Broadband data performance of third-generation mobile systems

Johan Sk ld, Magnus Lundevall, Stefan Parkvall and Magnus Sundelin

The rapid, widespread deployment of WCDMA and an increasing uptake of third-generation mobile systems (3G) services are bringing network performance into sharp focus. Besides efficiently supporting an increasing number of subscribers, network systems should also give end-users a high-speed experience. To solve this equation, with its seemingly conflicting components, we need to understand performance and how it is measured. Likewise, present-day and evolving 3G systems should include features for increasing system performance.

New high-speed services and greater end-user demand for performance are driving the evolution. WCDMA Evolved supports an enhanced broadband experience of WCMDA systems. WCDMA Release 99 (Rel-99) services have evolved into WCDMA Releases 5 and 6 (Rel-5, Rel-6), which will reach commercial deployment by year-end 2005. Systems based on CDMA2000 are going through a similar evolution.

The authors describe the path to WCDMA Evolved and how it affects performance for end-users and operators.

Performance of 3G services

There is no single universal measure of performance for a telecommunications system. Indeed, end-users (subscribers) and system operators define good performance quite differently. On the one hand, end-users want to experience the highest possible level of quality. On the other hand, operators want to derive maximum revenue, for example, by squeezing as many users as possible into the system.

Until now, performance-enhancing features could generally either improve perceived quality of service (QoS) or system performance. But now, with WCDMA Evolved (Ericsson's evolution of WCDMA for best-effort data services), one can potentially do both.

Mobile best-effort data services, such as web surfing and file downloads, have been available via packet data since the first release of WCDMA networks. They are a significant enhancement compared to 2G networks, and because the WCDMA specifications are evolving, packet-data support continues to improve. The first two phases of this evolution, commonly referred to as WCDMA Evolved, entail

- the introduction (in Rel-5) of high-speed downlink packet access (HSDPA); and
- the introduction (in Rel-6) of an enhanced uplink.

Compared to earlier releases of WCDMA, these changes yield better data rates and shorter delay; that is, they greatly improve the service experience and system capacity.

End-user perspective

Users of circuit-switched services are assured of a fixed bit rate. The quality of service in the context of voice or video telephony services is defined by perceived voice or video quality. Superior quality services have fewer bit errors in the received signal. By contrast, users who download a web page or movie clip via packet data describe quality of service in terms of the delay they experience from the time they start the download until the web page or movie clip is displayed.

Best-effort service does not guarantee a fixed bit rate. Instead, users are allocated whatever bit rate is available under present conditions. This is a general property of

BOX A, TERMS AND ABBREVIATIONS

1xEV-DO	CDMA2000 (single-carrier) evolution with data-only carrier	QoS	Quality of service
3GPP	Third-generation Partnership Project	Rel-5	Release 5 of 3GPP specifications
16QAM	16-level quadrature amplitude modulation	Rel-6	Release 6 of 3GPP specifications
ARQ	Automatic repeat request	Rel-99	Release 99 of 3GPP specifications (first release)
BWA	Broadband wireless access	TCP	Transmission control protocol
FDD	Frequency-division duplex	TDD	Time-division duplex
FTP	File transfer protocol	TTI	Transmission time interval
HSDPA	High-speed downlink packet access	UDP	User datagram protocol
IEEE	Institute of Electrical and Electronics Engineers	UMTS	Universal mobile telecommunications system
MIMO	Multiple input/multiple output antenna system	WCDMA	Wideband code-division multiple access
MWA	Mobile wireless access	WiFi	Wireless fidelity
OFDM	Orthogonal frequency-division multiplexing	WiMAX	Worldwide interoperability for microwave access

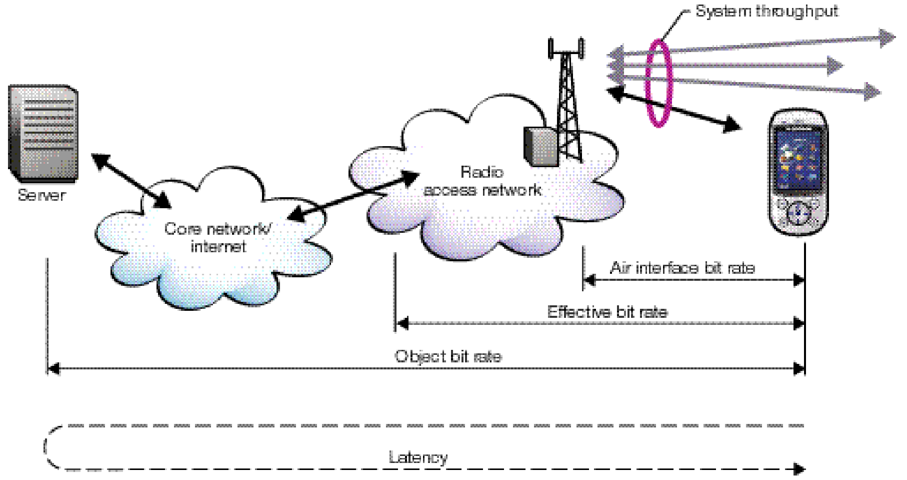

Figure 1
Definitions of bit rates, end-to-end (see also Box B).

packet-switched networks; that is, network resources are not reserved for each user. Given that delay increases with the size of the object to be downloaded, absolute delay is not a fair measure of quality of service.

A lone user in a radio network with good radio conditions may enjoy the peak bit rate of the air interface. But if radio conditions are less than optimum or there is interference from other users, the air interface bit rate will be less than the peak bit rate. In addition, some data packets might be lost, in which case the missing data must be retransmitted further reducing the effective bit rate as seen from higher protocol layers (such as IP). What is more, the effective bit rate diminishes even further as the distance from the cell increases (due to poorer radio conditions at cell edges).

The peak air interface rate and radio conditions are not the only factors that limit performance. Taking the radio network and core network as a whole all the way to the application server, one also encounters delays in various network nodes and protocols. This results in an object bit rate that is lower than the effective bit rate. The object bit rate, which is measured at the application level, takes into account all delays and is averaged over the objects transmitted to or from an end-user. It is the size of the object divided by total delay measured in kilobits per second (kbps).

The transmission control protocol (TCP)—the protocol at the transport layer—is commonly used together with IP traffic. But due to its slow-start algorithm, which is sensitive to latency in the network, it is especially prone to cause delay for small packets. The slow-start algorithm is meant to ensure that the packet transmission rate from the source does not exceed the capability of network nodes and interfaces.

Network latency, which in principle is a measure of the time it takes for a packet to travel from a client to server and back again, has a direct impact on performance with TCP. Therefore, an important design objective in WCDMA Evolved has been to reduce network latency. One other quality-

BOX B, DEFINITIONS OF BIT RATES

Air interface bit rate
Bit rate of the physical layer achieved under certain radio conditions with specific coding and modulation.
Peak bit rate
Peak bit rate of the air interface under ideal radio conditions.
Effective bit rate
Bit rate as seen from higher layers (IP). This rate is dependent on the bit rate of the air interface as well as protocol overhead, retransmissions and queing delays.
Object bit rate
Bit rate defined at the application level, end-to-end. It includes delays outside the radio network and delays from TCP flow control.
Latency
End-to-end round-trip time of a small packet.
System throughput
Total number of bits per second transmitted over the air interface (per sector).

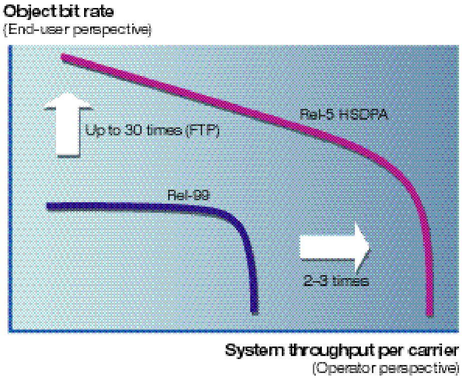

Figure 2
In the context of best-effort packet data, network load (number of users) steers the bit rate and system throughput on the downlink. WCDMA Evolved has the potential to improve bit rate and system throughput.

related criterion affects the setup time for initiating, for example, a web-browsing session.

Operator perspective

Radio resources need to be shared when multiple users are in the network. As a result, all data must be queued before it can be transmitted, which restricts the effective bit rate to each user. Notwithstanding, by scheduling radio resources, operators can improve system throughput or the total number of bits per second (bps) transmitted over the air interface. HSDPA and the enhanced uplink employ intelligent scheduling methods to optimize performance.

One important performance measure for operators is the number of active users who can be connected simultaneously. Given that system resources are limited, there is a trade-off, in terms of object bit rate, between number of active users and perceived quality of service.

WCDMA Evolved—the next step

To exploit or take full advantage of the bursty characteristics of packet data and rapid variations in radio conditions, WCDMA Evolved applies fast and dynamic resource allocation in both the uplink and downlink. More specifically, it employs hybrid automatic repeat request (ARQ) with soft-combining, scheduling, and for the downlink, fast link adaptation with higher-order modulation (Box C). Corresponding functionality is contained in the base station to allow for fast adaptation and low delays. In addition, the transmission time interval (TTI) has been reduced to 2ms to accommodate faster adaptation and reduce end-user delay.

Although the principles applied in the uplink and downlink are similar, certain fundamental differences have affected design choices. Most notably, for the downlink, the shared resource for power and codes is located in the base station. For the uplink, the power resource is distributed among the terminals. Soft handover solely applies to uplink transmissions.

Performance achievements

Performance analysis (by means of computer simulations) plays an important role in

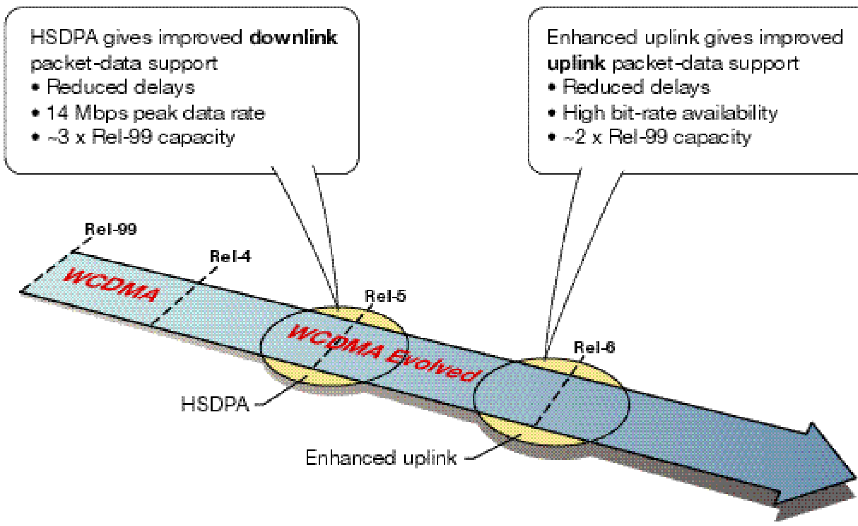

Figure 3 (see also Box C)
The evolution path of WCDMA Evolved.

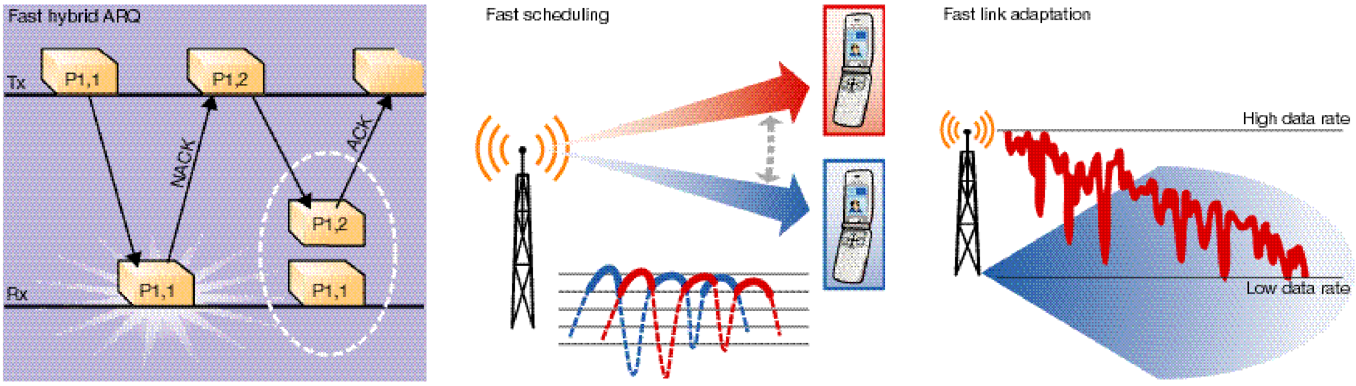

Figure 4 (see also Box C)
Basic principles of WCDMA Evolved.

the development of WCDMA Evolved. This analysis often relies on several assumptions, which although simplified, give a good indication of network performance, especially of the relative improvement for HSDPA and the enhanced uplink compared to WCDMA Rel-99. Field experience is also invaluable for obtaining the full picture of achievable performance.

End-user performance analysis

Below, using Rel-99 as a reference, we will demonstrate gains in performance from Rel-5 (HSDPA) and Rel-6 (HSDPA and enhanced uplink). The results were derived under the assumption that radio conditions do not limit the air interface bit rate. Further, it was assumed that the Rel-99 system provided radio bearer bit rates of 64kbps on the uplink and 384kbps on the downlink (denoted 64/384). The corresponding figures for Rel-5 and Rel-6 are 384/4,320kbps and 4,320/13,440kbps, respectively. The bit rates of Rel-5 and Rel-6 are considerably higher than those of Rel-99, but as we shall see they are not available over a larger part of the cell as is often the case for Rel-99.

Performance when transferring large files using TCP is determined by the bit rate of the bearer. For small files, latency is important. To highlight these aspects, the results illustrate TCP-based uploads and downloads of

BOX C, BASIC PRINCIPLES OF WCDMA EVOLVED

Fast hybrid automatic repeat request (ARQ) with soft-combining enables receivers to rapidly request the retransmission of erroneously received data blocks. In the downlink, the receiver is a terminal. In the uplink, the receiver is the base station. Before decoding a signal, the receiver combines information from the original transmission with that of subsequent transmissions. This procedure is called soft-combining. Fast hybrid ARQ with soft-combining is used in the uplink and downlink. Compared to earlier releases of the WCDMA specifications, it has the potential to substantially reduce delay and significantly increase capacity.

Fast scheduling is used in the uplink and downlink. The scheduling strategies for each may differ, however. Downlink resources (code and power), for example, are typically shared in a way that addresses a user with advantageous instantaneous channel conditions per time interval. Channel-dependent scheduling, as this strategy is called, exploits short-term variations in downlink radio conditions to increase capacity. In the uplink, the transmission power of a mobile terminal is substantially less than that of the base station. Therefore, a single user's transmission cannot use full system capacity. Multiple users are thus frequently scheduled in parallel. To control the overall level of interference in the cell, the scheduler controls when and at what rate each terminal should transmit.

Fast link adaptation applies to the downlink. In essence, downlink transmission power is held constant while the data rate is rapidly adjusted to adapt to varying radio conditions. This method is efficient for services that tolerate short-term variations in the data rate. Channel conditions permitting, spectral-efficient 16-level quadrature amplitude modulation (16QAM) can be used to further increase capacity and data rates.

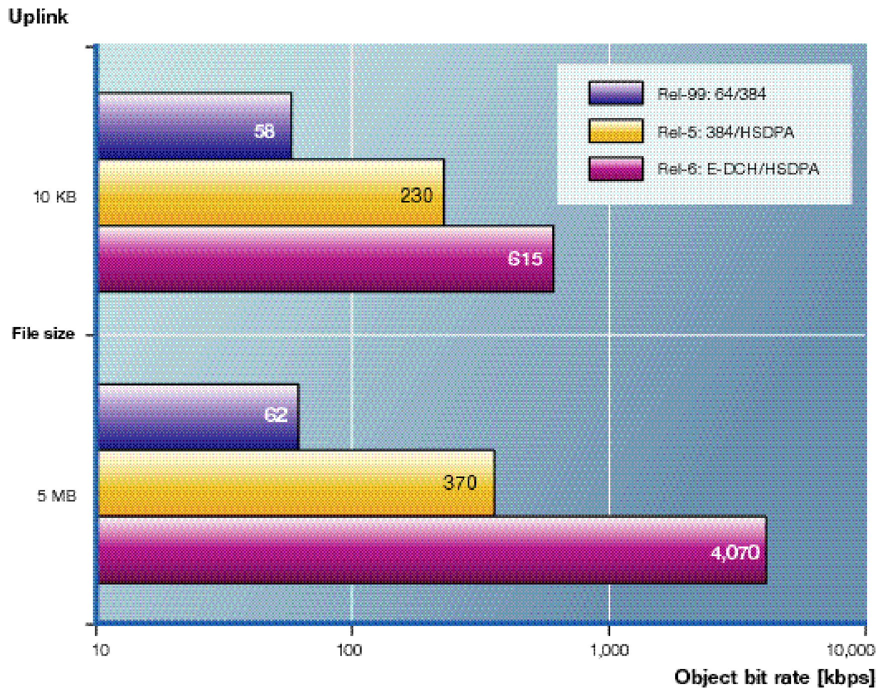

Figure 5
Evolution of WCDMA end-user bit rates for file upload (Note: A logarithmic scale has been used for the bit rates).

Figure 6
Evolution of WCDMA end-user bit rates for data downloads.

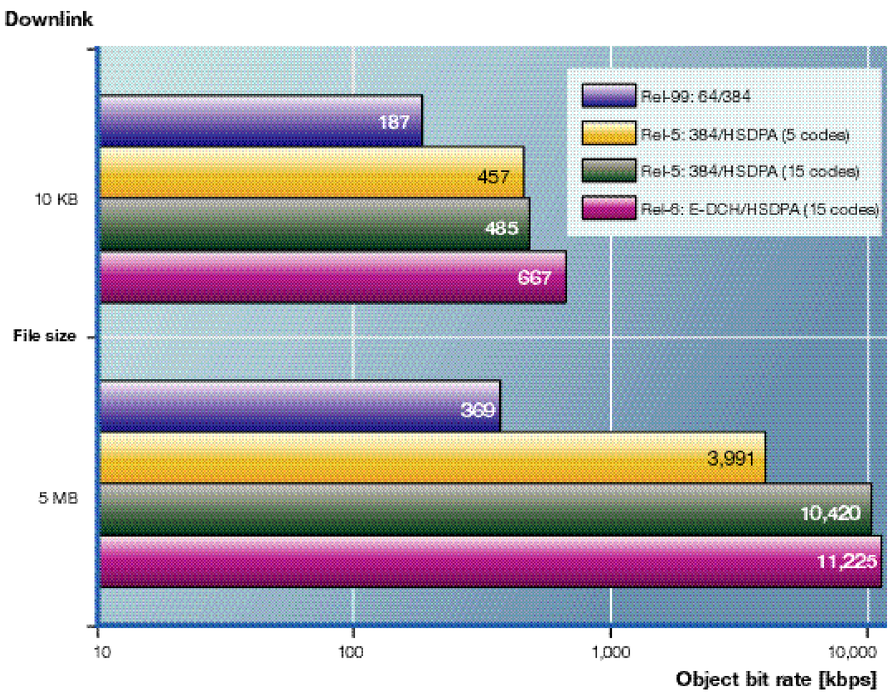

- a small, 10KB file (for instance, an e-mail message without attachment); and
- a large, 5MB file (for example, an MP3 file).

Finally, it was assumed that there is no loss of IP packets on the fixed network path between client and server. Packet loss would affect the results, but this impact has not been included.

Figure 5 shows upload performance. The gain from Rel-5 is due to the 384kbps uplink service. For large file transfers, the object bit rate approached the radio bearer bit rate, and the enhanced uplink in Rel-6 gave a significant improvement compared to earlier releases. For small file transfers, latency was a determining factor—one that made it impossible to reach the radio bearer bit rate. Rel-6 considerably increased the object bit rate, primarily by reducing latency.

Figure 6 shows TCP download performance. For large file transfers, the introduction of HSDPA (5 codes) increased the object bit rate by an order of magnitude (10^1) compared to Rel-99. Configuring HSDPA with the maximum of 15 codes further increased the object bit rate to 10Mbps. For small file transfers, performance was limited by TCP and network latency. As expected, Rel-5 improved the object bit rate compared to Rel-99. Especially interesting was the performance of Rel-6 compared to Rel-5. The enhanced uplink reduced latency, which in turn, improved TCP download performance.

These results demonstrate the capability of the air interface. However, radio conditions and network load influence the achievable air interface bit rate. Figures 7 and 8 show bit rate availability. The examples depict a single user. Bit rate availability is expressed in terms of coverage percentage; that is, the percentage of the cell area where a certain bit rate can be achieved. The modeled Rel-6 network has been deployed to provide an uplink bit rate of at least 64kbps with 95% probability. This means that the network can provide 64kbps in 95% of the cell area. Due to limited output power in mobile terminals, the uplink generally provides lower bit rates than the downlink. Heavy traffic load in the network increases interference, which reduces coverage. For both the uplink and downlink, 4Mbps can be achieved in more than half of the cell area without load; with load, more than 2Mbps (Figure 7). These results show the achievable bit rate when a user is allowed to transmit (uplink) or receive (downlink). Multi-

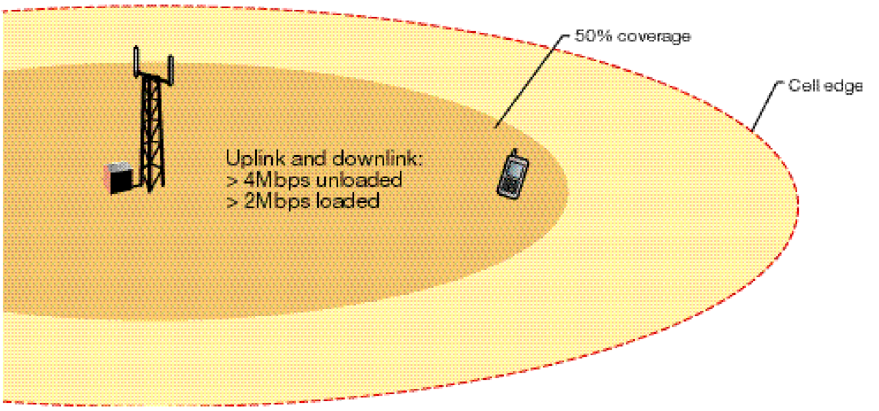

Figure 7
Bit rate availability for WCDMA Evolved.

ple users in the cell reduce the effective bit rate per user because the resources must be shared by means of scheduling.

System capacity analysis

Until new products become available and have been deployed in loaded networks, radio network simulations will be used to assess system capacity. Simulations are also used to better control the environment and conditions for performance analysis. These simulations include models of the cell layout, traffic behavior, radio propagation, and assumptions about the receiver performance of radio base stations and mobile terminals. Each of these parameters affects the results.

System capacity is defined as average system throughput at which perceived quality drops to an unacceptable level. Greater system throughput can be obtained by disregarding perceived quality and fairness among users. This measure of capacity is not dependent on traffic load generated per user, which varies from application to application.

Figure 8 shows the uplink and downlink capacity derived from simulations of a macro cellular network. Capacity intervals are given to illustrate that these figures are dependent on the models and assumptions

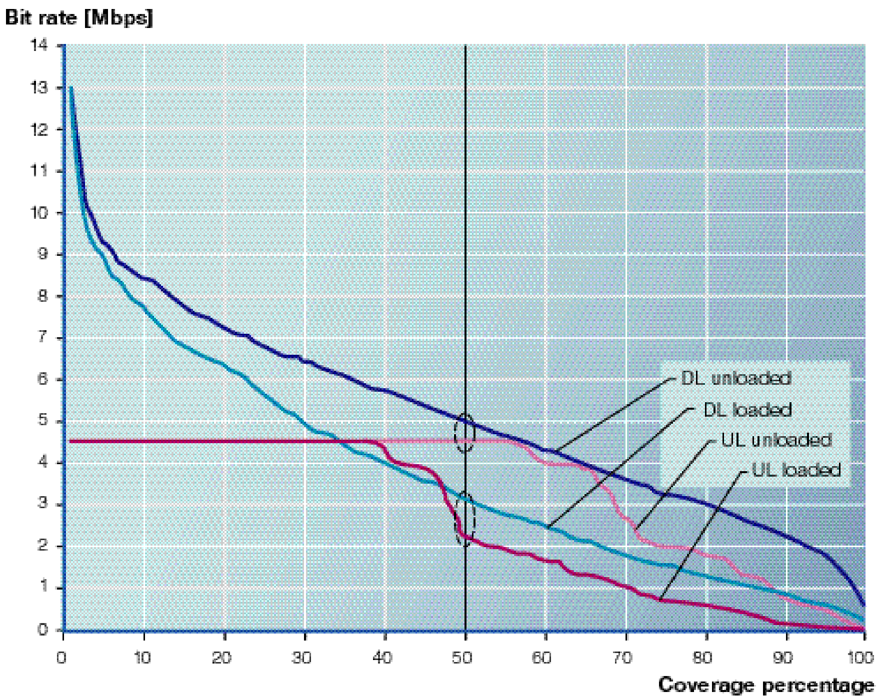

Figure 8
Heavy traffic load in the network increases interference, which reduces coverage.

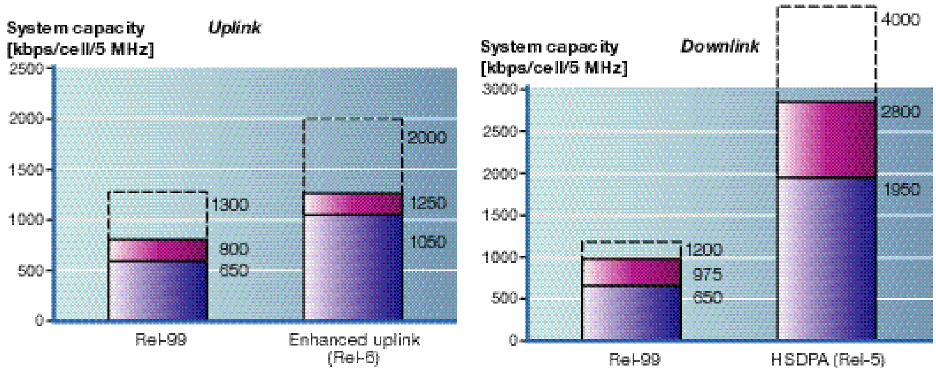

Figure 9
System capacity.

used. Compared to Rel-99, the enhanced uplink yields a 30-90% gain in capacity depending on whether hybrid ARQ has been optimized for latency (targeting few retransmissions) or capacity (targeting multiple retransmissions). In general, in terms of coverage or stability, maximum uplink capacity is determined by maximum tolerable interference.

Thanks to fast scheduling and link adaptation, HSDPA gives two to three times more capacity in the downlink than Rel-99. The capacity interval (marked in red, Figure 9) indicates that HSDPA capacity is dependent on the radio environment. Employing receiver diversity in combination with interference-suppression techniques in mobile terminals can further enhance capacity in the downlink.

The dashed bar (Figure 9) indicates capacity when more idealized assumptions have been used: the effect of TCP has been excluded, and scheduling and cell load have been optimized to permit as much data as possible to pass through each cell. This scenario, which is not a favorable choice of operating point for a 3G system, corresponds to operations at the lower right-hand corner of Figure 2.

System throughput, expressed in kbps per cell, is a key parameter for network dimensioning. Using assumptions about subscriber behavior (for example, data volume generated per month) one can translate system throughput into the number of subscribers per cell that the network can support. Appropriate margins should be applied to account for variations. Figure 10 describes the number of subscribers that can be supported when capacity per sector is 2800kbps.

Figure 10
Example of capacity of a mobile wireless access service. The capacity calculation (630 subscribers supported by the site) assumes that each subscriber uses 1GB per month, of which amount 0.6% is during the busy hour. This gives an average data rate of 13.3kbps per subscriber during the busy hour.

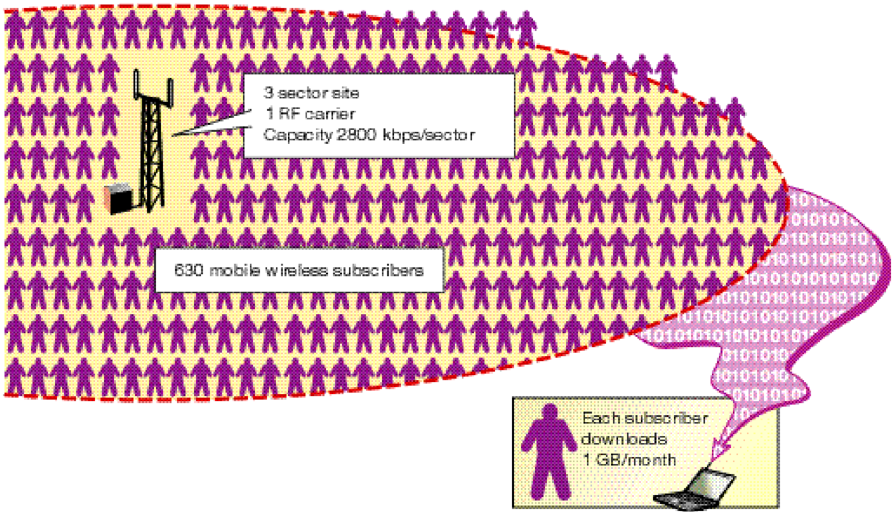

Field experience

Today, there is a substantial body of field experience from running large deployments of Rel-99 WCDMA equipment. WCDMA Evolved will have its first commercial launch later this year with HSDPA. CDMA2000 is moving along a parallel evolution path called 1xEV-DO. The first phase of CDMA2000 1xEV-DO has already been deployed commercially.

CDMA2000 1xEV-DO currently supports peak air interface bit rates of 2.5Mbps. In a subsequent version, Rev A, it will support 3.0Mbps. Field trials show that the high-speed version of CDMA2000 significantly increases object bit rates and system throughput. These improvements have the

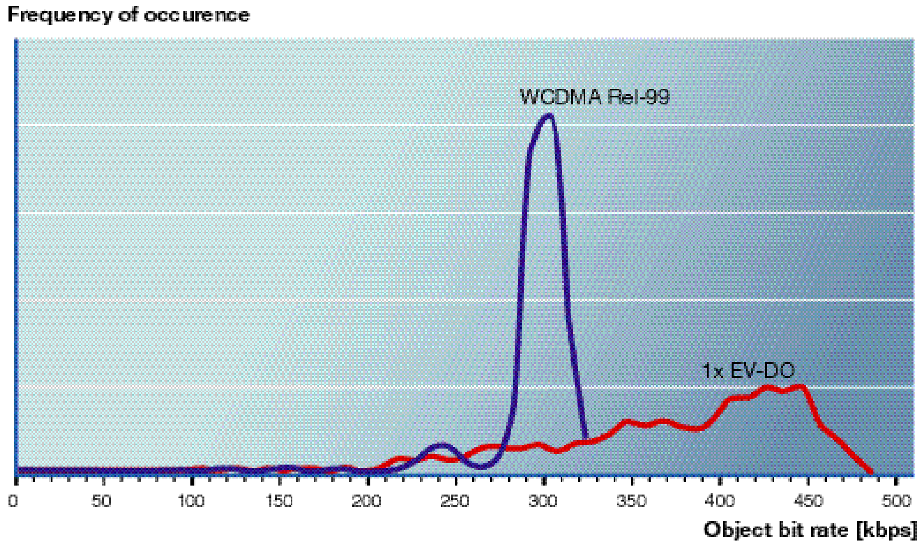

Figure 11
File download performance in a field test.

same magnitude and are based on the same principles as those employed by WCDMA Evolved.

Figure 11 shows a test of a 500KB FTP download over

- an EV-DO bearer; and
- a WCDMA Rel-99 bearer.

CDMA2000 has a radio bandwidth of 1.25MHz, whereas WCDMA has a radio bandwidth of 5MHz. Notwithstanding, we see that the EV-DO enhancement considerably improves downlink bit rates. Fast link adaptation adapts quickly to channel conditions, enabling greater object data rates. It also gives a larger spread of performance values than WCDMA Rel-99. This is because channel conditions vary over the test area. The EV-DO bit rate will thus vary with conditions. Thanks to its wider bandwidth, HSDPA will yield even greater downlink bit rates than 1xEV-DO.

End-user performance might be limited due to latency when TCP is used as the transport layer protocol for FTP download. EV-DO is a clear improvement but the latency (in the specific non-Ericsson 1xEV-DO deployment) limits performance. Ongoing work to improve latency will enhance the performance of EV-DO as well as for WCDMA.

Table 1 shows the potential (from test results) for further improvement. In addition to download performance using FTP/TCP, it shows download performance using the user datagram protocol (UDP). In this case, the gains from EV-DO are more obvious: UDP is not sensitive to latency, so the object bit rate for EV-DO is almost doubled.

TABLE 1. AVERAGE DOWNLOAD PERFORMANCE FROM A FIELD TEST.

	WCDMA Rel-99	CDMA2000 1xEVDO
FTP/TCP object bit rate	280kbps	320kbps
UDP object bit rate	240kbps	600kbps
Measured latency (ping time to a server)	170ms	300ms

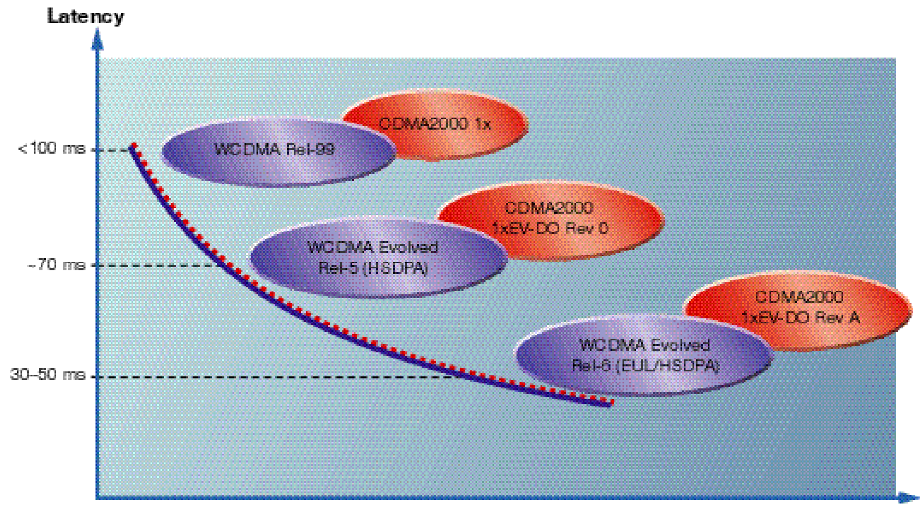

Figure 12
Latency road map for WCDMA2000.

These results stress the importance of low latency and indicate the potential of 1xEV-DO and HSDPA. Low latency is required to exploit the full performance potential of HSDPA and 1xEV-DO. Figure 12 shows the road map for latency and the targets that will enable improved end-user performance.

Figure 13
The bandwidth of WCDMA is wider than that of CDMA2000. As a consequence, WCDMA has a higher peak bit rate. The typical bit rate experienced in the field will thus also be higher for WCDMA Evolved than for CDMA2000 1xEV-DO.

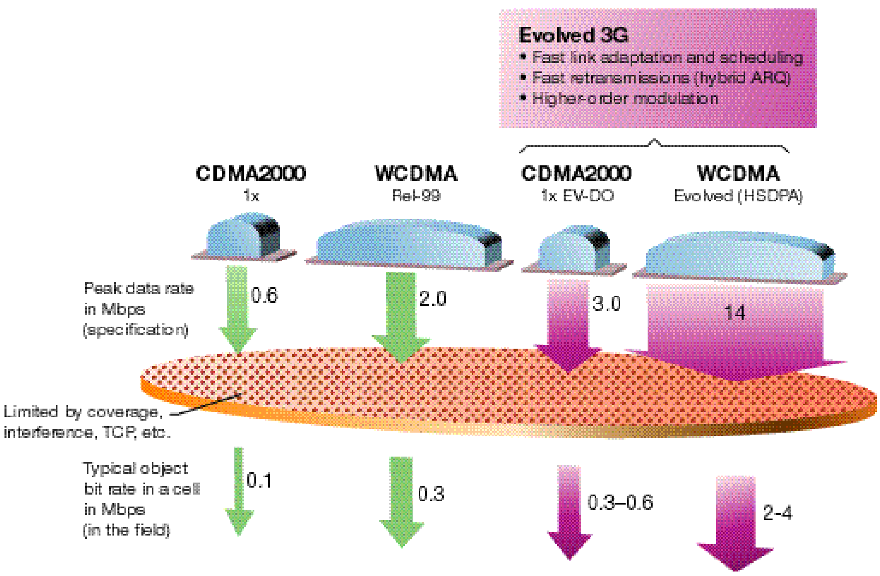

Ericsson has developed and built an experimental WCDMA HSDPA and enhanced uplink test bed that closely follows Rel-6 of the 3GPP specifications. The test bed, which is based on a commercial WCDMA Rel-99 network that has been updated with HSDPA and enhanced uplink functionality, can deliver peak data rates of more than 10Mbps in the downlink and 1.6Mbps in the uplink. The HSDPA and enhanced uplink test bed is operating over the air in Stockholm, Sweden. Its functionality and bit rates have been verified in the field. The user equipment (mobile terminal) is installed in a car. The RBS, RNC and core network are part of the Ericsson Experience Center in Stockholm. The test bed, which is used for customer demonstrations and performance measurements, has been in operation (with HSDPA functionality) since mid-2004.

Complementary technologies

Several complementary technologies are candidates for wireless broadband, including wireless LAN (WLAN), broadband wireless access (BWA), and short-range communications (such as Bluetooth). Each of these technologies has different properties in terms of peak bit rate, range, and mobility.

The IEEE 802 standards committee is working on several technologies. Of these, IEEE 802.16, driven by the WiMAX

Forum, is currently the BWA candidate with the broadest support.

WiMAX

The WiMAX industry forum has made IEEE 802.16 into an interoperable standard for broadband wireless access. Previous versions of the standard were designed for line-of-sight communication at higher frequencies. The first WiMAX products, based on published standard 802.16-2004, will be available in 2005. The 802.16e standard version (still under development) has broader support among vendors and will provide limited mobility. The first products for 802.16e are expected to arrive in 2007.

WiMAX can operate in FDD and TDD mode. It mainly addresses the 3.5GHz licensed and 5.8GHz unlicensed frequency bands. Unlike WCDMA and CDMA2000, WiMAX does not support full mobility. Instead, it will mainly support

- fixed or nomadic broadband wireless access as a complement to DSL when DSL is not available; and
- transmission backhaul for operators

WiMAX is defined for a range of bandwidths and can thus support numerous bit rates for the end-user. Line-of-sight (LOS) implementations give good coverage, but non-LOS implementations (such as indoor use or nomadic applications) limit the coverage as is true for any wireless technology. In similar deployments (LOS), WiMAX has similar coverage, bit rates and system capacity as WCDMA Evolved.

Conclusion

Third-generation system performance is dependent on numerous parameters. Deployment scenario, system load, propagation environment, and system features influence performance. To some extent, there is also a trade-off between end-user performance and operator performance (in terms of supporting many subscribers).

Field experience has shown that WCDMA can provide good performance for mobile broadband data, both for end-users and operators. WCDMA Evolved significantly improves the performance of best-effort packet data in WCDMA, with HSDPA providing up to 14Mbps in the downlink, and the enhanced uplink providing up to 5Mbps. Downlink bit rates of more than 10Mbps have been demonstrated in numerous field trials. A parallel evolution of CDMA2000 to 1xEV-DO gives the same kinds of improvement.

WCDMA Evolved improves the end-user experience by increasing peak bit rates and effective bit rates. It also improves, or reduces, latency. In addition, it supports more users thanks to greater system throughput per cell.

BOX D, COMPLEMENTARY TECHNOLOGIES

Flash-OFDM
Broadband wireless access technology developed by Flarion for IP communication. Designed to provide some mobility.
IEEE 802.11
Wireless local area network (WLAN) standard, mostly for home and office use. No mobility.
IEEE 802.16
Broadband wireless access (BWA) standard. Originally designed for transmission backhaul, now aiming at fixed/nomadic wireless access and limited mobility.
UMTS TDD
The "other" part of the UMTS standard designed for time-division duplex (TDD) spectrum. WCDMA has been designed for frequency-division duplex (FDD).
UWB
Ultrawideband. A short-range wireless technology for very high data rates. For applications similar to Bluetooth applications.
WiFi
Another name for 802.11 used by the WiFi Alliance.
WiMAX
Another name for 802.16 used by the WiMAX Forum.

REFERENCES

Parkvall, S., Englund, E., Malm, P., Hedberg, T., Persson, M. and Peisa, J.: WCDMA evolved—High-speed packet-data services. Ericsson Review, Vol. 80(2003):2, pp. 56-65

ACKNOWLEDGEMENTS

Eva Englund, Anders Furusk r, Per Beming, Jonas Wiorek and Janne Peisa

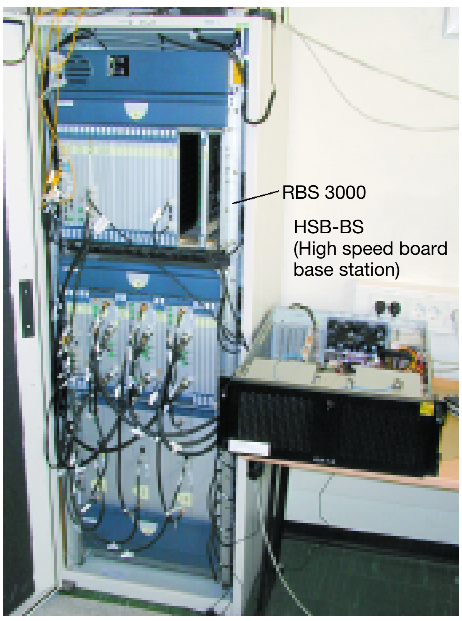

Figure 14
WCDMA Evolved test bed in Kista.