How DF Raptor Is Used in MBMS

March 2007

The Multimedia Broadcast/Multicast Service (MBMS) technical specification defines mobile multimedia services over GSM-based 3rd generation (3G) cellular networks. MBMS multimedia content is delivered via IP packets, either as a streaming service or as a file download service to the end user. In streaming services, a continuous data flow of audio and/or video is delivered to the end user's handset; in download services, a finite amount of data is delivered to the user in its entirety as a file.

The MBMS technical specification requires mandatory use of Digital Fountain's advanced forward error correction technology, DF Raptor™, in order to ensure that the 1-way point-to-multipoint transmission of multimedia content is delivered to end users efficiently and with the best possible quality. If MBMS did not intrinsically provide some means to counter the effects of packet loss by using DF Raptor, streaming services would suffer dropouts and distorted audio/video, and download services would incur delays and waste bandwidth in delivering files to end users. At the same time, however, the MBMS technical specification does not detail exactly how DF Raptor should be used, leaving it to the network operators and equipment vendors to choose the various DF Raptor parameters as they see fit.

MBMS Fundamentals

MBMS is a point-to-multipoint IP-based service carried by the 3G air interface of W-CDMA (UMTS Terrestrial Radio Access Network or UTRAN) or the 2.5G air interface of EDGE/GPRS (GSM EDGE Radio Access Network or GERAN). The W-CDMA air interface provides bearer rates up to 256 kbps or greater, while the EDGE/GPRS air interface provides bearer rates up to 128 kbps. ***Figure 1 Simplified MBMS architecture:***

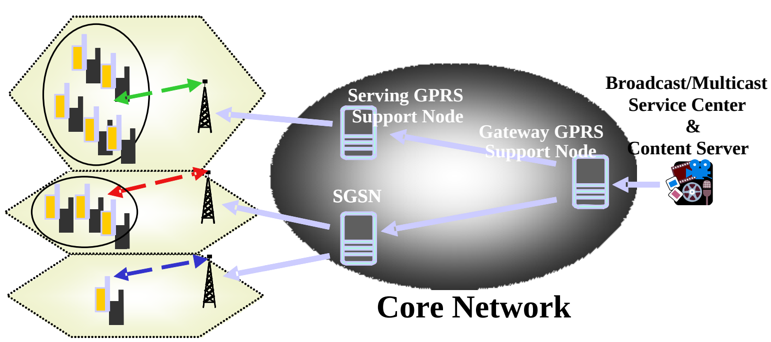

3G Radio Access Network

- UMTS Terrestrial RAN (W-CDMA)
- GSM/EDGE RAN

An MBMS session consists of three phases:

- An initial phase to announce and set up MBMS services, where either 2-way point-to-point communications over TCP/IP or 1-way point-to-multipoint data transmissions are employed
- A delivery phase to convey MBMS streaming and download services, where 1-way point-to-multipoint data transmissions are supported by protocols carried by UDP (User Datagram Protocol) over IP
- A post-delivery phase to optionally report the quality of content reception or ensure that a specific user has received a complete file download, where 2-way point-to-point communications over TCP/IP are employed

For the primary MBMS function of content delivery via point-to-multipoint data transmission, UDP provides a minimal 1-way transport service that supports data checksums and multiplexing by port number but cannot guarantee delivery of the data. Given that this MBMS transmission is 1-way, there is no feedback path available to ensure that an individual receiver receives the best possible signal. Unlike protocols that use 2-way communications, custom adjustments to the transmit power level, modulation type, or data rate are not possible, and a specific receiver may receive corrupted data. Irreparably corrupted UDP packets are discarded by the physical and transport layer protocols and thus are effectively "lost". Similarly, packets may be lost if the wireless signal is momentarily interrupted as the user enters an area of limited coverage or turns off the cell phone or if network congestion cause packets to be discarded before they are even transmitted over the air interface. Because lost packets cannot normally be retransmitted to each individual MBMS receiver, MBMS employs DF Raptor as a way for each receiver to recover its lost packets autonomously.

Forward Error Correction (FEC)

In general, forward error correction (FEC) technologies protect data by encoding the original source data at the sender such that redundant data is added. The FEC decoding algorithm then allows the receiver to detect and possibly correct errors in the original data based solely on the data that has been received. The error correction is "forward" in the sense that no feedback from the receiver to the sender or further transmission by the sender is required.

MBMS employs FEC technology in two complementary ways: first, at the physical layer, using the channel coding scheme of the underlying 3G air interface to correct bit errors, and, second, at the application layer, using DF Raptor to recover lost packets. While the FEC at the physical layer attempts to ensure that all the bits that are actually received are correct, the FEC provided by DF Raptor attempts to restore missing data.

FEC at the Physical Layer

To correct bit errors at the physical layer, FEC encoding and decoding algorithms are implemented in hardware as part of the W-CDMA or EDGE/GPRS air interface. On the send side at each base station, the IP packets to be transmitted by MBMS are first partitioned as required by the specific air interface into data blocks for transmission and then encoded according to the particular channel coding scheme. On the receive side at each mobile receiver, the decoding algorithm then attempts to correct bit errors that may have been induced by noise, interference, multipath, and other impairments over the wireless channel. If too many bit errors are present in a data block for the FEC decoding algorithm to correct, then the block is unrecoverable and the associated IP packet(s) are lost. The block error rate (BLER or block erasure rate) then represents the failure rate of the air interface FEC protection to recover the data blocks as partitioned at the physical layer.

FEC at the Application Layer -- DF Raptor

To provide packet-level protection at the application layer, DF Raptor is implemented at the MBMS server on the send side and on each mobile receiver. Because DF Raptor's algorithms are mathematically efficient, hardware support is not necessary, and the DF Raptor encoders and decoders are typically implemented in software.

DF Raptor acts to complement the bit-level FEC protection provided by the physical layer, recovering those IP packets that have been lost because the receiver's air interface FEC was unable to fully correct the underlying blocks of transmitted data. In addition to recovering packets that have been discarded because excessive physical layer errors were present, DF Raptor can also recover those packets that were lost because network congestion resulted in some packets being discarded prior to over-the-air transmission or because they were simply never received due to temporary outages, the receiver is turned off or the receiver is out of coverage.

At each MBMS server, the DF Raptor encoder processes a source block of data – for example, a complete data file to be transmitted or a segment of a continuous data stream – into a larger data block composed of the original source data and any desired amount of additional "repair" data. DF Raptor as adopted by MBMS is termed a systematic code because the encoding algorithm makes the original unencoded data available at the receiver in this way. How much repair data is added by the DF Raptor encoder depends on whether it is being used to protect an MBMS streaming or download service, the anticipated network conditions, the desired quality of delivery to be offered, and the amount of additional bandwidth or transmission time that can be used. The technical specifications for MBMS do not stipulate how much repair data to use, allowing considerable flexibility in the use of DF Raptor by individual equipment vendors and network operators.

At the receiver, the presence of packet loss along the transmission path means that not all the source and repair data may be successfully received. As long as “enough” data is received, however, the DF Raptor decoder can completely recover the original source data. If all the source data is received successfully, of course, then no decoding is needed at all. Otherwise, the amount of data that needs to be successfully received in order to fully recover the source data is only slightly greater than the original size of the source block. It makes no difference to the DF Raptor decoding algorithm whether source or repair data has been received -- as long as the total received source and repair data is only minimally larger than the size of the original source data, then the original data can be fully recovered with high probability. If, however, packet loss is such that insufficient data is available for DF Raptor to decode a source block, then only the parts of the source data block that was directly received will be available for that block.

Packet Loss Protection for Streaming Services

MBMS streaming services are delivered using RTP (Real-time Transport Protocol) as specified by RFC1889 of the Internet Engineering Task Force (IETF), operating over UDP/IP.

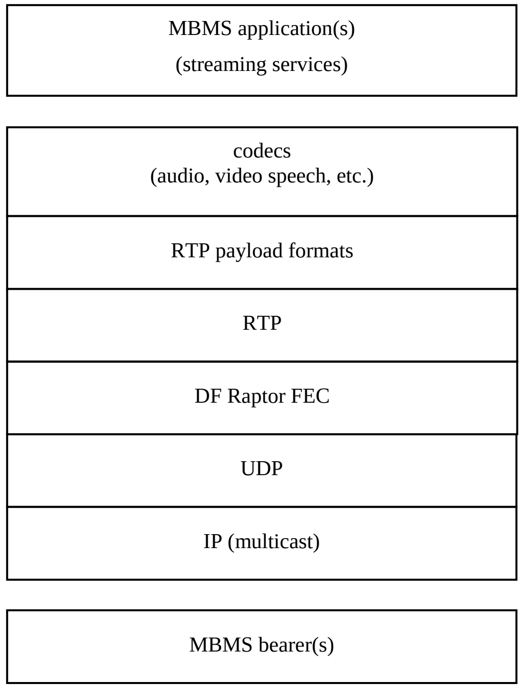

Figure 2 The MBMS protocol stack for streaming services

RTP is designed to carry real-time content, providing timestamps, identification of the type of payload, sequence numbering, and control mechanisms to synchronize different streams while respecting their relative timing properties. The sequence numbers included in RTP allow the receiver to reconstruct the sender's original packet sequence, but RTP does not guarantee delivery or prevent out-of-order delivery.

To improve the delivered quality of the multimedia stream, DF Raptor provides packet-level protection of the RTP data stream so that lost packets can be independently recovered by each receiver without requiring retransmission. At the sender – an MBMS server -- the original stream of RTP packets (or a bundled set of multiple content streams) is partitioned into consecutive source blocks of data that are treated independently. Each successive source block is encoded by DF Raptor, generating more packets than were originally present in the source block in order to compensate for potential packet loss. A new stream is then transmitted consisting of the original source packets followed by additional repair packets that are sent to a UDP destination port different from any of the original packets' destination port(s). By transmitting the source packets along with repair packets in this way, playback applications can make use of whatever source data has been received even if DF Raptor is unable to completely recover the full source block.

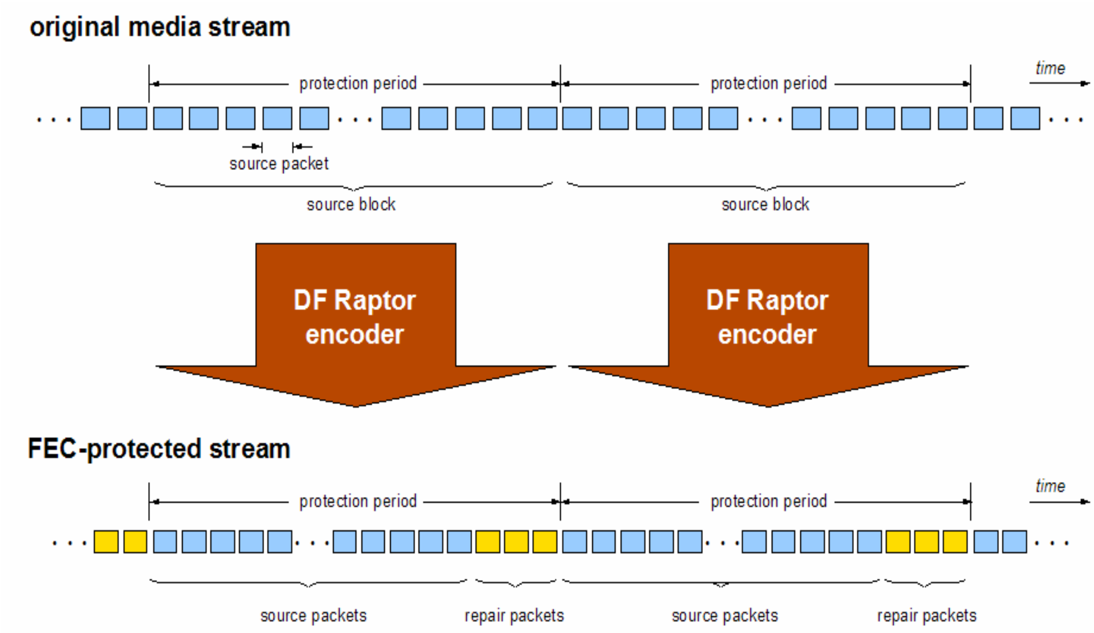

Figure 3 DF Raptor encoding of streaming media

At each receiver, the received stream of source and repair packets is processed as blocks and decoded by DF Raptor. It does not matter which of the original and repair packets have been lost – as long as enough data has been successfully received, then each original block can be fully recovered. Successive blocks of the received stream of source and repair packets are processed by the DF Raptor decoder in turn so that the original source stream is available for subsequent playback processing.

A number of operational considerations are involved in configuring DF Raptor to protect MBMS streaming services. For example, both the size of the source block and the number of additional repair packets produced from each source block are variables that can be set by the MBMS server as desired, directly affecting the performance of DF Raptor in protecting against packet loss. But trade-offs are associated with these parameters: the source block size directly affects latency, while the number of additional repair packets determines how much additional bandwidth is required to support the DF Raptor-protected stream.

Latency and Bandwidth Expansion

In general, each handset must buffer the first block of received source and repair packets in order to decode and obtain the first source block of the stream. If, for example, N seconds of the original data stream comprise each source block, then, in order to recover any lost source packets and avoid any interruptions to the stream presented for playback, the delay between receiving the first packet and the start of playback at the receiver is equal to the N seconds of buffering initially required at the decoder plus the negligible processing time required by the DF Raptor decoding algorithm.

The presence of latency in a 1-way MBMS broadcast can potentially affect the user experience by perceptibly delaying the initial playback of a stream. With MBMS, though, there are various techniques that can hide the effects of latency: for example, a stream that is scheduled to begin playback at a specific time can actually begin transmission seconds earlier so that the latency is not apparent to the user, or multiple streams can be bundled and delivered as a single stream so that switching between streams does not incur any latency. Even with such techniques, however, users will still experience a perceptible delay if an MBMS streaming service allows users to proactively join midstream. Minimizing latency by reducing the source block size is thus desirable.

At the same time, however, large source block sizes imply that the actual number of random lost packets over each block will more closely approach its average value and that any burst errors will affect a smaller percentage of data than for a small source block. On the one hand, then, large source blocks can potentially increase the perceived latency; on the other hand, large source blocks improve the ability of DF Raptor to protect against packet loss.

A similar performance trade-off exists in setting the number of additional repair packets produced by the encoding process for each source block. The overhead introduced by the repair packets directly translates into bandwidth expansion: to maintain the timing of the original stream, the source and repair packets associated with a source block must be transmitted over the same duration as the original data, requiring a faster data rate than the original stream in order to accommodate the additional repair packets. On the one hand, more repair packets imply the ability to tolerate more packet loss; on the other hand, more repair packets require more bandwidth in the form of higher data rates.

Representative Trade-Offs

The trade-offs involved in using DF Raptor for MBMS streaming services can be seen another way. In general, the quality of the delivered multimedia stream as protected by DF Raptor can be quantified by the probability that DF Raptor cannot recover an entire source block. If the source block size used to encode successive blocks of the original stream corresponds to, say, n seconds and the probability of decoding failure is, say, P, then the expected time between decoding failures will be n/P seconds. As long as the probability of decoding failure for DF

Raptor is below a desired threshold, then visible or audible artifacts in the playback of the video or audio stream will be few and far between. Even then, such decoding failures may not result in degraded quality as experienced by the end user – the available source data and the decoding algorithm employed by the multimedia stream may successfully mask the fact that the complete source block was not recovered. With quality assessed in these terms, however, the trade-offs among loss protection performance, latency, and bandwidth expansion then become:

- Given an anticipated packet loss rate, what DF Raptor source block size and what ratio of repair packets to source packets (overhead) are necessary to achieve a desired level of quality?
- And, assuming that the combined stream of source and repair packets fully utilizes the data rate allowed by the MBMS bearer channel, what does this required level of overhead imply in terms of the media data rate of the original stream?

To illustrate how these questions can be answered, Figure 4 shows simulation results for UMTS Terrestrial Radio Access Network bearer rates of 256, 128, and 64 kbps where DF Raptor has been used to protect source block sizes corresponding to either 5 or 20 seconds of the original stream. The curves show the maximum media data rates that can be supported as a function of the channel's block error rate so that DF Raptor, on average, fails to fully recover a source block once an hour. While not all such decoding failures necessarily lead to a perceptible artifact when the media is played back, the condition that the mean time between decoding failures is one hour (for these examples) provides a quantitative figure of merit for the stream quality as delivered by MBMS using DF Raptor.

Figure 4 assumes that DF Raptor is configured to generate repair packets so that the available bearer channel data rate is fully utilized. The maximum media rate shown in the figure thus corresponds to the minimum amount of overhead and bandwidth expansion necessary to provide the desired quality of one irrecoverable source block (decoding failure) per hour. If the media rate is actually less than the maximum shown here but DF Raptor is still used to generate repair packets expanding the original stream bandwidth up to the available bearer rate, then the mean time between decoder failures would be longer than one hour and the delivered quality would be accordingly better.

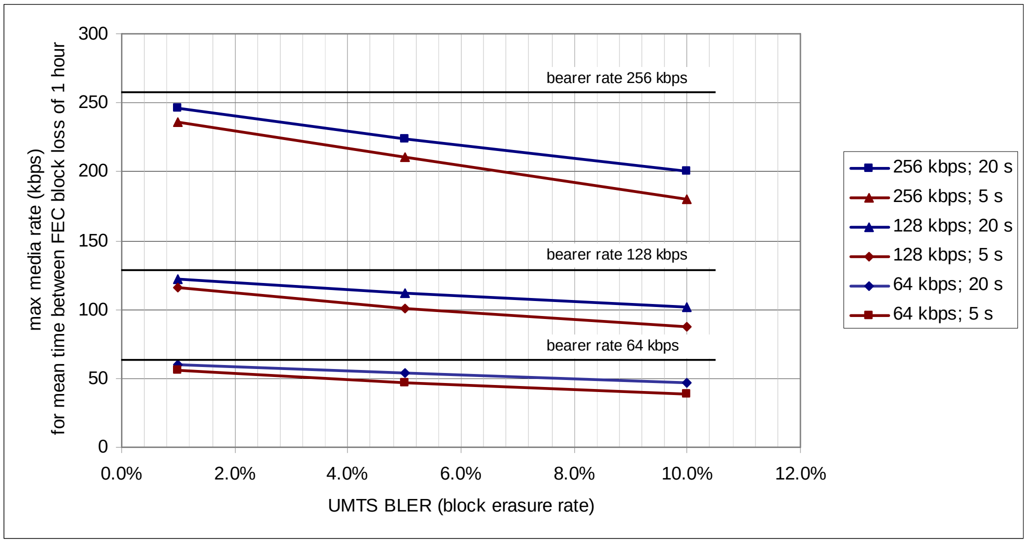

Figure 4 Maximum MBMS streaming media rates supported by DF Raptor for different UMTS bearer rates and channel conditions

The curves in Figure 4 also demonstrate the effects of different source block sizes -- as can be seen, the larger the source block is (i.e., corresponding to 20 seconds rather than 5 seconds), then the greater the maximum possible media rate because less overhead is required by DF Raptor to maintain a mean time of one hour between decoding failures.

Tradeoffs in the use of DF Raptor such as the ones illustrated here must be considered in order to protect MBMS streaming services optimally. Typically, an operator will choose a desired operating point in terms of the maximum block error rate to be protected, the desired level of quality, and the tolerable latency. When DF Raptor is used to protect MBMS streaming services in this way, the effects of packet loss can be minimized independent of the network characteristics, the stream characteristics, or the specific sources of packet loss, providing network operators with valuable flexibility in the engineering and operation of their networks but not compromising service quality.

Packet Loss Protection for Download Services

MBMS download services are based on the FLUTE (File Delivery over Unidirectional Transport) protocol as specified by IETF RFC3926, carried over UDP/IP.

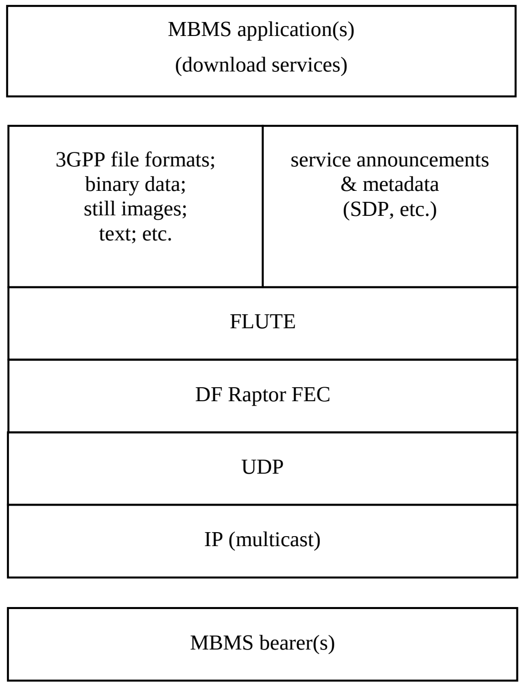

Figure 5 The MBMS protocol stack for download services

FLUTE is used to deliver both MBMS download services and MBMS service announcements via 1-way point-to-multipoint transmissions. At the sender, the MBMS server takes one or more data files (video clips, still images, software executables, etc.) and partitions them into data elements that can be constructed into FLUTE packet payloads. As part of the FLUTE transmission, the sender communicates the total length of the file(s), the length of a data element, and the overall structure of how the data is to be transmitted, thereby allowing each receiver to manage the reception and reconstruction of the data.

For 1-way transmissions, however, FLUTE does not in itself provide data reliability. MBMS tightly integrates DF Raptor into the FLUTE protocol in order to provide packet-level protection of the data so that any lost packets can be independently recovered by each receiver without requiring retransmission. At the MBMS server, each data file is considered to be one or more source blocks, with each source block encoded by DF Raptor and transmitted using FLUTE as source and repair packets. At each receiver, the received FLUTE packets are decoded by DF

Raptor. Again, it does not matter which of the source and repair packets have been lost – as long as enough data has been successfully received, then the original block can be fully recovered.

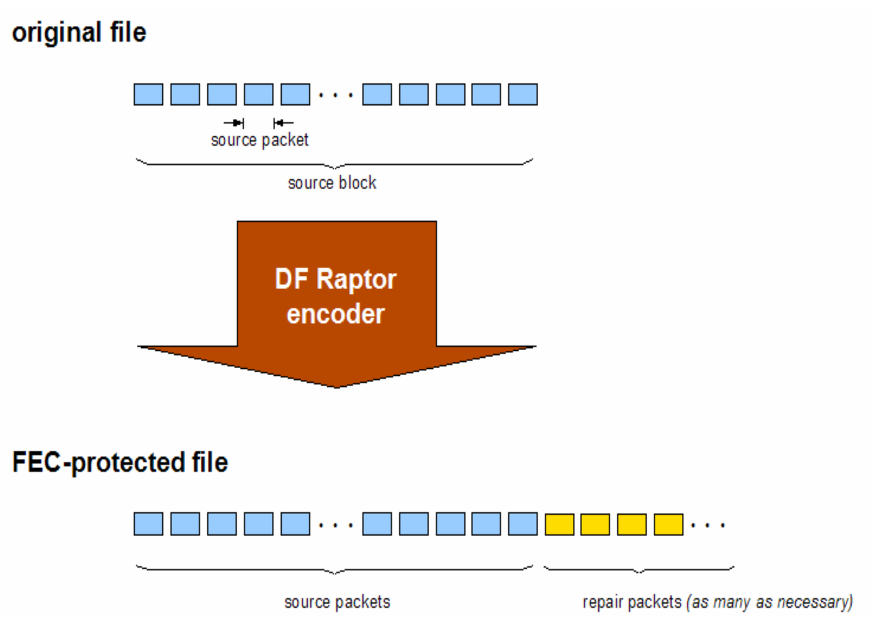

Figure 6 DF Raptor encoding of files

FEC Overhead

The number of additional repair packets generated by DF Raptor is a variable that can be set at the MBMS server as desired. A small number of repair packets may not allow most MBMS receivers to fully recover the source block, while a large number of repair packets will consume additional transmission time that could be used for the support of other file downloads or services. DF Raptor provides the flexibility to optimize file download services for efficiency by generating and sending just enough repair packets, transmitting additional repair packets as needed by individual receivers using post-delivery file repair.

Figure 7 illustrates the number of additional packets needed to deliver files of representative sizes over a UMTS channel as a function of the physical layer block error rate. The FEC overhead is here defined as the ratio of repair packets to source packets, assuming the PDU (Protocol Data Unit) size corresponding to UMTS Terrestrial Radio Access Network bearer rates of 128 and 256 kbps. The minimum FEC overhead required to realize a 99% probability of complete file recovery by a receiver is then provided as a percentage – if more repair packets are available for a given block error rate, then the probability of success will be higher.

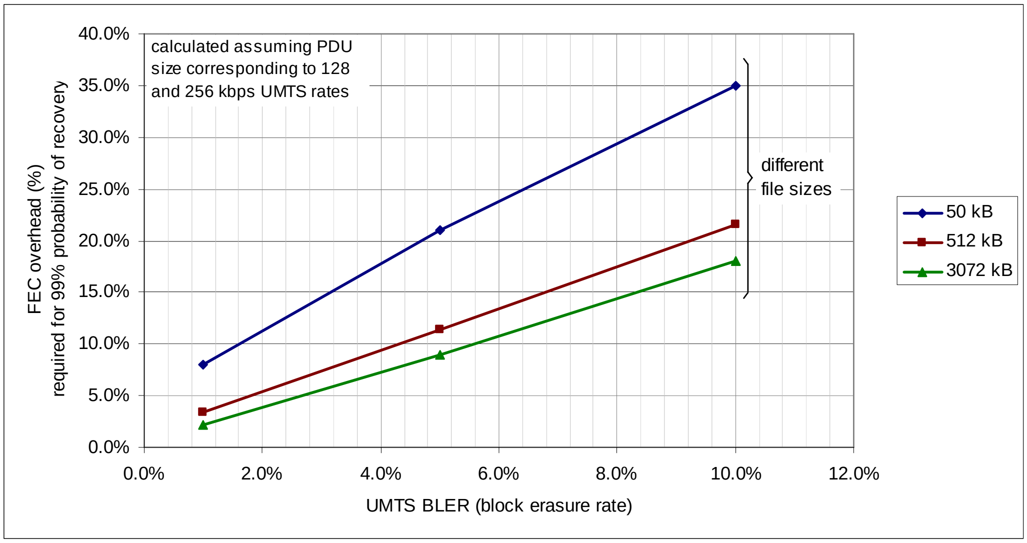

Figure 7 DF Raptor FEC overhead required for 99% decoding success as a function of different file sizes and UMTS channel conditions

Note that in Figure 7 small data files require higher levels of FEC overhead than large data files for a given block error rate. This performance reflects how randomly distributed block errors as simulated here are more likely to vary from the average rate when the source block size is small.

The results of Figure 7 indicate the need for relatively higher levels of FEC overhead in the presence of higher block error rates. The provision of MBMS services, however, involves a continual trade-off with voice and other services operating at each network base station as channel capacity and power levels are allocated among all the simultaneous uses. In certain cases, very high levels of FEC overhead may be warranted if it allows MBMS download services to be operated more efficiently from an overall system perspective (for example, by using reduced transmission power levels). Such flexibility is one of the key advantages of DF Raptor's use in MBMS.

Post-Delivery File Repair

Unlike streaming services where occasional data errors may be tolerated by the user, download services require that each transmitted file be obtained error-free and in its entirety at each receiver – a partially downloaded file typically is of no value to the user. If the FEC overhead provided by the MBMS server is sufficient and each receiver has obtained enough source and repair packets to fully recover the original data file, then no further action is required. Certain receivers, however, may not have successfully obtained enough packets by the end of the file delivery session, in which case they may each employ a point-to-point connection over TCP/IP to request a file repair session from the file repair server or if there are many such receivers then a point-to-multipoint repair session may be scheduled to deliver additional repair packets to the receivers via a new MBMS download service session.

Because *any* source or repair packets that have been successfully received can be used by DF Raptor to recover the original file, DF Raptor allows the number of needed additional packets to be minimized. Specific missing source packets do not need to be transmitted; rather, only additional repair packets need to be sent so that a receiver has the necessary number of total packets to fully recover the data file using DF Raptor. These repair packets can be used by all the receivers that did not fully recover the original data file. As a result, depending on the number of file repair session requests and on the available network resources, the additional packets needed for file recovery can be delivered to each receiver individually via TCP/IP or to the receivers as a group via a new MBMS download service session in order to conserve network resources.

As with streaming services, a network operator will typically choose a desired operating point in terms of the maximum block error rate to be protected. As long as the actual block error rates experienced by receivers is less than this operating point, DF Raptor will ensure the flawless delivery of download services. For those receivers that experience higher block error rates, the MBMS post-delivery file repair service provides the means to complete their download service efficiently. DF Raptor's ability to generate additional repair packets on demand without prior knowledge of the total number of repair packets needed is what makes this approach possible and effective.

Conclusion

MBMS promises to deliver a rich array of multimedia content to 3G subscribers. By employing DF Raptor to protect against packet loss caused by channel conditions, network congestion, or intermittent receiver outages, MBMS is able to deliver its streaming and download services with maximum quality and efficiency.