Winning the Bandwidth Battle

David Bettinger . Chief Technology Officer
iDirect Technologies

Advances such and DVB-S2/ACM and Group QoS enable network operators to address both the challenges of increasing bandwidth efficiency and improving bandwidth management.

Satellite capacity is an expensive resource. And network operators continually struggle to make the best use of their costly supply. Today, however, they can take advantage of two critical industry innovations that can lead to dramatically improved levels of bandwidth efficiency and new ways to allocate bandwidth across an increasingly complex range of customers and usage scenarios.

The first of these innovations is the much anticipated DVB-S2 communications standard powered by Adaptive Coding and Modulation (ACM). The adoption of DVB-S2 for the outbound carrier (hub to remotes) has the potential to deliver widespread technical and business benefits to satellite service providers.

ACM provides an additional and dramatic increase in bandwidth efficiency over the thirty percent already offered by DVB-S2 (over DVB-S) and allows far greater flexibility for network deployments.

With DVB-S2/ACM in place, network operators still need to pursue effective ways to maximize the distribution of their entire capacity. This is accomplished through the advancement of Quality of Service (QoS) by allowing the logical partitioning of shared bandwidth among different customers, groups of remotes and applications according to dynamic service plans that extend to the application level.

This article will examine how network operators can pursue both strategies to boost their competitive advantage.

Increasing Bandwidth Efficiency: DVB-S2/ACM

DVB-S2 is the second generation of the Digital Video Broadcasting Satellite standard used primarily for direct-to-home satellite broadcast. With enhancements such as more sophisticated modulation techniques and low-density parity-check error correction codes (LDPC), DVB-S2 promises a thirty percent bandwidth efficiency increase over existing DVB-S systems.

Most broadcast-orientated DVB-S2 systems provide much the same efficiency in like-for-like network profiles. However, when the DVB-S2 standard is implemented within two-way networks, additional capabilities can be brought into play.

Adaptive Coding and Modulation (ACM) is an enhancement to the DVB-S2 standard that dramatically improves its performance in the two-way VSAT environment, by dynamically optimizing the operating parameters of the outbound carrier.

ACM leverages the return channel to provide an assessment of channel conditions at each remote to determine the optimum link parameters based on satellite link performance, terminal RF Characteristics and local weather conditions.

The hub can then, on a site by site basis, adapt the specific modulation and coding scheme to account for any impairment of the outbound link to each terminal. Continual adjustments are made in real time without intervention by the network operator.

By changing modulation and coding according to current link conditions, DVB-S2 / ACM can provide an increase in bandwidth efficiency greater than fifty percent over non-ACM DVB-S2 systems.

ACM also allows far greater flexibility for network deployments. Traditionally, the specifications of VSATs (antenna size, BUCs, etc.) within a network tended to be fixed from the outset. Link operating parameters had to be maintained throughout the lifetime of services. This created an obstacle to various satellite applications.

Extending the Full Power of ACM

ACM is one of the most complex and extensive features ever implemented by the VSAT industry. It impacts nearly every feature of a satellite network, including data encapsulation, signaling overhead, timing/carrier recovery and real-time network monitoring and configuration.

As such, ACM must be approached as a total system design. If ACM is designed merely around a single element—such as a terminal chipset intended for the broadcast industry (CCM-based DVB-S2)—its performance will be significantly compromised.

The proper design of an ACM system must include all elements of the system (hub plus terminal) in order to ensure:

- Simplicity of network design and configuration for the network operator
- Complete and seamless integration of ACM into a QoS system
- Maximization of the theoretical ACM efficiency gain over a standard broadcast DVB-S2 carrier

In a QoS scenario, for example, the ACM system must be set up to provide information about the link condition at each terminal to the QoS system. This will enable the network to manage specific application requests for guaranteed bandwidth to a particular terminal—factoring both the condition of the link to that terminal and the QoS specifications that govern the specific activity on that terminal.

With ACM and QoS integrated, a VoIP call would continue to get the

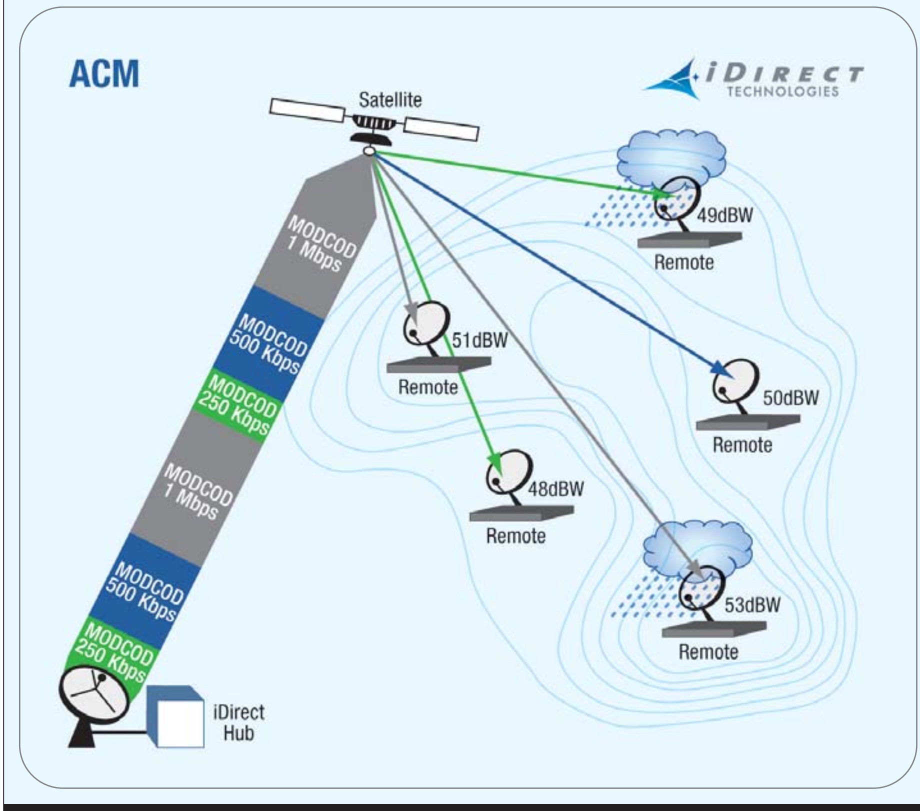

Figure 1. ACM enables each remote to achieve maximum data throughput by utilizing the most efficient coding and modulation scheme dependant upon the location within the satellite contour, antenna size, and clear sky conditions versus rain fade.

same IP data rate required to maintain the call even if the modulation and coding is changed several times during the call due to a passing rain storm.

Overcoming Rain-fade Degradation

One benefit of ACM will be particularly noticeable for network operators in Southeast Asia and other tropical zones where torrential rain challenges satellite links. Until now, network operators have been forced to balance the commercial imperative of making services economically viable for customers against engineering constraints required to maintain links during adverse weather conditions.

ACM automatically optimizes link performance, balancing efficiency and availability as link conditions change. See figure 1. The greater the difference between clear-sky and worst case conditions, the more the benefits of ACM become apparent.

Link margins previously required to survive tropical downpours can now be reassigned during the better conditions that prevail most of the time, to yield higher throughput from the same capacity. Given that rain-fade is often quite localized when averaged across a footprint, more customers can be served within the same capacity.

The availability of ACM for VSAT networks will substantially improve the engineering and commercial viability of Ku-band services, where only C-band was previously considered. This will enable Asia to enjoy the same economic and convenience benefits that Ku-band provides in the U.S. and Europe. ACM may even increase the adoption of Ka-band services, both in the region and across the planet.

Improving Bandwidth Management: Group QoS

However, DVB-S2, even with ACM, addresses only half the battle. It's one thing to maximize the efficiency of bandwidth based on the real-time conditions at every remote location. It's quite another to optimize the delivery of bandwidth across a network based on a unique and complex set of service requirements.

Establishing a QoS strategy to meet Service Level Agreements (SLAs) is a key requirement for network operators. Increasingly, the challenges presented to network operators to meet complex SLAs demand a flexible and powerful set of bandwidth control mechanisms.

Already, network operators can divide a network into logical groups with differing SLAs while sharing the same physical carrier. The network operator can even assign these logical bandwidth groups to different end customers or assign the group to different service providers or Virtual Network Operators (VNO).

A service provider or a VNO can further divide a bandwidth group into sub-groups referred to as "service groups" and assign each service group to a different customer. A service group could be used strictly to link remotes into sub-groups, or more typically, to differentiate groups by class of service. A Platinum, Gold, and Silver service, for example, could be defined as service groups under the same bandwidth group.

Representing a major development in expanding QoS, network operators can now define sub-groups by even more detailed criteria such as business processes, technology applications and data formats. We can call this innovation Group Quality of Service, or Group QoS.

Taking QoS Further

Groups QoS enables network operators to ensure SLAs on both inbound and outbound bandwidth. The ability to provide QoS on the upstream is crucial for real-time applications such as VoIP during periods of congestion. There are multiple scenarios that a network operator can now efficiently serve. Below are just a few examples.

Logical partitioning of shared bandwidth:

In this scenario, a satellite provider with six customer networks requiring a 3 Mbps outbound and 1 Mbps inbound each can configure a single 18 Mbps outbound carrier and multiple inbounds for a total of 6 Mbps that are shared by all six customers. The logical partitioning of the bandwidth allows the satellite operator to enforce a Committed Information Rate (CIR) of 3 Mbps for each customer on the outbound and a CIR of 1 Mbps on the inbound.

Furthermore, the satellite provider can configure a bursting capability beyond the 3 Mbps/1 Mbps CIR for

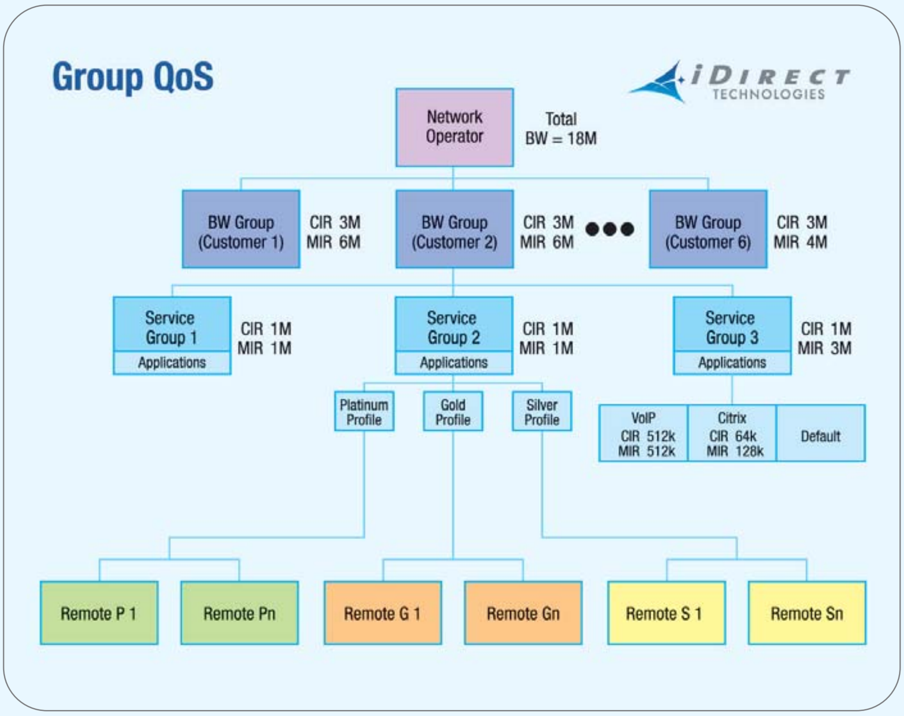

Figure 2. Group QoS allows for multiple configuration options increasing the flexibility for Network Operators when configuring by: Bandwidth Groups / Default Profile / Service Group / Additional Profiles of Remotes / Applications

some customers if desired. This allows customers to take advantage of capacity that is left unused by other customers that are sharing the same physical bandwidth. See figure 2.

Tiered service scenario:

In addition to CIR and MIR (Maximum Information Rate), priority can be used to provide a tiered service. A network operator with an 18 Mbps outbound, for example, can assign the highest priority for a Platinum Service with a bursting capability up to a MIR of 10 Mbps. That network operator can also establish a Gold Service with the next priority and a MIR of 12 Mbps, and a Silver Service, which would be a best effort service at a lower priority with a MIR of 18 Mbps.

In this example, the Platinum Service is guaranteed up to 10 Mbps, while the Gold Service is guaranteed 8 Mbps but could go as high as 12 Mbps. The Silver Service may use the entire 18 Mbps if there is no demand from the Gold or Platinum but is not guaranteed any bandwidth. Remotes in each tier will have different profiles that may include different applications with their CIRs as well as different MIRs for each application or total remote MIR. See figure 3.

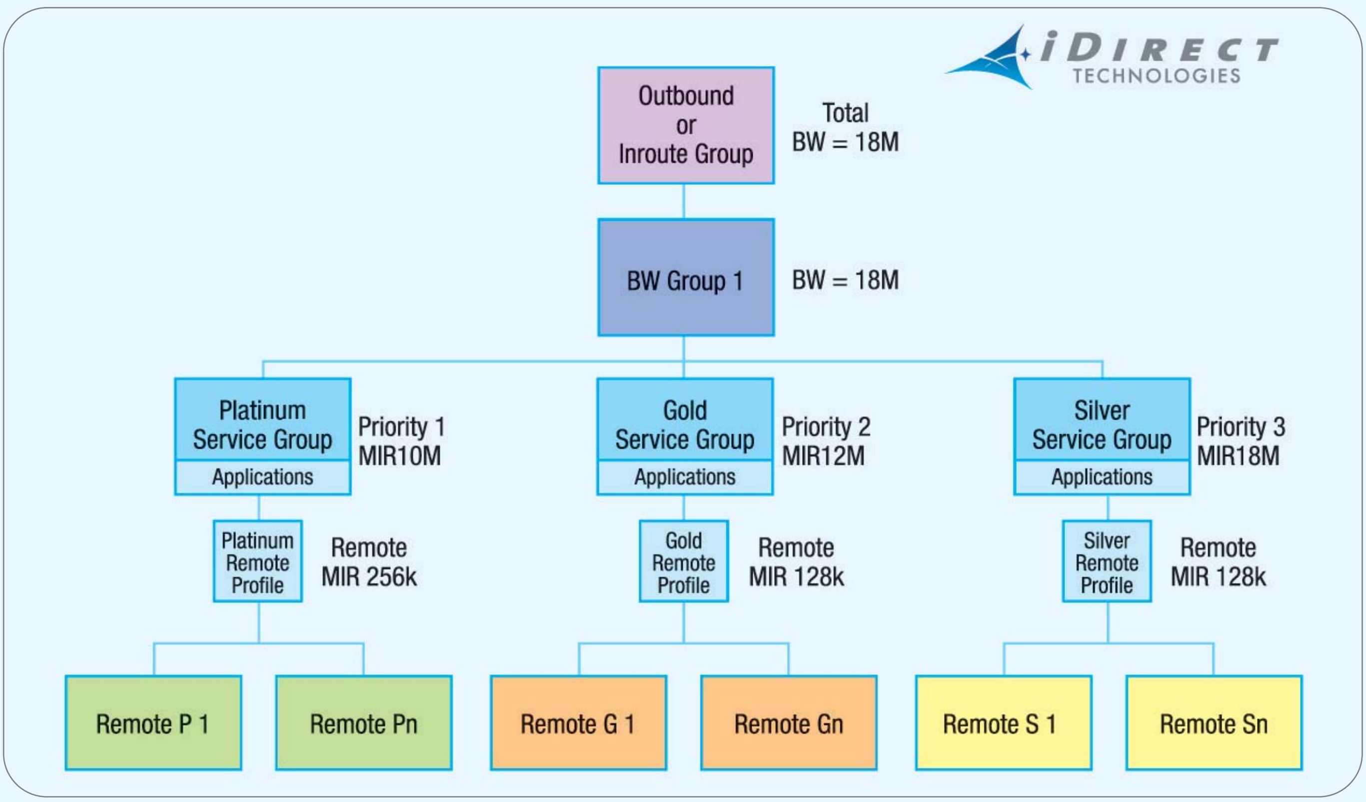

Figure 3. Tiered Service

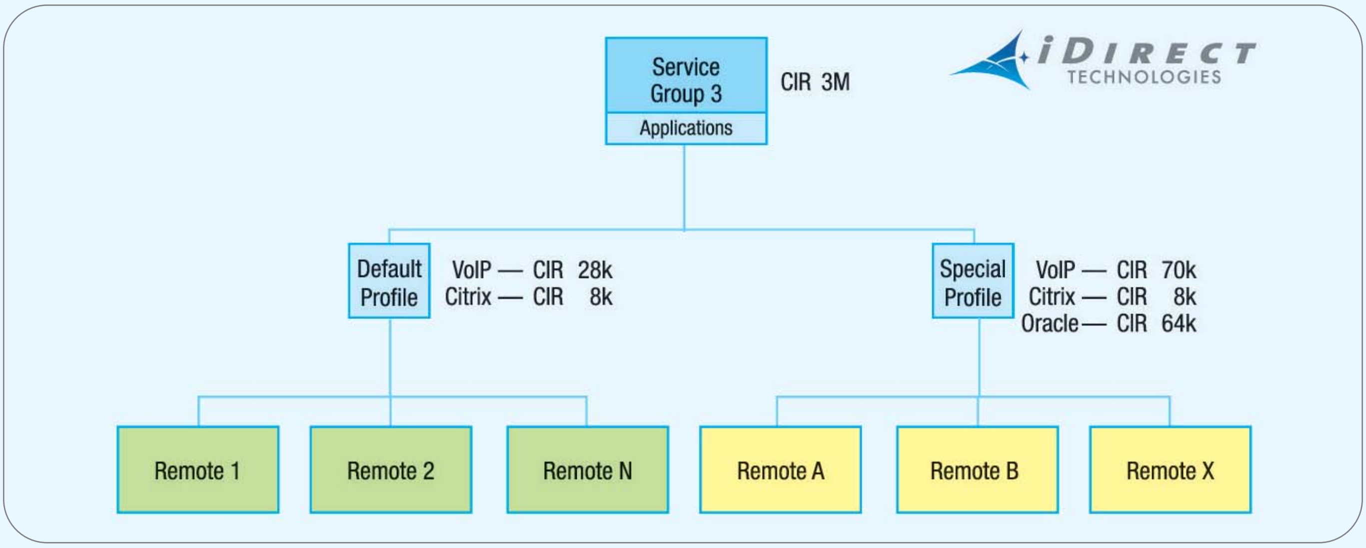

Figure 4. CIR by Application

CIR per application scenario:

In this situation, a customer with a 3 Mbps bandwidth on the outbound and 2 Mbps on the inbound may want to dedicate 1 Mbps for VoIP and allow each remote a CIR of 28 Kbps dedicated to VoIP to support two voice calls. The ability to configure a CIR by application, both at the network level as well as at the remote level, provides customers with this added flexibility.

Figure 4 shows an example of a customer configuring different applications on a 3 Mbps outbound by dedicating 1 Mbps for VoIP, 512 Kbps for Citrix applications and 1 Mbps for Oracle applications.

In this example, two different remote profiles are created. The first profile allows a CIR of 28 Kbps for VoIP, but no CIR for Oracle applications. The second profile allows a CIR of 70 Kbps for VoIP and a CIR of 64 Kbps for Oracle applications. Both profiles are configured for a CIR of 8 Kbps per remote for Citrix applications.

Multiple customers sharing a physical remote:

This scenario addresses many verticals where there are multiple tenant dwelling applications. Let's say a network operator provides service to oil rigs for two major oil companies. Many of the oil rigs service both oil companies. Instead of installing two separate remotes on oil rigs serving both companies, distinct VLANs can be used on shared remotes to separate traffic of the two companies.

QoS can be configured to allocate bandwidth by VLAN from two different logical service groups. This allows each company to have its own dedicated bandwidth and get billed separately by the network operator, while sharing the same physical remote.

Conclusion

As the demand for satellite communications continues to accelerate in Asia and globally, network operators must be do everything they can to increase bandwidth efficiency and improve bandwidth management. Advances such and DVB-S2/ACM and Group QoS enable network operators to address both these challenges. And they present a tremendous opportunity for network operators to decrease operating costs and improve their service offerings at a time when the satellite industry is experiencing renewed growth.

David Bettinger joined iDirect Technologies as the Director of Hardware Engineering in 1996 and took over responsibility of all hardware and software development as VP of Engineering in 2002. In his role as Chief Technology Officer he is responsible for the oversight of all technology decisions within iDirect and serves to drive the strategic direction for product development. Mr. Bettinger has been active in the satellite networking industry for over 15 years, and previous to iDirect, he was a senior member of the technical staff at Hughes Network Systems in the Satellite Networks Division. Mr. Bettinger is a graduate of Virginia Tech with a Masters of Science degree in Electrical Engineering.