

OVERVIEW OF ADAPTIVE TDMA IN AN IDIRECT EVOLUTION NETWORK

Introduction

The ST Engineering iDirect Adaptive Time Division Multiple Access (TDMA) delivers higher spectral efficiency and greater network versatility by optimally changing the return channel configuration to match the conditions of the over-the-air links and the constraints of the terminals. Network operators using this feature will be able to boost their network throughput, increase network availability, reduce terminal costs or combine these benefits to fit their business needs. With Adaptive TDMA, each Very Small Aperture Terminal (VSAT) can be assigned a different modulation, coding, symbol rate and spread factor, on a burst-by-burst basis to optimize the return channel. The network is then better able to accommodate:

1. Static differences among remotes like block up converter (BUC) size and antenna size.
2. Dynamic differences among remotes like the amount of rain fade over the hub and the amount of rain fade over each remote.
3. Differences like Gain over Temperature (G/T) and Saturated Flux Density (SFD), which are static for fixed VSAT, but dynamic for mobile VSAT.

Heterogeneous Inroute Group

Adaptive TDMA introduces the flexibility to change the inroutes within an inroute group, improving the carrier options available to a remote within that inroute group. In existing ST Engineering iDirect systems, inroute groups are groups of identical TDMA inroutes on which remotes are allowed to transmit. With Adaptive TDMA, the inroutes in an inroute group are no longer required to be identical. In fact, inroutes with different symbol rates, modulations, Forward Error Correction (FEC) coding rates, spread or unspread, as well as spread factors can now co-exist in the same inroute group. This new flexibility affords great latitude for a properly designed inroute group to provide inbound connectivity to VSATs with different terminal capabilities and under different link conditions. Figure 1 shows one example of one such heterogeneous inroute group.




Figure 1 – Heterogeneous inroute group

Short-term Adaptivity - Frequency Hopping and Uplink Power Control

Each remote can frequency hop from one burst to the next among all the inroutes within its frequency hopping group. A remote can therefore frequency hop to change its MODCOD, symbol rate and even spread factor on a burst-by-burst basis as necessary to optimize upstream port speed and bandwidth efficiency for varying link conditions. The ST Engineering iDirect hub allocates upstream bursts among remotes based on the demand for capacity and the Quality of Service (QoS) policies.

The upstream burst allocation algorithm maximizes aggregate return channel capacity by continuously measuring the Signal to Noise Ratio (SNR) of upstream bursts coming from all remotes. Using the latest SNR value and the symbol rate of the previous carrier, the instantaneous Carrier to Noise density value (C/N_0) is computed for each remote that determines into which subset of inroutes a particular remote is capable of bursting. This value is used by the hub to determine the “nominal carrier” for each remote, which is the most efficient carrier within each remote’s constraints, and the subset of inroutes that each remote can access.

Each remote uses uplink power control as much as possible to meet the SNR threshold of the carrier that the remote is accessing at any given time. When a remote, due to rain fade or some other reason, no longer has sufficient BUC power to achieve the SNR threshold of the current carrier, the hub will command the remote to hop down to another carrier with a lower symbol rate and/or a less spectrally efficient MODCOD with a lower SNR threshold. When a remote has more than enough BUC power to achieve the SNR threshold of the current carrier, the hub may command the remote to hop up to another carrier with a higher symbol rate and/or a more spectrally efficient MODCOD with a higher SNR threshold.

To illustrate how different rain fade conditions impact the subset of inroutes a remote can access, consider the example illustrated in Figure 2. An Operator manages a network for a large Enterprise customer with multiple remote offices. Throughput requirements at each of the remote offices are reasonably high, and the customer’s remote offices are geographically dispersed. Here, we consider a single remote office in that network suffering from increasingly adverse weather conditions.

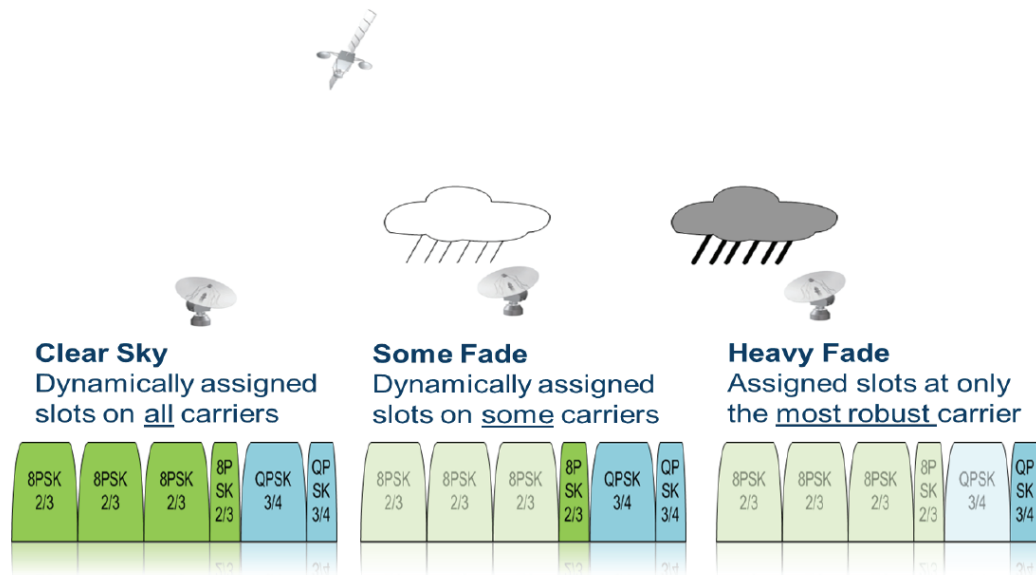


Fig. 2 – Rain fade condition affecting the subset of inroutes available to a remote

In ideal clear sky conditions, the remote in question is eligible to burst on all of the available carriers, from the largest and least robust 8PSK 2/3 inroute down to the smallest and most robust QPSK 3/4 inroute.

As the remote office starts to experience some fade, uplink power control will increase the transmit power in an attempt to keep it on the most efficient and highest bandwidth inroute. However, at a certain point, the remote will not be able to increase its transmit power any further; the hub will then remove the largest 8PSK 2/3 inroutes from those on which this remote will be dynamically assigned slots. The smaller 8PSK 2/3 inroute can still be an option, as its SNR threshold will be lower than the larger 8PSK 2/3 inroutes due to the size of the inroute.

As conditions at the remote office continue to worsen, and the site is in heavy fade, it is only on the most robust inroute in this network, the smallest QPSK 3/4 inroute, on which the remote could be assigned slots.

It is important to note that throughout this example, the composition of inroutes did not change: the system was able to accommodate fade at a single site without changing the MODCODs of the inroutes. However, in cases of widespread fade coding changes to the inroutes may be required as explained in the next section.

Medium-term Adaptivity – Modulation and Coding Changes to Inroutes

As the change of conditions of the remote population becomes more widespread, such as wider rain fade or mobile VSATs moving to different beam regions, the demand for some inroutes may eventually outstrip their availability under the current inroute group design. In iDX 3.2, the Adaptive TDMA Return Channel can change in real time the modulation and coding schemes (MODCODs) of the inroutes in an inroute group.

In any inroute group of Adaptive TDMA inroutes, there are always some upstream inroutes that are more desirable than others due to differences in symbol rate and MODCOD. Inroutes with higher symbol rates are more desirable because they offer faster upstream port speeds for the remotes that can access them. Similarly, inroutes with more spectrally efficient MODCOD are also more desirable because they maximize aggregate return channel capacity. However, more desirable inroutes require more BUC power, bigger transmit antennas, greater G/T, and less rain fade to access.

The ST Engineering iDirect system continuously quantifies the changing amount of traffic that must go upstream on less desirable inroutes. When there is a shortage of less desirable carrier capacity, the hub lowers the MODCOD of some upstream inroutes to create additional upstream capacity that more disadvantaged remotes can access. Conversely, when there is a large surplus of less desirable carrier capacity, ST Engineering iDirect increases the MODCOD of some upstream inroutes to create additional upstream capacity that is more spectrally efficient.

In an ST Engineering iDirect VSAT system, several different inroute group compositions of carrier MODCODs can be defined for any inroute group. An inroute group composition defines the MODCOD of each carrier in an inroute group. Every few tens of seconds, the exact interval configurable by the Network Operator based on the rate of change of conditions, the hub assesses the amount of contention among remotes and the aggregate throughput available within each inroute group composition, then chooses the composition deemed to best satisfy the network's current requirements.

Consider the Enterprise customer from the use case in the previous section. In addition to requiring reasonably high throughput at each site, the customer would also like to reduce any interruptions in connectivity at its offices to a minimum. The Network Operator, after considering service requirements and completing a link budget, might propose a configuration consisting of three inroute group compositions.

The first inroute group composition consists of all 8PSK 2/3 inroutes, which most likely corresponds to the composition used in clear sky condition. The two smaller inroutes can provide some rain fade protection to a small number of faded remotes. The middle inroute group composition is necessary when more rain fade protection inroutes are needed. If the hub is under fade, all remotes will suffer. In order to keep as many remotes in the network as possible, all the MODCODs are lowered further as shown in the third inroute group composition.

The Network Operator now is able to deliver the higher service levels demanded by the Enterprise customer, but also offer the protections required when adverse conditions affect parts of or the entire network.

condition. The two smaller inroutes can provide some rain fade protection to a small number of faded remotes. The middle inroute group composition is necessary when more rain fade protection inroutes are needed. If the hub is under fade, all remotes will suffer. In order to keep as many remotes in the network as possible, all the MODCODs are lowered further as shown in the third inroute group composition.




Figure 3 – Examples of inroute group compositions

Long-term Adaptivity - Inroute Group Composition Design

As illustrated in the previous section, the design of these inroute group compositions has direct impact on the total throughput, per remote data rates and the availability of the network. ST Engineering iDirect will provide enhancements to its existing network design tool to assist VSAT network designers in properly sizing the return inroutes in the inroute group and designing the inroute group compositions.

With the constant changes in most VSAT networks, either due to new service offerings or additional remotes in the networks, the design of inroute groups and their associated inroute group compositions should be reviewed and optimized from time to time. Using the new iMonitor tools, a network engineer can quickly analyze the composition usage distribution, now and over time, C/No distribution for each carrier, time plans over time as well as physical layer measurements of the remotes.

The following diagram depicts Adaptive TDMA

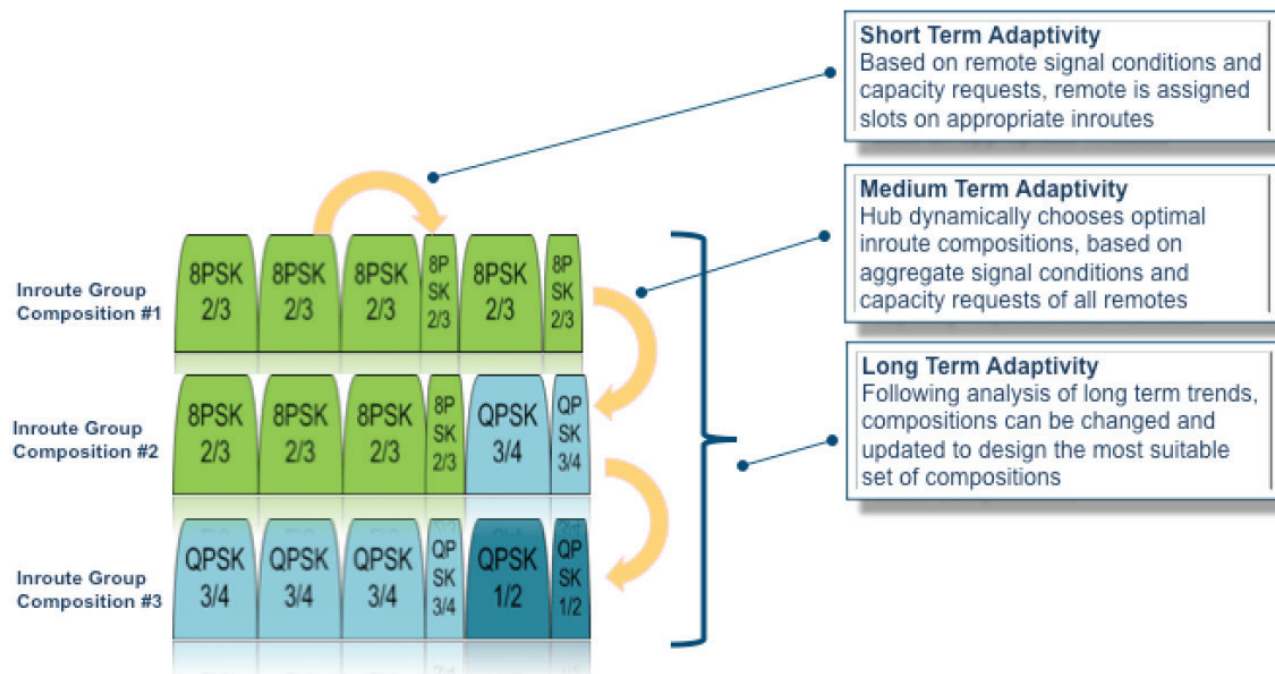


Figure 4 – Adaptive TDMA

Performance Advantage of Adaptive TDMA

Compared to traditional non-adaptive inroute group design, Adaptive TDMA will enable Evolution or Velocity networks to operate at more spectrally efficient MODCODs while providing far better protection for adverse link conditions.

Let's consider an existing network with a non-adaptive inroute group that uses five inroutes, each one at 1Mps. Let's also assume that the inroutes were designed to use QPSK 3/4 MODCOD and had 4dB in rain fade margin. Upgrading to iDX 3.2 with Adaptive TDMA will enable these inroutes to operate up to 8PSK 2/3 by not needing to maintain the sizable fade margin. Rain fade protection is, in turn, provided by having different carrier sizes in the same inroute group and the ability to change the MODCODs on the fly. As a result, under clear sky, both the total throughput and the per remote peak rate increase significantly while the fade protection also goes up. When fade becomes more widespread and two of the inroutes need to change to QPSK, the total throughput is still 16% higher and the fade margin increases to 7dB or 75% improvement. If fade occurs at the hub, all the MODCODs will be lowered to provide total fade margin of more than 12dB or 200% improvement. Total throughput will take a 10% hit but all the remotes will be able to stay in the network.

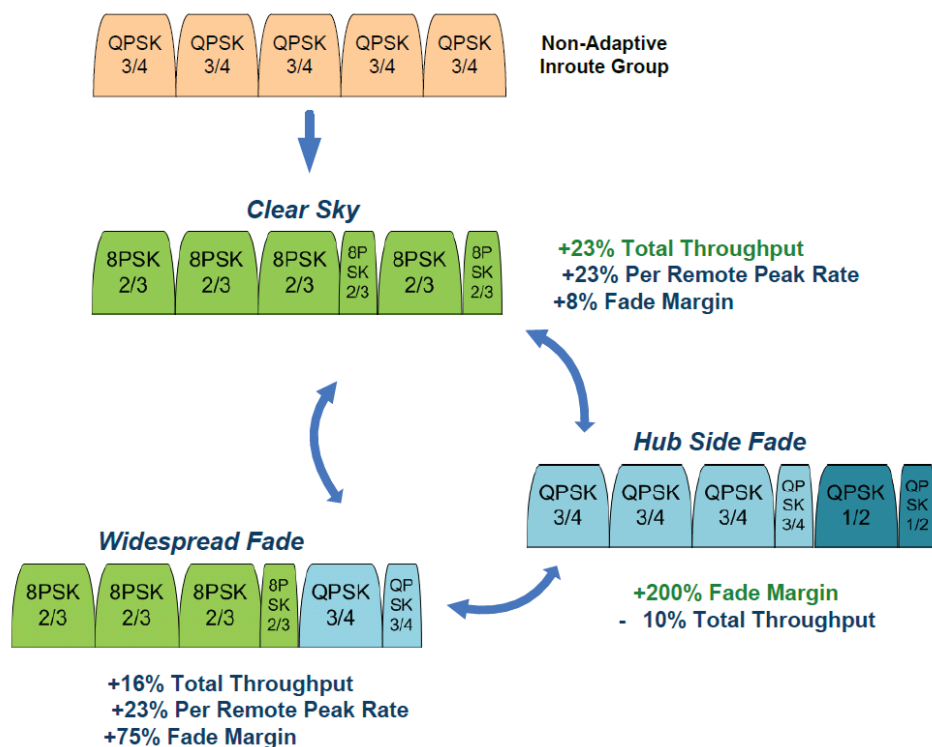


Figure 5 – Performance improvement using Adaptive TDMA

Adaptive TDMA and Group QoS

The ST Engineering iDirect Group Quality of Service (QoS) features function on an Adaptive TDMA inroute in roughly the same way as the Group QoS functions on a static inroute. With throughput gains and enhanced availability, Service Providers will be able to fulfill SLAs with confidence. First, all the existing GQoS features for packet prioritization are supported when working with Adaptive TDMA. The Committed Information Rate (CIR) and Maximum Information Rate (MIR) configured for each remote are honored. When a remote must hop to a slower carrier during rain fade it will be given more time slots on the slower carrier to keep its IP throughput constant even when its peak data rate is reduced.

As with Adaptive Coding and Modulation (ACM) on the downstream carrier, an Adaptive TDMA inroute cannot always guarantee that a remote gets its configured CIR or stays in the network. At some point, a sufficiently deep rain fade will force a remote to use a MODCOD or symbol rate that lacks the IP throughput capacity to deliver a remote's configured CIR. An even deeper rain fade may prevent the remote from staying in network altogether.

Despite GQoS' ability to apportion bandwidth fairly, especially when a remote is in deep rain fade, some network operators may wish to specify a link threshold beyond which the system will not attempt to satisfy a remote's configured QoS. This threshold instructs the system to ignore the configured QoS of remotes in sufficiently deep fade, easing the ability of the system to respect remaining QoS configurations. By default, the threshold is not enabled.

An Adaptive TDMA inroute is more resilient than a static inroute, yet Network Operators must still make policy decisions balancing equipment and bandwidth costs to achieve the desired level of link availability.

Systems Requirements

There are several requirements on the baseband equipment in order to implement Adaptive TDMA. First, Adaptive TDMA is supported only on ST Engineering iDirect X Series, 9 series and iQ Series remotes. Furthermore, only the 2D 16-State FEC codes can be used in an Adaptive TDMA inroute. The 2D 16-State FEC codes are always more efficient than the older Turbo Product Codes (TPC), making this constraint a desirable side effect. At the hub, ST Engineering iDirect's ULC, DLCs and XLC hub line cards can be used to implement Adaptive TDMA.

Additionally, each upstream carrier within an inroute group must use the same payload size. That is, each carrier may be configured with a different MODCOD, but all inroutes within an inroute group must use the same payload size (100, 170 or 438 bytes).

Conclusion

In summary, Adaptive TDMA enhances return channel performance and increases network availability under rain fade and link degradation.

An inroute group can support carriers with different symbol rates, MODCODs and spread factors dynamically adjusting to changing uplink conditions based on each remote's demand and the system's QoS configuration.

Service Providers have the flexibility to choose the business model that's right for them. With Adaptive TDMA they can gain up to 40% in throughput for fixed VSAT with the same satellite bandwidth, and even more for mobile networks. As a result, they can deliver higher data rates to end users eager to run bandwidth-intensive applications and critical operations over satellite.