5775 Morehouse Drive
San Diego, CA 92121
(858) 587-1121

Benefits of DF Raptor®
for Mobile File Broadcast

Introduction

Operators continue to search for more and better ways to provide desirable services to consumers, who crave high-quality, bandwidth-intensive content when they want it. Mobile file broadcasting enables this by allowing content to be automatically delivered to mobile devices and accessed at the subscriber's convenience, unlike streaming media where users are limited to tuning in at scheduled times. The cost and convenience of file broadcast have made it a compelling business proposition as service providers look to expand services and increase ARPU (Average Revenue Per User) by developing more content and services through efficient use of spectrum and monetizing the network during low usage hours.

File broadcasting offers an ideal mix of economy and convenience for sending the same content to multiple users. Broadcast is economically advantageous because operators can pay for infrastructure once and then amortize the cost over a growing subscriber base—as the subscriber base grows, the cost per subscriber falls. This is unlike most mobile unicast architectures where costs scale linearly with the number of subscribers.

As significant as the benefits of file broadcasting are, there are some major obstacles that must be overcome to ensure a successful deployment.

Mobile File Broadcasting is challenging for two key reasons.

1. Subscribers are mobile and disparate, providing for highly dynamic (often challenging) reception conditions.
2. The worst-case reception conditions must be supported to ensure the maximum number of subscribers receive all of the content.

These challenges oftentimes lead to errors and packet loss, which can reduce and even eliminate the usefulness of the content, so ***the key challenge to a successful mobile file broadcast deployment is eliminating loss at the client.***

From an operator's perspective, solutions are required to achieve the highest reception reliability with the shortest broadcast duration and/or lowest power consumption so as to provide the highest service levels while maximizing the amount of content that can be delivered.

The following sections list some of the primary solutions to recover from errors and losses in the mobile environment. Each has various advantages and disadvantages listed.

Physical Layer FEC

Physical layer forward error correction (FEC) is an essential component in almost every modern communication system. Physical layer FEC is done at the hardware layer, typically as part of the radio interface (e.g., EDGE/GPRS and WCDMA), to correct smaller disturbances at the bit level. Example issues include short time noise, fading and interference losses. It is widely deployed in many technologies and applications.

Physical layer FEC is accomplished by sending some amount of repair data for some amount of original data using either convolutional codes, Turbo codes, or other codes. The original data is broken into small blocks of bits (100's or 1000's), so its protection is limited in duration. For data transmission all content is protected equally, so any headers, flags, identifiers, synchronization cues, and padding are also protected. If burst losses are expected, interleaving is often used to provide additional time diversity by

spreading any errors across multiple blocks (at the expense of added latency and more hardware at the receiver), thus reducing the actual error experienced by any one block of bits, but the protection duration is still rather small (a few tens of milliseconds or less).
If the quality of the received signal over this protection period is below the correction capabilities of the physical layer FEC, the block cannot be recovered. Any IP or other packets included in this block will likewise be unrecoverable and eventually discarded resulting in packet loss at the receiver. In mobile broadcast environments, reception conditions can be quite poor, so it is necessary to use very low coding rates for physical layer FEC, on the order of 1/3 or 1/5, so most of the bits sent are for repair equating to repair overheads of 200% and 400%, respectively. Such high overheads significantly limit the available bit rate for content and yet still cannot guarantee error-free reception in poor conditions.

To provide the most effective protection against channel noise, the blocks of data protected by physical layer FEC should be as large as possible. This is because the channel noise is less variable averaged over longer time intervals than over shorter time intervals, and thus the relative amount of protection needed per block to provide relatively low block loss rates can be smaller for larger blocks than for smaller blocks. However, the maximum physical layer FEC protection provided against error is limited by available memory-on-chip, power, and decoding complexity, so for practical reasons blocks tend to be very small. For example, block interleaving can be used to provide some benefits, but block interleaving is especially constrained by the costs associated with the necessary memory-on-chip, and these costs increase rapidly with the higher bit rates associated with emerging services and technology, so block interleaving also tends to be over small time intervals. These are primary reasons why physical layer FEC is often combined with another protection mechanism to resolve even mild loss scenarios.

Furthermore, the hardware integration of physical layer FEC requires dedicated silicon, which impacts time to development and limits the flexibility of the solution rather significantly.

Notwithstanding, physical layer FEC has proven to be very useful in correcting very short noise and disturbances on the order of a millisecond, as demonstrated by its widespread use. But as physical layer FEC is limited in its ability to correct longer (multi-packet or multi-slice) loss events, it must be used in conjunction with other protection schemes.

Physical layer FEC has the following advantages:

- Pervasive—widely implemented
- Very efficient—high ratio of repair data sent vs. data that can be repaired
- Can be (and often is) easily implemented and combined with other protection schemes

Physical layer FEC has the following disadvantages:

- Must be combined with another protection method for higher error rates or longer outages
- Limited to small block sizes, so can only fix very short disturbances
- Implemented in hardware requiring longer development times and preventing flexibility
- Low coding rates typically used, i.e., lots of repair data sent relative to original content
- Both sender and receiver must support the same codes used

Link Layer Protection (Reed-Solomon Erasure Codes/MPE-FEC)

Reed-Solomon (RS) codes provide higher layer protection against larger and longer losses. By operating at the link layer, RS can provide protection for multi-packet or multi-slice loss events that physical layer FEC cannot, and it is used in conjunction with physical layer FEC to produce a more efficient system overall. The primary disadvantage of practical RS codes is the maximum effective block size, thereby producing low efficiencies for large files or high loss rates. As with physical layer FEC, larger block sizes provide greater efficiencies by allowing loss events to be averaged over longer time periods.

One of the more popular implementations of RS codes is MPE-FEC (Multi-Protocol Encapsulation FEC). As with most implementations of RS codes, MPE-FEC has pre-defined code parameters. MPE-FEC operates on frames, with each frame arranged as a 2-dimensional matrix of 255 columns and a variable number of rows (maximum 1024 rows). 64 of the columns are dedicated for repair data leaving a maximum of 191 columns for content, indicated as a code ratio of RS (255,191). Thus each frame can protect against up to 64 lost bytes in any one row of 191 bytes.

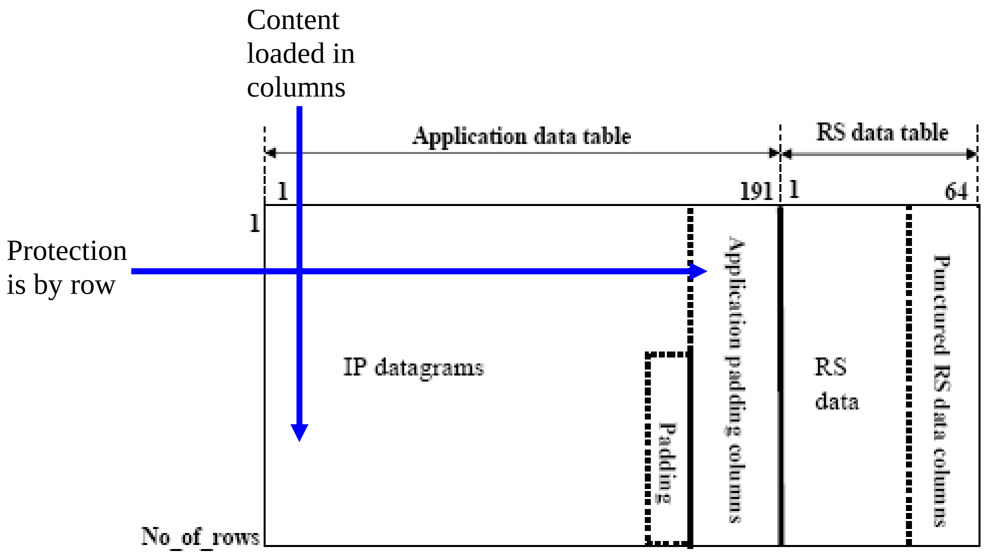

Figure 1. MPE-FEC frame structure [1].

Since each matrix position is 1 byte, this equates to a maximum size of just under 200kB (200 kilobytes) for the block, which means segmentation overhead is required for files larger than 200kB. In addition to these overheads, there is some added overhead for MPE-FEC administrative use (identifiers, flags, etc.).

It is possible to reduce (through “puncturing”) the amount of correction by sending fewer repair columns with the corresponding reduction in protection, thus reducing both the protection provided and the overhead required. An increase in protection is possible through shortening, whereby data bytes are padded with known good bytes and not sent thus changing the coding ratio from RS (255,191) to RS (214,151) or lower, similarly increasing both the protection provided and the overhead required. This has the downside of reducing the block size, which then means more segmentation and smaller averaging intervals.

Being a systematic code, MPE-FEC supports backwards-compatibility, which means receivers that do not support MPE-FEC can still receive the content when MPE-FEC is in use. It can also be enabled as needed for individual streams (e.g., audio or video) as part of a larger transport stream.

MPE-FEC is typically implemented in hardware, since Reed-Solomon codes are computationally very expensive. This produces similar disadvantages as physical layer correction in that the development times are extended, algorithms are fixed, and most changes require new hardware.

Due to limited memory, an MPE-FEC frame cannot span more than one time slice burst (for DVB-H). This makes it inefficient as time diversity cannot really be exploited. It is also not very efficient relative to other alternatives in that it requires much more repair data than it can correct, e.g., a 25% parity overhead is required to protect against a packet loss rate (PLR) of about 10%. This is due to the 2-dimensional nature of the MPE-FEC frame and the randomness of noise in that frame (i.e., more noise increases the likelihood of having multiple losses in any one frame’s row, refer to Figure 1). Since MPE-

FEC frames are limited to a maximum of 200kB, segmentation is required resulting in additional overhead. In scenarios where the loss rate within a single time slice burst is a bit higher (typically >10%), retransmission and/or carouseling are required when using MPE-FEC, since the probability of being unable to recover all of a frame's rows increases. Failure to recover any one row is likely to cause the entire frame to be lost. So while the overall loss rate might be much lower, time slices with higher loss rates cannot be dealt with MPE-FEC.

RS codes (MPE-FEC) have the following advantages:

- Configurable protection amount through puncturing or shortening
- Can protect against longer outages
- Backwards-compatible—receivers can ignore RS repair data

RS codes (MPE-FEC) have the following disadvantages:

- Small frame sizes (<200kB for MPE-FEC)—inefficient for large files or even mild outages
- 2-dimensional protection scheme is vulnerable to random losses
- Requires hardware to implement due to computational complexity
- Must be combined with data repetition (carousel) for higher loss rates
- Must be used with physical layer FEC and checksums (for error detection)

Carousel

File carouseling involves sending the same content over and over again to ensure all receivers get all of the data packets. It uses blind retransmission as a simple means to solve a difficult problem, i.e., knowing which receivers are missing what content without getting feedback from the receivers, but it is very inefficient. Carouseling is used in cases where the content is not being consumed soon, e.g., overnight file push or software updates.

Typically, a file is split into multiple segments (usually packets) and sent in its entirety. Once a cycle is complete, some delay may be added to send other content, and then the file is sent again in its entirety. This process is repeated until there is a reasonable statistical assurance that all receivers have received the content.

While common, carousels are highly inefficient. If a receiver misses the first packet of data, it must wait until the next file delivery to get an opportunity to receive that one packet. If it misses that packet again, it must again wait until the next file send. During this time, the receiver is forced to receive and process every packet being delivered, which wastes power and resources.

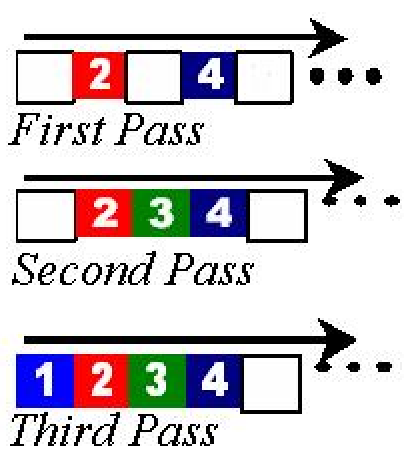

Figure 2. Demonstration of carousel.

Carouseling has the following advantages:

- Simple to implement and deploy
- Can send content as many times as necessary for all receivers to get it
- No limit to loss protection (can resend file forever)
- Backwards-compatible—receivers can ignore resent content

Carousel has the following disadvantages:

- Can require numerous repeats to ensure receivers get content
- Either send entire content or not at all—especially problematic for large files
- Content may never be received, e.g., in repetitive impulse noise conditions
- Receiver must receive all packets, even if only missing one
- Requires long on-times for faster battery drain

Retransmission

Retransmission is a standard means for unicast content delivery whereby the receiver tells the sender which pieces of content were received through acknowledgements and/or not received through negative acknowledgements and the sender or a repair server then resends any missed items. A protocol example would be TCP. Retransmission requires a return channel to a repair server and works best in cases where there are few receivers and the loss rate is low. As such, it tends to be a poor solution for large-scale applications where broadcast/multicast is desirable.

Retransmission requires that the receiver is able to signal to the sender that certain content must be resent. It also means the content cannot be consumed until this round-trip interaction can occur, thus adding latency and receiver buffering requirements that are dependent on the transmission method (many seconds for satellite). In addition, missed content is usually sent after having received a round of retransmission requests from all the receivers. As the number of receivers increases and/or as the channel loss conditions worsen, the fraction of content missed by at least one receiver becomes so large that effectively all the content must be sent multiple times, such that the delivery deteriorates into a carousel with multiple repeats of the content.

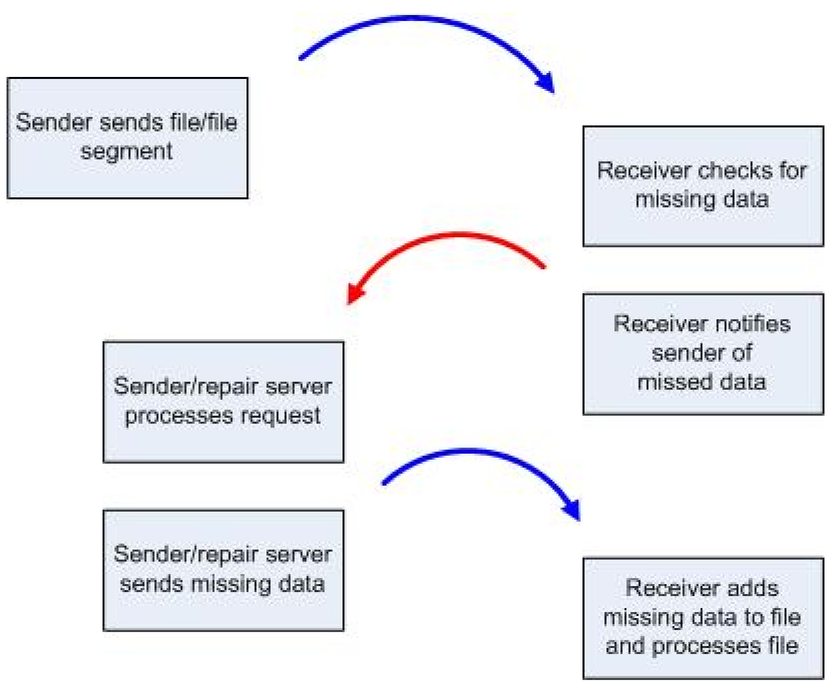

Figure 3. Retransmission negative acknowledgement process.

Retransmission can also place a significant load on the repair server (or sender, if also performing this function), as it must be able to process incoming retransmission requests. If retransmission requests are not sent to all receivers, each receiver missing content will send a retransmission request to the repair server even if it is the same content. In cases where there are numerous receivers and many or all of them do not receive the same content, the repair server can be overwhelmed by retransmission requests (feedback implosion). As such, this does not typically scale well in broadcast or multicast applications.

Retransmission has the following advantages:

- Only missed data is resent
- Backwards-compatible—only retransmission-enabled receivers will request missing content

Retransmission has the following disadvantages:

- Requires bi-directional transmission, e.g., uplink is needed
- Does not scale well when numerous receivers receive the same content
- Does not scale well when channel loss conditions worsen
- Requires sender or a repair server to buffer data packets

Application Layer Protection (AL-FEC/DF Raptor)

Application layer FEC intends to solve the problems not addressed by physical layer FEC in a manner superior to RS codes. As with RS codes, AL-FEC is combined with physical layer FEC to produce an overall superior solution. It is particularly relevant in situations where channel variations are long-term, i.e., the loss transients affect multiple packets or multiple time slices. AL-FEC works similarly to RS codes in that it can be used to offset the overhead required by physical layer correction resulting in a more efficient overall system. The specific implementation of AL-FEC by Digital Fountain using their DF Raptor™ solution also removes the need for carouseling and can be performed using software, so this solution provides added benefits over MPE-FEC and is addressed here. DF Raptor has been adopted in standards such as DVB-H IPDC and 3GPP MBMS for mobile and DVB-IPI for IPTV.

DF Raptor generates additional repair packets that can be used to reconstruct any missing source packets. Since the original source content is sent, the solution is also backwards-compatible. DF Raptor's notable difference with other techniques is that any repair packet can be substituted for any missing source packet, such that it is possible to reconstruct all of the source data solely from repair packets. Unlike systems that implement RS codes, DF Raptor performs optimally regardless of the loss pattern. Additionally, the number of repair packets that can be sent is virtually limitless, so there is no need to carousel—repair packets simply need be sent until enough are received to reconstruct the missing source content. The receiver will stop processing repair packets once enough are received.

DF Raptor also supports partitioning source blocks into sub-blocks, so large files can be segmented such that each segment has the same protection as an unsegmented file. This allows files of virtually any size to be sent as one source block to limit the amount of overhead required.

The associated overhead is two-fold: minor overhead for identifying blocks and packets and added overhead to process repair data (slightly more repair data must be received than data lost, on the order of 1% extra).

Since the solution is implemented in software, it can be reconfigured easily and distributed to receivers as a software upgrade allowing for simple deployment and changes. The small code size and mild CPU and memory requirements allow for integration into even the smallest devices.

AL-FEC (DF Raptor) has the following advantages:

- Recognized and adopted standard (DVB-H, 3GPP, IETF) that can be efficiently combined with file repair protocols
- Can protect against very high loss rates (99+%) and very long outages (seconds and longer), allowing for increased link margins of 3-5dB or more for increased tower reach
- Easy to configure—each file can have different protection parameters
- No carousel needed—any repair packet can reconstruct any source packet, resulting in reductions in file delivery times of 85% or more
- Efficient—only approx. 1% excess repair data needed for reliable reconstruction
- Software-only solution for easier integration and deployment
- Unlimited source block sizes to accommodate very large files (terabytes)
- Adaptable—FEC length can be set to the total file size for optimal time-averaging
- Backwards-compatible—only AL-FEC-enabled receivers process repair data

AL-FEC (DF Raptor) has the following disadvantages:

- Overkill for small file sizes (below 32kB)
- Must be used in combination with physical layer FEC and checksums (for error detection)

Benefits of AL-FEC (DF Raptor)

In comparing the relevant solutions to loss for mobile file broadcast, the following general conclusions about the DF Raptor solution can be made. DF Raptor offers

1. Highest efficiency mobile file broadcast delivery solution. DF technology (combined with physical layer FEC) optimizes existing networks, enabling operators to offer more content and thus increase revenue opportunities.
2. Improved coverage and range. Network operators can broadcast files over a broader geographic area and have a higher level of certainty of reliable reception.
3. Shorter broadcast duration. Operators can deliver more content over the same infrastructure. With any solution, whether DF or any other, operators have to broadcast the file for some amount longer than the duration of the file itself. DF technology reduces the broadcast duration.
4. Flexible correction. Different protection parameters can be set for each file sent so as to be adaptable to the loss conditions in effect or the priority of the file.
5. Reduced computational requirements. High efficiency code means network operators can potentially deploy lower cost devices and reduce power requirements. DF Raptor has been successfully integrated into a number of mobile handsets.
6. Proven, standards-based solution. Standard in both 3GPP MBMS, DVD-H IPDC, and IETF for file delivery, DF Raptor has been deployed by KDDI in Japan, licensed by XM, Sirius, and Honda in the US, will be deployed shortly in Europe, and is being developed for several others in Korea and elsewhere.

To further explain the reduction in broadcast duration, the following graph shows simulation results of the delivery time mapped against the required C/I, or carrier-to-interference ratio (CIR). Delivery time is indicated by time slices of several hundred milliseconds per DVB-H that roughly translate to physical time in cases where time slicing is not used. The various cases compare carousel (repetition), MPE-FEC, and/or DF Raptor (AL-FEC) with a 95% probability of reception for a 4MB file with low Doppler, e.g., a slow-walking person. Each case was run 10 times for reproducibility. Both Rayleigh and shadow fading were simulated. Details and additional results are available as part of the DVB-H IPDC specification [2]. Note that cases indicating 150 time slices indicate situations where the simulation likely did not terminate, i.e., delivery could not be 95% likely with that scenario.

The first case (No MPE-FEC + Raptor) is with just the DF Raptor AL-FEC solution. Note that it provides the best overall solution.

The next case (MPE-FEC 7/8 + Raptor) adds MPE-FEC at a rate of 7/8. It is interesting to note that adding MPE-FEC in this case <u>increases</u> the time required to receive the file. The extra time slices required are due to the extra overhead added by the MPE-FEC code.

The next case (MPE-FEC 3/4 + Carousel) uses carousel with MPE-FEC of rate 3/4. Note that with this protection scheme, lower C/I values provide no assurance of reception. The interesting note here is that the results are not significantly better than carousel alone.

The last case (No MPE-FEC + Carousel) shows carousel only. It is a case where simply resending the file does not produce acceptable results.

Required Broadcast Duration for MPE-FEC and Carousel vs. DF Raptor (large file)

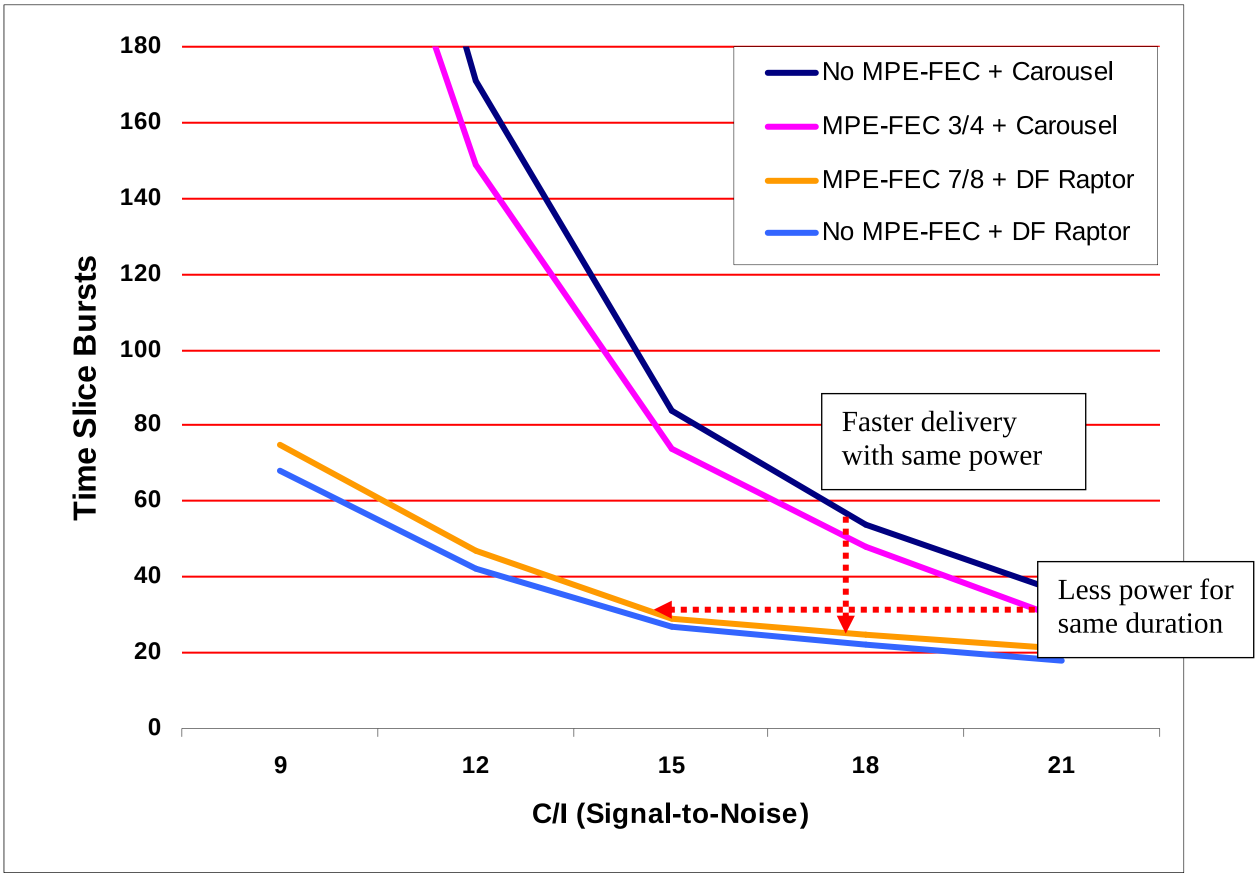

Figure 4. Broadcast duration for 4MB file with 1Hz Doppler and 95% probability of receipt.

These results clearly demonstrate the improvements possible with mobile file delivery by using the improved DF Raptor FEC scheme.

To translate these results to increased range, if the system were deployed today with MPE-FEC 3/4 and carousel with a given coverage area to provide a 21dB C/I at the receiver, a system with only DF Raptor could allow the coverage area to be increased to the 15dB C/I range resulting in a significant increase in range without changing any of the physical equipment or transmission parameters.

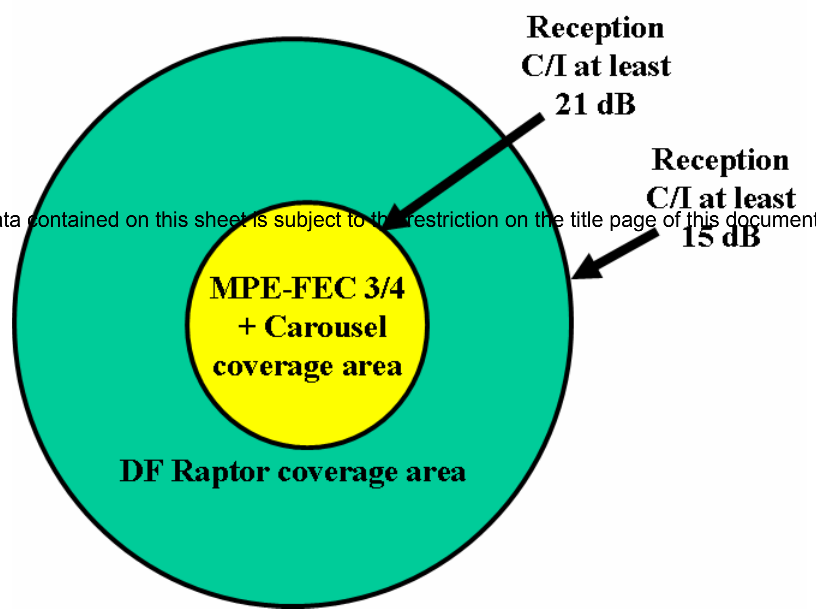

Figure 5. Demonstrated difference in coverage area with different protection schemes.

Similar results are achieved using a smaller file (128kB) with a higher Doppler (80Hz) simulating motion equivalent to traveling in a car. The MPE-FEC case quickly degrades to an undeliverable condition.

Required Broadcast Duration for MPE-FEC and Carousel vs. DF Raptor (small file)

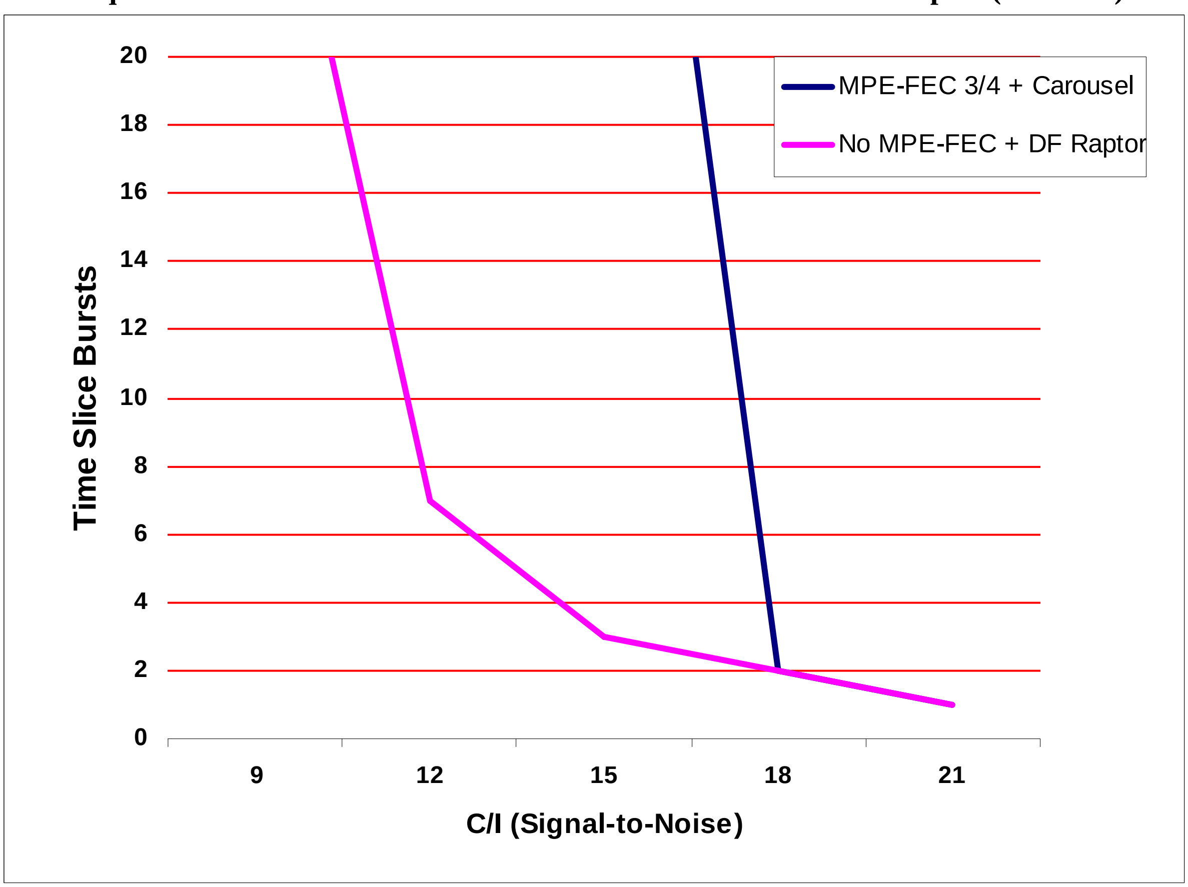

Figure 6. Broadcast duration for 128kB file with 80Hz Doppler and 95% probability of receipt.

Case Study: Tier 1 Mobile Operator in Asia

Digital Fountain technology was deployed last year by a tier 1 mobile operator in Asia (Operator A) to deliver a wide variety of mobile broadcast content. Broadcasting over BCMCS (Broadcast and Multicasting Services), Operator A offered new subscription services to their customers, including music, news, weather, sports, and entertainment. DF Raptor software was deployed across the network and on handsets. Since DF Raptor is application layer software, integration was quick and seamless.

Since much of the content Operator A was providing was time sensitive (current weather, scores, news, etc.), it was imperative that the content be delivered in its entirety in a very timely manner, despite the inherent challenges of mobile content delivery–bandwidth consumption, signal interference, mobility, handset proximity, and other factors. The most important of these challenges was their desire to reach as many potential customers using as little bandwidth as possible.

By deploying DF Raptor technology to deliver these services, Operator A was able to use **deliver files up to 7x faster, allowing them to deploy a vast array of incremental content and services. Further, the operator achieved a 3-5dB increase in link budget, allowing the operator to reach more subscribers with existing infrastructure.** Timely delivery was ensured using DF Raptor, so their customers were very happy with the service, and subscriptions for these services continue to climb.

Summary

Operators need to ensure files are delivered with the highest reception reliability and the shortest broadcast duration to maximize the amount of content that can be reliably delivered, ultimately leading to increased ARPU. A key mechanism to achieve these aims is error correction, and a better or stronger error correction scheme allows for

- higher data rates and/or faster file delivery for more ARPU and
- lower transmitter power, smaller receivers, and/or wider coverage area for lower costs.

As indicated above, the DF Raptor FEC solution, combined with some physical layer FEC, provides the best, strongest protection scheme with substantial additional improvements over the alternatives as summarized below.

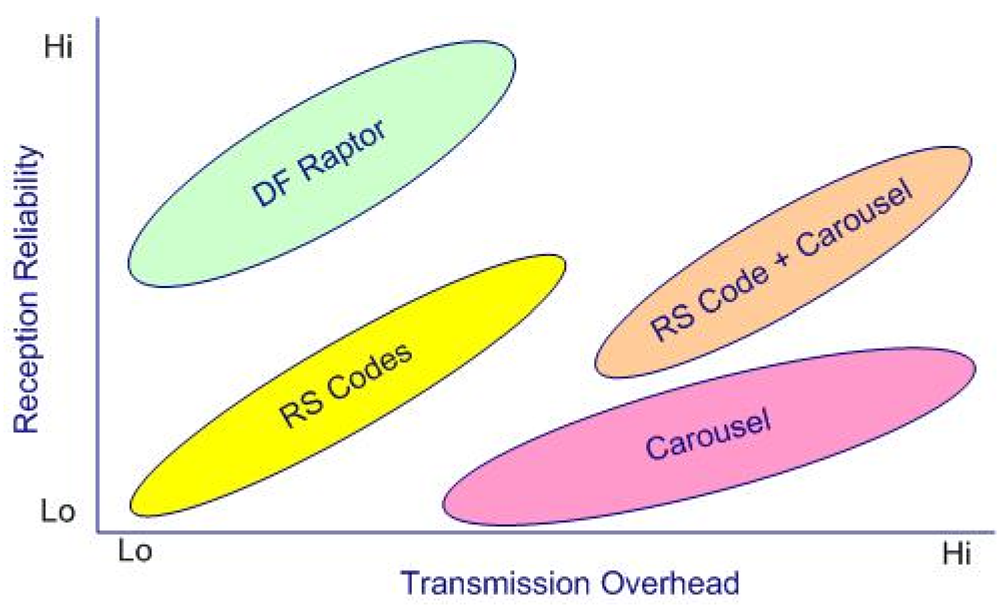

Figure 7. Comparison of various FEC alternatives.

References

[1] ETSI TR 102 377 v1.1.1: DVB-H Implementation Guidelines
[2] Draft ETSI TR: IP Datacast over DVB-H: CDP Implementation Guidelines