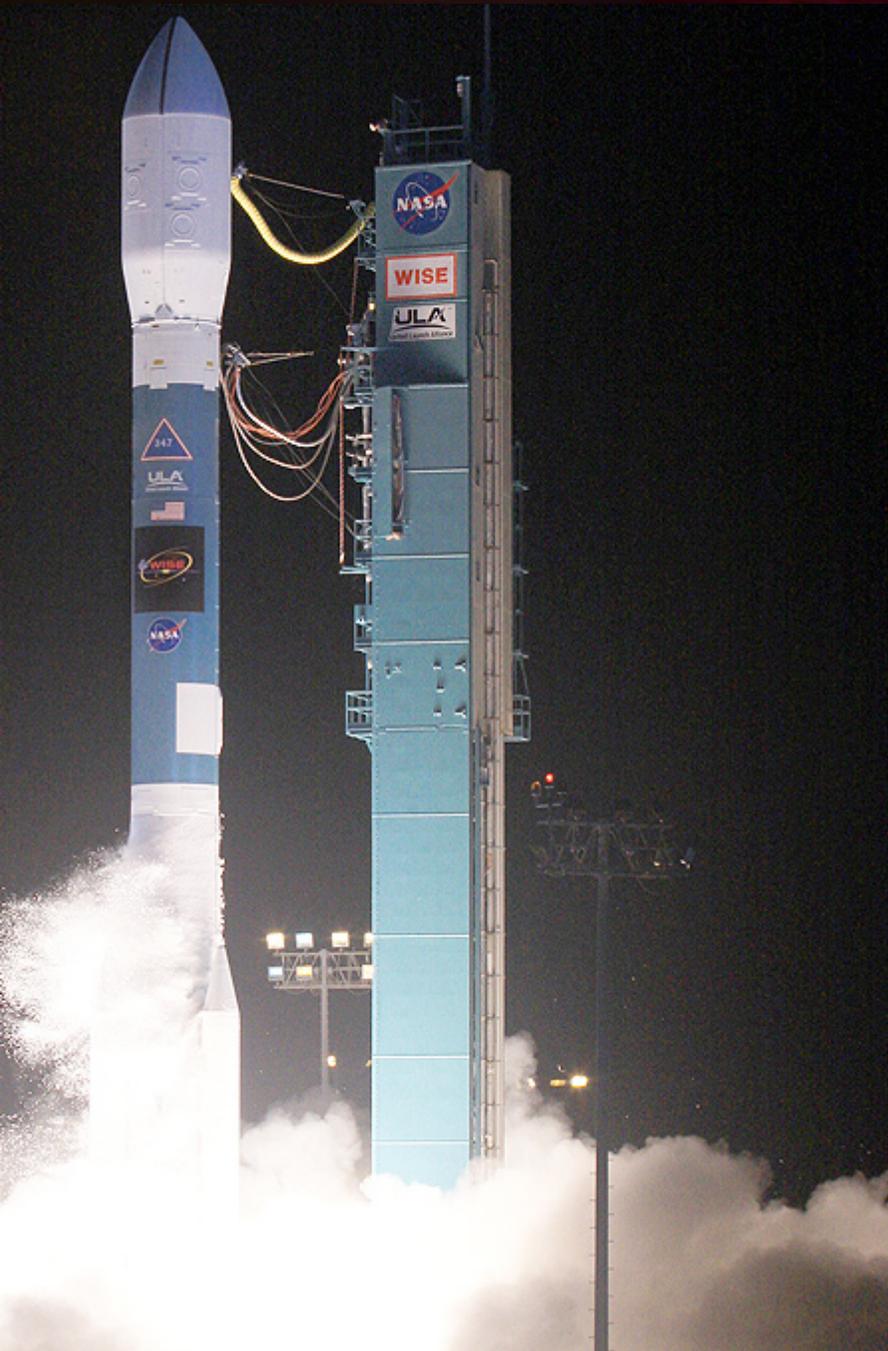
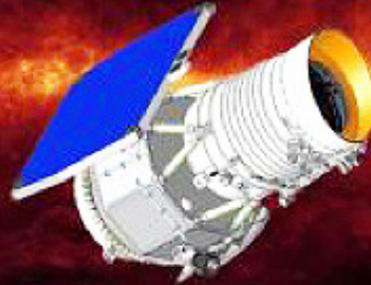
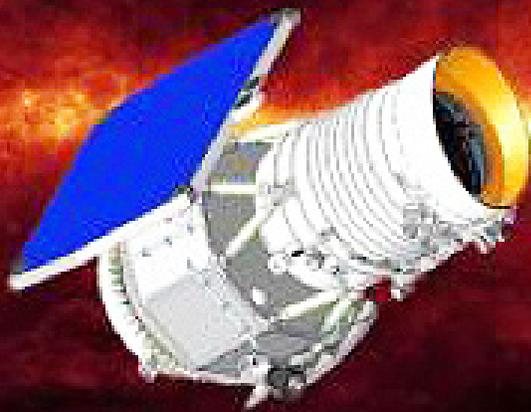


SatMagazine



THE U.S. SATELLITE MARKET

- BEAM
- A WISE PROJECT
- MONAHAN OF RT LOGIC
- NORTHERN SKY RESEARCH
- INSIGHT WITH FORRESTER
- O3B NETWORKS
- CABSAT MENA 2010
- PACTEL INTERNATIONAL
- COMBAT SAT INTERFERENCE
- WHAT IS NETWORK LATENCY?



This past year has been rather “interesting” for the U.S. satellite industry. While some companies have faired extremely well, others have felt enormous pressures descend upon them as global competition works its will upon the various SATCOM segments — most specifically the satellite manufacturing and launch markets — to say nothing of the decrease in capital spending across the globe.

We are not in the business of prognostication. We present fact and opinion — however, from all of our various contacts and analysis of accounts from within the industry, 2010 seems ready to offer far more in the way of positive, rather than a negative, business outlook for the U.S. satellite industry. Some of these observations are based upon the following facts...

- » *New technologies and ideas abound*
- » *New companies are arriving on scene with strong financial bases and talent*
- » *New thoughts and tactics are being employed to reduce costs*
- » *Companies are testing and entering new market segments, resulting in growth*
- » *Cooperation between competitors enables all to profit from shared resources*
- » *Increasing global competition for the U.S. satellite market is forcing a re-evaluation of business strategies to ensure prime supplier positioning*

We’ll be with you to highlight the pathways into 2010. Happy days are here again? Perhaps so, with cooperation comes... innovation through motivation!

— *The Editors*

Doing More With Less

Due to budgetary pressures by the US Government, there will be renewed interest in providing new, modernized capabilities through technical refreshment efforts. Upgrading a legacy ground system through a technical refresh offers many advantages over a new-start acquisition. Capabilities can be provided in a timely and cost effective manner. Risk can be reduced by taking a phased approach to adding new, modernized features.

BEAM

A more responsive system hardening capability is possible through on-going information assurance upgrades to counter evolving cyber threats.

Today's flexible, tailored-COTS systems are extremely powerful and enable software-defined solutions to perform the same missions that, in the past, required custom hardware. These upgradable software capabilities not only solve current problems, but also provide flexibility to meet future (currently undefined) needs without the corresponding cost or complexity associated with a hardware upgrade. This not only reduces overall hardware costs, but also simplifies logistics/sustainment maintenance of the fielded capability.

By taking a technical refresh approach, it also enables a new set of players to participate in the development activities. Without the programmatic burdens of a new acquisition, it encourages smaller, more agile technology companies to play a bigger role in solving legacy ground system problems. These companies tend to be more responsive to customer needs while offering greater flexibility and cost

effective solutions compared to the more established, bigger primes.

About the commentator

John Monahan is President of RT Logic, a wholly-owned subsidiary of Integral Systems. RT Logic has been supporting the space industry with cost effective, innovative, software defined solutions. Its products currently support 90 percent of the nation's space programs including GPS, SBIRS, MILSATCOM, TDRS, DMSP/NPOES and others. For more than 10 years, RT Logic has provided the space sustainment community with new ground system capabilities including: modems, receivers, front-end processors, gateway devices, red/black separation devices, and more. RTL has an unmatched reputation for solving the toughest obsolescence issues and developing and integrating next-generation capabilities into existing operational systems.



BEAM

Northern Sky Research



ProtoStar's Satellites – A Good Deal, Or No Deal?

The October 29th sale at auction of the ProtoStar-1 satellite for \$210 million to Intelsat raised a few eyebrows in the industry as some immediately questioned whether Intelsat had overpaid.

Protostar I and Ariane 5 launch vehicle images courtesy of Arianespace

NSR would question those who quickly jumped to this contention or, at a minimum, argue that **Intelsat** has made insufficient information publicly available to independently judge if they paid a fair price for **ProtoStar-1**.

There are of course many issues that must be taken into account when trying to estimate the value of a satellite asset, and the simple fact is that not all satellite operators will reach the same conclusion as to what is a “good deal” or not. Three of the key issues that must be assessed by any operator considering buying an in-orbit satellite follow.

Do You Have A Slot?

While a seemingly obvious first question, this issue will in fact eliminate the majority of potential bidders for an asset like **ProtoStar-1**. Especially given the history surrounding this satellite, any serious bidder has to have solid rights to an orbital slot that would match, as closely as possible, ProtoStar-1’s frequency plan. This includes both uplink and downlink frequencies and, in the ProtoStar-1 case, would also include the satellites transponder capacity that broadcasts in the extended C-band frequency range, as well as the Ku-band transponders, that are broadcasting both in the FSS and BSS segments.

While not exceptionally anomalous, the ProtoStar-1 satellite does not have

a straightforward frequency plan, and a satellite operator would only bid on this satellite if it already had rights to a slot and frequencies that covered, or nearly so, the entire broadcast range of the satellite.

Do The Frequency Coordination + Footprint Match?

Assuming an operator interested in an in-space asset has an adequate slot, then the next issue to address is if the slot and frequencies are coordinated with neighboring satellites. This is both a question of who has precedence to allocated frequencies and, specifically, if the coordination applies to the actual footprint of the satellite. For example, a slot may be coordinated for footprints covering Europe, but if it is not coordinated for a footprint over the Middle East and Africa, even if an operator has a slot, they could have problems selling capacity.

So again in the ProtoStar-1 case, the satellite has more inherent value to an operator who not only has a good slot, but has a good coordination (or precedence in its filings for eventual coordination negotiations) for the specific coverage areas that ProtoStar-1 would have at the location.

Is There Demand For The Capacity At The Location?

The third issue again seems simplistic, but is in fact critical. A satellite operator may

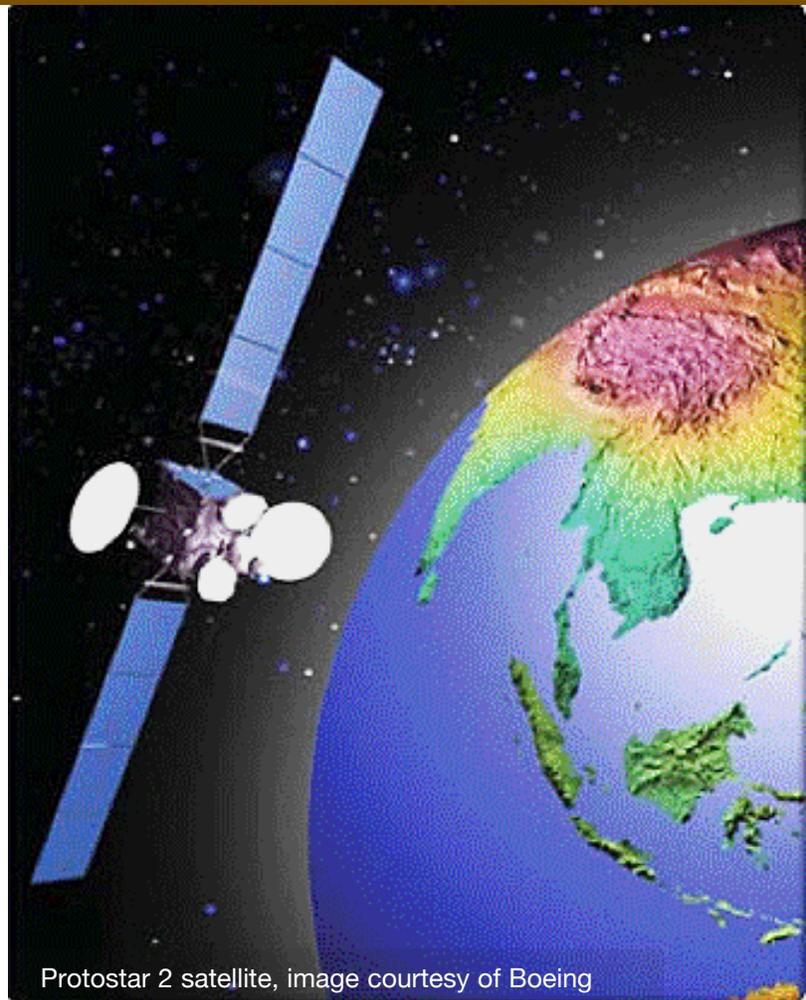
BEAM

have a good slot that is coordinated, but what if they believe there is only minimal demand for the capacity at the slot?

Again this all comes down to the inherent value of the satellite and coordinated capacity that is of no use if there is no one to sell it to. Conversely, if a satellite operator does see demand, then the potential value of the asset goes up. In fact, if we take the case of Intelsat and the reported Atlantic Ocean slot that they plan to relocate the satellite to, it would be reasonable (but not certain) to assume that the main target market is in fact Africa and the Middle East.

Assuming this is correct, ProtoStar-1 could be even more valuable to Intelsat, in net present value terms. This is because they can almost immediately begin serving demand in Africa and the Middle East instead of having to wait the typical two to three years that it would take to put a new satellite in the slot.

This is a good tactical move by Intelsat (again assuming Africa and the Middle East are in fact the target market) because there is currently substantial unmet demand in this region; however, a lot of new capacity is on its way within the next several years. The following chart from *NSR's Global Assessment of Satellite Demand, 6th Edition* study shows that regional average C-band and Ku-band fill rates could drop substantially in the coming years not because of weak demand (demand is growing at a good clip), but because supply in the 2010-



Protostar 2 satellite, image courtesy of Boeing

2012 period will climb up much faster than demand in Africa and the Middle East. By beating its competitors to market with the new capacity offered by ProtoStar-1, Intelsat can grab some of this demand at the prevailing higher capacity pricing before others reach the market and potentially start to push down on capacity pricing as supply increases.

What About ProtoStar-2?

While certainly not the only issues that Intelsat took into consideration when bidding for ProtoStar-1, the above three points do go a long way to explaining why Intelsat won the bid (one of the few

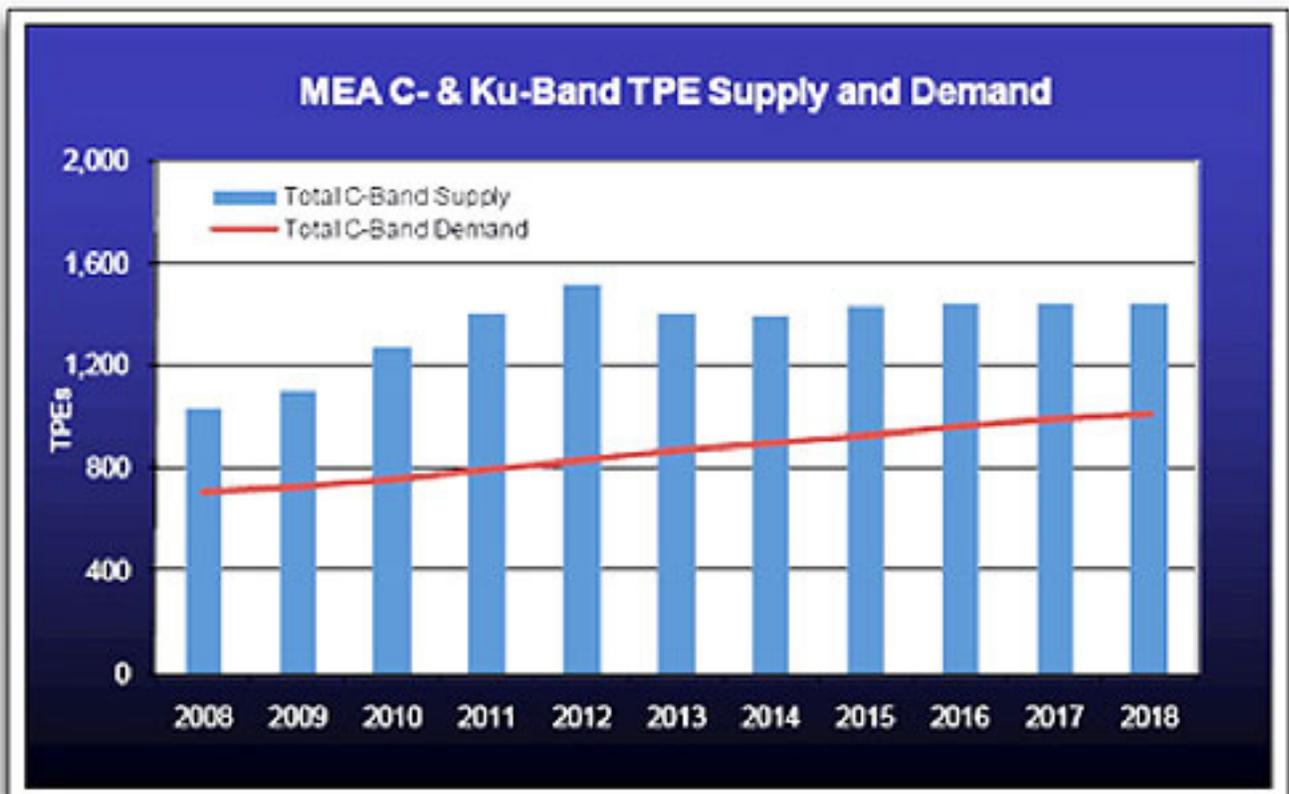
operators with available and coordinated slots) and paid the price it did (and can start tapping demand immediately). The above three points also prove interesting if they are applied to the upcoming **ProtoStar-2** auction. Again, the issue of having an available slot that is coordinated will eliminate a good number of potential bidders. ProtoStar-2 is also relatively unique in that it has an existing S-band client in the form of Indovision making use of the satellite for DTH services.

This complicates the financial evaluation of ProtoStar-2 because any slot too far away from the current location of **107.7 degrees East** essentially zeros out the net present value of the S-band payload since there is very little (if any) demand for FSS S-band capacity anywhere else in the world. And no doubt Indovision would be highly concerned about even a small

move away from 107.7 East because of the cost it would potentially entail if it had to go through a repointing exercise for over 500,000 subscribers.

The above seems to give an advantage to **SES World Skies** in the ProtoStar-2 bidding because, as is well known, they have precedence in Ku-band rights at the closest slot (**108.0 degrees East**). It is NSR's understanding that the **NS-11** satellite currently in the 108.0 degrees East slot does not have substantial, if any, overlap with the Ku-band frequencies used on the ProtoStar-2 satellite. On one hand, this implies that an eventual purchase of ProtoStar-2 by SES could be seen as a straightforward expansion of this operator's capacity at this slot.

Yet, how much demand does SES World Skies see for Ku-band capacity at this



BEAM

location? The fact that the company does not have expansion capacity under contract yet for this location (their issue with ProtoStar was not current interference, but potential future Ku-band interference if they exercised their rights to expand) implies that demand for Ku-band capacity might be low at the moment.

To further support this argument, **AsiaSat's** key satellites are only a few degrees away from the 108 East slot, and AsiaSat has left no doubts about its concerns about Ku-band demand weakness in East Asia. Low demand (if true) means that SES World Skies would likely assign a lower net present value to an eventual purchase of ProtoStar-2, and this would limit how much they might bid for the satellite. And again, moving the satellite away from 107.7 East would lower its value because of the impact on the S-band payload.

All of the above leads **NSR** to speculate that ProtoStar-2 might in fact fetch a lower price than ProtoStar-1. Of course, other issues not addressed here are applicable, and one never truly knows what will happen in auctions. Still, the December auction of ProtoStar-2 will no doubt be an interesting event to watch just to see which issues truly do dominate the purchase decision.

Information for this article was extracted from the NSR report entitled...

Global Assessment of Satellite Supply and Demand, 6th Edition

About NSR

Northern Sky Research (NSR) is an international market research and consulting firm specializing in telecommunications technology and applications. NSR primary areas of expertise include satellite and wireless networks, emerging technology, and media applications.

Our services and clients cover the entire globe. With extensive expertise in all geographic regions and a number of telecommunications sectors, NSR is a leading provider of in-depth market insight and analyses. Since each NSR product is based on future perspectives, our analyses allow our clients to stay a step ahead of the competition and plan for future opportunities in all markets.

Founded in 2000, NSR has gained a reputation as a premier market research and consulting firm. In fact the company is often quoted as an expert in many leading publications.

BEAM



>1

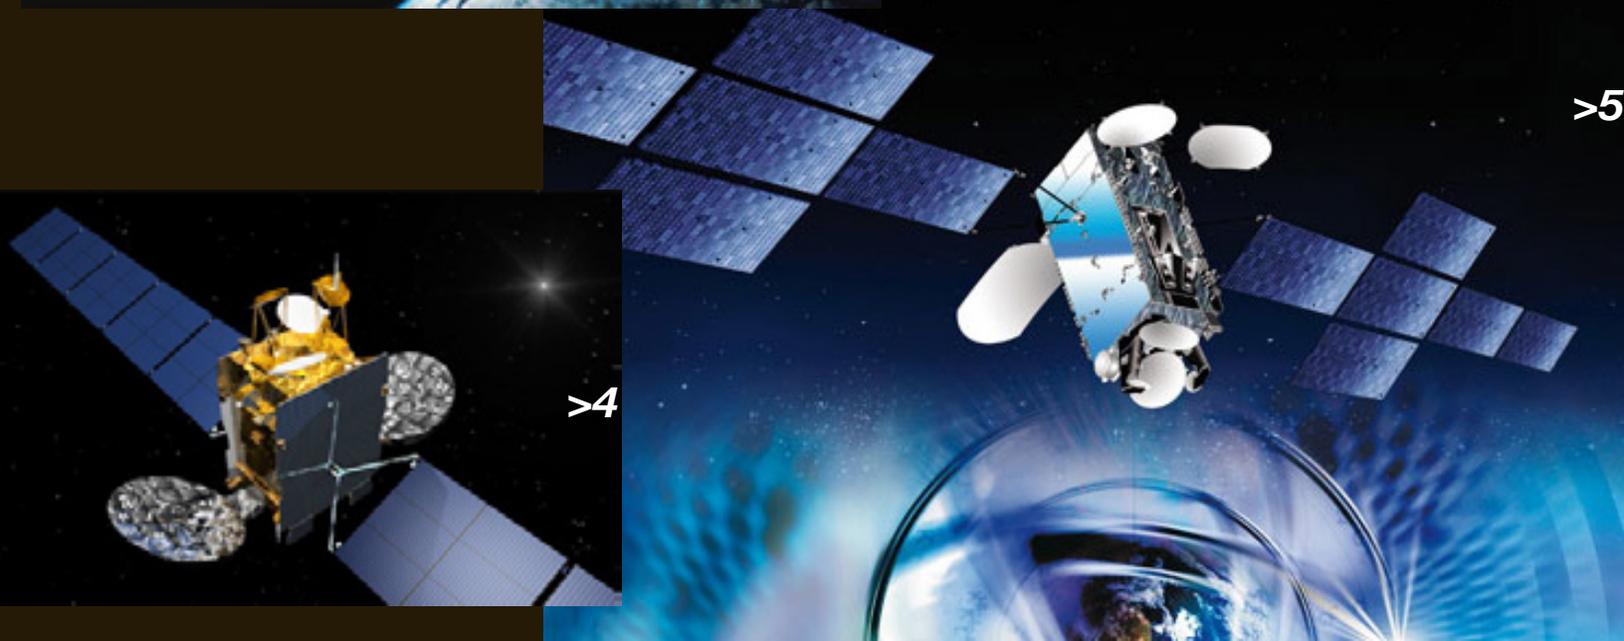
- >1 — Eutelsat fleet
- >2 — W2A
- >3 — KA-SAT
- >4 — W3B
- >5 — EUROBIRD 9A



>2



>3



>5



>4

INSIGHT FORRESTER

Being Global Is Not The Key

Eutelsat's new CEO, *Michel de Rosen*, is settling in well in his new role.

Former CEO *Giuliano Berretta* is now the Paris-based satellite operator's non-executive chairman following his retirement from the company's day to day responsibilities. Mr. *de Rosen*, who enjoyed a long career in the pharmaceutical sector, stressed that in some way he is at a disadvantage compared with 'incoming' chief execs.

INSIGHT

“They are usually hired because the company is in bad shape. Well, that absolutely isn’t the case here at Eutelsat.

looking closely at its options but now that the payload’s problems are fully understood they have been able to

Mr. de Rosen is on record as saying the Eutelsat must not be sucked into a comfortable state of self-satisfaction. “Business journals are littered with stories of companies that, in the belief they were the greatest, were toppled off their perch by someone else trying harder,” said de Rosen.

Giuliano has left us a strong company. But a company’s never good enough. My role and the management team’s role is to seek what we can do better, for instance, where customer service and operational excellence can be further improved and what new partnerships can be initiated. Our task is also to look to the future. Our challenge is to find ways to continue the great momentum we have built for the short and medium term and make decisions now in 2010 that will lead to revenue generators in 2025.”

He outlined a series of possibilities, although ‘being global’ or not was a topic he returned to more than once, of which he’ll address more later. One of his priorities is **Solaris Mobile**, Eutelsat’s joint-venture with **SES**, to see S-band exploited over Europe. The jv’s satellite payload is working aboard Eutelsat’s **W2A** satellite, although with some limitations.

Solaris has banked an insurance cheque for some 133 million euros as compensation for the payload’s S-band antenna failure to meet its specifications. Solaris Mobile is, says Mr. de Rosen,

implement some impressive work-arounds, and thus have time to fully consider options.

“The bankers often tell us ‘you must be global’. I ask them why?” — *Michel de Rosen, CEO, Eutelsat*



“What happened with the satellite was obviously a disappointment. It took time to fully understand the situation and clearly expose it to the insurance community for the claim. The agreement is that they will share a percentage of revenues that may flow from the S-band payload. Our short term intention is to showcase the system, and to continue working with regulators at national and European level to see the licences fully developed. Our friends at Inmarsat, to my knowledge, have yet to be active in the field. Obviously regulators want to know what Inmarsat and Solaris are doing and going to do. It is our common duty to answer clearly and convincingly these legitimate requests.”

Mr. *de Rosen* said that **Solaris** is now re-addressing the whole market, especially in regard to potential partners. “The rate of development of mobile video services in Europe has been slower than anticipated, and the fact that we ourselves have not moved faster than the market is not a handicap.”

But the technical result is now a tangible reality. The transmission system is beaming S-band/DVB-SH video and audio signals to some test vehicles. Customers — or more accurately, potential customers — have already seen the snazzy set-up, and the vehicle is now

INSIGHT

on a true road show to most of Europe's major cities, which will include Barcelona for the upcoming **3G Global Forum**, showcasing what a real system can look and sound like.

Solaris Mobile's next iteration will be a more powerful chip-set allowing potentially dozens of different users (on a bus or a train, for example) to pick and choose their own channel, or data stream. Indeed, it will be obvious to any operator that the concept's ability to handle data

China – Maybe – But Russia Now

Eutelsat and Intersputnik, the International Organisation of Space Communications, have signed a long-term agreement for 16 transponders at Eutelsat's 36 degrees East video neighbourhood. Valid for the entire on-orbit life of W7, which has been co-positioned with W4, it will guarantee long-term capacity for digital TV and broadband services in Russia.

The deal will provide new resources for the Gazprom Media Group, which operates the DTH platforms NTV Plus and Tricolor. With co-location of W4 and W7, Eutelsat has effectively doubled its broadcasting capability at 36 degrees East, which covers Russia and Ukraine.

And it is impressive. A test drive offered up a half-dozen video channels, a few radio services, and most importantly rock-solid images with zero video drop-out or other troubling artefacts.

Moreover, Solaris' clever technical solution beams a mini-WiFi signal into the car, enabling up to three different images to be viewed (including your author's *iPhone*). It was seamless to set up.

and other vehicle-based telematics is a key advantage.

Currently, the scheme assumes the satellite will beam its couple of Muxes to Europe's major regions with terrestrial repeaters filling in the local urban gaps but also complementing the core channel offerings with, perhaps, local traffic, weather or other personalized information.

Solaris is also working on a battery-driven ‘dongle’ receiver that will have a similar chip-based battery driven device that will transmit content to mobiles without an on-board receiver. The Paris test uses four, 30-Watt transmitters to cover the city. Solaris’ key target markets, remembering the problems with the antenna, are currently France, Germany, Italy and Spain.

Ondas Media is seen by most as being a strong contender in this space and already has contracts signed with **Renault**, **BMW** and **Nissan**. Interestingly, the specification also allows SMS-type short messages to be sent directly from a vehicle up to the satellite. Transmission power is just 1-Watt, but the technicians say it all works fine. Solaris is looking to see standards body ETSI establish an open standard for this aspect of the technology.

Mr. *de Rosen* has strong praise for former CEO Berretta’s time at Eutelsat. “My first steps into this building were in June last year, and it didn’t feel in any way

like a bureaucracy. I commend [former CEO] Jean Grenier and Giuliano for their work in transforming Eutelsat’s culture. They were really able to transcend the old vision focused on telephony into a vibrant, video-focused business with high productivity. Not all companies have this

enterprising spirit and it is a remarkable achievement. The commercial, financial and operational metrics of a private

da Vinci of the satellite world, you are better to build and rely on a strong team of professionals who are both experts in

And More Biz For Africa

Eutelsat signed a fresh contract with Multichoice Africa that covers 20 transponders on Eutelsat's W7 craft (36 deg E). It is one of Eutelsat's largest-ever deals, and represents a new 6-transponder extension to the current contract, and also a time extension. Multichoice already uses Eutelsat's W4 craft at 36 deg E, and services are being currently transferred to W7. This 15-year agreement extends the relationship to 2025.

With the move to W7, MultiChoice Africa has consolidated all of its Ku-band pay-TV activity outside Southern Africa onto one Eutelsat satellite. In addition to benefiting from better signal levels and improved coverage for all the regions already served from Eutelsat's 36 degrees East neighbourhood, MultiChoice Africa has extended its Ku-band service to additional countries, notably Ethiopia, which has been receiving DStv in C-band, requiring large domestic antennas.

company are well anchored at Eutelsat. Who would ever have spoken about debt/equity ratios all those years ago?"

"I have some management principles which include a strong belief in drawing on the skills and experience of people in the company. We are lean, dynamic and creative and I know already we can move fast to adopt a good idea. Unless you are Giuliano, whom I call the Leonardo

their field and leaders. And you want this group of people to be a real team, that is to say to work together like a rugby team. I do not want people to start focusing on us being the biggest. That's not necessary and can be a distraction. But I want us to be the best. Bankers often visit us and say 'you must become global and we can help you make that happen. I can see why that would be good for bankers, but in my six months at Eutelsat I have yet to hear

a compelling reason for Eutelsat to be global in terms of creating real value for the company and its shareholders.

He said Eutelsat would stay Number 1 in its large region and he wanted to remain as number one, even 'super number one'. "So I say to myself, beyond our core regions of western, central and Eastern Europe, the Middle East and Africa, where we are present, active and ambitious, where else should we go? We look at size of the market, growth rates, current prices, potential price evolution, competitive landscape and of course realistic opportunities. With such an approach, the East looks more interesting than the West. Obviously, we have more homework to do."

De Rosen emphasized that Eutelsat's number one business is and will be video, where he sees significant growth potential going forward with existing customers, new customers and new technologies such as HD and 3D. However, he spoke enthusiastically about the upcoming launch of **Ka-Sat**, later this year. He said what Eutelsat was increasingly hearing from the market is to believe in the potential of the Ka-band. He said their discussions with **Tooway** distributors and clients — and governments — show clear evidence that people need satellite resources for broadband.

"We know we have to be excellent technically and commercially with Ka-Sat. When I look at our situation today, I can tell you we are exactly where we wanted to be at this time. We are active in all our key markets, but I intend to increase the

pressure on our colleagues to do more, so that we are fully ready operationally when the satellite becomes available.

While being perhaps deliberately vague about his immediate plans geographically, *De Rosen* spoke enthusiastically about China, India and the Asian region generally. He and Mr Berretta had recently visited China together and been very well received, he added.

He commented this is the result of years of efforts and collaboration between Eutelsat top management and Chinese authorities. But he said there is more work to do to in order to decide what is the best way for Eutelsat to build a real presence in Asia and to then implement a plan. Eutelsat, amongst many, was interested in the **ProtoStar-1** auction, for example. He also anticipated consolidation from Asia, notwithstanding the sluggishness in many individual markets and surplus of capacity leading to low transponder rental prices.

Mr. *de Rosen* enters 2010 with Eutelsat in good shape. Last year's 940 million euros of revenues looks certain to pass through the 1 billion euros mark this year. Video revenues, as well as helping provide juicy 79.9 percent EBITDA margins, also represent almost 700 million euros of revenues and are still growing at 4.7 percent.

But data and broadband traffic is also steaming ahead, with 13.4 percent growth last year, helping bring in revenues of 173 million euros. From a more modest base the company's multi-

usage clients-generated revenues of 75.5 million euros and an impressive 30 percent growth rate.

But it is video that represents Eutelsat's core business, some 3,200 channels and growing by around 200 a year. That growth looks certain to continue, helped by HDTV and new hot topic — 3DTV.

The number of HDTV channels grew 76 percent over the year. "With the diversity of our businesses and geographic presence, we have many engines for growth". Eutelsat's current guidance is that its EBITDA margin would be above 77 percent for the next 3 years, to 2012. Further updated guidance is due next month.

Questioned on his optimism for Ka-band business ahead of the launch of Ka-Sat later this year to Eutelsat's 'hot bird' position, he remained firmly of the opinion that it would work in terms of business and generate solid revenues by the mid-term life of the satellite. Mr. *de Rosen* interestingly joked that even some of Eutelsat's rivals were privately more interested in Ka-band than they publicly admit. He is on record as saying that, as

with Solaris Mobile, there would be times when there was no need to own 100 percent of a satellite, and his statement included Ku-band as well as Ka-band.

(Editor's Note: The March issue of *SatMagazine* will offer a number of articles dealing with SatBroadcasting, including 3DTV, and Chris' exam of the latest 3D happenings at CES, which occurred earlier this year in Las Vegas.)

About the author

London-based Chris Forrester is a well-known entertainment and broadcasting journalist. He reports on all aspects of the TV industry with special emphasis on content, the business of film, television and



emerging technologies. This includes interactive multi-media and the growing importance of web-streamed and digitized content over all delivery platforms including cable, satellite and digital terrestrial TV as well as cellular and 3G mobile. Chris has been investigating, researching and reporting on the so-called 'broadband explosion' for 25 years.



INSIGHT

RADFORD

ew Or Used?

The Chronicles Of

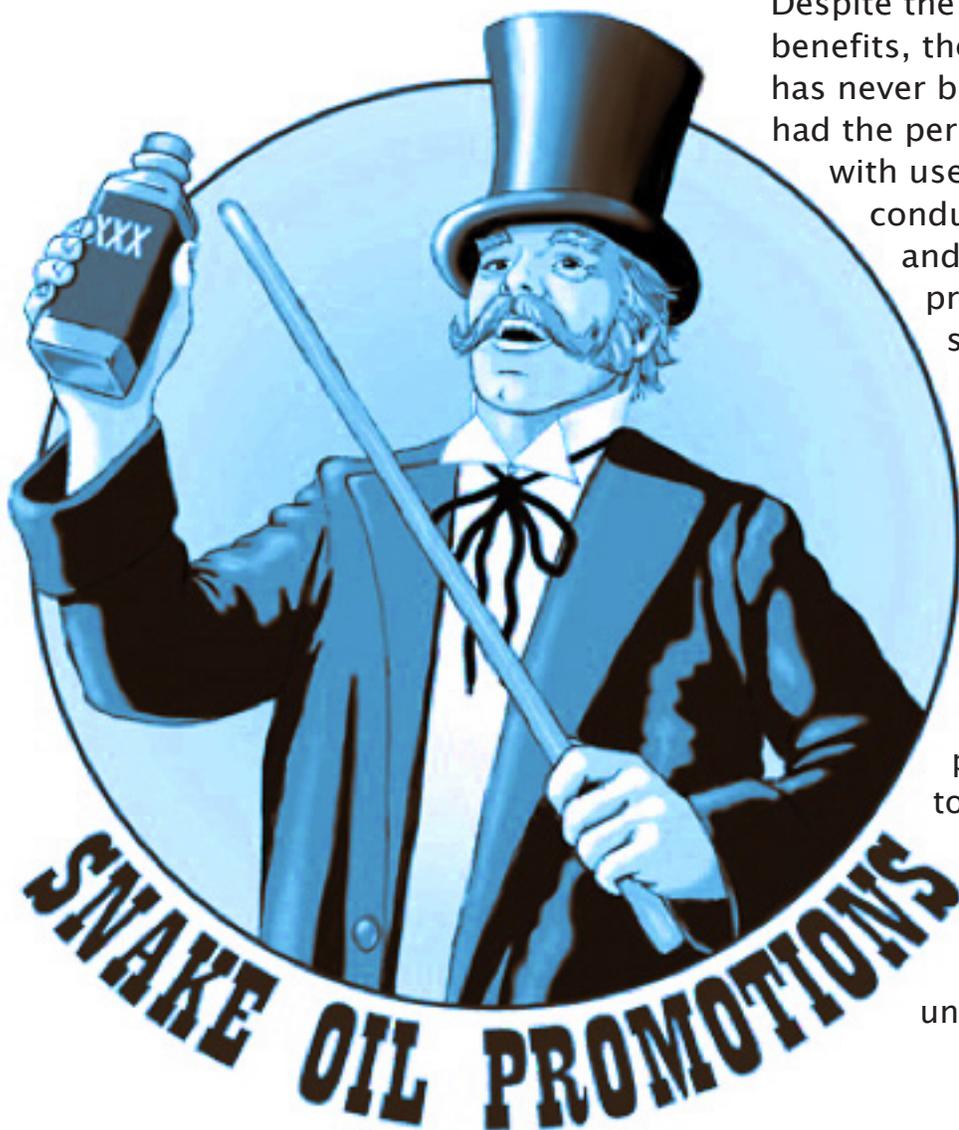
SATCOM

Long before the term “recycle” became the buzzword of the times, peddlers of pre-owned satcom equipment were scouring the planet in search of post-warranty treasure — bargain-priced system components scarred from years of faithful service, but with some operational time presumed to still be left on the clock.

INSIGHT

Vast inventories of potential harvests were created by the widespread demise of teleports unable to survive fierce competition and turbulent markets, and like hungry buzzards circling high above the parched carcass of a wandering desert steer doomed to a drink-less fate, resourceful brokers would lie in wait for the last motor, blower or fan to udder its final gasp — a death knell preceding the inescapable finality of liquidation.

But with MTBFs calculated to be in the millions of hours (by the manufacturers of course), satcom products had a propensity to outlive the life spans of the businesses for which they were originally purchased. This created the opportunity to resurrect dormant craftsmanship so it could once again generate precious revenue, breathe life into a new business plan and yield huge savings for architects desiring to build Earth stations on the cheap, or in days instead of months.



Despite the promise of these obvious benefits, the adage “buyer-beware” has never been more applicable. I’ve had the personal pleasure of working with used product brokers who conducted business with integrity and with the utmost in ethical practice, but history has shown that evil lurks in the shadows of this seemingly obvious value proposition.

Years ago, I fell victim to an unscrupulous broker, deceived, oversold, and ultimately betrayed, by a slick-talking troll I mistook for a charismatic purveyor of pre-owned bounty. This fork-tongued serpent possessed an inane talent for looking one straight in the eye and spewing words completely void of truth and rife with undeliverable promises.

Mesmerized by a sales pitch worthy of *Dale Carnegie* himself, this convincing demon lured me down a long dark path of despair — a path paved with financial disaster and emotional distress.

Being a system integrator at the time, I had the immediate need for a 9 meter antenna system that was on my list of deliverables for an Earth station we were constructing for a new customer. Facing schedule pressure to meet a critical uplink, the long lead-time for a newly-manufactured antenna threatened our committed timeline with impending disaster. Boy, was I was in trouble!

While perusing the advertisements in a popular industry magazine, I came across an ad promising huge inventories of quality, used equipment ready for immediate delivery. Being new to this concept, I was skeptical to say the least.

But a guided tour of the serpent's lair revealed a seemingly inexhaustible supply of de-commissioned Earth station components

piled about like Lego blocks ready to be assembled into instant systems. I could actually visualize the end to a history of being slave to equipment manufacturers' long deliveries and high prices.

INSIGHT

Like an oil-soaked puffin flapping helplessly in the wake of the Exxon-Valdez and drifting aimlessly into the arms of an awaiting PETA volunteer, I was saved from the scorn of an unhappy customer under immense pressure to get a satellite link established and identifying me as the only remaining hurdle. But was I truly saved? Hardly!



“Rescue A Relic”

Though delivery of the equipment to the site was expedient as promised, an 8-meter mount, 9-meter reflector and 11 meter feed equals one unusable pile of scrap metal. To quote the installers on site, “**the bolt-holes don’t line up!**” With tail-tucked and white flag in hand, we appealed to the antenna manufacturer’s sympathy and received an expedited, though deathly long delivery.

In the end, all was made right. The critical broadcast was facilitated by a contractor (at my expense), my customer received a new antenna, my company lost money

and the broker slithered back under the rock from whence he came, never to be heard from again. Though permanently scarred by the resulting experience, the trap into which I fell was indicative not of a market, but of an individual with a demonic soul sculpted by forces of evil. But as they say, any experience can create a positive result.

The only way my personal tragedy could be entered into the **Chronicles** was to have lived it. A testimonial intended to expose the potential mine field one might be forced to navigate should the necessary due-diligence not be done prior to taking the “used” route.

The moral of this story is simple. The availability and brokerage of quality used equipment has a distinct and quantifiable value to the industry. When purchased through a reputable broker, great deals can be had, schedules reduced and savings realized. Quality brokers are out there vying to keep the Earth’s landfills free of satcom’s pre-owned treasure. Do the world a favor — rescue a relic. After all, wasn’t analog really more fun?



Tony Radford



TRIAGNOSYS

The Future Of Air Traffic Management Communications

Notwithstanding the current economic climate, air transport is forecast to double by 2025 according to the latest Eurocontrol figures. And it is abundantly clear that current air traffic management (ATM) systems around the world will simply not be able to cope. To address this issue, as well as issues of safety and security, there are a number of

FOCUS

strategic government-funded projects currently running, whose remit is to develop the most appropriate and cost-effective new technologies and techniques for the future needs of ATM. These include the *Next Generation Air Transportation System (NextGen)* in the US and *Single European Sky ATM Research (SESAR)* in the European Union.

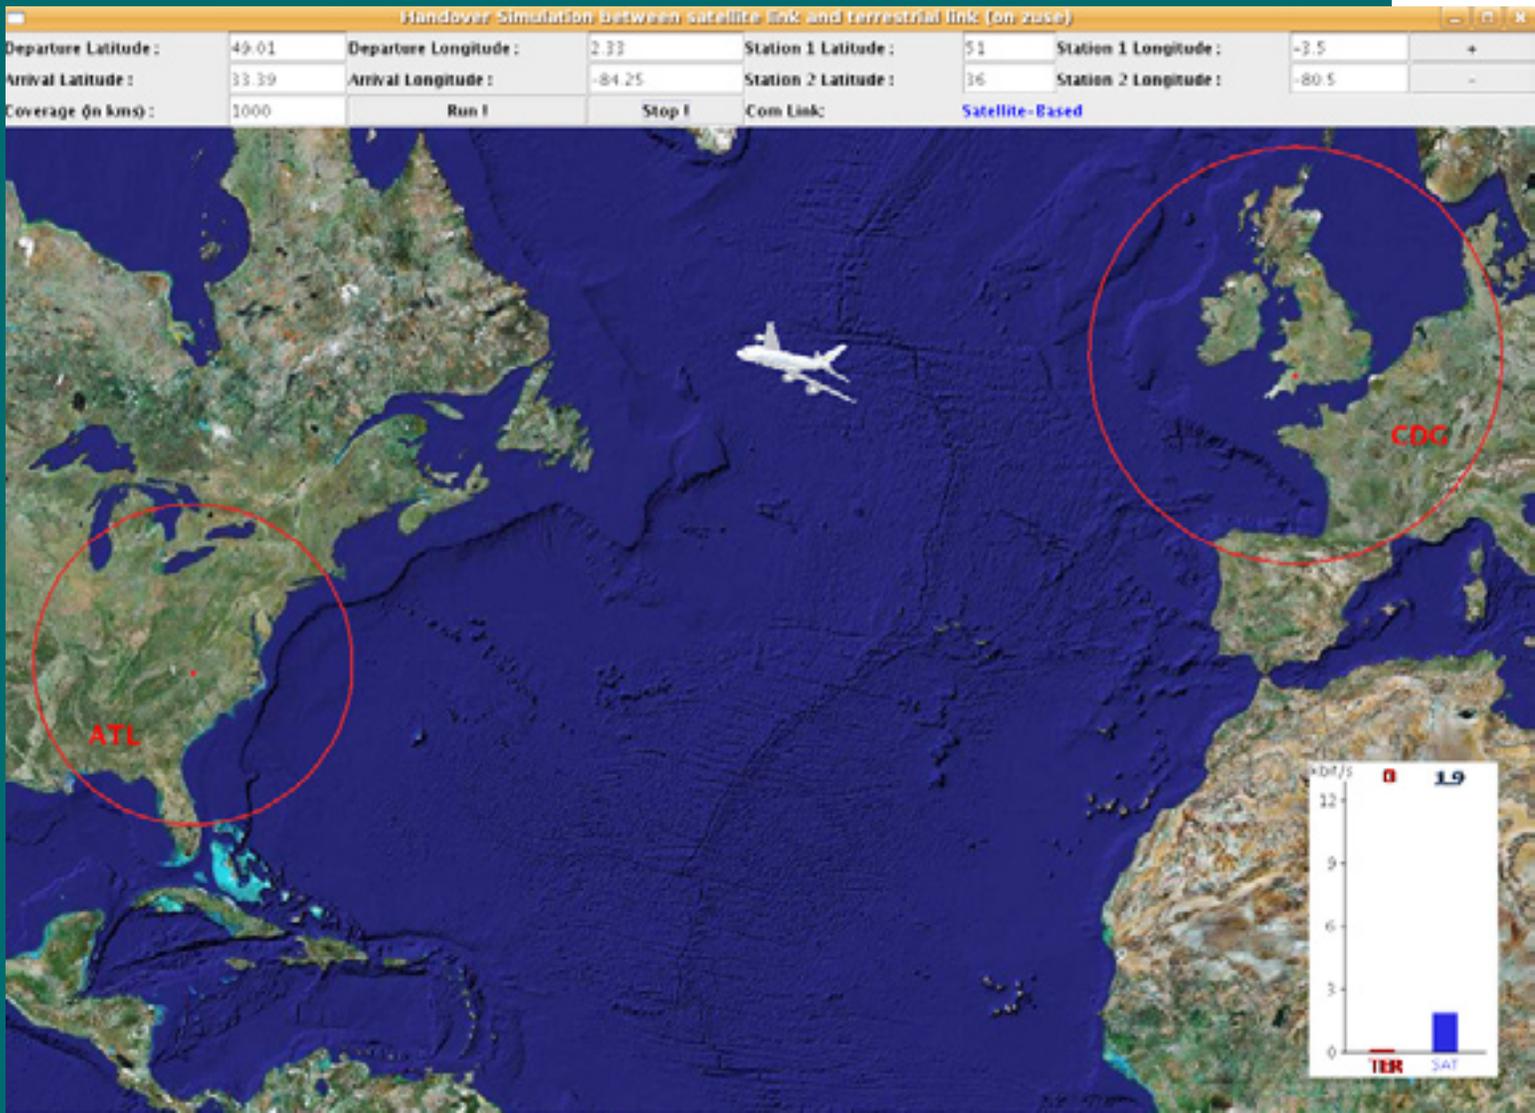
The overarching objective of these projects is the more efficient use of available airspace, without compromising on safety. A key enabler of many of the required changes is improved communications between air traffic controllers and pilots. That inevitably

will mean increased, and smarter, use of satellite communications.

Communication between controllers and pilots has, for many years, used VHF, or for longer range HF, voice communications. Both have limitations that will only be exacerbated as air traffic volumes increase. HF, for instance, does enable long-range communication, but certainly does not provide the comprehensive global coverage of which satellite communications is capable.

An additional limitation of the existing technology is that as controllers currently use a single frequency, and when a sector is very

Figure 1: Visualization of the NEWSKY seamless satellite/terrestrial link handover demonstration



busy, many pilots will be listening to the same signals, which can, and occasionally does, lead to confusion.

The latter limitation is already being addressed by the introduction of *Controller Pilot Data Link Communications (CPDLC)*. It enables air traffic controllers and pilots to communicate using pro forma data messages, to avoid confusion and save time, though of course the system also allows for free text entry when necessary.

NEWSKY — NexGen ATM

The next generation of ATM technology will be based on Internet technology, the use of both satellite and terrestrial communications, and the interface between them. This is an area in which the European Union is taking a lead, in particular through **NEWSKY**, an EU-funded research project that has recently come to a conclusion, and in which **TriaGnoSys**, the expert in satellite-based communications and information transfer, was responsible for the development of the relevant satellite and terrestrial communications solutions, and the handover processes between the two.

As proof of concept of the NEWSKY project, a simulated flight was conducted between Europe and the US at the project's final presentation in October 2009 at the German Aerospace Centre. The three phase simulation — over European, oceanic and North American airspace — involved testing both

voice and data connectivity between pilots and controllers, to deliver continuous, seamless communications throughout the three key phases of the flight. Weather maps and CPDLC-like data messages were also transmitted over the links.

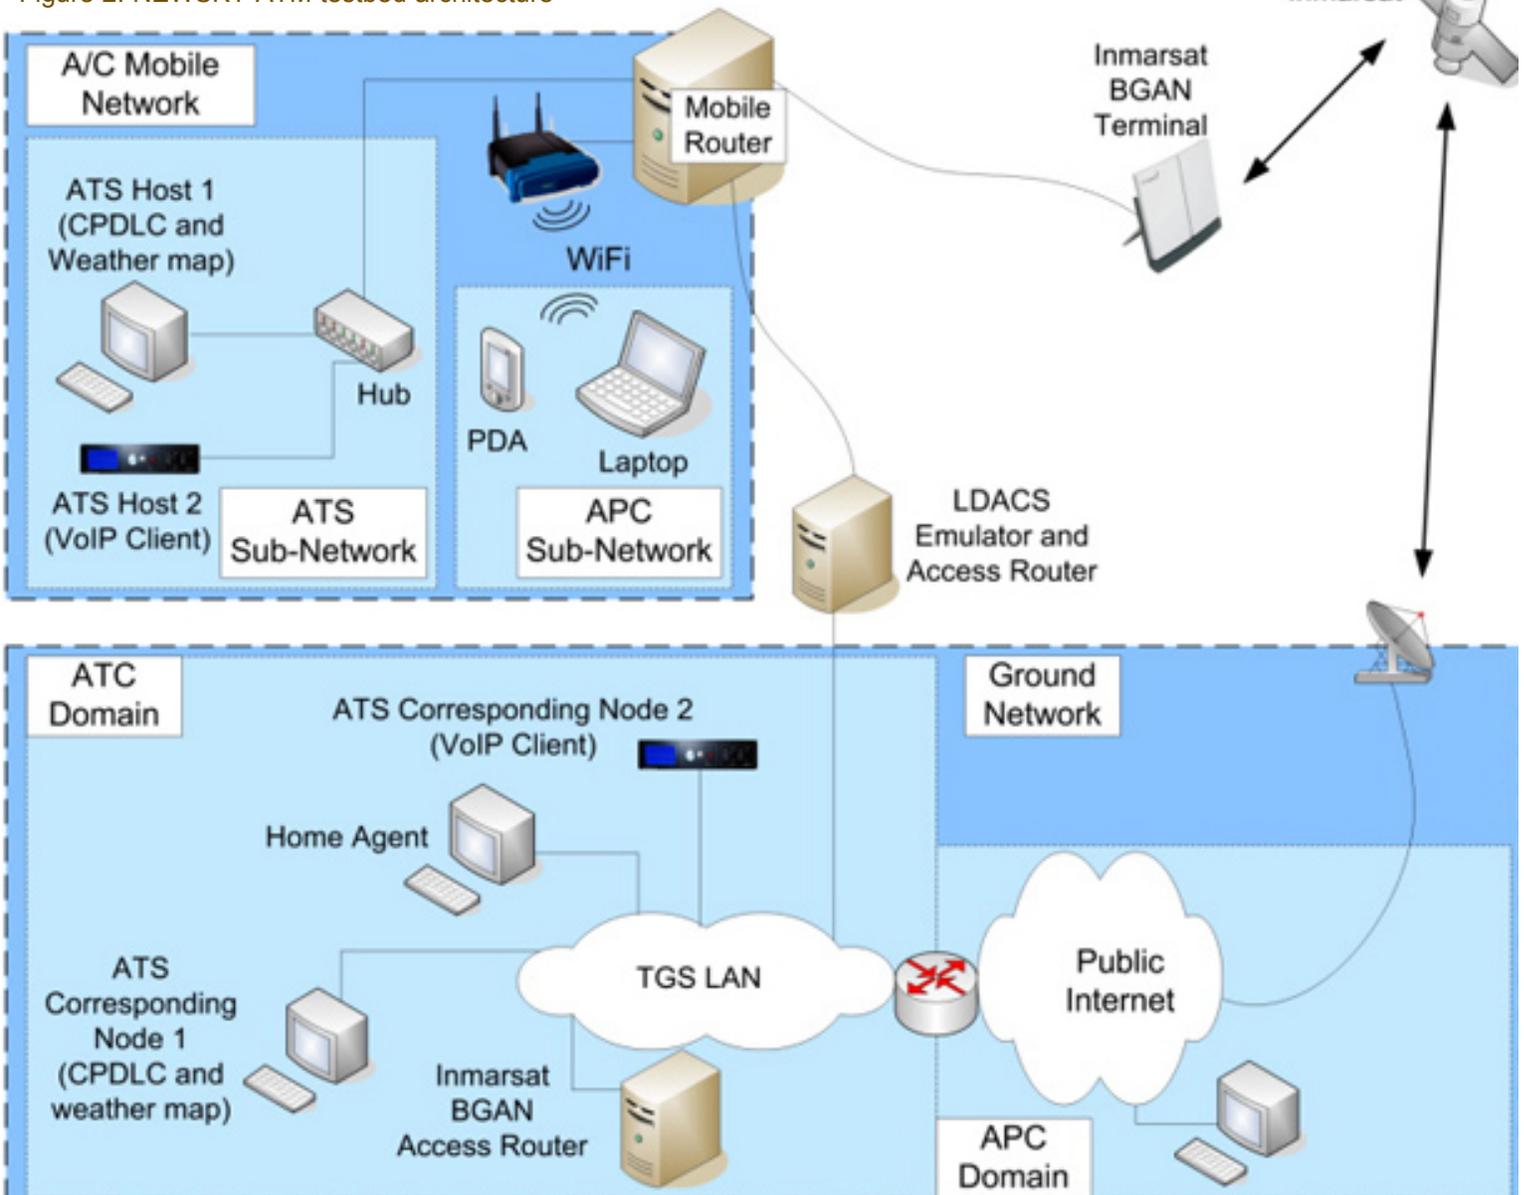
FOCUS

» *The first phase was in European airspace. Both terrestrial and satellite (Inmarsat BGAN, in this case) radio links were available and either could be selected, based on the programmed link selection rules. The rules are based on a number of factors, including the availability of a specific satellite link, efficiency of the terrestrial radio link — which depends on, among other factors, the weather — and indeed cost.*

» *The second phase of the flight was over oceanic airspace, where only satellite connectivity was available. The simulation demonstrated the seamless automatic handover from the terrestrial network without loss of session continuity.*

» *The third and final phase was over North American airspace. As with the first phase, the link selection rules governed which link to use, again with automatic handover.*

Figure 2: NEWSKY ATM testbed architecture



ATM technology is central to the NEWSKY research. The full name is '**Networking the Sky for Aeronautical Communications**', and the concept was the integration of a range of applications, using Internet Protocol version 6 (IPv6) as the basis. In addition to the key air traffic control (ATC) communications, other applications include airline operation communications, airline administrative correspondence, and passenger communications.

One of the key considerations for the use of IP is that it combines high reliability, cost savings, and an optimal alignment with the evolution of communication and security technologies through the use of *Commercial-Off-The-Shelf (COTS)* solutions. Also, the focus has been on a move away from proprietary software, to allow wide access. In addition, since this system is based on IPv6 for future-proofing, it is backwards compatible with IPv4. In short, every business uses IP, making it the obvious candidate for the next generation of ATM communications.

Cost remains a key factor in the development of air traffic technology, both for the ATMs and the airlines. Of course, it is not as important as reliability, and safety is the overriding consideration — but cost/benefit analysis is always part of the decision-making process in both the public and private sectors.

Until now, the benefits of using satellite communications for ATM have been outweighed by the costs. However, Voice Compression and Enhanced Multiplexing (VoCeM), developed by TriaGnoSys, which provides the most efficient satellite use of satellite payload, brings satellite usage costs into line with terrestrial radio links.

VoCeM, as the name suggests, cuts the bandwidth required for voice calls, be they Voice over IP (VoIP), or indeed GSM — and for NEWSKY, it is VoIP. To give an idea of the scale of savings provided by VoCeM, VoIP generally consumes 20 kbit/sec, and GSM, which has attracted so much attention, is actually even more inefficient with bit rates up to 30 kbit/sec. Using VoCeM, the requirement for VoIP is only 3.3 kbit/sec, including all overheads. To put it another way, VoCeM enables compression gains of between 500% and 1,000%, and does so through advanced header compression, voice transcoding and channel multiplexing techniques.

The three main elements of the technology are:

- » ***First, the voice payload is transcoded for transmission with a new AMBE2+ codec (Advanced Multi-Band Excitation), providing significant bandwidth efficiencies without compromising quality.***

Figure 3: Cost Savings Using VoCeM

	L-Band	Ku-Band
Call cost per minute uncompressed	US\$2.25	US\$0.675
Call cost per minute using VoCeM	US\$0.30	US\$0.10

FOCUS

- » ***Secondly, each header is compressed to the minimum (headers are the 'instruction manuals' that enable the receiving servers to reconnect the information in the correct way). VoCeM removes all redundant headers and compresses the useful ones.***
- » ***Finally, VoCeM uses multiplexing to bundle a number of voice payloads together, to form a single packet, further reducing the headers required.***

The key benefit this technology provides is cost saving. Based on current satellite costs, the table below shows the savings VoCeM can provide:

Furthermore, quality is not compromised. Telecoms providers use a scale of zero to five to measure Quality of Service (QoS). ISDN telephony typically achieves 4.2 or 4.3 and a full-rate GSM call is about 3.5, and VoIP 3.7. Using VoCeM, we achieve rates of 3.7, even though it demands only a fraction of the bandwidth.

Satellite communications will play an increasingly important role in ATM, particularly where terrestrial links are not available. As efficiency and safety demands increase, with it becoming unacceptable for aircraft to be out of touch at any point during the flight, and satellite costs are reduced, ATM services using satellite links are the future.

About the author

Dr. Markus Werner is one of the founders of TriaGnoSys. He received a Dipl.-Ing. Degree from Darmstadt Technical University in 1991 and a Ph.D. degree from Munich Technical University in 2002, both in electrical engineering. Markus received the Best Paper Award of ITG conference "Mobile Communications" in 1993. He has co-authored more than 120 publications, including two scientific textbooks and numerous scientific journal papers, and teaches satellite communications courses for telecommunications professionals.



EVENT

CABSAT MENA 2010

valuable + Free To Attend!

This important broadcasting conference — ‘Higher Value through Digital Broadcasting’ — runs from March 2nd through 4th and will provide high level information and real-life experiences for those involved in the digital HDTV, enhanced digital radio, and file based production industries. Speakers from MENA and around the world will offer attendees their invaluable expertise. This is a platform for broadcasters and other industry players to exchange views and experiences on launching and supporting HD services.

EVENT

Show highlights include...

- » *New Compression Technologies*
- » *Building Network Infrastructure*
- » *File-based Production Workflow*
- » *File-based Production for Radio*
- » *Implementing Digital Broadcasting: Network Planning*
- » *Higher Quality Services in Radio Broadcasting*
- » *HDTV Business Strategies in the Middle East*
- » *Implementing Content Protection Technologies*
- » *Status of Digital Broadcasting, DTH and IPTV in the MENA region*
- » *Transmitting Facility Infrastructure*

Also free to attend during this conference is the **GVF MENASAT Summit — The Middle East and North Africa Satellite Summit**, which is themed *Development Dynamics in the Broadband Satellite and Hybrid Wireless Applications Market*.

The GVF Summit, held on March 3rd, brings together industry leaders to identify and discuss the development of new opportunities in satellite and satellite-hybrid services and technologies. The Summit will offer interactive exchanges of ideas between communication technology users, operators, manufacturers and analysts, making this an unrivalled networking opportunity for the MENA region. The Summit highlights include...

- » *The MENA Communications Market Panorama*
- » *Understanding Unified Value Chain Opportunities*
- » *MENA Industry and Government as Satellite Service Customer*
- » *Application Focus*
- » *Regional End-User Vertical Market Focus*
- » *Technology Focus*

Since 1995, **CABSAT MENA** has become the third largest event in the world and the region's most important event for the digital media industry. CABSAT MENA is the ideal platform for companies wanting to gain an edge over their competitors, all the more so in the MENASA region, face to face relationships are critical to successful business dealings. The event is a showcase for the latest products and services from the entire spectrum of digital media from content creation to management to delivery.

Satellite MENA was launched in 2009 in response to the increasing demand of the satellite industry. **Satellite MENA** is the first ever definitive event of its kind in the region, serving broadcasters, commercial entities, telecoms operators, government, NGO's, military and many more.

Satellite MENA aims to be the industry networking platform to explore business opportunities, formulate strategies and create partnerships for video, voice, data and IP communications over satellite.

EVENT

Australasia Oil & Gas

Perth Premiere Industry Exposition

Staged in the city of Perth, AOG 2010 is held annually and has run for over a quarter of a century and is Australia's largest international petroleum expo, showcasing the latest technical challenges and solutions connected with oil and gas exploration, production, and transportation. The entire spectrum of the industry is covered, from exploration and production technology to gas processing, pipelines, communications and more.

EVENT

AOG is the place –

- » *To source new products + services*
- » *To compare new products and supplier offerings*
- » *To ensure that you select the correct products and services for your business*
- » *To keep abreast of the latest industry movements — if it's hot, it's here. From new product launches to free seminars on the show floor — AOG is a showcase for new innovations*

With more than 400 exhibiting companies in attendance this is your opportunity to assess all solutions in an impartial environment.

Plan to attend the *Subsea Australasia Conference, Flow Assurance, Health & Safety or Carbon Storage Seminars.*

Network with your peers at one of the many social functions or on the show floor in the Oyster Bar, Subsea Bar, or Bratwurst Bar

At **AOG 2010**, more than 7,000 like-minded individuals will meet to exchange ideas with their peers. One of the satellite communication companies participating at this exhibition and conference is **PacTel International**. The director of the company, *Andrew Taylor*, believes this is a great opportunity for the Company to present their range of telecommunications solutions for the Oil and Gas sector. Pactel International will be located at **Stand Y22**. Appointments can be made prior to the show to make the attendees' visit as efficient as possible.

Pactel International was co-founded in 2003 by Andrew Taylor and Steffen Holz, both with more than 20 years of experience in the satellite and telecommunications industry. Andrew worked in various commercial and engineering roles at PanAmSat, Comsys, Optus, and Telstra.

Steffen started his career in satellite communications in the Pacific Islands. In 2000, he founded Pacific IP Services in New Caledonia, which merged into Pactel International.

Pactel International focuses on satellite telecommunications and broadcast services within the Asia-Pacific Region. Company services include Internet backbone connectivity, VSAT data solutions, VOIP gateways, broadcast and streaming video solutions, international private leased line, equipment hosting as well as satellite ground system and network design.

The Company focus is on high quality, customized and affordable telecommunications solutions for the telecommunications carriers, ISPs, government organizations, Mining, Oil & Gas industries, broadcasters and businesses that have a need for dedicated solutions to meet their connectivity challenges.

With the head office in Sydney, Australia, and regional offices in the Pacific Islands, Pactel International offers support to all of their customers who live and work in similar time zones. Pactel International works closely with their customers and adopts a hands-on approach during the design, implementation, and delivery of all solutions.

The Company implements fully customized solutions to match their customers' specific telecommunications needs; from point to point links through to fully managed networks — all receive the highest quality solutions at extremely competitive rates. What sets Pactel International apart is their expertise in using satellite-based technologies to provide telecommunications solutions at highly competitive rates.



FOCUS

INTELSAT

ombattling Satellite Interference

For decades, the satellite industry has achieved steady growth by meeting end-user demand for reliable, cost-effective communications for a wide variety of applications. Each day, thousands of corporations, broadcasters, government agencies and other organizations worldwide depend upon immediate access to satellite networks for Internet

FOCUS

connectivity, video conferencing, LAN/WAN interconnections, distance learning, point-of-sale transactions and VoIP.



Over the past several years, *Radio-Frequency Interference (RFI)* has become a critical threat to satellite communications. The many RFI incidents reported each year create obstacles for daily operations. Ultimately, end users pay a high price as RFI disrupts television signals, data transmissions and other customer services — hindering business growth.

To promote an interference-free space environment, **Intelsat** recently launched the *Intelsat Interference Management Initiative (I3)*. As the world's leading provider of fixed satellite services,

Intelsat's I3 team is driving a far-reaching effort. The initiative will ensure technicians obtain proper training. It

will increase the quality of data and communications regarding RFI events within the satellite community. It will also guarantee proper alerts are implemented to reduce reaction times when RFI occurs. This initiative includes industry-wide activities to reduce RFI on a global level, as well as Intelsat-driven activities focused on enhancing customer and operation-center expertise.

The first step toward maintaining an interference-free space environment is to actively provide the means

to train the user community. This will ensure they have the technical knowledge to responsibly access global satellite networks, without creating unintended interference for other users. To meet this goal, Intelsat is endorsing programs to train VSAT installers and satellite newsgathering (SNG) field operators.

Training For VSAT Installers

For VSAT installers, Intelsat is endorsing training from the **Global VSAT Forum's (GVF)** training and certification program. The GVF courses educate technicians on proper equipment installation and operational parameters of VSAT networks.

Intelsat's goal is to train 1,200 customers through 2011, free of charge. Additionally, Intelsat negotiated a discounted rate to encourage all Intelsat-customer technicians to be trained.

Likewise, GVF is providing a 50 percent discount on training for technicians who work for organizations based in developing countries. The discounted rate is part of the *Andrew Werth-GVF Scholarship Program*.

The GVF's *VSAT Installation Certification* training program has served as the industry-standard since 2002. The program has been continually updated and is now delivered via a combination of online, interactive, simulator-driven training modules and formal hands-on skills testing — all managed through the GVF training portal at www.gvf.coursehost.com.

Hands-on skills testing and supplementary classroom sessions are supported by GVF instructors and regional training centers worldwide. The training program has recently been upgraded to specifically reinforce core

skills that field technicians need to prevent accidental uplink interference.

The new **GVF510** (*Core Skills for Satellite Technicians*) online course includes intensive instruction and practice using 3-D simulators. Upon completion of

FOCUS

the course, and demonstration of their hands-on skills, technicians will receive the new ***GVF Basic VSAT Installer Certification***. The certification indicates that a technician understands the importance of interference mitigation and has the skills to minimize interference when installing a VSAT.

Training For SNG Field Operators

As part of Intelsat's customer-driven initiative, we also have purchased a number of sessions on **BeaconSeek's SlingPath** SNG online training program. After evaluating a number of training programs against Intelsat's technical requirements, we selected SlingPath as one of the I3's approved training programs. Over the years, employees and contractors of Fox News, ABC News, Al Jazeera and Arqiva, among others, have successfully completed the SlingPath training. We believe SlingPath is ideally

suited to meet the complex technical requirements of SNG field operators. Similar to the GVF program, Intelsat will offer this training, free of charge, to our customers. We will also continue supporting our customers by facilitating this training course at a reduced rate.

Improving Industry-Wide QoS

In addition to comprehensive customer training, Intelsat's I3 strategy includes a wide range of industry initiatives. These include the establishment of a robust Carrier ID System, a Data Sharing program and a Satellite Operator RFI Alert Network.

Carrier ID System

Intelsat, along with other satellite operators and the ***Satellite User Interference Reduction Group (SUIRG)***, is working with hardware manufacturers to implement a Carrier ID system. The

system will ensure each signal uplinked to a satellite has an embedded identification “signature” — similar to an electronic fingerprint. This signature will provide the emergency contact information of the interference source to satellite operators, allowing satellite operators to quickly resolve the interference incident.

Data Sharing Program

The open exchange of operational data is imperative for numerous critical satellite operator procedures, including interference ID, analysis and RFI geolocation. The Data Sharing project will create a central depository where satellite operators can standardize, formalize and automate these procedures.

Satellite Operator RFI Alert Network

The Satellite Operator RFI Alert Network is designed to assist operators who have exhausted their preliminary investigation and suspect that the interfering signal may be from a source outside their particular network. The goal of the Satellite Operator RFI Alert Network is to reduce the duration of customer-impacting RFI events.

For the I3 initiative to be successful, satellite operators, manufacturers, industry organizations and end-users must prioritize interference mitigation within their own technical operations and ensure that field technicians and engineers have appropriate skills.

By sharing critical information on interference incidents and resolutions,

and properly training technicians, the satellite and end-user communities can gain significant ground in mitigating RFI. Together, we can ensure the effectiveness of end-user operations and continued strong growth for everyone involved in global satellite communications.

About the author

As Intelsat’s Senior Vice President & Chief Technical Officer, Thierry Guillemin is responsible for managing Intelsat’s terrestrial and space-segment operations encompassing Customer



Operations & Engineering, Space Systems Management and Planning, and Satellite Operations and Engineering. He has more than 25 years experience in all phases of spacecraft development, launch and operations.

Opening graphic courtesy of Grand Canyon University

The GVF Method

For many years, GVF has provided comprehensive skills training to thousands of VSAT technicians worldwide. To accomplish this, GVF uses online, interactive training to deliver a sequence of VSAT installation certification courses. The basic course teaches the core skills that are critical for satellite interference prevention. Optional, follow-on courses cover advanced theory, practical techniques, and specialized training for specific terminals.

Each student is assigned a private account on the GVF Learning Management System (LMS). After logging in, the student steps through the lessons at their own pace. Each lesson concludes with a quiz and, in certain cases, a simulator skills assessment. There is a comprehensive exam at the end of each course and downloadable reference materials are provided. Students who need extra time to study new materials are welcome to linger. However, if they encounter topics with which they are already familiar, they can move through the courses more quickly.

Flash-based simulators provide realistic, critical skills practice with features that go beyond what is possible in a classroom. For example, the antenna-pointing simulator uses

an interactive 3D model of an actual VSAT antenna. The student must learn the functions of all the bolts and fasteners, place a wrench on the right ones at the right times, and make real adjustments. The 3D graphic is updated as the student turns azimuth, elevation, etc. Behind the scenes, the software computes link budget updates, antenna beam patterns, mechanical backlash and lock-down shifts, and true meter behavior based on C/N. Not only is the simulation realistic, it can report details that cannot easily be seen on a real antenna — such as true pointing accuracy and graphical depiction of the beam and satellites in the sky.

No special Internet connection or PC hardware or software is required for these online courses. Certification concludes with a hands-on skills test, offered by GVF-qualified examiners worldwide. The screen shot below illustrates the dish-pointing skills exercise from the GVF Core Skills for VSAT Technicians course.

Exercise: az-el pointing start-to-finish

In this exercise you must preset the pol, find the satellite, perform the initial az/el peaking, final beam-balance peaking, and tighten all az/el locks.

Your location is 107 deg E, 32 deg N.

The satellite is at 134 deg E.

The VSAT will use H downlink polarization.

Pointing angles from your look angle calculator:

True azimuth = 136, Elevation = 43, Pol = 36

When you do the fine azimuth and elevation beam-balance peaking steps, you must fill in the following blanks:

Reference level:

Fine Azimuth:

Full span turns count: whole turns+ 12ths

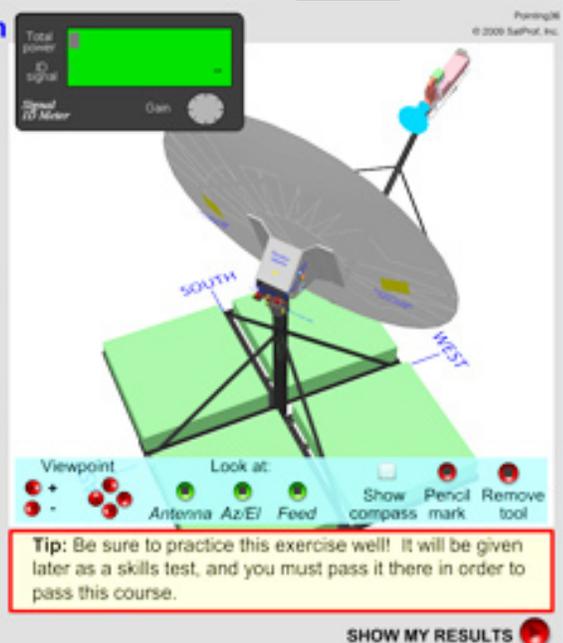
Half span turns count: whole turns+ 12ths

Elevation:

Full span turns count: whole turns+ 12ths

Half span turns count: whole turns+ 12ths

Go ahead and perform the complete procedure. When you are ready, or if you are stuck, click the SHOW MY RESULTS button to see how well you did.



The SlingPath Method

Basic SNG Operations is the initial SlingPath satellite newsgathering (SNG) course aimed at the absolute beginner. It takes the student through satellite basics, including how satellites are used in newsgathering. The course outlines the theory of critical SNG system component parts, various key technical parameters, and operational procedures with satellite operators — including the risk of causing interference, and how to work safely in the field. The presentation of the training is user-friendly, highly graphical, and interactive — with even the most complex subjects presented in an approachable fashion.

Once a student has enrolled, all the training is accessed through a web browser. There are no DVDs or training manuals sent out to the student. Anyone in the world with the most basic Internet connection

can access this state-of-the-art course. The student logs in with their unique username and password to access the course at any time to suit their schedule. It is comprised of seven modules, with a self-check test at the end of each module. Each time the student logs off, their position is stored. When the student logs back in, they can choose to resume where they last left off, or jump to a different module.

Most importantly, after completion of all seven modules (completion is defined as having attempted the self-check test at the end of every module), there is a multiple-choice exam the student must pass with a score of 80 percent in order to be awarded the Final Certificate. The student can attempt the modules and the examination as many times as they wish (though the questions randomly change each time) over the 12-month life of the training seat.

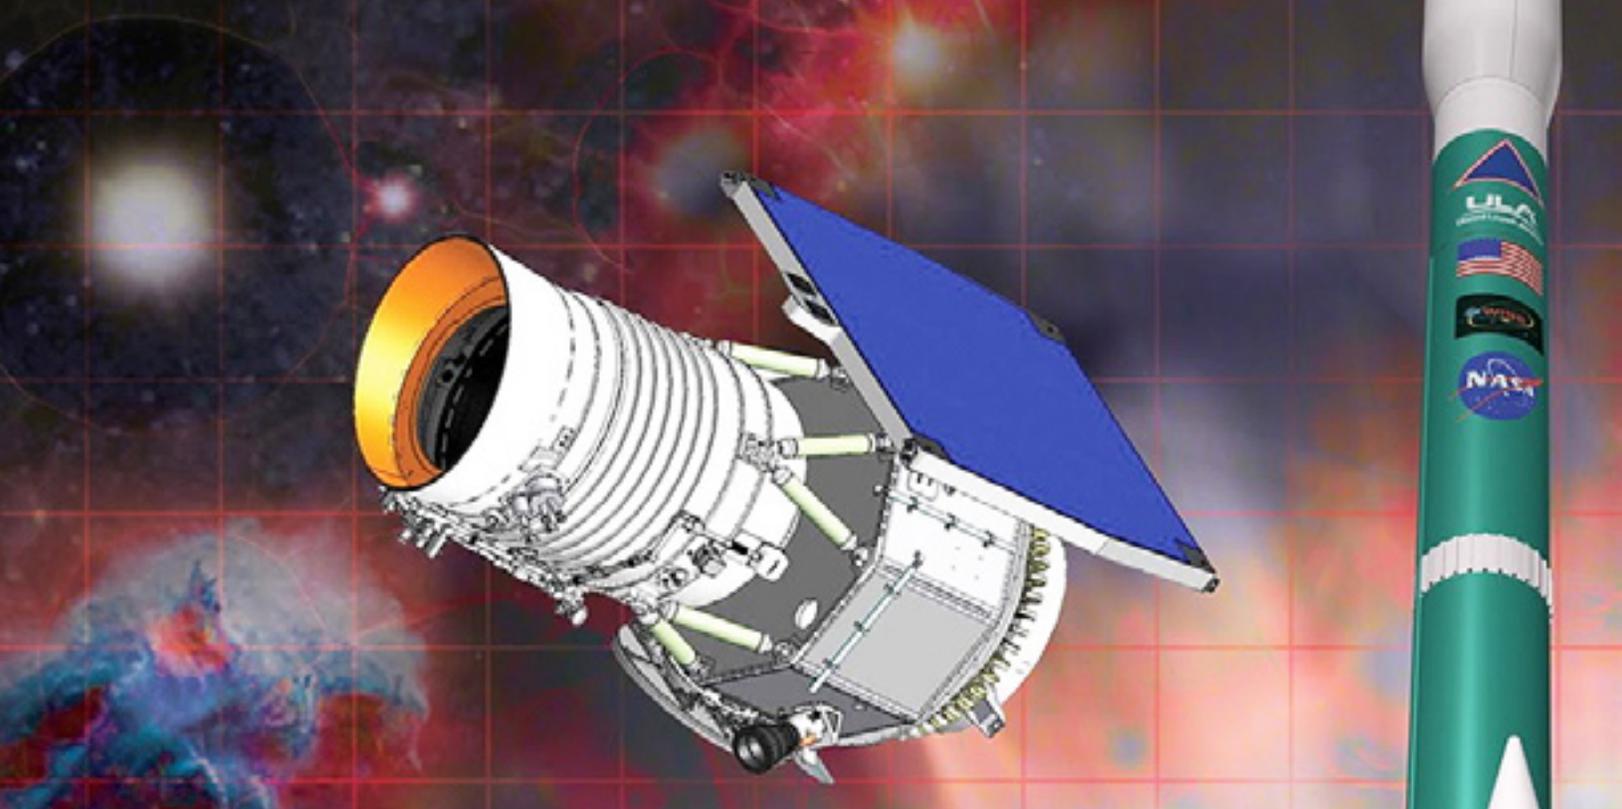
The normal price of the training seat is US\$595.

The screenshot to the left is from the SlingPath Basic SNG Operations training and is an overview summarizing the TV signal uplink process.

The screenshot displays the '3.5 Uplink Process Overview' section of the training. At the top, the header reads 'SLINGPATH BASIC SNG OPERATIONS'. Below the header, a text prompt says 'Click here to see the pathway a TV signal follows through the uplink process.' A diagram illustrates the signal path: 'Baseband Signal (audio/video)' flows into a 'Modulator Encoder' box, then to an 'Upconverter' box, and finally to a 'High Power Amplifier (HPA)' box. A dashed red line represents the signal path, with labels 'L-band or 70 MHz IF' between the Modulator Encoder and Upconverter, and 'SHP (Ku-band)' between the Upconverter and HPA. A satellite dish is shown at the end of the path. A 'Replay' button is located below the diagram. Below the diagram, text reads 'Let's take a closer look at each step in the process. We will begin with compression and digital encoding. Click **Next** when you are ready.' At the bottom of the interface, there are navigation buttons: 'Menu', 'BACK', a central indicator, and 'NEXT'.

WISE • DELTA II

Wide-field Infrared Survey Explorer



JET PROPULSION LAB

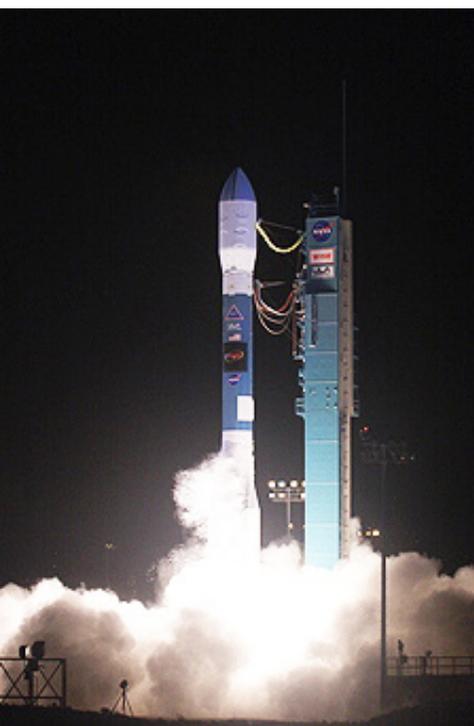
WISE Project

United Launch Alliance capped 2009 with the launch of a Delta II carrying NASA's Wide-field Infrared Survey Explorer (WISE) spacecraft at 6:09 a.m. PST, on Thursday, December the 17th. Rocketing from Space Launch Complex-2, the launch was the eighth Delta II during 2009 and represents the 37th successful mission launched by ULA in its first 36 months of operation.

FOCUS

WISE will scan the entire sky using an infrared telescope with sensitivity hundreds of times greater than ever before possible, picking up the glow of hundreds of millions of objects and producing millions of images. The mission will uncover objects never seen, including the coolest stars, the universe's most luminous galaxies and some of the darkest near-Earth asteroids and comets.

For the WISE mission, the spacecraft was launched on a **Delta II 7320-10C** configuration vehicle featuring a **ULA** first



stage booster powered by a **Pratt & Whitney Rocketdyne RS-27A** main engine and three **Alliant Techsystems** (ATK) strap-on solid rocket motors. An **Aerojet AJ10-118K** engine powered the second stage. The payload was encased by a 10-foot-

diameter composite payload fairing.

Mission Overview

The **Wide-field Infrared Survey Explorer**, or **WISE**, will scan the entire sky in infrared light, picking up the glow of hundreds of millions of objects

and producing millions of images. The mission will uncover objects never seen before, including the coolest stars, the universe's most luminous galaxies and some of the darkest near-Earth asteroids and comets. Its vast catalogs will help answer fundamental questions about the origins of planets, stars and galaxies, and provide a mountain of data for astronomers to mine for decades to come.

Thanks to next-generation technology, WISE's sensitivity is hundreds of times greater than its predecessor, the **Infrared Astronomical Satellite**, which operated in 1983.

WISE will join two other infrared missions in space — **NASA's Spitzer Space Telescope** and the **Herschel Space Observatory**, a European Space Agency mission with important NASA participation.

WISE is different from these missions in that it will survey the entire sky. It is designed to cast a wide net to catch all sorts of unseen cosmic treasures, including rare oddities.

The closest of WISE's finds will be near-Earth objects, both asteroids and comets, with orbits that come close to crossing Earth's path. The mission is expected to find hundreds of these bodies, and hundreds of thousands of additional asteroids in our solar system's main asteroid belt. By measuring the objects' infrared light, astronomers will get the first good estimate of the size

distribution of the asteroid population. This information will tell us approximately how often Earth can expect an encounter with a potentially hazardous asteroid. WISE data will also reveal new information about the composition of near-Earth objects and asteroids — are they fluffy like snow or hard like rocks, or both?

The next closest targets for WISE are cool “failed” stars called brown dwarfs. These Jupiter-like balls of gas form like stars but fail to gather up enough mass to ignite like stars. The objects are cool and faint, and nearly impossible to see in visible light. WISE should uncover about 1,000 in total, and will double or triple

the number of star-like objects known within 25 light-years of Earth. What’s more, if a brown dwarf is lurking closer to us than the closest known star, Proxima Centauri, WISE will find it and the little orb will become famous for being the “closest known star.”

The most distant objects that will stand out like ripe cherries in WISE’s view are tremendously energetic galaxies. Called *ultraluminous infrared galaxies*, or **ULIRGs**, these objects shine with the light of more than a trillion suns. They crowd the distant universe, but appear virtually absent in visible-light surveys. WISE should find millions of ultra-luminous

infrared galaxies, and the most luminous of these could be the most luminous galaxy in the universe.

Other WISE finds will include: newborn stars; disks of planetary debris around young stars; a detailed look at the structure of our Milky Way galaxy; clusters of galaxies in the far universe and more. The most interesting discoveries will lay the groundwork for follow-up studies with other missions, such as NASA's *Spitzer Space Telescope*, the *Herschel Space Observatory*, NASA's *Hubble Space Telescope*, NASA's upcoming *SOFIA* airborne telescope and NASA's upcoming *James Webb Space Telescope*. Powerful ground-based telescopes will also follow up on WISE discoveries.

As with past all-sky surveys, surprises are sure to come. For example, one of the most surprising finds to come out of the Infrared Astronomical Satellite mission was the discovery of excess infrared light around familiar stars like *Vega* and *Fomalhaut*. Astronomers soon determined that the excess light comes from pulverized rock in disks of planetary debris. The findings implied that rocky planets like Earth could be common. Today hundreds of astronomers study these debris disks, and Hubble recently captured an actual photograph of a planet orbiting *Fomalhaut* within its disk. WISE will orbit Earth at an altitude of 525 kilometers (326 miles), circling Earth via the poles about 15 times a day.

A scan mirror within the WISE instrument will stabilize the line of sight so that snapshots can be taken every 11 seconds over the entire sky. Each position on the sky will be imaged a minimum of eight times, and some areas near the poles will be imaged more than 1,000 times. About 7,500 images will be taken every day at four different infrared wavelengths.

The mission's sensitive infrared telescope and detectors are kept chilled inside a Thermos-like tank of solid hydrogen, called a cryostat. This prevents WISE from picking up the heat, or infrared, signature of its own instrument. The solid hydrogen, called a cryogen, is expected to last about 10 months and will keep the WISE telescope a chilly 12 Kelvin (minus 438 degrees Fahrenheit).

After a one-month checkout period, the infrared surveyor will spend six months mapping the whole sky. It will then begin a second scan to uncover even more objects and to look for any changes in the sky that might have occurred since the first survey. This second partial sky survey will end about three months later when the spacecraft's frozen-hydrogen cryogen runs out. Data from the mission will be released to the astronomical community in two stages: a preliminary release will take place six months after the end of the survey, or about 16 months after launch, and a final release is scheduled for 17 months after the end of the survey, or about 27 months after launch.

Mission Operations

The observatory's ground segment includes all facilities required to operate the satellite; acquire its telemetry; and process, distribute and archive data products. **JPL** is responsible for all ground operations that track and control the spacecraft.

WISE will survey the sky from a sun-synchronous polar orbit. It will circle around the poles at the line where day turns to night, called the terminator. The orbit is designed to precess, or shift, with time, so that it stays over this line. This ensures that WISE's solar panels continuously soak up the sun, while the telescope snaps images of the sky overhead, mapping out a 360-degree strip each orbit. As Earth orbits around the sun, these strips of images sweep across the whole sky in only six months. Adjustments are made every half orbit to keep the telescope pointed at the correct strip of sky, following commands that are planned out a week in advance.

Surveying is interrupted four times a day to transmit images to the ground via the *Tracking and Data Relay Satellite System*. These satellites first relay WISE data to a ground station at White Sands, New Mexico. From there, spacecraft telemetry is sent to JPL, while image

FOCUS

and some telemetry data are transferred to the *WISE Science Data Center* at the *Infrared Processing and Analysis Center* at **Caltech**.

Science Data Processing + Archiving

A specialized hardware and software system combines the raw science data with the spacecraft telemetry and converts them into high-quality images and a catalog of stars and galaxies detected by WISE. These products will be archived and distributed to the user community via the web-based services of the Infrared Science Archive at the Infrared Processing and Analysis Center.

Science Goals + Objectives

The primary goal of WISE is to scan the entire sky at infrared wavelengths with vastly improved sensitivity and resolution over past missions. All-sky surveys are essential for discovering new targets of interest, and, in some instances, have opened up entire fields of study. Many modern surveys have combed the sky using various wavelengths, but a gap remains at the infrared wavelengths WISE will observe. The best existing infrared all-sky survey at wavelengths beyond 10 microns is from the highly successful ***Infrared Astronomical Satellite***, operated in 1983. The only existing all-sky survey between 3 and 10 microns is from the ***Cosmic Background Explorer***, operated in 1989.

WISE will scan the sky with far better sensitivity and resolution. Its millions

of images will provide the astronomical community with a vast atlas of the infrared universe, populated with hundreds of millions of space objects. Like scanning the grains of sand on a beach with a metal detector, this infrared telescope will find rare gems buried in the vastness of space. As with past all-sky surveys, the mission's legacy will endure for decades to come.

In addition, WISE will scan much of the sky a second time. This will reveal even more asteroids, stars and galaxies, and catch objects that have changed brightness or position since they were last observed six months earlier. For example, if a cool star has moved noticeably, astronomers will know it is relatively nearby.

The science goals of the mission are:

- » ***To find the nearest and coolest stars***
- » ***To find the most luminous galaxies in the universe***
- » ***To find and study asteroids in our solar system***
- » ***To better understand the evolution of planets, stars and galaxies***

Other Goodies

As WISE is scanning the entire sky, it is going to see all kinds of cosmic wonders, both odd and expected. It will detect many varieties of stars and galaxies, including thousands of dusty, planet-forming disks swirling around stars, factories pumping out newborn stars, clusters of distant galaxies and more.

Surprises are also sure to come, as was the case with previous all-sky surveys. In the end, the data will help astronomers piece together the evolution of stars and galaxies, and gain a better understanding of how our own planet, sun and galaxy came to be.

WISE has already captured its first look at the starry sky that it will soon begin surveying in infrared light. WISE data will serve as navigation charts for other missions such as NASA's Hubble and Spitzer Space Telescopes, pointing them to the most interesting targets WISE locates. A new WISE infrared image was taken shortly after the space telescope's cover was removed, exposing the instrument's detectors to starlight for the first time. [***This direct link will reveal the picture, which shows 3,000 stars in the Carina constellation.***](#)

The image covers a patch of sky about three times larger than the full moon. The patch was selected because it does not contain any unusually bright objects, which could damage instrument detectors if observed for too long. The picture was

taken while the spacecraft was staring at a fixed patch of sky and is being used to calibrate the spacecraft's pointing system.

The Spacecraft

The WISE spacecraft is about the height and weight of a big polar bear, only

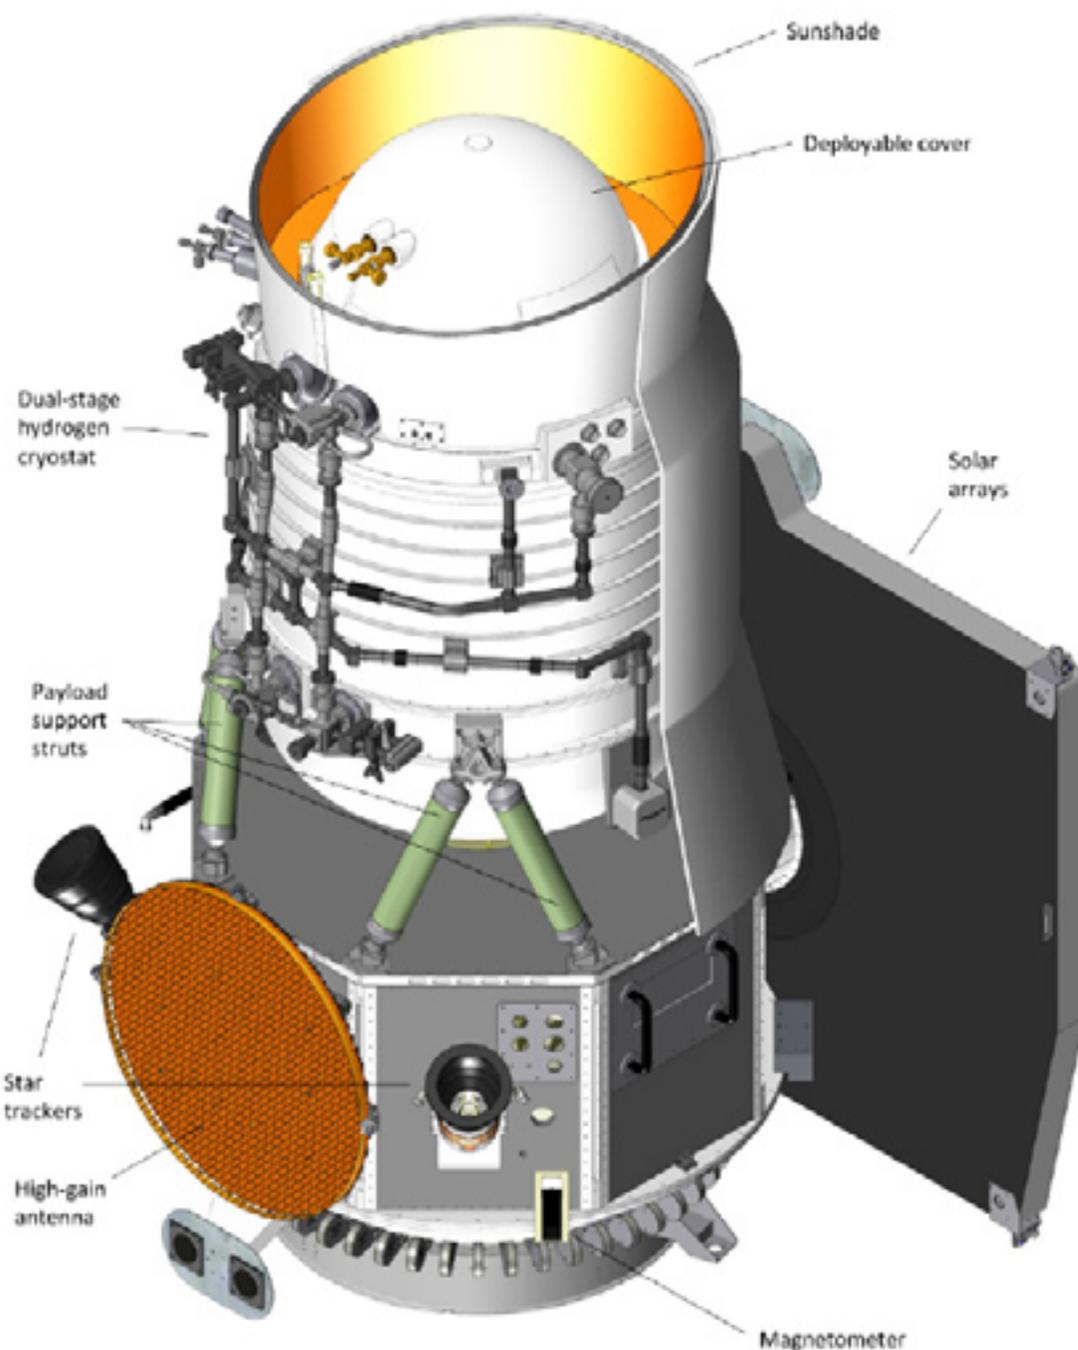
FOCUS

wider. It measures 2.85 meters tall (9.35 feet), 2 meters wide (6.56 feet), 1.73 meters deep (5.68 feet) and weighs 661 kilograms (1,457 pounds). It is composed of two main sections: the instrument and the spacecraft bus. The **Space Dynamics Laboratory** in Logan, Utah, designed, fabricated and tested the instrument,

which includes a 40-centimeter-diameter (16-inch) telescope and four infrared detectors containing one million pixels each, all kept cold inside an outer cylindrical, vacuum-tight tank filled with frozen hydrogen, called a cryostat.

Some say the whole assembly looks like a giant Thermos bottle, while others see a resemblance to the Star Wars robot R2-D2. After launch, the hydrogen vents on the cryostat are opened and the instrument cover is ejected. Once these events have occurred, a scan mirror in the telescope will be the only moving part in the instrument.

At the bottom of the instrument is a three-axis stabilized, eight-sided spacecraft bus that houses the computers, electronics, battery and reaction wheels needed to keep the observatory operating and oriented correctly in space. Two star trackers for precision



pointing are mounted on the sides of the spacecraft bus. A fixed solar panel that provides all the spacecraft's power is mounted on one side of the bus, and a fixed high gain antenna for transmitting science images to the ground is mounted on the opposite side. The bus structure is composed of aluminum honeycomb panels sandwiched between aluminum skin. It has no deployable parts — the only moving parts are four reaction wheels used to turn the satellite.

The base of the spacecraft structure includes a “soft-ride” system of springs to reduce stress from the rocket on the satellite. A metal clamp band attaches the second stage of the rocket to the base of the satellite, and is released to allow the spacecraft to separate from the launch vehicle once in orbit.

Command + Data Handling

The command and data handling system is the spacecraft's brain, responsible for monitoring and controlling all spacecraft functions. It consists of a

single box, called the *Spacecraft Control Avionics*, which was developed by the **Southwest Research Institute**, San Antonio, Texas. The box includes a single-board **RAD750** computer, memory, a command and telemetry interface, an instrument interface,

FOCUS

a flash memory card and spacecraft interface cards. This box can operate the spacecraft either with commands stored in its memory or via “real-time” commands radioed from Earth. It also handles engineering and science data to be sent to Earth.

Electric Power

WISE is powered entirely by a fixed solar panel. The panel is approximately 2 meters (80 inches) wide by 1.6 meters (61 inches) tall. To ensure that sunlight will hit the solar panel properly, the satellite is always oriented with its solar panel facing the sun and the instrument pointed 90 degrees away from the sun. The panel contains 684 solar cells manufactured by **Spectrolabs**, Sylmar, California. The maximum power produced is 551 Watts. Attitude Determination and Control The spacecraft attitude determination and control system is used to adjust the orientation of WISE in space. It consists of a complementary set of four reaction wheels used for maneuvering and three torque rods, which use Earth’s magnetic field to slow down the wheels.

Momentum build-up in the reaction wheels is periodically dumped via the torque rods as WISE flies over the poles. Two star trackers, a fiber-optic gyroscope, a magnetometer and 14 sun sensors provide measurements of the spacecraft’s position.

The star trackers are small telescopes with visible-light electronic cameras called *charge-coupled devices (CCDs)*. They

are capable of acquiring, tracking and identifying multiple stars in their fields of view. The positions of the observed stars are compared with an onboard star catalog to help orient the spacecraft. WISE will use both trackers simultaneously during its survey. The two trackers are pointed in different directions, so that, for example, if the moon interferes with measurements from one star tracker, the other one can be used. **Ball Aerospace & Technologies Corp.** manufactured the *CT-633* star trackers.

Telecommunications

WISE will communicate with Earth via NASA’s **Tracking Data Relay Satellite System**, a network of Earth-orbiting satellites. The spacecraft has one fixed high-gain radio antenna, which uses Ku-band frequencies to relay science data, as well as stored spacecraft health and telemetry, down to the ground at a rate of 100 megabits per second. The high-gain antenna is turned toward NASA’s tracking satellites four times per day, an average of 15 minutes each time, while WISE is orbiting over the poles. Because WISE observes the sky over the poles every orbit, the survey can be interrupted there without creating gaps in the all-sky map.

The NASA satellites relay the spacecraft health and science data to a ground facility in White Sands, New Mexico. From there the spacecraft telemetry is sent to JPL, and the science data to the *Infrared Processing and Analysis Center* at **Caltech**.

The NASA satellites also relay commands and spacecraft data between JPL and WISE using four omni-directional antennas on the spacecraft, which operate at S-band frequencies. Two of these antennas can receive commands from the ground at a speed of two kilobits per second; the other two transmit spacecraft health and telemetry data down to Earth at speeds of either four or 16 kilobits per second.

Flight Software

The observatory uses stored commands to perform its normal operations and also receives commands and sequences from Earth. The software on the flight computer translates the stored and ground commands into actions for various spacecraft subsystems. The flight software also gathers science data as well as engineering telemetry for all parts of the spacecraft and continuously monitors the health and safety of the observatory. The flight software can perform a number of autonomous functions, such as attitude control and fault protection, which involve frequent internal checks to determine whether a problem has occurred. If the software senses a problem, it will automatically perform a number of preset actions to resolve the problem or put the spacecraft in a safe mode until ground controllers can respond.

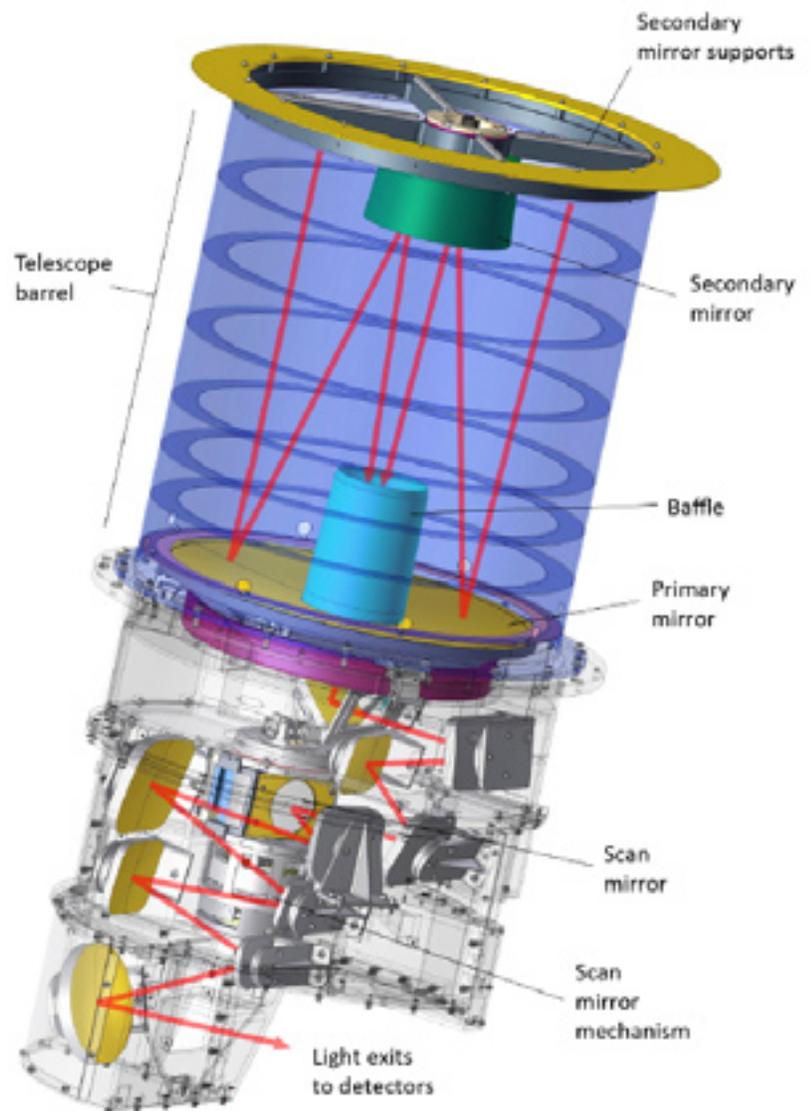
Thermal Control System

The spacecraft is protected by an onboard thermal control system that consists primarily of passive elements:

multilayer insulation blankets, radiator panels, thermal coatings and finishes. In addition, thermostatically controlled heaters provide precise temperature control where necessary for the instrument camera and scan mirror electronics, for example. Heaters are also used for maintaining electronics above survival temperatures in contingency modes.

Science Instrument

The WISE telescope has a 40-centimeter-diameter (16-inch) aperture and is designed to continuously image broad swaths of sky at four infrared wavelengths



FOCUS

as the satellite wheels around Earth. The four wavelength bands are centered at 3.4, 4.6, 12 and 22 microns. The field of view is 47-arcminutes wide, or about one-and-a-half times the diameter of the moon.

The telescope was built by **L-3 SSG-Tinsley** in Wilmington, Massachusetts. Its design uses a total of 10 curved and two flat mirrors, all made of aluminum and coated in gold to improve their ability to reflect infrared light. Four of the mirrors form an image of the 40-centimeter primary mirror onto the flat scan mirror. The scan mirror moves at a rate that exactly cancels the changing direction of the spacecraft on the sky, allowing freeze frame images to be taken every 11 seconds. The scan mirror then snaps back to catch up with the spacecraft as it continues to survey the sky.

The remaining mirrors form a focused image of the sky onto the detector arrays. Before reaching the arrays, the light passes through a series of flat “dichroic” filters that reflect some wavelengths and transmit others, allowing WISE to simultaneously take images of the same part of the sky at four different infrared wavelengths.

The image quality, or resolution, of WISE is about six arcseconds in its 3.4, 4.6 and 12 micron bands, meaning that it can distinguish features one six-hundredth of a degree apart. At 22 microns, the resolution is 12 arcseconds, or one three-hundredth of a degree.

This means WISE can distinguish features about five times smaller than the Infrared Astronomical Satellite could at 12 and 25 microns, and many hundred times smaller than NASA’s Cosmic Background Explorer could at 3.5 and 4.9 microns.

Detectors

Light gathered by WISE’s telescope is focused onto what are called focal planes, which consist of four detector arrays, one for each infrared wavelength observed by WISE. Each of the detector arrays contain about one million pixels (1,032,256 to be exact). This is a giant technology leap over past infrared survey missions. The Infrared Astronomical Satellite’s detectors contained only 62 pixels in total.

The 3.4- and 4.6-micron detectors convert light to electrons using an alloy made of mercury, cadmium and tellurium. The electrons from each of the million-plus pixels are measured on the spot every 1.1 seconds, and the result sent to the instrument electronics. These detector arrays, a type known as the **HAWAII 1RG**, were manufactured by **Teledyne Imaging Sensors**, Camarillo, California. They need to be warmer than the rest of the instrument to improve their performance. The 12- and 22-micron detectors sense light using silicon mixed with a tiny amount of arsenic. They have readout electronics specially developed for the low-temperatures of WISE and were manufactured by **DRS Sensors & Targeting Systems**, Cypress, California.

Cryostat

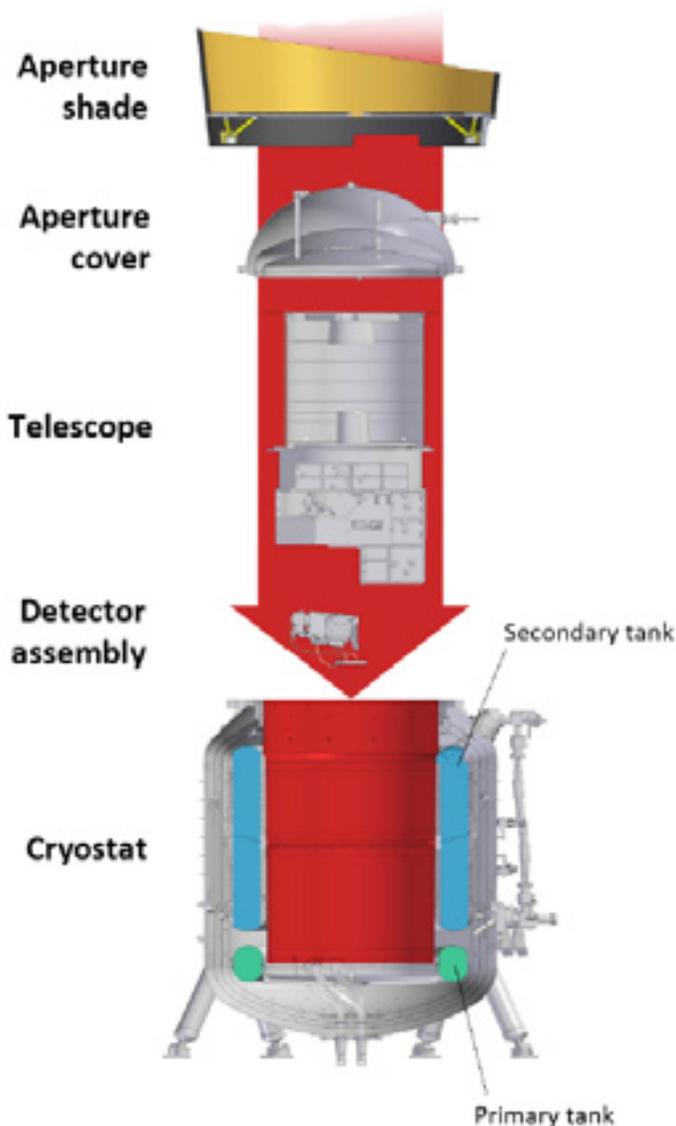
As WISE is designed to detect infrared radiation from cool objects, the telescope and detectors must be kept at even colder temperatures to avoid picking up their own signal. The WISE telescope is chilled to 12 Kelvin (minus 261 degrees Celsius or minus 438 degrees Fahrenheit) and the detectors for the 12- and 22-micron detectors operate at less than 8 Kelvin (minus 265 degrees Celsius or minus 447 degrees Fahrenheit). The shorter wavelength 3.4- and 4.6-micron detectors operate at a comparatively

balmy 32 Kelvin (minus 241 degrees Celsius or minus 402 degrees Fahrenheit).

To maintain these temperatures, the telescope and detectors are housed in a cryostat, which is essentially a giant Thermos bottle. The cryostat is extremely efficient at keeping heat away from the detectors -- its insulating power is equivalent to a home insulation rating of "R-300000."

The WISE cryostat, manufactured by **Lockheed Martin Advanced Technology Center**, Palo Alto, California, has two tanks filled with frozen hydrogen. The colder, or primary cryogen tank, the smaller of the two tanks, cools the 12- and 22-micron detector arrays. To achieve this low operating temperature, a larger 12-Kelvin secondary tank protects the primary tank from nearly all the heat from the outer structure of the cryostat, which is comparatively warm at about 190 Kelvin (minus 83 degrees Celsius or minus 117 degrees Fahrenheit). This secondary tank also cools the telescope and the 3.4- and 4.6-micron detectors. Small heaters are used to warm the 3.4- and 4.6-micron detectors from 12 to 32 Kelvin.

It is important to maintain a vacuum inside the cryostat when it is cold and on the ground; otherwise air would freeze inside it. It would become a giant popsicle. A deployable aperture cover seals the top of the cryostat while on the ground to prevent air from getting in. After WISE is safely in orbit, a signal is sent to eject the aperture cover. Three



pyrotechnic separation nuts will fire, and the cover will be pushed away from the spacecraft by a set of springs. An aperture shade is mounted at the top of the telescope to shield the open cryostat system from the sun and Earth's heat.

The expected lifetime of WISE's frozen hydrogen supply is 10 months. Since it takes WISE six months to survey the sky, this is enough cryogen to complete one-and-a-half surveys of the entire sky after a one-month checkout period in orbit.

NASA's Explorer Program

WISE was developed as a medium-class Explorer mission under NASA's **Explorer Program**. The Explorer Program is the oldest continuous program within NASA. It has launched more than 90 missions, starting with the **Explorer 1** launch in 1958 and including the Nobel Prize-winning **Cosmic Background Explorer (COBE)** Mission.

The early Explorer missions were managed by JPL for the U.S. Army. The objective of the Explorer Program is to provide frequent flight opportunities for world-class scientific investigations from space. Explorer missions are focused science missions led by a principal investigator and occur over relatively short periods of time. They are selected via a highly competitive announcement of opportunity process. The program currently administers only principal investigator-led heliophysics and astrophysics science investigations; in the past, it covered more fields of science.

The Explorer Program seeks to enhance public awareness of and appreciation for space science and to incorporate educational and public outreach activities as integral parts of space science investigations.

Individual Explorer missions are mutually independent, but share a common funding and NASA oversight management structure. The program is designed to accomplish high-quality scientific investigations using innovative, streamlined and efficient management approaches. It seeks to contain mission cost through commitment to, and control of, design, development and operations costs.

The Explorer Program is directed by the *Heliophysics Division* within NASA's *Science Mission Directorate*. The Explorer Program Office is hosted at NASA's **Goddard Space Flight Center** in Greenbelt, Maryland.

NASA's **Jet Propulsion Laboratory**, Pasadena, California, manages WISE for NASA's Science Mission Directorate. The mission's principal investigator, *Edward L. (Ned) Wright*, is at UCLA. The mission's education and public outreach office is based at the **University of California**, Berkeley.

Editor's Note: Our thanks to JPL and NASA for their invaluable contributions in authoring this article.



03b Networks

What Is Network Latency – Why Does It Matter?

Internet data is packaged and transported in small pieces of data. The flow of these small pieces of data directly affects a user's internet experience. When data packets arrive in a smooth and timely manner the user sees a continuous flow of data; if data packets arrive with large and variable delays between packets the user's experience is degraded.

TECH CHANNEL

What Is Latency?

According to *Wikipedia*, latency is a time delay between the moment something is initiated, and the moment one of its effects begins or becomes detectable. The word derives from the fact that during the period of latency the effects of an action are latent, meaning “potential” or “not yet observed”. Most people understand that it takes time for web pages to load and for emails to get from your outbox to the destination inbox and yes, this is a form of latency.

In order to understand why this happens, we need to think about latency at a lower level: Latency is a time delay imparted by each element involved in the transmission of data.

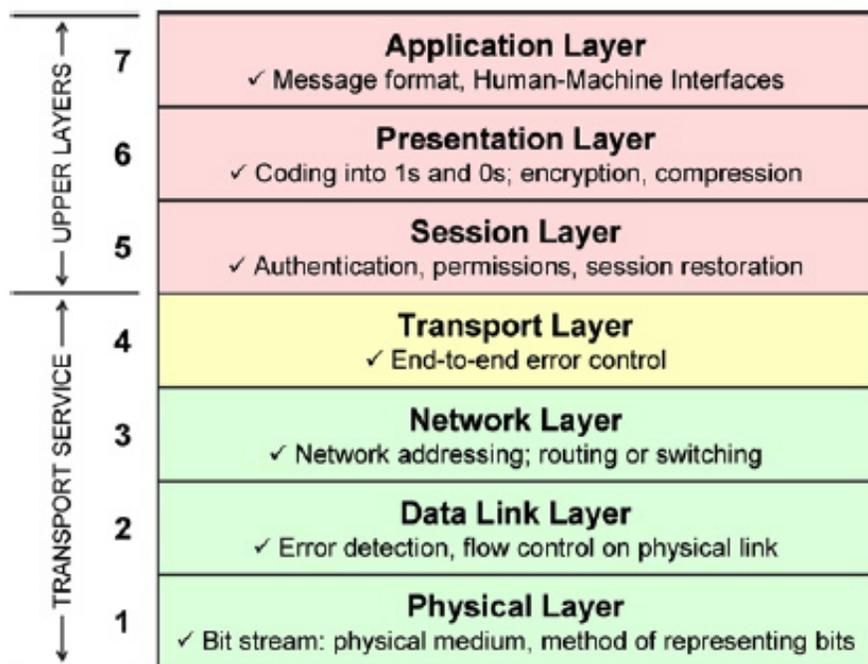
Networking 101

It's important to understand the basic elements of networking to properly grasp the latency issue. Early networking engineers anticipated the need to be able to handle thousands to millions of users on one cohesive network, and thus the TCP/IP networking model was developed.

The key design feature of the TCP/IP networking model is the concept of encapsulation, which is the idea of taking data and wrapping it in a common container for shipping. The container that was developed is called the **IP Datagram**, also known as an **IP Packet**.

The IP Packet is a simple thing: a **header**, followed by **data**. The Header contains

information used for routing the packet to the destination. The data can be any information that needs to be transported such as a snippet of streaming music or a portion of email traffic. The exact construct of the data portion of an IP Packet is defined by the data protocol that is being carried. Data protocols will be discussed later. To understand exactly where latency occurs, it's valuable to know how this most basic unit of networking data is built and transported. For this we turn to the **OSI Model**:



The OSI Model

The OSI model was created to describe the process of turning your application data into something that can be transported on the Internet. The upper layers of the OSI model describe things that happen within the applications that are running on the computer. These applications are web browsers, email programs, and so on.

The lower layers are where information to and from applications are turned into data for transport on a network. This is where data encapsulation occurs and our basic networking data element — the IP Datagram or “packet” is built.

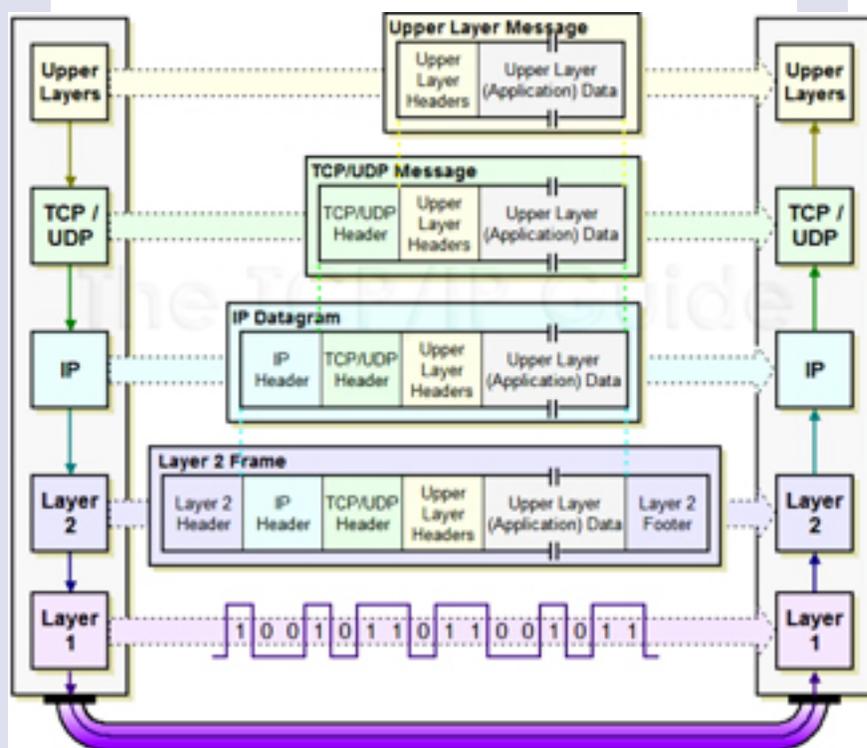
The following diagram shows the encapsulation process in what’s known as the **TCP/IP Stack**. The precise workings of the TCP/IP stack can be different between various computer operating systems. These differences may seem trivial as long as the protocols are implemented properly but when seeking the absolute highest levels of performance it’s important to know that the network stack implementation can be a significant cause of networking performance variability.

The transport of network data is a three step process:

1. *Data from a source application is passed down through the stack. During this process the application data is wrapped into IP Datagrams which are commonly called “packets”. Packets are then transmitted by the sending computer in the network.*
2. *Packets are passed along the network (purple line) until they reach the destination computer.*
3. *Packets are received from the network by the destination computer and are passed up through the stack. During this process the application data is extracted and then passed along to the destination application.*

The additional encapsulation at Layer 2 is called **framing**. This is the stage where the IP Datagram is turned into bits which are appropriate for a particular type of network.

Layer 1 is the physical network medium connection. This layer handles the conversion of the layer 2 bits into electrical, optical, or radio signals that can be transported. The network interface, often called the **NIC** or **Network Interface Card**, can be fiber-optic, copper wire, or a wireless radio interface.



TECH CHANNEL

What Causes Latency?

As described above there are many logical, electrical, and physical elements involved in computer networking. The OSI model identifies each of these elements with regard to specific functionality and delays, another name for latency, occur at every stage of the process.

Application Layer Latency

Layer 7, 6, 5 are the upper “application layers”. Regardless of the speed of the processor or the efficiency of the software, it takes a finite amount of time to manipulate and present data.

Whether the application is a web page showing the latest news, or a live camera shot showing a traffic jam, there are many ways in which an application can be affected by latency. One common source of application latency is the need to read and write data to a disk. There are also hardware limitations that affect application performance such as the amount of memory.

Serialization Latencies

The encapsulation of data, which occurs at the **Transport Layers** (1 through 4), is called **serialization**. Serialization takes a finite amount of time and is calculated as follows:

For example:

- » **Serialization of a 1500 byte packet used on a 56K modem link will take 214 milliseconds**
- » **Serialization of the same 1500 byte packet on a 100 Mbps LAN will take 120 microseconds**

Serialization can represent a significant delay on links that operate at lower transmission rates, but for most links this delay is a tiny fraction of the overall latency when compared to the other contributors.

Data Protocols + Latency

Protocols which operate in the transport layers range from simple to advanced. Many transport protocols use information obtained from the timing of the data flow to ascertain link performance and attempt to adapt to specific conditions present on the network. This is a key issue and will be addressed later in this article.

Routing + Switching Latencies

For a network to do its job packets have to be passed from Point A to Point B. This would be simple if the Internet was just two computers and two locations, but this is certainly not the case. In IP networks such as the Internet, IP packets are forwarded from source to destination

Serialization delay = packet size in bits / transmission rate in bits per second

through a series of IP routers or switches that are interconnected by links such as circuits. The IP routers use the destination address in the IP header to determine the next router in the path from source to destination. The IP routers utilize routing algorithms to continuously update their decision about which router is the best one to get the packet to its destination.

A router or circuit going down, or congestion along the path, can change the routing. Managing the wealth of Internet traffic induces delays (latencies) caused by the routing and switching process.

This refers to the amount of processing time for a router or switch to receive a packet, process it and transmit it on its way. Modern IP hardware interfaces have delays on the order of a few nanoseconds and are generally negligible when compared to network propagation delays which are discussed in the next section. High performance IP routers and switches each add approximately 200us (microseconds) of latency to the link due to processing and forwarding. If we assume that the average IP backbone router spacing is 800 km, the 200us of routing/switching delay is equivalent to the amount of latency induced by 40km of fiber. Therefore routing/switching latency tends to contribute only 5 percent of the end to end delay for the average Internet link.

Queuing + Buffer Management

Another issue which occurs within the transport layers is called “queuing

latency”. This refers to the amount of time an IP packet spends sitting in a queue awaiting transmission due to over-use of the outgoing link after the switching delay has been completed. While over-utilization of high-speed Internet backbone links tends to be rare, instances of congestion are common on lower speed access circuits.

When congestion occurs, routers use sophisticated data queue management algorithms such as **WRED** (*Weighted Random Early Detection*) to minimize data loss. Congestion can cause queuing delays to become infinite since packets are dropped when router buffers become full. Queuing algorithms use a variety of packet management schemes to ensure queuing latency is minimized; best common practice WRED configurations typically bound queuing latency at 20ms.

What Is Propagation Delay?

Propagation delay is a phenomenon in which the physical properties of the medium cause the transmitted information to slow down. The amount of slowing caused by the medium is called the **velocity factor (VF)**. Most people are surprised to learn that copper wire and fiber-optic cables have similar velocity factors. Fiber optic cables typically measure around 70 percent of the speed of light, whereas copper cable varies from 40 to 80 percent, depending on the construct.

Coaxial cable is commonly used and many types have a VF of 66 percent.

TECH CHANNEL

Satellite communication links use electromagnetic waves to propagate information. The information is converted from electrical signals to radio signals using a modem (modulator/demodulator). Once these radio signals leave the antenna, they travel at the speed of light. Let's calculate how long it will take an email to travel from New York to London assuming that we are the only user on a private communications' channel.

Ignoring the actual routes taken by undersea cables due to the ocean's floor, let's assume the path from New York to London is the great circle distance of 5458 km. These are the latencies caused only by propagation delays in the transmission medium. If you were the only one sending one single data bit and you had unlimited bandwidth available, the speed of the packet would still be delayed by the propagation delay. This delay happens without regard for the amount of data being transmitted, the transmission rate, the protocol being used, or any link impairment.

These are the latencies caused only by propagation delays in the transmission medium. If you were the only one sending one single data bit and you had unlimited bandwidth available, the speed of the packet would still be delayed by the propagation delay. This delay happens without regard for the amount of data being transmitted, the transmission rate, the protocol being used, or any link impairment.

Transmission Rate + Bandwidth

Transmission Rate is a term used to describe the nuMBER of bits which can be extracted from the medium. Transmission rate is commonly measured as the nuMBER of bits measured over a period of one second. The "maximum transmission rate" describes the fundamental limitation of a network medium — if the medium is a copper Local Area Network, maximum transmission rates are commonly 10, 100, or 1000 Megabits per second. These rates are primarily limited by the properties and construction of the copper

Propagation delay = distance/speed:

The email sent using a copper link:	$5458 / 197863.022 = 23.58 \text{ ms}$
The email sent using a fiber-optic link:	$5458 / 209854.720 = 26.01 \text{ ms}$
The email sent using a radio link:	$5458 / 299792.458 = 18.21 \text{ ms}$

wires but the capabilities of the network interface card are also a factor. Even the most inexpensive LAN cables can handle transmission rates of 100 MBps, but if the NIC only supports 10 MBps then the link will be rate limited to 10 MBps.

Fiber-optic transmission rates range from around 50 MBps up to 10 Gbps. While 10 Gbps is the highest existing standard, work is being done to extend fiber transmission rates to 100 Gbps. Unlike copper networks, the primary factor limiting fiber-optic transmission rates is the electronics which operates at each end of the fiber.

Wireless LANs and satellite links use a modem (modulator/demodulator) to convert the bits into a modulated waveform, and then on the other end a demodulator will then convert the signal back into bits. The limiting factor in radio-based links is the fact the signal which carries the data must occupy a limited bandwidth when compared to wire or fiber links.

Radio Bandwidth

Signals transmitted using radio waves occupy radio spectrum. Radio spectrum is not an unlimited resource and as such must be shared. To prevent radio interference between users the use of radio spectrum is controlled by nearly every government on the planet. The amount of radio spectrum occupied by any given radio signal is called its **bandwidth**.

The nature of radio spectrum use is beyond this paper but it's important to

understand that generally the occupied radio spectrum of a modem signal will increase with the data rate:

- » *Higher modem data rates cause the modem to occupy more radio bandwidth*
- » *Lower modem data rates will let the modem occupy less radio bandwidth*

As radio spectrum is a limited resource, the occupied radio bandwidth is an important limiting factor in wireless and satellite links.

Data Bandwidth

In data transmission, the data bandwidth is synonymous to the transmission rate being used. Bandwidth is important because it defines the maximum capacity of a data link.

- » *A 10 MBps copper LAN cannot sustain traffic flowing at a higher rate than 10 megabits every second*
- » *A satellite link using modems operating at a 600 MBps rate cannot flow any more than 600 megabits every second*

It's important to understand that data bandwidth is a maximum data flow obtainable over a given transportation segment over a given period of time.

Latency + TCP/IP

The final detail required for understanding why latency is important is an understanding what is going on in the transport layer. Recall that

TECH CHANNEL

the transport layer is the process of encapsulating application data into IP Packets suitable for transport. The protocol used in “data” portion of the IP packet defines the type of data exchange that is taking place. There are two types of network data exchanges:

- » **connectionless**
- » **connection based**

Connectionless data exchange is where data is simply pushed to the destination without regard for its well-being. A connectionless packet traverses the Internet bound for a destination computer but if anything happens to it along the way the sending and receiving computer are none-the-wiser. Transporting data in this manner seems risky but depending on the application it’s generally not detrimental if a few packets get lost along the way. This is a common method of transporting streaming music, and video, and **Voice over IP** traffic.

The transport protocol commonly used for connectionless traffic is called **User Datagram Protocol** or **UDP**. This protocol is popular as there is no overhead or connection management — the data is just sent along. There is no retransmission of lost packets because having them arrive late is not useful for voice or video that is being played out in real-time by the receiving computer.

Connection based data exchanges are far more complicated. These data exchanges rely on the establishment of a

“connection” which manages every packet which is transmitted. The reason for this is that both the sending and receiving computer applications are very interested in ensuring the integrity of every piece of data being exchanged. The transport protocol commonly used for connect-based traffic is called the **Transmission Control Protocol**, or **TCP**. TCP provides error free sequenced delivery of packets. If packets arrive out of order the TCP layer puts them back in order. If packets are missing, TCP asks for retransmission. The TCP protocol makes the Internet reliable.

To support the additional connection management features, TCP packets contain additional information in the header as well as many different packet “types”. TCP packet types are used throughout the establishment of the connection and are the key to providing control of the connection. TCP connections use a client/server model to describe the sender and receiver of data. The use of the term “server” does not necessarily mean a computer performing server duties, rather it’s a computer that is listening for TCP connections. TCP connections have three phases:

1. **Establish the connection**
2. **Send the data**
3. **Close the collection**

TECH CHANNEL

Phase 1

Establishing the connection requires 3 packets...

- » *the client sends a connect request SYN (synchronize) packet to the server*
- » *the server replies with a SYN-ACK (synchronize acknowledge) packet*
- » *the client confirms the receipt of the SYN-ACK by sending back an ACK (acknowledge)*

Phase 2

Once the link is established, the data transfer can start. During the TCP data exchanges, **ACK** (acknowledge) and **NACK** (negative acknowledge) packet types are used to tell the sender that packets have been properly received. If a packet is not received or it contains a bit error, the transmission of the exact same packet is repeated.

Phase 3

Upon completion of the data transport session, the connection will be closed by the following 3 packet exchange...

- » *the closing initiator sends a FIN (finish) packet*
- » *the other side of the link replies with a ACK*
- » *the closing initiator send a coMBination FIN/ACK to end the connection*

How does TCP know a link is operating poorly and what can it do about it? To protect the integrity of the data, TCP packets have several features...

- » *Sequence NuMBers*
- » *TTimestamps*
- » *Flow Control*
- » *Congestion Control*
- » *Checksums*

All of these features are used to guarantee the integrity of the data. They are also used by TCP to determine the quality of the link and to tune the flow of data to maximize the use of the available bandwidth.

An example of this behavior is the way TCP responds to congestion control. The TCP congestion control process uses timers to examine the data flow and subsequent ACK/NACK responses. When TCP detects that ACK/NACKs are taking longer than normal to respond, TCP assumes that the link is being congested somewhere and will slow down the release of packets using flow control. This vital step helps reduce the impact of congestion at routing and switching buffers as well as receiving computer data processing limitations.

By throttling the output of data when congestion is sensed, the TCP congestion control mechanism plays a key role in the flow of TCP/IP traffic.

TECH CHANNEL

Congestion control is generally a valid response to a lethargic network as slow response does often indicate that a portion of the link has a data bottleneck.

A Fictional Data Download

Let's examine the details of downloading a digital image which is 4 megabytes in size. For calculation purposes we need to use bits, so our 4 megabytes (4 MB) is actually 32 megabits (32 MB).

Our downlink will use TCP/IP but we cannot simply create a 32 MB packet to transport our image; such a large packet would be a very cumbersome to deal with. Internet traffic is made up of packets of variable sizes, but the **Maximum Transport Unit (MTU)** is generally only around 1500 bytes. Packets larger than 1500 bytes are considered "Jumbo" packets, but handling these is not yet commonplace for many parts of the Internet. To ensure our image file makes it, we'll stick with 1500 byte packets.

Our 1500 byte TCP packet has a header, which is required for transportation but not useful image data. The size of the header can vary in length from 20 bytes to 60 bytes depending on TCP packet

layer formatting of the data, but for this exercise we'll assume that the entire TCP payload is useful image data.

Our total file size divided by this maximum packet size tells us that transporting our entire image file will take 2777.7 packets. The last packet would normally become shortened instead of our full MTU but to keep the math easy we'll round up to 2778. Our assumptions:

- » ***We are downloading the file from a local computer using a 10 Mbps LAN***
- » ***There is no other traffic on that 10 Mbps LAN — we get the whole pipe***
- » ***The link is operating perfectly, no need to repeat any data during this transmission***
- » ***There is no appreciable latency. This is a copper LAN and a relatively short cable run which is not adding any appreciable propagation delay***

Notice that even for a link that is operating perfectly, the transfer of our 32 Mb image will actually take 34.6 Mb of data. The overhead of TCP added 7.8 percent to the total amount of data which needs to be transported. This overhead is an absolute worst case since we assumed all packets need to be acknowledged.

Activity #	packets x packet length	Total bytes	Total Bits (bytes X 8)
Connection setup	3 x 60 bytes	180	1440
Transmit the file	2778 x 1500 bytes	4,167,000	33,336,000
ACK packets	2778 x 60 bytes	166,680	1,333,440
Connection close	3 x 60 bytes	180	1440
Totals	5562 packets	4,334,040	34,672,320 bits

Now that we have the number of bits transmitted, we need to calculate how long that should take.

This is simple...

» **34.6 Mb/10 Mbps = 3.46 seconds**

Using our nice clean, short networking link we should be able to transport our image file in just over 3 seconds.

The Real World

We've made four assumptions in the above analysis which do not mimic the real world Internet at all.

Assumption 1: The file you want is available on a local computer.

Reality 1: The file is more likely to be some physical distance away. This is not necessarily a problem, but it means we need to traverse a much more complicated network path to get to our data.

Assumption 2: We get the whole 10 Mbps to ourselves.

Reality 2: This is feasible only for the local LAN. It is a certainty that once your little 1500 byte TCP/IP packet reaches the Internet backbone, it will be joined with millions of other packets working their way through the Internet. Your image file packets are going to be mixed in with other traffic such as emails, streaming music files, and so on. You simply don't get the Internet all to yourself.

Assumption 3: The link is operating perfectly.

Reality 3: Internet traffic is routed through an extremely complex collection of hardware which is scattered all over the Earth. The reality is that sometimes a fiber or copper cable is cut or is mistakenly disconnected. A piece of networking equipment such as a router or a switch can break, leaving some other path to pick up and route the extra traffic. When this happens, Internet traffic can start to fill up queues and bottlenecks occur. As mentioned earlier, queuing delays can become significant when the network is operating through a bottleneck.

Assumption 4: There is no appreciable latency present in our network.

Reality 4: The reality is that all of the earlier discussed sources of latency are genuine factors in real-world networks. The impact of latency starts to become noticeable when the latency is significantly longer than the transmission time for the data.

The previous example discussed the data transmission rate in terms of the number of bits per second. To understand how the user is exposed to the effects of latency, we need to convert transmission rate into its measure of time.

» **Bit transmission time = 1/(bits per second)**

TECH CHANNEL

It's easy to see that slower transmission rates take longer to transport packets of data. If the latency on a network is the same as the bit transmission rate, then the

- » *The 50 ms latency link took 3 sec*
- » *The 150 ms latency link took 5 sec*
- » *The 300 ms latency link took 11 sec*
- » *The 600 ms latency link took 17 sec*

Type of Link	Time required to transmit 1 bit	Time for one 1500 byte packet
14.4 Kbps telephone modem	69 microseconds	823 milliseconds
1 Mbps LAN	1 microsecond	12 milliseconds
10 Mbps LAN	100 nanoseconds	1.2 milliseconds
600 Mbps satellite channel	1.6 nanoseconds	19.2 microseconds

impact is very low since the IP packets can still be streamed very close to each other. If the latency on a network is several times longer than the bit rate, the impact will become much more noticeable because the latency spreads the entire data TCP/IP data exchange session over time.

The plots on the next page were made using a TCP/IP packet capture utility. These plots show the packet bit rate on the y-axis and time of day on the x-axis. The data being transmitted was the un-cached web-page reload of the content from the [CNN web page](#). The only condition changed during was the delay between packets — the transmission rate remained the same.

As you can see, the added network latency and its affect on the flow of TCP data spread the web page load over time.

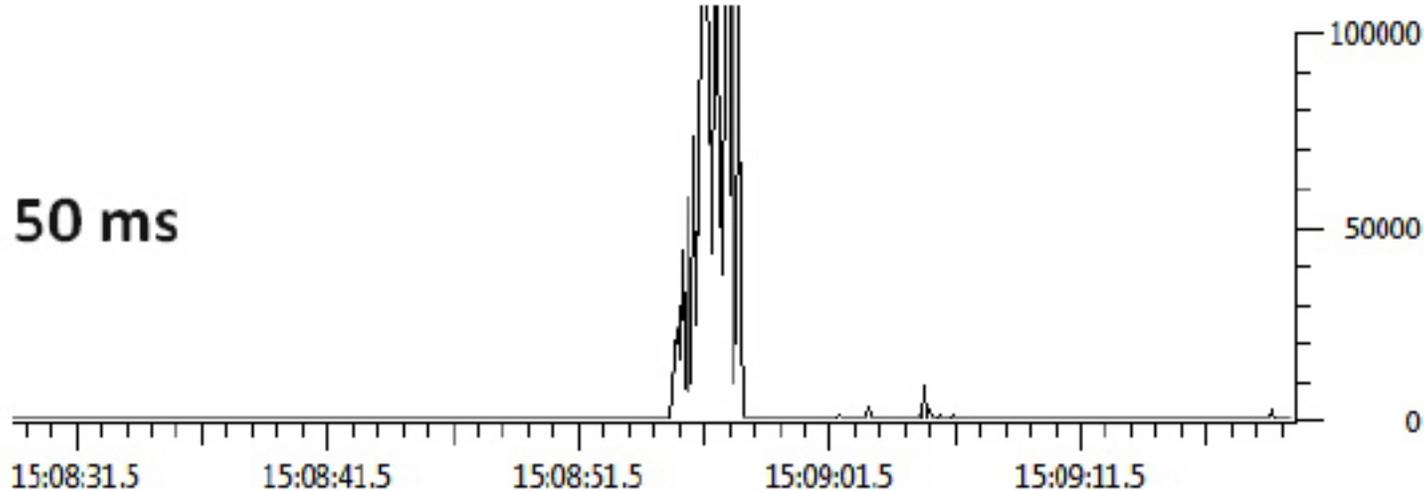
The spreading of network data over time reduces what's called the **Effective Bandwidth** of a link. Packets are still being transported at the same bit rate but due to latency it is taking much more time for all of the web-page packets to arrive.

It's this "spreading over time" behavior of high latency networks which becomes noticeable to the user and creates the impression that a link is not operating at a high speed.

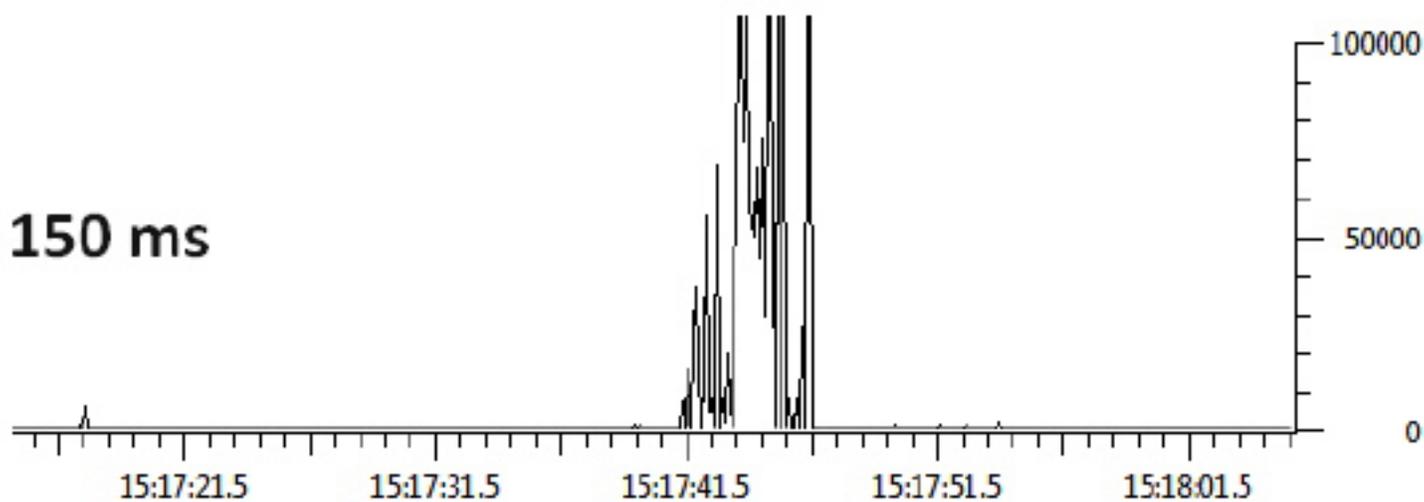
O3b recently conducted another demonstration of real-world effects of latency using the time to load a web page. This is a very common activity and clearly shows users that latency directly affects the way a user obtains data from the Internet.

The following plots show the effects of latency on the time to load the [Wall Street Journal](#) web page...

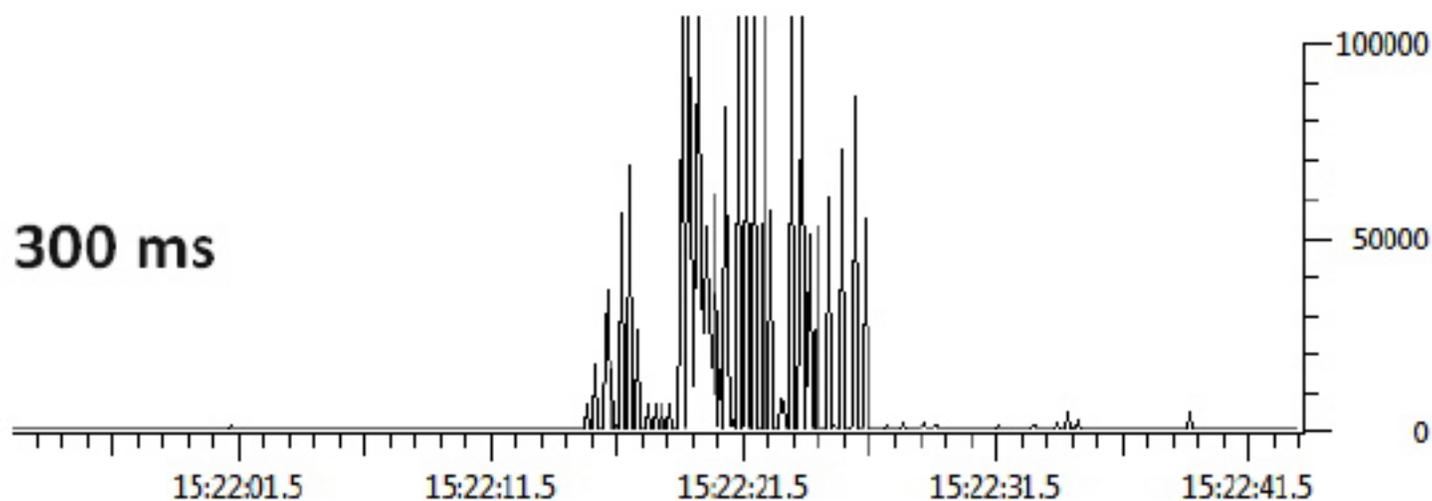
50 ms



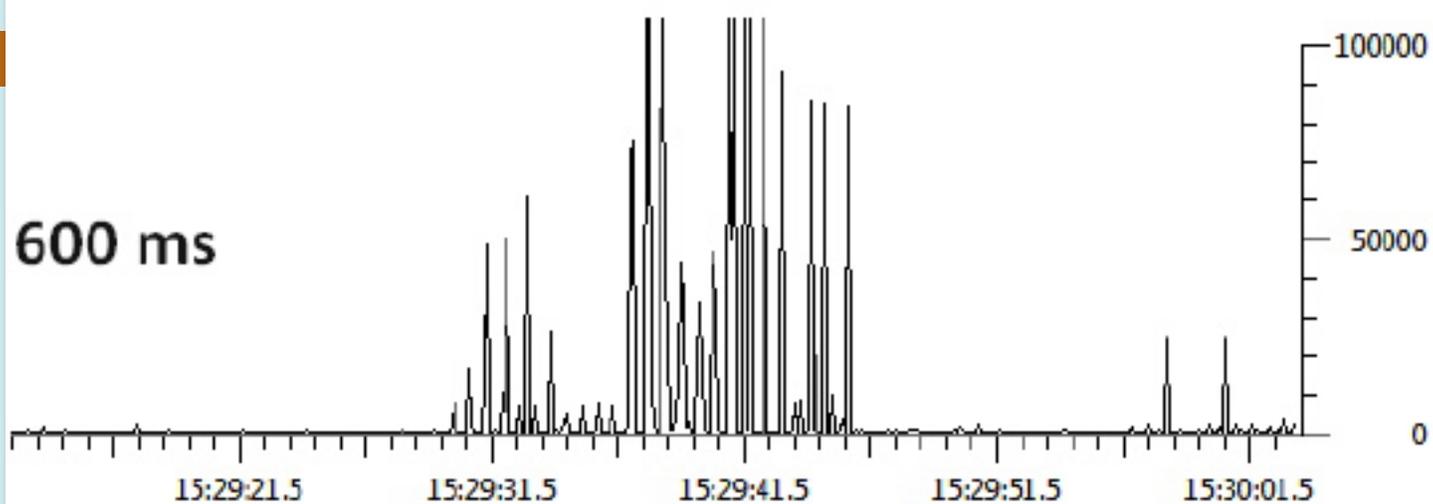
150 ms

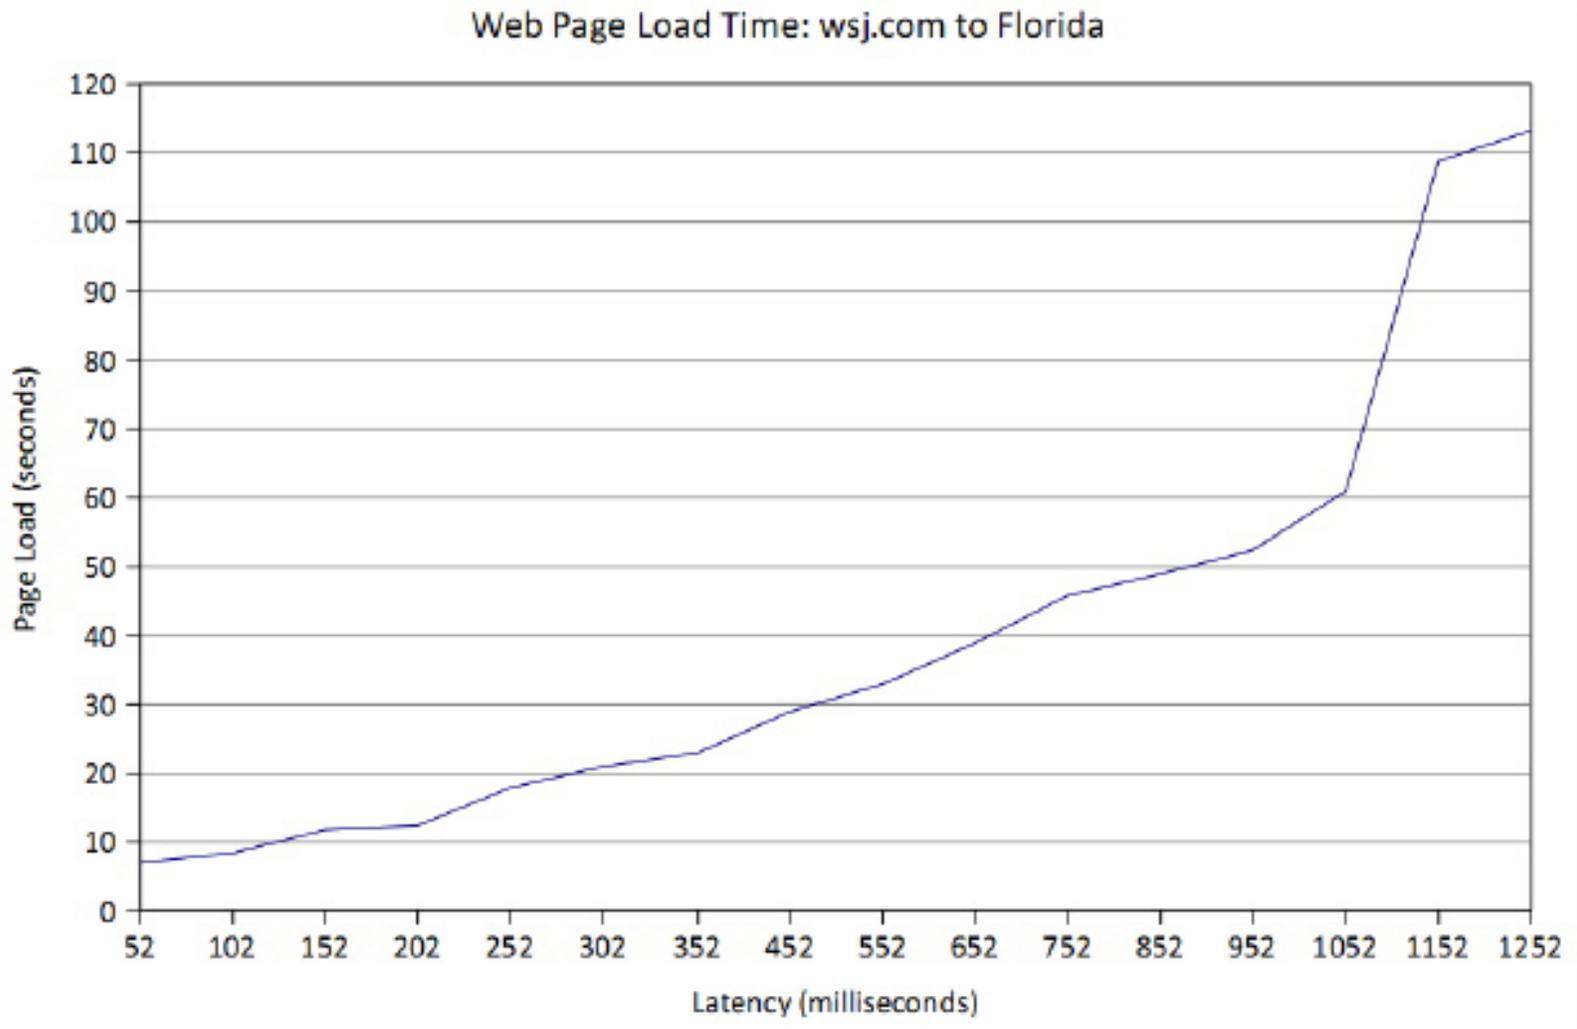


300 ms



600 ms





Satellite Link Latencies

Now that we know the effects of latency on real-world traffic, we'll discuss the latency differences in two satellite technologies. Satellite links can introduce larger latencies than most terrestrial networking segments due to long distances from the ground stations to the satellite. The table on the next page shows the latency caused by propagation delays from two types of satellite configurations...

- » *The O3b Networks MEO orbit constellation at an altitude of 8063 Kilometers*
- » *A geosynchronous satellite at 35,786 Kilometers*

It is important to understand that for satellites which operate as a bent-pipe, the propagation delays are doubled as the signal has to travel both up to the satellite and back down to the Earth before it reaches the next segment of the network.

Latency Calculations - Lagos, Nigeria to the Internet

	Teleport to Satellite (km)	Customer to Satellite (km)	TP to Sat (ms)	Sat to Cust (ms)	Sat RTT (ms)	Internet RTT (ms)	Total RTT (ms)	Data Request Cycle (ms)
O3b Networks								
8 satellites	Almeria	Lagos						
AOS	10427	9019	35	30	130	60	190	285
Max Elev.	10126	8135	34	27	122	60	182	273
LOS	10986	9019	37	30	133	60	193	290
Geosynchronous	Italy	Lagos						
	37923	35847	126	120	492	60	552	828

The table above shows that a ground station in Lagos, Nigeria using an O3b Networks MEO satellite to connect to a teleport in Almeria, Spain, will experience round trip time (RTT) ranging from 122 to 133 milliseconds. If we add in the average internet latency from the Almeria teleport to most Internet destinations in Europe (60 ms), we end up with an overall latency from Lagos to a European internet site of 183 to 193 milliseconds. This range of latency is caused by the change in distance to the ground sites relative to the moving O3b Satellites. The AOS is the Acquisition of Signal, and the LOS is the point at which the O3b system will perform the handover to the next rising satellite. The Maximum Elevation is the point at which the satellite is closest to the customer ground station which explains why this point has the lowest latency.

By comparison, the same Lagos customer site using a geosynchronous satellite to a European internet site using will have

to see latencies of 552 milliseconds. The last column in the chart shows the time required to make a data request and to start receiving the requested data. This data request time includes...

- » **The request packet from the user to the web server**
- » **The web server acknowledging the request**
- » **The web server pushing to requested data to the user**
- » **The data arriving at the user's computer**
- » **TCP also includes a returned ACK packet from the user to the web server but this time is not counted in the Data Request Cycle.**

TCP also includes a returned ACK packet from the user to the web server but this time is not counted in the Data Request Cycle.

Geosynchronous satellite users must wait almost 1 second before they start getting data, whereas the lower latency O3b

TECH CHANNEL

satellite link will receive it nearly 3x sooner. When looking at the basic latency numbers, it's easy to see that the O3b Satellite constellation will offer users a noticeably better Internet experience with more immediate feedback and quicker access to data.

Summary

We have described the structure of IP-based packet switched networks, the functions of the various protocol layers, and the causes of latency in packet switched data networks, such as the Internet.

Latency and overall throughput is dominated by two factors, the length of the route that the packets have to take between sender and receiver and the interaction between the TCP reliability and congestion control protocols and this path length. The O3b Networks satellite constellation in a much lower MEO orbit has significantly lower path length and therefore significantly lower latency than traditional geosynchronous satellites. Therefore, O3b's network latency and throughput approximate and, in some cases, exceed that of fiber based terrestrial networks.

O3b Networks will bridge the digital divide by building a global Internet backbone for telcos and ISPs in emerging markets. The Company is committed to connecting the networks of developed countries with the "other three billion" people who have limited Internet access, hence the name O3b. With investment and operational support from some of the world's largest communications and finance companies, O3b Networks is now building a distribution network to serve a population of several billion consumers and businesses in more than 150 countries. O3b Networks' headquarters are in St. John, Jersey, Channel Islands. Ground systems and technical development are managed through its wholly owned subsidiary in Englewood, Colorado.



Advertisers

AAE System	37
Advantech	47
CABSAT Mena 2010	39
Comtech EF Data	55
Futron	19
IDGA	23
Integral Systems	17
Miteq / MCL	59
NAB	29
Newtec CY	7
Northern Sky Research (NSR)	13
03b Networks	2
PacTel International	43
Paradise Datacom	21
Singapore Exhibition Services	67
SMi Group	61
Wavestream	27