
THE IMPACT OF MOBILE COMPUTERS AND SMARTPHONES ON CDMA2000 NETWORKS

**HOW CURRENT AND FUTURE TECHNOLOGY
ADVANCEMENTS CAN HELP CDMA2000
OPERATORS DEAL WITH THE EXPLOSIVE
GROWTH OF MOBILE DATA TRAFFIC**

JANUARY 2011

Paper developed for the CDMA Development Group

**SiGNALS**
Research Group
www.signalsresearch.com

On behalf of the CDMA Development Group (CDG), Signals Research Group, LLC researched the impact that mobile data and signaling traffic has on existing CDMA2000 networks and how operators can leverage proposed technology advancements to make the most efficient use of their network resources. The information presented in this whitepaper is based on numerous interviews and represents the collective views of the Signals Research Group and several leading CDG member companies.

In addition to providing consulting services on wireless-related topics, Signals Research Group is the publisher of the *Signals Ahead* research newsletter and *The Dollars and Sense of Broadband Wireless*, the first independent in-depth study of next-generation broadband wireless network economics (www.signalsresearch.com).

Table of Contents

Executive Summary	4
1.0 Introduction	7
2.0 The Growth of Mobile Data	8
2.1 Smartphone Mobile Data Usage and the Impact on Signaling	12
3.0 Potential Impact of Mobile Data Traffic on CDMA2000 Networks	15
3.1 The Common Constraints – Limited Spectrum and Cell Sites	15
3.2 The Impact on Network Quality and the User Experience	16
4.0 Mechanics of CDMA2000 Networks: How they Address the Challenges	18
4.1 Select Appropriate Dormancy Timer Settings	18
4.2 Utilize Dynamic Load Balancing	19
4.3 Deploy Femtocells, Remote Radio Heads and Repeaters	19
4.4 Leverage the EVRC-B Codec	19
4.5 Reduce Overhead and Optimize other Network Parameter Settings	19
4.6 Use Data over Signaling	20
4.7 Fully Leverage Available Spectrum and the Capabilities of the Air Interface	20
5.0 Maximizing Existing Network Resources through Advanced Technology Features	22
5.1 1X Advanced.....	22
5.1.1 <i>A More Efficient Voice Codec</i>	24
5.1.2 <i>Interference Cancellation</i>	25
5.1.2.1 <i>QLIC</i>	25
5.1.2.2 <i>Advanced QLIC</i>	25
5.1.2.3 <i>RLIC</i>	25
5.1.3 <i>Mobile Receive Diversity</i>	26
5.1.4 <i>Physical Layer Enhancements</i>	26
5.1.4.1 <i>Smart Blanking</i>	26
5.1.4.2 <i>Efficient Power Control</i>	26
5.1.4.3 <i>Frame Early Termination</i>	26
5.1.5 <i>Simultaneous Voice and Data (SVDO)</i>	27
5.1.6 <i>Market Opportunities and Implications</i>	27
5.2 EV-DO Revision B	28
5.2.1 <i>Market Opportunities and Implications</i>	31
5.3 DO Advanced	32
5.3.1 <i>Smart Networks</i>	33
5.3.1.1 <i>Network Load Balancing</i>	33
5.3.1.2 <i>Adaptive Frequency Reuse</i>	35
5.3.1.3 <i>Distributed Network Scheduler</i>	36
5.3.1.4 <i>Single Carrier Multilink</i>	38
5.3.1.5 <i>Smart Carrier Management</i>	39
5.3.2 <i>Enhanced Connection Management</i>	40
5.3.2.1 <i>Parameter Optimization and Implementation Enhancements</i>	40
5.3.2.1.1 <i>Asynchronous Control Channel Bundling</i>	41
5.3.2.1.2 <i>Multi-AT Page</i>	41

5.3.3	<i>Advanced Topology Networks</i>	41
5.3.4	<i>Device Enhancements</i>	42
5.3.5	<i>Market Opportunities and Implications</i>	45
6.0	Conclusions	46

Index of Figures

Figure 1.	Voice and Mobile Data Traffic Worldwide – 2007 through Q2 2010	8
Figure 2.	Forecasted Global Mobile Data Growth – 2008-2015	8
Figure 3.	Distribution of Mobile Data Applications	9
Figure 4.	Growth of Mobile Data Traffic Versus Data Connection Attempts	13
Figure 5.	Evolutionary Steps to Greater Network Capacity by Freeing up Spectrum	22
Figure 6.	1X Advanced Features and their Impact on Increasing Voice Capacity	23
Figure 7.	1X Advanced Tradeoff between Voice Capacity and Coverage	23
Figure 8.	1X Advanced Tradeoff between Voice and 1X Data Network Capacity	24
Figure 9.	Simultaneous Voice and Data	27
Figure 10.	Multicarrier EV-DO	29
Figure 12.	EV-DO Rev. B User Data Rate Improvements	30
Figure 11.	EV-DO Rev. B Performance Gain with Bursty Traffic	30
Figure 13.	EV-DO Rev. B Devices	31
Figure 14.	Typical Network Load Distribution During Busy Hour	32
Figure 15.	DO Advanced Features	33
Figure 16.	Network Load Balancing	34
Figure 17.	Adaptive Frequency Reuse	35
Figure 18.	Spatial View of Adaptive Frequency Reuse	36
Figure 19.	Distributed Network Scheduler	37
Figure 20.	Benefit of Distributed Network Scheduler	37
Figure 21.	Single Carrier Multilink	38
Figure 22.	Smart Carrier Management	39
Figure 23.	Enhanced Connection Management	40
Figure 24.	Advanced Topology Networks	42
Figure 25.	Enhanced Equalizer Performance Enhancements	43
Figure 26.	Mobile Transmit Diversity Performance Enhancements	44

Index of Tables

Table 1.	Smartphone Comparisons – 2003 versus 2010	11
-----------------	---	----

Executive Summary

Over the last decade the wireless industry has witnessed tremendous growth. During a recent one year period, voice traffic increased by twelve percent and mobile data traffic more than tripled worldwide. This growth trend is expected to continue, with some forecasts suggesting that worldwide mobile data traffic will surpass voice traffic by 2011 and increase ten-fold by 2015. While mobile operators may look forward to the revenue opportunities that mobile data presents, they must also be prepared to deal with the many challenges that it offers while continuing to deliver more efficient voice services to their customers.

As networks become congested with mobile data traffic and the underlying signaling traffic inherent to the “chattiness” of numerous smartphone applications, the user experience can degrade to unacceptable levels. Slow network response times, low data rates and dropped calls are just some of the undesirable characteristics of a network that is strained to its limits. Historically, operators had to add more network capacity by deploying new radio carriers in unused spectrum or build new cell sites specifically for capacity purposes. While these actions cannot always be avoided, they are not always practical, in particular if an operator has limited spectrum resources. Further, unless an operator has already fully optimized its network to maximize the amount of traffic that it can support, these actions can (and should) be avoided.

The CDMA2000 technology roadmap provides many options that operators can leverage to deliver considerably higher volumes of voice and data traffic over their existing networks.

The CDMA2000 technology roadmap provides operators with an assortment of options that they can leverage to deliver considerably higher volumes of voice and data traffic over their existing networks which will further enhance the user experience. To varying degrees these features require only modest changes to existing network infrastructure (e.g., software upgrades) or they are inherent features that are already being included in the chipset roadmap.

Operators should initially ensure that their networks have been fully optimized for the types of voice and data traffic that they are experiencing in their networks. In many cases, relatively simple changes, such as adjusting network parameter settings, pertaining to connection management and signaling overhead, can improve network performance and the user experience. Chapter 4 of this whitepaper takes a closer look at what operators can do today before considering some of the more advanced features of the CDMA2000 roadmap.

Once an operator has fully exhausted all of the optimization techniques at its disposal, there are numerous options which can offer even greater benefits with only modest changes to its existing network. These options, which are discussed in Chapter 5, can be grouped into two sets of solutions – 1X Advanced and DO Advanced.

According to various simulation studies, 1X Advanced allows an operator to increase voice capacity by up to four times, trade off this capacity improvement to increase network coverage by up to seventy percent, or triple 1X data network capacity. The performance gains associated with 1X Advanced are realized through a combination of mobile device enhancements and new base station modems. By taking advantage of these capacity gains, an operator can free up spectrum for new mobile broadband services. Since 1X Advanced is fully backwards compatible with CDMA2000, an operator has the flexibility to only deploy some of the 1X Advanced features or to introduce the features over time, knowing that the full benefits of 1X Advanced will not be realized until all of the devices and base stations are fully upgraded to 1X Advanced.

1X Advanced can quadruple voice capacity or trade off this capacity improvement to increase network coverage by up to seventy percent or triple the 1X data network capacity.

Sprint recently announced a major network upgrade that includes 1X Advanced as a key cornerstone behind its overall strategy. Likewise, MTS India has announced plans to deploy 1X Advanced to free up spectrum for additional voice and data services.

1X Advanced and DO Advanced fully maximize existing network resources to increase network capacity and improve the user experience in a relatively cost-effective manner.

Instead of adding data capacity across an entire network, DO Advanced dynamically allocates existing network capacity where and when it is needed.

Starting in 2011 CDMA2000 will support simultaneous voice and data sessions. This capability means that CDMA2000 users will be able to access the web, send emails, or use other data intensive applications while communicating on a phone call. DO Advanced is comprised of a variety of features which can be introduced in an EV-DO Revision A (Rev. A), Multicarrier EV-DO or EV-DO Rev. B network. DO Advanced leverages two basic characteristics of today's mobile data traffic, namely that the mobile data traffic across an operator's network is never uniformly distributed in space and time, and that many of the data sessions are relatively short in duration.

Instead of adding capacity across an entire network, DO Advanced allows an operator to dynamically allocate existing network capacity where and when it is needed the most. In most cases, the "Smart Network" features of DO Advanced only require software changes in the infrastructure with many of its benefits extending to legacy devices. Unlike any other network upgrade, the increase in network capacity and the improvement in user experience can be realized immediately without requiring a lengthy device migration strategy.

One of the Smart Network features, called Advanced Topology Networks, leverages macro-, micro-, pico-, and femtocells, as well as repeaters and remote radio heads, to move the transmitters closer to the user. This combination intelligently addresses the demand for data traffic by adding the appropriate amount of capacity where it is needed, while also improving the user experience.

DO Advanced also introduces features which can help optimize an operator's network to support a large number of mobile devices by minimizing the signaling overhead associated with these devices attaching and disconnecting from the network. To the extent this overhead is minimized it frees up even more network capacity. Plus, DO Advanced includes new features which will improve the performance of mobile devices, including faster data rates and extended battery life. Although DO Advanced does not require new devices, these features will eventually become enabled since they are an inherent part of the CDMA2000 chipset roadmap.

The worldwide base of CDMA2000 operators are at different junctures in their network technology evolution. Some operators, such as those in North America, who have access to large amounts of unused spectrum and capital, are deploying LTE. In Japan, KDDI also has plans to deploy LTE, but it recently launched Multicarrier EV-DO nationwide to maximize its existing network resources and to provide an enhanced user experience for its CDMA2000 subscribers.

Other CDMA2000 operators, such as those in Indonesia and Pakistan, have already deployed Multicarrier EV-DO Rev. B, and others in China, India, Russia, South East Asia, Africa and Europe are either deploying the technology or are trialing it. Further, nearly half of the 315 CDMA2000 operators worldwide have yet to deploy broadband data services due to a variety of factors, including a lack of demand, limited capital resources, or unavailability of spectrum. Regardless of their network evolution, all of these operators will continue to rely on their 3G networks to deliver voice, data and machine-to-machine (M2M) services well into the next decade.

Despite their differences, these operators share two things in common. First, they are addressing the constraints that the ever-increasing mobile data usage is placing on their networks. Second, these operators recognize that their existing CDMA2000 networks will need to support the preponderance of their network traffic for years to come.

The evolution of CDMA2000 addresses network traffic growth through incremental, affordable and backward compatible enhancements, while preserving existing investments.

By fully understanding the numerous options that are at their disposal, CDMA2000 operators will be able to prudently pick and choose the features that are best suited for their needs. For some operators, this may mean only changing a few network parameters. Other operators may take a more aggressive approach and introduce the more beneficial changes enabled by the CDMA2000 roadmap. In any case, the evolution of CDMA2000 enables operators to resourcefully address the tremendous growth in voice, data and M2M traffic through incremental, affordable and backward-compatible enhancements, such as Smart Network techniques, that will preserve their existing investments.

1.0 Introduction

This paper examines the numerous implications associated with the explosive growth of mobile data and signaling traffic due to the widespread and growing popularity of smartphones, tablets and other mobile computing platforms.

While the historical growth and projected future growth of mobile data usage has been widely documented, the implications that this usage has on an operator's 3G network have been largely ignored. More importantly, the steps that an operator can take to leverage its existing network, so that it can handle this traffic in the most efficient and cost effective manner, have not been explicitly outlined.

Too often the argument is made that next-generation technologies such as LTE will address these challenges. But the reality is that for most operators LTE remains in the pre-planning stages, if for no other reason due to the lack of capital and sufficient spectrum required to take full advantage of the technology. This statement is particularly true for major CDMA2000 markets, such as China and India, which have only recently auctioned spectrum for 3G technologies and where suitable and sufficient spectrum for LTE remains on the distant horizon. Furthermore, once these new networks are in place, these operators will continue to depend on their legacy 3G networks to provide rural coverage, sustain roaming, generate revenue and deliver valuable services to the majority of their subscriber bases for many years to come.

The purpose of this paper is to provide insight into the implications that data traffic can have on an operator's network and how to optimally address the demand for network traffic and signaling capacity. Its content is based on our own research and personal perspectives, supplemented by a series of interviews conducted with leading CDMA2000 vendors.

Today's explosive growth in mobile data traffic and connected devices can play havoc on the quality of a network and an end user's experience. Yet, this outcome can be avoided, frequently without requiring the deployment of an entirely new network. In some cases, the pain points are self-inflicted and easily addressed with network optimization. In other instances, the problem can be readily mitigated without jumping to a new technology or without introducing measurable changes to existing CDMA2000 networks.

Most of the solutions that the CDMA2000 vendor community is proposing to operators leverage existing standards and an operator's installed base of hardware. These changes are quite often software-based and many of them come naturally through the existing product roadmap of infrastructure and device chipsets.

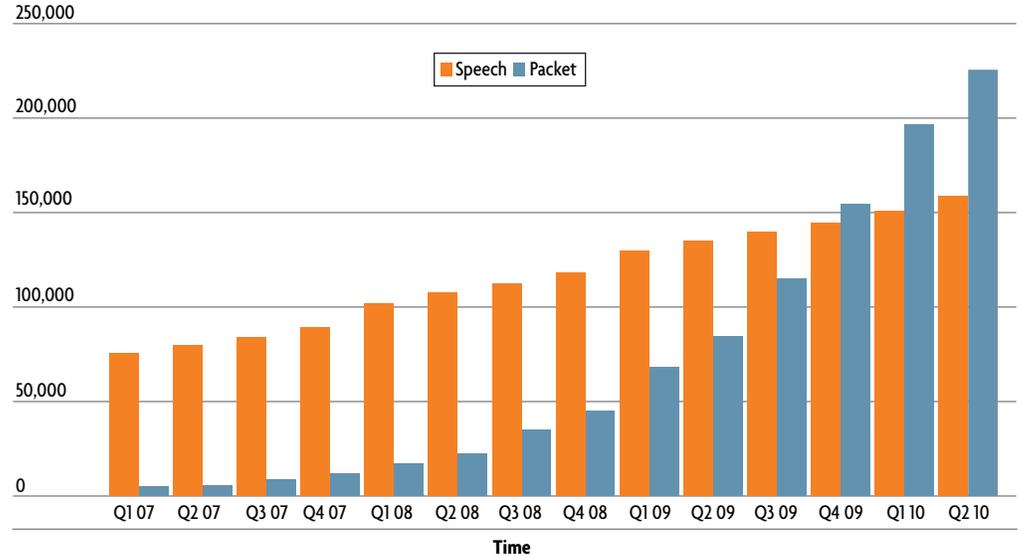
The CDMA2000 evolution path consists mostly of software upgrades that leverage existing standards, devices and infrastructure.

2.0 The Growth of Mobile Data

Over the last several years there has been a phenomenal growth in mobile data and signaling traffic. As illustrated in Figure 1, in a one year period, starting in the second quarter of 2009, mobile data traffic nearly tripled and exceeded voice traffic.

Figure 1. Voice and Mobile Data Traffic Worldwide – 2007 through Q2 2010

Total (UL + DL) traffic per month (TB/month)

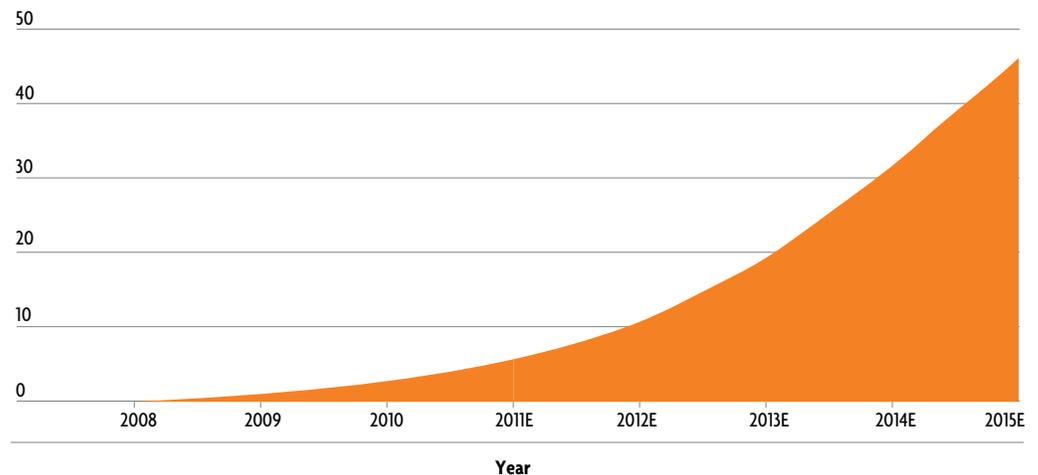


Source: Ericsson (recreated by SRG)

Further, industry forecasts suggest that there is no end in sight to this growth. Figure 2 provides one such forecast from Ericsson that predicts that by 2015 the worldwide total of mobile data traffic will increase by more than tenfold over today's levels.

Figure 2. Forecasted Global Mobile Data Growth – 2008-2015

Yearly data traffic (Exabytes)



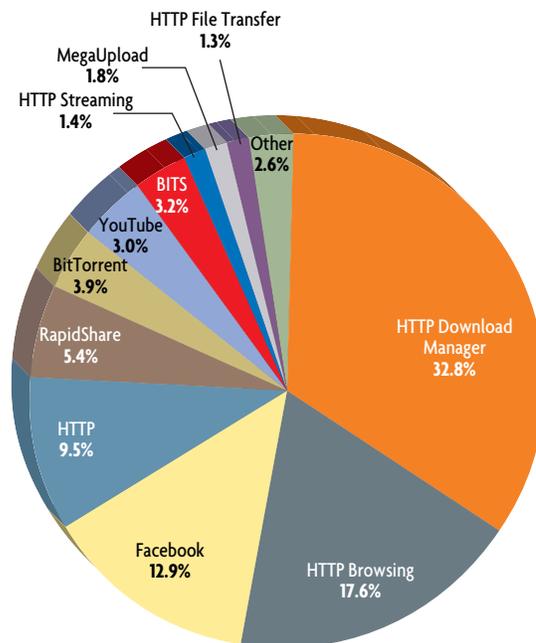
Source: Ericsson (recreated by SRG)

To address the growth, it is also helpful to understand the types of mobile data traffic that operators are seeing in their networks and to recognize that this growth is occurring on a worldwide basis. Figure 3 illustrates the five day distribution of mobile data traffic for Smart Telecom, a CDMA2000 operator in Indonesia and also the world's first operator to commercially launch Multicarrier EV-DO Rev. B services.

As indicated in Figure 3, downloading data was responsible for the highest percentage of total data traffic (32.8%), followed by browsing (17.6%) and Facebook (12.9%). Other well-known Internet applications such as YouTube, Skype, BitTorrent and iTunes are also present. Presumably, some of these applications, in particular those applications involving social networking, accounted for a higher percentage of total signaling traffic even though the percentage of data traffic may have been relatively modest. The reason why some applications generate more signaling traffic than other applications is discussed later.

Instead of there being a single reason for this phenomenal growth in mobile data and signaling traffic, there are actually several intertwined factors that must be considered.

Figure 3. Distribution of Mobile Data Applications



Source: Smart Telecom (recreated by SRG)

More aggressive rate plans. Over the last decade, mobile operators have become far more aggressive in their rate plans, thus making it more affordable for subscribers to consume large amounts of data each month.

When operators in North America first launched 3G services with CDMA2000 1X (circa 2002) they frequently offered plans that charged for data usage on a per minute basis, which meant that even under the most optimistic scenarios consumers were paying at least \$0.35 per

The cost of an unlimited data plan has fallen by 60% in the last decade and data packages are now available which do not require an annual service contract.

MB, but substantially more if they remained idly connected to the Internet while browsing the Internet or reading their email. Later in the same year, when these operators started offering per-Megabyte pricing plans, an unlimited usage plan was \$99.99 per month or \$75 for a mere 40MB of data usage per month – that’s \$1.88 per Megabyte.

When 3G services using EV-DO were launched a few years later the price plans became more affordable, with unlimited usage plans dropping from \$99.99 per month to \$79.99 per month. Now, however, it is common to pay only \$40 to \$60 for unlimited usage, although some plans are now capped at a still impressive 5GB per month. There are also various pay-as-you-go plans that do not require a service contract, lowering the barrier to entry for infrequent users of mobile data that may not commit to an annual or multi-year contract.

Separate from the unlimited data plans, which usually involve data cards or USB dongle modems, operators are now starting to require that subscribers sign up for a data plan whenever they purchase a subsidized handset. For example, one North American operator requires its subscribers to sign up for an unlimited data plan whenever they purchase a smartphone. Also, with many of its multimedia feature phones, the operator’s subscribers must commit to at least a 25MB per month data plan, if not an unlimited plan.

The implications are two-fold. First, subscribers who would have otherwise not selected a data plan are now compelled to use mobile data since they are paying for it anyway. Second, to get the most value out of these plans, subscribers are incentivized to consume as much data as possible, or at least up to the threshold associated with their rate plan. With unlimited plans this means seemingly unlimited usage.

Higher performance networks. Initially, the user experience associated with data plans was not very compelling, in particular relative to the capabilities of today’s 3G networks. Not only was the consumer paying nearly \$100 per month, but he or she was experiencing an average data rate of 40-80 kbps and latency on the order of a few hundred milliseconds. For the end user, these limited capabilities meant long delays to access web pages, download emails and transfer files. Yet, it prepared the way for EV-DO, the first commercially-successful mobile broadband technology.

With the widespread availability of EV-DO Rev. A, operators now advertise download data rates of 600 kbps to 1.4 Mbps. Plus, the average reverse link data rate, or the speed at which data is sent from the mobile device to the network, is almost as fast as the average forward link data rate, while the latency has fallen below 100ms. In the case of smartphones, low latency is just as important as high data rates

Lower latency is just as important as high data rates in improving the user experience.

Given the improved capabilities of the network a typical user is now far more likely to spend more time consuming mobile data. Most importantly, with the improved user experience subscribers are consuming more mobile data than previously possible, due to a combination of the increased time spent in a mobile data session and the amount of data that can be sent during any given period of time.

Compelling devices offering simple and easy access to services. Separate from the networks being capable of delivering more bandwidth and lower latency, the features and usability of handsets have greatly advanced over the last decade. While impressive at the time of introduction, some of the industry’s first smartphones would today be considered a basic or entry-level

feature phone in most markets. Table 1 compares the features of a high-end smartphone from 2003 with today's entry-level and advanced smartphones.

Table 1. Smartphone Comparisons – 2003 versus 2010

Feature	2003 Advanced Smartphone	2010 Entry-Level Smartphone	2010 Advanced Smartphone
Display Size	208 by 320 pixel screen	480 by 800 pixel screen	960 by 640 pixel screen
RAM	16MB RAM	256MB RAM	512MB RAM
Camera	VGA camera	3MP camera	Dual-cameras: 5MP + 3MP
Color Display	12 bit 4,096 color display screen (TFT/TFD)	262,144 color display screen (TFT)	16,777,216 color display screen (LCD)
Memory	12MB internal memory and 16MB external memory stick included	8GB internal memory and external support for 4GB included	32GB internal memory and external support for memory card (not included)
Weight	158grams	158grams	137grams
Size	117 x 59 x 27mm	116 x 56 x 16mm	115 x 59 x 9mm
Audio	MP3, WAV	MP3, AAC, AAC+	MP3, WAV, AAC, AAC+
Video	View movie clips at 176 x144 pixels	MPEG4, WMV, 3GP, 3GW	Play/record video (H.264 up to 720p, 30fps, M-JPEG) and view HD-video
Position Location	N.A.	A-GPS	A-GPS
Messaging	SMS, MMS	AIM, Windows Live, MSN, Yahoo, Google Talk, Facebook pre-installed	Supports virtually all IM and social networking applications
Connectivity	Bluetooth	Bluetooth and Wi-Fi (b/g/n)	Bluetooth and Wi-Fi b/g/n
Access Method	Stylus and push button	Full keyboard (QWERTY) and touch screen	Multi-Touch touch screen
Additional	N.A.	Accelerometer, magnetometer, compass	Accelerometer, compass
Processor	156MHz processor	600MHz processor	Dual 1.2GHz processors
Talk Time	5 hours talk time	7.5 hours talk time	7 hours talk time
Cost	\$700 (unsubsidized)	\$370 (\$50 subsidized)	\$700 (\$300 subsidized)

Source: Various websites and Signals Research Group, LLC

Smartphones in 2010 cost almost half as much as their predecessors from 2003, yet they have considerably more processing power

Entry-level smartphones in 2010 cost almost half as much as their predecessors from 2003, yet they have considerably more processing power to support advanced features, the ability to provide a far more compelling multimedia experience, and easy-to-use touch-screen displays. Compared to the 2003 models that typically supported SMS, web browsing, email and MMS, the 2010 models have built-in support for many of the more popular social networking applications, 3D massively multiplayer online role-playing games (MMORPG) and IM clients that drive a lot of network signaling and mobile data traffic.

Access to a vast number of compelling applications. When 3G mobile data services were first launched, the only real applications available were games, email and Internet browsing – be it on the phone's browser or with a notebook computer tethered to a PC card. Mobile operators also virtually forced their subscribers to purchase applications from within their walled gardens. These walled gardens, which were accessed with a basic WAP browser, featured a

limited number of applications and downloadable content – customized ringtones and wall-paper generally topped the list.

Over the last few years the market for smartphones has exploded. According to Gartner Research, smartphone shipments grew by 96% over the past year, as of Q3 2010, and accounted for 19% of global handset sales. By 2013 the firm estimates that smartphones could account for 43% of global handset sales

Smartphone shipments grew by 96% over the past year and accounted for 19% of global handset sales.

In addition to the advanced capabilities of the devices and networks, consumers now have more choices for what they can do with their smartphones. Apple currently advertises “250,000 ways to make the iPhone better,” while subscribers with an Android-based platform can choose from more than 140,000 applications. These applications can easily be downloaded directly to a phone or sideloaded after downloading onto a computer. Similar online stores are available for Symbian, Windows and Palm webOS smartphones that may not offer quite as many applications, but arguably far more applications than what was available only a few short years ago.

Social networking is at the forefront of modern society. There has been a dramatic paradigm shift in how people interact and stay in contact with each other. Prior to the Internet, people actually had to put pen to paper to keep in touch. If someone wrote a letter home once a week then that individual would have been considered a frequent communicator.

With the introduction of Internet-enabled electronic communication applications (e.g., SMS and email) it became much easier and faster to communicate. What was once a weekly or monthly activity became a daily activity more often than not.

A recent *Wall Street Journal* article shed some light into this social trend. According to the article, which leveraged research conducted by the media research firm, Nielsen Co., the average North American teen sends and receives nearly 3,400 messages a month – with their parents also responsible for a meaningful amount of messaging traffic. Adults aged 45 to 54 send and receive more than 320 messages a month. In large part, this behavior can be attributed to social networking. One major 3G operator had 400 million text alerts for various social networking services on its network in October 2009. By September 2010 that figure doubled to more than one billion. Twitter claims that its service generates more than two billion tweets per month.

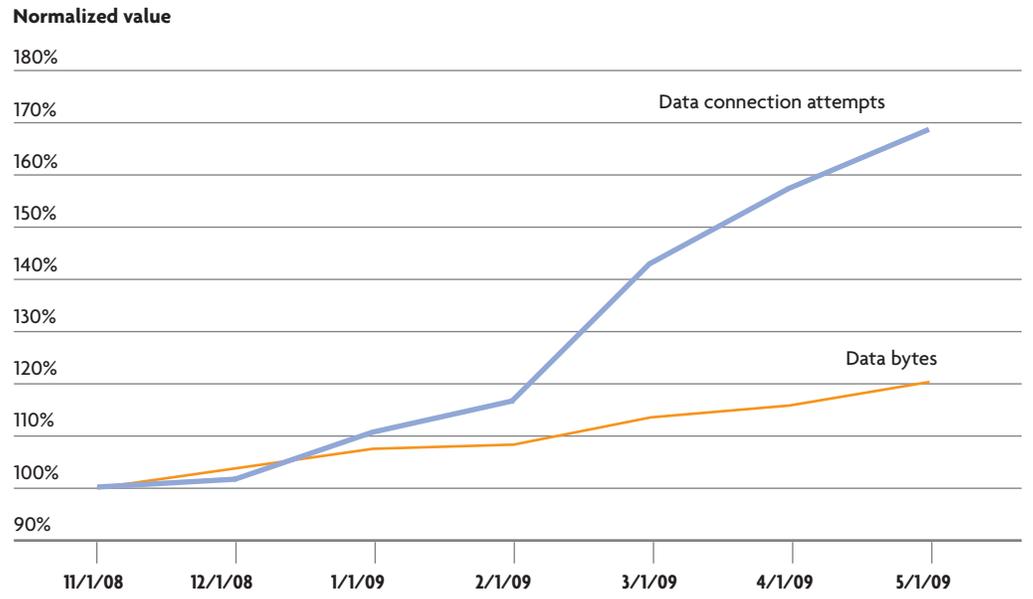
Social networking sites and services, such as Twitter and Facebook, are responsible for a large percentage of the signaling traffic that occurs on an operator's network.

Social networking sites and services, such as Twitter and Facebook, have completely revolutionized the concept of keeping in touch with friends and family. Popular celebrities and famous politicians have tens of thousands, if not hundreds of thousands of followers who anxiously wait for the next update. Each update, sent via a simple message, is instantly communicated to everyone signed-up for the service. These messages, sent mostly to mobile devices, generate a massive amount of messaging traffic over the wireless networks.

2.1 Smartphone Mobile Data Usage and the Impact on Signaling

While the demand for mobile data usage around the world is doubling each year, the signaling traffic associated with this data traffic is growing at even a faster pace. This relationship is shown in Figure 4, which is based on real world data that Airvana collected from a North American operator. During the six-month period that it examined, mobile data traffic grew by an impressive 20%, but the signaling traffic associated with setting up and tearing down those data sessions grew at an even faster 70% rate.

Figure 4. Growth of Mobile Data Traffic Versus Data Connection Attempts



Source: Airvana (recreated by Signals Research Group, LLC)

There are several reasons why signaling traffic is outpacing data traffic. First and foremost, smartphones require eight times as many connection attempts than data cards and similar form factor devices. According to the Airvana study, data cards generated twenty-five times more data than smartphones, but only connected to the network a third as often as smartphones.

Given the small screen size of a smartphone and the likelihood that a consumer is far more likely to download a multi-megabyte file to a notebook computer than to a handheld device, it isn't surprising that smartphones consume less data traffic. However, the relative abundance of signaling traffic does require an explanation.

A typical smartphone user “snacks” on mobile data, meaning lots of short data sessions, while someone with a 3G-enabled notebook computer has less frequent, albeit more lengthy data sessions.

A typical smartphone user “snacks” on mobile data, meaning the subscriber engages in a large number of very short data sessions, for example while standing in line at a bus stop, versus a lengthy data session that someone with a notebook computer might engage in at a coffee shop. Further, unlike desktop email that is received on a periodic basis, most mobile emails are received the instant they are sent – just like text messaging. Each one of these short data sessions requires signaling messages to establish and tear down the session, regardless of the actual length of the session or the amount of data that is sent and received. All things being equal, and ignoring mobility-related signaling messages that are not reflected in the previous figure, a twenty-second web browsing session on a mobile phone generates the same amount of signaling messages as a multi-megabyte download data session or syncing Outlook. Current web-page designs are also inherently interactive. After HTTP data is received by the client, it triggers more HTTP GETs until the web page is fully loaded.

Even more problematic is that the aforementioned social networking connections are by nature very frequent and of very short time duration. Each time someone sends or receives a tweet, or is informed that a friend has updated their Facebook account, a data session is setup and torn

down. Further, a single IM session involving multiple back and forth messages between friends could actually equate to several data sessions – one for each message – depending on how the operator has set some of its network parameters. Each IM message generates the equivalent call set up and tear down signaling messages of a more lengthy data session in which large volumes of data traffic are sent. The same applies to email which requires each data session to be setup and torn down.

Finally, there is “always-on” and the “hidden” signaling traffic associated with “keep alive” messages. These messages are used by various social networking services or other applications, such as Virtual Private Networks (VPNs) and healthcare applications. Anyone who has used an application such as Skype and received pop-up messages indicating that a contact is online should be all-too-familiar with these messages.

Although “keep alive” messages only generate a miniscule amount of data traffic, they can generate an inordinate amount of signaling traffic without the user being aware.

While these messages generate miniscule amounts of data traffic and the intended purpose of these “keep alive” messages can provide beneficial information to the user, they can generate an inordinate amount of signaling traffic. Some applications can generate these “keep alive” messages as often as every couple of minutes. They occur without the knowledge of the user and take place anytime the phone is turned on and the application is launched. Given that many social networking applications launch automatically when the phone is turned on, this means that most smartphones inherently generate a large amount of signaling traffic whenever they are turned on and even when they are not in the hands of the subscriber.

In addition to having an impact on battery life and causing additional radio interference, a phone’s radio must wake up to transmit/receive each “keep alive” message. Again, the amount of signaling traffic required to establish and tear down each session is equivalent to the amount of signaling traffic required to establish and tear down a more lengthy data session in which meaningful amounts of data are sent and received.

The overall implication is that operators need to take into consideration network capacity and optimization that not only addresses the data traffic on its network, but the underlying signaling traffic associated with the data traffic. This issue will be revisited several times in the next three chapters.

3.0 Potential Impact of Mobile Data Traffic on CDMA2000 Networks

The future success of mobile broadband is not predicated on a lack of demand, but by a lack of supply.

Unlike most industries, the future success of mobile broadband is not predicated on a lack of demand, but by a lack of supply. Regardless of an operator's chosen radio technology or set of technologies, it will face numerous constraints, including a lack of sufficient spectrum, backhaul capacity, switching capacity, indoor coverage and cell sites in densely-populated urban settings. Besides these constraints, an operator may deal with overloaded networks, which degrade the user experience, frustrate the consumer and increase churn, as well as exacerbate the economics associated with delivering large amounts of data at ultra-competitive subscription fees.

In this chapter we will look in more detail at these implications, thus setting the stage for Chapters 4 and 5, which look at how the existing CDMA2000 roadmap can address many of these challenges.

3.1 The Common Constraints – Limited Spectrum and Cell Sites

Operators face two common constraints that inherently limit the amount of mobile data traffic that they can support on their networks. First, operators have a limited amount of spectrum at their disposal. Once that spectrum is fully utilized they must resort to cell splitting or deploying additional cell sites to carry the increased network traffic.

Smartphone sales are not waiting for new network deployments. Therefore, operators will continue to deploy additional CDMA2000 resources to deal with the expected growth, even as they look to deploy LTE.

Mobile operators around the world have varying amounts of spectrum. Some North American operators have an abundance of new and unused spectrum, providing them an opportunity to deploy next generation technologies such as LTE. Still, these operators face constraints within their pre-existing spectrum allocations that support the growth of CDMA2000 voice and data. It is technically and economically challenging for them to move their subscribers onto a new network at a fast enough pace to offset the ongoing growth of new data subscribers and their ever-increasing data usage. Smartphone sales are not waiting for new network deployments. Therefore, these operators are continuing to deploy additional CDMA2000 radio channels to deal with the expected growth over the next several years, even as they deploy LTE.

Outside of North America, most operators are not as fortunate. Operators in India represent an extreme example, since they must support tens of millions of voice and data subscribers with no more than 2x5MHz of spectrum. In any given market, the mix of 1X and EV-DO radio carriers varies, but in all cases these operators have to judiciously manage their networks and squeeze as much capacity as possible from their limited spectrum to support both voice and data traffic.

While cell-splitting is a positive indicator in the sense that the operator is seeing strong demand for its services, it is also an expensive proposition to undertake and doesn't yield much of an improvement in dense urban areas where cell sites are already in close proximity to each other. Most of an operator's network delivery cost is due to site-specific factors and not the actual hardware that it deploys. Connecting a new cell site to an electrical power grid, installing the necessary heating and/or cooling elements, provisioning for backhaul, building a tower in some cases, and the ongoing site leasing expenses are only a few of the major cost drivers that an operator incurs when it deploys a new cell site. Further, the operator must secure the proper zoning permits which can take years to obtain due to the NIMBY (not in my backyard) mentality that exists in many local markets.

It is best for an operator to meet its targeted coverage requirements and traffic demand expectations with the fewest number of cell sites.

Therefore, from a pure economics point of view, it is best for an operator to meet its targeted coverage requirements and traffic demand expectations with the fewest number of cell sites. Once an operator has used up its entire spectrum allocation, they must begin using their existing spectrum resources more efficiently without impacting their installed network and disrupting existing services. As discussed in the next two chapters, many of the proposed advancements within the CDMA2000 roadmap are designed with this in mind.

3.2 The Impact on Network Quality and the User Experience

Prior to running out of spectrum or the ability to deploy additional capacity cell sites, excessive mobile data traffic can have numerous impacts on an operator's network. Likewise, there are many factors that an operator must take into consideration which can impact the user experience, especially with mobile data applications.

There are several KPIs (key performance indicators) that operators use to evaluate the ongoing quality of their networks, and more are being introduced with DO Advanced. For voice calls, operators monitor the call completion rate, which defines the percentage of successful call attempts and the number of dropped calls. Both KPIs are of critical importance as they are very evident to the typical consumer. For data calls, knowing webpage download times and the duration of data session connections are of equal importance. The number of KPIs and their ratios are expected to evolve over time to capture operational metrics that can improve the user experience. The Smart Network techniques introduced with DO Advanced are accompanied with several new KPIs. For instance, new DO Advanced KPIs will allow an operator to know whether the sector is connection-capacity limited or data-capacity limited, as well as the number of existing sector users that a new user would have to compete against for service.

When networks become congested due to excessive voice or data traffic, there may not be enough capacity in the serving cell or adjacent cell to support the call or data session. Consequently, the attempted call or data session does not go through or an active call is dropped or the data session is delayed. The latter is most likely to occur when the mobile device attempts to complete a handover into a congested cell, but can occur even if there is available capacity in the air interface (e.g., between the base station and the user).

Data traffic does not adversely impact voice call performance since CDMA2000 networks handle voice traffic exclusively across CDMA2000 1X radio channels and data traffic across separate EV-DO radio channels.

With regards to a phone call, this outcome could be due to congestion within the CDMA2000 1X base station controller (BSC)¹ that serves as the controlling element for all of the base stations connected to it and is responsible for handling radio resource management. In this case, the BSC could be overburdened by excessive processing requirements that it was not designed to handle and which could be due to smartphones generating large amounts of signaling traffic. Unlike other network technologies, it is important to note that CDMA2000 networks handle voice traffic exclusively across CDMA2000 1X radio channels and data traffic across separate EV-DO radio channels. Therefore, data traffic across EV-DO radio channels does not adversely impact voice call performance across 1X channels, and vice versa.

Separate from these KPIs there are other factors which an operator needs to take into consideration which could impact the quality of its network.

¹ CDMA2000 1X uses a BSC to control the individual base stations assigned to it while CDMA2000 EV-DO uses a Radio Network Controller (RNC) to perform similar functions, albeit for the EV-DO radio channels only.

Slow Response Times. A slow network response time, or high latency, impacts the user in a couple of different ways. First, a slow response time causes the user to wait for the requested action to take place. Second, high network latency can impact the quality of real-time applications, such as VoIP or video conferencing. Finally, high latency indirectly impacts the achievable throughput since the TCP protocol requires confirmation for each data packet that is sent before it sends the next packet. With today's 3G network technologies, which are capable of sub-100ms round trip time (RTT) pings, a slow response time generally indicates network congestion either within the backhaul, air interface, BSC and/or Radio Network Controller (RNC).

Low Average Throughput. Users may not understand what a megabyte is, but for some reason they all understand the meaning of Mbps, or at least they know that higher is better. Further, even if a consumer is unaware of the actual data rate, a slow connection is readily apparent. Microsoft Outlook is an inherently chatty application that suffers from low throughput and high latency and is susceptible to long synch times. With low data rates, web pages take a long time to load, and bandwidth intensive applications such as streaming video deliver a poor user experience. In the case of streaming video, users spend as much time waiting for the video to buffer as they do watching the video. Like all other broadband data connections (e.g., cable and DSL), users share a common resource. The more users sharing the bandwidth means less network resources are allocated to each user, causing a lower average user data rate.

Assuming the user is not in an RF-challenging environment (see next bullet) and that the operator has properly dimensioned the backhaul to support the capabilities of the air interface, low throughput is likely a sign of network congestion. It's worth noting that the average achievable throughput in an operator's network is somewhat at the discretion of the operator to determine versus the capabilities of the radio technology itself. An artificially low target value could suggest that an operator is not willing or able to add capacity until it reaches a higher threshold of congestion. Conversely, an operator could advertise and deliver on higher promised data rates, but the operator would then be obligated to have more network capacity.

Poor Coverage. A poor user experience, such as the inability to connect, a dropped call, or low data rates, is usually due to poor network coverage and not to network congestion. Operators face numerous constraints which limit their ability to deploy a cell site wherever they desire. Further, the effective range of an RF signal has to adhere to the basic laws of physics, so the further a mobile device is from the serving cell site, the poorer the signal quality, and the lower the achievable data rate. Conversely, a mobile device that is operating from an indoor environment, yet being served by an outdoor cell site, faces greater challenges since RF signals propagate much further in free space than they do through concrete, brick and mortar.

Lastly, while not necessarily a coverage issue, a mobile device at the edge of a cell normally detects signals from multiple cell sites. In this situation there is a higher amount of interference adversely impacting network capacity and the user experience in the form of lower data rates.

Rapid Battery Consumption. When a device checks, connects, transmits, receives and acknowledges the receipt of packets of information from a network it consumes more energy and its battery life is reduced. With smartphones, increased battery life has become a key factor leading to consumer satisfaction. Performance indicators or KPI ratios, such as the duration of a connection relative to the total time required to download/upload the desired data, are becoming more important to optimize battery life.

4.0 Mechanics of CDMA2000 Networks: How they Address the Challenges

One of the unique attributes of CDMA2000 networks, as it pertains to supporting voice and data services, is that there are actually two closely intertwined networks in place – one network supporting voice and basic data services (CDMA2000 1X) and the other network supporting the majority of the data traffic (CDMA2000 EV-DO). The CDMA2000 1X voice traffic and EV-DO data traffic use dedicated radio channels. The EV-DO traffic channel is optimized for data traffic – thus “EV-DO” stands for Evolution - Data Optimized.

Since 1X and EV-DO traffic travels across separate air links, an operator can optimize its network for data traffic without impacting the quality of its voice services, and vice versa.

Within the core network, CDMA2000 1X and EV-DO traffic also follow somewhat separate paths, although they often have common touch points. For example, there is a separate radio network controller (RNC) that is used to control the individual EV-DO radio carriers at each cell site and to assign and control how data traffic is allocated within the EV-DO air interface. This nuance has important implications, since it means that the operator can separately dimension the capacity and processing power of its EV-DO RNCs and CDMA2000 BSCs. Further, if an operator faces data congestion, there will not be any impact on the quality of the voice network.

Also, when it comes to dealing with chatty smartphone applications, EV-DO requires minimal signaling to establish and tear down a data session since EV-DO data sessions are constantly maintained within the CDMA2000 core network. This feature has important implications, since numerous smartphone applications that generate only modest amounts of data traffic can overload the network (e.g., RNCs and BSCs) from a signaling perspective.

Before deploying a new radio technology, there are several prudent steps that an operator can take to fully leverage its existing network resources. These initial steps, described below, share several important characteristics, most notably they only require a minimal amount of effort (e.g., changing network parameter settings) and they are fully compliant with existing standards.

4.1 Select Appropriate Dormancy Timer Settings

An operator must select appropriate network timer settings to balance the tradeoff between freeing up network capacity for data traffic and increasing the signaling traffic due to frequent network reconnections.

Operators can adjust various network parameter settings and one of the most important settings is the dormancy timer setting. This timer determines when an inactive device (e.g., smartphone) should drop its network connection, thus freeing up resources for other devices.

If a smartphone begins using the network to transfer data and then ceases this activity, there is a chance that the smartphone will resume the transfer and once again require network resources.

However, at a certain point in time the probability of a device connecting back to the network is substantially reduced. As a result, it no longer makes sense for the device to maintain its network connection since it is tying up network resources that could be used by other devices, causing additional radio interference and power consumption to keep the connection.

Conversely, if the dormancy timer setting was too short then the smartphone would frequently connect and disconnect from the network, and this constant back and forth activity would be inefficient from a signaling perspective. Consequently, an operator must select appropriate network timer settings that take into consideration user behavior in its network to balance these two tradeoffs. There is also a strong correlation between dormancy and battery life.

Recent analysis suggests that after two seconds of inactivity the probability that a smartphone will immediately reconnect to the network is extremely low; hence, operators should consider a dormancy timer setting in the range of two seconds. Infrastructure suppliers may also allow the operator to dynamically adjust these settings based on network load. For example, an operator could select a very low threshold for areas of its network that have a high concentration of traffic (e.g., a ballpark), but revert to a higher threshold in other parts of its network. Such an approach would ensure that any increase in signaling traffic associated with the lower setting is kept to a minimum while simultaneously freeing up network resources where they are needed the most.

4.2 Utilize Dynamic Load Balancing

Although DO Advanced offers sophisticated load balancing improvements, there are some relatively simplistic load balancing techniques that can be implemented on existing networks. Typically a mobile device selects a radio carrier based solely on the forward link channel quality, so long as the reverse link channel quality is acceptable. While this approach may correctly identify the radio carrier with the best channel quality, it does not take into consideration the amount of traffic in each carrier. Instead, the device could use a selection process that also took into consideration the slot occupancy or utilization of each radio carrier, thereby selecting the best radio carrier based on a combination of the number of active users and the amount of available/unused capacity.

4.3 Deploy Femtocells, Remote Radio Heads and Repeaters

To varying degrees, femtocells, remote radio heads and repeaters take advantage of the principle that a device which is closer to the serving cell site will deliver a better user experience; and by serving that device more efficiently with a higher bandwidth connection it will free up network resources for other devices. In the case of femtocells, the traffic is completely offloaded from the operator's macro radio access network, backhaul network, and eventually even its core network, thus providing an even greater benefit to operators.

In addition to increasing network capacity these solutions can be used to extend coverage into hard to reach locations, like inside residential homes or enterprises (e.g., femtocells) and at the edge of the macro cell or in lightly-covered regions of an operator's network (e.g., remote radio heads and repeaters). In the case of remote radio heads and repeaters they have the added benefit of providing coverage and capacity less expensively than deploying a large macro cell site.

4.4 Leverage the EVRC-B Codec

One of the key features of 1X Advanced is the EVRC-B (Enhanced Variable Rate Codec – Version B) voice codec. This codec improves voice quality without compromising network capacity, or it can be used to increase voice capacity without sacrificing quality. EVRC-B has already been commercialized in devices and chipsets, but not all operators have deployed the necessary software in their network infrastructure (e.g., 1X BSC) to enable its benefits. By deploying this software and leveraging the full capabilities of EVRC-B (along with Quasi-linear Interference Cancellation) an operator can increase its voice capacity by up to 50%.

4.5 Reduce Overhead and Optimize other Network Parameter Settings

When operators first deployed their networks they typically used the default settings that the industry believed were most appropriate at the time. While these various settings may have been appropriate given the available information, the industry has since learned a lot about how these various parameters can impact network performance. Further, user behavior has dramatically

Femtocells, remote radio heads and repeaters move the serving cell site closer to the user, thus improving the user experience and freeing up network resources for other devices.

By deploying 1X BSC software and leveraging the full capabilities of EVRC-B an operator can increase its voice capacity by up to 50%.

changed over the last few years (e.g., social networking) and this behavior was not taken into consideration when these settings were first determined.

For example, with the Quick Paging Channel feature, the mobile device wakes up for the first slot of the Paging Channel message to see if the network is trying to send a message. If no messages are in the queue, the mobile device goes back to sleep until the next Quick Paging Channel message is sent. Although this feature doesn't reduce the amount of signaling traffic or otherwise improve network performance, it does have a positive influence on battery life.

Each operator should work closely with its infrastructure partner to identify network parameters that it should modify to maximize the performance of its network.

Operators can also adjust how often a mobile device reports its Data Rate Control (DRC), which contains information about the quality of the channel to assign bandwidth or the appropriate data rate on the forward link. If this information is updated less frequently (e.g., when a device is not actively receiving data), then the amount of overhead traffic in the reverse link is reduced, freeing up capacity for other types of traffic.

Likewise, with the appropriate parameter settings, a mobile device may be able to reduce its transmit power when setting up a connection, thus reducing the interference associated with the access channels that are carrying data. Similarly, the network could reduce the interference from a connected mobile device by using optimized power control settings.

In the forward link, it is also possible to combine certain information that is sent in the control channels to free up network resources for additional traffic. Additionally, vendors are introducing features which limit the amount of signaling traffic that is being sent. For example, instead of sending a paging message to a device over a large number of cells and potentially even multiple radio carriers, the transmitted paging message could be isolated to only a select number of cells and to only a single radio carrier (e.g., where the mobile device was last present).

Each operator should, therefore, work closely with its infrastructure partner to identify network parameters that it should modify based on the capabilities of the infrastructure and the specific user behavior that the operator is experiencing in its network.

4.6 Use Data over Signaling

CDMA2000 1X and EV-DO support a feature known as Data over Signaling. This feature allows a very modest amount of data to be carried over a common signaling channel without establishing a traffic channel connection. This connectionless transmission mechanism is frequently used with PTT (Push-To-Talk) to improve the user experience – faster connection time. But it can also be used in other situations, such as transmitting short burst data. In addition to freeing up some network resources that would otherwise be carrying the data that is now being carried over the signaling channels, this feature can improve the user experience by making at least some content readily available to the user while the bulk of the data is still being sent in the background. Further, to the extent that only a limited amount of data needs to be sent (e.g., a “keep alive” message), the device would not have to establish an active connection, thereby reducing signaling traffic and increasing battery life.

4.7 Fully Leverage Available Spectrum and the Capabilities of the Air Interface

While perhaps an obvious solution to addressing the ever-growing amount of data traffic, operators should take full advantage of their available spectrum. If an operator has used its entire spectrum, it becomes impossible to deploy additional radio carriers without cell splitting. Until that point is reached, an operator should put to use the spectrum that it has at its disposal.

It is also possible that operators could have deployed ample radio network capacity, but underdimensioned their network in other areas, like the backhaul. Due to its significant expense, operators are sometimes hesitant to deploy sufficient backhaul capacity. Yet, unless the amount of bandwidth in the backhaul network is commensurate with the capabilities of the radio network, an operator will not be able to fully leverage its network.

5.0 Maximizing Existing Network Resources through Advanced Technology Features

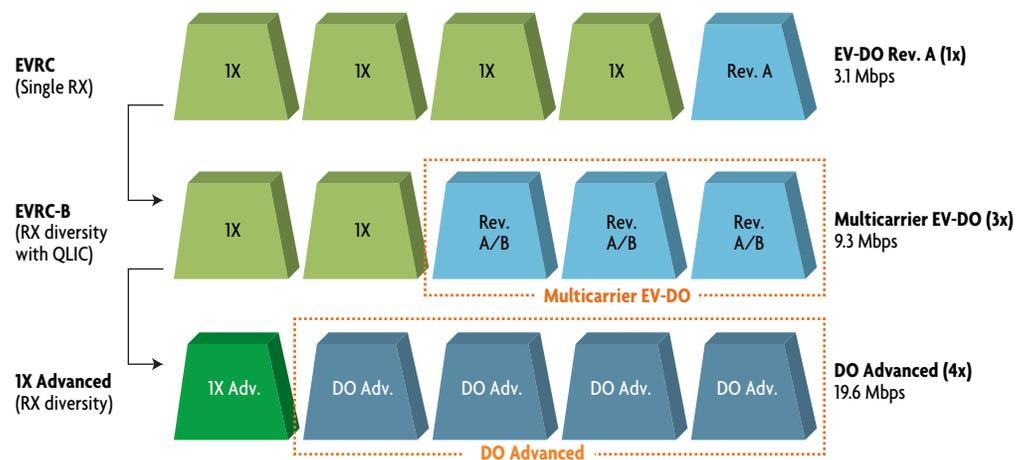
The CDMA2000 roadmap allows an operator to efficiently manage network capacity with backward compatible enhancements that minimize the need for new hardware and abide by existing standards. Most of the enhancements are enabled through software upgrades or they are introduced through new device and cell site baseband chipsets. These last two decisive features are critically important since the capacity gains associated with other technologies require new devices/chipsets or in some cases entire new networks (e.g., LTE).

5.1 1X Advanced

1X Advanced allows an operator to more efficiently deliver its voice traffic, thus freeing up additional spectrum for data traffic channels.

1X Advanced includes a collection of enhancements that can theoretically increase voice capacity by up to 100 Erlangs, or a factor of four over legacy 1X networks.² To put this into perspective, 1X Advanced can support more than 80 times the number of calls than older analog networks. Some of the enhancements are standards-based (e.g., CDMA2000 1X Release E), while others are not. While seemingly a “voice only” feature, 1X Advanced has meaningful implications for data traffic, femtocells, coverage, and 1X data network capacity. As illustrated in Figure 5, once fully implemented in the network, an operator can consolidate voice traffic on fewer carriers and repurpose that spectrum to support additional EV-DO carriers or possibly LTE, if additional contiguous spectrum is available.

Figure 5. Evolutionary Steps to Greater Network Capacity by Freeing up Spectrum



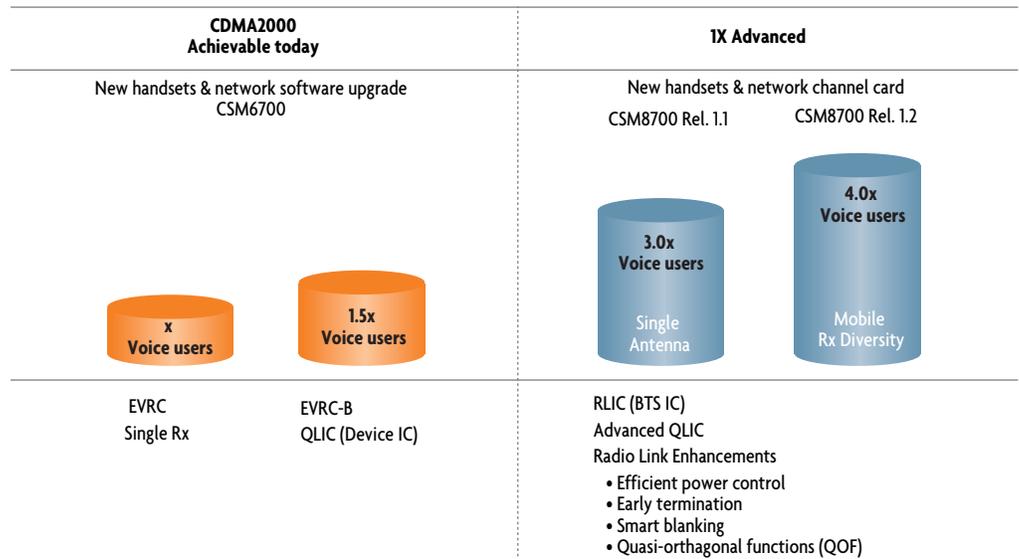
Source: Various CDG member companies

1X Advanced is based on several discrete technology enablers which, as shown in Figure 6, can enable a fourfold gain in voice capacity when combined.³

² All performance claims presented in this whitepaper are based on simulation studies conducted by various CDG member companies.

³ Based on simulations conducted by various CDG member companies.

Figure 6. 1X Advanced Features and their Impact on Increasing Voice Capacity

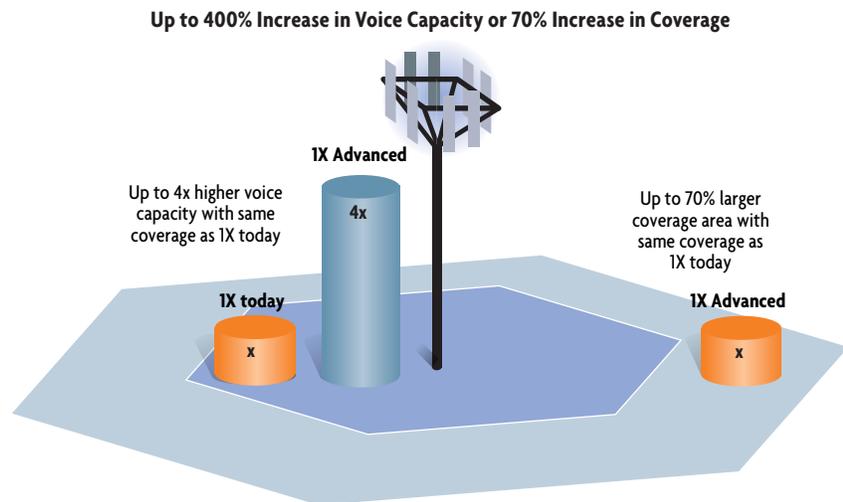


x = Today's baseline capacity for CDMA2000 1X; QLIC = Quasi-linear Interference Cancellation; RLIC = Reverse Link Interference Cancellation (RLIC), also known as Base Transmission Station IC

Source: Various CDG member companies

1X Advanced also contains certain interference cancellation features to improve the performance of femtocells in a 1X network. Furthermore, 1X Advanced can trade off voice capacity to increase 1X coverage by up to 70% as shown in Figure 7.⁴

Figure 7. 1X Advanced Tradeoff between Voice Capacity and Coverage



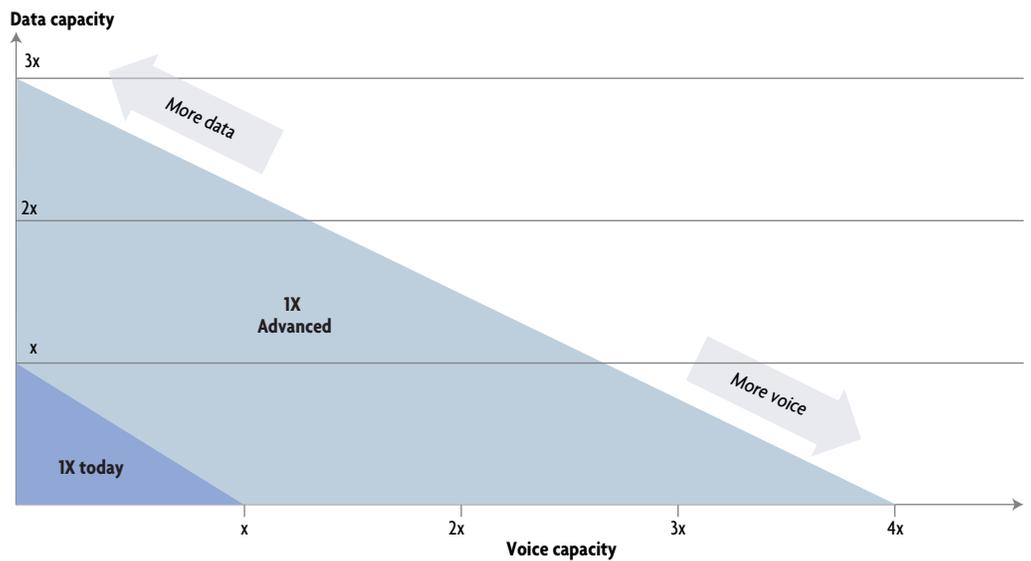
Note: Based on Qualcomm simulations, coverage is defined as the maximum area with less than 1% of the users in outages
Assumptions: 3GPP2 simulation frame work, embedded sector, with all the 1X Advanced features considered including MRD, FL and RL Interference Cancellation, new RC

Source: Various CDG member companies

⁴ Ibid.

If additional voice capacity is not necessary, it can be traded off for additional 1X data network capacity. In effect, the 1X channel is used to transfer packets of data instead of packets of voice traffic. Increasing the data capacity of a CDMA2000 1X network by up to a factor of three⁵ is beneficial for rural/low data demand areas in emerging markets, where broadband connectivity is not deployed or economical. Based on their needs, operators can choose to dimension their networks such that the network operates at a point which gives both coverage and capacity improvements.

Figure 8. 1X Advanced Tradeoff between Voice and 1X Data Network Capacity



Thus, 1X Advanced can substantially improve the competitiveness of an operator by enabling:

- More efficient use of limited spectrum
- Lower cost per call or more available minutes of use
- Frees up spectrum for broadband data
- Up to 70% greater 1X coverage (tradeoff)
- Up to a threefold increase in 1X data network capacity (tradeoff)
- Improved femtocell performance

5.1.1 A More Efficient Voice Codec

1X Advanced includes the use of a more advanced coding scheme (codec), called EVRC-B, to digitize the voice communications before they are carried over the network. EVRC-B offers more capacity with the same voice quality as EVRC. And, if a portion of the network becomes congested, the operator has the flexibility to surrender some voice call quality to further increase voice capacity.

⁵ Ibid.

EVRC-B requires a new chipset in the devices and a software upgrade at the base station controller to support the new codec. EVRC-B is now an inherent feature in many commercially available CDMA2000 chipsets; therefore, the increased voice capacity promised by this feature is almost a “freebie” other than the software requirement at the BSC.

5.1.2 Interference Cancellation

There are several types of interference cancellation (IC) that are associated with 1X Advanced.

5.1.2.1 QLIC

QLIC (Quasi-linear Interference Cancellation) reduces the interference in the forward link, or directional path between the serving base station and the mobile device. The interference is due to other users within the cell or in adjacent cells, as well as the interference that is created by the pilot channel, or that portion of the transmitted signal that contains information that the mobile device needs to communicate with the network and to decode the transmitted information.

By reducing the amount of interference at the mobile device, the serving base station can transmit at a lower power and still maintain the voice call. In turn, the power savings can be used to support additional voice calls within its serving area.

5.1.2.2 Advanced QLIC

Increasing voice capacity requires additional Walsh codes, beyond the existing 128 orthogonal Walsh codes. Whenever the forward link is limited by Walsh codes, 1X Advanced makes use of Quasi-orthogonal functions (QOF). However, QOF creates additional interference among users. Therefore, Advanced QLIC is used in 1X Advanced to improve the performance of the forward link by canceling the additional interference from QOF and the interference from base stations in the neighbor set.

QLIC is included in most existing CDMA2000 chipsets and both QLIC and Advanced QLIC are an inherent part of all future CDMA2000 chipsets. Depending on an operator’s requirements, it can select what percentage of its installed base should take advantage of the feature or only enable the feature with its heaviest voice users.

5.1.2.3 RLIC

1X Advanced also introduces interference cancellation in the reverse link, or the directional path between the mobile device and the serving base station. Reverse Link Interference Cancellation (RLIC), also known as BTS Interference Cancellation, requires a new channel card in the base station that supports the feature, but it will be a standard feature in all future cell station modem (CSM) chipsets.

In addition to increasing network capacity, RLIC can be used to improve network coverage. Typically, the effective range of a cell site is limited by the reverse link since the transmit power of the mobile device cannot exceed a certain threshold due to regulatory restrictions as well as the desire to extend the life of the battery. With RLIC the improved signal that is received at the cell site means that the device can transmit from a greater distance or in a harder to reach location (e.g., indoors) and still maintain the connection. In other words, the capacity gains associated with RLIC can be used to increase capacity or improve coverage at current capacity levels. The full benefits of this feature can also be realized immediately and equally across both 1X Advanced and legacy CDMA2000 handsets, since RLIC is introduced at the cell site.

Upon installing new 1X Advanced channel cards, the full benefits of RLIC can be realized immediately and equally across both 1X Advanced and legacy CDMA2000 devices, since the feature is introduced at the cell site.

With Mobile Receive Diversity (MRD), the base station is able to transmit at a lower power level and still maintain the voice connection, thus freeing up network resources for additional voice calls and data sessions.

5.1.3 Mobile Receive Diversity

Mobile Receive Diversity (MRD) leverages two mobile device antennas to boost the quality of the received signal for a given transmit level at the serving base station. As such, the base station can reduce the amount of power required to support the connection, thus freeing up power that can be used to support additional voice calls and data sessions.

MRD is presently available in most 1X and EV-DO chipsets, although it is at the discretion of the device manufacturer and/or mobile operator whether or not the feature is enabled. MRD is widely deployed in data centric devices, such as USB dongles, and it is starting to appear in high-end phones to support higher voice quality, not to mention better data performance. It is also believed that some operators will mandate its use in CDMA2000 handsets and begin seeding the market so they can realize up to a 2.3x voice capacity gain after installing the new channel cards and assuming all the handsets support MRD. Based on independent chipset testing, the effect that it can have in improving device performance under challenging network conditions is well documented – in some cases more than doubling the throughput over a single antenna configuration.

MRD can also help increase battery life since downloads can occur more quickly, enabling the transmitter to be disabled sooner. Although a second receiver can modestly increase power consumption, the transmitter causes most of the drain on the battery. We note that a mobile device must transmit acknowledgement (ACK) messages to confirm that it has received the transmitted packets. Therefore, as long as the ACK packets are being sent the transmitter must remain on. It is difficult to quantify the improvement in battery life since the savings is application specific. For example, since the benefit is only realized when the length of the download time is decreased, there wouldn't be any benefit associated with watching a streaming video or maintaining a voice call for a fixed duration of time.

5.1.4 Physical Layer Enhancements

Separate from the new codec and various interference cancellation mechanisms, 1X Advanced introduces new changes to the radio configuration or Physical Layer. These changes, which were introduced in CDMA2000 1X Release E (September 2009), are supported in all new 1X Advanced CSM chipsets. The following features further optimize the air interface by leveraging the inherent characteristics of a voice call to reduce its signaling and traffic overhead.

5.1.4.1 Smart Blanking

“Smart blanking” eliminates the transmission of background noise that typically consumes air interface resources, unless it changes. To the extent the background noise remains constant – a likely scenario in most voice calls – the amount of data that the voice call requires is reduced.

5.1.4.2 Efficient Power Control

1X Advanced reduces the overhead associated with the power control information that is shared between the network and the mobile devices. By reducing how often and how this information is exchanged, additional resources are made available to support more voice calls.

5.1.4.3 Frame Early Termination

By introducing Frame Early Termination (FET) in both the forward and reverse link, the base station (mobile device) does not have to transmit an entire frame if the mobile device (base station) has already successfully decoded the information and sent an acknowledgement (ACK)

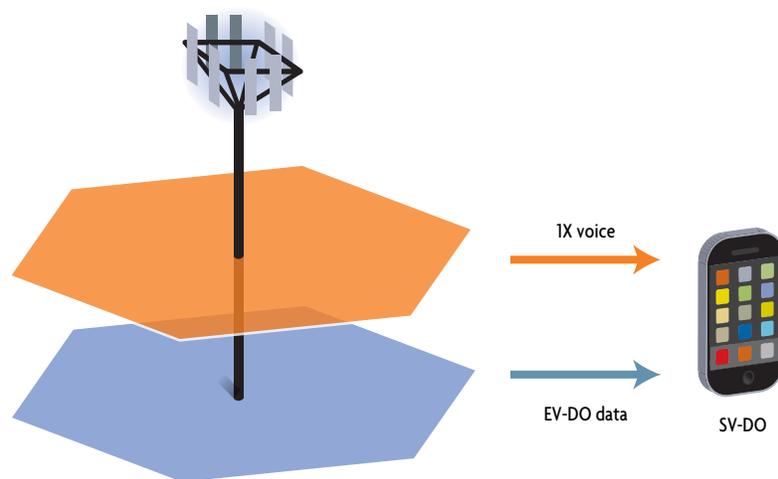
1X Advanced reduces signaling and traffic overhead by minimizing the transmission of “blank” voice packets, power control information and acknowledgement receipts.

receipt. Similar to the gains achieved by the H-ARQ feature in EV-DO, FET reduces the amount of transmit power over a duration of time, thus increasing overall system capacity.

5.1.5 Simultaneous Voice and Data (SVDO)

Simultaneous voice and data sessions will be enabled on CDMA2000 networks starting in 2011. This capability means CDMA2000 users will be able to access the web, send emails, play games, view photos, watch videos and use data intensive GPS navigation options like updating maps or downloading traffic info while simultaneously communicating on a phone call. Previously, if a subscriber accepted a voice call while using the data connection (e.g., while viewing a video), the data session would be suspended.

Figure 9. Simultaneous Voice and Data



Source: Various CDG member companies

Although not directly tied to the introduction of 1X Advanced, from now on EV-DO chipsets will support SVDO through the use of a separate transmitter and receiver for the active 1X voice and EV-DO data traffic. This feature doesn't necessarily increase network capacity, but it is a compelling feature that is enabled through new handsets and does not require changes to the standard or existing infrastructure. Dual-antenna handsets supporting SVDO will become available in the first half of 2011.

5.1.6 Market Opportunities and Implications

While generally considered as a technology to increase voice capacity or coverage, 1X Advanced can free up spectrum for EV-DO to support additional data traffic.

While generally viewed as a technology to increase voice capacity or coverage, 1X Advanced is attracting operator interest and commitment due to its ability to free up 1X radio channels (e.g., spectrum) and repurpose them to support additional data traffic.

North American CDMA2000 operators, such as Sprint, who are facing ever-increasing data usage on their networks, are looking to use 1X Advanced to free up spectrum and support the continued growth voice and data traffic.

Given the inherent need for wider bandwidth channel allocations (e.g., 10 MHz or greater of contiguous spectrum) to obtain the capacity benefits of LTE, some operators are challenged

to deploy LTE in their existing bands.⁶ Their limited spectrum position combined with ever-increasing voice traffic further complicates matters. Yet, the introduction of 1X Advanced can go a long way toward freeing up existing spectrum or at least limiting the need to utilize even more spectrum for future voice traffic.

CDMA2000 operators in India are perhaps in the best position to benefit from the capabilities of 1X Advanced. Due to regulatory restrictions limiting the amount of available spectrum, these operators can only deploy a very limited number of radio carriers (e.g., no more than 3-4 radio carriers); yet, they want to dedicate one, if not two, carriers to EV-DO to offer a competitive mobile broadband data service. In the case of MTS India, they intend to leverage 1X Advanced to increase the voice capacity in their network and free up spectrum for EV-DO.

Operators who want to achieve the full capabilities of 1X Advanced must fully populate their network with 1X Advanced devices. Therefore, operators will need to aggressively push 1X Advanced handsets to realize their full benefit, while also recognizing that the capacity gains will take place gradually over time, or as the penetration rate of 1X Advanced handsets increases.

Also, given the competitive market environment, these operators are considering SVDO to support simultaneous 1X voice and EV-DO data services.

5.2 EV-DO Revision B

EV-DO Revision B (Rev. B) is implemented in two phases. The first phase, known as Multicarrier EV-DO, is a software upgrade that leverages existing EV-DO Rev. A channel cards by aggregating multiple carriers to triple the data rate in 5 MHz of FDD spectrum or double the network capacity, especially for bursty smartphone traffic.⁷ The increased data network capacity is due to the trunking efficiency associated with logically combining multiple carriers. The term “logical combining” is used since each original EV-DO Rev. A carrier remains unchanged with Multicarrier EV-DO.

This feature also ensures backwards compatibility with legacy devices and infrastructure. With Multicarrier EV-DO, unique data streams can also be simultaneously sent on different EV-DO Rev. A carriers, thus increasing the throughput to and from an MRD device.

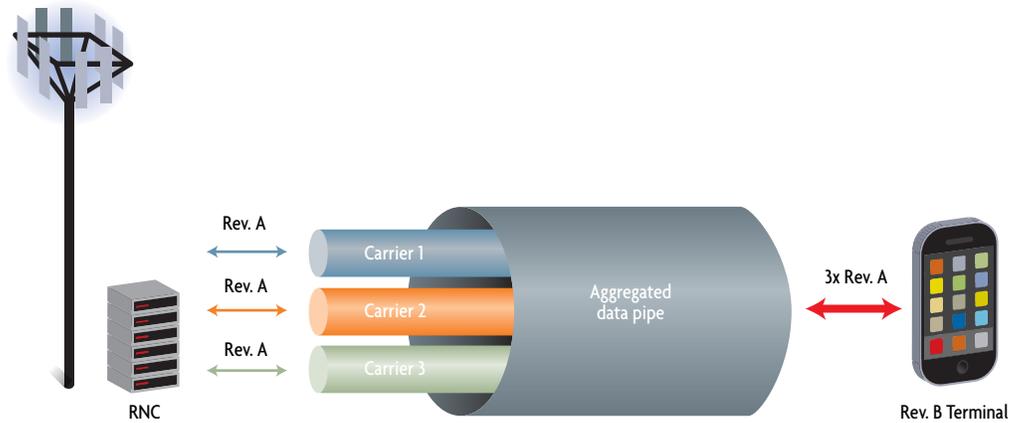
As shown in Figure 10 (following page), Multicarrier EV-DO combines up to three EV-DO carriers, thus increasing the forward and reverse link peak data rates to 9.3Mbps and 5.4Mbps, respectively within 5MHz of FDD spectrum (3x1.25MHz carriers, plus guard band), while decreasing latency for delay-sensitive applications. Presuming the operator has already deployed two or more EV-DO Rev. A carriers in their base stations, the upgrade path to Multicarrier EV-DO only requires a software upgrade at the base stations and at the RNC. Note that the EV-DO carriers do not need to be adjacent to each other. And while the performance enhancements associated with Multicarrier EV-DO require new EV-DO Rev. B devices, legacy EV-DO Rev. A devices will also witness an improvement in performance since the overall network will be less loaded.

Multicarrier EV-DO can combine up to three EV-DO carriers, thus increasing the forward link and reverse link peak data rates to 9.3Mbps and 5.4Mbps, respectively within 5MHz of FDD spectrum, or it can maintain existing Rev. A data rates and potentially double network capacity.

⁶ While OFDM-based technologies, such as LTE, can be deployed in increments of 1.4 or 5MHz, the full benefits are only seen when carriers are deployed in increments of 10 or 20MHz.

⁷ Based on simulations conducted by various CDG member companies.

Figure 10. Multicarrier EV-DO



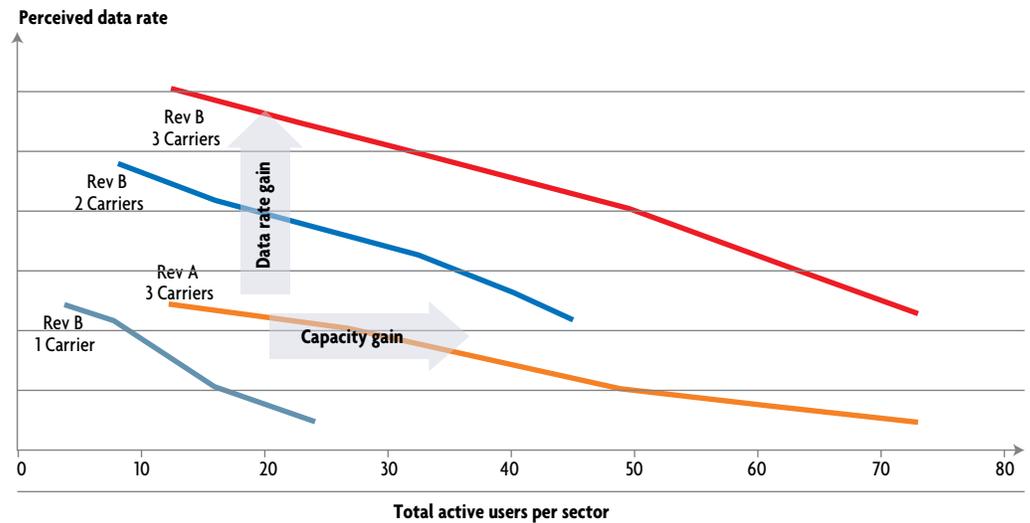
Source: Various CDG member companies

Multicarrier EV-DO can also have a favorable impact on battery life which improves the utility of the device and the overall user experience. When a given amount of data is transferred faster with Multicarrier EV-DO, the radio is able to return to its dormancy state much quicker than with EV-DO Rev. A. This capability is especially important since a radio's transmitter requires substantially more battery consumption than its receiver. Similarly, to the extent a Multicarrier EV-DO device no longer requires network resources; it frees up additional resources for single-carrier EV-DO Rev. A devices, thus indirectly benefiting all devices in the network in terms of both user data rates and battery life.

Multicarrier EV-DO is commercially available today. According to the CDG, up to eighteen operators have or are planning to deploy Multicarrier EV-DO Rev. B, and seven EV-DO Rev. B devices from four manufacturers are commercially available.

The second phase, the full implementation of EV-DO Rev. B, introduces even higher data rates by enabling a higher order modulation scheme, 64QAM, to increase the forward link peak data rate to 14.7Mbps (4.9Mbps per 1.25MHz carrier) while the reverse link data rate remains the same. The full upgrade requires an EV-DO Rev. B (CSM 6850) channel card and new handsets. The new CSM has more processing power and supports more individual carriers than its predecessor, resulting in base station cost savings and enabling an advanced interference cancellation solution that improves reverse link capacity and provides improved efficiencies for typical smartphone applications where the reverse link overhead is a significant percentage of the total interference (since the amount of data transferred per connection is small). Therefore, the CSM delivers additional benefits beyond the higher data rates, including reduced battery consumption.

Figure 11. EV-DO Rev. B Performance Gain with Bursty Traffic

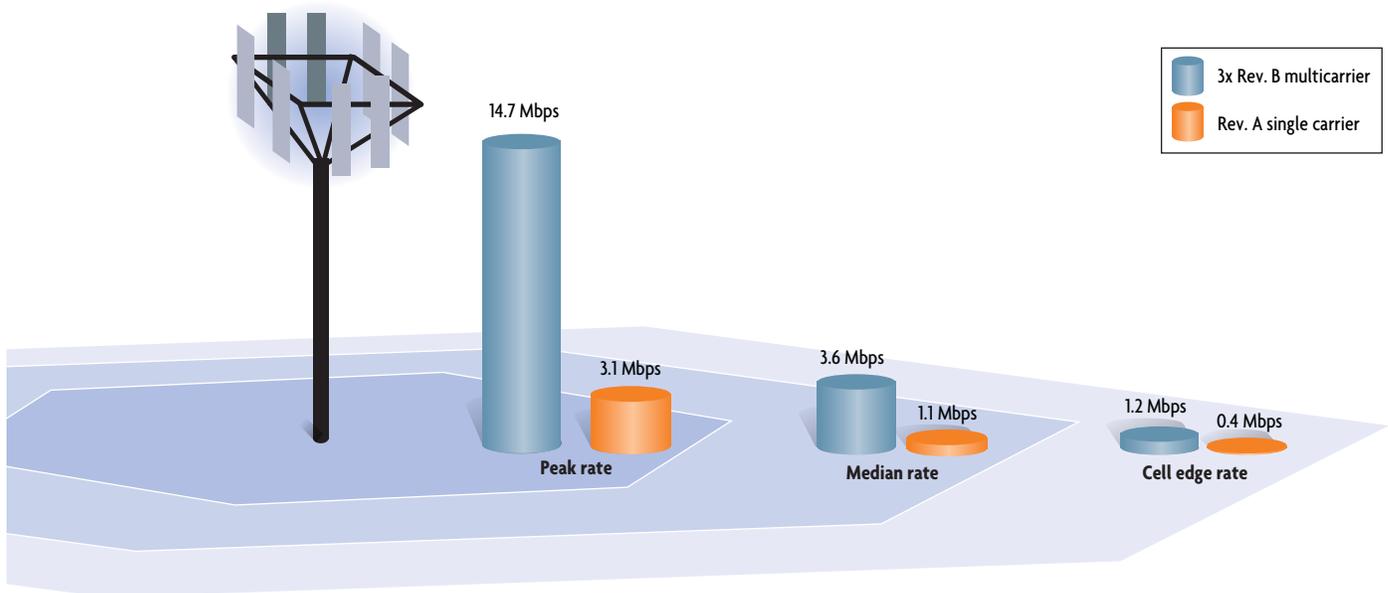


Source: Various CDG member companies

In addition to higher data rates, EV-DO Rev. B is well suited for bursty smartphone applications.

Finally, the use of multiple carriers benefits users at the cell edge, as shown in Figure 12. In this instance the average user data rate could be double or triple the speed that it would otherwise have been. This feature is important, since it helps deliver a more consistent user experience across the entire network and especially inside of buildings. With a multicarrier network, improved robustness (i.e., fewer service interruptions) can be achieved, which in turn leads to better performance of popular applications such as video chat and web browsing.

Figure 12. EV-DO Rev. B User Data Rate Improvements



Note: Peak data rates are as defined by the standard. The median and cell edge rates are based on simulations using mixed channel model with EV-DO Rev. B devices supporting 64QAM
Source: Various CDG member companies

5.2.1 Market Opportunities and Implications

The first Multicarrier EV-DO network was launched by Smart Telecom in Indonesia, in January 2010, earlier than the industry’s first dual-carrier HSPA network. Pakistan Telecommunication Company Ltd. (PTCL) and KDDI in Japan followed Smart Telecom’s lead. KDDI’s commercial launch is notable since it has an installed subscriber base of more than 32 million subscribers and included the introduction of several new EV-DO Rev. B handsets. Figure 13 illustrates some of the EV-DO Rev. B devices that are or will be available.

Figure 13. EV-DO Rev. B Devices



Source: Smart Telecom, Pakistan Telecom Company and KDDI

Major CDMA2000 operators in China, India, Russia, Africa, South East Asia, and Europe represent some of the operators around the world who are likely to follow KDDI and deploy EV-DO Rev. B.

Besides these operators, there is also interest in EV-DO Rev. B in China, India, Russia, Africa, South East Asia, and Europe. China Telecom, with an installed base of nearly 86 million subscribers at the end of October 2010, has conducted trials with its infrastructure partners. Indian and Chinese CDMA2000 operators could create a much-needed ecosystem of devices that can support launches by smaller operators who lack the ability to create an ecosystem on their own. For the moment these larger operators have not publicly committed to launching commercial EV-DO Rev. B services.

The larger North American operators who have access to lots of spectrum and capital do not have plans to deploy EV-DO Rev. B, although they continue to aggressively deploy additional EV-DO Rev. A CSM 6850-based channel cards (that support EV-DO Rev. B) to meet the current and expected growth in 3G mobile data traffic. Given their comments and their strong interest in deploying LTE in new frequency bands, it is unlikely that these operators will deploy EV-DO Rev. B, although they may have interest in certain DO Advanced features.

5.3 DO Advanced

The basic tenet of DO Advanced is to make existing EV-DO networks and devices perform as efficiently as possible by exploiting unevenly loaded networks, and to do so without substantial changes to existing infrastructure. The software-upgradeable Smart Network techniques enabled by DO Advanced appear to be fairly easy to implement, and their impact on overall network performance can be quite dramatic. And to the extent an operator takes advantage of newly enhanced devices, additional gains can be realized.

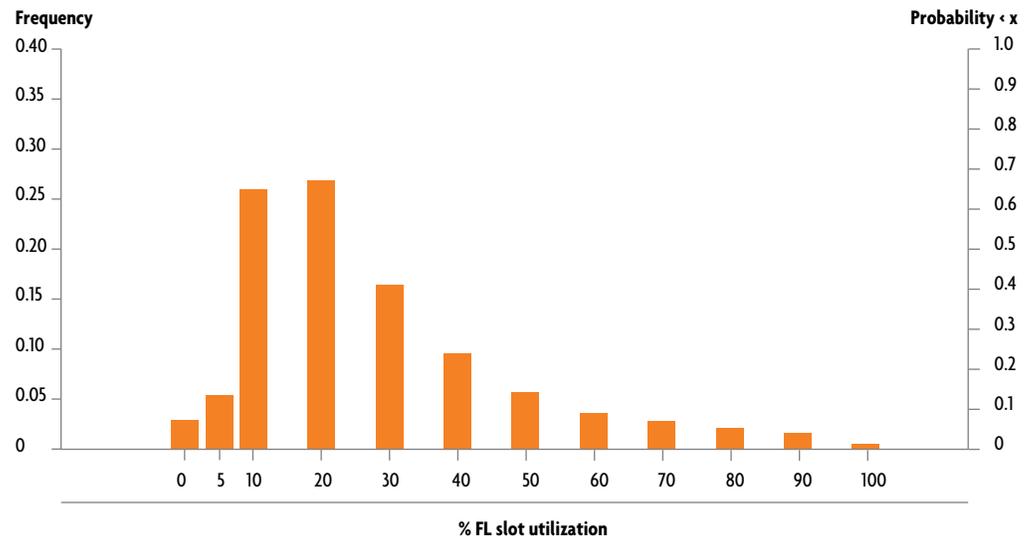
Since network loading continuously changes with time and location, data traffic within a network is difficult to predict and it is never uniformly distributed. When an operator designs a network, it bases its network planning on the peak busy hour, or the maximum amount of traffic that occurs within a given day. Historically, the busy hour was during morning or evening rush hour, when voice traffic dominated data traffic. More recently, data traffic dwarfs voice traffic and the peak busy hour for data traffic is more flat, meaning that the peak usage occurs over a longer period of time.

DO Advanced introduces software-upgradeable Smart Network techniques that exploit unevenly loaded networks to improve network capacity and performance without requiring new channel cards or devices.

However, one thing remains unchanged. The distribution of traffic within the network is never uniformly distributed since people are mobile and connect to the network whenever they desire. Therefore, it is quite common for a heavily-loaded cell to be surrounded by lightly-loaded cells, or at least cells that contain less traffic and the location of the heavily-loaded cells are frequently changing throughout the day. As an example, a major North American operator commented that the portion of its network covering a major college football stadium witnesses on average a 330 percent increase in voice traffic, a 160 percent increase in data traffic, and a 40 percent increase in text messaging when a football game is played in the stadium. Had there been more capacity available in the surrounding network, the percentages would have been even higher.

Figure 14 illustrates the typical network load on a per sector basis for an operator's network during busy hour. As shown, the distribution of the traffic is far from being uniform across the network and only a small percentage of the sectors are close to being capacity constrained.

Figure 14. Typical Network Load Distribution During Busy Hour



Source: Various CDG member companies

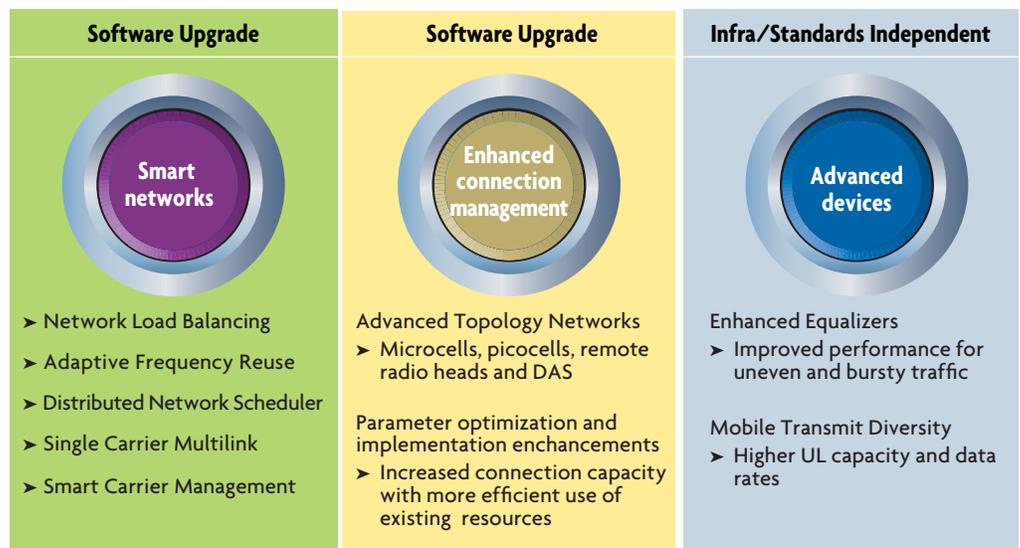
The unevenness of the traffic has two implications. First, since most of the available capacity in the network is seldom utilized, the operator has excess capacity that is going to waste. Second, during those brief periods of time when the traffic with any given cell site is reaching its designed peak load, the user experience within the cell (e.g., throughput and latency) is negatively impacted. For example, at the aforementioned football stadium, fans readily acknowledge difficulties using their mobile phones once the game starts and the stadium is full.

Ideally, an operator would be able to dynamically reposition network resources where and when it is needed throughout the day – much like a general repositions his troops on the front line of a battle when the enemy breaks through. Thus, DO Advanced is quite dynamic in how it addresses uneven network loading with fixed network assets.

DO Advanced enables an operator to cost-effectively add data network capacity when and where it is needed.

DO Advanced is comprised of three basic building blocks: Smart Networks, Enhanced Connection Management and Advanced Devices. Each of these blocks, shown in Figure 15, is described in the next sections.

Figure 15. DO Advanced Features



Source: Various CDG member companies

5.3.1 Smart Networks

The term Smart Networks refers to five unique features that can substantially increase the data capacity of an existing network with existing devices, without deploying more cell sites or using more spectrum. Many of these features do not require the deployment of EV-DO Rev. B, and they can be applied to networks that have multiple EV-DO Rev. A carriers, yet in some cases there is an additional performance gain in multicarrier configurations or if EV-DO Rev. B devices are also present in the network.

5.3.1.1 Network Load Balancing

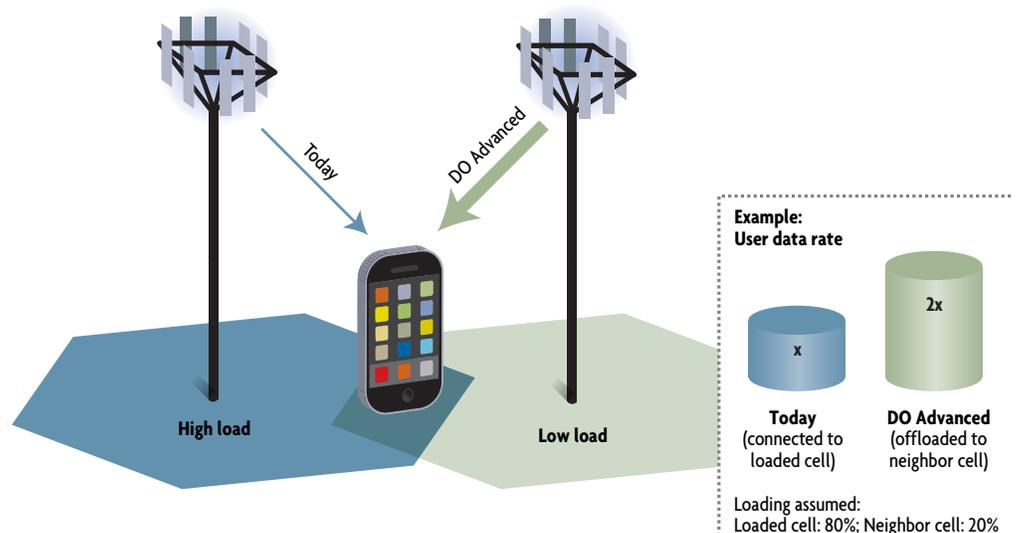
Network Load Balancing utilizes unused network capacity of lightly-loaded neighboring cells. This Smart Network technique opportunistically reassigns mobile devices from heavily-loaded sectors to lightly-loaded sectors, even if the quality of the radio channel is actually better in the

Network Load Balancing utilizes unused network capacity of lightly-loaded neighboring cells to achieve more than double the data rate under certain loading conditions.

heavily-loaded sector. The end result is increased network capacity, reduced backhaul bottlenecks and improved data rates for both the offloaded users and the users in the loaded cell.

Imagine a mobile device that can connect to two different cell sectors. With existing implementations the device will always select the sector (carrier frequency) that has the highest Data Rate Control (DRC) value, since the higher the DRC, the higher the potential data rate that the mobile device can receive. Now suppose the carrier frequency with the highest DRC is also heavily loaded with traffic such that there is insufficient bandwidth to support the reported DRC data rate value. If so, the mobile device may be better off being connected to a sector with a lower DRC value, as long as the data rate associated with that DRC value is greater than the available bandwidth (e.g., fraction of time allocated to a user) in the capacity-constrained sector.

Figure 16. Network Load Balancing



Source: Various CDG member companies

As illustrated in Figure 16, the mobile device that is located between the two cells/sectors prefers the highly loaded sector since it offers a higher DRC value – indicating that it should be able to receive the highest data rate from that cell. Yet, unbeknownst to the mobile device, the sector is heavily loaded and has insufficient capacity to support that data rate, so the device achieves a much lower data rate. Had the mobile device been assigned to the adjacent lightly-loaded sector – the sector with the lower DRC value – its data rate would have increased two-fold.

The Network Load Balancing feature requires software changes in the infrastructure and it is compatible with EV-DO Rev. A and EV-DO Rev. B. Legacy EV-DO devices support the feature without any changes. However, newer devices that are enhanced to better support Network Load Balancing (e.g., perform server (sector) selection taking the channel condition and load into account) can respond to changes more quickly.

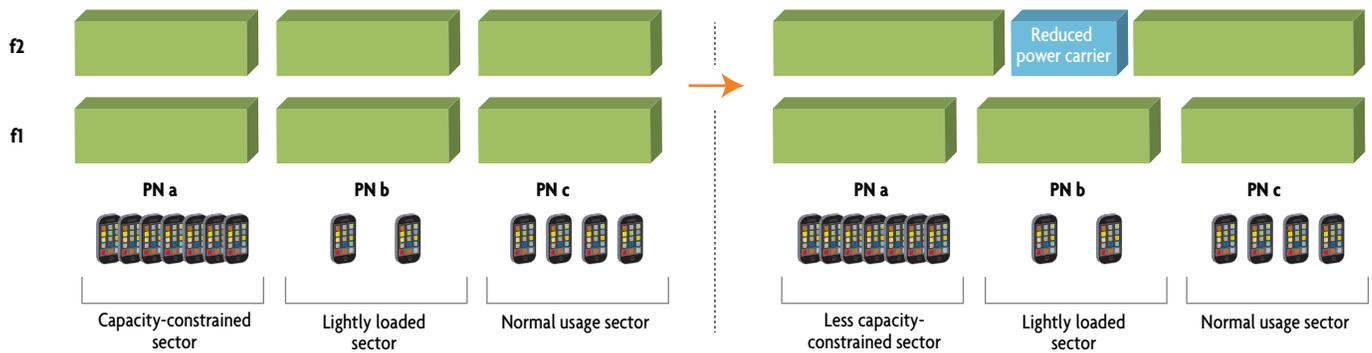
Adaptive Frequency Reuse reduces interference by adjusting the transmit power of lightly-loaded cells to increase overall network capacity and improve data rates.

5.3.1.2 Adaptive Frequency Reuse

Adaptive Frequency Reuse, also known as Demand Matched Configuration, reduces the interference in a capacity-constrained cell by adjusting the transmit power of the adjacent lightly-loaded cells. This increases overall data network capacity and improves data rates.

Within a typical network deployment of multiple EV-DO Rev. A carriers, all carriers across all cell sites are transmitting at maximum power, if for no other reason, to ensure that coverage persists between cells. The left-hand side of Figure 17 illustrates this concept.

Figure 17. Adaptive Frequency Reuse



Source: Various CDG member companies

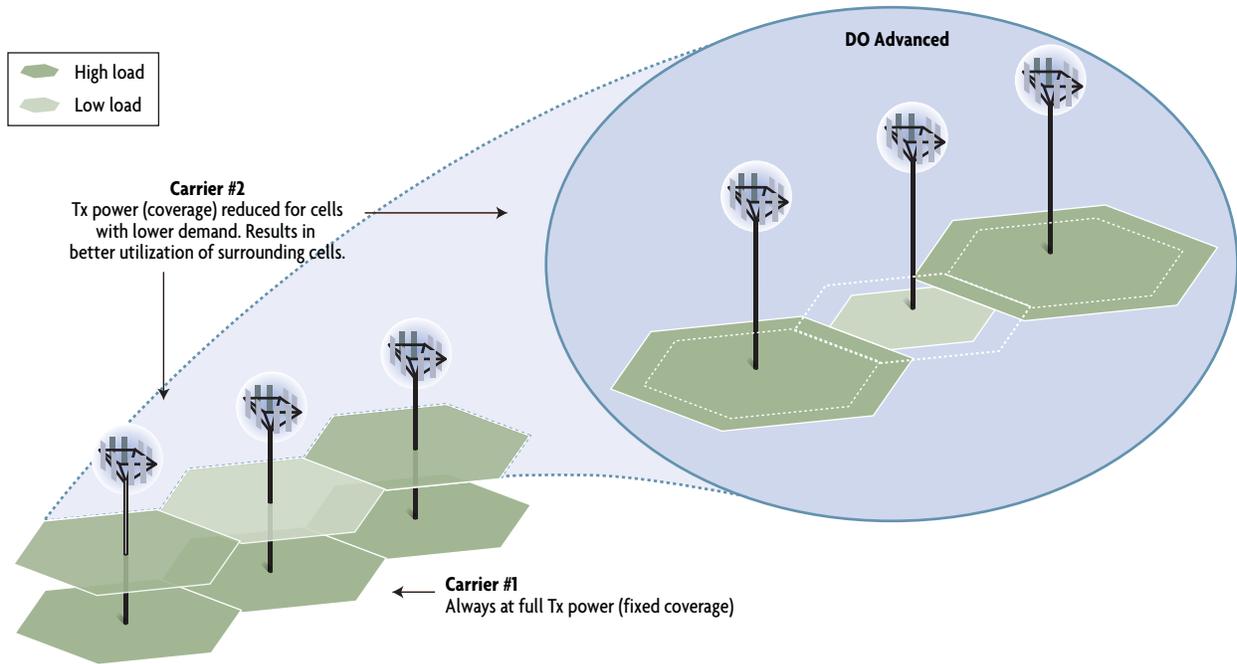
However, in reality, only one of the carrier frequencies in each cell needs to be transmitting at full power to ensure ubiquitous coverage while the remaining carriers can be transmitting at a lower power level, as shown on the right-hand side of Figure 17. In the case of a capacity-constrained cell that is adjacent to lightly-loaded cells, the secondary carrier frequencies in the lightly-loaded cells could transmit at a lower power level, thus reducing the detectable interference in the capacity-constrained cell and increasing its available capacity.

Another way to look at the benefit of the Adaptive Frequency Reuse feature is to consider the spatial view shown in Figure 18 (following page).

In effect, the network coverage remains the same since the primary carrier frequency at each cell site is being transmitted at the maximum power level, while the secondary carriers dynamically adjust their forward link transmit levels according to the amount of traffic that is being generated in each cell site. As the secondary carrier frequencies adjust their transmit power levels, the effective radius shrinks or expands, but since these carriers are being used for capacity reasons and the primary carrier is being used for coverage purposes, the operator is able to maximize available network resources while preserving its original network coverage requirements.

Adaptive Frequency Reuse is compatible with EV-DO Rev. A and EV-DO Rev. B. It does require software changes in the network, but there are no changes required to any devices.

Figure 18. Spatial View of Adaptive Frequency Reuse



Source: Various CDG member companies

5.3.1.3 Distributed Network Scheduler

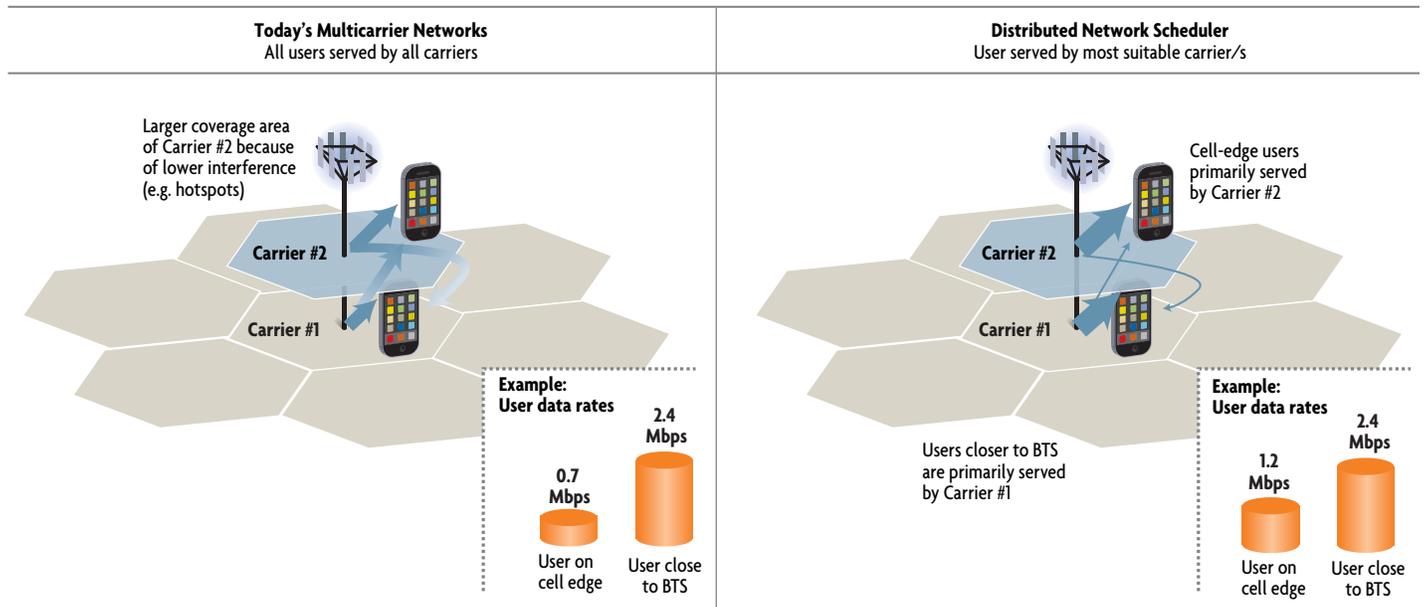
As the name suggests, the Distributed Network Scheduler is a multicarrier feature that prioritizes and allocates bandwidth (e.g., time allocated) to multiple users at the network level. In a typical mobile broadband network, bandwidth is distributed to all mobile devices using a proportional fair scheduler. With such a scheduling algorithm the mobile device with the best channel conditions (e.g., reported DRC) receives the most bandwidth, while the mobile device with the worst channel conditions receives the least amount of bandwidth. In other words, the bandwidth is assigned to each mobile device in proportion to the channel conditions of each device.

Given that channel conditions often fluctuate, it is highly likely that the conditions reported by a set of mobile devices can change in the next scheduled slot such that a mobile device which previously had poor channel conditions now has relatively favorable channel conditions and therefore receives more bandwidth, and vice versa. If a mobile device remains in an area with poor channel conditions, such as at the cell edge, the *fairness* part of the algorithm takes effect.

The Distributed Network Scheduler maximizes network capacity by prioritizing and allocating bandwidth across multiple carriers, serving sectors and cells.

With the Distributed Network Scheduler, the distribution of available bandwidth (e.g., allocated time) to each user is prioritized and allocated across multiple carriers, serving sectors and cells within the active set based on the relative RF environment (SNR), thus maximizing the overall efficiency of the network and taking the concept of a proportional fair scheduler to a whole new level, especially in hotspots. Proportional fair scheduling is no longer limited to a sector-carrier level. With this Smart Network technique, it is now extended to a network-level.

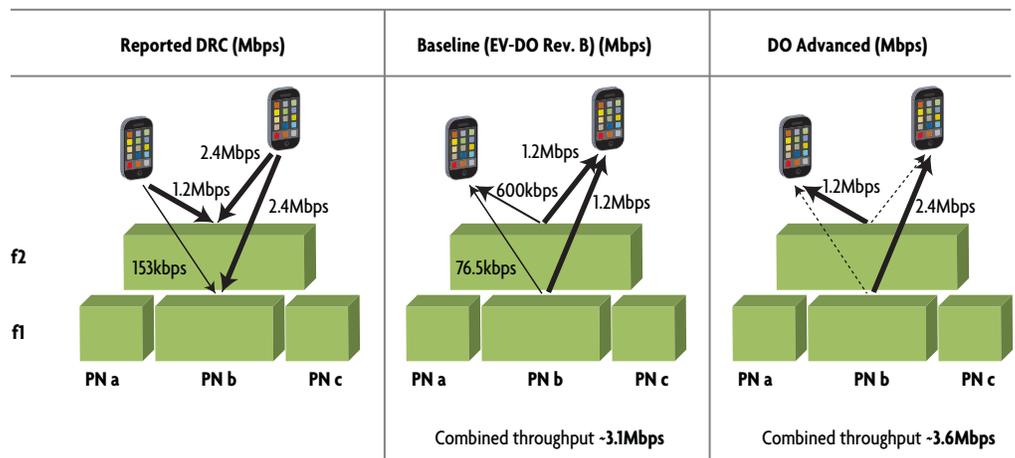
Figure 19. Distributed Network Scheduler



Source: Various CDG member companies

Figure 20 illustrates the concept of a Distributed Network Scheduler. In the top illustration there are two EV-DO Rev. B devices connected to the network. Based on the reported DRC values, one mobile device could theoretically support a full 2.4Mbps from each EV-DO carrier (f1 and f2) while the other mobile device, located in a more challenging RF environment, can only support 1.2Mbps (f2) and 153kbps (f1) from the two carriers. With the current implementation of EV-DO Rev. B, the first mobile device gets served 1.2Mbps from each EV-DO carrier and the second device gets served 600kbps and 76.5kbps, respectively. The total combined throughput for both mobile devices is 3.1Mbps.

Figure 20. Benefit of Distributed Network Scheduler



Source: Various CDG member companies

With the Distributed Network Scheduler feature, the bandwidth is assigned to each carrier such that it maximizes the total throughput of the network. In this somewhat simplistic example, the first mobile device gets a full 2.4Mbps from a single EV-DO carrier (f1) and 0Mbps from the second carrier (f2), while the second mobile device gets 1.2Mbps from the second EV-DO carrier (f2) and 0Mbps from the first carrier (f1). As a result, the combined throughput increases from 3.1Mbps to 3.6Mbps.

The Distributed Network Scheduler feature requires a Multicarrier EV-DO or EV-DO Rev. B network and an additional software upgrade to the infrastructure. There are no device implications as long as the MRD devices support EV-DO Rev. B.

5.3.1.4 Single Carrier Multilink

The Single Carrier Multilink feature of DO Advanced is an interesting twist on Multicarrier EV-DO. With Single Carrier Multilink, an EV-DO Rev. B device (that can process two independent data streams simultaneously) can achieve the benefits of a multicarrier network, albeit in the hand-off regions of a single carrier environment.

Typically, with Multicarrier EV-DO, the separate EV-DO carriers that are serving a given mobile device are originating from the same cell/sector or different cell/sectors, although using different carrier frequencies. With Single Carrier Multilink, the same frequency can be used for both carriers that are serving the mobile device as long as the carriers originate from different cells/sectors. Figure 21 illustrates the concept of Single Carrier Multilink.

Figure 21. Single Carrier Multilink



Single Carrier Multilink extends the benefits of multicarrier EV-DO to single carrier networks.

The Single Carrier Multilink feature provides better network load balancing and increases cell-edge data rates, especially for multicarrier devices. The end result is greater data network capacity, reduced backhaul bottlenecks and improved data rates for all existing devices.

The Single Carrier Multilink feature requires EV-DO Rev. B devices that support multicarrier connectivity. Additionally, a software upgrade is required in the network, although every

base station doesn't have to support EV-DO Rev. B. This feature is best suited for an operator that has deployed EV-DO Rev. B in a portion of its network, but in other areas of its network, such as in rural areas where there is less traffic, the operator has only deployed a single EV-DO carrier. The performance of multicarrier devices will be enhanced in these single-carrier hand-off areas.

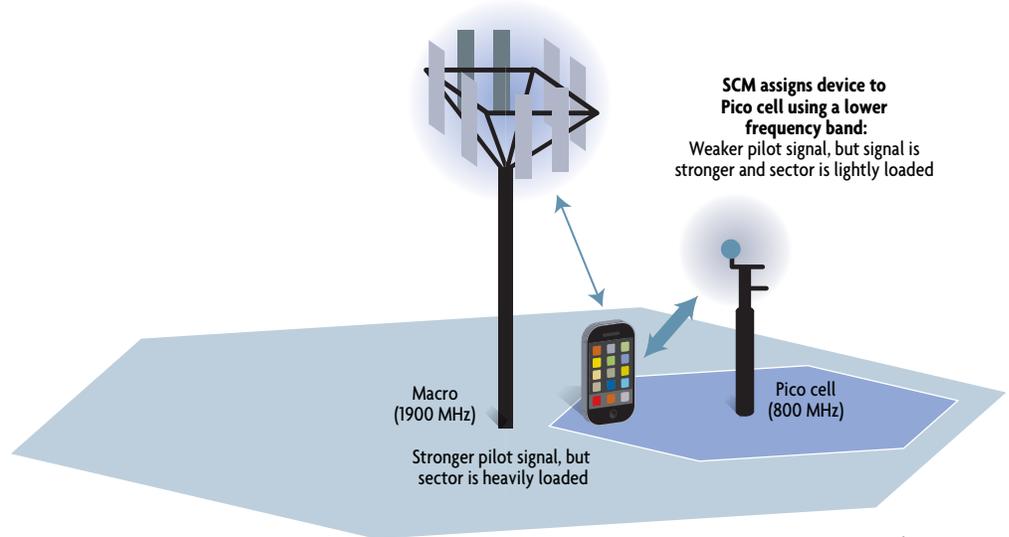
5.3.1.5 Smart Carrier Management

Normally, an EV-DO device accessing the network is assigned one or more carriers depending on its capabilities and the network load. For an EV-DO device that only supports one carrier (e.g., Rev. A), the network assigns a single carrier to the device based on the carrier it used to send the access probes.

Smart Carrier Management uses the signal strength and load on each carrier to assign the optimal combination of carriers for each device.

With Smart Carrier Management (SCM), the network, as shown in Figure 22, can assign a different carrier based on signal quality and the network load. For example, if the device tries to access the network from Carrier A, but the network determines that the device would receive a stronger signal from Carrier B, the network would assign Carrier B. This scenario could occur in a hotspot deployment where the carrier used by the hotspot has a stronger signal due to less interference from neighbor sectors. Another example is dual-band deployments where the carrier in the lower band (e.g., 850MHz) has better propagation characteristics and results in a stronger signal than the higher band carrier (e.g., 1900MHz).

Figure 22. Smart Carrier Management



For a multicarrier device (e.g., EV-DO Rev. B), the network can also decide on the number and specific carriers to assign to the device. If a network assigns the maximum number of carriers that the device can support, its trunking efficiency increases and the length of the data transfer decreases.

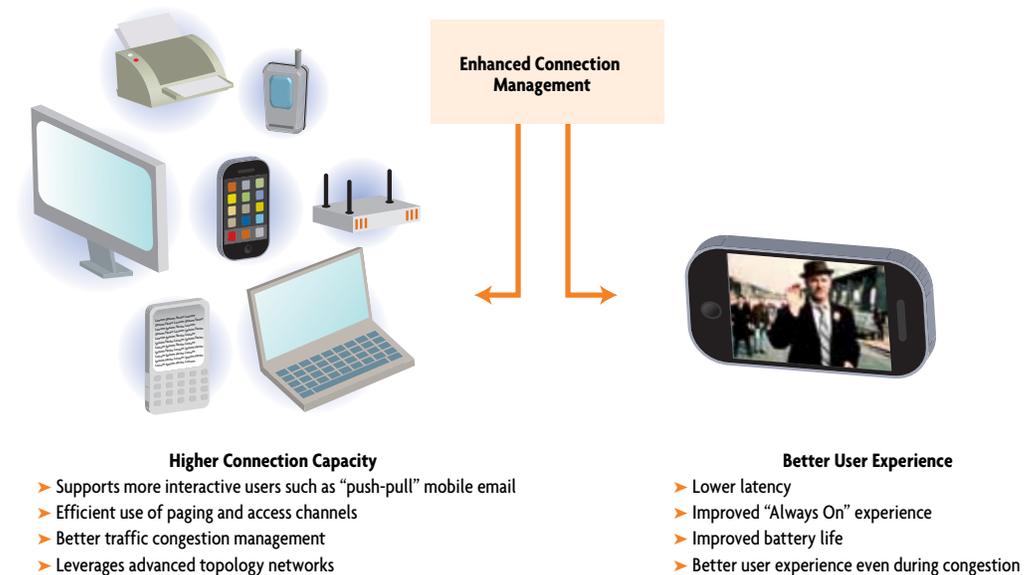
There is, however, a tradeoff between optimizing the user experience and managing the load on the carriers. The network must take into consideration the additional overhead due to the device

transmitting reverse link pilot and overhead channels and the use of channel elements and MAC indices on each carrier that the device is using. SCM balances this tradeoff by monitoring the load on each carrier. At low to moderate load, the network assigns the highest number of carriers to maximize the user experience. As the load increases, the network assigns fewer carriers to each device while shortening the dormancy timers to quickly remove inactive devices from the network.

5.3.2 Enhanced Connection Management

The Enhanced Connection Management feature of DO Advanced is designed to increase the number of connections, manage bursty applications better and improve the user experience using parameter optimization, implementation enhancements and Advanced Topology Networks.

Figure 23. Enhanced Connection Management



Source: Various CDG member companies

Enhanced Connection Management addresses the need for additional connection capacity, signaling capacity, lower latency and longer battery life by connecting devices to the network in an optimal way that is based on the type of application being used.

5.3.2.1 Parameter Optimization and Implementation Enhancements

Unlike some of the other DO Advanced features, which are designed largely to improve overall network capacity or data throughput, the parameter optimization and implementation enhancements that are part of the Enhanced Connection Management feature deal with how these devices establish and tear down their network connections based on the type of application being used. This is particularly useful for services, such as “Push” email and Instant Messaging that generate huge volumes of short and bursty traffic, creating lots of signaling overhead and congestion. For such applications, the Enhanced Connection Management feature optimizes the allocation and de-allocation of resources and improves the usage of access and paging channels to support a larger number of simultaneously connected devices. The resulting benefits include improved efficiencies from a signaling perspective, longer battery life, and a better user experience since the latency associated with connecting a device to the network is reduced.

To a large extent, many of the features discussed in Section 4.5 and Section 4.6 fall under the category of Enhanced Connection Management features. For example, the Data over Signaling feature can be used for chatty smartphone applications that generate lots of signaling traffic, but not very much in the way of data traffic. To the extent this feature is used, the operator benefits from improved network efficiencies, while the consumer benefits from an increased battery life. Yet, there are a few other new enhancements that will be implemented with DO Advanced.

5.3.2.1.1 Asynchronous Control Channel Bundling

Asynchronous Control Channel Bundling makes more efficient use of the overhead channels, thus freeing up capacity for more data traffic or allowing more users to access the network. For example, when there are a large number of mobile devices trying to access the network simultaneously, overhead messages, such as the Access Channel Acknowledgement message (ACAck) or the Traffic Channel Assignment message (TCA) that are sent from the base station to the mobile device, can be combined using Asynchronous Control Channel Bundling such that the ACAck or TCA messages for multiple devices are sent in the same packet.

5.3.2.1.2 Multi-AT Page

Conversely, operators can reduce loading on Synchronous Control Channels by employing a feature called Multi-AT Page which can reduce the overhead by up to three times, removing redundant information by combining paging messages to multiple devices into a single packet.

5.3.3 Advanced Topology Networks

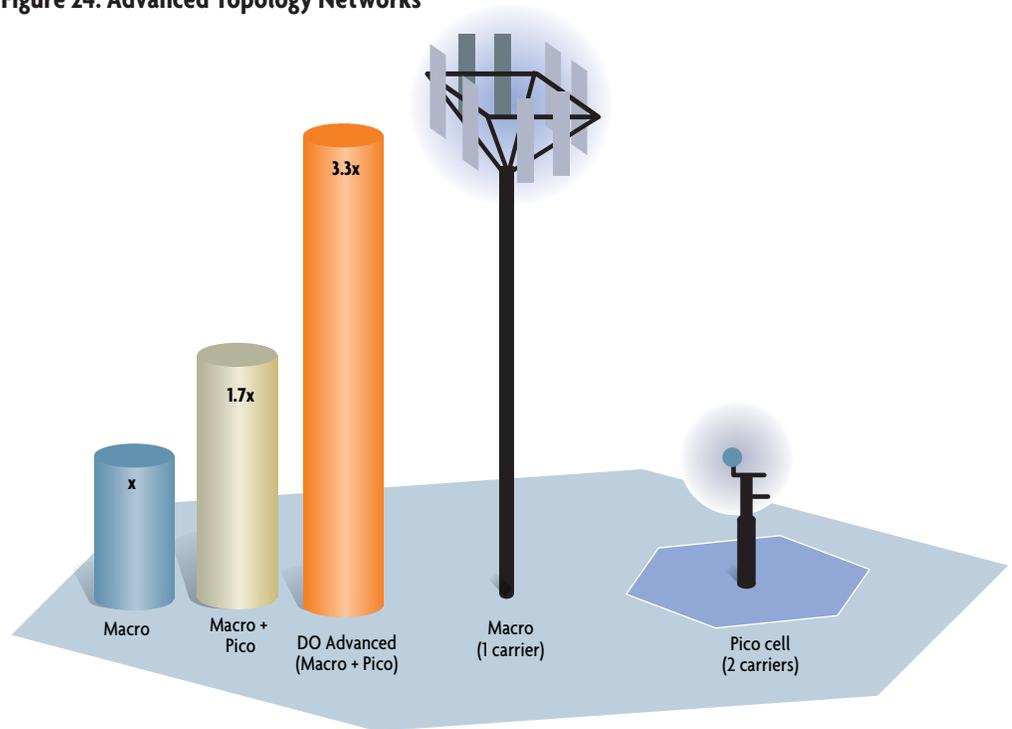
Advanced Topology Networks, as shown in Figure 24, are heterogeneous networks that include macro-, pico-, micro-cells, repeaters, remote radio heads and femtocells, all of which move the transmitters closer to the user to increase bandwidth. The same Smart Network techniques enabled by DO Advanced can be applied to all network deployment scenarios, including Advanced Topology Networks, to significantly improve overall network performance. Implementing DO Advanced across Advanced Topology Networks offers the following benefits:

- Unified operations across macro-, pico-, micro-, repeaters and remote radio heads
- Increased bandwidth in high-traffic areas
- Higher network capacity
- Improved coverage
- Lower latency
- Less signaling traffic
- Improved user experience

Implementing smart network techniques across Advanced Topology Networks adds a whole new dimension to wireless networks. The combination intelligently addresses the demand for network signaling and data traffic by adding the appropriate amount of capacity exactly where it is needed. For example, the capacity can be added inside of buildings and in high-traffic venues such as stadiums, airports and convention centers. *Network capacity* is defined as the total capacity of an area or a cluster of cell sites. This definition is different from *spectral efficiency*, the conventional term used to measure capacity. The difference is that spectral efficiency considers a fully loaded network, whereas *network capacity* considers uneven loading, which is more typical of actual network conditions.

With the introduction of self-organizing networks, these enhancements are further optimized.

Figure 24. Advanced Topology Networks



Example: Improvement with DO Advanced Pico cell deployment

Note: Simulations assume 1 single carrier macro, with 2 double carrier picocells. Pico-cells are randomly placed in the network. A data loading ratio of 2:1 is assumed between high-demand and low-demand areas.

Source: Various CDG member companies

5.3.4 Device Enhancements

As is the case with any intricate system, the quality and performance of a cellular network is only as good as its weakest link. All things being equal, a network will be more efficient and support more traffic if the mobile devices perform at their optimum level. In other words, if a mobile device can maintain a given data rate at a lower signal level or support higher throughput for a given signal quality, the overall network benefits. Obviously, with higher data rates the user experience is also improved. Likewise, if a mobile device can transmit data faster, or maintain a given reverse link data rate using less transmit power, there are additional benefits to be gained, including, surprisingly enough, the ability to achieve higher data rates in the forward link.

In the context of DO Advanced there are two device enhancements which are used to achieve these objectives: Enhanced Equalizer and Mobile Transmit Diversity. While these enhancements require new chipsets in the devices – something that an operator can realize through a natural device migration and replacement strategy – there do not have to be any changes introduced to existing network infrastructure.

The Enhanced Equalizer feature that is included in all new CDMA2000 device chipsets is specifically designed to cope with the chatty nature of mobile data traffic. A mobile device determines the condition of the radio link, and in turn the maximum data rate that it can

The Enhanced Equalizer takes advantage of uneven and bursty traffic channel conditions caused by the chattiness of smartphone applications to increase cell-edge data rates and forward link sector capacity.

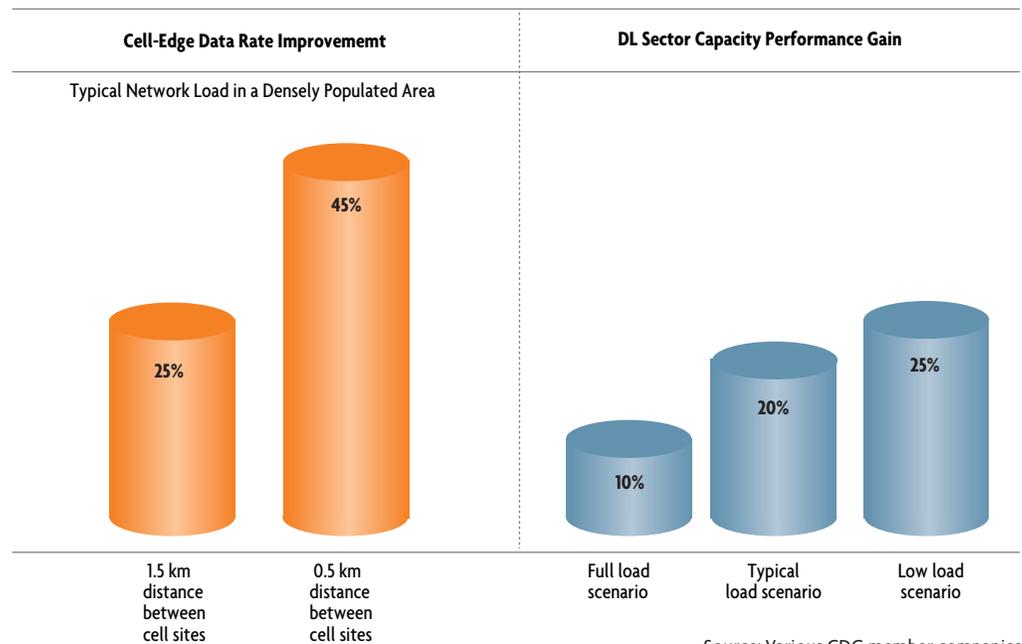
support, based on a variety of factors including the level of interference that it detects. This interference can change every slot (e.g., 1.66msec). The Enhanced Equalizer is better equipped than the legacy equalizers to handle this very dynamic change in interference.

For both pilot and traffic channel interference, the impact is most noticeable when the mobile device is in handoff mode (e.g., the active set of cell sites is greater than one). In these situations, which occurs almost two-thirds of the time in a typical network, a mobile device with a legacy equalizer may not be able to adjust quickly enough to deal with the rapidly changing interference levels. Therefore, the device could request a lower data rate in comparison to the higher data rate that it could actually support when the uneven and bursty traffic channel interference subsides.

With EV-DO, the pilot signal is always transmitted at full power. Therefore, interference between pilots is relatively high during pilot bursts. During traffic transmissions, the interference level depends on whether or not the neighboring sectors are transmitting. Legacy equalizers take this traffic interference into account, but they are not able to adapt to the rapidly changing interference from neighboring sectors due to bursty traffic.

The Enhanced Equalizer, which is part of the DO Advanced feature set, detects and analyzes the interference from both the pilot and traffic channels, thus allowing it to take advantage of those instances when data traffic on adjacent cells subsides. Given the chattiness of smartphone applications, the receiver with the Enhanced Equalizer will frequently achieve a higher data rate than a legacy receiver that only analyzes the interference from the pilot channel while ignoring the widely fluctuating traffic channel interference.

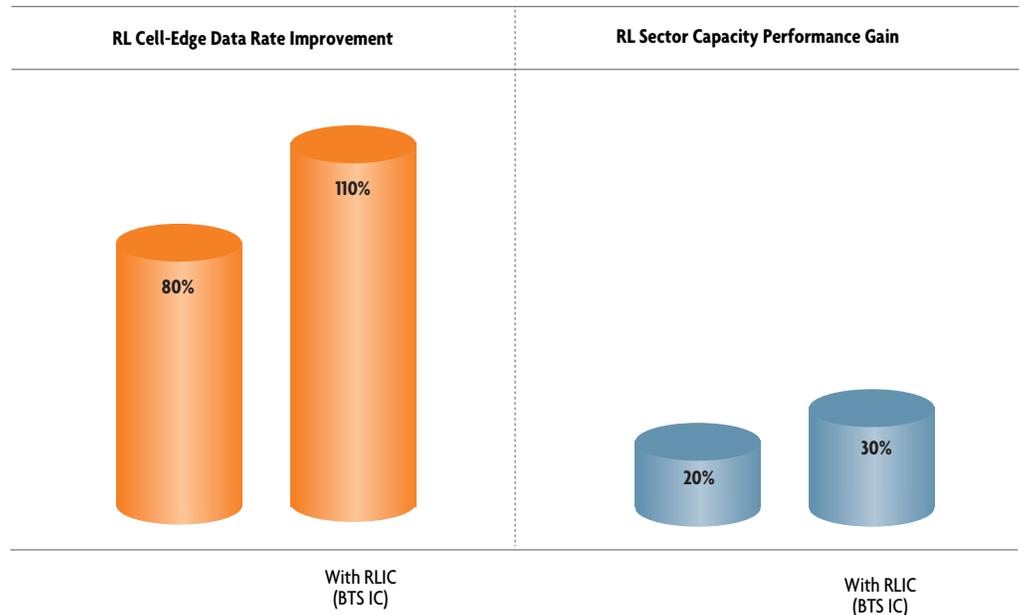
Figure 25. Enhanced Equalizer Performance Enhancements



Source: Various CDG member companies

The impact on network capacity due to the Enhanced Equalizer is a function of the penetration rate of the devices that support the feature. The increase in cell-edge rates can be as high as 45 percent, under typical network load, in dense urban areas where site-to-site distances

Figure 26. Mobile Transmit Diversity Performance Enhancements



Source: Various CDG member companies

are as short as 0.5km, and up to 25 percent for distances of 1.5km.⁸ The forward link sector capacity improvement largely depends on the relative load of the neighboring sectors. As shown in Figure 25, the increase can be as high as 25 percent for low-load scenarios, up to 20 percent under typical loads, and up to 10 percent for full load situations.⁹ Further, to the extent that these new devices make more efficient use of network resources, it frees up these resources for legacy devices, thus benefiting to varying degrees the entire population of EV-DO devices.

With mobile transmit diversity the dual antenna devices can transmit data at a faster rate, especially at the edge of the cell, while forward link data rates also increase.

Mobile Transmit Diversity (MTD) is the reverse link companion to Mobile Receive Diversity (MRD), a multi-antenna feature that was discussed in Section 5.1.3. With MTD the mobile device uses an additional transmit chain and antenna to send data. This feature improves performance across the network, including at the cell-edge where the end-user data rates can increase by as much as 80% in the reverse link with no changes to the network infrastructure, or even 110% with the separate introduction of interference cancellation at the base station (RLIC).¹⁰ With respect to reverse link sector capacity, the performance gain can be as much as 20%, or 30% with interference cancellation at the base station (RLIC).¹¹

While a bit counter-intuitive, a higher reverse link data rate also benefits the forward link performance of mobile devices, specifically, the forward link data rate with certain protocols, such as HTTP that require sending confirmations of received data via the reverse link (e.g., HTTP GET messages initiate data downloads and the received data triggers additional HTTP GETs, resulting in a significant amount of traffic.). To the extent that these HTTP GETs can be sent more quickly with mobile transmit diversity, it can reduce the total webpage download time and improve the user experience.

⁸ Ibid.
⁹ Ibid.
¹⁰ Ibid.
¹¹ Ibid.

Another important observation is that any improvement in the reverse link data rates significantly improves latency of applications like Web browsing on the forward link, especially for users at the cell edge. Hence, with MTD, the user experience of bursty applications is greatly enhanced.

While MTD does require a new chipset, the impact on new components in the handset is somewhat mitigated by the presence of the second antenna that is already being used for receive diversity.

5.3.5 Market Opportunities and Implications

It is difficult to provide a single cohesive view on the market opportunity for DO Advanced since there are so many different technology options to consider. Some of the features, such as those that fall under the Enhanced Connection Management and Advanced Device categories, are likely to achieve widespread adoption since they are an inherent part of the chipset roadmap (e.g., Mobile Transmit Diversity and Enhanced Equalizer) or, in the case of Enhanced Connection Management, many of the underlying features are already being offered by the infrastructure vendors.

It is also likely that there will be a direct correlation between some of these features and the expected data usage of the mobile device. For example, while Mobile Transmit Diversity benefits all mobile data devices, it is more likely going to be included with a mobile device, such as a USB dongle or high-end smartphone that generates a lot of mobile data traffic each month.

It is foreseeable that even North American operators could consider certain DO Advanced features, even if they are not currently considering EV-DO Rev. B.

With respect to the Smart Network features of DO Advanced, Network Load Balancing, Smart Carrier Management and Adaptive Frequency Reuse seem to be the most promising since they do not presuppose the existence of multiple EV-DO carriers and EV-DO Rev. B devices. Therefore, it is foreseeable that even North American operators are considering these features, even if they are not currently considering upgrading their 3G network to EV-DO Rev. B.

Since the Distributed Network Scheduler and (to a lesser extent) the Single Carrier Multilink features require that the network supports Multicarrier EV-DO, their addressable market opportunity is somewhat limited. That being said, these two features appear to be quite compelling since they allow the operator to advertise and deliver the higher data rates of Multicarrier EV-DO while also allowing the operator to fully maximize network resources due to the unevenness of network traffic. Additionally, with Single Carrier Multilink it offers “EV-DO Rev. B-like” services in areas of the network where a second carrier does not exist.

Most importantly, the basket of DO Advanced features give CDMA2000 operators flexibility in how they plan their network evolution and how they cope with the tremendous growth in mobile data and signaling traffic. Instead of taking the traditional route of adding more capacity – deploying more carriers and/or deploying more cell sites – CDMA2000 operators have the opportunity to use existing network resources and devices in a more efficient manner. And to the extent these operators add more hardware (new carriers, new cell sites) to their network for capacity purposes, the gains would be realized throughout the network instead of being isolated to a single cell site.

6.0 Conclusions

This whitepaper provides a comprehensive look into how mobile data usage and signaling traffic impacts a mobile network and what CDMA2000 operators can do to handle this traffic.

With regards to the unprecedented growth in traffic, we know that today's smartphones and tablets have enhanced features which make them compelling to use. Meanwhile, higher performing network technologies are enabling us to do things that were previously unimaginable; a vast number of applications are becoming more popular; and data pricing plans tend to favor high usage, even though operators are beginning to introduce usage caps. Yet, while the rapid growth in mobile data traffic is well recognized, it is also important to recognize where the growth exists, when it overloads the network and how to best address the overall and localized demand for voice, data and signaling traffic.

It is critical to understand and appreciate that the mobile data traffic in a network is never uniformly distributed. This phenomenon means that at any given location and point in time a network may have reached its maximum capacity in some cells while surrounding cells may have unused and available capacity. If an operator is able to take advantage of this unevenness, then it can continue to deliver a compelling user experience without having to purchase additional spectrum and deploy additional network resources that would involve cost, disruption and complexity.

The CDMA2000 technology roadmap provides operators with great flexibility and a large number of technology features that can be leveraged to support existing voice and mobile data traffic, as well as future capacity requirements. In addition to optimizing their networks to better handle the characteristics of today's mobile data traffic, operators can also consider more comprehensive solutions, such as 1X Advanced and DO Advanced, to further improve the user experience and cost-effectively leverage existing network resources and devices. These incremental, affordable and backward compatible enhancements allow operators to fully capitalize their existing investments and enable them to judiciously add network capacity where and when it is needed.

Beyond potentially quadrupling voice capacity, 1X Advanced can be used to either increase coverage by as much as 70%, increase 1X data capacity by up to three fold, improve femtocell performance or free up spectrum for Multicarrier EV-DO or 4G data services. For many operators, this latter ability is paramount to their mobile data strategy since they may not have unused spectrum at their disposal.

In a similar fashion, DO Advanced maximizes existing network resources by introducing Smart Network techniques that take advantage of unevenly loaded networks and bursty applications. Most of its features can be implemented with a simple software upgrade, without having to purchase new infrastructure or devices. This flexibility, and the realization that existing 3G networks will carry the preponderance of network traffic for years to come, provides CDMA2000 operators with the resources that they need to make strategic decisions that are in their best interest.

For some operators, this may mean that they only consider changing a few basic network parameter settings and foregoing many of the enhanced capabilities that the CDMA2000 roadmap offers. For other operators, that have limited access to capital and spectrum, they may need to take a more aggressive approach and implement all of the beneficial changes to their network infrastructure and device strategy. Regardless, the CDMA2000 technology roadmap offers a variety of solutions to help operators address the ever-growing demand for data well into the future.

