission. It is expected that STR will play an important role in 5G systems, in both evolution and integration of WiFi networks, and enhancements to D2D communications.

CONTEXT AWARENESS

In order to provide a high QoE for services, 5G systems will need to be context-aware, utilizing context information in a real-time manner based on network, devices, applications, and the user and his/her environment. This context awareness will allow improvements in the efficiency of existing services, and help provide more user-centric and personalized services. For example, networks will need to be more aware of application requirements, QoE metrics, and specific ways to adapt the application flows to meet the QoE needs of the user. There will need to be new interfaces between the application layers and network layers to efficiently adapt both the application source and networking resources to deliver the best QoE for the most users (capacity).

The context-based adaptations discussed above take into account:

- Device-level context, including battery state, CPU load, and device characteristics
- Application context, such as video, web browsing, gaming, or interactive cloud-based applications; QoE metrics; and video-specific parameters such as on-demand vs. real-time streaming, bit rate and resolution
- User context, such as user-specific preferences on quality, user activity, user location, and user level of distraction
- Environment context, which includes motion, lighting conditions, and proximate devices
- Network context, such as congestion/load, airlink and backhaul quality, available timely throughputs, and alternative network/spectrum availability

In the 5G Era, new ways to abstract and efficiently generate context information are needed, as well as new ways to share context information between the application, network, and devices.

DEVICES IN THE 5G ERA

Looking beyond network improvements, the mobile industry's transition from 4G to 5G will require devices capable of taking a greater role in sharing contextual information and managing all the aforementioned connections. In the years to come, devices — such as the one illustrated conceptually in Fig. 2 — will evolve in size, form, and function, but they will not leave the past behind. Even as 5G services eventually achieve global deployment and availability, legacy 4G and 3G networks will continue to operate on a widespread basis for many years. Devices supporting a new 5G RAT will thus likely also support at least LTE, wideband code-division multiple access (WCDMA)/high-speed packet access (HSPA), WiFi, and Bluetooth, although in some cases the supported air interface technology will have evolved significantly compared to today's deployments, especially LTE and WiFi.

For device chipset and platform suppliers, the need to support an increasing number of RATs intersects a trend toward a highly diverse set of device form factors, with wearable and machine-type devices joining tablets and smartphones. This progressive increase in form factor diversity is driving multiple levels of platform support and capability, making transceiver complexity a key challenge for 5G devices.

Transceiver complexity is rooted in part in the need to support an ever increasing number of bands, a problem illustrated in Table 1, which shows a total of 41 LTE bands — for frequency-division duplex (FDD), time-division duplex (TDD), and unpaired (UNP) variants (e.g., 3GPP Band 29) — defined in the five-year period between 3GPP Release-8 to Release-12.[1] Effectively, this means the definition of two new bands each calendar quarter. This effect is compounded by rapid growth in the number of second-order (CA2) and third-order (CA3) carrier aggregation combinations — which now exceed 50 — defined in the three-year period following Release-10. Deployment of new spectrum access modes such as licensed shared access (LSA) will only increase this challenge.

Given the corresponding pressure to support global device operation in a single stock keeping unit (SKU), in the 5G Era, the number of radio ports supported by each transceiver will increase well beyond the 10 or more multi-RAT ports defined today and will include ports capable of supporting one or more new 5G RATs. Correspondingly, the maximum bandwidth aggregated by LTE will increase from 30 MHz+ in 2013 to over 100 MHz and beyond by 2020, and this will heavily influence the initial radio configuration of 5G RATs in the sub-4-GHz region. It will also provide an incentive for implementation in higher frequency bands where larger chunks of spectrum (> 500 MHz) could be available.

The need to enhance cell edge performance, and to support higher-order spatial multiplexing, will also drive devices to increase the number of supported antenna ports from two to four, for both LTE and any new 5G RAT. The severe limitation on available device volume, made worse by wearable form factors and the slow progress in battery chemistry and supportable energy densities, will accelerate the trend toward wideband antennas shared between multiple RATs and will drive still further the integration of wideband active antenna impedance matching over today's designs. The number of RF observation ports supported by the transceiver and dedicated to impedance matching, active envelope tracking, and power amplifier pre-distortion processing will multiply correspondingly.

A key intersection between transceiver and

Looking beyond network improvements, the mobile industry's transition from 4G to 5G will require devices capable of taking a greater role in sharing contextual information and managing all the aforementioned connections.

---

[1] *3GPP bands are often, of course, Release independent. Here we refer to the Release during which a band was first defined.*

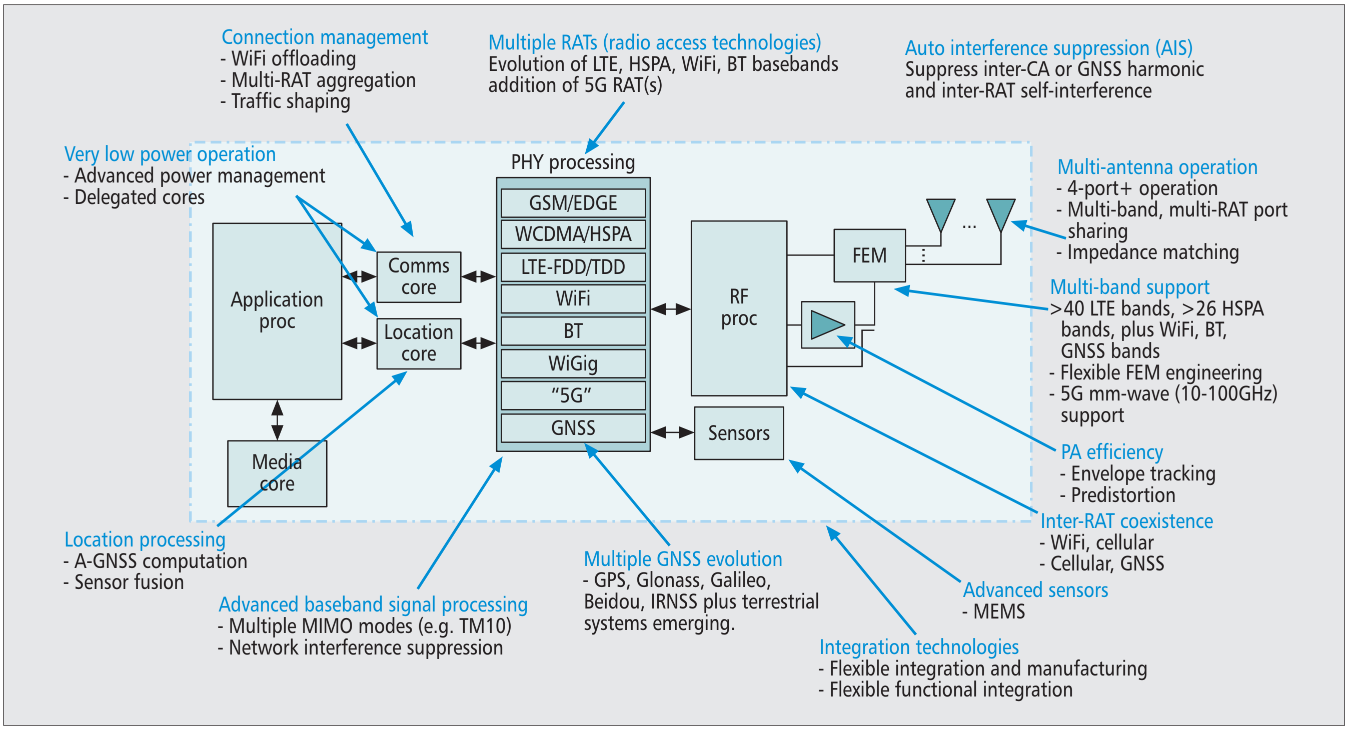

Figure 2. *Device technologies for 5G era.*

baseband capabilities for 3G, 4G, and 5G technologies will arise in the domain of interference management. From the perspective of network capacity, this intersection will require devices to autonomously identify, characterize (interferer modulation and spatial multiplexing mode), and suppress co-channel interference from other cells, as well as suppress increasing levels of in-band interference from other devices operating in D2D modes such as WiFi and LTE Direct, or their 5G equivalents. Such sources of interference external to the device will be exacerbated by self-interference, and by the needs to: operate in multiple aggregated bands, operate multiple RATs simultaneously (including a 5G RAT), share antennas between RATs, manage the resulting coexistence scenarios, and relax the requirements and complexity of multi-band passive radio front-ends (RFEs). As a result, and in addition to other-cell or other-device interference suppression, integrated active self-interference cancelation (SIC) will likely be deployed in widespread fashion as a means of suppressing both direct linear and nonlinear sources of self-interference. In HetNet deployments, clever radio resource management techniques across RATs can help manage interference, such as by selecting WiFi RAT for heavily interfered cell edge users in LTE systems [17].

Advanced signal processing in the device will not be limited to interference suppression. Support for massive MIMO or evolved MU-MIMO, at least in the sub-4-GHz bands, as well as interference-aware coding and new options in non-orthogonal access are candidates for adoption. For any 5G RAT operating in mm-wave frequencies, many of the fundamental problems in antenna array physics and high-speed transceiver design and equalizer design have already been demonstrated by WiGig/802.11ad devices. The principal challenges for mm-wave integration remain cost and low-power operation. However, the possibility of dynamically scheduling the WiGig/802.11ad interface or a new mm-wave 5G RAT to operate only when the resulting energy expenditure per bit is optimal will be enabled by increasingly sophisticated device connection managers.

5G-Era connection managers will not only execute network-based policies controlling RAT activation for power optimization, per-application routing (including D2D routing), and traffic shaping to preserve network control plane resources, but they will also make decisions about which RATs to aggregate for peak-rate enhancement. To achieve truly differentiated 5G-Era peak rates, the industry will need advanced device connection managers that aggregate resources between different networks of the same RAT (e.g., LTE public safety network resources or LSA-aggregated spectrum), combine resources from different RATs (e.g., via multipath IP), and opportunistically exploit mm-wave resources. As we approach realistic 300 Mb/s (Cat. 6) LTE device deployment in 2014, 1 Gb/s peak-rate support in 5G Era devices by 2020 now appears well within reach.

5G Era devices will increasingly support a rich array of device location capabilities. These capabilities will range from processing, with ephemeris assistance data, an ever-maturing set of GPS and GLONASS satellites to include the Beidou, Galileo, and even IRNSS constellations. They will also require support for new frequency bands (e.g., the GPS L2, L5 bands) and new codes and pilot signals that will require 3G-like channel and time-delay estimation and intra-GNSS interference suppression. Critically, indoor location with resolution better than 1 m will be introduced in 5G-Era devices. This will be based on the integration of increasingly sophisticated sensor arrays — including barometric sensing to support altitude estimation in multi-story buildings — combined with WiFi. Additionally, it is expected that Bluetooth-based ranging and RF "fingerprinting" techniques will be deeply embedded into a holistic location engine that incorporates geo-fencing and crowd sourcing techniques.

Power consumption optimization will enable much of the location fusion processing and even connection manager processing to be handled by dedicated, low-power processors integrated into the device system on chip (SoC). Low-power operation will require a dedicated SoC focus on dynamic voltage and frequency scaling techniques and careful clock tree, network on chip, and SoC subsystem state management, as well as continued commitment to high-performance low-leakage silicon processes. As the industry moves beyond 28 nm designs in 2014, we can look to the deployment of sub-10 nm silicon processes as the vehicle for devices supporting a new 5G RAT.

SPECTRUM IN THE 5G ERA

5G systems will need to provide significant improvement in cell capacity to accommodate the rapidly increasing traffic demands. Although 5G will introduce a bevy of new technologies that enable networks and devices to make better use of scarce spectrum resources, more efficient use of current spectrum resources will not be sufficient to keep pace with the mobile data usage increase. 5G systems are expected to provide data rates on the order of gigabits per second, anytime and anywhere. This could only be realized with much more spectrum than that currently available to IMT[2] systems through the International Telecommunication Union's (ITU) process.

Frequency bands currently in use by IMT systems are fragmented with varying degrees of availability and amount of bandwidth across bands, countries, and regions, leading to problems such as roaming, device complexity, lack of economies of scale, and harmful interference. Some technologies for aggregation of IMT bands have been developed, but they have limitations in meeting the broader bandwidth needs of future systems. Therefore, contiguous and broader frequency bands are needed to provide gigabit data rate services in the future.

All spectrum currently available to cellular mobile systems, including IMT, is concentrated in bands below 6 GHz due to the favorable propagation conditions in such bands. For the same reason, these frequencies are in high demand by other services, including fixed, broadcasting, and satellite communications. As a result, these bands have become extremely crowded, and prospects for large chunks of new spectrum for IMT below 6 GHz are not favorable.

Recent advancements in mobile communication systems and devices operating at frequencies around 60 GHz, combined with advancements in antenna and RF component technologies, have opened the gates to using non-conventional bands for cellular applications. Such advancements will help enable dense small cell deployments over a diverse set of spectrum that includes higher unconventional bands. Such deployments will be an important 5G usage scenario as there will be continued need to meet exponential growth in traffic demand and address the requirement for gigabit data rates everywhere, including at the cell edge. It is expected that small cells operating over spectrum not traditionally used by cellular systems (e.g., mm-wave bands, such as bands around 30–50 GHz) will be deployed indoors or outdoors to meet this requirement toward "edgeless" 5G networks.

Release	FDD	TDD	UNP	CA2	CA3
8	11	2	0	—	—
9	16	5	0	—	—
10	23	11	0	4	—
11	27	12	1	26	—
12	28	12	1	41	9+

Table 1. *3GPP LTE band definitions vs. year.*

Another potential means to make large amounts of spectrum available to future IMT systems is through novel schemes such as LSA, whereby certain underutilized non-IMT spectrum could be integrated with other IMT spectrum in a licensed pre-determined manner following mutual agreement among the licensees. In this way, LSA has the potential to expand some existing IMT bands and make them suitable to support channel bandwidth larger than is possible today.

As the mobile industry continues its efforts to open new technological frontiers, governments worldwide can also play a role in meeting growing data demand by adopting sensible, innovative, and technology-neutral spectrum policies to facilitate more efficient use of spectrum, and thus create new economic opportunities.

The combination of network innovations, new device capabilities, and support from key ecosystem stakeholders will help pave the road to 5G, and will progressively enrich our mobile user experience throughout the 5G Era.

REFERENCES

[1] GSMA Intelligence, "Mapping Worldwide 4G-LTE Network Launches," Aug. 2013, http://gsmaintelligence.com/analysis/2013/08/dashboard-mapping-worldwide-4g-lte-network-launches-august-2013/399.

[2] Informa Telecoms & Media, *LTE Spectrum Strategies and Forecasts to 2018*, 3rd ed., Aug. 2013, http://www.informa.com/Media-centre/Press-releases—news/Latest-News/Informa-Telecoms-and-Media-Huawei-and-Ericsson-dominate-LTE-contracts-as-deployments-accelerate.

[3] 3GPP TR 36.932: "Scenarios and Requirements for Small Cell Enhancements for E-UTRA and E-UTRAN," v. 0.2.0, May 2013.

[4] China Mobile, "C-RAN Road Towards Green Radio Access Network," *C-RAN Int'l. Wksp.*, Beijing, Apr. 2010, http://labs.chinamobile.com/report/34516.

[5] 3GPP TR 36.814: "Further Advancements for EUTRA Physical Layer Aspects," v. 9.0.0, 2010.

Although 5G will introduce a bevy of new technologies that enable networks and devices to make better use of scarce spectrum resources, more efficient use of current spectrum resources will not be sufficient to keep pace with the mobile data usage increase.

---

[2] *International Mobile Telecommunication (IMT) is the generic name adopted by ITU administrations to represent terrestrial mobile cellular systems and their enhancements and evolutions. IMT encompasses IMT-2000 (commonly referred to as 3G) and IMT-Advanced (commonly referred to as 4G).*

The combination of network innovations, new device capabilities and support from key ecosystem stakeholders will help pave the road to 5G, and will progressively enrich our mobile user experience throughout the 5G Era.

[6] 3GPP TR 36.819, "Coordinated Multi-point Transmission for LTE Physical Layer Aspects," v. 11.1.0, 2011.

[7] Alcatel-Lucent, "Light-Radio Portfolio: Technical Overview," technical white-paper, 2011.

[8] J. Andrews, "Can Cellular Networks Handle 1000x the Data?" http://users.ece.utexas.edu/~bevans/courses/realtime/lectures/Andrews_Cellular1000x_Nov2011.pdf, 2011.

[9] A. Zemlianov and G. de Veciana, "Capacity of Ad-hoc Wireless Networks with Infrastructure Support," *IEEE JSAC*, 2005.

[10] F. Rusek *et al.*, "Scaling up MIMO: Opportunities and Challenges with Very Large Arrays," *IEEE Sig. Proc. Mag.*, Jan. 2013.

[11] D. Gesbert *et al.*, "From Single User to Multiuser Communications: Shifting the MIMO Paradigm," *IEEE Sig. Proc. Mag.*, vol. 24, no. 5, Oct. 2007, pp. 36–46.

[12] T. L. Marzetta, "Non-Cooperative Cellular Wireless with Unlimited Numbers of Base Station Antennas," *IEEE Trans. Wireless Commun.*, vol. 9, no. 11, Nov. 2010, pp. 3590–3600.

[13] A. Pitarokoilis, S. K. Mohammed, and E. G. Larsson, "On the Optimality of Single-Carrier Transmission in Large-scale Antenna Systems," *IEEE Wireless Commun. Letters*, May 2012.

[14] J. Hoydis, S. ten Brink, and M. Debbah, "Massive MIMO: How Many Antennas Do We Need," *Proc. IEEE Allerton Conf.*, Urbana-Champaign, IL, Sept. 2011.

[15] Y. Lin*et al.*, "Wireless Network Cloud: Architecture and System Requirements," *IBM J. Research and Development*, vol. 54, no. 1, Jan./Feb. 2010, pp. 4:1–12.

[16] M. Jain *et al.*, "Practical, Real-Time, Full Duplex Wireless," *Proc. Mobicom*, 2011.

[17] A. Y. Panah *et al.*, "Utility-Based Radio Link Assignment in Multi-Radio Heterogeneous Networks," *Proc. IEEE GLOBECOM Wksp. LTE and Beyond 4G Technologies*, Dec. 2012.

BIOGRAPHIES

BOYD BANGERTER directs the Wireless Communications Lab within Intel Labs. A 20-year veteran of the company, he has worked extensively on wireless programs, driving Intel's research into next-generation WLAN and WWAN technologies. Under his leadership, the Wireless Communications Lab performs extensive research and development of next-generation mobile wireless systems and technologies for improving application performance over wireless networks. He holds an M.B.A. from Brigham Young University and received his B.S.E.E. with honors from the California Institute of Technology.

SHILPA TALWAR is a principal engineer in Intel's Wireless Communications Laboratory and leads a research team focused on improving the wireless service experience of 5G subscribers. Prior to Intel, she held several senior technical positions in the wireless industry, working on algorithm and system design for 3G/4G/WiFi networks, satellite communications, and GPS technologies. She graduated from Stanford University with a Ph.D. in applied mathematics and an M.S. in electrical engineering. She is credited with 60+ technical publications and patents.

REZA AREFI is the director of spectrum strategies, radio communication standards, and related global public policy efforts for Intel's Standards & Advanced Technology group. With more than 20 years' telecom industry experience, he has been an active contributor to and leader of various wireless systems standards and industry groups, including ITU-R, IEEE, and the WiGig Alliance. His current focus is on enabling spectrum for 5G cellular systems. He holds numerous patents spanning various wireless communications technologies.

KENNETH (KEN) STEWART is an Intel Fellow and the chief wireless technologist for Intel's Wireless Product R&D organization. In these roles, he contributes to Intel's advancements in wireless system-on-chip solutions, cellular and connectivity transceiver solutions, codecs, radio access offloading and routing, and location systems. A recognized mobile innovator and alumnus of various executive leadership positions at Motorola and TE Connectivity, he holds more than 80 issued patents and is a graduate of the University of Strathclyde, where he is also a visiting professor.