30th Space Symposium, Technical Track, Colorado Springs, Colorado, United States of America
Presented on May 21, 2014

THE IMPACT OF MASS MARKET SATELLITE BROADBAND ON US DOD SATCOM STRATEGIES

Steve Gardner

ViaSat, Inc. steve.gardner@viasat.com

ABSTRACT

Data communication by satellite has traditionally been viewed as a service of last resort, primarily serving customers that can't be reached any other way. Examples of markets that fit this niche are maritime shipping, resource extraction, and the military. Because the market has been defined as one without alternatives, the pressure for technology advancement has been low, and for nearly a quarter century, two-way satellite technology has generally moved at a leisurely pace compared to other communication fields. The advent of High Capacity Satellites (HCS) in commercial markets has disrupted this paradigm, and the implications for the DoD are profound. One of the best indicators of the magnitude of the change is that today a traveler with a coach class ticket on Jet Blue airlines can get an Exede by ViaSat© broadband connection that's faster than any head of state can use on their government aircraft.

In an independent study based on actual measurements of performance experienced by end users, the Federal Communication Commission (FCC) found that well-designed broadband access provided on High Capacity Satellites can outperform the majority of terrestrial consumer broadband offerings. This development vaults satellite into a world in which terminal designs must achieve mass market price points, design cycles must shrink to a year or two, and reliability must be "4 nines". Although DoD applications may seem disconnected from the mass market, commercial HCS satellites will serve many of these applications with far better performance and capability, far greater affordably, and far better resilience than would have seemed possible just a few years ago. Commercial High Capacity Satellites provide a viable path for the DoD to meet its future capacity needs.

Index Terms—High Capacity Satellites, Ka-band, satellite broadband, resilience, commercial satellites.

INTRODUCTION

In 2008 ViaSat announced that in three years it would launch its ViaSat-1 satellite to provide consumer broadband in the United States. The satellite was designed to be able to serve over a million broadband customers with internet access at speeds as good as the best DSL and better than many cable offerings. The capital cost (other than consumer terminal cost) to reach any one subscriber is about $500, which is economically competitive with terrestrial approaches.

After a year of operation, ViaSat's Exede Satellite Internet service became the first satellite offering ever to be included in the annual FCC benchmarking study [1] of broadband speeds (a study that mostly includes DSL, cable and fiber offerings), and when the report presenting study findings was published in February 2013 Exede took the top spot for performance with respect to advertised speeds (Figure 1). Moreover, based on FCC data [2], Exede's advertised speeds of 12 Mbps to the home and 3 Mbps return are very competitive with most terrestrial offerings (Figure 2). In one blow, the concept of satellite as a service of last resort has been shattered.

This result will impact every area of satellite communications profoundly, because the economies that result from the mass market will roll back into the smaller, high value markets that satellite has traditionally served. Satellite is now in the mainstream of broadband delivery services, and the industry – service providers, consumers, and the DoD – must start viewing itself in this new light.

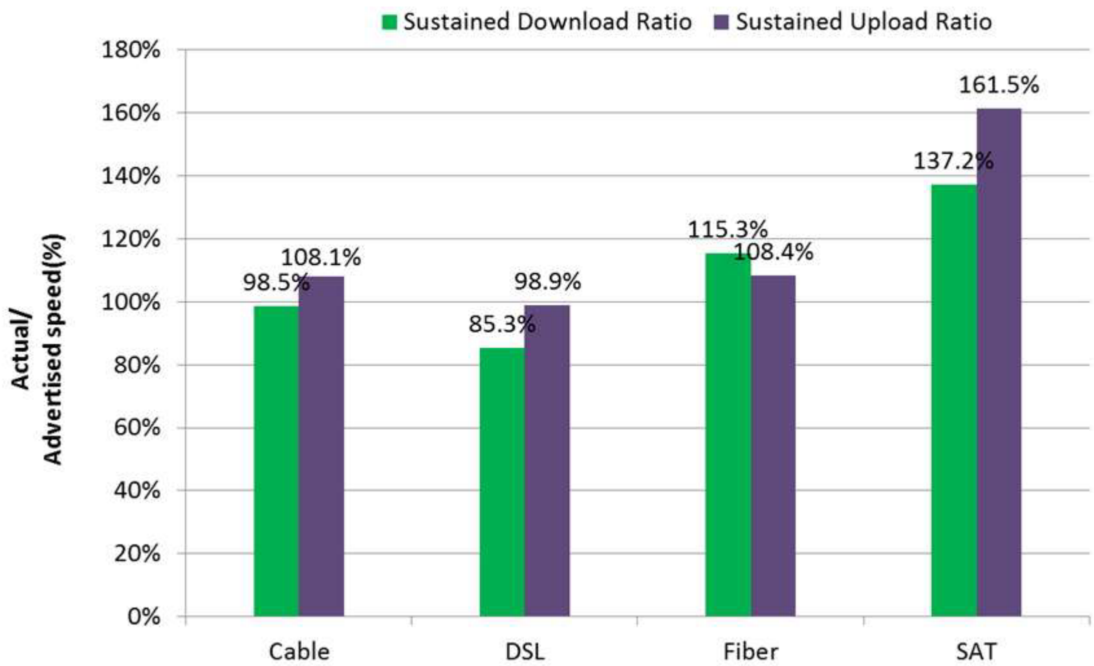

Figure 1: HCS Performance vs Advertised Speeds as Measured by the FCC

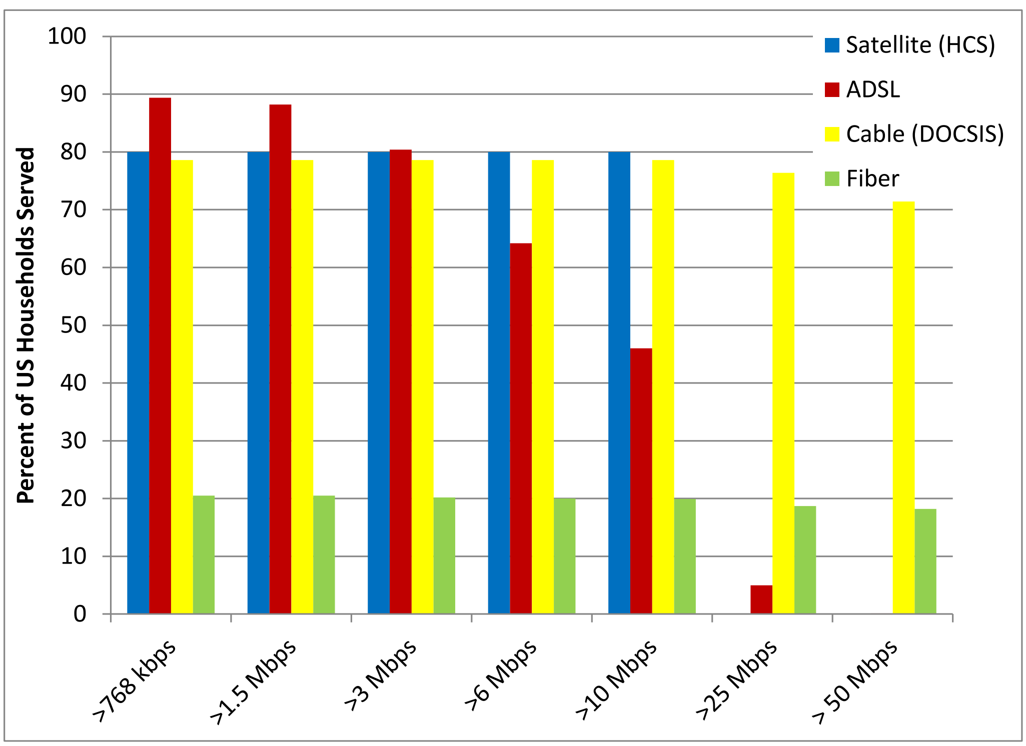

Figure 2: Percentage of US Population Served by Various Broadband Technologies and Speeds

THE COMMUNICATION SATELLITE AS A SYSTEM COMPONENT

In most fields of communication, the system design considers the functionality, performance and cost of both ends of the link as part of a total system optimization problem. For example, in cellular telephony networks, architectural design considers the economics of handsets, base stations, spectrum, and performance, and proceeds by trading each against the other in an attempt to find the best system outcome. A cellular architecture that uses more base stations and thus increases infrastructure cost might enable cheaper handsets, and given the volume of handsets, this might be a good trade. Increasing handset cost might use spectrum more efficiently and allow more subscribers, thus increasing revenue in a way that overcomes the handset cost. Cellular evolution in the past two decades has served to enable new and more demanding applications to fit into existing bandwidth resources without disrupting cost models. This evolution would be impossible without an end-to-end system design encompassing everything from the application to the waveform.

But in satellite data communications it has been commonplace for the space segment design to be isolated from the ground segment design and the application. Too often the satellite design doesn't include terminal performance and the waveform in the trade space, and there is even less opportunity to provide optimization needed to expand the application space. This separation of roles is a crucial problem for the satellite industry in both commercial and government sectors and has been a primary impediment to satellite's ability to move into the mass market.

THE BURDEN OF LEGACY EQUIPMENT

Of course, the implication of an end-to-end approach to the satellite system design is that a new ground segment is required, and this impacts user terminals, gateways, and terrestrial networks. Substantial investment in the installed base of legacy terrestrial equipment creates a very high psychological barrier to new satellite approaches. But this is a coin with two sides, since the history of technology is littered with examples where a legacy installed base has been more of a curse than blessing. Consider the period around 1982-83 when the PC first began to make substantial in-roads into enterprise computing. Companies with a large infrastructure of mainframe computers were faced with a dilemma: build upon the base, or throw everything overboard and start from scratch? Those who over-valued their installed base often paid severely over the next few years in higher costs and lower capability.

The danger in asserting the value of legacy equipment is that the question is really about total economic value and not just about the sunk cost of existing equipment. If new equipment costs a fraction of legacy equipment, enables far lower service cost, and provides a higher degree of capability, then the legacy equipment will be a burden, not a boon. This is in fact the situation the satellite community, including the military, has entered with the advent of HCS.

One of the most notable architectural differences between HCS offerings and traditional commercial SATCOM is that HCS service offerings are based on selling bits and not Hz. The tradition of selling raw bandwidth is another fundamental difference between the existing satellite industry and other communication fields. It is almost universally recognized that sharing the channel medium is critical to efficiency of communication systems. Yet with only isolated exceptions the satellite communications industry, including the DoD, has failed to embrace this tenet. The DoD relies far too heavily on dedicated connections rather than taking advantage of the well understood benefits of a shared medium. To gain some perspective, imagine if the cell phone industry couldn't persuade its customers that they didn't require a dedicated frequency channel on a 24/7 basis to ensure access to the network. A mass market cell phone industry would be completely impossible in this situation. Almost every other field of communications has come to recognize that, far from impeding performance, sharing resources generates substantial improvement.

The DoD clearly can't be expected to turn on a dime with respect to legacy technologies. But they can recognize that continued support for the legacy equipment is in opposition to moving to a much better future and consider new proposals for terminal investments with a view of whether these investments are consistent with achieving that future.

WHAT KIND OF BROADBAND SATCOM DOES THE DOD NEED?

It's very hard to predict exactly what broadband SATCOM capabilities the DoD is going to need in the future. But some generalized required attributes can be asserted with confidence, and these are useful qualifiers to keep in mind.

Capacity available everywhere missions can occur

The DoD can potentially need to conduct operations nearly everywhere on the surface of the earth that can be reached by geosynchronous SATCOM. WGS uses steerable beams to provide an "anywhere" capability, allowing a limited number of platforms to have communication anywhere on the visible earth. But future HCS satellites will have "everywhere" capability, meaning that large numbers of platforms scattered everywhere on earth visible from a geo-synchronous orbital slot will be supported.

Concentrated capacity in a theatre

The "everywhere" requirement demands relative thin capacity available globally. Satellites like Inmarsat's I5 do this today, although the cost per bit of these satellites is not competitive with true HCS. But to respond to a conflict requires an ability to concentrate capacity in a theatre, and the location of the theatre can't be identified in advance of the conflict. Future conflicts of the scale of OIF or OEF will require many tens of Gbps of capacity in a region the size of the state of Kentucky. Thin capacity satellites can't provide this, but HCS designs can.

Highest possible speeds to low cost terminals

Access to information is like access to oxygen in today's world. Individual warfighters will need access to a small, low cost terminal that provides a user experience (including data rates) comparable to or better than what they have in civilian life. At higher echelons, broadband connectivity comparable to commercial enterprise service levels will be demanded. High Capacity Satellites have enabled the development of low cost terminals with consumer price points for a broad range of applications. For example, residential terminals operate outdoors in Alaska winters and Texas summers yet cost under $500. Similarly, quick-deploy terminals for satellite news gathering (SNG) costing a small fraction of a WIN-T terminal are used by a number of US media outlets to stream HD video back to studios.

The speeds provided by these terminals are impressive. A terminal with a 70 cm reflector achieves rates of many tens of Mbps. These rates are made possible by the high G/T and EIRP of the satellite, which result from the high gain antennas used to make their small beams.

Low cost of capacity

The DoD has the same problem that consumer markets have. Increase in capacity demand is unrelenting, driven by the same force for both markets: video streaming. The wind-down of OIF and OEF may provide a temporary lull in demand growth as the number of DoD consumers decreases, but consumption per user can be guaranteed to keep increasing. To be ready for the next conflict, the DoD should be planning to meet this capacity demand. Only an approach based on HCS can do this.

Resilience

DoD SATCOM requires jam resistance and secure transmissions. HCS can deliver both. Consider anti-jam performance. HCS has two of the most important attributes that a jam resistant system should have. First, it has very wide transponder bandwidths (typically 500 MHz or more). Since the power a jammer needs is directly proportional to the available bandwidth, this means that HCS has about 10 dB better jamming resistance than a traditional satellite with a typical 54 MHz transponder bandwidth. Secondly, High Capacity Satellites have small beams, which greatly limits the geographic area in which the jammer can be successful. HCS can be as effective as AEHF in jam resistance.

HCS can provide good transmission security as well. DoD consumers can use end-to-end HAIPE encryption, but the SATCOM link can also provide good privacy. Standard techniques to defeat traffic analysis or provide cyber security are important in the commercial world. These techniques exist to some extent today and can be expected to strengthen over the next few years.

HOSTED PAYLOADS OR CONDOSAT?

The use of hosted payloads by the DoD has received substantial attention of late. While such payloads may make sense in some scenarios (e.g., remote sensing), it's hard to envision a hosted payload as an effective basis for broadband capability. The most significant reason for this is that minimum cost per unit capacity is achieved when the payload is as big as the satellite bus can accept. A payload design providing respectable capacity economics requires a host that is severely under-utilizing the available bus resources, which is an unlikely event. Secondly, a broadband payload that provides low cost per bit requires a lot of TWTAs, a lot of waveguide plumbing, and considerable reflector deck real estate, none of which lend themselves to a modular design. This means that the hosted payload design has to be specific to the host, and it thus has to be specified and priced when the hosting opportunity becomes available. This creates many procurement headaches: the commercial world moves from RFP to contract in a few months, far faster than the government typically can react.

A CondoSat approach could be more sensible. In this scenario, the government can buy part of the capacity of a commercial HCS. The term "CondoSat" arises from the real estate concept. The owner of a condominium secures rights to a portion of a building, but what exactly is his property is not always easy to tell. The owner has right of use of common areas such as walkways, grounds, pools, and tennis courts. He owns his internal walls and improvements, but he may share walls or roofs with his neighbors. A CondoSat is similar in that the owners share the bus, the power system, the reflectors, and possibly even the waveguide and TWTAs, but there is agreement that each party has access to certain shares of the satellite capacity in certain geographic regions.

A CondoSat approach makes participation in HCS much easier. The DoD demand for capacity is far smaller than what a typical HCS can provide. This means that the DoD can participate in the economics of HCS – benefiting from the statistical advantage of sharing a large amount of capacity among a large number of users. The DoD doesn't have to build the large number of gateways that HCS requires: they instead can ride on top of the gateways built to support the underlying commercial satellite network.

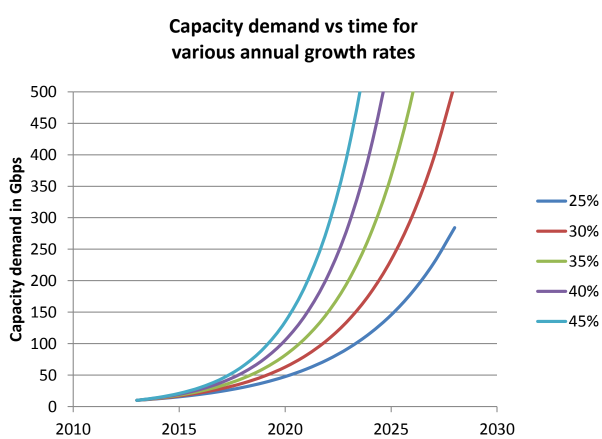

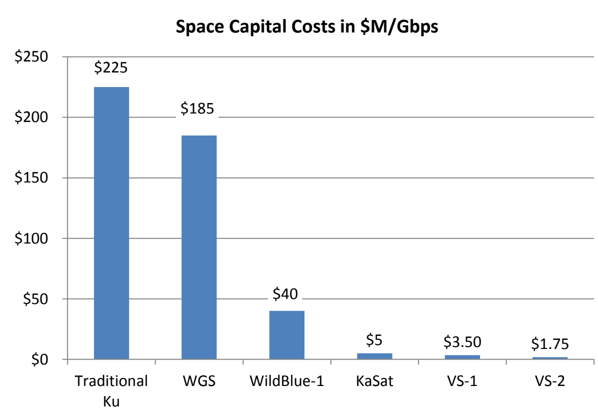

Figure 3 Potential capacity demand growth and actual evolution of space capital costs

The true High Capacity Satellites on orbit today cover regions that are relatively small by the standards that the DoD demands, but the real advantage of participating in the HCS-spawned mass market will come with future satellites that will eventually paint a global picture. Satellite broadband has proven itself capable of meeting today's consumer market requirements. Estimates of the historic growth in consumer broadband speeds vary from doubling every eighteen months to four years, and no one forecasts that this will slow any time soon. Yet the price that consumers will be satisfied in paying will not increase appreciably, which means that for a service provider to have a successful business plan, he must see a path to reducing costs at a rate that inverts the capacity demand growth. The Omnibus Broadband Initiative Technical Report No. 4 entitled "Broadband Performance" [3] states that in the United States per-person usage growth was 30-35% annually in 2009. The first chart in Figure 3 gives an example of what these kinds of growth rates mean in a scenario in which 10 Gbps is needed to serve a market today: *by 2024 it must be possible to provide 250 Gbps of capacity at the same cost as 10 Gbps today*. The second chart in Figure 3 shows that the evolution of commercial HCS satellites has been moving in the right direction. Taken together, the two charts illustrate the futility of being constrained by a perceived value of legacy satellites and equipment.

Figure 4 ViaSat-2 coverage footprint

Satellite providers naturally view the details of their future plans as highly valuable and will be reluctant to explain their technologies or business strategy publicly, but it is clear that continued success in the mass market requires innovative architectures and thinking in new ways. More important than the specific details is that the pressure of the demand growth resulting from the entry of HCS into the mass market has created completely new dynamics in the satellite industry that require a much greater pace of innovation than has been the case previously.

Announced in May 2013, the ViaSat-2 satellite provides an indication that such an innovation path exists. Prior to this announcement, the general perception in the industry was that high capacity required small beams, and that use of small beams in turn would limit coverage to small regions. Thus it has been widely believed that the choice was between architectures similar to Inmarsat's I5 constellation that provide modest capacity over a large area, or satellites like ViaSat-1, that provide large capacity over a modest area. ViaSat-2 breaks this paradigm by providing both high capacity and broad coverage area (Figure 4) through the use of a unique satellite system architecture. ViaSat-2 provides twice the capacity economics of ViaSat-1 and seven times the geographic coverage. This architectural evolution is motivated by the desire to layer high performance mobile broadband services onto the underlying (and ever more demanding) residential broadband business – growing the application space in a manner analogous to the cellular industry approach described above. While targeted primarily at commercial markets, the ViaSat-2 architecture has the potential to bring the power of the mass market to many DoD applications.

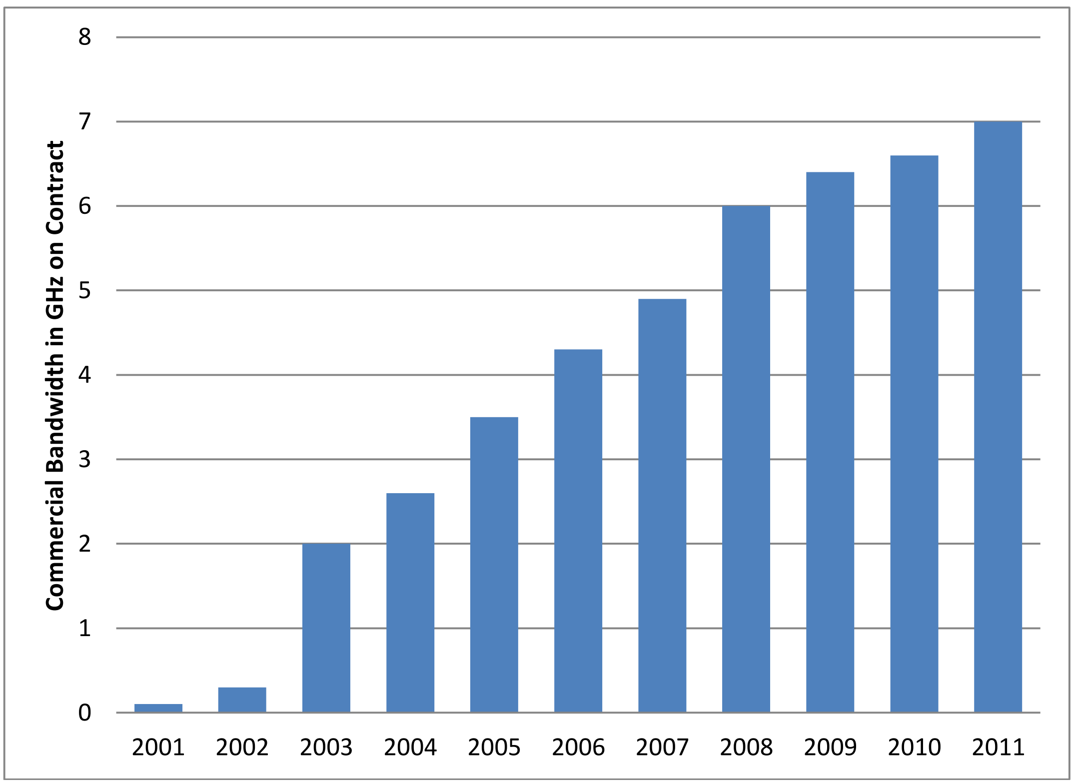

Figure 5 DISA Commercial SATCOM Bandwidth Demand Growth

THE IMPACT OF MASS MARKET SATCOM ON THE DOD

Throughout the period of the Iraq and Afghanistan wars, DoD demand for satellite bandwidth has risen to levels that were never previously conceived (Figure 5). SATCOM-enabled applications such as airborne ISR have provided vital new capabilities, but have propelled bandwidth demand dramatically. Interestingly, the growth in total DoD broadband demand closely matched the growth in consumer demand in this period, and there are good reasons to assume that the two growth curves will remain closely aligned in the future. Innovative as WGS satellites were in their time, their architecture doesn't scale in the way needed to follow this kind of growth curve in an affordable way. Moreover, government procurement methodology separates the definition and purchase of ground and space segment, making the kind of total system optimization used in HCS design difficult or impossible while creating an inflated perception of the value of legacy equipment.

Thus there are roadblocks in both space and ground. One path around is a more whole-hearted adoption of commercial technologies. While this is a big change, it should be thought of in the context of other commercial services used by the DoD. For example, the DoD doesn't run its own phone company – it buys service from commercial carriers, and uses handsets designed for the commercial marketplace. As another example, commercial smart phones have vital capabilities and a tempo of innovation driven by a mass market whose sheer size makes it far more of a change force than any amount of DoD business could ever be. As a result, the armed services are beginning to use these commercial devices in tactical environments.

Pursuing the mass market trend in satellite communications can also have substantial advantages for the DoD, and this is true both with respect to leasing of satellite capacity and in terminal acquisition. Consider the question of leased capacity first. There is increasing awareness in the DoD that capacity demand is outpacing the rate at which government owned satellites can be acquired and launched, and that leasing of commercial capacity will be required. Some commercial bandwidth providers have expressed concern that the current government acquisition approach that only allows lease commitments for a year at a time makes it economically impossible for them to have the capacity in place, and they have argued for a new acquisition approach with long commitments. This argument makes some sense when a "service of last resort" paradigm exists. But with satellite communications in the mass market and subject to the downward cost pressures that this entails, the DoD should be wary about long term agreements. It could find in a few years that it is trapped by commitments costing more than an order of magnitude beyond current market prices and providing far less capacity than the DoD needs.

DoD strategies should be built around the supposition of a not-too-distant future in which access to low cost capacity on commercial HCS is available everywhere that's of interest. Informed by this view, DoD strategy might also consider ways to facilitate moving that future closer to the near term rather than supporting legacy technologies through long term leases that favor the persistence of services that the DoD should not really want.

With respect to DoD terminals, today these are often acquired under costly programs of record that take years to complete and the terminals themselves are often expensive, large and of limited capability. Once in place, the installed base of terminals is regarded as a highly valuable capital investment that needs to be defended and maintained in place for a long time.

While it is commonplace to assign the blame for the cost of these terminals to lack of competition and to ruggedization requirements, the lack of economies of scale in manufacturing is a far bigger contributor. Moreover, because of the cost of new development, terminals are used well past the point of obsolescence. No one would expect the DoD to use computers or cell phones that were designed 15 years ago. The tempo of innovation in these devices is far too great, so that their capabilities become woefully inadequate after the passage of a few years, and the mass market use of the devices makes them so inexpensive that a frequent wholesale replacement of the installed base is economically justifiable. With the movement of satellite broadband into the mass market, the same situation is beginning to exist for satellite terminals, which cost as little as $500 today.

The DoD has worked diligently and thoughtfully to evolve terminal procurement processes with the goal of controlling costs. But the best processes conceivable can never match the power of a mass market in driving down unit cost. Contrast the development of a government funded DoD terminal to that of a terminal developed for an HCS system. In the HCS terminal development, the waveform is developed jointly with the satellite with the intent of maximizing the return on investment. In doing this, the terminal design team considers the total cost of terminals – development, production, and operating cost. The trade between development and production cost drives decisions about what design features should be included to improve manufacturing and test cost – for example, enabling a high degree of automation or increasing the level of integration by designing custom integrated circuit devices. Since the volumes of terminals in the consumer market are in the millions, substantial investment in automation and integration is warranted.

In a traditional DoD terminal program, development is funded under one program that is typically awarded to the lowest bidder, and other contracts are then subsequently awarded for production (sometimes the development contract is accompanied by an initial production run, but this is usually small). The separation of the development from the total production run provides a disincentive for designing features that reduce unit cost, but even without this disincentive, DoD volumes are never sufficient to motivate the levels of integration and automation needed for consumer price points. It is often assumed that lack of competition drives terminal cost. While this may be somewhat true, the lack of volume is a much bigger factor.

Ruggedization certainly adds something to terminal cost, but its impact is vastly over-rated in comparison to volume. If the electronics design of a DoD SATCOM terminal can draw heavily from a mass market design, then the price of a rugged terminal will benefit greatly from spreading the cost of system and technology development across a broader base of customers.

Thus strong consideration should be given to a path forward for DoD SATCOM that follows commercial models. Rather than specifying waveforms and attempting to create open standards with multiple sources, the DoD could specify capabilities and cost points and accept terminals that can meet these needs.

HCS ADVANTAGES FOR PROTECTED COMMUNICATIONS

Another consideration for DoD applications is that HCS architectures have interesting features for protected communications. One of the first considerations in a contested environment is the amount of bandwidth available, and HCS satellites typically have transponder bandwidths with 10 to 50 times the bandwidth of traditional satellites, providing the potential for 10 to 17 dB more processing gain.

Secondly, transponders in a traditional satellite typically cover large geographic areas, and an interference source can be effective when operating anywhere in this area. The small transponder coverage area of HCS satellites like ViaSat-1 greatly reduces the options for employing a jammer.

Third, the gateway in an HCS system is typically not co-located with the user beam. This means that to interfere with the forward link requires a physical presence in the gateway beam, and interfering with the reverse link requires a presence in the user beam.

Thus it is typically extremely difficult to interfere with the forward link of an HCS system, and the options for interfering with the reverse link are much more limited than for traditional satellites. And when interference does occur, HCS provides much greater mitigation through processing gain.

SATELLITE ROAMING – A MIGRATION PATH TO THE FUTURE

A complaint about HCS technology is that it isn't universally available. This is true – today. But the fact that HCS is not available everywhere does not have to be a reason for waiting until it is. Terminals capable of roaming across satellite networks can be developed at reasonable cost using a model similar to a cell phone designed for international use. A roaming cellphone has a preferred service provider, or perhaps a list of preferred providers with some set of selection criteria. These providers don't necessarily use the same waveforms, they may not use the same frequency bands, and they might offer different functionality and performance in different areas. The cell phone determines which networks are available at any given time and place, and it chooses the waveform and frequencies suitable to the preferred network. All of this is transparent to the user.

These same kinds of roaming capabilities are now being developed for satellite networks, and roaming agreements are being negotiated between satellite operators so that the entire process is transparent to the end customer. This concept could be extended to include government terminals that can operate on either WGS or on commercial networks. Such terminals might provide substantially higher levels of capability when operating on commercial HCS service, and over time as the footprint of commercial HCS expands, these greater capabilities will be available in more and more places. Thus multi-waveform terminals with satellite roaming can provide an ability to begin realizing the benefits of HCS well ahead of worldwide coverage.

CONCLUSION

Satellite data communications has entered a new era. The model of the past quarter century, in which satellite was the communications service of last resort for applications with no other options, provided a comfortable but limited market in which a slow pace of innovation was tolerated. But with the advent of High Capacity Satellites, satellite data communication has entered the mass market and has proven that it can compete in capability and economics with terrestrial data communication systems. However, success in the mass market over the long term demands a sharply increased tempo of innovation. Satellite systems must be viewed as fitting an economic model where the cost per unit capability falls by a factor of two every few years. There are providers who believe that they can achieve this, and the evidence to date says that they are right.

The "service of last resort" sector of the industry, which includes the DoD, should expect to benefit hugely from innovations made for the mass market, with substantial improvements in costs of both terminals and service accompanying significant upgrades in capability. Commercial High Capacity Satellites will provide what is almost certainly the only way for the DOD to escape being caught between the hammer of budget constraints and the anvil of rising broadband satellite capacity demand. The DoD needs to align its SATCOM strategies with this situation.

REFERENCES

[1] Measuring Broadband America - February 2013, FCC Report online at http://www.fcc.gov/measuring-broadband-america/2013/February

[2] Broadband Statistics Report: Access to Broadband Technology by Speed, January 2013 FCC report online at http://www.broadbandmap.gov/download/Technology%20by%20Speed.pdf

[3] Broadband Performance – Omnibus Broadband Initiative Technical Paper No. 4, August 2010 available online at http://www.broadband.gov/plan/broadband-working-reports-technical-papers.html

