ARTEMIS

pCell

Wireless Reinvented

An Introduction to pCell

White Paper, February 2015

Steve Perlman
Founder & CEO

Antonio Forenza, PhD
Co-Founder & Principal Scientist

An Introduction to pCell

pCell is a revolutionary new approach to wireless, dramatically increasing the spectral efficiency of LTE and Wi-Fi systems, while remaining compatible with standard LTE and Wi-Fi devices. This White Paper describes pCell technology in detail, including vision, history, architecture, applications, performance and comparison to conventional wireless technologies.

To contact Artemis Networks, please visit www.artemis.com.

U.S. patents 7,418,053, 7,451,839, 7,599,420, 7,633,994, 7,636,381, 7,711,030, 7,885,354, 8,160,121, 8,170,081, 8,307,922, 8,428,162, 8,469,122, 8,542,763, 8,571,086, 8,654,815, 8,971,380, international patents, patents pending. pCell, pWave, Artemis and the Artemis logo are trademarks of Artemis Networks LLC.

Copyright © 2015, Artemis Networks, LLC. All Rights Reserved.

Table of Contents

1 pCELL: 5G Now

Artemis™ pCell™ technology is a radical new approach to wireless that increases the capacity of LTE and Wi-Fi networks by over an order of magnitude, while dramatically improving Quality of Service (QoS) and maintaining off-the-shelf compatibility with unmodified LTE and Wi-Fi devices, such as iPhone, iPad and Android devices. pCell meets proposed 5G performance targets today, while remaining compatible with 4G devices.

pCell technology accomplishes this through an entirely new approach to wireless: rather than *avoiding* interference like cellular or Wi-Fi systems, pCell *embraces* interference, utilizing interfering transmissions to synthesize a tiny *personal cell*, a "pCell", around each individual user device, enabling every user to utilize the full capacity of the spectrum at once. Instead of many users sharing the limited capacity of one cell, resulting in steadily declining data rates as new subscribers join the network, with pCell technology, each user gets a personal cell. So, no matter how many users are sharing the same spectrum, <u>each user is able to experience the full capacity of the spectrum concurrently with other users.</u>

pCell has a wide range of advantages over conventional wireless technologies:

- **Leapfrog in spectral efficiency.** LTE networks today achieve a spectral efficiency (SE)[1] of **1.7 bps/Hz**[2]. Practical pCell systems today achieve average spectral efficiency of **59 bps/Hz**, a 35X leapfrog with *unmodified standard LTE devices*, such as iPhone 6/6 Plus, iPad Air 2 and Android devices, as well as Wi-Fi devices. pCell's spectral efficiency scales indefinitely, while remaining compatible with legacy devices.

- **Consistent, reliable data rate.** Cellular or Wi-Fi data rate drops off rapidly from cell center (e.g. 100 Mbps) to cell edge (e.g. 1 Mbps), resulting in highly variable and unreliable service quality. With pCell the data rate remains uniformly near peak throughout the coverage area, including vertically in tall buildings, enabling QoS service offerings, such as 4K UHD video.

- **Low deployment, operations cost.** Cellular radios must be carefully placed in *specific* locations within a "cell plan" (reconfigurable by SON[3] or not) with antennas carefully aimed to avoid interference and dead zones. pCell deployment is fast with minimal real-estate and fronthaul costs because inexpensive pWave™ radio heads can be placed *anywhere* in the coverage area, allowing carriers to choose low-cost real estate locations that have line-of-sight (or low-cost fiber) paths for fronthaul.

- **Easily scalable capacity.** Adding radios to a cellular system to increase capacity requires complex and *precise* placement so as not to interfere with the existing cell plan. pCell capacity scales easily by simply adding new pWave radio heads into the coverage area *wherever* it is inexpensive and convenient. Each added pWave increases the aggregate data capacity of the pCell system.

- **Highly robust for public safety.** During a crisis, if a cell tower is disabled, or if there is a surge in demand that overwhelms the cell, mobile service in the coverage area of the tower will be unavailable. With pCell, many pWave antennas overlap throughout the coverage area, so even if several pWave antennas are disabled or if there is a sudden surge in demand, mobile service is maintained.

- **Physically secure communications.** Since cellular transmissions can be received from anywhere within the coverage area, anyone who gains access to the cryptographic keys can intercept communications[4]. With pCell, the transmitted data only physically exists at the point of reception, providing physically secure communications.

- **Lower power consumption at the user device.** Uplink signals are received by many pWave antennas at once, enabling user devices to transmit at lower power and save battery life.

- **Lower cost user devices.** While pCell is compatible with unmodified standard LTE devices, the same pCell infrastructure simultaneously supports far lower cost and lower power user devices utilizing pCell-optimized protocols in the same spectrum as LTE, enabling carriers to offer far lower cost entry points and support "Internet of Things".

- **Low latency cloud hosting.** Under heavy load, pCell maintains <10ms latency with LTE devices, and <1ms latency with Wi-Fi and pCell-optimized protocols. Fast storage and real-time computing resources are available in pCell data centers with direct access at these latencies.

- **Accurate location positioning**: pCell can provide highly precise 3D location information, whether indoors (including tall buildings) or outdoors, enabling location-based services (e.g. E911 for public safety[5]) where GPS is unavailable or vertical position data is needed.

pCell is a major step forward with wireless, but it is also a very different way of thinking about wireless. Details of these pCell advantages are described in the sections that follow, along with perspective about how pCell stands to transform the communications landscape worldwide.

2 Background

Artemis began pCell development over a decade ago as a project in the Rearden Companies incubator with the goal to establish wireless Internet as a global ubiquitous utility for the 21st century, much as electricity became a ubiquitous utility (in the developed world) in the 20th century.

A decade ago, mobile data traffic was minimal, but given the simultaneous trends of (a) video moving from broadcast/physical media to Internet streaming, and (b) Internet usage moving from wired to mobile[6], we projected mobile data demand to roughly double each year, or about 1000X growth per decade. It quickly became apparent that the biggest challenge we would be facing was the physics limit of wireless capacity (i.e. spectral efficiency or "SE").

Given projected growth, it was quite clear that mobile data demand would far exceed the SE limits of all known wireless technologies, even if all usable spectrum[7] was allocated to mobile.

The severity of the problem was only widely acknowledged recently and dubbed the "Spectrum Crunch"[8]. In 2010, the FCC released the data graphed in Figure 1, projecting U.S. mobile demand exceeding capacity in 2013, and rapidly getting worse thereafter[9].

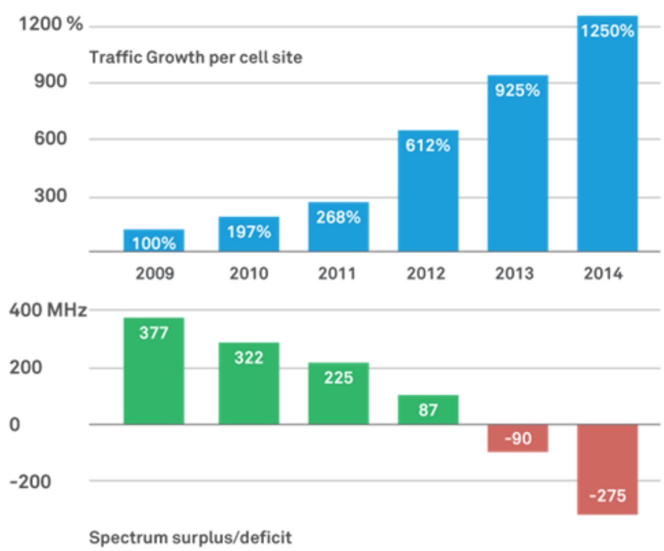

Figure 1: U.S. Mobile data demand vs. spectrum availability [FCC 2010]

The FCC projections were accurate. In 2013, U.S. mobile statistics should have been enviable…

a. The U.S. had more LTE subscribers than all other countries in the world combined[10].
b. LTE represented 24.5% of all U.S. mobile connections, compared to 2.9% globally[11].

c. At 65% penetration, the U.S. was the only region of the world with mostly smart devices. Smart devices devour mobile data, accounting for 88% of global mobile traffic[12].
d. Most of U.S. mobile traffic was video[13], enabled by LTE's speed and smart device apps.

…but instead of enviable performance, U.S. mobile demand far exceeded spectrum capacity, causing ***U.S. average LTE download speed to plummet from 2013 to 2014 by 32% to 6.5Mbps,*** **by far one of the slowest LTE speeds in the world[14]**, per Figure 2. As a comparison, 3G HSPA achieves faster speeds than 6.5Mbps[15].

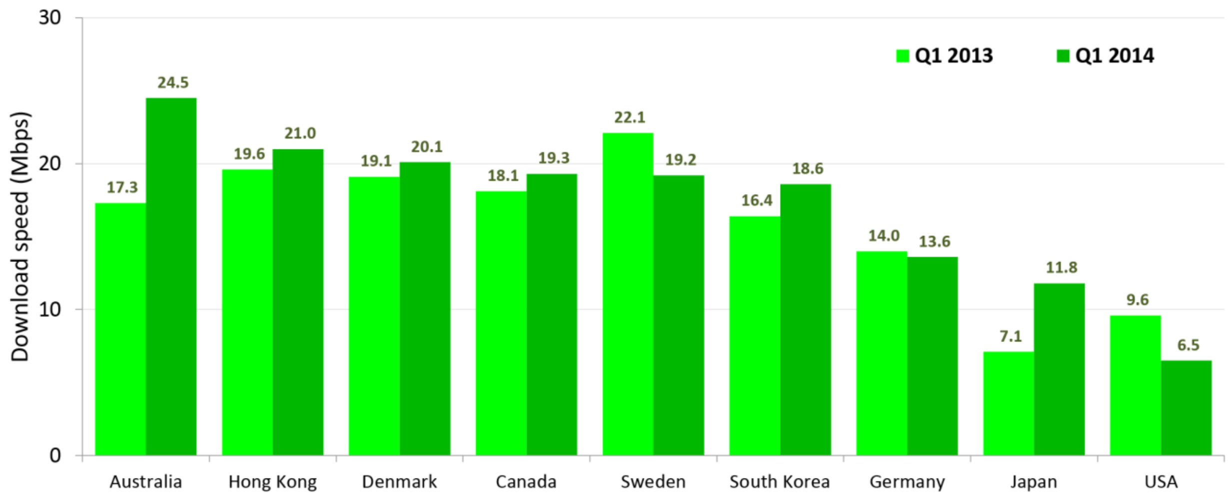

Figure 2: Average LTE download speed by country 2013 vs. 2014 [*OpenSignal 2014*]

Ironically, slow U.S. LTE data rates are a consequence of LTE's market *success.* The same result is seen worldwide: Japan's NTT DOCOMO launched LTE two years before competitors and holds a wide market lead. The consequence of its "success" is an average download speed half that of later Japan LTE market adopters that have far smaller subscriber bases[16], almost as slow as U.S. LTE. The unfortunate reality of cellular technology is *market success results in declining service quality*, ultimately reaching the point where service becomes almost unusable. Such results are already seen in dense cities like New York and Chicago[17]. New spectrum deployments may briefly mitigate congestion[18], but given the inexorable growth of mobile data demand,[19] data rates ultimately decline with market success[20]. Verizon CEO Lowell McAdam put it bluntly to investors in 2013, "…it's the *physics* that breaks it…you just run out of gas."[21]

Costly efforts to increase cell density with small-cells (or tightly packed Wi-Fi access points) have been unable to mitigate the performance decline. Small-cells suffer from increased inter-cell interference and handoff overhead[22], ultimately exceeding capacity gains which establishes a practical upper limit for cell density, even disregarding the economic considerations of backhaul, power, physical access, etc. So, as observed by the FCC and Verizon, the only known option to significantly increase capacity is to allocate more spectrum for mobile use.

Unfortunately, the world is almost out of mobile spectrum. Only a narrow range of frequencies that can efficiently penetrate obstacles are suitable for mobile[23]. Even if all of these frequencies were allocated for mobile, it would only accommodate three years of mobile data growth[24]. After that, all mobile spectrum would be gone for decades, with mobile congestion getting worse every year. As noted in an October 2013 Wall Street Journal op-ed commemorating the 40th anniversary of cellular technology, "…wireless engineers will have to come up with a better way to use the finite amount of spectrum they already have. If they don't, soon enough your smartphone will remind you of the dial-up speeds of the 1990s."[25]

Any doubts as to the severity of the mobile spectrum crunch have been brushed aside by the sky-high bids in the recent FCC AWS-3 spectrum auction. US$44.9 billion was bid for 65 MHz of U.S. mid-band spectrum, more than two times the US$22 billion that the most bullish analysts had projected for the auction[26]. US$44.9 billion almost equals the total raised in *all* prior FCC spectrum auctions—including auctions for far better low-band spectrum[27]—over the last 20 years[28].

Cell phone pioneer Martin Cooper[29] and FCC Commissioner Jessica Rosenworcel[30] underscored the urgency of the spectrum crisis in a September 2014 op-ed[31], noting spectrum is a finite natural resource that is quickly being exhausted:

> *Spectrum is the basis of our new wireless world. But the laws of physics being what they are, we cannot create more. So we need to find ways to use the airwaves we have more efficiently.*

They proposed to award 10 MHz of mobile spectrum (worth billions of U.S. dollars) for the first to develop a practical technology that increases current spectral efficiency by 50 times.

As industry and governments are finally coming to the inescapable conclusion that 4G is unable to meet mobile data demands today, let alone catch up to growing demands in the future, they are urgently turning to new 5G technologies to overcome the spectrum crunch. In the last year several heavily-funded 5G research efforts were initiated, with projected 3GPP standardization efforts starting in 2016[32] and deployment dates ranging from the early- to mid-2020s[33]. All are preliminary R&D efforts, exploring a wide range of directions, with no practical prototypes operating today[34]. So, it is uncertain that 5G—whatever it ends up being—will be deployable by the mid-2020s, let alone be capable of overcoming what will then be a massive spectrum deficit relative to mobile data demand. Even if a viable 5G solution does arrive at some point in the 2020s, between now and then, mobile congestion may well become more than 100X worse than it is today.

2.1 pCell: Meeting 5G requirements today, in compliance with 4G devices

In the early 2000s, we came to the same inescapable conclusion that most of the wireless industry has only recently recognized: all known wireless technologies hit SE upper limits far below what is required to meet skyrocketing wireless data demands. Fundamental R&D and a complete rethink of wireless are necessary to achieve the required leapfrog in SE.

So, while most wireless research worldwide over the last decade focused on *evolutionary* improvements of known wireless technologies (resulting in standards like LTE-A and 802.11ac), we had concluded evolutionary approaches would fall far short of SE requirements, and were entirely focused on *revolutionary* approaches that would leapfrog the SE limits of known wireless technology and continue to scale to not just meet projected traffic demands, but continue to keep up with growth. The outcome of our research, pCell technology, did indeed require a revolutionary approach to wireless, and not only exceeded the original goals of the research, but opened up new avenues of inquiry that led to many advantages in performance, power and cost over conventional wireless technology in addition to pCell's leapfrog in SE.

In fact, pCell's *actual* performance today conforms closely to proposed 5G performance for the 2020s. In 2014 the ITU's IMT-2020 (5G) Promotion Group presented a comprehensive vision for what 5G would look like in the 2020s[35], which is largely consistent with other recent 5G visions[36] and establishes an expert consensus on goals for the next generation of mobile services. Table 1 lists some IMT-2020 requirements for 5G in the 2020s.

Target Spectral Efficiency	Traffic Density per km^2	Connection Density per km^2	User Density	QoS	Min Per-user Data Rate	Latency	Min Device Cost	Min Device Power
45 bps/Hz[37]	**100 Tbps**	**1 million devices**	**Subway, Stadium**	**Reliable, consistent**	**100 Mbps[38]**	**<1ms**	**IoT[39] scale**	**IoT power**

Table 1: IMT-2020 proposed requirements for 5G [*IMT 2020*]

At the heart of the proposed IMT-2020 goals is a huge leapfrog in average SE to 45 bps/Hz, 15X beyond the IMT-Advanced average SE of 3 bps/Hz that LTE-A is targeting.[40]

The IMT-2020 SE goals reflect the inescapable conclusion that the world has almost exhausted its supply of practical mobile spectrum, and that the only hope to achieve such a leapfrog would be a revolutionary approach to wireless. Related to SE are IMT-2020 goals for traffic-, connection- and user-density as well as reliable and consistent QoS to make such performance predictable. The minimum per-user data rate reflects the expectation of a cellular architecture with cell-edge performance of at least 100Mbps, so that at least 100Mbps is available consistently throughout the coverage area.

pCell has been extensively tested with unmodified LTE devices in bands 38, 39, 40 and 41 as well as with lab LTE devices in 900 MHz and 400 MHz bands, both indoor and outdoor. pCell can meet all of the *light* green IMT-2020 targets in Table 1 today on off-the-shelf LTE devices. pCell is compatible with unmodified LTE Release 8 and above devices such as iPhone 6, Galaxy S5, iPad Air 2, and LTE dongles. pCell achieves consistent and reliable SE throughout the coverage area in excess of the 5G target SE in Table 1, and can maintain a consistent and reliable per-user data rate in excess of 100 Mbps with LTE devices capable of that speed[41]. pCell far exceeds all of the Density requirements of Table 1, not only supporting devices at stadium/subway densities, but supporting densities of devices separated by only a few centimeters.

pCell supports new protocols concurrently in the same spectrum as LTE and LTE-A devices such as protocols that would meet the *dark* green IMT-2020 targets in Table 1, which are not achievable within the LTE protocol[42]. As such, pCell future-proofs spectrum to allow concurrent use with yet-to-be finalized (or yet-to-be-conceived) standards.

Current 5G R&D efforts, such as millimeter waves and massive MIMO, are in early prototyping stages and face numerous technical challenges before they can be considered for actual deployment. Initial experimental testbeds are far from achieving IMT-2020 spectral efficiency, traffic and device density, and reliability goals.[43,44]

To our knowledge, pCell is the only technology deployable in this decade, let alone in the next year, that can keep pace with mobile data demand or come close to achieving IMT-2020 5G performance goals. Despite this enormous performance advantage over LTE networks, pCell is actually less expensive and much faster to deploy, and far lower cost to operate than LTE networks, while maintaining compatibility with LTE devices. How pCell not only leapfrogs LTE performance, but also dramatically lowers CAPEX and OPEX is detailed in subsequent sections of this white paper.

The next section explains why current wireless systems have hit upper limits in SE and Density and why a radical approach to wireless such as pCell is necessary to overcome these limitations.

3 The Limits of Cellular Architecture

One of the greatest revolutions in wireless history began on April 1973 when the first cellular phone call was made by the inventor Martin Cooper at the original Motorola Corporation. The first cellular systems began commercialization in the U.S. a decade later in October 1983, ultimately defining the modern era of wireless. In just three decades, cellular has connected half of the people on Earth[45], evolving from exclusively voice traffic thirty years ago to primarily video traffic today[46]. But cellular's success has also exposed its limitations.

Within the scope of this white paper we will discuss three fundamental limitations to cellular:

1. Inconsistent data rates due to rapid decay of RF power with distance
2. Limited SE gain, relying upon multi-paths for spatial multiplexing (micro-diversity limits)
3. Poor cell-edge performance due to inter-cell interference (macro-diversity limits)

3.1 Inconsistent data rates

Cellular networks are planned by design so that power within each cell drops off by the cell edge, and as power drops, SINR[47] and data rate drop. Ericsson[48] shows the practical implications of this effect upon 4G cellular networks in Figure 3: 10 Mbps is available at cell center[49], at mid-cell it drops by 10x to 1 Mbps, and by cell edge it drops another 10x to 0.1 Mbps. Thus, cellular architecture has a 100:1 data rate ratio from cell center to cell edge.

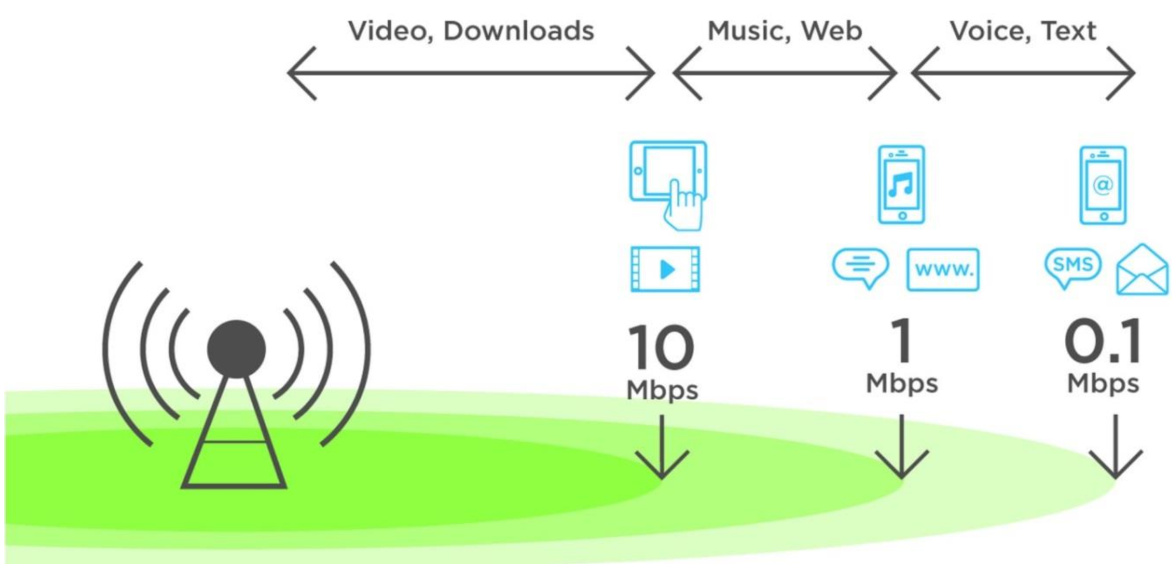

Figure 3: Steep decline of cellular data rates [*Ericsson 2013*]

The large dynamic range of 100:1 results in enormous data rate inconsistency throughout the coverage area. This is visually illustrated in Figure 4 as a 3D signal-to-interference-plus-noise ratio (SINR) heat map showing cellular base stations in an ideal layout[50]. Transmit power attenuates proportionally to the *n*th power of the distance resulting in a "volcano" SINR shape for each cell. The center of each volcano has peak (red) SINR and the edge has minimum (blue) SINR.

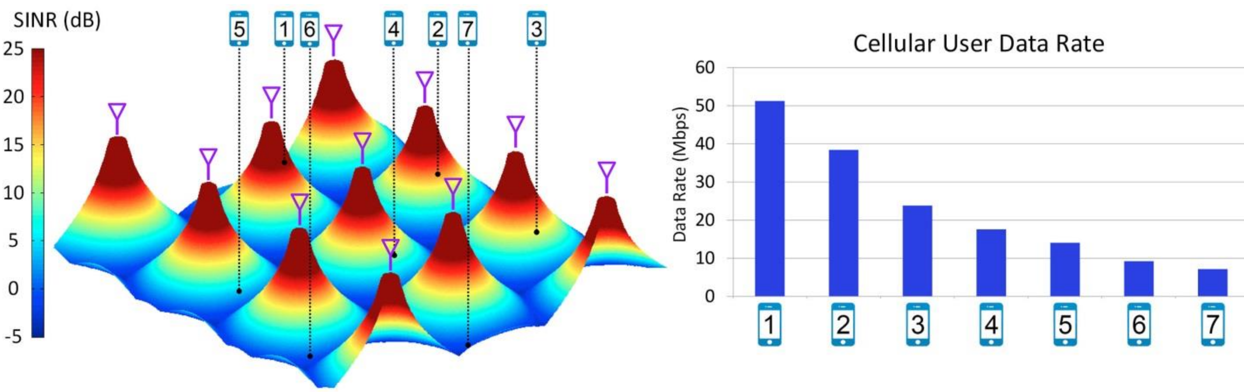

Figure 4: Cellular data rate distribution: Hexagonal base station layout, Sparse users

The user data rate (depicted in a histogram to the right of Figure 4) drops with SINR. Depending on the location, some users experience data rates close to the peak data rate of 56 Mbps[51] (e.g., user 1 at the cell center), whereas others experience only 7 Mbps (e.g., user 7 at the cell edge). This cellular heat map has seven users in a sparse arrangement, with one user per cell in a random location. As can be seen, even with only one user per cell and ideal conditions, cellular data rate is enormously variable.

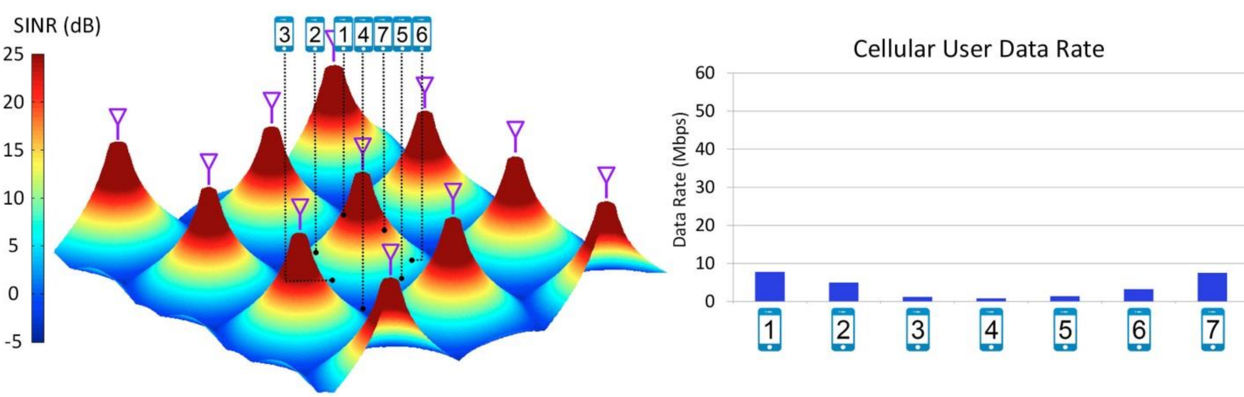

Figure 5: Cellular data rate distribution: Hexagonal base station layout, Clustered users

The same ideal cellular base station layout is depicted in Figure 5, but showing all seven users clustered within one cell. Now, the data rate for each user is not only impacted by how far it is from the cell center, it is also affected by the large number of users sharing the same spectrum and data capacity available in the one cell. The result is extremely low data rates for all users, as shown in the histogram to the right. This is why areas with high densities of users, such as stadiums, airports, tourist attractions, etc. suffer from such poor data rates per user.

An *arbitrary* base station placement with a sparse user arrangement is depicted in Figure 6 to illustrate how poorly cellular performs when base stations are not placed in accordance with an ideal cell plan, for example, to accommodate practical real estate restrictions. Cellular base stations typically limit transmit power to avoid interference with adjacent cells, so inconsistent spacing results in cells with rapid drop-off in SINR/data rates and large dead zones between cells. For example, users 1, 5 and 6 are all in large dead zones. Although some are fairly close to base stations, they have suboptimal performance because of the rapid SINR drop-off; e.g., at the same distance from cell center in the larger cells of Figure 4, user 1 would have higher data rate. Figure 6 clearly illustrates both the sensitivity of cellular architecture to base station placement and how inter-cell interference limits SINR performance.

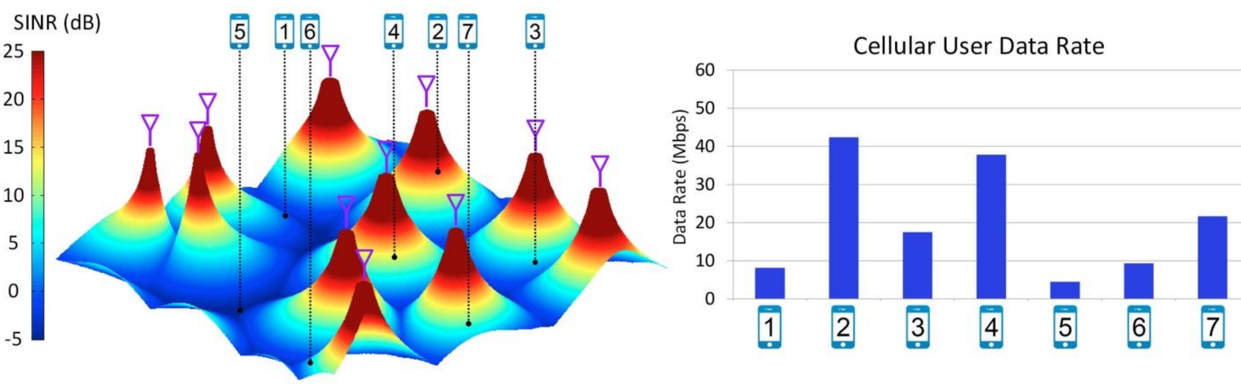

Figure 6: Inefficient Cellular data rate distribution: Arbitrary base station layout, Sparse users

In summary, cellular data rate has an extremely variable 100:1 dynamic range, performs poorly in areas with high user density, and is highly dependent on specific base station placement.

Cellular 100:1 data rate variability is very inefficient for video traffic, which represents more than 50% of U.S. mobile traffic and is growing at 69% CAGR globally[52]. Video requires high and consistent data rates[53] throughout the coverage area, resulting in extremely disproportionate cell capacity consumption toward the cell-edge. For example, 5G specifications, such as IMT-2020 in Table 1, in planning for very heavy video traffic, recognize the need for minimum performance of 100 Mbps throughout the coverage area. With cellular, that mandates a 100 Mbps cell edge requirement, and thus a 100 * 100 Mbps = 10 Gbps cell-center requirement[54] which not only must be supported by the cell base station, but by every user device near cell center. Thus, a 100 Mbps consistency requirement becomes a 10 Gbps performance requirement, an extremely high performance and power burden, particularly for user devices.

3.2 The limits of cellular micro-diversity

Current cellular systems utilize multiple antennas at transmit and receive sides of a communication link to increase SE by exploiting multi-paths in the propagation channel (i.e., space diversity, or in particular micro-diversity[55]) via multiple-antenna spatial processing. The simplest form of spatial processing originated at the beginning of 20th century with initial experiments on phased arrays[56] and later in the 1980s with digital beamforming[57], utilizing multiple transmit antennas to focus wireless energy to combat signal fading and reduce interference. In 1992, pioneering work by Kailath and Paulraj[58] enabled transmission of multiple independent data streams over multiple-input multiple-output (MIMO) links via spatial multiplexing[59], followed in 2001 by the first commercial MIMO-OFDM system by Iospan Wireless, founded by Paulraj.

Multi-antenna techniques are used in cellular systems either to improve coverage (via beamforming or diversity schemes) or to increase SE (via spatial multiplexing schemes)[60,61]. Figure 7[62] shows the average SE achieved through an evolution of increasingly efficient standard protocols and technologies that exploit micro-diversity. Note that the primary approach to improve SE in current and future LTE releases is to increase the MIMO order (i.e., number of antennas) at the base station and user devices.

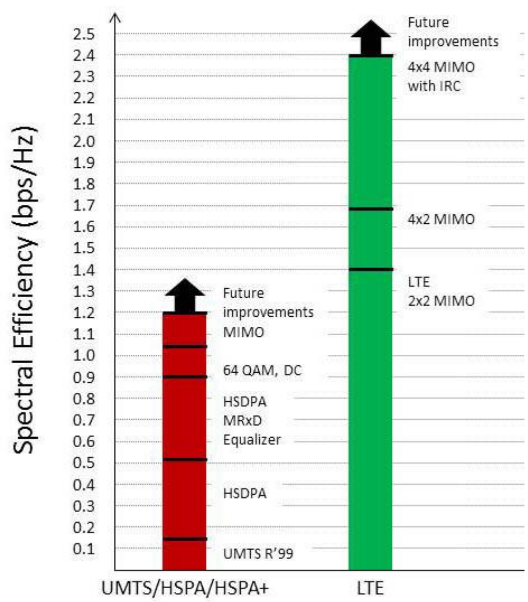

Figure 7: Evolution of average DL SE through cellular standards [Rysavy 2014]

While—in theory—SE scales linearly with MIMO order[63], in practice MIMO multiplexing gain is achieved only in a high SINR or high SNR regime[64] (e.g., close to the base station) and in *space-selective* channels[65]. Channel *space selectivity* refers to statistical variations of wireless signal amplitude at different points in space due to constructive and destructive interference of radio waves as they propagate through multi-path environments[66]. Space selectivity depends on the characteristics of the antenna arrays (e.g., spacing, polarization, radiation pattern, etc.) and the multi-path channel (e.g., number of paths, angles of departure/arrival of the radio waves) [67]. Since the array geometry is constrained by the limited real estate available for commercial base stations and user devices, MIMO performance mostly relies on the multi-paths available in the propagation channel (or "resolvable paths"). In general, for a given MIMO order, the number of resolvable paths defines the channel space selectivity and is proportional to the number of independent data streams that can be sent concurrently over MIMO links (or "multiplexing gain")[68].

In outdoor tests, the 2003 3GPP study in Figure 8 found between 1 and 6 resolvable paths with an average of between 2 and 4 resolvable paths, depending on the environment. These results are consistent with an Ericsson 2011[69] outdoor LTE MIMO 8x8 study showing at most a 4x gain due to MIMO multiplexing, despite having eight antennas that theoretically could yield up to 8x multiplexing gain if eight resolvable paths existed[70].

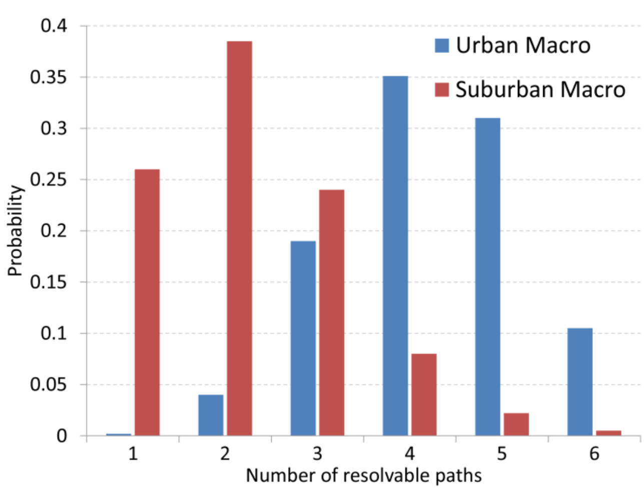

Figure 8: Probability of resolvable propagation paths [*3GPP 2003*][71]

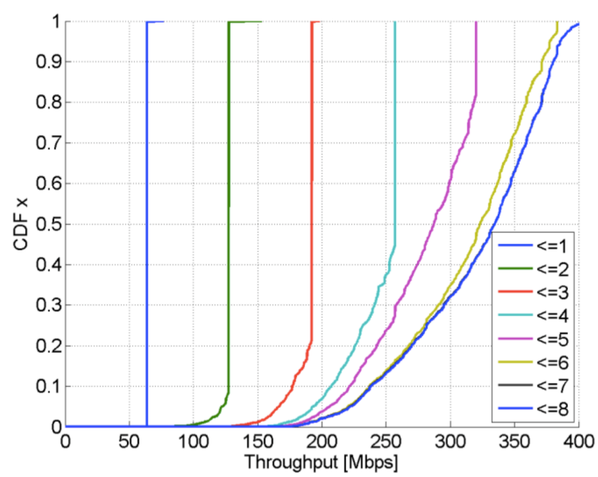

Figure 9: CDF of MIMO data rates [*Ericsson MIMO 2013*][72]

In indoor environments, a 2004 IEEE 802.11n study[73] found between 2 and 6 propagation paths. These results are consistent with an Ericsson 2013 LTE MIMO 8x8 indoor study in Figure 9 which, despite having eight antennas (that theoretically could yield 8x multiplexing gain if eight resolvable paths existed), achieved only an average 4x multiplexing gain with a peak of 6x[74]. Commercial Wi-Fi MIMO 4x4 systems have shown to only achieve up to 2x indoors, despite a theoretical peak of 4x, with performance dramatically degrading with distance from the access point[75]. Therefore, in practical MIMO systems, multiplexing gain does not scale linearly with the number of antennas due to limited number of resolvable paths.

MIMO has other limitations such as: (a) highly variable performance throughout the cell (e.g., cell center vs. edge), so MIMO cannot be relied upon for services requiring sustained high data rates, such as video, which today make up the bulk of mobile traffic[76], (b) MIMO cost grows rapidly with MIMO order, as each antenna requires a closely-spaced, but isolated, RF chain and computational complexity grows dramatically, and (c) performance degrades due to Doppler effects and channel aging from user and environment motion[77]. Despite these limitations, currently MIMO remains the best available solution to increase *average* SE of cellular systems, albeit limited to average multiplexing gains of at most 4x in practical systems.

A number of other approaches to MIMO have been explored to improve space selectivity and achieve higher multiplexing gains. One solution is to separate the receiving antennas as in multi-user MIMO (MU-MIMO) systems. MU-MIMO links are established between a base station antenna array and multiple users equipped with one or multiple antennas. The theory behind MU-MIMO systems was formulated in the seminal work by Caire and Shamai[78] followed by Yu and Cioffi[79,80] and Goldsmith et al.[81,82] based on the idea of “dirty paper coding” (DPC)[83,84,85]. One of the first commercial systems employing MU-MIMO technology was designed by ArrayComm using space division multiple access (SDMA) techniques to form individual focused beams to different users[86]. The MU-MIMO scheme is part of the LTE standard, limited to a maximum of four users with only low-resolution channel state information (CSI) available[87].

Another technique that has emerged in the last three years is “massive MIMO”[88,89]. The basic concept of massive MIMO is to have far more base station antennas than users and exploit the excess antennas to increase space selectivity and create independent spatial channels to multiple concurrent users via beamforming. Massive MIMO is still in early stages of academic research and only recently a limited number of propagation studies have been published to show its performance with two types of array configurations (i.e., linear and cylindrical arrays)[90]. Those measurement campaigns verified experimentally that indeed the limited space selectivity achievable in MIMO channels (due to collocation of antennas within limited real estate) can be compensated by a very large number of excess antennas. But it is yet unclear whether the multiplexing gain achievable through massive MIMO can scale linearly with the number of user antennas to achieve high spectral efficiency required in next generation wireless systems. Other practical limitations are: i) highly complex base stations equipped with many tightly-packed, but isolated RF chains, increasing design costs and power efficiency requirements; ii) degradation due to pilot contamination as the technology is implemented within a cellular framework; iii) degradation due to Doppler effects and channel aging from user and environment motion; and iv) undefined interoperability with existing cellular networks and devices that may delay practical deployments[91,92,93].

Because MIMO performance is inherently unpredictable in practical deployments, MIMO can only be deployed as an "as-available" enhancement to baseline SISO (single-input single-output) performance. Thus, while MIMO increases the peak and average data rate of a wireless network, cellular multi-antennas systems only marginally increase the minimum data rate (through beamforming providing higher SINR), which is still defined by SISO performance (typically at the cell edge). Services reliant on consistent data throughput, such as streaming video—the majority of data traffic today—cannot rely upon MIMO enhancements.

In summary, practical MIMO systems can achieve an average multiplexing gain up to 4x, peaking at 6x, which is not sufficient to meet the target SE of next generation wireless systems. Moreover, MIMO performance is highly variable and unpredictable, arbitrarily determined by the characteristics of objects in the environment, the distance from cell center and user and environment motion. While MIMO increases the peak and average data rate of a wireless network, the minimum data rate is still defined by SISO data rate, which limits MIMO's benefit for services reliant on consistent data rate, such as streaming video.

3.3 The limits of cellular macro-diversity[94]

In cellular networks, as a user device moves away from its serving cell, it is handed off to the adjacent cell's serving base station for uninterrupted service. Virtually all mobile networks utilize this basic cellular network architecture (to the point where the terms "mobile" and "cellular" are used synonymously). While cellular architecture has served the world well for over thirty years, as networks have become denser, cellular is approaching inherent capacity limits.

Cellular systems seek to constrain RF propagation within 2D geometric layouts to minimize interference between adjacent cells. From left-to-right, Figure 10 illustrates the difference between conceptual and real-world depictions of cellular layouts. The leftmost diagram shows an ideal base station placement (blue dots) with the ideal propagation that would occur with free-space path-loss. When the same carrier frequency is used for adjacent cells with universal frequency reuse[95] (as in current LTE systems), inter-cell interference occurs in the overlapping regions. The effect of inter-cell interference becomes more severe in real-world scenarios characterized by irregular cell shapes in the middle diagram, due to shadowing from obstacles in the propagation environment. Finally, on the far right is a diagram that approaches reality, where the base stations are in sub-optimal locations and there are real-world obstacles, resulting in not only irregular cell shapes, but different cell sizes and considerably more severe inter-cell interference effects. Further, the rightmost diagram only shows a snapshot in time; real-world shadowing is variable, producing changing cell shapes and interference patterns[96].

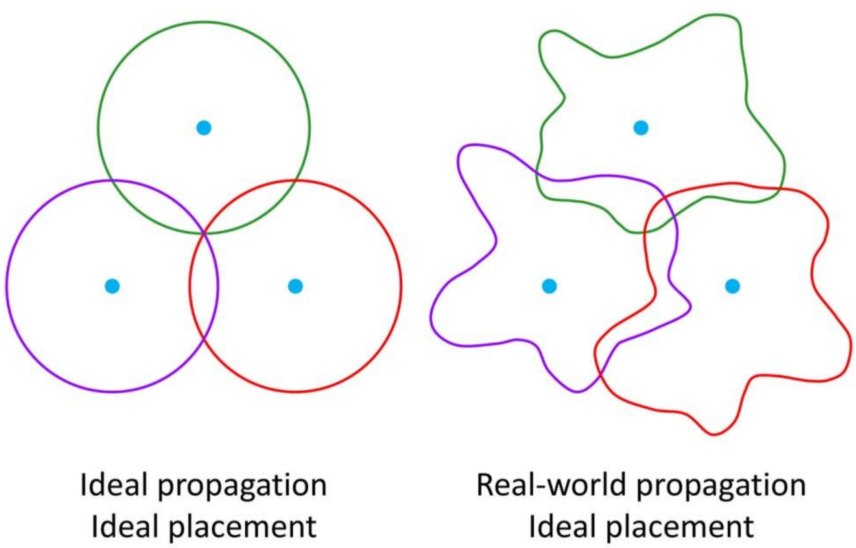

Figure 10: 2D cell geometry: Conceptual to Real-world

3.3.1 Small-cells and inter-cell interference coordination

Inter-cell interference effects exacerbate when cell density increases to support more subscribers and heavier usage per unit area, as wireless operators seek to increase their network capacity. For example, in densely populated areas (e.g., downtown areas, shopping malls, airports and stadiums) current cellular networks have been testing small-cells (e.g. pico-cells and femto-cells) and Heterogeneous Networks ("HetNets", where small-cells are overlaid within the macrocell umbrella)[97]. Since small-cells are far more dense than macrocells, the number of overlapping regions with inter-cell interference depicted in the rightmost diagram in Figure 10 increases substantially[98]. Further, small-cell base stations are typically installed near street level (e.g. lamp posts), which makes propagation highly unpredictable (unlike conventional macrocells) with large signal variations between LOS paths (or urban "street canyons"[99]) and NLOS paths (e.g., outdoor-to-indoor propagation through buildings)[100], yielding cell geometries that are far more irregular than those shown in the rightmost diagram of Figure 10.

In addition to inter-cell interference, these complex propagation patterns cause unnecessary handoffs, which make mobility in small-cell networks highly inefficient. Frequent handoffs may trigger undesired radio link failures (due to dragging effects, when the network waits too long before initiating handoff) or ping-pong effects (when handoff is initiated multiple times unnecessarily between two adjacent cells)[101,102]. Depending on the speed of mobile subscribers, the rate of handoff failure and ping-pong effects can be as high as 60% and 80%, respectively[103], yielding a large amount of control overhead that reduces cell throughput. These issues are exacerbated by the highly-degraded small-cell edge performance caused by extensive inter-cell interference[104,105,106].

The LTE standard attempts to solve these issues through self-organizing networks (SON), consisting of one centralized unit self-optimizing the configuration of adjacent cells to enable load balancing, reduce handoff overhead and mitigate inter-cell interference. SON uses inter-cell interference coordination (ICIC) as one method to mitigate inter-cell interference[107,108] through "cell-autonomous" schemes employing different frequency reuse patterns (e.g., full frequency reuse, hard frequency reuse, fractional frequency reuse) or coordinated techniques enabling cooperation between base stations to coordinate the allocation of time/frequency resources. ICIC does not remove interference; it only *avoids* interference by frequency coordination, and consequently throughput gains are only limited. Field trials with SON have demonstrated only marginal average throughput improvements on the order of 10%[109].

In summary, as wireless operators have turned to small-cells and HetNets to increase subscriber density, increased hand-off overhead and inter-cell interference quickly has reduced the capacity gains from higher cell density, ultimately hitting a limit to the capacity achievable in a given area, despite the use of SON and ICIC techniques. The impracticality of small-cells is reflected in weak market acceptance: despite widespread promotion of small-cells as a solution to mitigate congestion, small cell adoption by mobile operators has been very limited, falling far short of projections[110].

3.3.2 Network MIMO and CoMP

An alternative approach to mitigate inter-cell interference is to coordinate transmissions from multiple base stations in so-called "network MIMO", "distributed MIMO" or "CoMP" systems. Network MIMO was first proposed in 2005 with the goal of improving SINR in cellular systems to increase SE via MIMO spatial multiplexing[111] (which only works in a high SINR regime, as described in the Section 3.2). Theoretical analysis of network MIMO in ideal complex Gaussian channels[112] and simulations in indoor channel models[113] showed significant gains in SE over conventional cellular systems, although these gains do not grow linearly with the number of base stations in the network and are limited[114]. While theoretical analysis of network MIMO shows up to 5x gain over conventional systems using power control to mitigate interference[115], recent lab prototypes have demonstrated only up to a 4x gain in SE[116], comparable to conventional MIMO. Base station coordination has become part of the LTE-Advanced standard for 4G cellular systems with so-called coordinated multi-point (CoMP)[117] schemes, namely "joint processing" (JP) and "coordinated scheduling/beamforming" (CS/CB)[118]. CoMP only enables cooperation among adjacent cells and its performance degrades when LTE-compliant limited feedback mechanisms are employed, due to coarsely quantized CSI[119,120]. Further, CoMP is highly sensitive to Doppler effects and channel aging from motion by the user and the environment[121]. Practical field tests have demonstrated at most 30% gain in average downlink throughput with CoMP[122,123,124,125].

3.4 Other limitations of cellular systems

The following subsections describe other practical limits encountered in the deployment of current and next generation cellular systems.

3.4.1 3D environments

Although we've thus far illustrated 2D horizontal cellular layouts, the real world is 3D with a vertical dimension, resulting in additional challenges for cell planning, particularly in urban environments, such as the example depicted in Figure 11. An indoor user in a high-rise building might well experience LOS from many cell base stations, with no base station having much different power than any other. Beyond the challenge of deciding which cell should serve the user device, the combined downlink transmissions from interfering cells results in low SINR to the user, while the user uplink transmission interferes with whichever cells are not chosen to be the serving cell[126]. Given that 80% of mobile Internet use is indoors[127] and 53% of people live in urban areas (growing about 0.5% every year)[128], challenges from mobile usage in tall structures will have an increasing impact on global mobile capacity and, accordingly, it has been considered as a case study by 3GPP LTE Rel. 13[129].

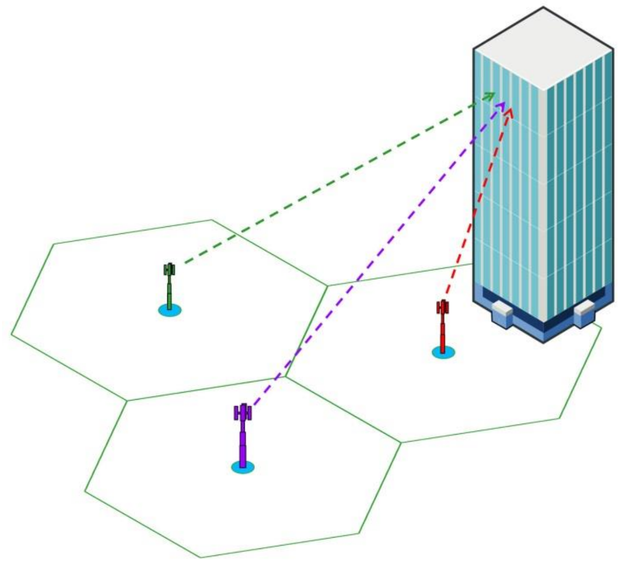

Figure 11: 3D high-rise cellular challenges

In general, cellular provides only approximate user location information, particularly in dense urban areas with large 3D structures, where accuracy errors can exceed 100 meters[130]. Accurate 3D location information (e.g. identifying the correct floor and room in a building) is critical for first responders in emegency situations and also can be useful for personal and commercial applications.

3.4.2 Crisis situations

As payphones and landlines have rapidly vanished from public locations, homes and businesses[131], the world has become almost entirely reliant upon mobile communications in crisis situations. Because cellular architecture inherently divides the coverage area into largely non-overlapping cells, if a single cellular base station fails for any reason (e.g. loss of power or backhaul, or if it is damaged or destroyed), the entire cell loses coverage. Power loss can be mitigated by batteries or generators, but if an incident results in physical harm to a base station or its connectivity, the base station's coverage area will lose connectivity. Where operators share mobile infrastructure[132], damage to a shared cell tower or its backhaul can result in the loss of all mobile service in the surrounding area, as shown in Figure 12.

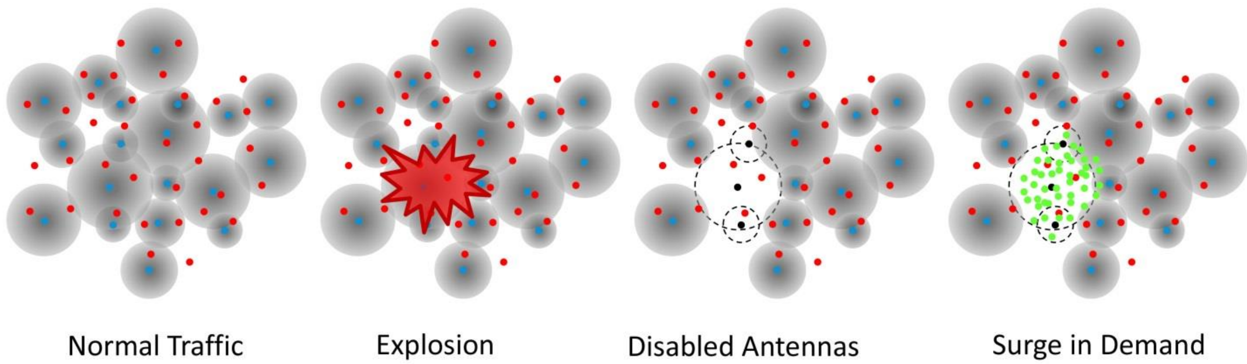

Figure 12: Cellular coverage loss in a crisis situation

A crisis situation at a given location can result in a vast number of people in the affected area attempting to initiate calls, and once the news of the incident is widely known, a vast number of people attempting to call into the affected area. First responders urgently need communications to and from the affected area, potentially requiring video teleconferencing so remotely-located specialists can make visual assessments of the situation. Even if the crisis situation does not damage the cellular infrastructure, a given cellular base station capacity is sized for typical peak usage which assumes a small fraction of people within a cell are concurrently engaged in communications. Although there have been proposed SON technologies to theoretically balance loads between adjacent small cells, a sudden surge of a large percentage of people in the area concurrently attempting connections will overwhelm *all* base stations serving the area, regardless of how loads are balanced, potentially resulting in little or no service for anyone. For example, in the wake of the Boston Marathon bombing in 2013, the sudden surge in mobile traffic in the area overwhelmed all major U.S. carriers resulting in no service at all[133].

As previously noted, cellular networks are only able to provide very coarse location information, so even when base stations are not overwhelmed, they often are unable to provide useful location information to find people who need help, particularly if they are located in tall buildings. In contrast, a landline call can typically pinpoint the location of the caller at the installation address of the landline, whether an apartment in a tall building in a dense city or a farmhouse in a rural area. Cellular technology is unable to come anywhere close to this level of precision, particularly in densely populated areas.

3.4.3 Security vulnerabilities

Transmissions by conventional wireless systems, including cellular, can be intercepted by an eavesdropping device within range of the transmitter. Current cellular systems utilize key-based cryptography to encrypt data transmissions. Even with the strongest encryption system, if the key is compromised, it is possible to intercept and decrypt data transmissions. For example, it was recently reported the encryption keys of potentially most SIM cards in use today have been compromised[134]. Because cellular transmissions can be easily intercepted, this means that possibly most cellular communications in the world have been compromised, vividly illustrating the inherent security vulnerability of cellular communications.

3.4.4 Millimeter wave performance in cellular networks

Millimeter waves are defined to be at frequencies much higher than those of conventional cellular systems (e.g., between 30 GHz and 300 GHz of the electromagnetic spectrum), and have been proposed for future 5G cellular networks. The principle benefit of millimeter waves is the vast amount of unused millimeter wave spectrum available, but the fundamental drawback is millimeter waves incur high attenuation from most natural objects (including heavy rain and fog) and man-made objects other than clear glass[135], behaving much like light waves. Experiments with high-gain directional antenna arrays (analogous to a sweeping spotlight) have been used to sweep through coverage areas to demonstrate that many non-obstructed areas can be reached by millimeter waves[136]. But millimeter waves, like light, are unable to reach locations only reachable by penetration of non-transparent objects, requiring extensive antenna placement throughout the coverage area[137] (analogous to installing enough spotlights to illuminate every outdoor and indoor location in the coverage area on a dark night).

Millimeter wave research thus far has largely been focused on propagation studies and to date, little work has been done to assess the unique propagation characteristics of millimeter waves within a cellular architecture. While millimeter waves, even as a focused beam, are largely blocked by walls, they propagate through air and clear glass very efficiently, resulting in highly irregular cell shapes. For example, a millimeter wave small-cell within a room might not penetrate the walls of a room, but could extend very far through the open windows of the room or through an open door, interfering with cells serving the areas outside of the room. If shades are drawn or the door is closed, the millimeter waves would be blocked, dramatically and dynamically changing the shape of the cell.

It is unknown how long it will be before millimeter wave systems can be deployed on a commercial basis or what form such networks and user devices will take. Whether cellular architecture can be used efficiently with millimeter waves remains an open area of research.

3.5 Summary: Cellular architecture has reached capacity and reliability limits

While cellular has served the world well for over 30 years, cellular has reached its limits, given today's requirements for wireless networks in terms of capacity, reliability, consistency, density and public safety.

We summarize the fundamental limits of cellular architecture as follows:

- Inconsistent data rate throughout the coverage area (100:1 cell center to cell edge)
- Poor cell-edge performance further degraded by inter-cell interference
- Poor performance in high-density user scenarios due to bandwidth sharing
- Inability to scale capacity through small-cells due to uncontrolled inter-cell interference
- Large handoff overhead, particularly in small-cell deployments
- Limited and inconsistent capacity gains through MIMO (only up to 4x over SISO)
- MIMO limited to 4x average SE gain in ideal scenarios, no gain in unfavorable scenarios
- Spatial processing (MIMO, CoMP, beam-forming) highly sensitive to Doppler/mobility
- Inflexible antenna installation, requiring specific and expensive base station locations
- Poor vertical 3D performance in high-rise buildings
- Inability to accurately determine user location
- Crisis situations result in loss of coverage and/or severe congestion
- Cellular transmissions are vulnerable to interception and decryption

Despite these inherent limitations, cellular is still taken as a given as the architecture for mobile communications, and virtually all mobile wireless development is constrained within a cellular framework and its limitations, resulting in marginal improvements at best.

The only way to overcome the above limitations—and all can be overcome—is to transition from the traditional cellular architecture to a new architecture for wireless, designed from the outset to serve today's requirements for wireless networks, thus delivering scalable, consistent capacity in even high-density scenarios, and supporting public safety needs. And, on a practical level, ideally this new wireless architecture would remain compatible with standard LTE and Wi-Fi devices and use readily-available backhaul infrastructure so that it can be immediately deployed and put into use.

This new wireless architecture is pCell.

Figure 13: pCell technology can concurrently provide every user in this photo the full capacity of the same spectrum, including users in tall buildings and moving vehicles

4 pCell Technology

pCell is a new approach to wireless that—rather than avoiding interference like conventional wireless systems—exploits interference. pCell technology synthesizes a tiny personal cell (a "pCell") for each user device in the coverage area. Since each pCell is an independent radio link, each device in the coverage area is able to concurrently utilize the full capacity of the same spectrum, dramatically increasing wireless capacity.

pCell systems inherently have exceptional space selectivity, synthesizing pCells that are physically very small (a fraction of a wavelength in diameter at practical mobile frequencies[138]) with three-dimensional polarization[139], thereby maintaining independent data links with user antenna spacing on the order of millimeters at practical mobile frequencies[140]. As a result, pCell achieves orders of magnitude higher spectral efficiency than cellular, even at very high device density.

Although pCell technology is radically different than cellular or other interference-avoidance technologies, pCell is compatible with off-the-shelf LTE and Wi-Fi devices. pCell accomplishes this by synthesizing LTE or Wi-Fi protocols within each device's pCell. From the device's point of view, the received waveform is the same as a conventional LTE eNodeB or Wi-Fi waveform. In fact, because pCell synthesizes an SINR peak for each device antenna, a pCell-synthesized waveform has the same quality as a *cell center* waveform in cellular systems regardless of where the device is located in the pCell coverage area. pCell has no cell edges and no hand-offs. So, device throughput remains consistently near peak without interruption throughout the coverage area, a result that cannot be achieved with cellular or Wi-Fi.

As shown in Figure 14, pCell is deployed in a C-RAN [141] architecture using IP fronthaul [142] to either single-antenna pWave™ radio heads (the gold-colored devices in Figure 14) or to Artemis Hubs for distributed antennas (described in Section 5.5.2). All of the processing, down to baseband physical layer, is implemented in proprietary real-time Software-Defined Radio (SDR) within the pCell Data Center (in the middle of Figure 14).

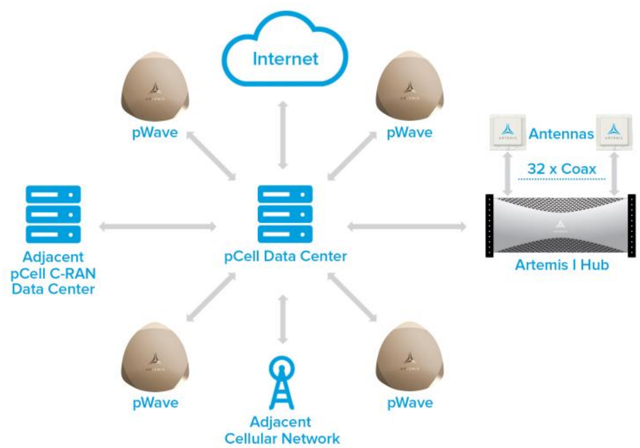

Figure 14: pCell C-RAN Architecture

While cellular deployments follow a precise cell plan to avoid interference between base station transmissions, pCell deployments are instead the opposite: pWave radio transmissions are intended to interfere with each other, so pWaves can be placed in any location that is inexpensive and convenient. Thus, pWave placement is typically determined by where pWaves are least expensive to deploy (e.g. where there is inexpensive fronthaul, such as locations with a Line-of-Sight (LOS) route back to the data center, and inexpensive rent). Consequently, pCell infrastructure is far less expensive to deploy and operate than cellular[143].

To understand pCell technology, it is helpful to start by comparing pCell to cellular architecture. There are many dimensions that need to be considered, including coverage, SE, data rate, latency, mobility, scalability, location accuracy, power requirements, deployment cost, operating cost, crisis resilience, position location and security. We will discuss each in turn.

4.1 Uniform coverage

Figure 15[144] shows cellular vs. pCell RF power distribution[145]. In both cases we show real-world placement of blue dots for base stations (in the case of pCell, pWave remote radio heads or Artemis Hub antennas, collectively "pCell antennas"), but for the sake of illustration, we show free-space propagation resulting in circles around each base station. The red dots are users.

As can be seen in Figure 15, cellular layout is carefully planned to *avoid* interference, choosing locations and transmit power to minimize transmission overlap. But overlap and dead zones are impossible to avoid, especially as cells get smaller and shapes become increasingly irregular and variant. Because cellular is inherently a 2D architecture there are even further challenges with 3D propagation (e.g. tall buildings).

In contrast, pCell *exploits* interference and, as a result, requires no planning, in either 2D or 3D. pCell antennas are placed in arbitrary locations wherever convenient, transmitting at arbitrary power levels with arbitrary transmission overlap. For the sake of illustration the transmissions are shown as free-space circles, but since the interference overlap is arbitrary, the transmission shapes are arbitrary 3D shapes that are constantly changing.

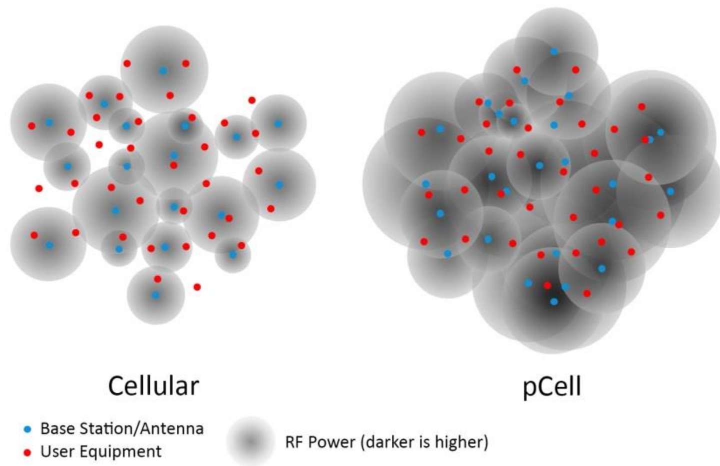

Figure 15: Cellular vs. pCell RF power distribution

Bear in mind that Figure 15 only shows *distribution of RF power*, not *SINR*. Figure 16 shows the exact same base station and user layout as Figure 15, but instead of the gray shading indicating *RF power distribution* as in Figure 15, the gray shading in Figure 16 indicates *SINR*, where there is a receivable signal and determining user coverage.

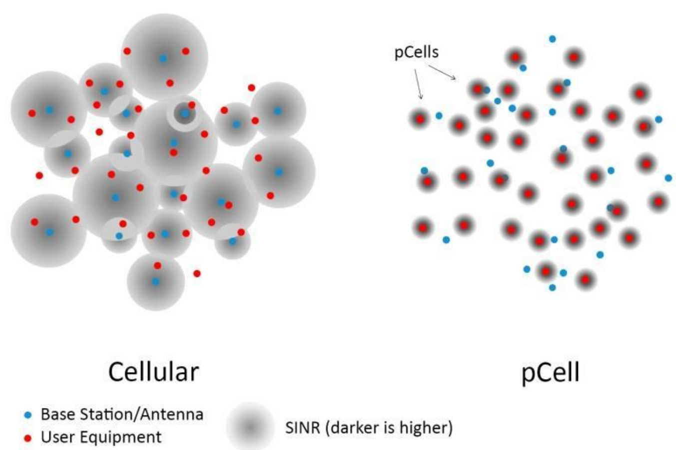

Figure 16: Cellular vs. pCell SINR

With cellular, RF power and SINR distributions are nearly identical, except where there is inter-cell interference (e.g. where one cell overlaps another).

In contrast, pCell *RF power* shown in Figure 15 and pCell *SINR* shown in Figure 16 are *drastically* different. At the precise location of each user antenna (red dot), all of the interfering signals add up to the exact waveform intended for that user (e.g. an eNodeB waveform for LTE) with high SINR. Each of these high SINR waveforms only exists in a very small volume around each user device antenna—what we call a “pCell”. Each pCell is shown as a small gray circle around each red dot. All of the pCells can be synthesized at once, and each is an independent radio link, enabling every user to concurrently experience the full capacity of the channel.

4.2 Consistent data rate

As discussed in Chapter 3, cellular users experience *widely varying* performance, due to a wide range of factors, including their distance from the base station, pathloss, shadowing, multi-paths, inter-cell interference, MIMO schemes, user density, Doppler effects, hand-off overhead, 3D propagation, etc. Average performance within an LTE cell is less than one-fourth of device peak performance[146], and cell edge performance is 1/100th that of cell center performance[147].

In contrast, pCell users experience *highly consistent, near peak* throughput regardless of their location in the coverage area. There are no cell edges and no hand-offs. Users experience a high-SINR pCell consistently throughout the coverage area.

We will first consider how pCell data rate is consistent, regardless of location among pWaves in the coverage area.

The top section of Figure 17 shows the same SINR heat map and blue data rate histogram as Figure 4 in Chapter 3: an ideal hexagonal cellular base station layout with sparse users (one per cell) in ideal channels including only pathloss without shadowing (for the sake of illustration). The bottom section of Figure 17 shows pCell’s performance using the same antenna layout and user positions, but in the pCell case a 3GPP pathloss and shadowing model[148] is used that is far less favorable than the cellular pathloss-only model, but reflects real-world conditions.

The cellular SINR heat maps show SINR peaks centered on the base station antenna locations with a “volcano-shaped” decay, resulting in the highly variable user data rate, depending on how far the user is from the cell-center, as shown in the blue histogram. In contrast, a pCell is synthesized at the precise location of each user device, creating a very narrow peak of SINR for each user. The result is a near-peak data rate for each user throughout the coverage area.

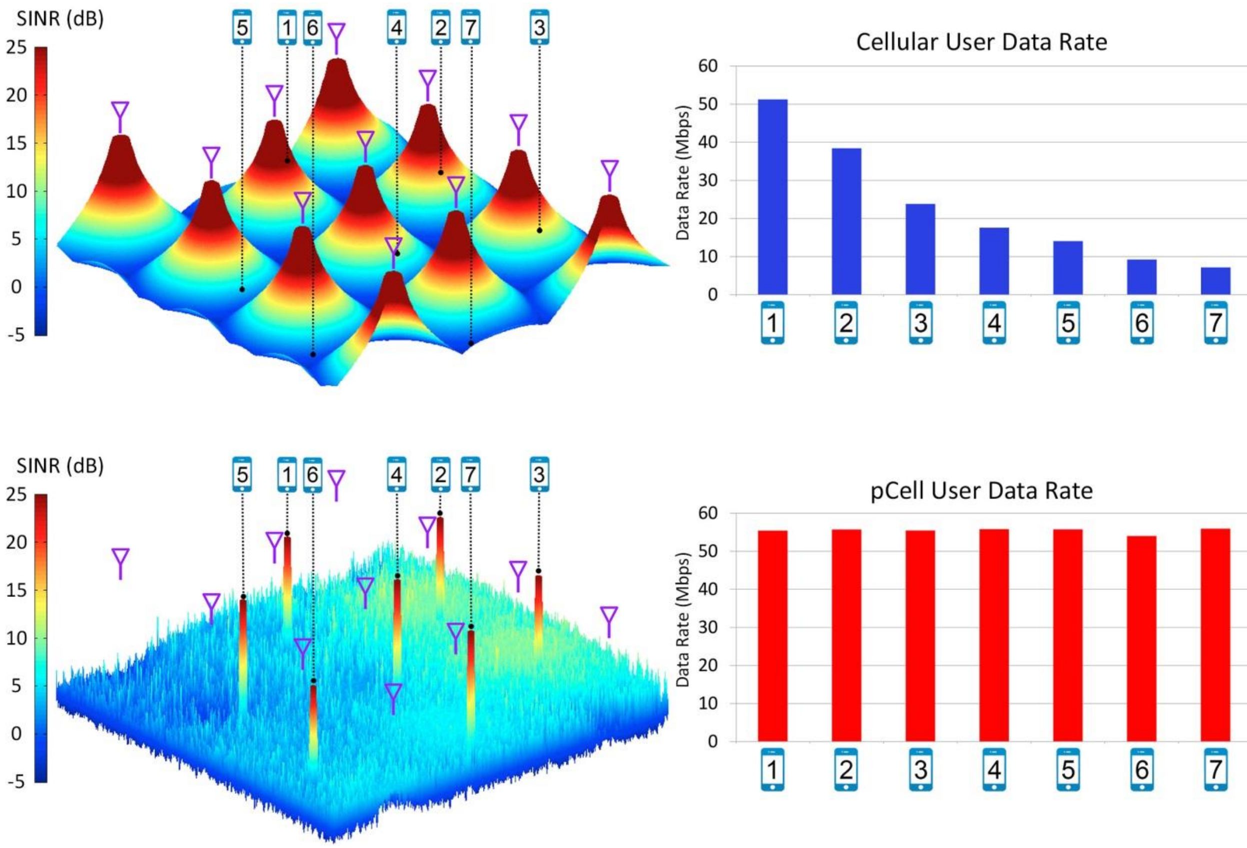

Figure 17: Cellular vs. pCell data rate distribution: Sparse users

The hexagonal antenna layout and channel assumptions of Figure 18 are the same as Figure 17 but in this case all seven users are clustered in one cell. Not only do the cellular users have variable data rate due to their distance from cell-center, but they are all sharing the same spectrum capacity of a single cell. The result is very poor data rate for all cellular users, as shown in the blue histogram.

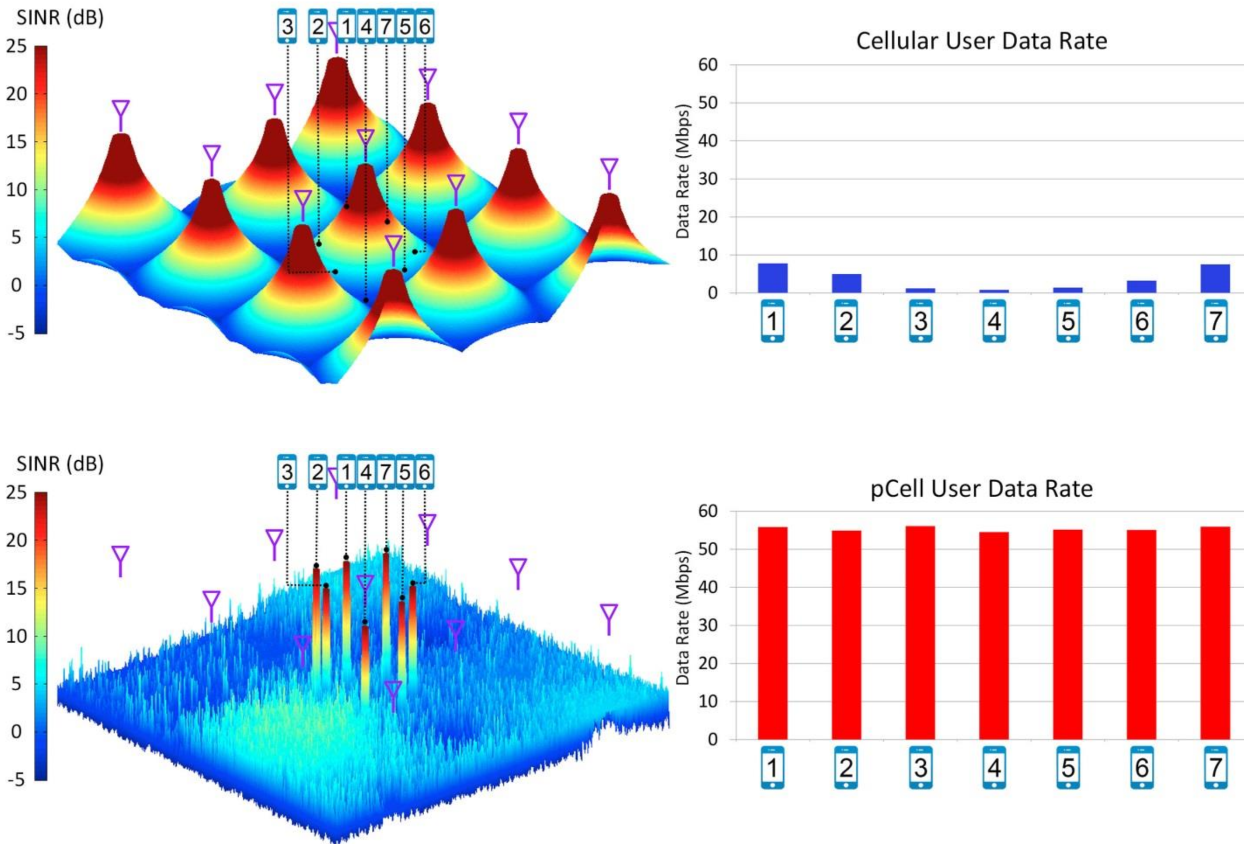

Figure 18: Cellular vs. pCell data rate distribution: Clustered users

In contrast, despite the fact that users are clustered, a pCell is synthesized at the precise location of each user device, creating a very narrow peak of SINR for each user. The result is a high data rate for each user[149].

The top section of Figure 19 shows the same SINR heat map and blue data rate histogram as Figure 6, which is an arbitrary cellular antenna layout with arbitrary user placement. The bottom section of Figure 19 shows pCell's performance in the same configuration. The assumptions on the propagation channels are the same as Figure 17 for both cellular and pCell.

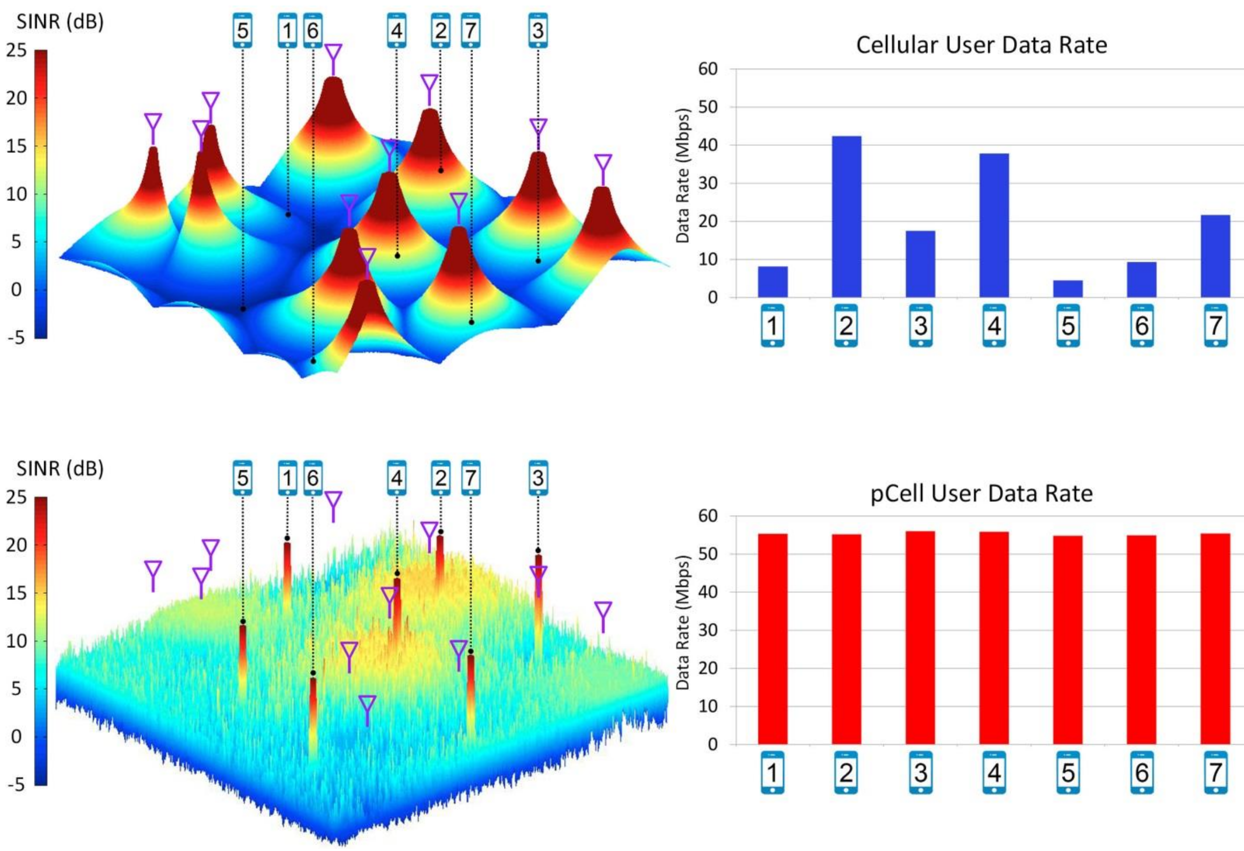

Figure 19: Cellular vs. pCell data rate distribution: Arbitrary antenna placement

Cellular performs poorly because the antennas are close together, resulting in very steep SINR decay and leaving wide dead zones. Users 1, 5 and 6 are in dead zones with very little throughput, and even users that are close to antennas, such as user 1, have suboptimal performance because of the rapid SINR drop-off of small-cells; e.g., at the same distance from cell center in the larger cells of Figure 17, user 1 would have higher data rate. In contrast, a pCell is synthesized at the precise location of each user device, creating a very narrow peak of SINR for each user, despite the fact the antennas and users are in arbitrary locations. The result is a high data rate for each user.

4.3 Indoor field trials

The consistent data rate to multiple concurrent users shown in the pCell simulation results of Figure 17, Figure 18 and Figure 19 conforms to real-world field trials. pCell has been tested both indoor and outdoor, with detailed indoor testing completed thus far, the results of which are presented as follows.

Although pCell supports arbitrary antenna placement, for the purpose of the SE testing 32 pCell antennas were placed in a regular grid with roughly 2.5 meter spacing, with aligned polarization[150], at a uniform height (except in low-ceiling corridors) and pointing downward, resulting in a mix of LOS and NLOS paths, some through walls and some through free space. A room in the pCell SE test environment is shown in Figure 20. The antennas (2”x2” patch antennas, 8dBi, HPBW=75°) are circled in red. Every antenna transmits one LTE spectral mask-compliant waveform at 1 mW average power, with signal bandwidth of 5 MHz and carrier frequency of 1917.5 MHz. The cart carrying the 1 m^2 plexiglass sheet with the user devices being tested is circled in blue in the back of the room. The Artemis I hub used to drive the antennas is circled in blue in the middle of the room. It is connected by 10 Gig Ethernet fiber to three servers in a nearby data center which compute the baseband waveforms in real time using Artemis’s pCell software-defined radio (SDR). The Artemis I Hub and Artemis pCell SDR are described in Section 5.5.2 and Section 5.3, respectively.

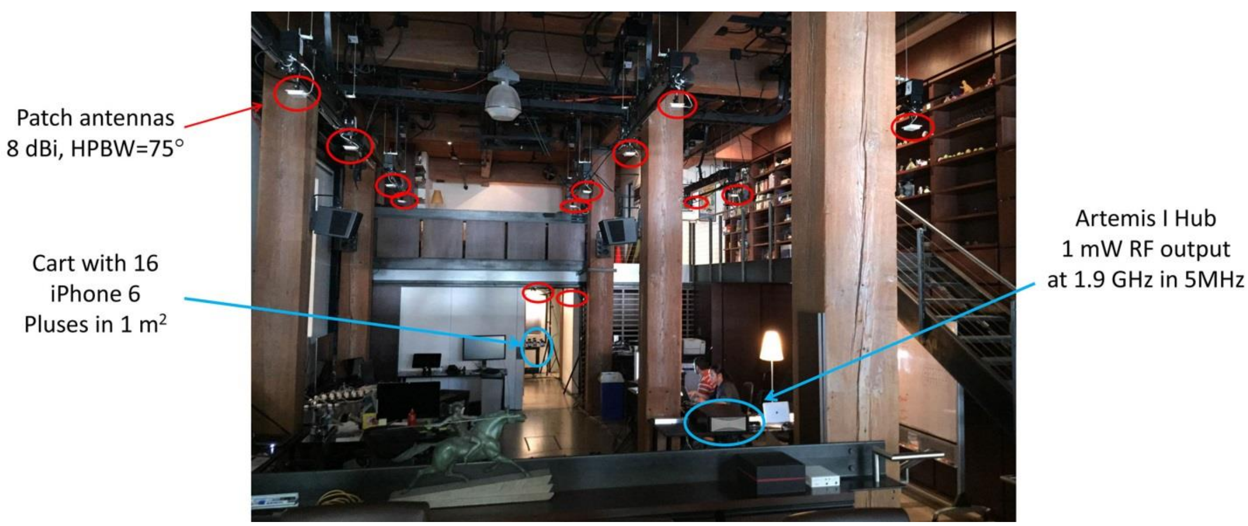

Figure 20: A room in the pCell SE test environment

The pCell LTE waveforms are compliant to the TDD frame configuration #2 (i.e., DL:UL ratio 3:1) with S-subframe configuration #7. pCell SE was determined by measuring the aggregate SE of a group of user devices (iPhone 6 Pluses were used for these measurements) all concurrently receiving data from the pCell antennas in the coverage area. The iPhone 6 Plus devices were placed in a uniform pattern on a 1 m^2 plexiglass table as shown in Figure 21[151]. The 1 m^2 plexiglass sheet was moved throughout the coverage area in 75 cm increments.

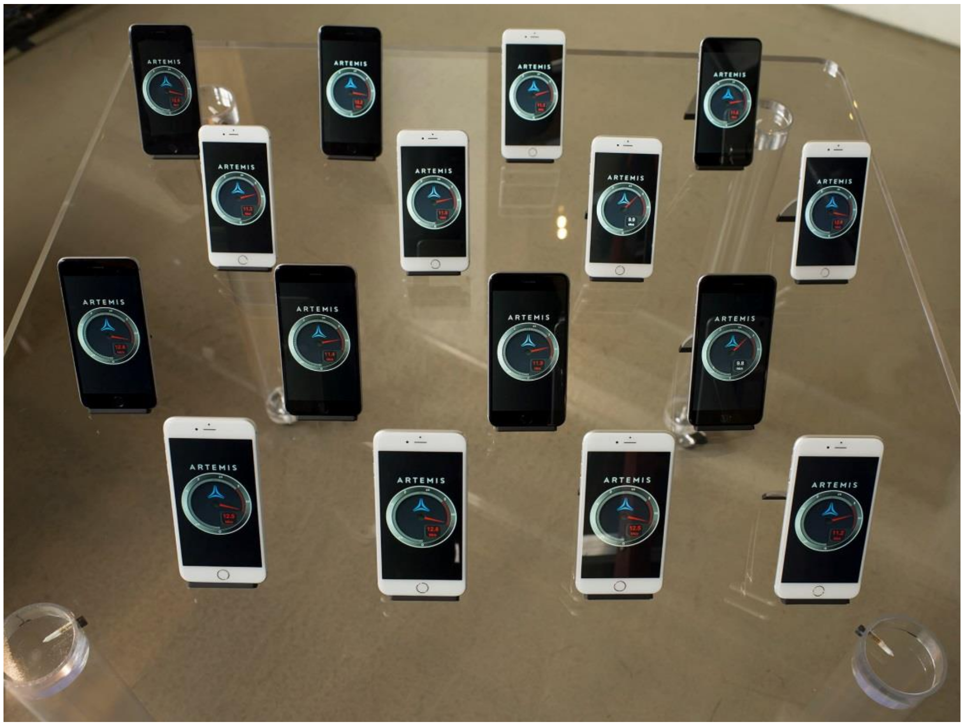

Figure 21: 1 m^2 plexiglass table with 16 iPhone 6 Plus devices

The aggregate LTE MAC layer downlink (DL) throughput to a given number of user devices within the area of the 1 m^2 plexiglass table was used to calculate the aggregate DL SE of each location. A heat map for the aggregate SE for the 16 iPhone 6 Pluses in Figure 21 throughout the coverage area is shown in Figure 22, along with the layout of 32 pCell antennas (white squares with the blue Artemis logo). The average SE across all locations was 59.3 bps/Hz, the peak was 59.8 bps/Hz and the 5% outage[152] SE was 58.1 bps/Hz mostly due to locations in the upper right corner (obstructed by several walls).

The average SE of 59.3 bps/Hz is 35x cellular's average SE of 1.7 bps/Hz with 2-antenna devices, such as the iPhone 6 Plus used in our testing. But direct comparison between pCell SE and cellular SE requires careful examination since pCell and cellular architecture are radically different. This is discussed in the following section.

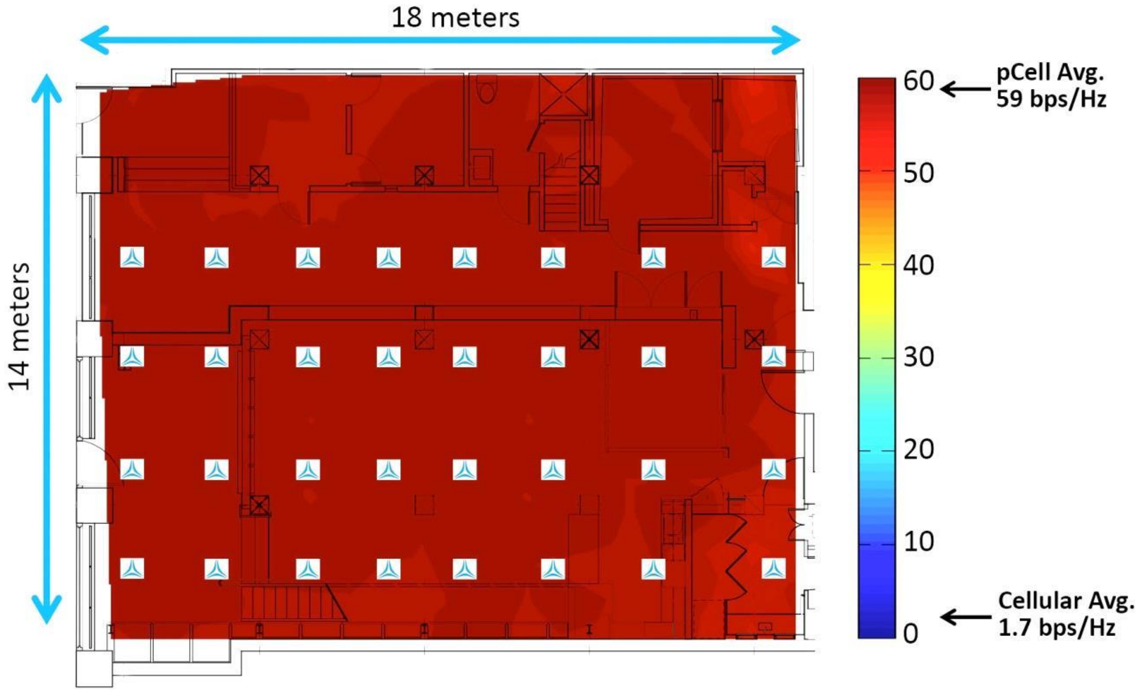

Figure 22: Aggregate DL SE heat map of coverage area

4.3.1 pCell and cellular performance compared

pCell is a user-centric technology that maintains consistent, near-peak SE for each user in the coverage area, regardless of the user's location, whether clustered or sparsely distributed relative to other users.

Cellular is a base-station technology with 100:1 cell-center to cell-edge SE variability within each cell depending upon user location, and further divides up the SE of each cell among all users in the cell.

Because the two architectures are very different, to compare the two systems requires the assessment of a range of scenarios.

To start with, consider pCell and cellular in the case of small numbers of network antennas.

4.3.1.1 Performance with 2 or 4 network antennas

4GAmericas.org contributors[153] recently reported their consensus view of the average SE of deployed cellular networks configured as LTE 5+5 MHz FDD LTE downlink (DL)[154] with either 2 or 4 network antennas per cell serving 2-antenna LTE devices[155] using MIMO 2x2 and 4x2, respectively.

pCell coverage areas can be served by far more than 2 or 4 antennas, but we have conducted extensive indoor surveys with just 2 or 4 pCell antennas serving 2-antenna LTE devices.

The SE results are shown in Table 2 below[156]:

Network Antennas	Cellular LTE DL			pCell LTE DL		pCell vs. Cellular SE
	Avg. SE bps/Hz	5+5 MHz FDD Mbps	5 MHz TDD Mbps	Avg. SE bps/Hz	5 MHz TDD Mbps	
2	1.4	6.3	4.7	7.5	25	5x
4	1.7	7.7	5.7	15	50	9x

Table 2: 5 MHz Cellular LTE vs. 5 MHz pCell LTE Downlink

The pCell SE shows actual indoor measurements using off-the-shelf iPhone 6 Plus LTE devices in 5 MHz TDD LTE DL. At each location in the coverage area, 2 or 4 network antennas served 2 or 4 iPhone 6 Plus LTE devices, respectively, within 1 m^2.

As can be seen in Table 2, 2 and 4 pCell antennas deliver 5x and 9x higher average SE than 2 or 4 cellular antennas, respectively. In fact, 2 pCell antennas deliver over 7.5/1.7=4.4x higher average SE than 4 cellular antennas.

Of course, the pCell antennas are each at different locations, while the cellular antennas are all at the same location in a MIMO array. But, even if we compare multiple cellular base stations placed in as many locations as the pCell antennas, pCell still far outperforms cellular, even assuming optimal conditions.

If we divide the cellular coverage area optimally into 4 cells, each with a 4-antenna base station, and assume the users are optimally distributed evenly among the 4 cells. It's still the case that the aggregate cellular SE of the 4 cells is 4 * 1.7 bps/Hz = 6.8 bps/Hz served by 4-antenna base stations at *4* locations is less than the pCell SE of 7.5 bps/Hz SE served by a single pCell antenna at each of *2* locations. Further, the 4-location SE of 6.8 bps/Hz is *less than half* of the pCell SE of 15 bps/Hz of a single pCell antenna at each of 4 locations.

And, achieving such results with cellular is far from trivial: it requires uniform spacing of base stations at specific locations, requires users to be evenly distributed among the 4 cells to achieve the aggregate SE result, and suffers from inter-cell interference as cells get smaller. In contrast, the pCell antennas can be placed arbitrarily throughout the coverage area, and the SE will be achieved regardless of whether the users are clustered or sparsely distributed through the coverage area.

The remaining difference drawn between LTE's 4-antenna SE of 1.7 bps/Hz and pCell's 4-antenna 15 bps/Hz is that the cellular numbers are from full mobile deployments that are a mix of indoor and outdoor, while pCell is from indoor-only measurements. Although 80% of mobile traffic is indoors[157], it is still the case that while pCell has been tested outdoors, we do not yet have extensive outdoor surveys. That said, the peak SE achievable using cellular with 2-antenna LTE devices is 7.6 bps/Hz[158] under any conditions, so pCell's 4-antenna 15 bps/Hz is almost double cellular's peak SE, even if it was only measured so far in the indoor conditions that constitute 80% of cellular traffic.

Thus, in any scenario with at least 2 pCell antennas, pCell delivers higher average SE than cellular, even if cellular base stations are located in as many locations as pCell antennas.

4.3.1.2 Performance with increasing numbers of user devices

While pCell scales to very large numbers of antennas and concurrent devices, LTE does not.

Although LTE-Advanced supports up to 8 antennas per user device, as shown in Figure 9, above, even in rich multi-path conditions, there is no gain beyond 6 antennas and the average gain is only 4x. And, as noted above, we are unaware of any high-volume 4-antenna LTE devices that utilize the MIMO 4x4 capability of standard LTE, let alone the MIMO 8x8 capability of LTE-Advanced (which on average would achieve at best a 4x gain).

So, the highest practical LTE SE used as a point of comparison is MIMO 4x2 serving 2-antenna user devices, which is 1.7 bps/Hz. We will use 1.7 bps/Hz as our reference for cellular average SE with 2-antenna devices to compare against pCell average SE with 2-antenna devices.

Table 3 shows an increasing number of 2-antenna user devices (iPhone 6 Pluses) which, in each case, were clustered on a 1 m^2 plexiglass table and tested throughout the coverage area as described above in Section 4.3 and shown in Figure 20, Figure 21 and Figure 22.

User Devices	pCell LTE DL		pCell SE vs. Cellular SE
	Avg. SE bps/Hz	5 MHz TDD Mbps	
2	7.5	25	5x
4	15	50	9x
8	30	100	18x
12	45	150	26x
16	59	198	35x

Table 3: pCell average indoor DL SE with user devices clustered in 1 m^2

Despite the increasing number of user devices all concurrently using the same spectrum, and the increasing density of user devices within 1 m^2, average pCell network SE grows almost linearly with number of devices throughout the coverage area, and device SE is highly consistent at peak or near-peak SE for each device, achieving as high as an average of 59.3 bps/Hz with 16 devices, 35x higher average SE than the cellular average SE of 1.7 bps/Hz.

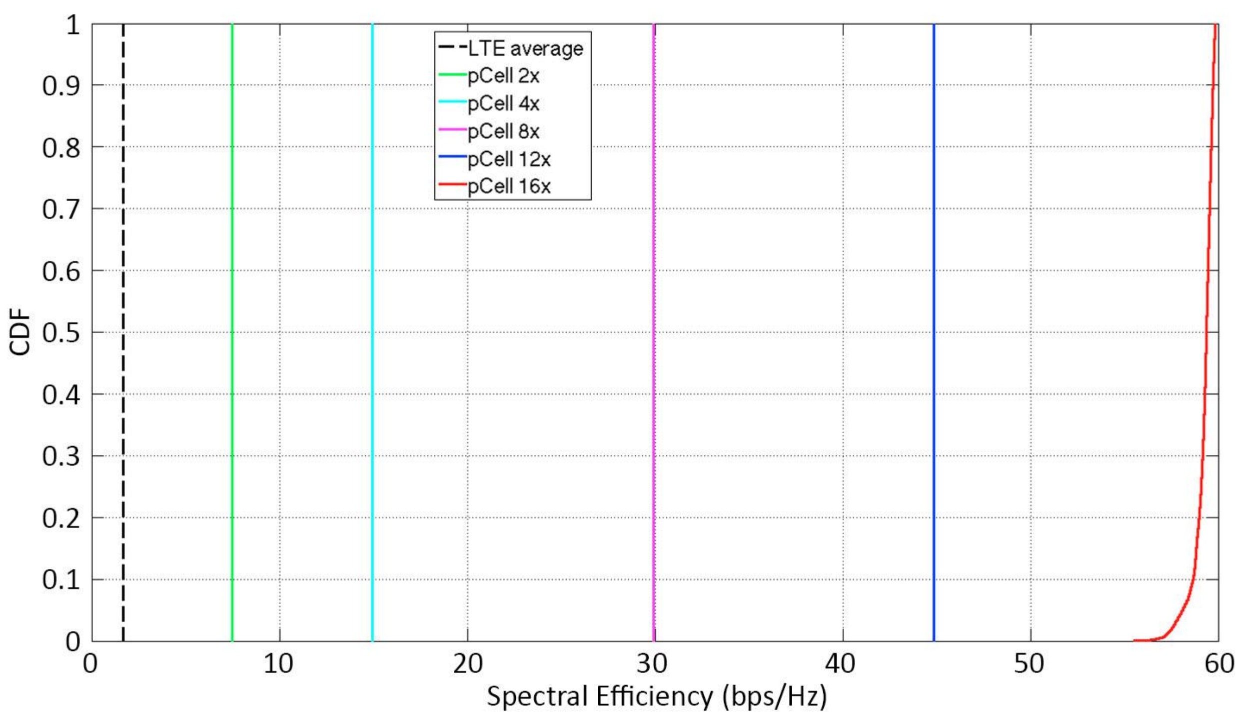

Figure 23: CDF of aggregate DL SE for different pCell orders

Figure 23 shows the cumulative distribution function (CDF) obtained from SE data captured throughout the coverage map in Figure 22 with different pCell orders (i.e., 2x, 4x, 8x, 12x, and

16x iPhone 6 Plus devices within 1 m^2 plexiglass table) compared against cellular LTE average SE of 1.7 bps/Hz. With 2 through 12 devices in 1 m^2, average SE is perfectly at peak, showing linear SE growth with devices. With 16 devices, the average SE is 1% below peak. The reason 16 devices show a 1% drop from linear growth is because such a large number of devices require contributions from a proportional number of pCell antennas, which are increasingly farther away from clustered user devices. The Artemis I Hub that was used for this test has an average power output per antenna of 1 mW. While this is adequate power for normal indoor densities of users, it is not quite enough power for all of the required antennas to reach the 1 m^2 table with adequate power to achieve 100% SE on all devices. But, 100% SE with 16 user devices is certainly achievable with pCell, by using higher power RF chains, or by distributing at least some of the 16 user devices elsewhere in the coverage area. In fact, if you look closely at the heat map in Figure 22, you see that in the central area under the antennas, all 16 devices are at peak. The only below-peak locations are around the edges of the room, because of the distance from the antennas.

pCell performs better when devices are distributed because they are within reach of more antennas and additional spacing yields higher space diversity. Clearly, the density of devices tested is far beyond the density of any real-world scenario. But, we have found that the 1 m^2 table serves as a good “stress test” for pCell. With a large number of devices, there are far too many possible distributed arrangements to exhaustively test them all and present them succinctly in this whitepaper. Furthermore, since we know that tightly-clustered user devices represent a worse case than any distributed arrangement of user devices, we can conduct worst-case surveys by moving the 1 m^2 table throughout the coverage area.

4.3.1.3 Performance with more user devices than network antennas

For the purpose of illustration, thus far we’ve always shown every user device in the pCell coverage area demanding the maximum data rate available. In a real-world scenario, per-user data rate demands varies enormously, with some users demanding a steady stream of 5 Mbps for an HD movie, other users demanding the maximum data rate for a brief interval for a download, other users requiring very small sporadic data requests to send texts or check email. But, of course, it is unrealistic that every single user in the network would be demanding the maximum data rate available all at once and all the time.

The examples shown thus far illustrate the capacity of the pCell system for a given pCell “order”, which is the number of user devices that can be served at full data rate concurrently. For example, in Table 3 at pCell order 12, there are listed 12 users, all within 1 m^2 that are running concurrently at 12.5 Mbps, which is peak DL data rate in a 5 MHz TDD channel, for an aggregate data rate of 12.5*12 = 150 Mbps (listed in the "5 MHz TDD Mbps" column).

This aggregate capacity can be divided among any number of users, using TDMA and OFDMA. For example, at pCell order 12, the pCell system could serve 24 user devices in the same 1 m^2, and if the aggregate capacity were divided equally among them, the data rate per user would be 12.5/2 = 6.25 Mbps. Of course, the data rate normally would be divided based upon individual user demand, resulting in a highly variable allocation, including zero allocation for users requiring no data at all at a given moment. So long as the total data rate demand is no more than the aggregate capacity (150 Mbps for pCell order 12 in 5 MHz of TDD), then the data demands of all users will be met. If demand is higher than the aggregate capacity of the order, then the pCell scheduler will seek to increase the order (as detailed Section 6.3 pCell clusters, below). If the order cannot be increased further, the pCell scheduler will limit data traffic, in accordance with the scheduling policies in place.

Note that, regardless of the aggregate capacity, each pCell device is limited to the maximum data rate that can be delivered within the channel bandwidth to that user. In the case of 5 MHz TDD, as used in the examples of this section, the peak data rate per user is 12.5 Mbps.

Table 4 shows a range of examples in 20 MHz of TDD bandwidth (commonly used in LTE deployments) for different pCell orders (listed in accordance with number of users who could receive full data rate at once), the SE for each order, and the data rate[159] per user at each order if the aggregate capacity were allocated equally among all users. The user data rates that are limited by the peak data rate each user device can receive are shown in dark blue.

pCell Order	pCell LTE DL		pCell SE vs. Cell SE[160]	Per-user Mbps for Number of Connected Users							
	Avg. SE bps/Hz	20 MHz TDD Mbps		1	2	4	8	16	32	64	128
2	8.4	113	5x	56	56	28	14	7	4	2	1
4	17	226	9x	56	56	56	28	14	7	4	2
8	34	452	19x	56	56	56	56	28	14	7	4
16	67	905	37x	56	56	56	56	56	28	14	7
32	134	1,810	75x	56	56	56	56	56	56	28	14

Table 4: 20 MHz TDD per-device data rate per pCell order for increasing numbers of connected users

Table 4 shows the actual achievable data rates if all of the connected users were sharing the aggregate data rate available for every pCell order through TDMA and OFDMA. For example, at pCell order 32 with 128 connected users, every user would receive a constant 14 Mbps, an adequate data rate for 4K UHD video streaming.

With carrier aggregation, even higher data rates would be sustained, both in the aggregate and per user device. For example, with five 20 MHz channels aggregated into 100 MHz, the aggregate data rate and the maximum per-user data rate would increase by 5x, resulting in over 9 Gbps aggregate data rate at pCell order 32, and a maximum of 280 Mbps per individual user, consistently served to devices, whether packed within 1 m^2, or distributed throughout the coverage area.

4.3.2 pCell SE consistency

As illustrated in Figure 3 cellular SE declines by a factor of 100 from cell center to cell edge[161], with average SE throughout the cell of roughly 20% of the peak SE with 2 network antennas and roughly 25% of the peak with 4 network antennas.

pCell has no cell edges for a direct comparison with cellular, but per the ITU[162], the cell edge corresponds to the 5% outage point in a cell, so the spread is between the 5% outage and the maximum data rate.

pCell SE is far more consistent. With up to 12 clustered users at 1 mW RF power, the 5% outage/peak spread is 0%. And, even with 16 user devices clustered in 1 m^2 at 1 mW, as can be seen in the Figure 22 heat map, throughout the coverage area, the 5% outage/peak spread is about 9%, with an average SE throughout the coverage area at 99% of the peak SE.

pCell's consistency not only provides a better user experience, but enables applications reliant on Quality of Service (QoS), such as streaming video, to operate with confidence that the network performance experienced upon initial connection will continue to be available on an ongoing basis. For example, streaming video today typically starts with a long period of buffering, where the video server preloads tens of seconds, if not minutes, of video in advance of starting the video so that video is cached in the event of a large drop in data rate (e.g. if a user in cell center moves to cell edge). Using pCell, streaming video can begin instantly with little or no buffering, since the connection will be continuously maintained with a highly consistent data rate.

4.4 Outdoor performance

While pCell technology has been tested in both outdoor urban and rural environments, thus far it has not been tested in detailed outdoor surveys. However, we have applied our measurements of pCell real-world behavior against standard outdoor models for simulations. In the past, when we have used the same simulation methodology, results have accurately modeled indoor behavior confirmed by indoor surveys. The following simulations compare cellular and pCell performance in urban outdoor non-line of sight (NLOS) environments.

4.4.1 Hexagonal and arbitrary layouts

Figure 24 simulates the same conditions shown in Figure 17 and Figure 18 of antennas uniformly spaced every 100 m with 200 mW maximum power, 20 MHz bandwidth and 64-QAM maximum modulation order[163]. The left simulation of Figure 24 shows an ideal hexagonal cellular base station layout with sparse users (one per cell) in ideal channels including only 3GPP pathloss without shadowing (for the sake of illustration). As in Figure 17 and Figure 18, 7 full cells shown in the middle of a larger field of cells. The right simulation of Figure 24 shows pCell's performance using the same antenna layout, maximum power and modulation order and user positions, but in the pCell case a 3GPP pathloss and shadowing model is used that is far less favorable than the cellular pathloss-only model. As in Figure 17 and Figure 18, 7 concurrent users are modeled in this simulation.

While the Figure 17 heat maps show *SINR*, the Figure 24 heat maps show user *data rate*.

In the cellular simulation on the left of Figure 24, the heat map shows what the data rate would be at each point in space if there were exactly one user per cell. At the center of a cell, the single user would experience a peak of about 56 Mbps, and on the cell edge the data rate would drop to below 7 Mbps (i.e., the 5% outage[164]), an 8:1 ratio. As noted in Section 3.1, in a real-world cell with shadowing effects (not included in these cellular simulations for the sake of readability), the data rate drop from cell center to cell edge is even more extreme, on the order of 100:1[165]. Further, if there were more than one user per cell, the data rate per user would drop proportionately relative to the number of users per cell due to bandwidth sharing among users.

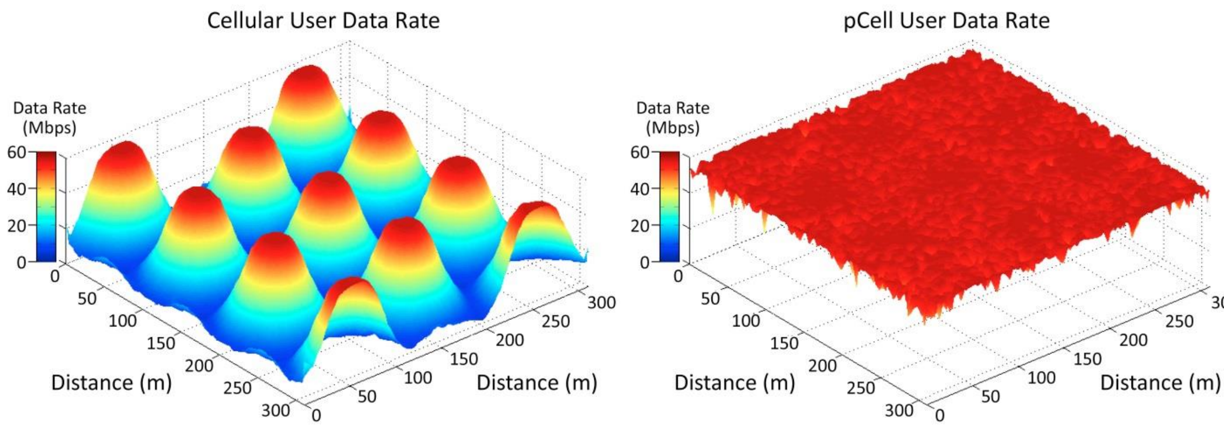

Figure 24: Data rate throughout coverage area: 200 mW, 100 m antenna spacing, 64-QAM max

In the pCell simulation on the right of Figure 24, the heat map shows what the data rate would be for each of 7 concurrent users distributed throughout the coverage area in any arrangement. The simulation positions the 7 users through one thousand iterations of 7 random locations each, with all or some of the users close to each other, all or some users far apart from each other, etc. throughout the coverage area. For example, the sparse arrangement of 7 users shown in Figure 17 and the clustered arrangement of 7 users in Figure 18 are just 2 of the one thousand user positions that were simulated.

As can be seen, regardless of how 7 users are arranged in the coverage area, the performance is highly uniform. The peak data rate is 56 Mbps and the average is 54 Mbps. The 5% outage is 48 Mbps. So, pCell's data rate spread is 56 to 48, or less than a 15% spread. Thus, with pCell, including modeling for pathloss and shadowing, concurrent users can expect performance within 15% of peak data throughout the coverage area. This is a far better result than can be achieved by a single user in a cell, let alone 7 users clustered or scattered in any arrangement throughout the cellular coverage area.

Figure 25 simulates the same conditions shown in Figure 19, with the same maximum power output of 200 mW per antenna as Figure 24, but with arbitrary antenna spacing. As with Figure 25 the data rate per user is shown, not the SINR.

In the cellular case on the left, there is only one user per cell. Because of the arbitrary antenna arrangement, there are large dead zones with very low data rates and, where there are cells close together, there are sudden drops in throughput, resulting in a highly variable and unpredictable data rate for the user. And, of course, in realistic scenarios if more than one user shares a cell, the data rate of the cell is divided among the users proportionately.

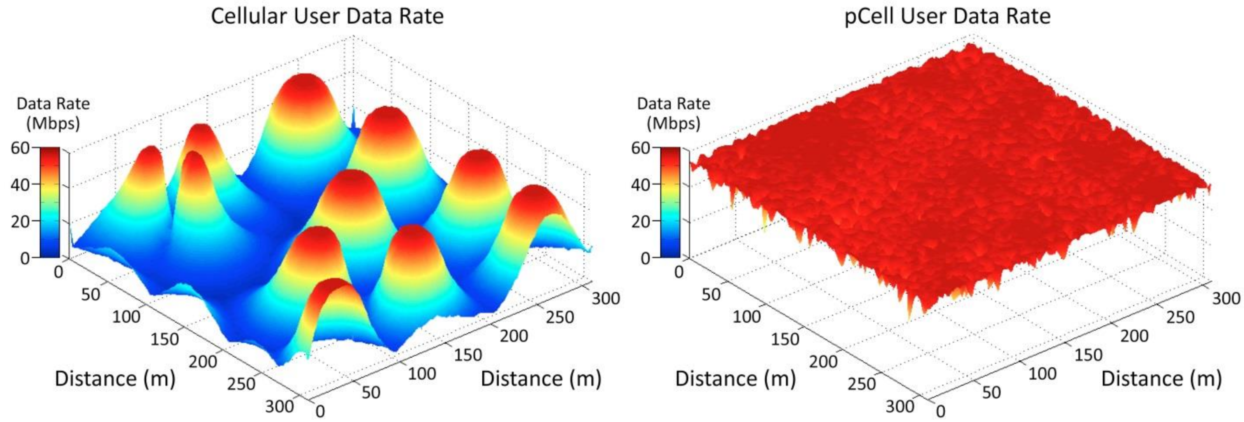

Figure 25 Data rate throughout coverage area: 200 mW, Arbitrary antenna spacing, 64-QAM max

In the pCell simulation on the right of Figure 25 the heat map shows what the data rate would be for each of 7 concurrent users distributed throughout the coverage area in any arrangement, positioning the users though one thousand iterations of 7 random locations each, just as in Figure 24, but in this case with the same arbitrary antenna arrangement used for the cellular simulation on the left.

As can be seen, the results are the same as they were for the uniform hexagonal antenna arrangement of Figure 25: regardless of how the 7 users are arranged in the coverage area, the performance is highly uniform: The peak data rate is 56 Mbps and the average is 54 Mbps. The 5% outage is 48 Mbps. So, pCell's data rate spread is 56 to 48, again, less than a 15% spread. Thus, with pCell, including modeling pathloss and shadowing, and modeling both uniform and arbitrary antenna arrangements, concurrent users experience a highly uniform data rate, within 15% of peak performance.

4.4.2 256-QAM performance

Future LTE-Advanced Rel. 12 devices will support 256-QAM DL[166], in addition to the 64-QAM and lower DL modulation supported by current LTE devices, increasing peak data rate by 25%. Because 256-QAM can only be demodulated in high SINR conditions, it can only be utilized at cell center where SINR is at its peak. Throughout the rest of the cell, modulation will fall back to 64-QAM and lower, resulting in little benefit.

Because pCell maintains high SINR throughout the coverage area, unlike cellular, pCell can maintain 256-QAM modulation throughout the coverage area, achieving a 25% improvement in data rate everywhere.

This is illustrated in Figure 26. The conditions are the same as they were in Figure 25 except that the peak power per antenna has been increased from 200mW to 500mW. Because of inter-cell interference, increasing the peak power for cellular results in little SINR benefit, but in the case of pCell, it increases the SINR throughout the coverage area.

Also, while Figure 25 simulates a maximum modulation order of 64-QAM, Figure 26 simulates a maximum modulation order of 256-QAM. Note that the data rate axis of Figure 26 has a maximum of 75 Mbps, while the data rate axis for Figure 25 had a maximum of 60 Mbps.

In the cellular case of Figure 26, the peak data rate of 75 Mbps is only achieved at the center of each cell. The 5% outage cell edge data rate is similar to that of Figure 25, at 6 Mbps.

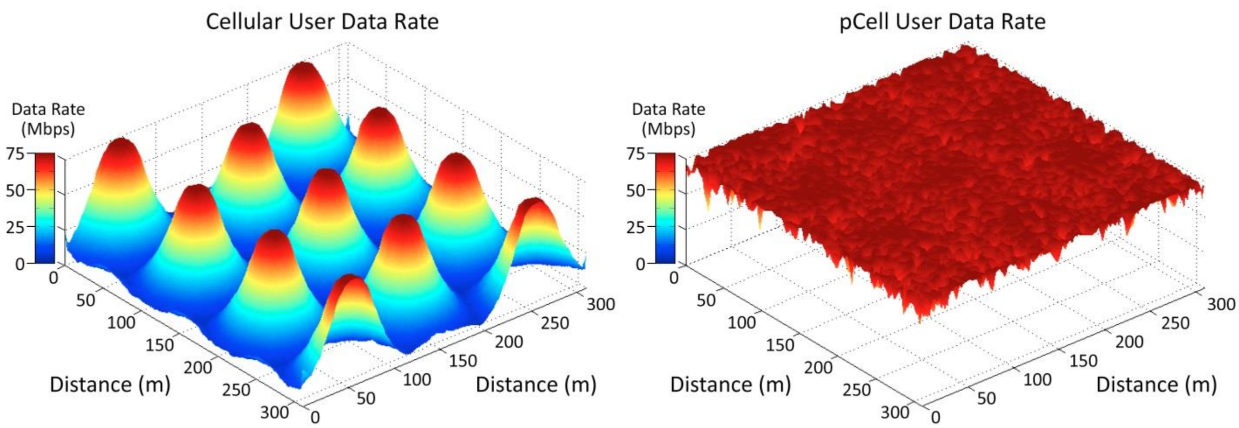

Figure 26: Data rate throughout coverage area: 500 mW, 100 m antenna spacing, 256-QAM max

But, in the pCell case of Figure 26, there is a dramatic benefit from 256-QAM. The data rate is maintained near the 75 Mbps peak most of the time for all 7 users, regardless of their arrangement in the coverage area, with a 5% outage of 62 Mbps, which is within 17% of peak data rate, and higher than the maximum 56 Mbps data rate achievable with 64-QAM. Thus, even at 256-QAM, pCell is able to maintain concurrent user data rates within 17% of peak throughout the coverage area.

4.5 pCell vs. Cellular

In Section 3.5, we summarized fundamental limitations of cellular architecture and noted that pCell overcomes all of the limitations. We list these limitations here again, and contrast pCell with cellular in each case. Some of the comparisons refer to pCell technologies that are discussed in later Chapters, and they are noted as such.

Issue	Cellular	pCell
Data rate	100:1 ratio cell center to cell edge	Consistently near-peak throughout
Cell-edge performance	Poor, due to low SINR	No cell edges, high SINR throughout
High user density performance	Poor, due to bandwidth sharing	Near-peak per user
Capacity scalability	Limited, due to small-cell inter-cell interference	Unlimited, scales uniformly as antennas are added
Capacity gain from spatial multiplexing	Up to 4x	35x achieved today, theoretically unlimited
Handoffs	High overhead	No handoffs
Antenna installation	Inflexible, expensive real-estate and backhaul	Arbitrary placement, inexpensive fronthaul[167]
3D environments	Poor vertical performance	Consistent performance[168]
Indoor location positioning	Inaccurate	Accurate[169]
Crisis scenarios	Loss of coverage, severe congestion	Consistent coverage, handles demand surges[170]
Security	Interceptable, vulnerable to compromised keys	Physically secure, even if keys are compromised[171]

Table 5: Cellular vs. pCell comparison

4.6 pCell ubiquitous connectivity

As discussed in Chapter 2, the convenience and simplicity of wireless connectivity has already displaced wireline connectivity, first with Wi-Fi as a short-distance broadband data solution, and increasingly with mobile, as a long-distance broadband data solution. But, what was lost in the transition from wireline to wireless is the consistency and reliability of wireline connectivity.

As shown in Section 3.1 and 4.4, cellular *inherently* cannot deliver anything close to uniform data rates to users even under ideal conditions, let alone in real-world conditions, where the peak to 5%-outage spread is much larger. Because cellular performance is so inconsistent, it limits the scope of its applications to ones that are usable under extremely variable network conditions (e.g., stuttering linear video/audio with large buffers, inconsistent Web page loading, etc.).

Effectively, pCell provides ubiquitous connectivity throughout the coverage area and delivers consistency and reliability close to wireline connectivity with the convenience and simplicity of wireless, scalable to any number of concurrent users in the same spectrum.

To fully exploit these propagation properties and achieve the SE gains showed in the field trials described in Section 4.3, pCell employs advanced signal processing techniques in the data centers as well as an innovative software-defined radio Cloud-RAN architecture compatible with the LTE and Wi-Fi standard that can be efficiently and cost-effectively deployed. These practical implementation aspects of pCell technology are described in the following Chapters.

5 pCell Deployment

5.1 pCell compatibility with current standard mobile devices

pCell was designed to be immediately deployed into the existing LTE and Wi-Fi wireless device ecosystems, without requiring any hardware or software modifications to current LTE and Wi-Fi devices. When an LTE or Wi-Fi device is located in a pCell coverage area, it experiences what appears to be a conventional LTE or Wi-Fi signal (at high SINR), both as the device attaches and as it continues to maintain the connection and moves throughout the coverage area.

5.1.1 pCell LTE

pCell is compatible with off-the-shelf Release 8 and above LTE user devices, with no software installation nor hardware modifications. Currently, pCell supports up to 20 MHz TDD.

pCell works with operator SIM cards, multi-operator SIM cards, and also SIM cards for private LTE networks.

pCell will support LTE-Advanced features of Carrier Aggregation as well as 256-QAM. In early years, legacy LTE devices will far outnumber LTE-Advanced devices, limiting overall cell capacity and mitigating much of the benefit of LTE-Advanced devices. With pCell, each LTE-Advanced device has its own pCell, operating at full performance concurrently with any number of slower legacy LTE devices sharing the same spectrum, resulting in immediate user benefit.

pCell is compatible with all mobile frequencies[172]. Further, pCell supports LTE protocols in unlicensed spectrum as well as in white spaces.

5.1.2 pCell Wi-Fi

In a future release, pCell will also be compatible with off-the-shelf Wi-Fi devices, both in 2.4 GHz and 5 GHz. pCell Wi-Fi SE grows linearly with the number of network antennas comparably to pCell LTE SE growth. In particular, pCell increases the capacity of the limited 2.4 GHz Wi-Fi spectrum, eliminating congestion in high-density situations where 5 GHz is unavailable because of pathloss or 2.4 GHz-only devices.

pCell Wi-Fi tests have verified pCell operates compatibly with Wi-Fi devices and scales linearly. Once robust testing has been completed, results will be detailed in a future white paper.

5.1.3 Voice support and internet connectivity

pCell will implement VoLTE for mobile devices that support VoLTE. Currently, pCell is implementing VoLTE protocol that is interoperable with Verizon and AT&T VoLTE,[173] but pCell SDR is flexible enough to support other VoLTE implementations as required by other operators.

pCell will also support 3G fallback for voice on an operator-specific basis, per the implementation requirements of the operator.

As shown at the top of Figure 27, pCell SDR attaches directly to the Internet at each C-RAN data center location, whether the pCell C-RAN data center is located remotely from the pCell deployment or within a venue, such as a stadium or business. pCell SDR can route traffic as desired by the operator onto one or more backbones so as to minimize hops and/or peering cost.

pCell SDR can also interface with an operator's stack at any level the operators requires if it chooses to route data to the Internet itself.

5.2 pCell Cloud-RAN architecture

The pCell system is implemented entirely in proprietary real-time software-defined radio (SDR) in a Cloud-Radio Access Network (C-RAN) architecture. The current version of the pCell system is implemented entirely on general-purpose x86 CPUs, so the pCell C-RAN data center consists of generic x86 servers and 10 GigE switches; there is no special-purpose hardware nor are there any dedicated baseband line cards as are typically found in a conventional C-RAN data center (frequently called a "Broadband hotel"). All of the communication in and out of the data center is through general-purpose non-synchronous digital links, such as 10 Gig E. As a result, any data center in the vicinity of a pCell deployment can be used or, if preferred, the pCell servers and switches can be located in an on-premises data center, for example, in a stadium.

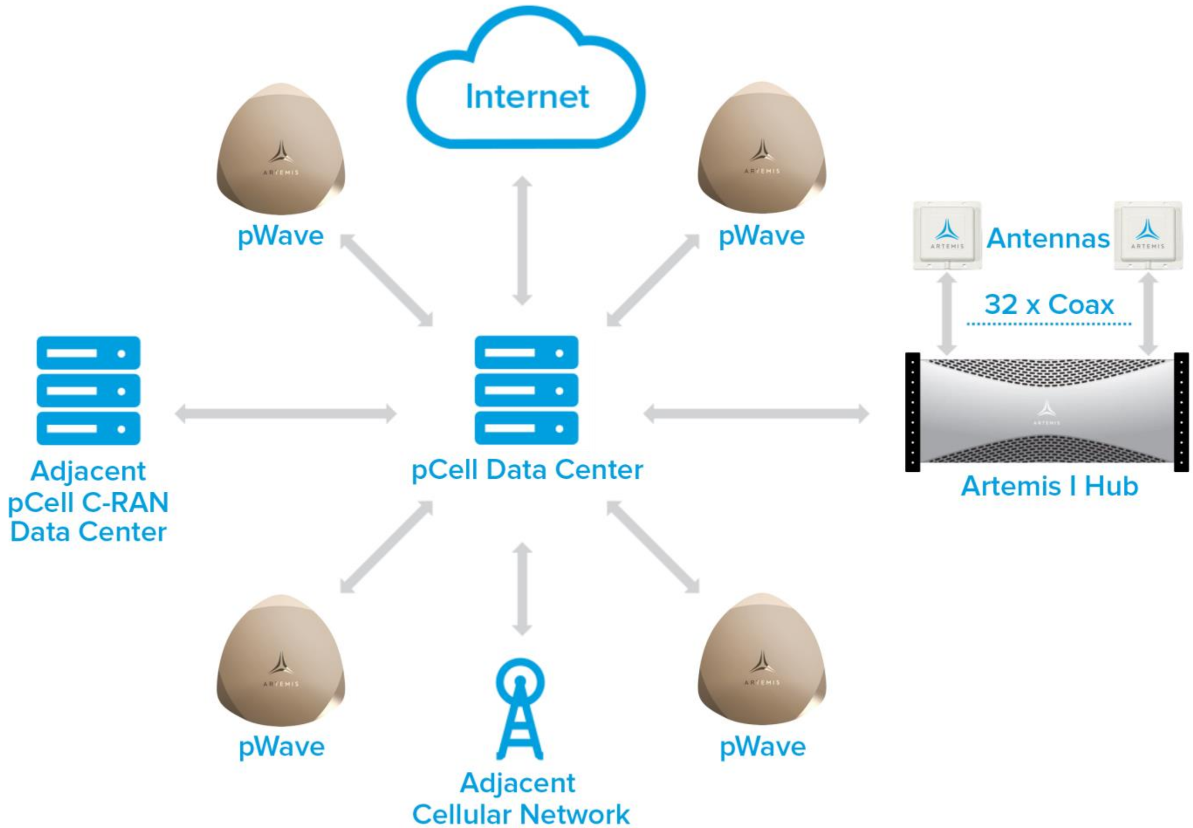

Figure 27: pCell C-RAN architecture

As illustrated in Figure 27, the pCell C-RAN data center is connected over fronthaul to pWave remote radio heads (RRHs), each with a single integrated antenna, and Artemis Hubs, each with multiple coax outputs driving multiple antennas. The pCell data center generates baseband waveforms and receives baseband waveforms in software in real-time from the pWaves and Hubs. For downlink (DL), the pWaves and Hubs receive the baseband waveform from the data center over fronthaul, synchronize it (as described in Section 5.5.3), convert it from digital to analog, modulate the carrier frequency, amplify it and transmit it through a single antenna. On uplink (UL), a single antenna receives the waveform, it is downconverted to baseband, digitized, synchronized and transmitted over fronthaul to the data center. The only exceptions are a few control messages in the LTE and Wi-Fi protocols which require a very fast response. If the fronthaul latency is too high to meet the response deadline, then a local DSP in the pWave or Hub receives the message and quickly generates a response, but this constitutes a tiny percentage of the overall traffic. The vast majority of the baseband processing is implemented in pCell SDR in the data center.

The pCell SDR implements both LTE and Wi-Fi protocols, but currently only TDD LTE is available for release. It is described in this section.

Each user device in a pCell network has its own pCell for the duration of the time it is attached to the pCell network, whether it is milliseconds, seconds, minutes hours or days, and whether for the duration of the link the user device remains in the same location, within the range of the same pCell C-RAN, or within the range of multiple pCell C-RANs. Thus, a complete real-time user state is maintained for each user in the pCell SDR within one or more pCell C-RANs that is created when the user first attaches until the user finally detaches.

The start of the attach procedure for an LTE user device first connecting to a pCell LTE network is the same as the attach procedure to a cellular LTE network: The pCell SDR DL provides standard LTE signaling information at a specified carrier frequency throughout the coverage area, the LTE user device detects this signaling information and initiates an attach procedure through the UL.

The pCell SDR receives the user device attach request and immediately spawns a new eNodeB instantiation in SDR for that user, and then synthesizes a pCell for that user. The new user then continues the attach procedure to the newly created eNodeB via its own pCell, independent of all other users concurrently using the same spectrum. From that point forward, until the user detaches from the pCell network, that eNodeB will remain as a real-time SDR process that supports the full LTE protocol stack for that user through the user's pCell. Regardless of how many other users happen to be sharing the spectrum at once, and regardless of how many eNodeB processes (or other protocol stacks) are concurrently executing the pCell SDR, that user device will experience what seems to be its own private LTE eNodeB, as if it is the only user in a cell. Further, the user will consistently experience high SINR, as if it is always near cell center.

Because each user device has its own eNodeB and a consistent pCell throughout the coverage area there are a number of differences between pCell LTE and cellular LTE. First, as the user moves through the pCell coverage area, there are no hand-offs; the user device experiences the same eNodeB wherever it goes, even if it travels over 1,000 km of highway on a pCell network. Second, there is no upper limit to the number of users that can be concurrently connected to the pCell network. Cellular LTE systems establish an upper limit to the number of users that can concurrently be attached to an eNodeB, both because the eNodeB implementations typically have capacity limits, but also because aggregate protocol overhead of such large number of users eventually swamps the capacity of the network. With pCell, even if tens of thousands of LTE user devices are concurrently connected to a pCell network in a dense area, such as a stadium or a downtown area, not only will the pCell network enable them to attach concurrently, but the protocol overhead associated with each user is confined to each user's pCell and as such, as more users are added the protocol overhead does not increase.

Thus, even at very high densities of LTE users, not only can all users remain attached to the network, but the users do not incur the overhead normally incurred by large numbers of users sharing one eNodeB in a cell.

When a user detaches from a pCell network, any user state information that needs to be retained by the network is stored, and the user's eNodeB instantiation is terminated.

5.2.1 Traversing adjacent pCell Cloud-RANs

So long as latency constraints are met, pCell C-RANs can be located wherever it is convenient, relative to the pCell network deployment. Thus, a pCell C-RAN might be located in a venue such as a stadium or hotel, or it might be in a data center serving a number of pCell networks that are within a 25 km radius. But, for large coverage areas, more than one pCell C-RAN will be needed to provide continuous pCell service.

As shown on the left in Figure 27, a pCell C-RANs can be interconnected with an adjacent pCell C-RAN to enable them to cooperate in maintaining pCells for users located in areas served by antennas from both C-RANs. When a user's pCell cluster is served by pCell antennas from more than one pCell C-RAN, the servers in one C-RAN execute the instantiation of the user's eNodeB and handle the DL and UL data flow to and from the user, but the servers in both C-RANs cooperate in synthesizing the user's pCell from a combination of pCell antennas from both C-RANs.

If the user moves and is served by antennas from only one C-RAN, then if it is the first C-RAN that had been executing the user's eNodeB instantiation, then that first C-RAN will continue to do so. If the user's new location is served only by the antennas of the second C-RAN, then the first C-RAN will transfer the state of the eNodeB instantiation to the second C-RAN, which will spawn a new eNodeB instantiation which will begin executing exactly where the first C-RAN's eNodeB instantiation left off. The user device will experience an uninterrupted link to the same eNodeB, despite the execution of the eNodeB being transferred from the pCell SDR in one C-RAN to the pCell SDR in an adjacent C-RAN.

5.2.2 Hand-off to adjacent cellular networks

As shown on the bottom of Figure 27, pCell C-RANs can be interconnected with adjacent cellular networks so that when a user moves beyond the edge of the pCell coverage area (or from a cellular network into the pCell coverage area), the user is handed-off to or from an adjacent cell.

pCell SDR is quite flexible and can implement any hand-off protocol that is required, but the hand-off process is both vendor- and operator-specific, so Artemis would work with the operator and vendor in each case to implement and test a protocol that interoperates with the operator's cellular network.

Although pCell SDR implements LTE protocol within its own network, it can hand-off to any cellular or other protocol required by the operator, including 3G, 2G or Wi-Fi.

5.3 pCell software-defined radio (SDR) architecture

pCell is the world's first end-to-end cloud-based commercial real-time software-defined radio (SDR) system. The entire LTE protocol stack down to baseband processing is implemented in pure SDR on general-purpose x86 CPUs.

Achieving both high performance and high efficiency SDR on general-purpose CPUs required entirely new approaches to SDR processing and performance optimization, resulting in an SDR system architected unlike any SDR system which has preceded it. But, the end result is a highly versatile and inexpensive SDR system, with extremely fast turnaround time for adapting the pCell system to be compatible with operator requirements.

The pCell system was architected from the outset as a pure SDR system that could be scaled indefinitely to accommodate increasing user density and data demands. pCell CPU requirements scale linearly with the number of pCell antennas in the coverage area, and further the linear scaling is achieved through independent parallel systems, with linear growth of parallel network interconnects. pCell was also architected to minimize connectivity requirements between adjacent pCell C-RANs to only exchange the data required to support users that are served by pCell antennas of both C-RANs.

In summary, as the number of user devices and demand grows, not only is it easy to deploy additional pWaves and Hubs, as detailed in previous sections, but it is easy to expand pCell computing capability in the pCell C-RAN to accommodate the added pWaves and Hubs. Additionally, connectivity costs for network traffic between adjacent pCell C-RANs are limited only to user traffic that spans both C-RANs.

5.3.1 pCell CPU architecture

Although the number of CPU cores required for a pCell SDR implementation varies depending on a number of factors, the current pCell SDR system running in trial now requires roughly two

3.4 GHz Intel Xeon cores per pCell antenna operating in 5 MHz of TDD LTE spectrum, or 8 cores per pCell antenna in 20 MHz of spectrum. These numbers include the full end-to-end implementation, from the baseband processing to the entire LTE eNodeB protocol stack and the Internet gateways.

For both operating cost and energy conservation reasons, pCell SDR code has been optimized to minimize CPU power consumption, resulting in pCell SDR CPUs consuming less than half of their rated power, requiring smaller heatsinks and less cooling, suitable for telephone closets as well as data centers. The pCell SDR code has also been optimized to minimize RAM and disc requirements, resulting in low cost. Further, because pCell SDR servers and switches are based entirely on standard components, they can be leased, resulting in little up-front capital for their deployment.

Although pCell SDR is currently implemented on x86 architecture, it is by no means tied to x86 and can be recompiled for other architectures, such as ARM or ARM+DSP, to the extent such systems offer additional cost or power-saving advantages.

5.3.2 pCell virtual radio instances (VRIs)

Section 5.2 discussed how the pCell SDR allocates one eNodeB for every user for the duration of the connection. Every eNodeB is implemented in the pCell SDR running in a pCell data center as a software instantiation called a "virtual radio instance" (VRI). The VRI is spawned as soon as the user begins the attach procedure and remains operational throughout the duration of user's connection, maintaining the user's active state. When the user detaches from the network, the VRI saves any user state relevant to future connections and the VRI instantiation is released. Each LTE VRI implements the entire protocol stack of the eNodeB. More generally, VRIs can implement any standard and proprietary protocol stack (e.g., WiFi, GSM, HSPA+, or pCell-specific protocol). VRIs executing different standard protocols coexist within the same pCell SDR system and can operate concurrently in the same spectrum.

Figure 28 illustrates a pCell data center with pCell SDR software subsystems relevant to discussing VRI functionalities shown. In this example, there are eight users and eight pCell antennas, and all eight users are demanding the maximum data rate concurrently. For each user 0-7, one VRI 0-7, respectively, is associated with it.

Each VRI in this example is an implementation of the full eNodeB (abbreviated eNB) protocol stack, starting from the connection to the Internet gateway on top of the VRI to the baseband waveform at the bottom of the VRI.

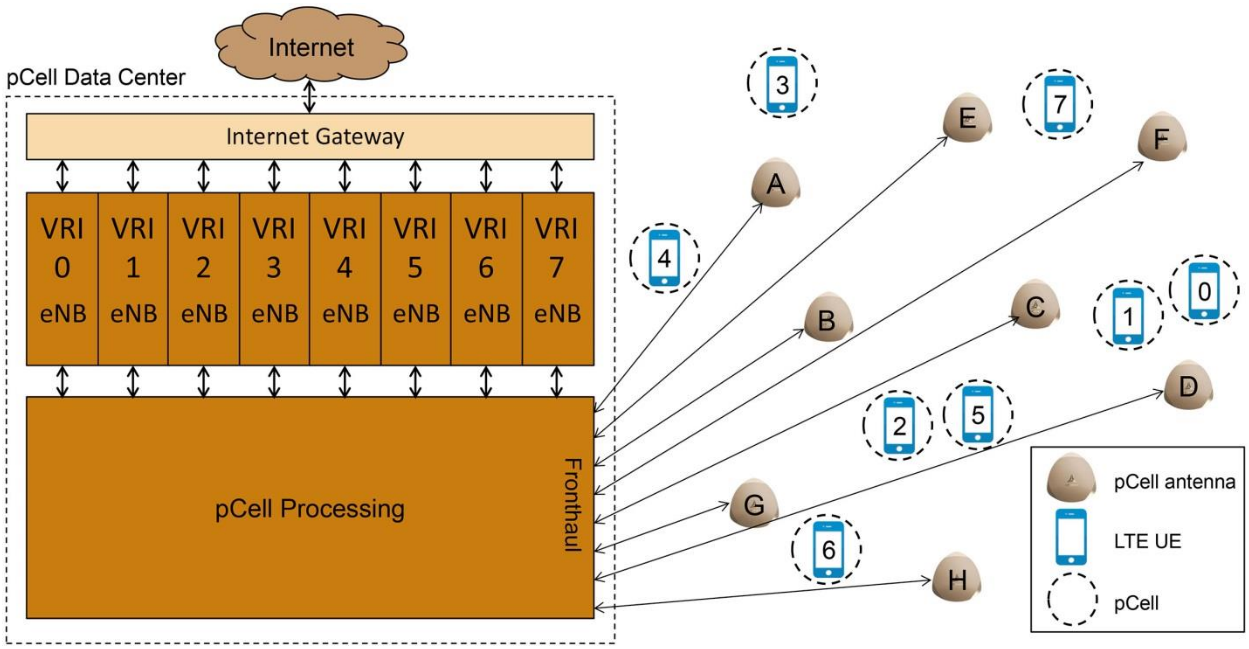

Figure 28: 8 VRIs, 8 pCell antennas and 8 user devices

The baseband waveforms of all of the VRIs feed into the pCell Processing, which combines the waveforms using pCell technology and produces the eight complex waveforms that are sent over fronthaul to the eight pCell antennas A-H. The eight pCell antennas concurrently transmit the waveforms, and the waveforms propagate through the environment and constructively interfere with each other at the exact location of each LTE user equipment device (UE). The combination of these waveforms at the location of each user device results in the synthesis of the baseband waveform that had been output by the VRI associated with that user device.

All the users receive their respective waveforms within their own pCells concurrently and in the same spectrum.

Of course, as users move throughout the coverage area, pCell Processing will keep the user's pCell locked to the user device so the user device will experience an uninterrupted, high-SINR link with its own VRI for the duration of the connection.

For the sake of illustration, we have shown eight users all demanding the full capacity of the spectrum concurrently. As previously discussed, this is a highly unlikely scenario. Users will have different data demands at different times, whether requiring steady streaming of HD at 5 Mbps for a 2-hour movie, or briefly downloading a file at maximum data rate, periodically sending text or emails at very low data rates, or being completely idle for a period of time. The pCell

SDR system can schedule VRI use of the aggregate pCell capacity using OFDMA, to divide up aggregate capacity by frequency and/or TDMA, to divide up aggregate capacity by time.

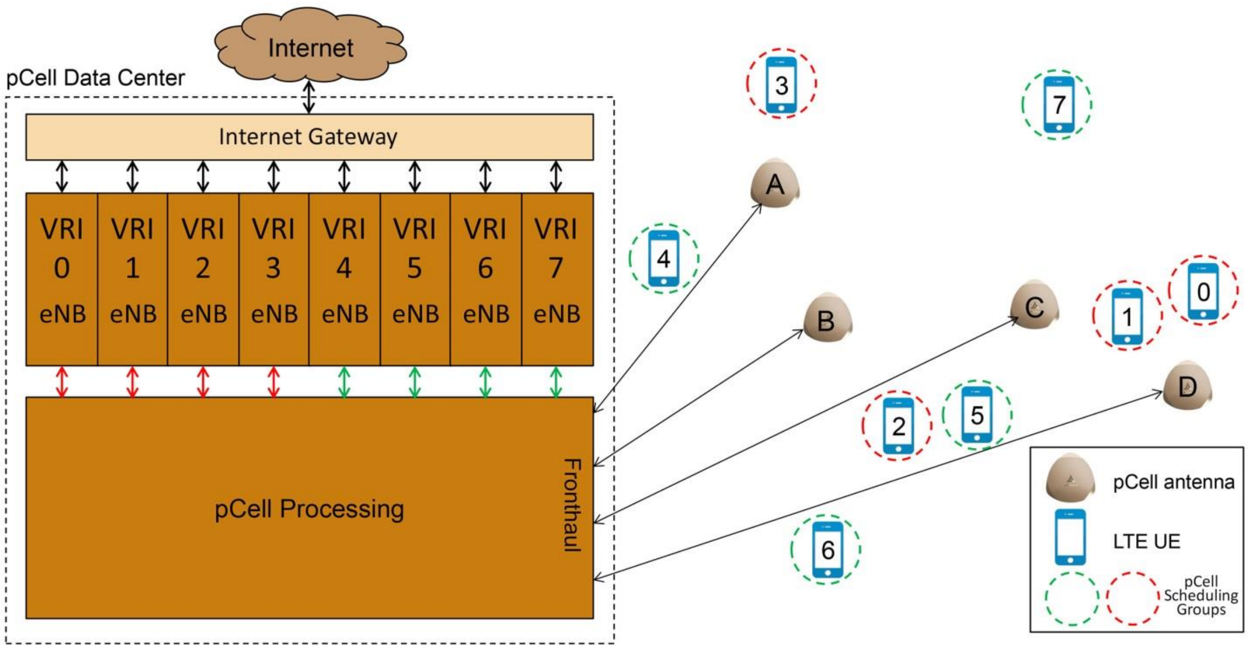

Figure 29: 8 VRIs using OFDMA and/or TDMA with 4 pCell antennas and 8 user devices

Figure 29 illustrates how pCell uses OFDMA to share aggregate capacity with more users (0-7) than pCell antennas (A-D). For example, the pCell SDR scheduler can decide to allocate the upper portion of the LTE resource grid to VRIs and users 0-3 (red dashed lines), whereas the lower portion to VRIs and users 4-7 (green dashed lines). Then, the pCell Processing computes the waveforms for the full resource grid and sends the waveforms to the four pCell antennas. In this way, all eight user devices concurrently have a link with their respective VRIs, with half of the spectrum capacity allocated to each.

Similar mechanisms apply for TDMA. In this case, the pCell SDR scheduler can allocate different groups of users to different LTE subframes while using the full spectrum at all time. For example, in Figure 29 VRIs and users 0-3 (red dashed lines) are scheduled to operate in one set of subframes, whereas VRIs and users 4-7 (green dashed lines) are scheduled in a different set of subframes.

Of course, TDMA can divide aggregate capacity among any number of users by subframes, and TDMA can be combined with OFDMA to divide up aggregate capacity by both time and frequency. The pCell SDR system is completely configurable as to what policies are specified by the operator in dividing up aggregate resources and whether OFDMA and/or TDMA is used.

Note, however that pCell scheduling is *quite* different than cellular scheduling. With cellular, there is enormous variance in SINR and user data rates, and complex scheduling decisions must be made as to how much capacity to allocate to each user. For example, if a user is streaming an HD video at 5 Mbps, the user may be consuming a relatively small percentage of a cell's capacity at cell center. But, at cell edge, the user may be consuming all of the cell's capacity since the achievable data rate is so low. The scheduling policy must decide whether to interrupt the user's video while the user is located at cell edge so that the other users in the cell are not "starved" of data, or to allow the user to continue to watch the video uninterrupted. Often there is no good scheduling solution and a sub-optimal solution is used.

Because all pCell users in the coverage area have high SINR and are capable of receiving near-peak data rate all of the time, pCell scheduling not only is simpler, but the outcome for users is far better since scheduling decisions are not based on the inconsistent user's SINR (as in cellular), but just on data demand and aggregate data capacity.

Uplink (UL) traffic utilizes VRIs analogously to downlink (DL) traffic. As would an eNodeB in a conventional cellular deployment, each eNodeB in a VRI receives UL requests from the associated user device and grants UL transmissions. Aggregate UL capacity is scheduled by allocating frequency and time slots using SC-FDMA and TDMA, respectively.

As the user moves through the coverage area, its VRI instantiation will remain an active software instantiation in the pCell SDR system, running in the pCell data center serving the user's coverage area. The user will go in and out of the range of different antennas (as detailed in Section 6.3), and the pCell SDR system will continue to direct the waveform from the user's VRI through the appropriate pCell Processing to reach the antennas within range of the user.

As the user moves across coverage areas supported by different pCell C-RAN data centers, its current pCell data center will "teleport" (i.e. send the complete, live execution state information) of the VRI to the adjacent pCell data center, which will continue executing the VRI without interruption. The VRI teleportation can occur in less than an LTE subframe time, resulting in uninterrupted connectivity for the user. Similarly, as the user moves in or out the pCell coverage area, the pCell C-RAN data centers supports hand-off to adjacent cellular systems operated by any carrier.

Although pCell VRIs are used for supporting an eNodeB protocol for each user and are also used for scheduling aggregate data capacity among users, pCell VRIs also enable concurrent operation of different wireless protocols to different user devices in the same spectrum. Figure 30 shows eight VRIs concurrently instantiated and executing in the same pCell SDR system, with different protocols executing in different VRIs. VRIs 0 and 4 are executing LTE Rel. 8 eNodeBs, VRI 5 is executing an LTE Rel. 10 eNodeB, and VRI 2 is executing an LTE Rel. 12 eNodeB. These are all LTE or LTE-Advanced protocols with different features. For example, LTE Rel. 12 supports 256-QAM, while Rel. 8 only supports up to 64-QAM.

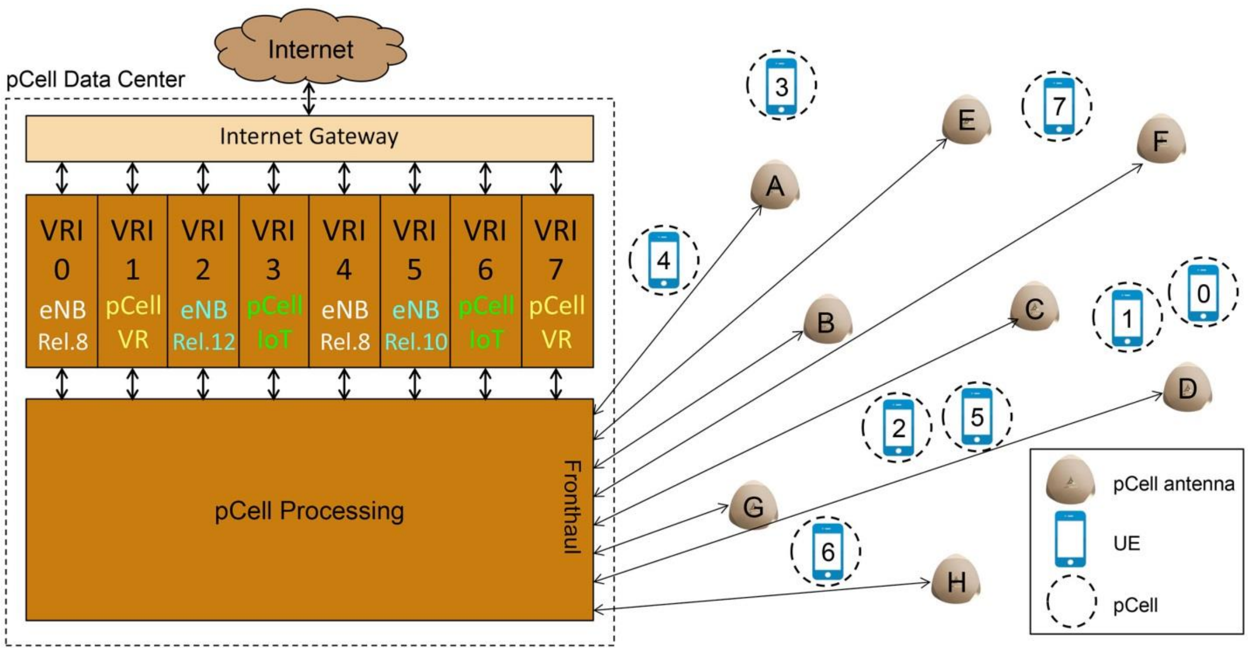

Figure 30: 8 VRIs with multiple protocols

Notably, VRIs 1, 3, 6 and 7 are *not* executing LTE protocols. VRIs 1 and 7 are executing pCell VR protocol and VRIs 3 and 6 are executing pCell IoT protocol. Both of these protocols are proprietary protocols optimized for particular devices which have requirements that could not be met by LTE. For example, pCell VR is optimized for sub-millisecond latency traffic for Virtual Reality and Augmented Reality headsets, which is a far lower latency than can be achieved by LTE protocol, and pCell IoT is optimized for very low power, low cost “Internet of Things” devices whose cost and power requirements cannot be met by LTE protocol. (These protocols are described in more detail in Section 7.2.)

All of the devices would be sharing the same spectrum concurrently used by conventional LTE devices, thus allowing LTE-Advanced devices to utilize LTE spectrum with higher performance, without being slowed down by the initially larger number of standard LTE devices unable to support higher speeds. And, it allows devices requiring specialized protocols that LTE can’t support to utilize the same spectrum as LTE devices.

5.4 pCell infrastructure and fronthaul

pCell was designed to dramatically reduce the capital expense (CAPEX) and operating expense (OPEX) of infrastructure deployment while providing far more complete and reliable coverage.

5.4.1 Cellular infrastructure is expensive to deploy and operate

Half of the OPEX of a cellular network is the backhaul and real estate rental[174]. A major reason for this high OPEX is the requirement to locate cellular base stations at *specific* locations to conform to a cell plan. As shown in Figure 17, above, cellular layout is carefully planned to *avoid* interference, choosing locations and transmit power to minimize transmission overlap. As shown in Figure 19, if cellular base stations were placed in arbitrary locations, the base stations that were too close to each other would have to reduce power to avoid interference, and with too little power, base stations that are too far apart would leave dead zones between them. Further, because an entire cell relies on transmissions from a single centralized location, the chosen location needs to be in a good vantage point relative to the coverage area, often on a tall tower, so as to minimize shadowing and dead zones.

This leaves cellular operators with few real estate choices for locating base stations, often resulting in high rental costs from the owners of these ideal locations. Although LOS backhaul is usually much less expensive than fiber, both for installation and for monthly cost, because cellular base stations have to be located in a narrow range of locations to conform to a cell plan, it is purely chance if a base station location happens to have an LOS view to a backhaul source.

Small cells are even more expensive to place since they have similar backhaul requirements as macro cells (e.g. both deliver the same peak data rate in a given bandwidth), but the range of installation locations not only narrows proportionately to cell size, but is also complicated by obstructions in the environment and inter-cell interference considerations. Given the limited choices of street-level locations, even if efficient small cell locations are available, provisioning backhaul to such specific sites can be extremely expensive.

Also, because service in a base station's coverage area is lost if the base station fails, base stations require some form of power backup, either from a battery or generator. This adds to the cost, size and maintenance requirements for base stations, particularly small cells.

5.4.2 pCell infrastructure is inexpensive to deploy and operate

Figure 31 shows a deployment comparison between cellular backhaul/fronthaul and pCell fronthaul. Because of the requirement to locate a cellular base station (BTS) or remote radio head (RRH) in accordance with a cell plan, cellular backhaul/fronthaul generally requires fiber[175]. If, by chance, some locations happen to be situated with an LOS view to a fiber backhaul/fronthaul location, then LOS[176] can be used, or in the case of indoor/venue small cells copper Ethernet can be used, but in general, cellular deployment requires fiber infrastructure.

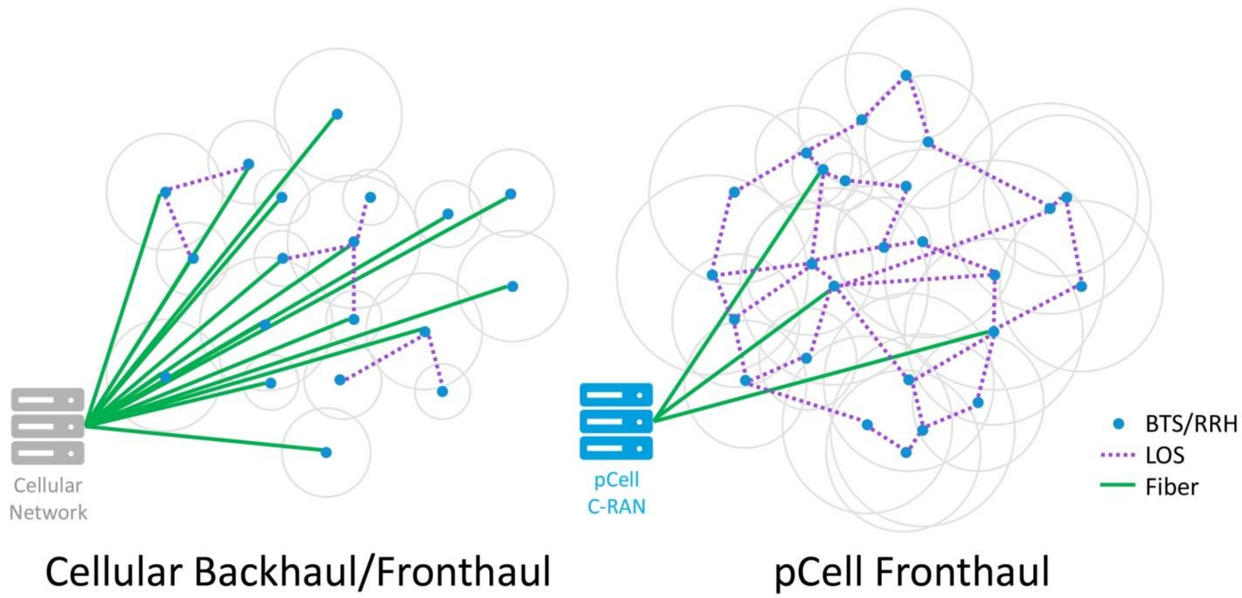

Figure 31: Cellular backhaul/fronthaul vs. pCell fronthaul

pCell RRH or Hub antenna (collectively "pCell antenna") placement simply requires overlap—in any pattern, at any power levels—throughout the coverage area. As such, pCell antenna locations are arbitrary; instead of being limited to a narrow range of locations that conform to a cell plan, virtually any location can be utilized, whether close or far from other pCell antennas, whether outdoor or indoor, whether at street-level or on a rooftop, and whether in free space or in an area full of obstructions.

Because of this flexibility, pCell antenna locations can be chosen based upon where it is convenient and least expensive to place them, for example, where there is inexpensive rent and an LOS view to a location with fronthaul back to the pCell C-RAN. As shown in Figure 31, an LOS fronthaul mesh network can be created, served by only a small number of fiber feeds. LOS radios are available with exceptionally low latency (e.g. <10 microseconds[177]), so the cumulative latency through the LOS mesh can be very low. Also, with routing redundancy, LOS mesh architectures can be highly robust against single-link failures.

pCell can also use fiber (or copper Ethernet) links where such links are available (e.g. fiber-connected buildings, indoors, venues, etc.) and can use coax in a DAS. But, unlike cellular, pCell antennas can be located where such links are already available or easily installed, without the constraints of conforming to a cell plan or the complexities of small cell inter-cell interference. Also unlike cellular, the vast majority of pCell RRHs in a pCell network do *not* require backup power from generators or batteries. Because pCell antenna transmissions overlap, if some of the RRHs lose power, there will still be coverage from other RRHs; the only impact will be reduced aggregate capacity until power is restored. To provide coverage during a regional power failure, a subset of pCell RRHs and fronthaul can be deployed with backup power to provide basic overlap throughout the coverage area. Although the network would have diminished capacity, a regional power failure would result in reduced data demand since many high data rate devices in the coverage area, e.g., TVs, become inoperative without power.

pCell infrastructure is easily expanded as user demand increases over time. In a given area pCell capacity scales linearly with the number of pCell antennas within range of users, so increasing capacity is a matter of adding pCell antennas in high-demand areas. Unlike cellular, which requires complex cell planning and interference testing to subdivide existing areas into smaller cells, pCell antennas can be placed in arbitrary locations in the general vicinity of the target area and there are no modifications to existing pCell antennas. Literally, when a new pCell antenna is activated, within milliseconds it is operational with the rest of the pCell system, adding capacity to the target area.

5.4.3 pCell fronthaul data rate comparable to cellular backhaul data rate

Another economic and practical consideration is the efficiency of the backhaul/fronthaul relative to the average data rate delivered to users over their wireless links.

Cellular can be provisioned via either backhaul to a BTS or via fronthaul from a C-RAN to an RRH. Although C-RAN/fronthaul offers advantages in flexibility, fronthaul data rate per RRH can be 9x or more higher than cellular backhaul data rate per BTS[178], and further, fronthaul typically requires specialized fiber, such as CPRI[179], that carries synchronization and clock signals. Thus, relative to the average data rate delivered to users, cellular backhaul is much more efficient than cellular fronthaul.

pCell is provisioned via fronthaul from a C-RAN to pWave RRHs or to Artemis I Hubs. Relative to the average data rate delivered to users over the wireless link, pCell fronthaul (unlike cellular fronthaul) operates within 10% of the efficiency of cellular backhaul. Also unlike cellular fronthaul, pCell fronthaul uses conventional, non-synchronous connectivity, whether fiber, LOS or copper Ethernet. Yet, pCell experiences the advantages and flexibility of C-RAN architecture.

For a direct comparison of cellular backhaul and pCell fronthaul, consider the highest data rates for both cellular and pCell, assuming TD-LTE (3:1) frame structure and 5 MHz bandwidth as in Table 2.

The highest cellular data rate in Table 2 is with 4 network antennas. The average data rate to users is 5.7 Mbps, but the cellular backhaul to the single 4-antenna base station must be provisioned to sustain the peak data rate of 25 Mbps (for two-antenna devices), e.g. if a user is near cell center. Thus, provisioned cellular backhaul data rate compared to average user data rate is 25/5.7 = 4.4x.

A comparable 4-antenna pCell network in Table 2 has an average data rate to users of 50 Mbps, and the pCell fronthaul requirement is 236 Mbps (59 Mbps per pCell antenna). Thus, provisioned pCell fronthaul data rate compared to average user data rate is 236/50 = 4.7x.

These results are summarized below:

	5 MHz TDD LTE		
	Cellular Backhaul	Cellular Fronthaul	pCell Fronthaul
Total provisioned data rate (Mbps)	25	227	236
Average delivered data rate (Mbps)	5.7	5.7	50
Provisioned vs. average delivered data rate	4.4x	40x	4.7x
Provisioned vs. delivered relative to backhaul	1x	9.1x	1.07x

Table 6: Cellular fronthaul/backhaul vs. pCell fronthaul

In conclusion, relative to average delivered data rate, the pCell fronthaul uses only 1.07x (7%) higher data rate than cellular backhaul, and both utilize conventional non-synchronous IP infrastructure. In contrast, cellular fronthaul has over 9x higher data rate than cellular backhaul and requires specialized synchronous, clocked infrastructure. Thus, pCell benefits from the advantages and flexibility of a C-RAN architecture with less than a 10% higher data rate cost compared to cellular backhaul.

5.5 pCell radio deployment options

pCell provides a number of options for radio deployment as illustrated in Figure 32. Fiber fronthaul is routed from the pCell SDR servers in the pCell data center that can be:

1. connected to pWave remote radio heads (RRHs) configured for fiber interface,
2. connected to Line of Sight (LOS) radios which connect to pWaves configured for gigabit Ethernet (GigE),
3. connected to a gigabit Ethernet switch which connects to pWaves configured for GigE,
4. connected to an Artemis Hub with up to 32 radios that connect to antennas through coaxial cables

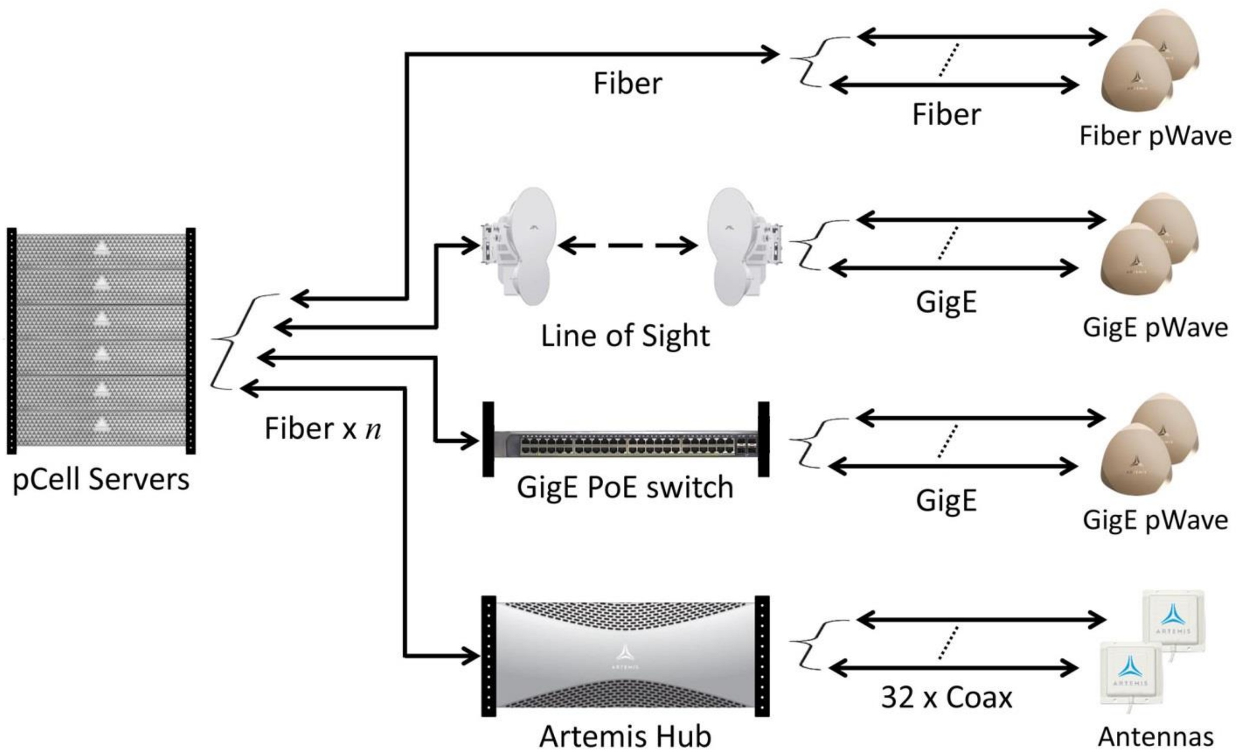

Figure 32: pCell radio and antenna deployment options

5.5.1 pWave remote radio heads

Each pWave is a single antenna, single RF chain device. pWaves can be manufactured at any power level (e.g. 100mW to 20W or more), with the size of the enclosure scaling in accordance with the heat dissipation requirements. A 5W pWave unit is shown in Figure 33. pWaves can be manufactured with GigE copper or fiber interfaces, using PoE with copper GigE, if the power requirements are within PoE limits, or otherwise using an external power source. Each pWave has a GPS receiver and a GPS local oscillator built in, and it has a bi-directional external sync connector. Each pWave can be configured to slave to the GPS receiver, its GPS local oscillator, to an in-band synchronization signal (detailed in Section 5.5.3), or to a wired GPS signal. Also, each pWave can be configured to output an in-band synchronization signal or a wired GPS signal. pWaves intended for outdoor operation are waterproof.

Figure 33: 5W pWave RRH

5.5.2 Artemis Hubs

The Artemis Hub is a multi-radio hub with multiple coaxial cable outputs for indoor and venue installation. The Artemis Hub has a 10 GigE fiber interface for fronthaul. Each Hub has a GPS local oscillator built in, and it has a bi-directional external sync connector. Each Hub can be configured to slave to its GPS local oscillator, to an in-band synchronization signal, or to a wired GPS signal. Also, each Hub can be configured to output an in-band synchronization signal or a wired GPS signal. The first member of the Artemis Hub family is the Artemis I Hub, which is available for trials. The Artemis I Hub, shown in Figure 34

Figure 34: Artemis I Hub front/back with antenna

has 32 coaxial cable outputs with a maximum power output of 1mW per antenna. Artemis Hubs for commercial deployment will begin at 100mW per antenna.

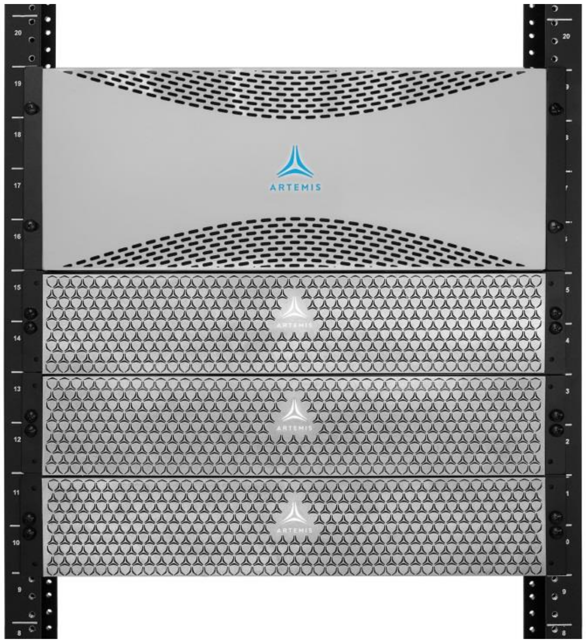

Figure 35: Artemis I Hub rack-mounted with three pCell SDR servers

All Artemis Hubs and low-wattage pWaves are frequency-agile from 600 MHz to 6 GHz to cover all LTE bands worldwide, Wi-Fi (ISM) bands as well as future bands yet to be allocated for mobile. Thus, when pWaves and Artemis Hubs are deployed they are both band-neutral and carrier-neutral. Also, the Artemis I Hub supports baseband signal bandwidth up to 20 MHz per coax output. Multiple coaxial outputs can be combined to single antennas for wider bandwidths (e.g. 40 MHz or 60 MHz) for carrier aggregation and other wideband applications.

Any number of pWaves and Artemis Hubs at varying power levels and using any form of fronthaul can be combined in the coverage area. As operators need to grow, additional pWaves or Artemis Hub antennas can be added to the coverage area wherever it is convenient and simply turned on. The pCell SDR system will immediately make use of the new radio and add capacity to the coverage area. Similarly, if a pWave or Artemis Hub antenna needs to be deactivated (e.g. for maintenance, etc.), the pCell SDR will immediately cease using the radio and the only impact will be that capacity will be reduced in the coverage area.

5.5.3 pCell radio synchronization

Artemis I Hub and pWave RRH radios (collectively "pCell radios") transmit and receive synchronously. The pCell system was designed to operate within the timing precision of GPS clock signals, whether received from GPS satellites or synthesized from a local GPS clock generator.

pCell radios receive their synchronization reference from either:

1. A direct GPS reference
    a. pWave radios each have a GPS receiver and slave to the global GPS clock
    b. All radios are capable of either:

        i. Slaving to an external GPS clock
        ii. Synthesizing a free-running GPS signal, if designated as the clock master

2. An in-band timing signal synchronized to a GPS reference that is non-interfering to LTE
    a. One radio is designated to be a cluster master to radios within its transmit range
    b. The cluster master transmits an in-band timing signal
    c. The cluster master receives its GPS timing reference from either:
        i. A GPS receiver, if outdoors
        ii. Slaving to an external GPS signal
        iii. Synthesizing a free-running GPS signal, if designated as the clock master
        iv. Reception of an in-band timing signal from another cluster master

In the case where two radio clusters are indoors and the spacing is too large for an in-band timing signal to be transmitted between them, then either:

1. A GPS signal must be wired between a radio in one cluster to a radio in the other cluster
2. Both radio clusters must slave to a synchronized external GPS reference
3. The two radio clusters will have independent time references and will not be able to concurrently support users that are within range of both clusters.

Since pCell radios receive their synchronization from one of the above GPS clock sources, they remain synchronized, despite utilizing non-synchronous fronthaul.

6 pCell Technologies and Features

pCell is a radical departure from conventional wireless technologies, utilizing new technologies not previously used in wireless systems, and also enabling new capabilities and features that would be difficult or impractical to implement in conventional wireless systems. This chapter discusses several new pCell technologies, and also details new features enabled by pCell.

The list of new technologies and features is by no means exhaustive, but it serves to illustrate how differently pCell operates than conventional wireless, and the wide range of applications pCell enables.

6.1 Exploiting wireless propagation effects

As shown in Chapter 4 through practical and simulated measurements, pCell delivers highly consistent wireless performance throughout the coverage area, even in the case of arbitrary antenna placement. Its SINR consistency is such that 256-QAM can be utilized throughout the coverage area, providing a uniformly higher data rate per user (e.g. ideal for real-time video streaming or video teleconferencing).

Further, pCell capacity gain scales linearly with the number of users: as more users are added to the coverage area (even at extreme densities of 16 users within 1 m^2) all users continue to receive peak or near-peak data rates throughout the coverage area.

pCell achieves these gains by exploiting two key properties of wireless propagation channels that have been mostly underutilized by cellular networks over the years: space selectivity and RF interference, as described hereafter.

6.1.1 Space selectivity

Space selectivity is an essential property of wireless propagation channels needed to obtain high capacity gains via spatial multiplexing schemes as described in Section 3.2. In general, space selectivity is achieved by placing transmit/receive antennas far apart and/or by operating in propagation environments consisting of large number of scattering objects that create multi-paths.

Commercial cellular MIMO base stations have antennas clustered only a few wavelengths apart. As such, cellular networks rely highly on multi-paths in the propagation environment to operate with spatial multiplexing schemes. But multi-paths are typically limited, unpredictable and highly dependent on the type of environment. Further, since all antennas of MIMO arrays are placed at a centralized location, their transmit waveforms are all subject to similar pathloss and shadowing effects as they propagate through the channel. Thus, when the user is in an unfavorable location, the links from all centralized antennas may degrade all at once, possibly causing the connection to drop. This “incidental” (i.e., conditions that might, by chance, exist) micro-diversity[180] in MIMO systems can provide only limited (e.g., at most 4x) SE gains in typical scenarios.

By distributing antennas far apart, pCell achieves highly diverse angular directions of propagation of radio waves over the wireless links, thereby synthesizing a larger number of independent multi-paths from the distributed antennas (rather than relying on limited multi-paths from the environment as in point-to-point MIMO links) yielding higher degrees of space selectivity. Further, pCell’s wireless links from distributed antennas undergo independent pathloss and shadowing effects, resulting in highly diverse propagation paths and benefits from macro-diversity[181]. Thus, pCell creates a high degree of “structural” (i.e., conditions that are very likely to exist by design) macro-diversity, which is used to achieve consistent multiplexing gain throughout the coverage area, scaling linearly with the number of users.

6.1.2 RF interference

Interference from adjacent coverage areas has been the bane of wireless system planning since the first wireless deployments, requiring careful antenna placement, aiming, and power management. Enormous effort is made in cellular systems planning to avoid inter-cell interference, particularly as cells become smaller[182]. As noted in Section 3.3, current cellular standards utilize inter-cell coordination or cooperation schemes to mitigate interference, particularly at the cell-edge, but performance gains through field trials so far have been only marginal.

Inter-cell interference in cellular networks has grown past the point where it is simply a challenge in wireless deployments; it indeed establishes an upper limit for capacity.

pCell exploits RF interference from multiple distributed antennas. By exploiting interference, pCell not only eliminates cell planning, pCell leapfrogs cellular’s upper limit for capacity, it seeks to maintain a high degree of overlap throughout the coverage area, thus providing consistently high SINR throughout.

6.2 Channel state information

To synthesize a pCell precisely at the location of each user device in space, pCell utilizes channel state information (CSI) obtained through channel sounding or feedback transmissions from the device, depending on whether TDD or FDD mode is used. Since FDD transmit and receive bands are different, to provide CSI feedback, FDD devices must measure the DL channel and transmit

CSI on the UL channel. With TDD systems, the UL and DL channels are reciprocal, and as a result UL channel sounding transmissions—which have far less overhead than CSI feedback—can be utilized for characterizing the DL channel.

Although pCell has been implemented in both FDD and TDD modes, the first commercial release of pCell only implements TDD. TD-LTE operation is as follows:

LTE devices are configured to transmit various UL channel sounding signals, such as SRS or DMRS. When an LTE device in the pCell coverage area transmits a sounding signal, the signal is received by the pCell antennas within range of that device. Each antenna fronthauls its received sounding signal from that device back to the pCell data center, and the signals are processed together to determine precisely what signals would need to be transmitted from each of the pCell antennas such that when they sum together at the location of the user device, they will synthesize the desired waveform for that user device.

pCell exploits channel *reciprocity*: the property for which the same physical propagation channel in a DL transmission is experienced over a UL transmission. Although real-world RF propagation environments are inherently reciprocal, real-world RF devices are not: the user UL transmitter and pCell radio UL receiver RF chains are dramatically different than the pCell radio DL transmitter and user DL receiver RF chains. Further, RF chains vary between different models of user devices, and even due to manufacturing variances between units of the same model. RF chain characteristics vary with temperature, the state of the battery power, antenna coupling, adjacent RF transmissions, etc. In short, the UL and DL RF chains are completely different and constantly varying.

As a result, conventional reciprocity techniques, which require calibrated UL and DL RF chains, cannot be utilized.

pCell overcomes these issues with a new approach to exploiting reciprocity that derives highly precise DL channel information from a UL transmission, despite the fact the UL and DL RF chains have different, time-varying RF characteristics and further, a wide range of highly different RF chains are transmitting into the same spectrum at once.

Consequently, the very brief UL channel sounding signals transmitted by a user device are sufficient for the pCell network to quickly derive extremely high-resolution DL channel information, which it then uses to calculate the DL waveforms to precisely combine at the location of the user device and synthesize a precise waveform for that device.

For UL data transmissions, user devices all concurrently transmit the UL signals, which interfere with each other so that a different combined waveform is received at each pCell antenna within range of a user device. The same UL sounding signals are used by the pCell data center to separate these combined waveforms into an individual UL waveform for each user device.

Thus, by using only basic UL sounding signals and a new approach to exploiting reciprocity, pCell is able to concurrently deliver independent DL waveforms to each user device and concurrently receive independent UL waveforms from each user device in the same frequency band.

The overhead for UL sounding signals, even with a very large number of concurrent users, is small. For example, if there are 32 concurrent users with overlapping transmissions, the LTE channel sounding overhead is under 3%.

6.3 pCell clusters

Since base station and user transmissions have a limited range, conventional wireless systems partition the coverage area into non-interfering local zones so that the spectrum can be concurrently reused in different local zones. In cellular systems, the local zones are called cells. In Wi-Fi, local zones are the coverage areas of the access points that dynamically schedule transmissions at different times to avoid interference.

pCell also subdivides the coverage area into local zones. But, because pCell exploits RF interference, these local zones deliberately interfere with each other, overlapping in space, frequency and time. These overlapping local zones are called “pCell clusters”.

Conventional wireless systems rely on *base station-centric* wireless architectures: each user has a physical link to a *single* base station[183] until handed-off to establish a new physical link with another base station. All users attached to a base station share the data capacity of that base station.

pCell is a *user-centric* wireless architecture, where each user has a continuous physical link to a pCell cluster of pCell antennas (Artemis Hub and/or pWave RRH antennas), concurrently with all other users. A pCell cluster is the group of antennas that are within the transmission range of a given user for one DL or UL transmission interval (e.g. 1 ms in the case of LTE). As the user moves or the RF environment changes (e.g. there is an obstruction to an RF path), different antennas drop out of range or come into range of the user, and the pCell cluster adapts dynamically[184].

Figure 36 illustrates pCell clusters. The black dots in Figure 36 represent pCell antennas[185]. In the left diagram, the red dot represents one user. The red shaded region shows that user's transmission range, which reaches the seven antennas that are in the red shaded region[186]. These seven antennas form the "pCell cluster" for that user at this particular DL or UL transmission interval.

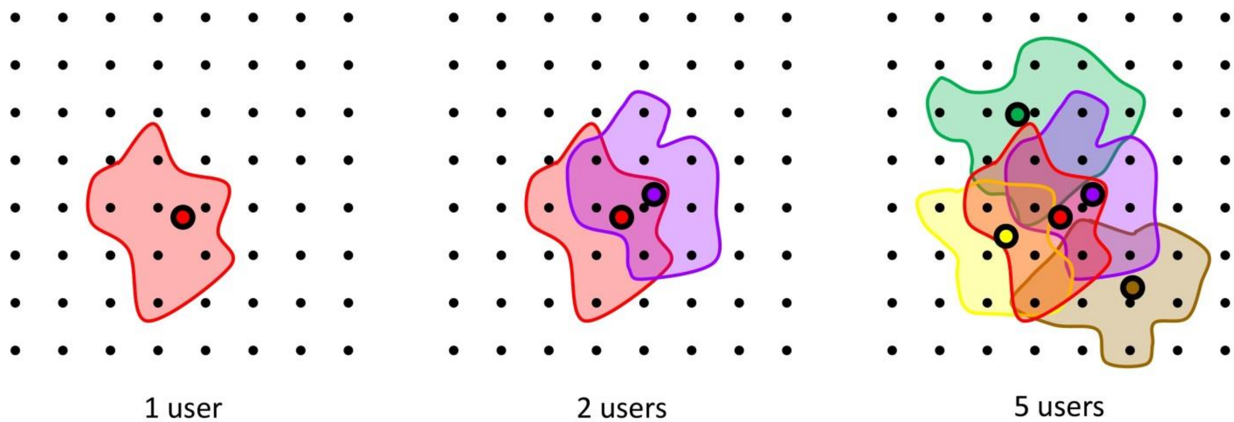

Figure 36: pCell clusters, uniform antenna pattern

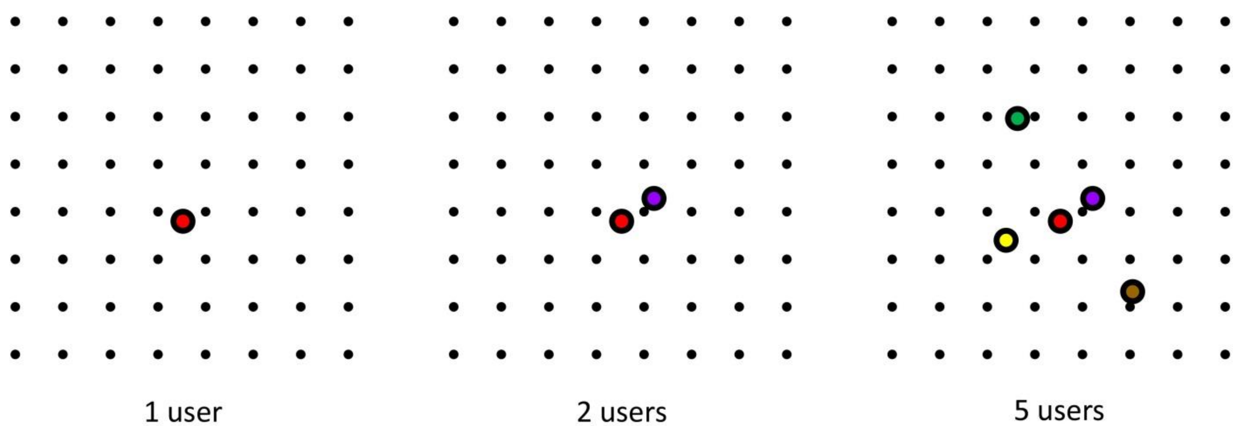

Figure 37: pCells formed by pCell clusters

For example, to transmit DL data, each of the seven antennas in the shaded red region would transmit waveforms that would overlap the red user. The sum of these seven waveforms at the location of the red user's antenna would add up to the desired waveform and synthesize a pCell (represented by the black circle around the red dot) at the location of that user, as shown in Figure 37.

The middle and right diagrams of Figure 36 show the pCell clusters for two and five users, respectively. All antennas in every pCell cluster associated to each user contribute to synthesize the pCell in Figure 37 around that user. Note that the waveforms transmitted by the antennas in the overlapping regions of different pCell clusters contribute to the pCells for multiple users.

The right diagram of Figure 36 adds 3 more users for a total of 5. Note that some antennas are part of only one pCell cluster and others are part of 2 or more pCell clusters. The waveforms transmitted by the antennas in each user's shaded areas sum to synthesize the pCell (the black circle) for that user as shown in Figure 36 with the shaded transmission areas, and in Figure 37 showing just the user dots surrounded by pCells, without the shaded transmission areas.

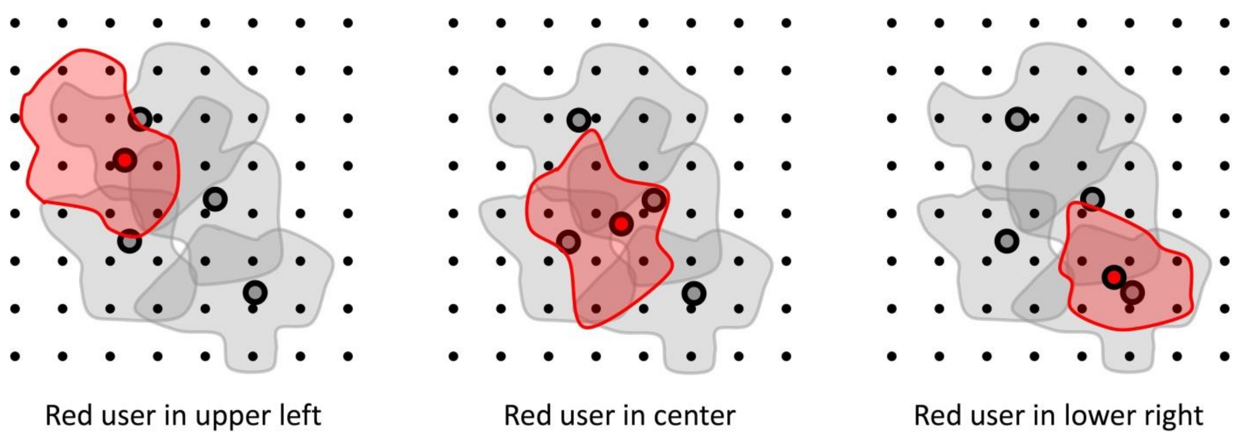

Figure 38: pCell cluster mobility; Red user in motion

Figure 38 shows three consecutive snapshots describing how a pCell cluster adapts when the red user is in motion and all other users (and their propagation environment) are not moving. Note that as the red user moves, the shape of the user cluster changes dynamically. This is due to the fact that the user's RF environment (e.g. obstacles, the device orientation) changes significantly[187], affecting which pCell antennas are within the user's transmission range.

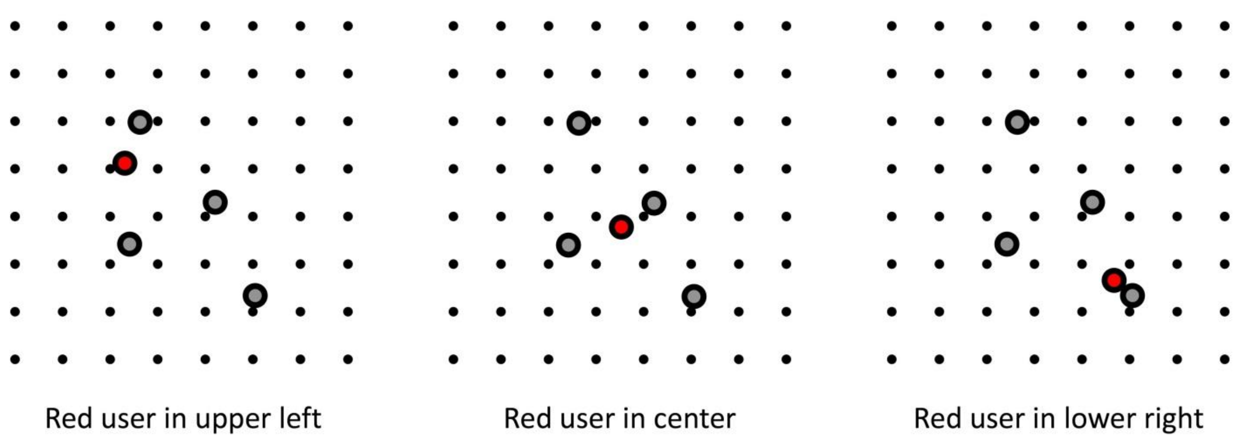

Figure 39: pCells formed by pCell clusters; Red user in motion

Figure 39 shows the pCells formed by the red pCell cluster and the gray pCell clusters around their respective user dots. The red pCell tracks the red user as it moves, while the gray pCells remain in place because their users are stationary. Note that even if the shape of the pCell cluster varies dynamically due to changes in the RF environment, the red pCell still tracks the red user as it moves, and the gray pCells would remain in place around their users.

The previous examples showed a simplified 8x8 grid with equally spaced antennas, for the sake of illustration. But pCell works with any arbitrary 3D layout of antennas, since pCell clusters are formed dynamically based on which antennas are within range of a user. For example, Figure 40 shows the same five users of Figure 36, but with antennas in an arbitrary pattern and more densely placed. Regardless of the antenna pattern or pCell cluster shape, pCells are still consistently synthesized by the antennas around each of the five users, as shown in the right diagram.

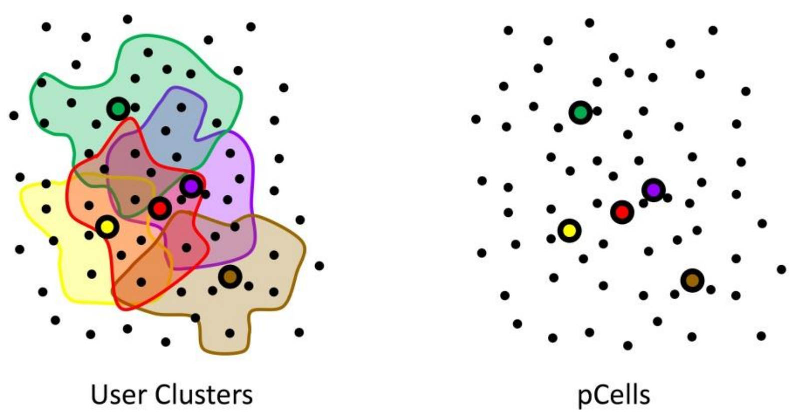

Figure 40: pCell clusters and pCells, arbitrary antenna placement

In all of the above illustrations there are far more antennas in each pCell cluster than necessary. The size of the pCell clusters can be dynamically changed by adjusting transmit power at the users and pCell antennas. Further, as operators add more antennas to increase pCell network capacity, less transmit power is needed to reach the same number of antennas in every cluster. So, as pCell networks scale in capacity, user battery life gets increasingly longer.

A key benefit of cellular technology is, by virtue of hand-offs, to provide largely uninterrupted service for mobile users throughout an arbitrarily large coverage area. But, while cellular maintains a link throughout the coverage area, the data rate varies by a factor of 100 to 1 from cell center to cell edge[188], and the data rate also varies due to network congestion or changes in RF environment resulting in highly variable and unpredictable service quality.

pCell clusters achieve the same goal—of continuously maintaining a link for mobile users throughout an arbitrarily large coverage area (e.g. a large city or a 1000 km-long highway)—but with far more consistent data rate. pCell clusters maintain a high-SINR pCell for each user device constantly, regardless of changes of the RF environment, user density or data demands. From the perspective of the user device, it has a consistent, uncongested, high-SINR connection throughout the coverage area with reliability approaching that of a wireline connection.

pCell clusters can also support scenarios where there is sudden peak demand in a specific location in the coverage area, such as during public events or crisis situations with high data demand from a large number of users that are in close proximity.

In this case, if there are not enough antennas in the immediate vicinity to meet the aggregate data demand of all users, but there are antennas further away that have available capacity, the pCell system can increase transmit power temporarily so that transmissions encompass additional antennas that are further away and increase aggregate capacity. Once peak demand subsides, transmit power can be reduced so as to minimize user device power consumption.

6.4 User mobility

As discussed in Section 6.2, the pCell data center uses channel sounding from every user to every antenna within its pCell cluster to estimate CSI and create one pCell around every user. If the user is stationary, as it is the case for the 80% of mobile data users[189], the CSI does not change much over time and channel sounding can be operated at lower update rates. For the remaining 20% of mobile users in motion (or if the RF environment is changed by motion) the propagation channel between the user and the pCell antennas may vary (due to “Doppler effects”[190]) such that the CSI becomes outdated by the time it is used to transmit/receive DL/UL data. This effect is referred to as “channel aging”[191].

Although Doppler effects are caused by both user motion and environmental motion, user motion accounts for the primary Doppler component causing channel variations. Doppler effects, as perceived by a given user, vary widely depending on the direction of motion of the user relative to the antenna. In cellular systems using MIMO technology, the antennas are placed in one centralized location at the base station, and as such, all antennas experience similar Doppler effect relative to the user motion. But in pCell systems, the antennas are distributed such that the Doppler components can be very different at different antennas for the same user. This extra degree of freedom is exploited by pCell technology to mitigate the effect of channel aging and support high user mobility scenarios.

To achieve higher link robustness in mobility scenarios and minimum data latency (i.e. "ping") to the Internet, clearly it is desirable to have minimum fronthaul latency. Although the LTE protocol itself adds a minimal round-trip latency of about 5-10 ms, pCell native protocols support 0.5 ms round-trip latency or less, assuming the fronthaul latency is itself well below 0.5 ms.

But, in many situations, fronthaul with sub-millisecond latency is not readily available. Although 5 ms or less fronthaul latency is recommended, pCell was designed to accommodate fronthaul latencies as high as 10 ms. In such high-latency scenarios, the few LTE or Wi-Fi messages requiring fast response are handled locally by pWaves or Artemis Hubs.

6.5 3D propagation

Cellular is inherently a 2D architecture, and as discussed in Section 3.4.1, cellular performs poorly in 3D environments like cities with tall buildings.

pCell is inherently a 3D architecture. pCell exploits interference among pCell antennas regardless of their position relative to one another, whether they are located in a horizontal dimension like cellular antennas, in a largely vertical dimension such as antennas installed on multiple floors in a tall building, or in a complex mixture of pCell antennas that are both horizontal and vertical. The only thing that matters is that the transmissions from the antennas overlap with each other.

The fact that pCell is inherently 3D not only makes pCell far simpler to install, it enables pCell to provide highly consistent service even at the top floors of very tall buildings, without compromising service at mid-floors or ground level.

6.6 Location positioning

Wireless user location positioning has become an increasingly important capability of wireless networks, both for public safety and commercial applications. In the wireline telephony era, the location of a wireline phone could be determined to the accuracy of the residential or business address of the wireline installation. As cellular telephony has increasingly displaced wireline infrastructure, location information provided by mobile wireless is increasingly the only location information available to first responders to in emergencies. Further, precise location positioning has many commercial applications where GPS is not available (e.g. indoors), or in multi-story structures, where 3D location information is needed to determine the customer's position.

Based on responses from vendors as to what information precision is commercially achievable, the FCC has recently made recommendations for "E911" emergency service location accuracy requirements of eventually locating indoor mobile callers horizontally within 50 meters with 3 meters vertically 80% of the time.[192] These are challenging goals for cellular systems because user devices are (ideally) served from a single base station location, limiting options for triangulating position. Further, as noted in section 3.4.1, cellular systems have challenges providing mobile service to tall buildings at all, let alone determining the position of a user who has made the call.

While 50 meter horizontal precision is helpful for emergency services (e.g. narrowing position down to one or two buildings), it is far from ideal, which would be identifying a user's precise location within a building at all times. Many commercial applications, such as delivering food to a user's seat location in a stadium, or offering a promotion to a user walking past a store in a shopping mall, require far more precision, and may require more reliability to be commercially viable.

Because pCell users are served by multiple pCell antennas concurrently and pCell is constantly calculating very precise user channel state information about the path from each antenna to each user, the pCell system can constantly determine user position both horizontally and vertically with high accuracy and reliability. As the user moves (e.g. if the user is trying to find an exit from a burning building), the user position will be updated constantly, enabling emergency services personnel to guide the user in real-time. Or, in a commercial application, if a user orders food from a stadium seat and moves to another seat, the food can be delivered to the user's new location. Further, determining such location information requires no modification to existing mobile devices, nor does it require GPS to be operative.

Of course, appropriate privacy policies will have to be put into place regarding the use of location information, much as they are in place today for GPS usage. But, just as GPS in mobile devices has opened up a vast range of potential applications through precise 2D outdoor positioning, pCell opens up a range of applications through precise positioning in 3D, whether indoor or outdoor.

Thus, pCell enables high-accuracy, reliable and continuous user location positioning with no modification to existing devices, both for emergency services and for commercial applications.

6.7 Crisis support

As noted in Section 3.4.2 cellular performs poorly in many crisis situations. pCell was designed to be highly robust, both in the event of damage to pCell infrastructure and in the event of a surge in traffic in affected areas.

Figure 41 shows base stations (or pCell antennas, in the case of pCell), their coverage areas, and "normal" users (i.e. users when the network is under a normal load) for both cellular and pCell. For the sake of illustration, far fewer users are shown. As previously described, cellular avoids transmission overlap while pCell exploits extensive transmission overlap.

The black dots show cellular and pCell base stations and antennas that are knocked out by an explosion (e.g. damaged themselves, or disabled due to damaged infrastructure and/or power source), and the coverage area that was previously served by each knocked-out base station or antenna is shown as a dashed circle. As can be seen in the case of cellular, there is no coverage in the area of the dashed circles since other cellular base stations are not in range. But in the case of pCell, there remains extensive coverage from other overlapping antennas in the area of the dashed circles. Thus, with cellular, there is no service at all when base stations are disabled. With pCell, when antennas are disabled, generally there is full coverage, just reduced capacity.

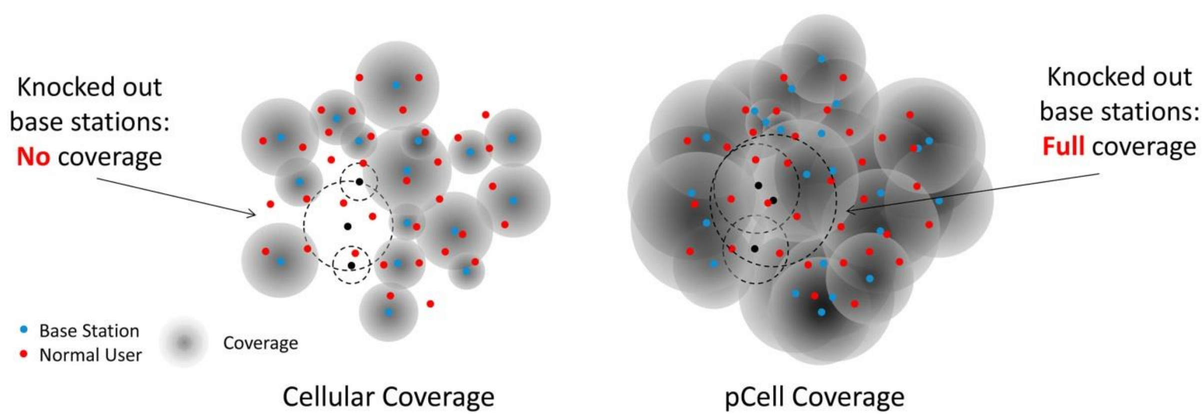

Figure 41: Disabled cellular and pCell base stations and antennas

Frequently, immediately after a crisis situation there is a surge in demand, from users in the crisis area trying to initiate calls, first responders and people who may have heard reports of the crisis trying to call in. Figure 42 illustrates this situation with light green dots representing "surge users", users that suddenly place demands on the network far in excess of the normal user load. In the case of cellular, even if the cellular infrastructure is operational, surge users can overwhelm the cells serving the affected area, resulting in almost no service for anyone, which is what happened after the Boston Marathon bombing. Thus, if the cellular infrastructure has been knocked out, then the surge users will have no service at all[193]. In contrast, with pCell, whether pCell infrastructure has been knocked out or not, there will be service from surrounding antennas. Further, the sudden high demand can be served by utilizing as many surrounding antennas as necessary to provide additional capacity. As noted in Section 6.3, pCell clusters will expand their range in peak demand scenarios, with the only consequence being that user devices may use more transmit power for the brief period of the surge until demand drops down closer to normal levels.

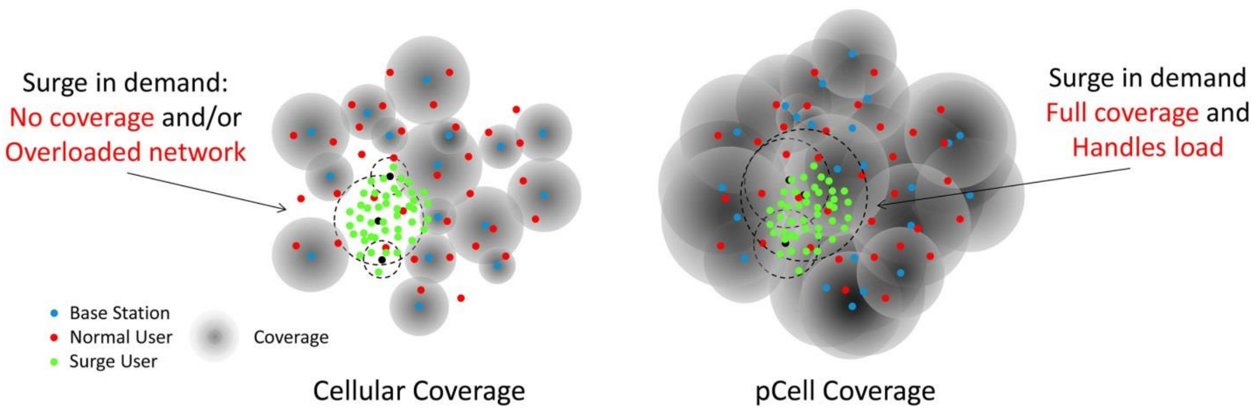

Figure 42: Demand surge handling by cellular and pCell

Thus, pCell is not only robust in providing continuous coverage despite knocked-out infrastructure, but pCell is able to accommodate surges in demand that may occur in the wake of a crisis.

6.8 Physical data security

Conventional wireless systems such a LTE and Wi-Fi utilize key-based cryptography to encrypt data transmissions. Although the encryption systems are strong, if the key is compromised by any means, it is possible to intercept and decrypt data transmissions[194].

pCell data transmissions are *physically* secure: even when over-the-air RF signals or fronthaul transmissions are intercepted, the data in the transmissions cannot be recovered, even if the device intercepting the signal has access to any key used in encrypting the data being transmitted.

The reason for this is illustrated in Figure 43. In general, within a conventional wireless system there is only one transmission occurring at a given time and at a given frequency and it can be intercepted from anywhere within the sector of the cell where the user is located. If the eavesdropper possesses the encryption key that is being used (which is stored in each user's

SIM card as well as the in the wireless carrier's authentication servers), the user communication can be intercepted and decrypted.

With pCell, there are many overlapping transmissions occurring at any given time and the only place in the coverage area where, a distinct downlink transmission can be recovered is within the pCell of the intended user for that transmission. Everywhere else in the coverage area outside of every user's pCell, there is a large number of randomly overlapping waveforms. *A given user's data simply does not exist at any single location in space except within the user's pCell.* In fact, even if a fronthaul connection to a pWave RRH were intercepted, no user's data would be recoverable from the fronthaul.

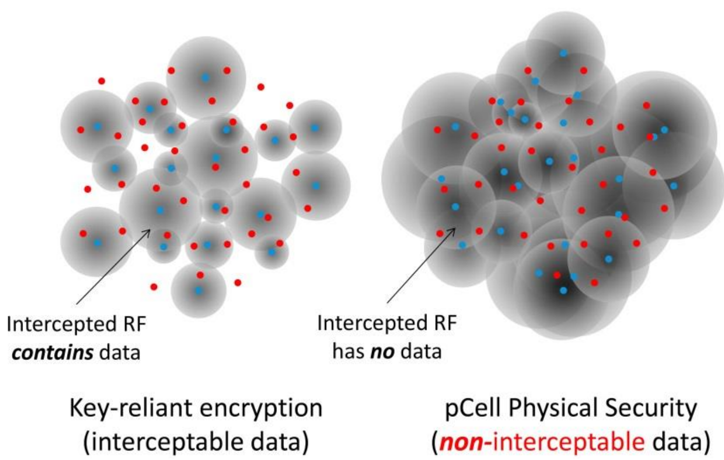

Figure 43: Cellular key-based encryption and pCell physical security

In the UL, if an intercepting device was very close to a user device and there were no other devices nearby, conceivably it could intercept that user device's UL transmission where the signal is strong enough above the noise floor. But, because the DL is physically secure, a constantly-changing key used by the user device for the UL can be sent over the physically secure DL, resulting in secure UL encryption.

Thus, pCell can be used for applications which require physical security. And, in an era when there is heightened awareness of communications security breaches, pCell provides exceptionally high security for consumers and businesses.

7 Future pCell Applications

7.1 256-QAM support

All deployed pCell systems will support 256-QAM once LTE-Advanced Rel. 12[195] devices are deployed that support higher order modulation. As noted in prior sections, because pCell maintains a consistent, high-SINR signal throughout the coverage area, it is able to utilize high-order modulation far more efficiently than cellular systems, which can only achieve the high SINR needed for high-order modulation near the cell center.

Specialized protocols that are native to pCell, such as those described below, all are designed to support 256-QAM from the outset. Also, pCell can dramatically simplify both RF chain and baseband processing designs at the user device relative to that of a standard LTE cellular device (e.g., no need for multiple RF chains, or receive/transmit diversity techniques, or MIMO schemes, etc.). Thus, the reduced complexity of pCell-native user devices can be utilized for more linear RF chain or higher processing power within the same device real estate to support higher-order modulations like 256-QAM.

7.2 pCell compatibility with specialized wireless devices

While pCell dramatically improves aggregate capacity and data rate consistency experienced by standard LTE devices, other parameters essential to certain specialized applications, e.g. sub-millisecond latency, are not achievable within the LTE framework.

New specialized protocols can be defined for new spectrum, but it is inefficient to allocate new spectrum and new infrastructure to support specialized protocols. And, given the dire scarcity of spectrum, it is unlikely that there will be much, if any, spectrum available at mobile frequencies[196] by the time such specialized protocols are standardized and introduced.

pCell overcomes this limitation by concurrently supporting multiple protocols—standardized and specialized—in the same spectrum. As previously discussed, pCell supports both 64-QAM legacy LTE devices and 256-QAM LTE-Advanced devices concurrently in the same spectrum, but pCell can concurrently support new protocols using non-LTE protocols in the same spectrum by simply instantiating VRIs for new protocols that operate concurrently with LTE protocols, as described in Section 5.3.2. The same pCell network infrastructure and pCell antennas that support LTE/LTE-Advanced protocols would concurrently support these new protocols in the same spectrum. Thus, devices for specialized applications can be supported efficiently and inexpensively.

7.2.1 pCell VR

pCell VR is a sub-millisecond, low UL-power protocol designed to support the extreme performance demands of virtual reality (VR) and augmented reality (AR) applications (referred collectively below as "VR/AR"), as well as other low-latency computing or low power devices.

Recently announced VR/AR systems on the market or in development include Oculus VR's Oculus Rift [197], Microsoft's HoloLens[198], Samsung Gear VR[199], Google Cardboard[200] and Magic Leap's light field technology[201]. All of these technologies deliver extremely high-resolution, low-latency visual experiences to the user, potentially as a pre-recorded 3D stereoscopic video experience (e.g. Jaunt's cinematic VR technology[202]) or as fully interactive 3D experiences synthesized in real time, many of which are in early stages of development[203].

VR/AR experiences require extremely low-latency responsiveness to create the illusion of reality and, in the case of AR, to keep the AR image locked to the real-world view. Also, low latency is essential to prevent "simulator sickness" (a seasick feeling caused by lag in responsiveness). Total response time, from the moment a user's head moves to the time a new image is presented to the user ideally is on the order 7 ms or less.[204] State-of-the-art motion sensors, high-performance GPUs in high-end computers, and high-resolution display systems are hard-pressed to deliver a realistic 3D world with high detail in only 7 ms. In fact, simply delivering pre-recorded "cinematic VR" is challenging, given the amount of data that must be transferred from storage to the VR system's screen or light field display.

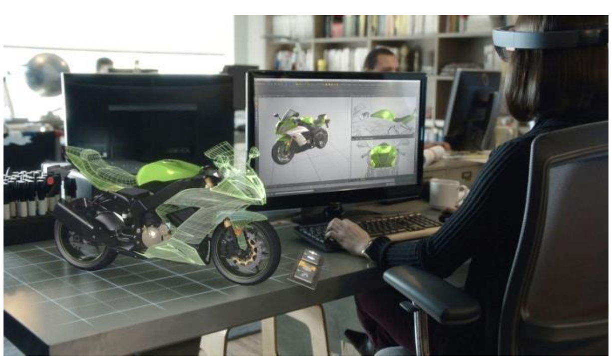

While there is no doubt these challenges can be overcome, delivering realistic, low-latency experiences requires exceptionally powerful computing resources for each user, far beyond the performance of current video game consoles, let alone those of lightweight and low-power electronics that could readily fit within a wearable headset. Such powerful computers are typically expensive, physically large, and have noisy fans for cooling. And, given the rapid rate of evolution of VR/AR technology such systems will likely be subject to frequent hardware and software updates.

pCell VR is a protocol that enables VR/AR systems to place such powerful computers in pCell Cloud-RAN data centers so that the only local capability needed in the VR/AR system are input systems (e.g. head or hand tracking and voice input) and output systems, e.g. display (visor, light field, etc.), audio (e.g. earphones, 3D sound), and haptic feedback[205] (e.g. vibration, tactile feedback).

For example, consider a VR/AR visor using pCell VR:

The VR visor would have a pCell VR radio built into it with 1 or more antennas, depending on desired DL data rate. The pCell VR radio has a round-trip latency of 500 microseconds (0.5 ms) and supports highly asymmetric TDD. For example, the UL data rate in each 500-microsecond (µsec) frame time can be extremely small compared to the DL data rate, even 100:1.

At the start of the VR/AR session, a high-performance graphics-capable server (a “VR server”) is allocated for the user in the pCell C-RAN data center. A VR server is capable of generating high-resolution stereoscopic 3D scenes at very low latency, e.g. 3 ms or less.

The pCell VR protocol stack is simple and shallow, with a hard-real-time, pipelined path to and from VR servers, minimizing network latency between a VR server and the pCell network to 10s of microseconds, but also supporting a real-time, pipelined flow of data.

An example round-trip scenario would start with the VR/AR visor’s input system detecting head motion. The VR/AR visor would transmit a very small packet containing the head motion data through pCell VR UL, which would immediately route it to an allocated VR server. The VR/AR server would render a new stereoscopic 3D image in 3 ms, pipelining the frame out of the GPU and routing it to the pCell VR DL, which would transmit the pipelined stereoscopic image to the VR/AR visor, which would then display the image to the user. Assuming fast head tracking and fast display refresh, the entire round trip from head motion to updated image would be less than 7 ms anywhere in the pCell coverage area (e.g. even in a moving car). And, the VR/AR visor would consume a minimal amount of power in the process.

The network overhead of pCell VR wireless and Cloud-RAN data center network would be only 500 µsec of latency to the entire round trip, resulting in VR server latency of 3.5 ms instead of 3 ms. This is only slightly higher latency than a fast wired network connection from the VR visor to a local computer, but the user would have the benefit of an exceptionally powerful, always up-to-date, VR server that would be impractical for most users to have in their home, and certainly impractical for mobile use. Beyond that, the VR server would be shared with others, e.g. using it at different times in the day or different days of the week, dramatically reducing the cost per user.

pCell VR minimizes cost and power consumption in local devices (whether VR/AR visors or other devices), despite potentially multi-gigabit DL data rates. In the above round-trip example, only low data rate input data, such as head tracking data, is transmitted, requiring very brief transmissions and consuming minimal UL power. The DL data rate scales with the available bandwidth and number of device antennas. 1 antenna can deliver up to roughly 100 Mbps in 20 MHz, 200 Mbps in 40 MHz, or 2 antennas can deliver 200 Mbps in 20 MHz. Unlike MIMO systems, pCell data rate scales linearly with the number of antennas with practical device spacing. As an extreme—but achievable—example, if 10 antennas are distributed around the perimeter of a VR/AR visor, in about 200 MHz of spectrum[206], roughly 1 Gbps * 10 = 10 Gbps would be delivered with 500 µsec latency to reach a C-RAN-based VR server, anywhere in the pCell coverage area, effectively delivering the equivalent of a 10 Gbps fiber connection to the AR/VR visor from a data center-class visualization server. This performance would be achieved while the spectrum is concurrently shared with other users, including LTE and LTE-A devices.

While multi-antenna pCell devices would have multiple RF receive chains, the added cost would be substantially less than that of current multi-antenna devices. In a pCell network, each device antenna is allocated its own pCell in which it receives an independent, high-SINR waveform, requiring only a straightforward RF transceiver per antenna, not complex multi-antenna transceivers as in MIMO devices, which are expensive and consume a great deal of power at very high orders.

Thus, pCell VR enables real-time VR/AR, with low-latency, high DL data rate, with minimal RF power consumption in the VR/AR device, concurrently sharing spectrum with LTE devices.

While the most demanding application currently envisioned for pCell VR is its namesake, VR, pCell VR enables a wide range of other applications. 500 µsec roundtrip latency to C-RAN-based resources enables any application requiring fast-response cloud computing or storage, and minimal network overhead for Internet access outside of the C-RAN. While the VR/AR example above is highly asymmetric, for applications requiring high UL data rates, pCell VR can be configured with more symmetric UL.

In fact, pCell VR can be thought of as not simply a communication protocol, but as an always-available hard-real-time local computing resource available to any device in the pCell coverage area.

7.2.2 pCell IoT

While pCell VR is capable of delivering low-latency, high data rate services, the same low device power attributes, implemented with lower data rates, result in a protocol that is highly suitable for devices requiring minimal connectivity, such as Internet of Things (IoT) devices from wearables to LED bulbs. This protocol is called pCell IoT

pCell IoT's attach procedure is far simpler than LTEs, enabling brief, minimal-power transmission for quick "hit and run" connections. And, like all pCell protocols, pCell IoT maintains high, uniform SE while the user or device is in the pCell coverage area, even if a pCell IoT device has only one brief uplink and the pCell network has only one brief downlink.

pCell IoT supports millions of small devices in the pCell coverage area, all sharing the same spectrum as LTE, pCell VR and other specialized protocols that can now co-exist within a pCell network.

8 Acknowledgments

pCell has been a journey of over a decade from its first conception to the announcement of its first commercial deployment, coinciding with the release of this white paper. Despite the fact that enormous gains and advantages of pCell have been demonstrable during most of its development, it was such a huge leapfrog beyond conventional wireless that few people who saw pCell working recognized that what appeared to be “too good to be true” not only was very real, it wasn’t a research project; it was practical as a commercial offering.

We, Steve Perlman and Antonio Forenza, who each have been with the project for a decade or more, would like to both thank and congratulate our team and friends of our team, who stuck through every challenge with pCell, creating a technology that stands to transform society, far beyond what we can imagine today:

Roger van der Laan, Cindy Ievers, Fadi Saibi, Mario Di Dio, Allan Ievers, Ben Jirasut, Lynne Freeman, Eric Peltier, Jane Anderson, Kelly Tunstall, Jonathan Carta, Gabi Schindler.

9 Endnotes

[1] Spectral efficiency measures the amount of bits per second that can be transferred over a wireless link per unit of electromagnetic spectrum (e.g. Mbps per MHz of spectrum). For example, one technology is more spectrally efficient than another if it can wirelessly transfer data at higher bit rate within the same amount of spectrum. The calculation of all LTE spectral efficiency numbers published throughout this white paper are calibrated against Tables 16.1 in: [*3GPP TR-36.912*] 3GPP, "Feasibility study for further advancements for E-UTRA (LTE-Advanced)", TR 36.912, v 12.0.0, (36912-c00.doc) Sep. 2014, http://www.3gpp.org/ftp/Specs/archive/36_series/36.912/36912-c00.zip

[2] [*Rysavy 2014*] Rysavy Research, "Beyond LTE: Enabling the Mobile Broadband Explosion", Aug. 2014 http://www.rysavy.com/Articles/2014-08-4G-Americas-Mobile-Broadband-Explosion.pdf, Fig. 29, p. 71.

[3] SON: Self-organizing networks. Described in Section 3.3.1.

[4] [*Cellular security*] E.g., A. Mizroch, L. Fleisher,"Digital-Security Firm Gemalto Probes Alleged U.S., U.K. Hack", Fortune, Feb. 20, 2015, http://www.wsj.com/articles/dutch-firm-gemalto-investigates-hacking-claim-1424423264, R.B. Reilly, "Mystery cell towers host mobile ID catchers that stealthily intercept calls", VentureBeat, Sep.18 2014, http://venturebeat.com/2014/09/18/the-cell-tower-mystery-gripping-america-has-now-been-solved-or-has-it/

[5] [*FCC E911 2014*] "Wireless E911 Location Accuracy Requirements", FCC, PS Docket No. 07-114, February 20, 20114. http://www.fcc.gov/document/proposes-new-indoor-requirements-and-revisions-existing-e911-rules

[6] These projections proved to be quite accurate. Excluding managed IP traffic, most global IP traffic was wireless in 2013, with Wi-Fi at 55% and mobile at 4%. Mobile is growing at 61% CAGR and Wi-Fi at 25% CAGR, so that by 2018, mobile will be 15% of global IP traffic with Wi-Fi at 61% of IP traffic. Video became the majority of global mobile traffic starting in 2012 and is projected to grow at 69% CAGR, to reach 69.1% of global mobile traffic by 2018. By 2018 mobile video traffic alone will be 8X as much as *all* mobile traffic in 2013. [*Cisco VNI 2014*] Cisco, "Visual networking index: global mobile data traffic forecast update, 2013-2018", Feb. 2014

[7] [*Mobile Bands 2014*] The global spectrum allocated for mobile systems ranges between 700 MHz and 3800 MHz, of which about 2000 MHz of spectrum is mapped for LTE deployments worldwide (3GPP TS 36.101, "E-UTRA: UE radio transmission and reception", Release 11, version 12.4.0, document 36101-c40_s00-07, Oct. 2013, http://www.3gpp.org/DynaReport/36101.htm. Additionally, in 2014 the FCC approved the 600 MHz band for LTE deployment in the U.S.

[8] FCC, "Spectrum crunch", http://www.fcc.gov/encyclopedia/spectrum-crunch. While the existence of the spectrum crunch had been debated in past years, actual declines in mobile data rates in the U.S. and practical limitations of reducing cell size, as described in this document, confirms it. Still, a recent paper (draft), A. Mehta and J.A. Musey, "Overestimating wireless demand: policy and investment implications of upward bias in mobile data forecasts", Social Science Research Network, Aug. 2014 *http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2418364*, questions whether mobile data forecast overestimations, such as the highly-cited Cisco's Visual Networking Index (VNI), are biased and whether its future projections should be trusted for policy decisions regarding the spectrum crunch. We repeated Mehta, et al.'s analysis and compared each year of Cisco VNI's mobile data projections against actuals, starting from when Cisco began making projections in 2008, and found that, at most, the projections were off by about one year by the end of their long-term forecasts. For example, the Cisco 2008 projection for 2012 mobile data rates was actually reached in 2013. In successive years after 2008, Cisco appropriately lowered its 2012 forecast. For example, Cisco's 2009 forecast for 2012 was almost spot-on (within one quarter-year). One-year offsets in long-term predictions do not materially impact either public policy analysis or the analysis in this white paper.

[9] [*FCC 2010*] FCC, "Mobile broadband: the benefits of additional spectrum", FCC Staff Technical Paper, Oct. 2010, http://transition.fcc.gov/national-broadband-plan/mobile-broadband-paper.pdf, Exhibit 10.

[10] Statista, "Ranking of countries/territories by LTE mobile subscribers 2013 (in millions)", 2014 http://www.statista.com/statistics/309599/lte-mobile-subscribers-by-country/

[11] [*Cisco VNI 2014*] Table 7
[12] [*Cisco VNI 2014*]
[13] Verizon reported in April 2013 that 50% of mobile traffic was video, growing to 67% by 2017. http://www.fiercewireless.com/story/verizon-ceo-50-our-wireless-traffic-video/2013-04-10
[14] *[OpenSignal 2014]* S. Westwood, "Global state of LTE report", OpenSignal, Feb. 2014. http://opensignal.com/assets/pdf/reports/2014_02_opensignal-state-of-lte-report.pdf
[15] HSPA supports data speed of 7.2 Mbps in the U.S. http://techcrunch.com/2010/01/05/att-hspa-72-iphon/
[16] *[OpenSignal 2014]* and A.C. Nichols, "We anticipate the iPhone will slow DoCoMo's subscriber losses", Jun. 2014 http://analysisreport.morningstar.com/stock/research?t=DCM&region=usa&culture=en-US&productcode=MLE
[17] R. Cheng, "Verizon admits network faces traffic pressure in big cities," CNET, Nov. 2013 http://www.cnet.com/news/verizon-admits-network-faces-traffic-pressure-in-big-cities/ Verizon at Wells Fargo Technology, Media & Telecom Conference, Nov. 12, 2013 (edited transcript): http://www.verizon.com/about/file/1317/download?token=EHI45LJk
[18] K. Fitchard, "Verizon quietly unleashes its LTE monster, tripling 4G capacity in major cities," Gigaom, Dec. 2013. https://gigaom.com/2013/12/05/verizon-quietly-unleashes-its-lte-monster-tripling-4g-capacity-in-major-cities/
[19] [*Cisco VNI* 2014] and [*FCC 2010*]
[20] Verizon average LTE speed fell from 9.5 Mbps to 8.1 Mbps from 1Q13 to 1Q14, OpenSignal, "The state of US LTE", Mar. 2014, http://opensignal.com/reports/state-of-lte/usa-q1-2014/
[21] R. Cheng, "Verizon CEO: unlimited data plans just aren't sustainable", CNET, Sep. 2013 http://news.cnet.com/8301-1035_3-57604368-94/verizon-ceo-unlimited-data-plans-just-arent-sustainable/ Verizon at Goldman Sachs Communcacopia Conference, Sep. 24, 2013 (edited transcript): http://www.verizon.com/about/file/1315/download?token=f30Bc66s
[22] [*Rysavy 2014*] Fig. 19, p. 41.
[23] [*Mobile Bands 2014*]
[24] 3200 MHz of spectrum (between ~600-3800MHz) is suitable for mobile (cfr. [*Mobile Bands 2014*]) of which 547 MHz was available for cellular as of 2010 (cfr., FCC, "Connecting America: the national broadband plan", p.84 http://download.broadband.gov/plan/national-broadband-plan.pdf). Assuming all other spectrum users are displaced, at most 4.9 times the 547 MHz remains in the 3200 MHz, i.e. (3200-547)/547 = 4.9. Since wireless traffic is projected to grow by 5.2X in the three-year period from 2013-2016 (cfr. [*Cisco VNI 2014*]), this is sufficient to cover at most three years of traffic growth.
[25] Townsend, "Smartphones to monitor insulin and smell flowers: the wireless industry will be transformed by 2023—if it can overcome a lack of spectrum", The Wall Street Journal, Oct. 28, 2013, http://online.wsj.com/news/articles/SB10001424052702304520704579129730041631274
[26] The AWS-3 auction was for 1700 MHz + 2100 MHz "mid-band" spectrum, which has poorer propagation than "low-band" 600-900 MHz UHF spectrum, but better propagation than "high-band" 2500-3500 MHz spectrum. Analyst projections ranged from US$11 to 22 billion. P. Goldstein, "AWS-3 spectrum auction primer: What you need to know before the bidding starts", Nov. 12, 2014. http://www.fiercewireless.com/special-reports/aws-3-spectrum-auction-primer-what-you-need-know-bidding-starts
[27] "…the 700 MHz [FCC] auction in 2008 raised $18.9 billion…" *Ib.*
[28] "Consider that, in nearly two decades, [the FCC's] spectrum auctions have raised over $50 billion…" FCC Commissioner J. Rosenworcel, testimony to the U.S. Senate Appropriations Committee, Sep. 2013, p. 9. http://www.appropriations.senate.gov/sites/default/files/hearings/09_11_13%20FSGG%20FCC%20Budget%20GPO%20Record_0.pdf
[29] http://dynallc.com/martin-cooper/
[30] http://www.fcc.gov/leadership/jessica-rosenworcel
[31] http://www.mercurynews.com/opinion/ci_26597034/marty-cooper-and-jessica-rosenworcel-heres-how-expand
[32] B. Bertenyi, "3GPP system standards heading into the 5G era", Spring 2014 http://www.3gpp.org/news-events/3gpp-news/1614-sa_5g

[33] E.g., the EU Commission has funded and co-funded a variety of projects to promote 5G research and development, including: €700M to the 5G Public and Private Partnership (5G-PPP, with members including Alcatel-Lucent, Ericsson, NSN, Orange and SES, http://5g-ppp.eu/contract/) through the Horizon 2020 program (http://ec.europa.eu/digital-agenda/en/towards-5g), $77M to universities in UK and Germany (http://au.ibtimes.com/articles/542541/20140310/uk-germany-5g-mobile-network-funding-internet.htm#.U_-EXBZOiRN), €16M to METIS 2020 coordinated by Ericsson (https://www.metis2020.com), €3M to CROWD coordinated by Intecs (http://www.ict-crowd.eu), €2.5M to 5GNOW (http://www.5gnow.eu). In South Korea, the Ministry of Education, Science and Technology (MEST) will be investing $1.5B to fund 5G research and development (http://www.ibtimes.com/south-korea-will-invest-15b-building-5g-network-will-be-1k-times-faster-existing-speeds-allow-1), while Huawei plans to invest €444M for a similar goal (http://www.huawei.eu/events/huaweis-5geurope-summit-2014).

[34] For example, Samsung recently demonstrated 5G tests that achieve peak data rates up to 7.5Gbps using millimeter waves operating at 28 GHz in a controlled environment. Y. Kang, "Samsung 5G test hits speeds 30 times faster than LTE", Oct. 15, 2014 http://www.pcworld.com/article/2834072/samsung-claims-data-speed-test-30-times-faster-than-lte.html But the performance of that testbed in practical urban environments or through practical outdoor-to-indoor propagation (real-world adverse conditions for millimeter waves due to propagation loss) has not yet been demonstrated.

[35] [*IMT-2020*] IMT-2020 (5G) Promotion Group, "IMT Vision towards 2020 and Beyond", Feb. 2014. http://www.itu.int/dms_pub/itu-r/oth/0a/06/R0A0600005D0001PDFE.pdf

[36] Other visions for 5G performance include: [*NSN 2014*], Nokia Siemens Networks "5G use case and requirements", Apr. 2014 (http://networks.nokia.com/sites/default/files/document/5g_requirements_white_paper.pdf) and [*METIS 2014*] A. Osseiran, et al., "Scenarios for the 5G mobile and wireless communications: the vision of the METIS project", May 2014 (https://www.metis2020.com/wp-content/uploads/publications/IEEEComMag_Osseiran_et_al_METIS_overview_scenarios_201405.pdf) and [*GSMA 5G*] GSMA, "Understanding 5G: Perspectives on future technological advancements in mobile", Dec. 2014 (https://gsmaintelligence.com/files/analysis/?file=141208-5g.pdf). All visions have similar goals to those of IMT-2020 (indeed, GSMA references IMT-2020). Notably, all propose a cell center data rate range of 10 Gbps. IMT-2020, NSN, and METIS assume a 100:1 drop off to cell edge (consistent with known cellular technology), for a 100 Mbps minimum data rate, while GSMA implicitly assumes new technology will be developed that improves cell edge performance by 10X to maintain a minimum 1 Gbps data rate. While all three recognize the requirement for much higher traffic density, NSN projects these density gains will be achieved specifically by large transmission bandwidths at high frequencies with limited propagation range (see page 13 in the NSN paper), while IMT-2020, METIS and GSMA are open to exploring many approaches to achieve higher traffic density. All recognize the need for low latency (<1ms for NSN, IMT-2020, and GSMA; <5ms for METIS), for low power device support, and wireline (99.999%) reliability.

[37] 45 Mbps/Hz spectral efficiency is equivalent to 15X the IMT-Advanced target average spectral efficiency of 3.0 bps/Hz [*ITU 2008*] Report ITU-R M.2134, "Requirements related to technical performance for IMT-Advanced radio interface(s)", 2008, http://www.itu.int/dms_pub/itu-r/opb/rep/R-REP-M.2134-2008-PDF-E.pdf

[38] Contrary to other 5G proposals, [*GSMA 5G*] proposes a 1 Gbps minimum per-user data rate and a 10 Gbps peak data rate in a cellular architecture, which implies a 10:1 drop off from cell center to cell edge. Known cellular technologies incur a 100:1 drop off from cell center to cell edge (e.g. see Figure 3 and its [*Ericsson 2013*] citation) and [*GSMA 5G*] does not detail how this 100:1 drop off would be improved to 10:1. We are unaware of any technology in development that achieves this 10X increase in cell edge efficiency. Assuming 100:1 drop off, 1 Gbps at cell edge would require 100 Gbps at cell center, which is 10X beyond the performance of any 5G proposal, including GSMA. As such, we've listed 100 Mbps as the minimum user data rate for 5G.

[39] Internet of Things (IoT) refers to any devices able to utilize a connection to the Internet. pCell can support IoT devices (e.g., LED "Smartbulbs", wearables, etc.).

[40] 3 bps/Hz requires LTE-Advanced MIMO 8x8 capability [*ITU 2008*], both in the network and devices. To the best of our knowledge, this has not yet been deployed in any major market. Current LTE average spectral efficiency is up to 1.7 bps/Hz for current 2-antenna devices, and up to 2.4 bps/Hz for future 4-antenna devices [*Rysavy 2014*], Fig. 29, p. 71. IMT 5G's 45 bps/Hz spectral efficiency is a 26X leapfrog beyond today's 1.7 bps/Hz.

[41] In 40 MHz of spectrum with devices capable of carrier aggregation of two 20 MHz bands.

[42] E.g., the minimum LTE latency is several milliseconds and many aspects of the LTE standard (e.g. the requirement of 2 antennas, a complex attach protocol) preclude minimal cost and power devices.

[43] [*Rappaport 2013*] T. Rappaport, et al. "Millimeter wave mobile communications for 5G cellular: it will work!", IEEE access, vol.1, pp.335-349, May 2013. http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=6515173.

[44] [*Larsson 2014*] E. Larsson, O. Edfors, and T.L. Marzetta, "Massive MIMO for next generation wireless systems", IEEE Comm. Magazine, pp.186-195, Feb. 2014.

[45] There were 3.4 billion cellular subscribers in 2013 (cfr., GSMA, "The mobile economy 2014" http://www.gsmamobileeconomy.com/GSMA_ME_Report_2014_R_NewCover.pdf). The world population is about 7.2 billion.

[46] Video streaming today accounts for more than 50% of mobile data traffic [*Cisco VNI 2014*]

[47] Signal-to-interference-plus-noise ratio (SINR), indicates signal quality and determines achievable data rate.

[48] [*Ericsson 2013*] "Ericsson mobility report. On the pulse of the networked society", Jun. 2013, Fig. 21. http://www.ericsson.com/res/docs/2013/ericsson-mobility-report-june-2013.pdf

[49] For example, if the LTE base station could support 100 Mbps shared among 10 users, every user at the cell center would experience on average a 10 Mbps connection speed.

[50] [*LTE simulator*] The SINR and data rate results are obtained from Artemis' LTE system level simulator using a carrier frequency of 1910 MHz, system bandwidth of 18 MHz (assuming LTE spectrum allocation for 20 MHz) and universal frequency re-use 1 (i.e., all eNodeBs operate at the same carrier frequency as in practical LTE networks). The eNodeB antenna radiation patterns are omnidirectional. The network geometry assumes a regular hexagonal grid with fixed inter-eNodeB distance of 100 meters. "A hexagonal cell, having equal distance with all adjacent cells therefore provides the best coverage model with the least number of required base stations to serve a particular area.", cfr., p. 4 in:[*Agbinya 2013*] *4G Wireless Communication Networks: Design Planning and Applications*, edited by Johnson I. Agbinya, J. I., Aguayo-Torres, M.C., Klempous, R., 2013, River Publishers. The cell layout consists of one eNodeB surrounded by two tiers of eNodeBs (the first tier with 6 eNodeBs and the second tier with 12 eNodeBs) adding up to total of 19 eNodeBs (though only 11 eNodeBs are shown in the for the sake of illustration). Every user device (UE) is equipped with only one antenna, such that SISO links are established between eNodeBs and UEs. While we recognize that MIMO technology can be used to increase peak data rate of cellular systems (as described in a later section of this white paper), the purpose of these graphical results is only to illustrate the exponentially decaying performance of cellular and its limits due to sharing bandwidth among users in the same cell, which can be explained even with single antenna devices. The average transmit power per eNodeB is 200 mW. The channel is generated as uncorrelated Rayleigh fading process with pathloss and shadowing from the LTE urban microcell NLOS model described in:[*3GPP-SCM*] 3GPP, "Spatial channel model for MIMO simulations", TR 25.996, v12.0.0, Sep. 2014, http://www.3gpp.org/ftp/Specs/archive/25_series/25.996/25996-c00.zip To simplify the graphics of the heat maps, for the case of cellular, it is assumed there is no shadowing. Based on this assumption, the results for cellular are more optimistic than in real-world scenarios. For example, our simulations show only 10:1 cell center-to-edge data rate ratio (rather than 100:1 as in real world scenarios) because of lack of shadowing. The data rate is computed at the PHY layer using the SE values of the LTE MCSs reported in the CQI tables 7.2.3-1 and 7.2.3-2 in: [*3GPP TS-36.213*] 3GPP, "E-UTRA: physical layer procedures", TS 36.213, v.12.4.0, Rel.12, Dec. 2014 http://www.3gpp.org/ftp/Specs/archive/36_series/36.213/36213-c40.zip The data rate is adjusted to account for 26% loss due to the TD-LTE configuration 3:1 (i.e., TDD frame configuration #2 with S-subframe configuration #7) and additional 24% loss due to LTE PHY layer overhead. Link adaptation is employed to switch between different MCSs depending on SINR performance. The link adaptation switching thresholds are the SINR values for which the block error rate (BLER) is at 10% in AWGN channels, as reported in Figure 2 at: Nokia, NSN, "On CQI/MCS/TBS table design for 256QAM", 3GPP TSG-RAN WG1 Meeting #76, R1-140555, Feb.2014

http://www.3gpp.org/ftp/tsg_ran/WG1_RL1/TSGR1_76/Docs/R1-140555.zip

[51] 56 Mbps is the peak data rate for SISO TD-LTE in 20MHz with a 3:1 DL:UL ratio (i.e., TDD frame configuration #2) and the S-subframe configuration #7).

[52] [*Cisco VNI 2014*]

[53] High-definition (HD) video is about 5X the data rate of standard-definition (SD) video, so even at the sacrifice of quality, video data rate can only adapt to declining data rate by about 5:1. From cell-center to cell-edge the data rate declines by 100:1, resulting in 1/100th the capacity at cell-edge than cell-center. Since, at best, video can adapt to 1/5th data rate, then at a minimum, video will consume 20X more cell capacity at cell-edge than at cell-center. Note that some video streams are non-adaptive; In such cases, video will consume 100X more cell capacity at cell-edge than at cell-center.

[54] E.g., [*IMT-2020*], [*NSN 2014*] and [*METIS 2014*]

[55] [*Micro-diversity*] is one type of space diversity obtained when multiple antennas are placed at a relative distance shorter or in the order of the wavelength.

[56] K. F. Braun, "Electrical oscillations and wireless telegraphy," Nobel lecture, Dec. 1909

[57] J.H. Winters, "Beamforming Techniques," in *Space-Time Wireless Systems: From Array Processing to MIMO Communications*, Cambridge University Press, 2006, ch2. pp. 23-43.

[58] T. Kailath and A.J. Paulraj, "Increasing capacity in wireless broadcast systems using distributed transmission/directional reception (DTDR)", U.S. Patent 5,345,599, Feb. 1992

[59] Although the term "MIMO" can be used to refer to any multi-antenna scheme, including beam-forming and diversity schemes, in accordance with usage within the LTE and Wi-Fi standards, "MIMO" is used in this white paper to only refer to schemes involving spatial multiplexing.

[60] A. Paulraj, R. Nabar, and D. Gore, Introduction to Space-Time Wireless Communications, Cambridge University Press, 40 West 20th Street, New York, NY, USA, 2003

D. Gesbert, M. Shafi, D. Shiu, P.J. Smith and A. Naguib, "From theory to practice: an overview of MIMO space-time coded wireless systems", IEEE Journal on Selected Areas on Communications, vol.2, n.3, pp.281-302, Apr. 2003

L. Zheng and D. N. C. Tse, "Diversity and multiplexing: a fundamental tradeoff in multiple-antenna channels," IEEE Trans. Info. Th., vol. 49, no. 5, pp. 1073–1096, May 2003

[61] [*Forenza 2006*] A. Forenza, "Antenna and algorithm design in MIMO communication systems exploiting the spatial selectivity of wireless channels", PhD dissertation, May 2006

[62] [*Rysavy 2014*], Fig. 29, p. 71.

[63] G. J. Foschini and M. J. Gans, "On limits of wireless communications in a fading environment when using multiple antennas", Wireless Personal Commun., Kluwer Academic Press, no. 6, pp. 311-335, 1998.

E. Telatar, "Capacity of multi-antenna Gaussian channels", Eur. Trans. Tel., vol. 10, no. 6, pp. 585-596, Nov. 1999.

[64] "MIMO can be used when S/N (Signal to Noise ratio) is high, i.e. high quality radio channel. For situations with low S/N it is better to use other types of multi-antenna techniques to instead improve the S/N, e.g. by means of TX-diversity, see figure 5." Wannstrom, Jennifer, for 3GPP, "LTE-Advanced", June 2013, Fig. 5. http://www.3gpp.org/technologies/keywords-acronyms/97-lte-advanced

[65] D.-S. Shiu, G. J. Foschini, M. J. Gans, and J. M. Kahn, "Fading correlation and its effect on the capacity of multielement antenna systems," IEEE Trans. Comm., vol. 48, no. 3, pp. 502–513, Mar. 2000.

[66] B. H. Fleury, "First- and second-order characterization of direction dispersion and space selectivity in the radio channel", IEEE Trans. on Info. Theory, vol.46, n.6, Sep. 2000

[67] [*Forenza 2006*]

[68] Although in practice other channel and array characteristics may also affect the multiplexing gain as described in [*Forenza 2006*]

[69] Ericsson, "LTE Advanced: mobile broadband up to 10 times faster", Jun. 2011, http://www.ericsson.com/news/1526485

Ericsson, "LTE-Advanced demo in Kista, Sweden", Jun. 2011, http://www.youtube.com/watch?v=oouTy9Prn_M

[70] Outdoor 12X gain in SE with MIMO 12x12 testbed has been achieved by NTT DOCOMO, but only under controlled propagation conditions, including experimental user equipment with a large antenna array (impractical

for commercial mobile devices) in NLOS and proximity to the base station yielding high SNR (i.e., favorable conditions for spatial miultiplexing):
H. Taoka, K. Higuchi, “Experiments on peak spectral efficiency of 50 bps/Hz with 12-by-12 MIMO multiplexing for future broadband packet radio access”, Proc. 4th Intern. Symp. On Comm., Mar. 2010
Another scenario that has achieved full rank with MIMO 2x2 is LOS microwave links, obtaining 2X gain with very specific antenna configurations and type of receive processing:
Ceragon, “Boosting microwave capacity using line-of-sight MIMO”, Technical brief
R. Valtolina, “Making the most of MIMO”, Alcatel-Lucent Techzine, Sep. 19, 2013
http://www2.alcatel-lucent.com/techzine/making-the-most-of-mimo/
[71] 3GPP TR 25.996, “Spatial channel model for multiple input multiple output (MIMO) simulations”, Rel. 12, v12.0.0, Sep. 2014, (see Figure 5.4). http://www.3gpp.org/ftp/Specs/archive/25_series/25.996/25996-c00.zip
[72] [*Ericsson MIMO 2013*] K. Werner, H. Asplund, D.V.P. Figueiredo, N. Jaldén, B. Halvarsson, “LTE-Advanced 8x8 MIMO Measurements in an Indoor Scenario”, Ericsson AB, Stockholm, Sweden. April 16, 2013. http://www.ericsson.com/news/130416-lte-advanced-8x8-mimo-measurements-in-an-indoor-scenario_244129228_c?categoryFilter=conference_papers_1270673222_c
[73] V. Erceg et. al.,“TGn channel model”, IEEE 802.11 Wireless LANs, doc.: IEEE 802.11-03/940r4, May 2004
[74] Indoor 8X gain in SE has been achieved with a MIMO 8x8 testbed, but only in very controlled lab environment with transmit and receive antenna arrays at very close distance, as demonstrated in the following video:
National Instruments, “8x8 MIMO (LTE Advanced)”, Aug. 2011, http://www.youtube.com/watch?v=wklxfXGQ_7s
[75] A. Goldsmith, “Designing reliable Wi-Fi for HD delivery throughout the home”, Quantenna Communications, white paper, Mar. 2009 http://www.quantenna.com/pdf/DesigningReliableWiFi.pdf
[76] [*Cisco VNI 2014*]
[77] F.J. Martin-Vega, I.M. Delegado-Luque, F. Blanquez-Casado, G. Gomez, M.C. Aguayo-Torres, J.T. Entrambasaguas, “LTE performance over high speed railways”, IEEE Conf. on Vehic. Tech., pp.1-5, Sep. 2013
[78] G. Caire and S. Shamai, “On achievable rates in a multi-antenna broadcast downlink,” in Proc. 38th Annual Allerton Conf. Comm., Control, Computing, pp. 1188–1193, Oct. 2000
[79] W. Yu and J. Cioffi, “Trellis precoding for the broadcast channel,” in Proc. Of Global Comm. Conf., pp. 1344–1348, Oct. 2001.
[80] W. Yu and J. Cioffi, “Sum capacity of Gaussian vector broadcast channels,” IEEE Trans. Info. Th., vol. 50, pp. 1875–1892, Sep. 2004
[81] A. Goldsmith, S. A Jafar, N. Jindal, S. Vishwanath, “Capacity limits of MIMO channels,” IEEE Journal on Select. Areas in Comm., vol.21, n.5, Jun. 2003
[82] S. Vishwanath, N. Jindal, and A. Goldsmith, “Duality, achievable rates, and sum-rate capacity of Gaussian MIMO broadcast channels,” IEEE Trans. Info. Th., vol. 49, pp. 2658–2668, Oct. 2003.
[83] M. Costa, “Writing on dirty paper,” IEEE Transactions on Information Theory, vol. 29, pp. 439–441, May 1983.
[84] P. Viswanath and D. Tse, “Sum capacity of the vector Gaussian broadcast channel and uplink-downlink duality,” IEEE Trans. Info. Th., vol. 49, pp. 1912–1921, Aug. 2003
[85] P. Viswanath and D. Tse, “On the capacity of the multi-antenna broadcast channel,” in Proc. DIMACS workshop on Signal Processing Wireless Communications, DIMACS Center, Rutgers Univ., Oct. 7–9, 2002.
[86] ArrayComm, “Essential concepts of multi-antenna signal processing”, Jun 2011 http://www.arraycomm.com/technology/mas-essentials/
[87] 3GPP TS 36.211 V11.5.0, “Evolved Universal Terrestrial Radio Access (E-UTRA); Physical channels and modulation”, Jan. 2014
[88] [*Larsson 2014*]
[89] Shepard, H. Yu, N. Anand, L. E. Li, T. L. Marzetta, R. Yang, and L. Zhong, “Argos: Practical Many-Antenna Base Stations”, in Proc. ACM Int. Conf. Mobile Computing and Networking (MobiCom), Aug. 2012
[90] [*Larsson 2014*]
[91] Massive MIMO or “Full-dimension MIMO” has only recently been considered for discussion in 3GPP LTE Rel.13:

[*LTE Rel. 13*] D. Flore, "Initial priorities for the evolution of the evolution of LTE in Release-13", 3GPP RAN, Sep. 20, 2014. http://www.3gpp.org/news-events/3gpp-news/1628-rel13

92 Massive MIMO has been considered for future 5G standardization. METIS, "First performance results for multi-node/multi-antenna transmission technologies", Deliverable D3.2, doc. ICT-317669-METIS/D3.2 https://www.metis2020.com/wp-content/uploads/deliverables/METIS_D3.2_v1.pdf

93 [*Truong 2013*] K.T. Truong, R.W. Heath, "Effects of Channel Aging in Massive MIMO Systems", July 2013. http://arxiv.org/pdf/1305.6151.pdf

94 [*Macro-diversity*] is one type of space diversity obtained when multiple antennas are placed at a relative distance much longer than the wavelength.

95 The term "frequency reuse" indicates how often the same carrier frequency can be reused across adjacent cells and it is employed to mitigate interference between cells. For example, frequency reuse 3 indicates that three adjacent cells operate at three different carrier frequencies, frequency reuse 1 (also called "universal frequency reuse") means all cells share the same carrier frequency. Lower frequency reuse enables higher network capacity at the expense of higher inter-cell interference. In 1G and 2G cellular systems, typical frequency reuse was 3, 4, 7, 9 or 12. In 3G and 4G systems, universal frequency reuse is employed, with all cells in the network using the same carrier frequency.

96 [*Fujitsu 2011*] Fujitsu, "Enhancing LTE cell-edge performance via PDCCH ICIC," 2011. http://www.fujitsu.com/downloads/TEL/fnc/whitepapers/Enhancing-LTE-Cell-Edge.pdf

97 J.G. Andrews, "Seven ways that HetNet are a cellular paradigm shift", March 2013 http://users.ece.utexas.edu/~jandrews/pubs/AndrewsHetNets_CommMag2013.pdf

98 [*Gigaom 2015*] Fitchard, K., "[The small-cells will] be spaced near enough that Verizon has to be careful they aren't too close, otherwise their signals might interfere with one another, Verizon's director of network engineering and operations Jake Hamilton told me." Gigaom, Feb. 20, 2015. https://gigaom.com/2015/02/20/verizon-is-laying-down-400-tiny-cells-in-sf-to-boost-lte-capacity/

99 A.F. Molish, "A generic model for MIMO wireless propagation channels in macro- and microcells", IEEE Trans. on Sign. Proc., vol.52, n.1, pp. 61-71, Jan. 2004 http://www.merl.com/publications/docs/TR2004-013.pdf

100 [*RealWireless 2012*] RealWireless, "An assessment of the value of small cell services to operators", Virgin Media trials, Oct. 2012 http://www.realwireless.biz/realwireless/wp-content/uploads/2012/10/Small-Cells-Value-for-operators-based-on-Virgin-Media-Trials-V3.1.pdf

101 [*Fujitsu 2011*]

102 4G Americas, "Self-optimizing networks: the benefit of SON in LTE", Jul. 2011 http://www.4gamericas.org/files/2914/0759/1358/Self-Optimizing_Networks-Benefits_of_SON_in_LTE-July_2011.pdf

103 3GPP, "Mobility Enhancements in Heterogeneous Networks (Release 11)," TR 36.839, v11.0.0, Sep. 2012 http://www.3gpp.org/ftp/Specs/archive/36_series/36.839/36839-b10.zip

104 H-S Jo, Y.J. Sang, P. Xia, and J.G. Andrews, "Heterogeneous cellular networks with flexible cell association: a comprehensive downlink SINR analysis," IEEE Trans. Wireless Comm., vol.11, n.10, pp.3484-3495, Oct. 2012

H. Wang, X. Zhou, M.C. Reed, "Coverage and throughput analysis with a non-uniform small cell deployment.", IEEE Trans. on Wireless Comm., vol.13, n.4, pp.2047-2059, Apr. 2014.

105 Mobile operators listed interference and mobility management as the biggest challenge for widespread adoption of small cells. Backhaul was the second largest concern. [*RealWireless 2012*] Figure 4.

106 "The Hype and Reality of Small Cells Performance", Rayal, F., Xona Partners. February 12, 2013. https://communities.cisco.com/community/solutions/sp/mobility/blog/2013/02/12/the-hype-reality-of-small-cells-performance

107 X. Mao, A. Maaref, K. H. Teo "Adaptive Soft Frequency Reuse for Inter-Cell Interference Coordination in SC-FDMA Based 3GPP LTE Uplinks," Globecom, 2008

V. Pauli, J. D. Naranjo, E. Seidel, "Heterogeneous LTE Networks and Inter-Cell Interference Coordination," Nomor research, Dec. 2010

[108] LTE standard Release 10 onward deals with extreme interference conditions by adding enhanced inter-cell interference coordination (eICIC) techniques so as to protect vulnerable LTE control channels, thus avoiding radio link failure (RLF), and handle extreme interference on data channels with schemes such as the Almost-Blank Subframe (ABS).
S. Deb, P. Monogioudis, J. Miernik, J. Seymour, "Algorithms for Enhanced Inter Cell Interference Coordination (eICIC) in LTE HetNets", Feb. 2014, IEEE/ACM Trans. On Networking, pp. 137-150
L. Lindbom, R. Love, S. Krisnamurthy, C. Yao, N. Miki, V. Chandrasekhar, "Enhanced Inter-cell Interference Coordination for Heterogeneous Networks in LTE-Advanced: A Survey", Dec. 2011,
http://arxiv.org/pdf/1112.1344.pdf
Release 11 defines further-enhanced ICIC (FeICIC) to recapture some of the lost capacity within the ABS scheme by employing Reduced-Power ABS (RP-ABS) and to protect vulnerable DL reference signals.
3GPP TSG-RAN WG1 Metting #66bis, R1-113482, Oct. 2011.
http://www.3gpp.org/ftp/tsg_ran/wg1_rl1/TSGR1_66b/Docs/R1-113482.zip
[109] "AT&T deploying SON technology as part of latest network upgrades", Feb. 2012
http://www.cellular-news.com/story/Operators/53297.php
[110] "In the past three years in particular, market forecasts of outdoor small cells have been high – just look at the many press releases and news bits released from analysts covering this space. Yet, to date, only a fraction of the forecasted numbers has materialized." *Small Cells Progress Report – Challenges and Opportunities*, F. Rayal, Xona Partners, June 6, 2014. http://frankrayal.com/2014/06/06/small-cells-progress-report-challenges-and-opportunities/ Frank Rayal, Nov. 17, 2013 http://frankrayal.com/2013/11/17/what-the-future-bodes-for-small-cells-an-operator-perspective/
[111] G. J. Foschini, M. K. Karakayali, and R. A. Valenzuela, "Coordinating multiple antenna cellular networks to achieve enormous spectral efficiency," *Proceedings of the IEEE*, vol. 153, no. 4, pp. 548–555, Aug. 2006
[112] [*Network MIMO*] M. K. Karakayali, G. J. Foschini, R. A. Valenzuela, and R. D. Yates, "On the maximum common rate achievable in a coordinated network," *Proc. of the Int'l Conf. on Communications (ICC'06)*, vol. 9, pp. 4333–4338, Mar. 2006.
M. K. Karakayali, G. J. Foschini, and R. A. Valenzuela, "Network coordination for spectrally efficient communications in cellular systems," *IEEE Wireless Comm Magazine*, vol. 13, no. 4, pp. 56–61, Aug. 2006.
Y. Liang, R. Valenzuela, G. Foschini, D. Chizhik, and A. Goldsmith, "Interference suppression in wireless cellular networks through picocells", ACSSC, pp.1041-1045, Nov. 2007
[113] S. Venkatesan, A. Lozano, and R. Valenzuela, "Network MIMO: overcoming inter-cell interference in indoor wireless systems", Proc. of Asilomar conf., pp.83-87, Nov. 2007
S. Venkatesan, H. Huang, A. Lozano, and R. Valenzuela, "A WiMAX-based implementation of network MIMO for indoor wireless systems", EURASIP Journal on Advances in Signal Processing, Sep. 2009
[114] M. K. Karakayali, "Network coordination for spectrally efficient communication in wireless networks", PhD dissertation, Jan. 2007, Figure 2.10
A.Lozano, R.W.Heath, J.G.Andrews, "Fundamental limits of cooperation", IEEE Trans. On Info. Theory, vol.59, n.9, pp.5213-5226, Sep. 2013
[115] [*Network MIMO*]
[116] H. V. Balan, R. Rogalin, A. Michaloliakos, K. Psounis, G. Caire, "Achieving high data rates in a distributed MIMO system," Proc. Of Mobicom, pp.41-52, 2012
R. Rogalin, O. Bursalioglu, H. Papadopoulos, G. Caire, A. Molisch, A. Michaloliakos, V. Balan, and K. Psounis, "Scalable synchronization and reciprocity calibration for distributed multiuser MIMO", Oct. 2013,
http://arxiv.org/pdf/1310.7001.pdf
H. V. Balan, R. Rogalin, A. Michaloliakos, K. Psounis, and G. Caire, "AirSync: Enabling Distributed Multiuser MIMO with Full Spatial Multiplexing," Aug. 2012, http://arxiv.org/pdf/1205.6862.pdf
[117] S. Parkvall, E. Dahlman, A. Furuskar, Y. Jading, M. Olsson, S. Wanstedt, and K. Zangi, "LTE-Advanced – evolving LTE towards IMT-Advanced", IEEE VTC, pp.1-5, Sep. 2008

D. Gesbert, S. Hanly, H. Huang, S. Shamai, O. Simeone, W. Yu, “Multi-cell MIMO cooperative networks: A new look at interference,” IEEE Journal on Selected Areas in Communications, vol. 28, no. 9, pp. 1380-1408, Dec. 2010.
I. F. Akyildiz, D. M. Guterrez-Estevez, E. C. Reyes, “The evolution to 4G cellular systems: LTE-Advanced”, Physical comm., Elsevier, pp.217-244, 2010
A. Barbieri, P. Gaal, S. Geirhofer, T. Ji, D. Malladi, Y. Wei, and F. Xue, “Coordinated downlink multi-point communications in heterogeneous cellular networks”, Info. Theory and App. Workshop, pp. 7-16, Feb. 2012
[118] 3GPP TR 36.819, “Coordinated multi-point operation for LTE physical layer aspects (Release 11)”, Dec. 2011, http://www.qtc.jp/3GPP/Specs/36819-b10.pdf
D. Lee, H. Seo, B. Clerckx, E. Hardouin, D. Mazzarese, S. Nagata, K. Sayana, “Coordinated multipoint transmission and reception in LTE-Advanced: deployment scenarios and operational challenges”, IEEE Comm. Mag., pp.148-155, Feb. 2012
[119] S. Annapureddy, A. Barbieri, S. Geirhofer, S. Mallik, A. Gorokhov, “Coordinated Joint Transmission in WWAN,” Qualcomm, May 2010
[120] V. Jungnickel, K. Manolakis, W. Zirwas, B. Panzner, V. Braun, M. Lossow, M. Sternad, R. Apelfrojd, nd T. Svensson, “The role of small cells, coordinated multipoint, and massive MIMO in 5G”, IEEE Comm. Magazine, pp.44-51, May 2014.
[121] K. Manolakis, S. Jaeckel, V. Jungnickel, V. Braun. “Channel prediction by doppler-delay analysis and benefits of base station coordination”, IEEE Conf. on Vehic. Tech., Jun. 2013
http://www.mk.tu-berlin.de/publikationen/objects/2013/manolakis2013a/datei
[122] Samsung Smart LTE Networks Awarded Best Mobile Infrastructure and Outstanding Overall Mobile Technology, Feb. 2013. http://global.samsungtomorrow.com/?p=22526
[123] ArrayComm, “A-MAS for LTE and WiMAX”, http://www.arraycomm.com/products/a-mas-software/
[124] Fraunhofer, “Coordinated multipoint (CoMP) Transmission”, May 14, 2013. http://www.youtube.com/watch?v=6aBL7BUqxjI
[125] China Mobile Research Institute, “C-RAN, the road towards green RAN”, Oct. 2011
http://labs.chinamobile.com/cran/wp-content/uploads/CRAN_white_paper_v2_5_EN.pdf, p, 22
[126] [*Fujitsu 2011*]
[127] [*Pellegrino 2012*] G. Pellegrino, K. Griffin, L. Buckalew, “Service delivery and consumer cloud”, Cisco IBSG, Jul. 2012. (slide 7)
https://www.cisco.com/web/about/ac79/docs/clmw/CLMW-Service-Delivery_CrossCountryComparison.pdf
Cisco IBSG, “Connected life market watch”, Aug. 2011. (Slide 28)
https://www.cisco.com/web/about/ac79/docs/clmw/CLMW_Service_Delivery_US_Short.pdf
[128] The World Bank, “Urban population (% of total)”, 2013 http://data.worldbank.org/indicator/SP.URB.TOTL.IN.ZS/countries?display=graph
[129] [*LTE Rel.13*]
[130] [*FCC E911 2014*], Table 1.
[131] D. Andreatta, “As pay phones vanish, so does lifeline for many”, Democrat & Chronicle, December 2013. http://www.democratandchronicle.com/story/news/local/2013/12/16/bell-tolls-for-area-pay-phones/4045131/
J. Sparshott, “More people say goodbye to their landlines: Younger Americans drive trend as cellphone ownership soars”, Wall Street Journal, September 2013. http://on.wsj.com/1E9VFgN
[132] GSMA, “Mobile infrastructure sharing”, April 2008. http://www.gsma.com/publicpolicy/wp-content/uploads/2012/09/Mobile-Infrastructure-sharing.pdf
[133] [*Farrell 2013*] M.B. Farrell, “Cellphone networks overwhelmed after blasts in Boston”, Boston Globe, April 2013. http://www.bostonglobe.com/business/2013/04/16/cellphone-networks-overwhelmed-blast-aftermath/wq7AX6AvnEemM35XTH152K/story.html
[134] [*Cellular security*]
[135] [*Rappaport 2013*]
[136] [*Samsung 2013*] A. Bleicher, “Millimeter waves may be the future of 5G phones”, IEEE Spectrum, Jun. 13, 2013. http://spectrum.ieee.org/telecom/wireless/millimeter-waves-may-be-the-future-of-5g-phones

[137] [*Rappaport 2014*] T.S. Rappaport, W. Roh, K. Cheun, "Smart antennas could open up new spectrum for 5G", IEEE Spectrum, Aug. 28, 2014

[138] [*Mobile Bands 2014*]

[139] [*3D Polarization*] pCells exhibit x, y and z polarization in both line-of-sight and non-line-of-sight environments.

[140] E.g., Using 2.6 GHz spectrum, our demos show iPhone 6 devices can be stacked on top of each other, each maintaining an independent channel and concurrently streaming HD video. The iPhone 6 is 6.1 mm thick.

[141] C-RAN stands for "Cloud-Radio Access Network". It is a wireless network architecture where the RF baseband waveforms are generated in a centralized data center and distributed digitally over an IP network to "remote radio heads" (RRHs) throughout the coverage area. Each radio head contains an interface to the IP network, and one or more analog front-ends (Analog-to-Digital and Digital-to-Analog converters), power amplifiers and antennas. The IP network from the data center to the radio heads is called "fronthaul".

[142] Many cellular C-RAN networks rely upon synchronous transport such as CPRI. CPRI is not necessary for pCell. Conventional asynchronous IP connections can be used, such as fiber, Line-of-Sight and Ethernet. The pWaves self-synchronize through RF.

[143] [*ALU Network Cost*] J. Segel and M. Weldon, "lightRadio White Paper #1: technical overview" 2011, Fig. 4 shows that Site rental and Transmission make up 49% of network costs. http://webform.alcatel-lucent.com/res/alu/survey/alu2CustomForm.jsp?cw=alu2CorpDocDownload&LMSG_CABINET=Docs_and_Resource_Ctr&LMSG_CONTENT_FILE=White_Papers/LightRadio_WP1_Technical_Overview.pdf

[144] As noted in Figure 10 in the Chapter 3, radio waves do not propagate with precise circular patterns in real-world scenarios, but for the sake of explanation we will show them as such in Figure 15 and Figure 16. Irregular propagation patterns are addressed later in this section.

[145] RF power distribution indicates how the RF power being transmitted from the base stations propagates in space. For the sake of illustration, Figure 15 only depicts the areas where the RF power is high enough above the noise floor at the user device to demodulate data successfully (e.g., while meeting a predefined target error rate). In other words, it is the area where the received signal-to-noise ratio (SNR) meets predefined target performance, without accounting for interference from other base stations.

[146] Most LTE devices are 2-antenna devices, with a peak SE of about 8 bps/Hz (20 MHz FDD, MCS 28, 2x2 MIMO) LTE 4x2 averages 1.7 bps/Hz [*Rysavy 2014*], Fig. 29, p. 71. 1.7/8= 0.21: less than one-fourth.

[147] [*Ericsson 2013*]

[148] [*pCell assumptions*] The simulator that generated these results uses the same assumptions as those described for the case of cellular systems at [LTE simulator]. One difference is that for pCell, the power levels at different antennas are dynamically adjusted, and to have fair comparison with cellular, the sum power of all active network antennas for pCell is constrained to be the same or lower than the sum power for cellular. Moreover, since the power of pCell may vary for different network antennas (and to avoid distortions from RF power amplifiers due to high peak-to-average ratio in practical scenarios), we limit the peak power of every network antenna for pCell to 200 mW. While we are aware this condition disfavors pCell since it reduces its sum power over cellular, at least it is representative of practical deployable scenarios. Further, unlike cellular, the 3GPP LTE shadowing model is included in the channel simulation for pCell. While shadowing adversely affects performance for pCell, it reflects more realistic propagation conditions

[149] Some pCell users show a lower data rate than the others because pCell is modeled with the same average power as the cellular pathloss model, but pCell is also modeled with a 3GPP shadowing model that includes regions with high attenuation, such as indoor regions. At slightly higher average power, appropriate for the 3GPP shadowing model, all pCell users would experience near-peak data rate.

[150] While network antenna cross-polarization can increase space-selectivity in certain channel scenarios, to show that pCell achieves high space selectivity without antenna cross polarization, all antennas were aligned with the same polarization.

[151] While device-to-device antenna cross-polarization can increase space-selectivity in certain channel scenarios, to show that pCell achieves high space selectivity without device-to-device cross polarization, all devices were aligned with the same polarization.

[152] [5% outage] To compute the cell-edge performance, we used the cell-edge definition by IMT-Advanced in [ITU 2008], which is the 5% point of the cumulative distribution function (CDF) of the spectral efficiency or throughput

[153] Operators and technology companies providing LTE service in North, Central and South America. http://www.4gamericas.org/

[154] [Rysavy 2014], Fig. 29, pp. 70-71.

[155] Almost all LTE devices on the market today are limited to 2 concurrent antennas for downlink. Although [Rysavy 2014] reports that MIMO 4x4 achieves 2.4 bps/Hz with 4-antenna LTE devices, we are unaware of any high-volume consumer 4-antenna LTE devices on the market today, and if they are introduced, for many years they will represent a small percentage of LTE devices sharing a given cell. If a cell were limited to only 4-antenna devices, the achieved 2.4 bps/Hz is only 41% higher than 1.7 bps/Hz at the cost of 2X the number of antennas. This is consistent with [Ericsson MIMO 2013] and Figure 9, showing rapidly diminishing SE gains with higher-order MIMO, even in a rich multipath environment, and no SE gain beyond 6 antennas at all. [Rysavy 2014], Fig. 29, p. 71. shows a maximum of 4 network antennas for 2-antenna LTE devices, using MIMO 4x2.

[156] While this table shows actuals for 5 MHz DL for both cellular and pCell, the SE for 10 and 20 MHz would be 3% and 5% higher, respectively, and the data rate would be 2.06X and 4.2X higher, respectively. (Figure 10 in [Rysavy 2014] shows 5 MHz SE to be 95% the SE of 20 MHz and 10 MHz to be 97% the efficiency of 20 MHz.) Cellular TDD has about the same SE as FDD, so the same SE is used for both ([Rysavy 2014] states that TDD is within 1-2% the SE of FDD.) pCell is currently only available for TDD.

[157] [Pellegrino 2012]

[158] The peak SE of 20 MHz LTE to a 2-antenna device is 8 bps/Hz computed based on [3GPP TR-36.912]. 5 MHz LTE is 5% lower SE as at Figure 10 in [Rysavy 2014], or 7.6 bps/Hz.
SE, or 7.2 bps/Hz.

[159] SE and Data rate numbers in Table 4 are projected numbers for 20 MHz bandwidth obtained using the calculation described at [LTE simulator]. Cellular SE is also scaled from 5 MHz to 20 MHz accordingly for comparison with pCell SE.

[160] Cellular SE at 20 MHz bandwidth is derived from the cellular SE of 1.7 bps/Hz at 5 MHz, accounting for 5% loss at 5 MHz due to higher PHY layer percentage overhead (cfr., Figure 10 in [Rysavy 2014])

[161] [Ericsson 2013] Fig. 21

[162] [5% outage]

[163] The same Artemis LTE system level simulator and the same assumptions as in [LTE simulator] and [pCell assumptions] are used for the current simulated results.

[164] [5% outage]

[165] See Fig.21 in [Ericsson 2013] Also see the definition of target peak SE (15 bps/Hz) and cell-edge SE (0.075 bps/Hz for outdoor microcell) by IMT-Advanced in [ITU 2008] resulting in 200:1 ratio.

[166] [3GPP TS-36.213]

[167] See Section 5.4.2.

[168] See Section 6.5.

[169] See Section 6.6.

[170] See Section 6.7.

[171] See Section 6.8.

[172] [Mobile Bands 2014]

[173] "Verizon and AT&T Move Toward Voice over LTE (VoLTE) Interoperability", Verizon, November 3, 2014. http://www.verizonwireless.com/news/article/2014/11/verizon-and-att-move-toward-voice-over-lte-volte-interoperability.html

[174] [ALU Network Cost]

[175] The fiber paths in Figure 31 shown as straight lines emanating from a central hub are conceptual, and do not represent real-world geometries of fiber runs (e.g. in conduits under streets, airborne across utility poles, etc.). Also, not all cellular system are backhauled to a single location, although remote radio heads in an area are typically fronthauled to a single location in a C-RAN architecture.

[176] Unlike the fiber paths in Figure 31 the LOS paths shown as straight dashed lines between base stations or RRHs are representative of real-world geometries: real-world LOS paths are straight lines between locations as shown.
[177] Many vendors offer LOS radios with <10µsec repeating latency. The following is just one example: “For an end-to-end link, the EXM-2LL+ adds about 6.5µs total latency, measured from first byte in on ingress at the near-side of the link to first byte out on egress of the far-side”. http://www.escapecom.com/pdf/low_latency_white_paper.pdf
[178] “C-RAN architecture and fronthaul challenges”, Orange Labs Networks, June 25, 2013 http://www.slideshare.net/zahidtg/day-11450annapizzinatorangebackhaulsummit, page 8 shows three sectors, LTE 20 MHz FDD 2x2 MIMO requires CPRI 3 x 2.457 Gbps. One sector requires 2.457 Gbps, and scaling down to one sector of LTE 5 MHz, is 2.457/4= 614 Mbps. The FDD to TDD DL ratio for TDD configuration #2 (3:1) and S-subframe configuration #7 is 5:3.7, so TDD fronthaul DL is 614*3.7/5 = 455 Mbps. Cellular TDD backhaul is 25 Mbps, so the ratio of fronthaul to backhaul data rate is 455/25 = 18:1. Fronthaul can be compressed 2:1, see e.g., “Front-haul compression for emerging RAN and small cell networks“, IDT, April 29, 2013, https://www.idt.com/document/whp/front-haul-compression-emerging-c-ran-and-small-cell-networks. So, achievable fronthaul to backhaul ratio is 9:1.
[179] Common Public Radio Interface. http://www.cpri.info/
[180] Defined in [*Micro-diversity*]
[181] Defined in [*Macro-diversity*]
[182] [*Gigaom 2015*]
[183] With CoMP, an edge user does connect to three cooperating base stations at once, but during this time the three base stations cease serving other users in their individual cells independently, effectively acting like one base station serving one large cell. As such, CoMP is still a base station-centric technology.
[184] The rate at which a given pCell cluster changes depends on many factors, including the device protocol, the wavelength, the speed and/or direction of user or environment motion, and can range from sub-millisecond to seconds..
[185] Note that, for the sake of illustration, the antennas are in a uniform 8x8 grid pattern of 64 antennas, and only 5 users are shown. In the real world, pCell antennas would be in an arbitrary pattern and 64 antennas would be enough capacity for extremely large number of concurrent users (e.g. with 20 MHz TDD LTE, 64 antennas would provide multi-gigabit capacity). So, this illustration is grossly underutilizing pCell capacity for the sake of illustration.
[186] The transmission range is an irregular shape because RF propagation in real-world environments is highly non-uniform, subject to obstacles, device antenna polarization, etc. And, as noted previously, propagation patterns become increasingly irregular at lower power as antenna placements become denser.
[187] For the sake of illustration, the assumption is the RF environment is completely static, so the gray pCell clusters do not change. In the real world, given that the red user must have some non-negligible physical dimensions, the red user would affect the RF environment at least by a small amount, which in most scenarios would change the shape of the gray pCell clusters. Real-world RF environments are extremely sensitive to environmental change.
[188] [*Ericsson 2013*]
[189] [*Pellegrino 2012*]
[190] C. Jakes, Microwave Mobile Communications, IEEE Press, 1974.
[191] [*Truong 2013*]
[192] [*FCC E911 2014*] ¶38
[193] [*Farrell 2013*]
[194] [*Cellular security*]
[195] [*3GPP TS-36.213*]
[196] [*Mobile Bands 2014*]
[197] https://www.oculus.com/
[198] http://www.microsoft.com/microsoft-hololens/en-us
[199] http://www.samsung.com/global/microsite/gearvr/
[200] https://www.google.com/get/cardboard/

[201] http://www.magicleap.com/

[202] http://www.jauntvr.com/technology/

[203] Oculus VR's share site (https://share.oculus.com/) has several VR examples. Microsoft produced a video showing a future vision for Hololens AR (https://www.youtube.com/watch?v=aThCr0PsyuA).

[204] "I can tell you from personal experience that more than 20 ms is too much for VR and especially AR, but research indicates that 15 ms might be the threshold, or even 7 ms." *Latency—the sina qua non of AR and VR*, Michael Abrash, Valve Software blog, December 29, 2012. http://blogs.valvesoftware.com/abrash/latency-the-sine-qua-non-of-ar-and-vr/. Also see J*ohn Carmack Delivers Some Home Truths On Latency*, John Carmack, Oculus Rift blog. http://oculusrift-blog.com/john-carmacks-message-of-latency/682/

[205] http://en.wikipedia.org/wiki/Haptic_technology

[206] E.g., 200 MHz could be the result of 10 carrier-aggregated 20 MHz bands or fewer larger bands, such as Band 41, which is 196 MHz. With pCell, this usage of 200 MHz would be concurrent with other usage (e.g. LTE devices using just 20 MHz in sub bands), so such a large block of spectrum would preclude other usage.