IETE Technical Review
Vol 21, No 5, September-October 2004, pp 305-315

An Overview of High-speed Packet Data Transport in CDMA Systems

PARVATHANATHAN SUBRAHMANYA, RAJESH SUNDARESAN
Qualcomm Inc., 675 Campbell Technology Parkway, Campbell, CA 95008, USA.
email : {subra, rajeshs}@qualcomm.com
AND
DA-SHAN SHIU
Graduate Institute of Communication Engineering, National Taiwan University, Taiwan
email : dsshiu@cc.ee.ntu.edu.tw

Rapid growth of data services in wireless communication systems requires an efficient means for high speed data transport. Modern CDMA systems have evolved to support this need. This paper presents an overview of the key techniques for improving high speed data transport in the following standards: 1xEV-DO (IS-856 Rev A), 1xEV-DV (cdma2000 Rev D), and HSDPA & EUL (WCDMA FDD Release 6). The emphasis is on identifying the common themes for high speed data transport in all these systems: adaptive modulation and coding, incremental redundancy and hybrid ARQ, smart scheduling to exploit channel variations. limited downlink macro-diversity aided by fast sector switching, and the use of shared channels to improve perceived latency.

INTRODUCTION

DATA services are growing rapidly in wireless communication systems. Examples include web-browsing, email, streaming video and music, multi-player games, instant messaging, mobile commerce, video telephony, and location based services. The increased usage of wireless data necessitates an efficient means for high speed data transport. Modern CDMA systems have evolved to support this need. This paper presents an overview of the key techniques for improving high speed data transport in modern CDMA systems. The emphasis is on identifying common themes for high speed data transport in all these systems.

Voice and packet data differ in many ways. See Table 1 for a list of key differences. Second generation (2G) wireless systems such as IS-95 and GSM/GPRS were optimized for voice transmission. They do not take advantage of the bursty and delay-tolerant nature of packet data to achieve greater data transport capacity. The bursty nature of packet data encourages a solution in which users share system resources and transmit/ receive data for short durations of time at high data rates. Delay tolerance allows smart scheduling algorithms to take advantage of variations in channel quality and schedule data transmission under the most favorable channel conditions. It also permits efficient retransmission techniques to recover data at the physical layer.

The main techniques involved in improving packet data transport efficiency are:

- rapidly adapting the modulation and coding scheme based on channel quality;
- adding redundancy incrementally through retransmissions until the packet is successfully received. At any stage, the receiver uses the coded information in all the (re)transmissions received up to that point in time;
- scheduling data transmission to and from a user at those time instants when that user's channel quality is relatively good;
- limiting the use of macro-diversity on the downlink; transmitting to each user only from the sector with the strongest instantaneous channel quality to that user.

The systems we will specifically discuss are: 1xEV-DO (IS-856 Rev A) [1-3], 1xEV-DV (cdma2000 Rev D) [4], and HSDPA [5-7] & EUL [8] (WCDMA FDD Release 6[1]). Table 2 provides a glossary of these terms. The specific revisions IS-856 Rev 0, cdma2000 Rev C and WCDMA Release 5 included techniques to improve the efficiency of packet data transmission on the downlink whereas IS-856 Rev A, cdma2000 Rev D and WCDMA Release 6 added techniques to improve the efficiency of packet data transmission on the uplink.

The rest of the paper is organized as follows. We discuss adaptive modulation and coding, incremental redundancy, scheduling, and limited macro-diversity &

[1]Standardization of WCDMA FDD Release 6 is still work in progress at the time of writing this paper. We will present some of the ideas being considered for standardization.

TABLE 1 Contrast of characteristics between voice and data

VOICE	PACKET DATA
Circuit-switched[2]	Packet-switched
Requires only a low data rate (<15 kbps)	Higher data rate improves user experience (100s of kbps to several Mbps).
Continuously active	Bursty
Delay-intolerant (requires latency less than a few hundred milliseconds)	Some packet data services can tolerate large delays (up to several seconds).

fast sector switching in the following sections. We then explain how the use of shared channels improves perceived latency. Finally, we indicate how voice can coexist with packet data in systems optimized for packet data transport. We conclude by summarizing our presentation.

ADAPTIVE MODULATION AND CODING

In this section, we discuss qualitatively how the use of adaptive modulation and coding technique can improve performance. We then briefly discuss the kind of information that needs to be exchanged between the transmitter and the receiver, and finally see how these techniques are implemented in the three standards.

Consider the transmission of information under varying channel conditions, as is typical in a mobile wireless environment. Suppose that the transmitter knows the instantaneous signal-to-noise ratio (SNR) at the receiver. Furthermore, suppose that the SNR does not vary for a certain interval. Then, for that time interval, the transmitter should choose a coding rate and modulation format based on that SNR so that the data can be transmitted as fast and as reliably as desired by the application.

For low SNR, QPSK transmission suffices for operation close to the information-theoretic channel capacity. For high SNR, higher order modulation techniques get us closer to capacity. Similarly, low coding rates are required for low SNR while higher coding rates and therefore fewer parity bits are sufficient for high SNR. These principles form the basis of adaptive modulation and coding techniques.

The transmission time intervals, the available channel coding rates and the modulation formats are given in the following Tables 2 and 3.

A consequence of adaptive coding and modulation is that there is significant fluctuation in the instantaneous data rate that a user receives. This fluctuation parallels the fluctuation in the channel SNR. Adaptive modulation and coding is therefore suited for applications that are tolerant to delay and rate fluctuations.

To support adaptive modulation and coding, the receiver needs a means to inform the transmitter of the

TABLE 2 Allowed physical layer configurations for the downlink

Specification	Number of physical layer configurations	Range of data rates	Transmission times	Coding rates	Modulation formats
1xEV-DO	12	38.4 kbps to 3.1 Mbps	1, 2, 4, 8 or 16 units of 1.66 ms.	1/5 to 2/3	QPSK, 8-PSK, 16-QAM
1xEV-DV	18	81.6 kbps to 3.1 Mbps	1, 2, or 4 units of 1.25 ms	0.0379 to 0.7188	QPSK, 8-PSK, 16-QAM
HSDPA	1890	Up to 13.9 Mbps	2 ms	1/3 to 1	QPSK, 16-QAM

[2]Voice can also be packetized and transported over packet data networks (Voice over IP). When transported thus, voice will be considered as a packet data service.

TABLE 3 Allowed physical layer configurations for the uplink

Specification	Number of physical layer configurations	Range of data rates	Transmission times	Coding rates	Modulation formats
1xEV-DO	11	4.8 kbps to 1.8 Mbps	4, 8, 12 or 16 units of 1.66 ms	1/5 to 2/3	BPSK, QPSK
1xEV-DV	11	19.2 kbps to 1.8 Mbps	10 ms	1/5 to 3/5	QPSK and 8-PSK
EUL	Under discussion	Up to 2-4 Mbps	2 ms or 10 ms	1/3 to 1	QPSK

instantaneous channel quality information, and the transmitter needs a means to signal the chosen data rate and modulation format to the receiver. These are accomplished in different ways. We go over some of these in the rest of this section.

For the downlink, the channel quality is estimated in all systems using downlink pilot transmission. This estimate is fed back on the uplink as soon as possible and as frequently as needed. See illustration in Fig 1. The feedback is a representation of the highest data rate the receiver can support for the measured SNR. The scheduler on the transmitter side chooses a data rate and modulation format based on the feedback. The chosen format is either the one requested by the receiver (in 1xEV-DO) or is signaled to the receiver via a downlink control channel (in 1xEV-DV and HSDPA). The details are summarized in Table 4.

The channel quality feedback takes up some capacity on the uplink. Therefore, in the standards covered in this paper, this feedback is carefully optimized to be represented using as few bits as possible. The same consideration holds for downlink configuration signaling.

On the uplink too, it is desirable to adapt the modulation and coding format to the instantaneous channel condition. Instead of relying on explicit signaling from the base-station, a terminal obtains the receiver SNR information in the following implicit fashion. A

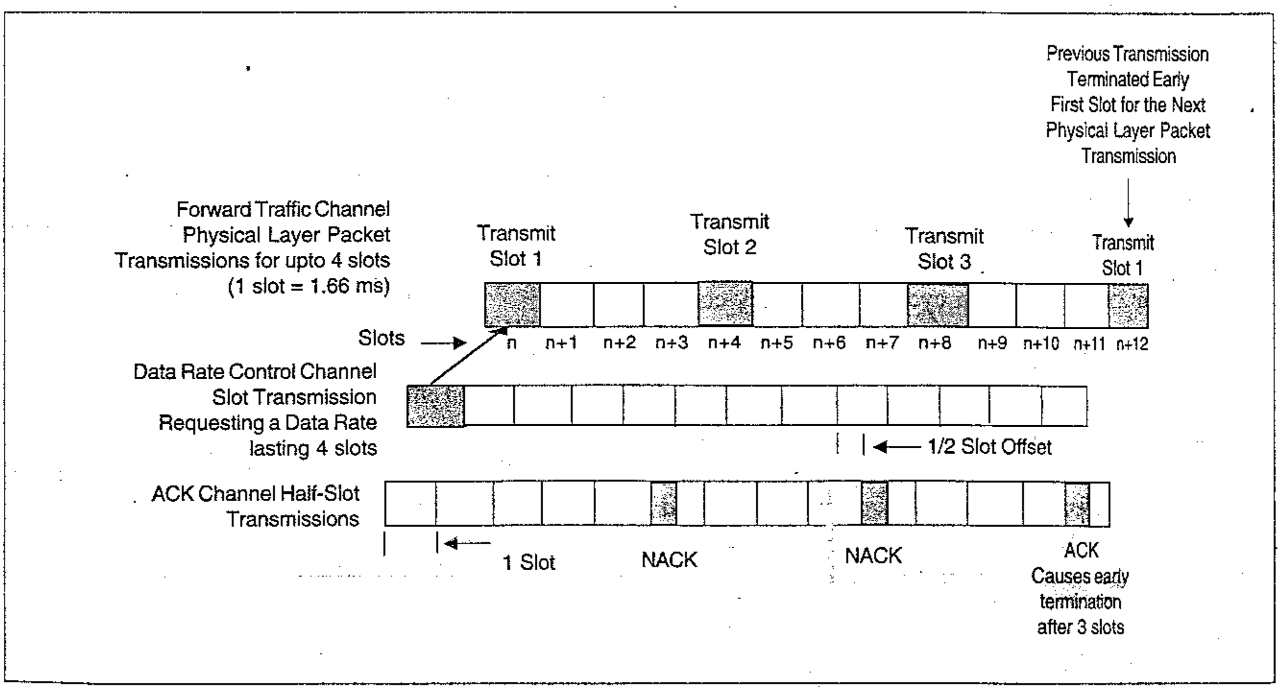

Fig 1 Downlink in 1xEV-DO: The terminal first makes an SNR measurement and requests a certain data rate lasting 4 slots on the Data Rate Control Channel. The transmitter schedules a transmission with the requested rate. The first two transmissions of the packet result in a CRC failure, and a NACK is issued. The third transmission's CRC succeeds. The receiver sends an ACK to terminate the transmission of this packet early. A new packet is then transmitted

TABLE 4 Channel Quality Feedback and Downlink Configuration Indication

Specification	Channel estimation	Channel quality feedback	Feedback information coding rate	Downlink signaling overhead
1xEV-DO	Using time-multiplexed pilots	4 bits / 1.66 ms	1/4 with repetitions of 1, 2, 4, or 8.	Only to identify recipient.
1xEV-DV	Using code-multiplexed pilots	4 bits / 1.25 ms or (4 bits / 20 ms + 1 bit / 1.25 ms)	1/12	Identifies recipient and chosen configuration
HSDPA	Using code-multiplexed pilots	5 bits/ (4/3 ms)	1/4	Identifies recipient and chosen configuration

terminal continuously transmits a reference uplink pilot channel. Such a channel is power-controlled by the base-station so that it is received at a constant SNR. The weaker the channel condition, the higher is the transmitted pilot power. Therefore, the instantaneous channel condition is implicitly communicated in the power control procedure[3]. A terminal can choose an appropriate coding rate and modulation format based on the transmitted pilot power. The chosen rate is signaled along with the payload. Additional peak power controlling mechanisms may be employed by the base-station. For details on the exact mechanism, refer to the section on Scheduling.

INCREMENTAL REDUNDANCY AND EARLY TERMINATION

The physical layer link reliability can be increased using incremental redundancy transmission. Suppose that the physical layer is asked to transmit a certain payload. The physical layer adds a few parity bits to enable error detection via a cyclic redundancy check (CRC). It also adds some form of coding to enable error correction. Suppose further that the code rate is 1/5, that is, 4 parity bits are added to every information bit. Depending on the number of physical channel bits available, and the channel condition, a few of these bits may be punctured or repeated to get to the desired coding rate. Let us assume that 2 of the 4 parity bits are punctured to get the coding rate to 1/3.

After decoding the payload, the receiver performs the CRC check. If it passes, it transports the payload to a higher layer. If not, it informs the higher layer of the error. The higher layer may request a retransmission (via the automatic repeat request or ARQ mechanism). Classical communication systems discard the original failed transmission. The retransmission must be received correctly on its own merit.

Better reception can be obtained if the retransmission is combined with the original transmission before the error correction mechanism (or decoding) is performed. A step further is to recognize that the retransmission need not be a replica of the original transmission, but could be additional parity bits; in particular, those that were initially punctured. As a consequence, after one or more retransmissions, the resulting code rate can be brought back to the base code-rate, thereby achieving the full coding gain afforded by the rate 1/5 code. This is transmission with incremental redundancy. Such a technique significantly improves the link efficiency.

Wireless channels exhibit rapidly-varying and unpredictable channel condition fluctuations. It is quite common for the transmitter's notion of the channel condition to differ significantly from the actual channel condition. Without incremental redundancy, the transmitter must employ a sizeable link margin to ensure reliable transmission. Incremental redundancy enables the transmitter to gradually lower the coding rate. In the fortuitous situation when the channel SNR is high, fewer parity bits suffice, and the packet transmission can thus be terminated early.

Incremental redundancy requires storage of symbols for combining with retransmissions. The storage requirements are typically several bits per symbol. To

[3]Strictly speaking, when the mobile is in soft handoff, the relationship between channel quality and power control holds only for the base-station to which the mobile has the strongest uplink.

keep this memory requirement manageable, the turn-around time for retransmission is made quite small, typically on the order of a few tens of milliseconds. This requires the receiver to make a quick decision on whether the decoding succeeded and to send the indication back to the transmitter as soon as possible. Similarly, the transmitter needs to react to the success or failure indication in a timely fashion.

All three standards use the "stop-and-wait" protocol on the downlink. Reception of acknowledgements (ACK) by the base-station will result in the transmission of a new packet. Reception of negative acknowledgements (NACK), or the mere absence of ACK, can result in the retransmission of the packet via a new redundancy version. The stop-and-wait protocol forces the transmitter to wait before making a new transmission. In order to make better use of the wait time, parallel processes are established, so that while one process is waiting for a retransmission, another transmission is in progress. In all three systems the feedback on the status of a demodulation is transmitted on the uplink at a fixed time offset from the corresponding downlink transmission. All these concepts are illustrated in Fig 2.

In 1xEV-DO and 1xEV-DV downlink, the multi-slot configurations support transmission with incremental redundancy. However, certain high data rate configurations last only 1 slot and hence do not support incremental redundancy. In HSDPA, on the other hand, all configurations are allowed retransmissions. There is no limitation on the number of retransmissions. On the uplink, incremental redundancy transmission is supported in all three systems: 1xEV-DO, 1xEV-DV and EUL.

SCHEDULING

The bursty nature of packet data suggests a solution that shares system resources to take advantage of statistical multiplexing gains. On the downlink, it is natural to think of time-slots, channelization codes, and transmit power as system resources to be shared. On the uplink, multiple users access the channel simultaneously using different pseudo-noise codes. It is natural to think of time-slots and transmit data rate (equivalently transmit power) as system resources in this context because mutual interference between users can be decreased if different users transmit at different time-slots or at lower data rates (using less transmit power).

Resource sharing raises the following question. How best can limited system resources be shared among competing users? If maximizing system data throughput

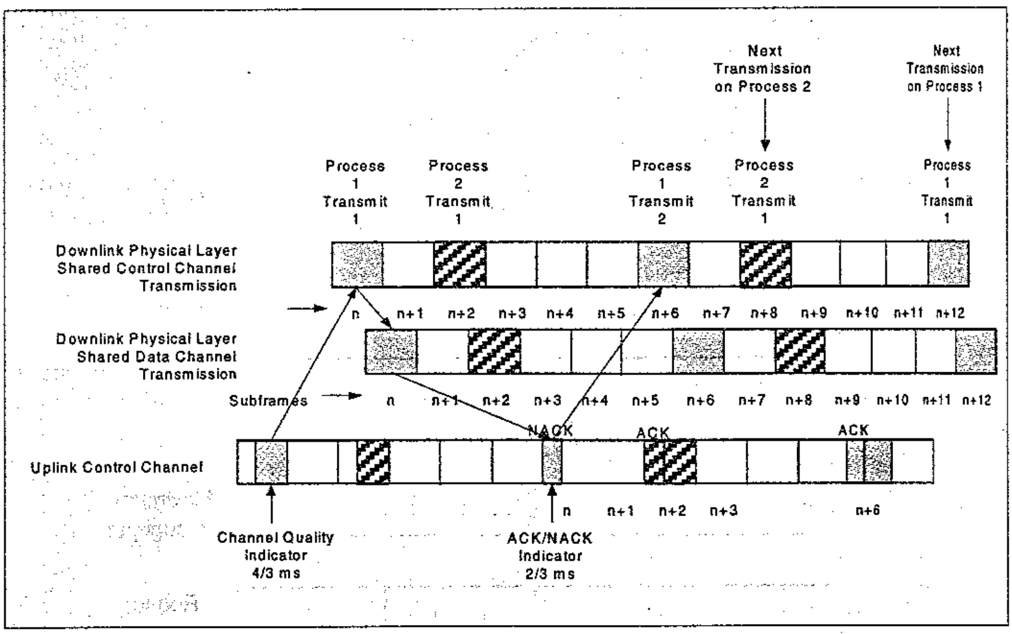

Fig 2 WCDMA Release 5 HSDPA: The terminal first provides information on the data rate it can support given the current SNR, using the channel quality indicator (CQI) field. This lasts 4/3 ms. A subframe is 2 ms long. The base-station schedules a transmission in subframe n on the process 1 (the shaded transmissions). Observe that the control channel indicating the chosen configuration and the data or traffic channel are themselves offset by 4/3 ms. This allows the receiver to prepare in advance for the demodulation. Particularly, the chosen modulation format is known prior to traffic demodulation. In this example, the receiver then decodes the received symbols, and sends a NACK in the uplink channel. This ACK/NACK indicator lasts 2/3 ms and is time-shared with the CQI channel. The base-station makes a second transmission in subframe $n+6$ in response to this NACK. Meanwhile, the base-station schedules another transmission (the hatched transmissions) on process 2 in subframe $n+2$. This increases the peak throughput to the terminal by utilizing the subframes lying in-between successive transmissions on the same process. The standard allows up to 8 processes to transmit concurrently to a terminal

s the goal, system resources at each time instant must e allotted to the user with the best channel quality. This pproach, however, does not treat users fairly; in Fig 3 (Maximum Throughput), user 3 is not allocated any resources for the entire time period shown. This can lead to unacceptably large delays in the transmission of data to or from such users and calls for a more equitable distribution of system resources. However, any scheme that enforces fairness towards users necessarily sacrifices system throughput.

Loosely speaking, we can achieve a sense of fairness and yet maintain a high system throughput by scheduling each user's transmission around the local maxima of his channel quality, as long as the transmission occurs within the user's delay requirements. Since different users' channels undergo independent time variations, the local maxima of different users' channel quality variations are likely to occur at different time instants, leading to an effective data rate that matches the peak channel qualities of all the users. The throughput gain that can actually be realized by such a scheduler depends on the time-scale of channel variations relative to the user's service delay constraint. If the channel varies significantly over the period within the delay constraint, it is likely that the channel quality will achieve its local maximum within this period, leading to high system throughput.

An example of a scheduler that trades-off fairness and throughput is the proportional-fair scheduler (PF). The proportional-fair scheduler tracks the short-term average served data rate of each user. It defines a metric that increases with increasing channel quality and decreases with increasing served data rate. The user with the largest metric at any given instant is allocated system resources. One example of such a metric is the relative instantaneous channel quality (RICQ) metric. It is the ratio of the instantaneous achievable data rate to the average served data rate over a suitable averaging interval. Since RICQ metric increases with improvement in channel quality, users with better instantaneous channel quality will be favored. At the same time, if a user with poor channel quality is not served for some period, his average served data rate diminishes. Hence, that user's RICQ metric will increase, causing him to be served

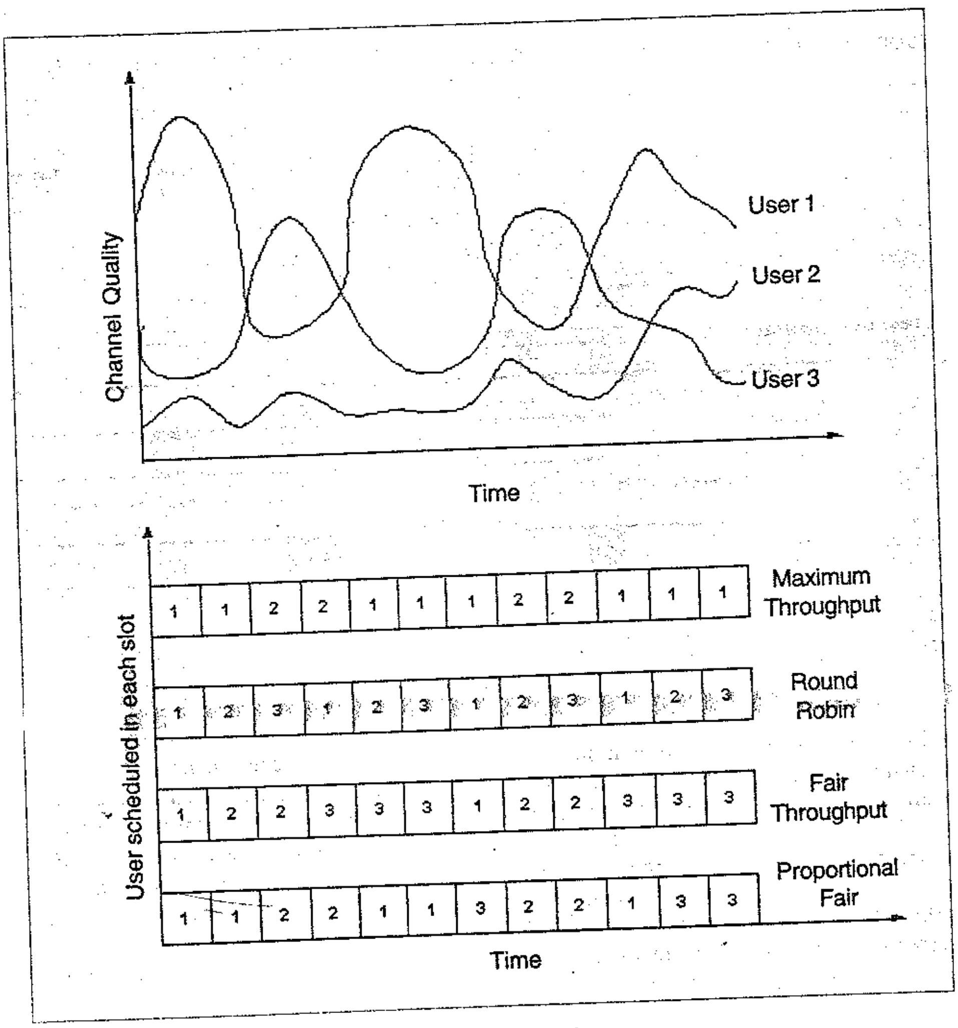

Fig 3 Schedulers

preferentially. In this sense, PF enforces a degree of fairness while striving to maximize system throughput.

A real-life scheduler will also need to consider several other factors such as relative priorities of different types of user data based on QoS requirements, buffer occupancies of different users, differential treatment of users based on pricing plans, and so on. The scheduling algorithm itself is not specified in any of the standards and is left open to implementation. Scheduling algorithm design continues to be an area of active research.

The downlink scheduler uses channel quality measurements reported by the mobile on the uplink to schedule transmission for one or more users during the upcoming time-slot. Some relevant downlink scheduler characteristics are listed in Table 5.

Uplink scheduling is initiated based on requests received from mobiles. The request could indicate the type of service that generated the data, the amount of data to be transmitted and the amount of power headroom at the mobile. The scheduler responds with a "grant" based on requests from multiple mobiles and other information that it has access to. The grant may specify a set of times during which the mobile may transmit, and the maximum power ratio it can maintain between the uplink data channel and the uplink pilot channel during this period. This ratio is referred to as the traffic-to-pilot ratio (T/P). A higher data rate typically requires a higher T/P; thus, a limit on T/P is equivalent to a limit on the highest data rate the mobile is allowed to transmit. The network may control the transmission by a user even after a grant has been issued by instructing the mobile to increase or decrease the maximum allowed T/P ratio in steps. This procedure is illustrated in Fig 4.

For system stability, the ratio of the total base station received power to the thermal noise level (referred to as rise-over-thermal or RoT) must be kept within predefined limits. RoT is one factor the scheduler may consider in determining the maximum T/P for a user since a larger T/P leads to a larger contribution to RoT. Other factors the scheduler may consider include the downlink channel quality reported by a mobile, the uplink power control SIR target, and the mobile's soft-handoff status. A poor downlink channel quality, a high SIR target, or a mobile in soft-handoff are all indications that the mobile's transmission is likely to cause a higher level of interference on the uplink. The scheduler could manage this interference by decreasing the T/P in the grant or by deferring a grant until channel conditions improve.

LIMITED DOWNLINK MACRO-DIVERSITY AND FAST SECTOR SWITCHING

Macro-diversity or soft handoff (or handover) is a key feature of CDMA systems. While a mobile is in soft handoff with two or more sectors[5], each sector transmits to the mobile on the downlink, typically at close to the same power as all other sectors. Each sector also receives the mobile's uplink transmission. Establishment of a link with a new sector before tearing down the link with an older sector improves the reliability of sustaining a call through the handoff process. Further, since the channel conditions between the mobile and the multiple sectors vary independently in the presence of fading, soft handoff provides a form of diversity gain.

On the uplink, soft handoff is always beneficial because demodulation of uplink transmission by additional sectors improves the chances of successfully decoding the payload. On the downlink, however, macro-diversity may not be desirable. First, the total transmit power from all the sectors is minimized if the sector with the best instantaneous channel is the only one transmitting to the mobile. Second, soft-handoff transmission requires the allocation of code resources and/or time-slots at

TABLE 5 Downlink scheduler parameters

Specification	Time-slot	Fraction of code space allotted to shared channel	Number of instances of shared channel
1xEV-DO	1.66 ms	1	1[4]
1xEV-DV	1.25 ms	Up to 28/32	1 to 2
HSDPA	2 ms	Up to 15/16	1 to 15

[4]Although there is only one instance of the shared channel, data from multiple users can be multiplexed on to the same physical layer packet. This improves packing efficiency for low data rate, delay-sensitive services.

[5]The mobile may be in soft handoff with sectors belonging to one or more cells. If all the sectors belong to the same cell, it is referred to as softer handoff.

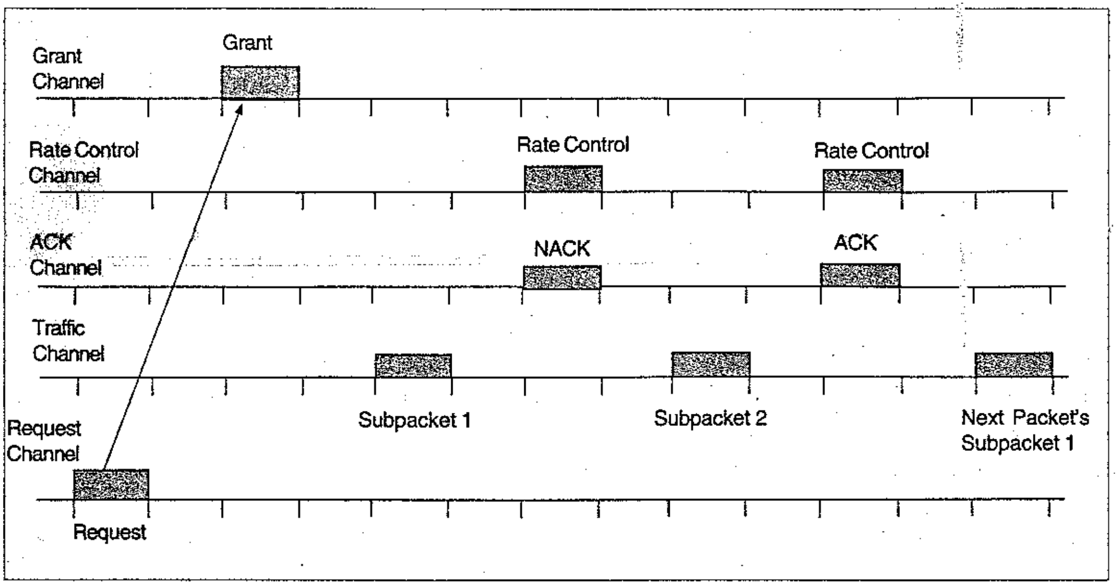

Fig 4 Uplink in 1xEV-DV: The terminal makes a request by providing relevant information on QoS and data rates. The base-station responds with a grant that contains information on allowed rate configurations. The terminal then starts transmission of subpackets. Each new subpacket is a new version of the original transmission. In the above example, the first subpacket transmission fails the CRC, and the second one passes. During the transmission of this packet, the base-station provides additional rate-control information

multiple sectors. Third, coordination of transmission from multiple sectors complicates the execution and limits the flexibility and performance of scheduling algorithms. In light of these issues, all standards covered in this paper favor the approach in which downlink transmission occurs from only one sector. They all provide mechanisms that allow for fast re-pointing to the best sector.

In 1xEV-DV and 1xEV-DO, a mobile indicates the sector with the strongest downlink channel quality explicitly in its uplink transmission. Multiple sectors monitor this uplink transmission concurrently. Once a sector identifies that it has the best downlink to the mobile, it assumes the role of the serving sector. This mechanism is accomplished through fast physical layer signaling, and hence the name fast sector switching. Sector switching is achieved in HSDPA using a slower procedure based on radio resource control (RRC) signaling.

QUALITY OF SERVICE IMPROVEMENT DUE TO SHARED CHANNELS

In traditional CDMA standards, data transfer is achieved by first assigning a code and a data rate to a mobile. Then the transmitter (either a base station or a mobile) is allowed to transmit, continuously in time, at that data rate until it exhausts its data buffer. Because all such communication takes place continuously in time, each communication link is allocated a data rate that is a fraction of the aggregate data rate associated with the base station.

This mechanism – multiple parallel dedicated narrow pipes – is suitable for voice traffic. However, for bursty data traffic with a central control node, the opposite mechanism – few shared fat pipes – can yield a superior user experience. We explain this concept with the following example. Suppose all the mobiles in a sector request to send some equal-sized blocks of data at virtually the same time. We define latency as the time elapsed between the transfer request and the completion of data transfer. The shorter the latency, the better is the user perception. With the dedicated channel approach, each mobile is given the same data rate, and hence every mobile sees the same latency. On the other hand, with the shared channel approach, the data requests are serviced on a first-come, first-serve basis. Hence, every mobile sees a different latency. Nevertheless, even the worst-case latency with the shared channel approach is the same as the best case latency with the dedicated channel approach. Furthermore, since it is highly likely that the data requests from different mobiles are initiated at staggered time instants, the service latency can be much lower using shared channels. This is illustrated in Fig 5.

The modern CDMA standards covered in this paper all employ shared channels for packet data transport instead of dedicated channels. This gives them an advantage of improved latency over traditional standards. Even though it is possible to achieve the shared channel effect in traditional CDMA standards by quickly re-configuring/re-allocating dedicated channels, the signaling interfaces are not designed to support such frequent and rapid reconfigurations in an efficient fashion.

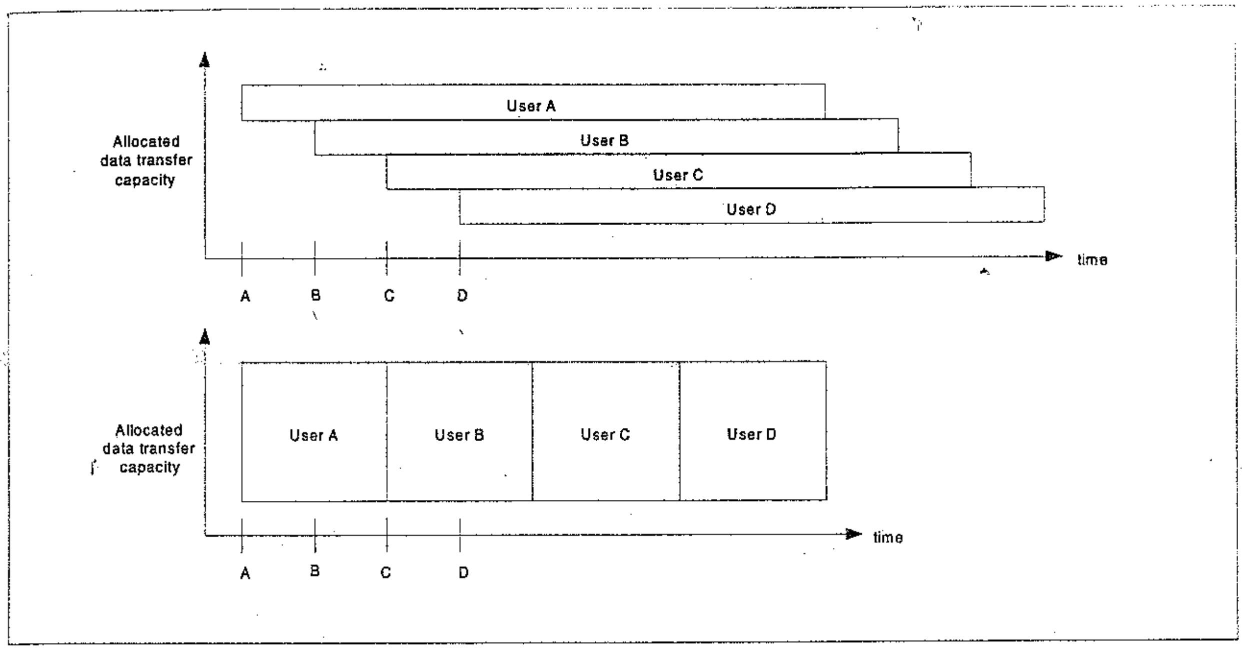

Fig 5 Example to show the difference in latency between the dedicated channel approach and the shared channel approach. Suppose user A requests data transfer at time A, user B at time B, and so on. In (*a*) each user is given an equal fraction of the total data transfer capacity. As a result, all four users experience the same worst-case latency. In (*b*) each user is served in succession. The median latency in (*b*) is less than half its counterpart in (*a*)

CO-EXISTENCE OF PACKET DATA USERS AND VOICE USERS

Thus far, we described the advantages that modern CDMA systems offer over traditional CDMA systems for supporting bursty data traffic. We now address approaches that enable voice traffic over such systems.

The stringent latency requirements of voice applications may not allow the benefits described in this paper to be realized to the fullest extent. Tight delay constraints may force the scheduler to schedule a user's transmission under suboptimal channel conditions. Furthermore, high speed packet-data transport requires a channel reconfiguration overhead per packet. For instance, the base station must alert a user every time a downlink or an uplink resource is allocated to the user. Additional signaling is needed to support adaptive modulation and coding and incremental redundancy. Voice applications typically transport approximately 200 bits every 20 ms – a low data rate. The signaling data rate is therefore a significant fraction of the voice data rate. Moreover, reliable communication of these signaling channels requires a significant power overhead as well.

One approach to carry voice in such systems is to transmit voice over shared channels with some methods to reduce the associated overheads. The alternative approach is to transmit voice over dedicated low data-rate channels. The former approach is employed by 1xEV-DO, while the latter is employed by 1xEV-DV and HSDPA.

In 1xEV-DO, downlink voice data for several users can be multiplexed onto a single larger physical layer packet. This reduces the signaling overhead per user and increases the packing efficiency on physical layer packets, while still preserving the coding efficiency advantage of a larger physical layer packet. The modulation and coding rate for the packet must be chosen such that all the users with data in the multiplexed packet can receive it with the desired reliability. Subject to this constraint, multi-user diversity can be exploited wherever possible.

In 1xEV-DV and HSDPA, each user's voice data is transported over a dedicated channel. To cope with channel quality variations, the service provider typically dedicates downlink power resource so as to provide guaranteed quality of service for each user under a wide range of channel conditions. This leaves a significant fraction of the base-station power unused most of the time. This excess resource can be allocated to a shared channel that can carry delay-tolerant packet data. The transmit power allocated to this shared channel can be varied on the time-scale of milliseconds. This interval is small enough that the excess power resource available can be considered roughly a constant over the interval. The excess power when available is therefore better utilized to improve system capacity. This idea is illustrated in Fig 6.

TABLE 6 Glossary

IS-95	The first CDMA-based cellular system (primarily serving voice users)
cdma2000	An evolution of IS-95 supporting higher data rates and greater system capacity. The downlink improvements described in this paper were added in Rev C while the uplink optimizations were added in Rev D.
1xEV-DV	1x Evolution – Data and Voice. A name commonly used to refer to cdma2000 Rev C and Rev D
IS-856	A CDMA system optimized for data transmission. Also referred to as CDMA/HDR. IS-856 Rev 0 optimized downlink performance while IS-856 Rev A added uplink optimizations.
1xEV-DO	1x Evolution – Data Optimized. A name commonly used to refer to IS-856
WCDMA	Wideband CDMA. The downlink improvements described in this paper were added in Release 5 while the uplink improvements are expected to be added in Release 6.
HSDPA	High Speed Downlink Packet Access. Refers to the downlink packet data improvements added in WCDMA Release 5.
EUL	Enhanced Uplink. Refers to the uplink packet data improvements added in Release 6 of WCDMA.
Downlink or Forward Link	The link from the base station to the mobile
Uplink or Reverse Link	The link from the mobile to the base station
Cell	A base station in a cellular system. It may comprise of multiple sectors. Also referred to as a node B in WCDMA terminology.
Sector	A part of a cell typically covered using directional antennas. A sector is referred to as a cell in WCDMA terminology. In this paper, we consistently refer to it as a sector.

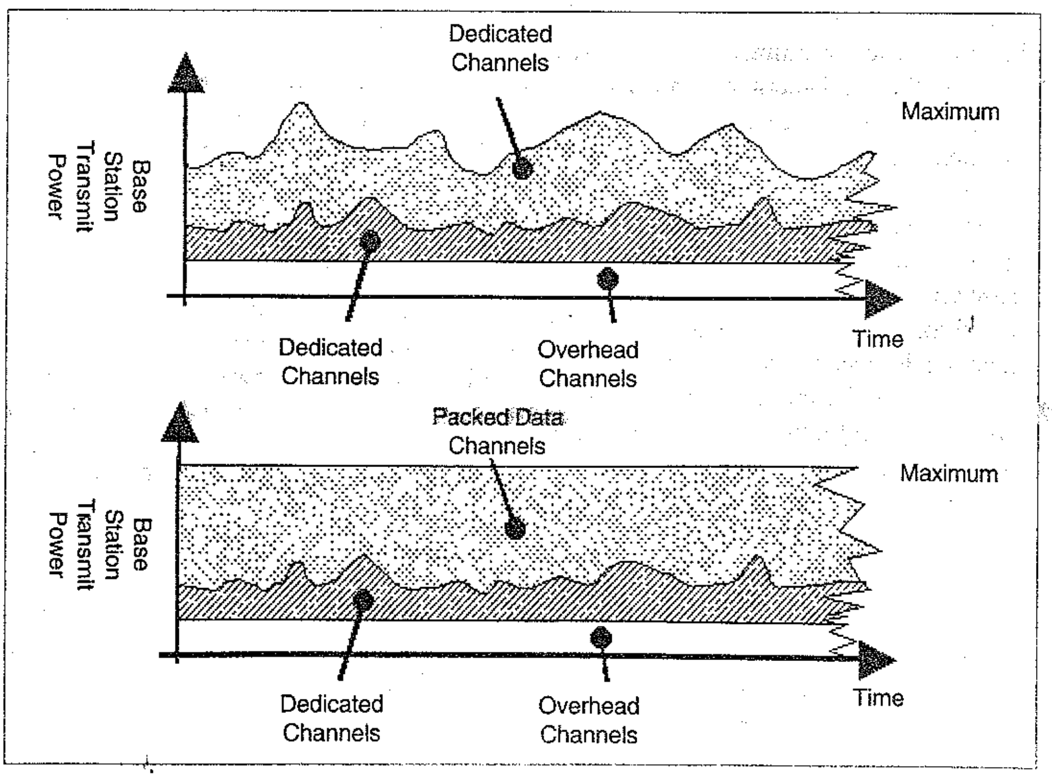

CONCLUSIONS

In this paper, we have seen how modern CDMA systems have evolved to support the need for high speed data transport. In particular, we covered the following standards: 1xEV-DO (IS-856 Rev A), 1xEV-DV (cdma2000 Rev D), and HSDPA & EUL (WCDMA FDD Release 6).

The key concepts for improving high speed data transport in these standards are adaptive modulation and coding, incremental redundancy transmission, smart scheduling to exploit channel variations, limited downlink macro-diversity aided by fast sector switching, and the use of shared channels to improve perceived latency. We indicated how these concepts matched well with the characteristics of bursty data services. In addition to describing the common concepts, we also provided examples, where appropriate, of implementation in the standards. We also saw how voice can coexist with packet data in modern CDMA systems.

It is expected that these systems with greatly improved packet data performance will enable economical wireless data access over wide-area as well as metropolitan-area networks, thereby accelerating the growth of wireless data services.

REFERENCES

1. TIA/EIA/IS-856 cdma2000 High Rate Packet Data Air Interface Specification, *Telecommunications Industry Association, (www.3gpp2.org)*
2. P Bender, P Black, M Grob, R Padovani, N Sindhushayana, & A Viterbi, CDMA/HDR: A Bandwidth-Efficient High-Speed Wireless Data Service for Nomadic Users, *IEEE Communications Magazine*, July 2000.
3. N T Sindhushayana & P J Black, Forward Link Coding and Modulation for CDMA2000 1xEV-DO (IS-856), pp. 1839-1846, PIMRC 2002.
4. TIA/EIA/IS-2000.2 Physical Layer Standard for cmda2000 Spread Spectrum Systems, *Telecommunications Industry Association, (www.3gpp2.org)*
5. WCDMA FDD Specifications Release 5, *(www.3gpp.org)*
6. S Parkvall, E Englund, P Malm, T Hedberg, M Persson, & J Peissa, WCDMA Evolved – High speed packet data services, Ericsson Review No. 2, 2003.
7. T E Kolding, K I Pedersen, J Wigard, F Frederiksen, P E Morgensen, High Speed Downlink Packet Access: WCDMA Evolution, IEEE Vehicular Technology Society News, Feb. 2003.
8. TR 25.896 – Feasibility Study for Enhanced Uplink in UTRA FDD, *3rd Generation Partnership Project (www.3gpp.org)*

Authors

P Subrahmanya received his BTech degree in Electronics and Communication Engineering from IIT Madras in 1991 and his MS and PhD degrees in Electrical Engineering from Cornell University in 1994 and 1997 respectively. He is currently with Qualcomm Inc., where he leads the system design effort on one of Qualcomm's ASICs implementing HSDPA. His interests include terrestrial and satellite wireless communication systems, VLSI signal processing for communications, and cross-layer system design.

* * *

Rajesh Sundaresan received the B.Tech. degree in electronics and communication engineering from Indian Institute of Technology, Madras, in 1994, the MA degree in electrical engineering from Princeton University, Princeton, New Jersey, in 1996, and then PhD degree, also from Princeton University, in 1999. Since 1999, he has been with the CDMA technologies division of Qualcomm Inc., Campbell, California. His current interests include wireless communication systems and networks, signal processing, information theory, and statistics.

* * *

Da-shan Shiu received the BSEE degree from National Taiwan University, Taipei, Taiwan, ROC in 1993, and the PhD degree in Electrical Engineering and Computer Sciences from the University of California at Berkeley in 1999. From 1999 to 2004, he was with Qualcomm Incorporated, where he was involved in the development of signal processing algorithms for WCDMA/HSDPA products. In August 2004, he joined the Graduate Institute of Communication Engineering of National Taiwan University as an Assistant Professor. His current research interests include MIMO space-time wireless communication techniques, next generation LAN, cellular, and wireless broadcast systems, wireless multimedia applications, information theory, and wireless optical communications.

* * *