Implementing the Interplanetary Internet

differing approaches

Lloyd Wood
Surrey Space Centre guest lecture
Tuesday 17 February 2009.

How did it all begin?

- Vint Cerf announces start of effort over ten years ago, in July 1998.
- Collaborates with Adrian Hooke of **NASA Jet Propulsion Lab (JPL)** – who leads CCSDS (Consultative Committee for Space Data Systems), an ISO subgroup that sets *standards for space*.
- Space probes predate computing; tape recorder bitstream mindset. Want to move them towards packets and networking.
- Long propagation delays difficult; can't work with protocol timers.

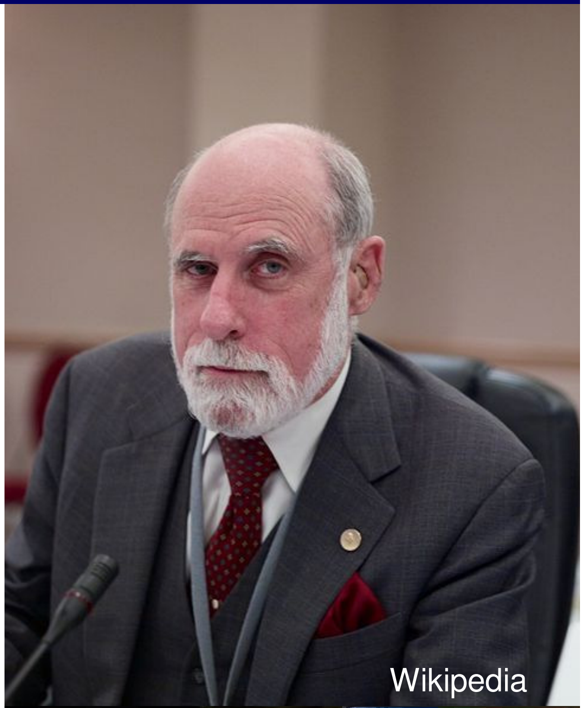
Wikipedia


Associated Press

Vint sets up an Internet Society SIG…

- IPN Special Interest Group (IPNSIG).
- Then a short-lived IRTF 'Interplanetary Internet' group (IPNRG) and a couple of internet-drafts. 2000/2001.
- Problem scope widens to 'Delay Tolerant Networking' (Kevin Fall) and bundles are created, 2002/2003.

- IRTF DTN research group set up. (Kevin introduces DTNRG at IETF 56, March 2003.)
- DARPA *Disruption-Tolerant Networking* proposers' day, January 2004. (Lots of funding.)

Problem scope was consistently widened

- First, let's solve *interplanetary* networking for the long delays of deep space.
- Then, let's solve *delay-tolerant* networking for intermittently-connected ad-hoc networks.
- Then, let's solve *disrupted* ad-hoc military networks under battlefield conditions.
- Increased the interest/attention/funding.
- A way for NASA JPL to attract experts and approval, at low cost to JPL – but will it still solve the original problem?

Two different problem spaces

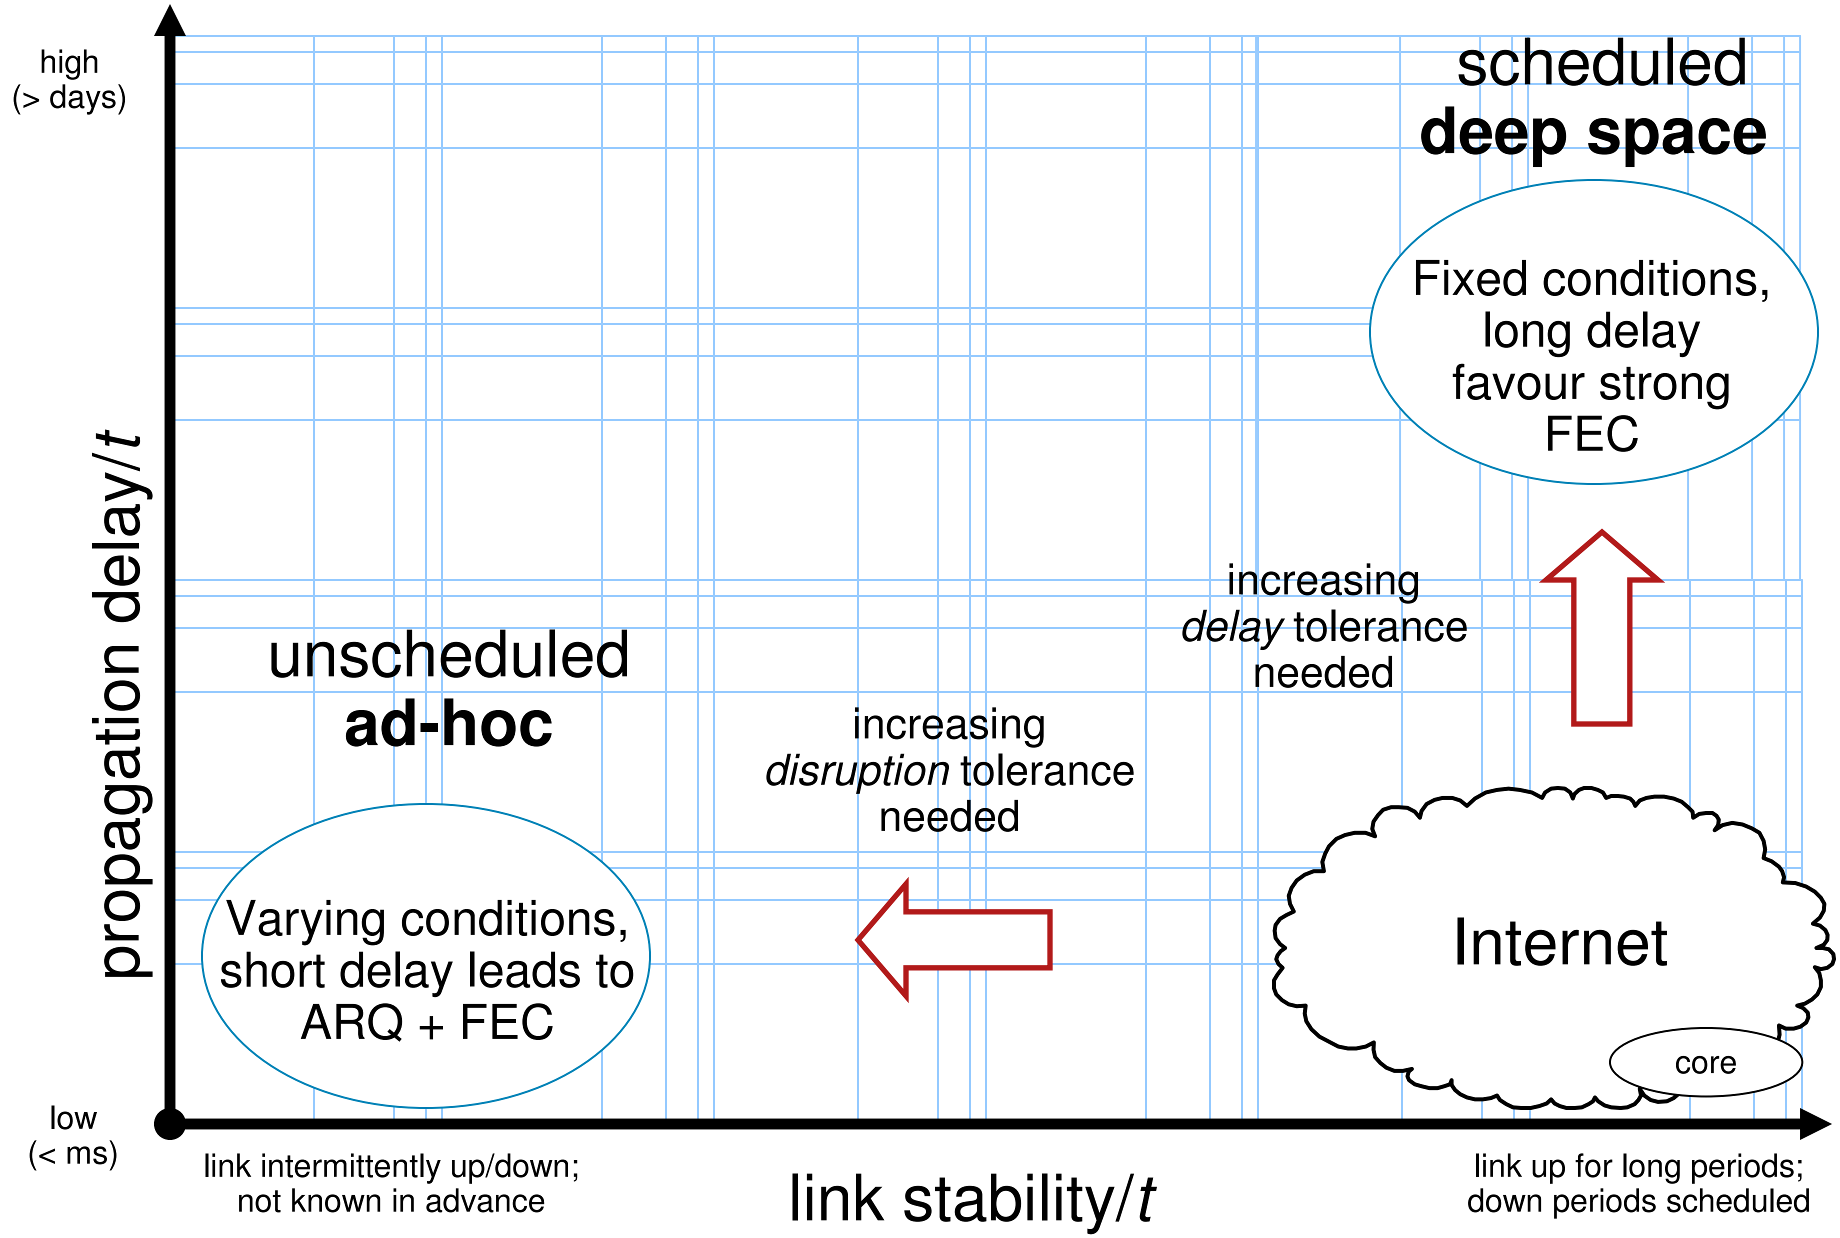

Terrestrial fixed Internet little need for resends between or checking at nodes when resends can easily and quickly be done end-to-end over the whole path instead

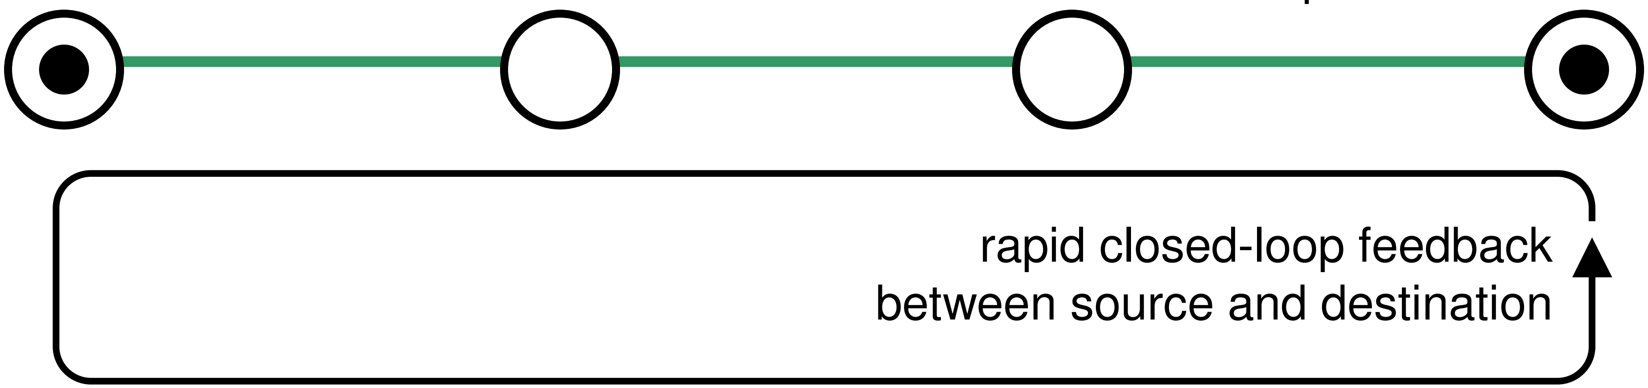

Delay-tolerant network more reliance on separate closed loops between each pair of nodes with local checking for *e.g. custody transfer* and to increase throughput

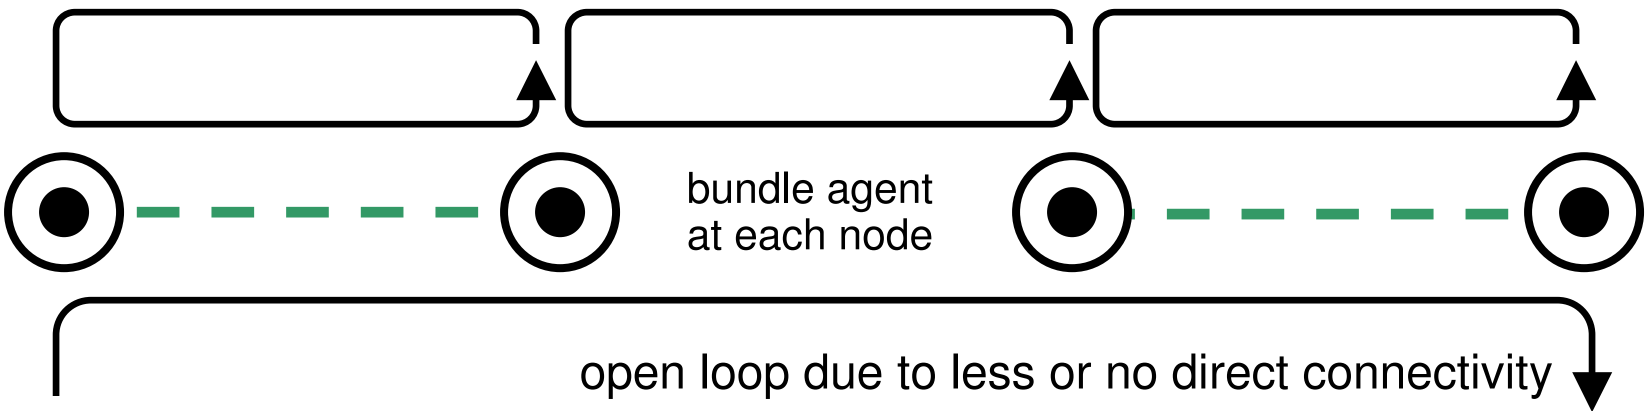

open loop due to less or no direct connectivity between source and destination; no end-to-end loop; no permanent path

What is the Bundle Protocol?

- Basically layer over different *internets*, just as the Internet Protocol layered over different *networks*.
- late binding of Bundle *endpoint identifiers* to a local network address.

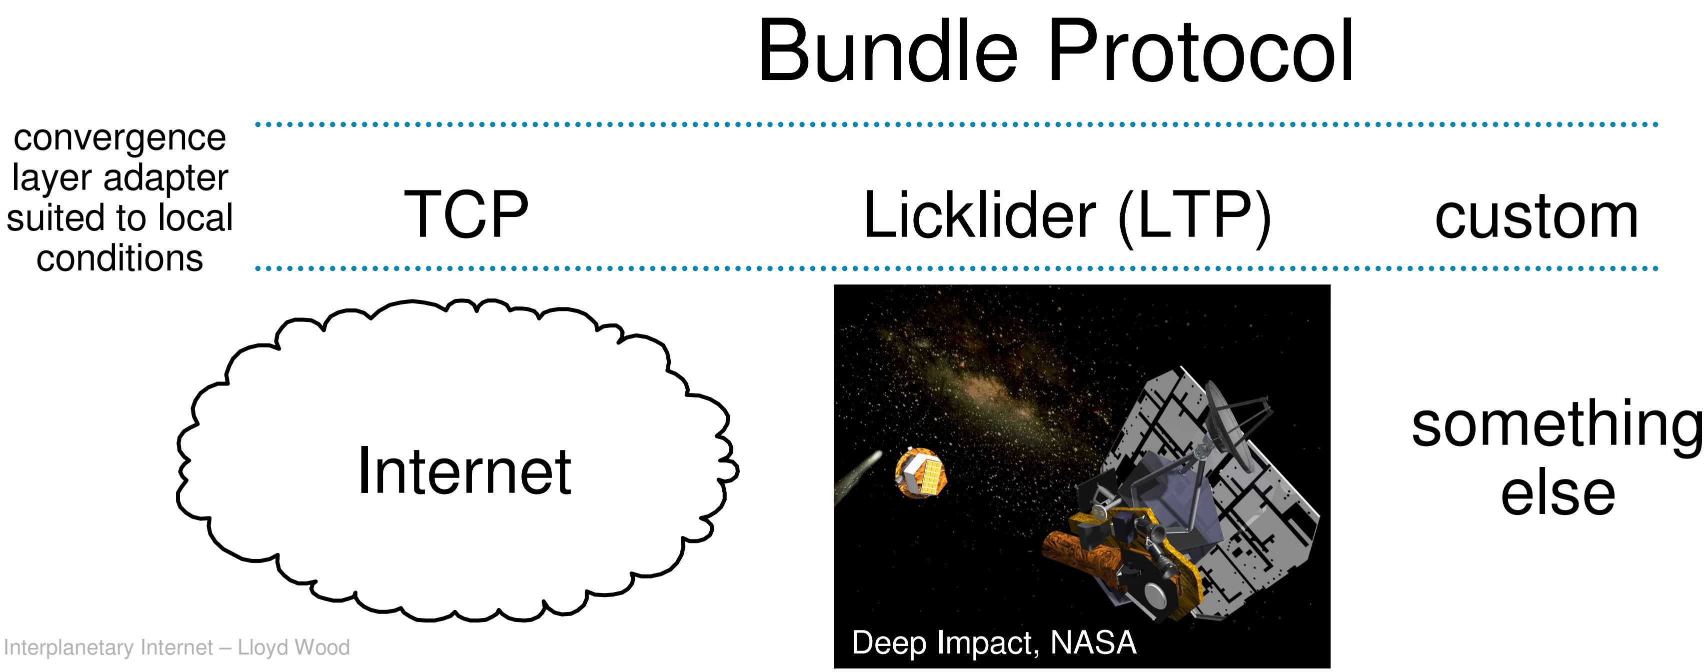

Deep Impact, NASA

Basic Bundle structure – blocks.

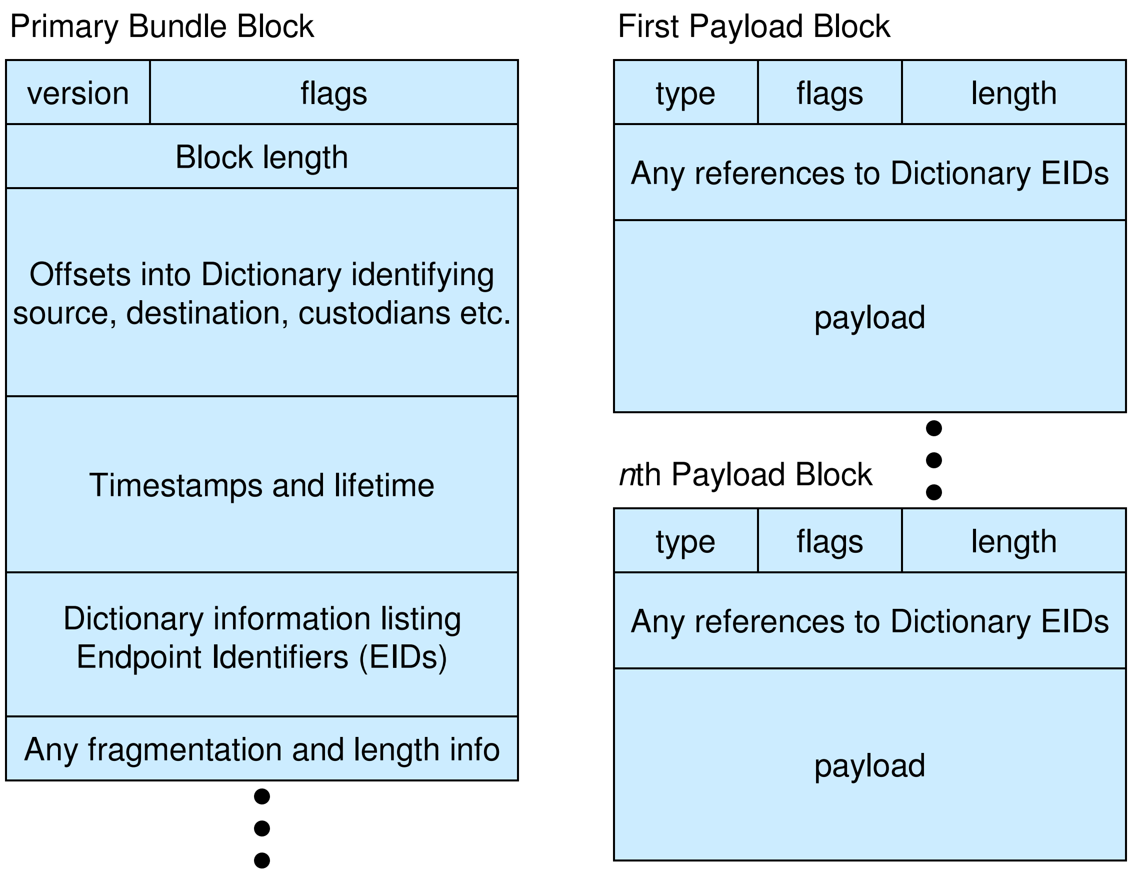

Most fields use SDNVs (Self-Delimiting Numeric Values, like ASN.1) and are not fixed-length.
No checksums.

Bundle Protocol really a container format.

- Multiple blocks, following a primary block with a dictionary. Blocks can be encrypted.
- *Mutable canonicalisation* – idea that block ciphers can cover and protect some different metadata (header) primary fields, similar to IP pseudo-header.
- Custody transfer allows handing over responsibility of delivery.
- **But no end-to-end reliability**. Custody transfer doesn’t check bundle has been copied correctly!
- Variable-length SDNVs are like ASN.1 – last bit indicates continuation. If that bit gets corrupted…

Bundle Protocol could keep CCSDS and Internet separate and self-contained….

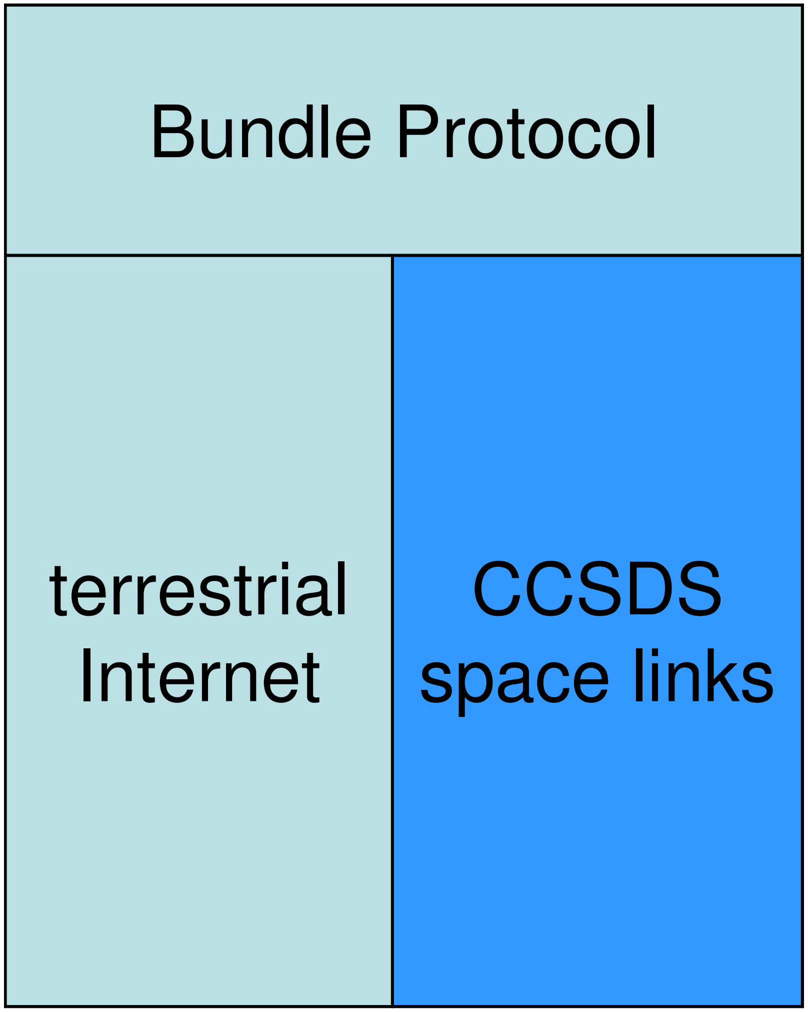

Existing terrestrial and space stacks

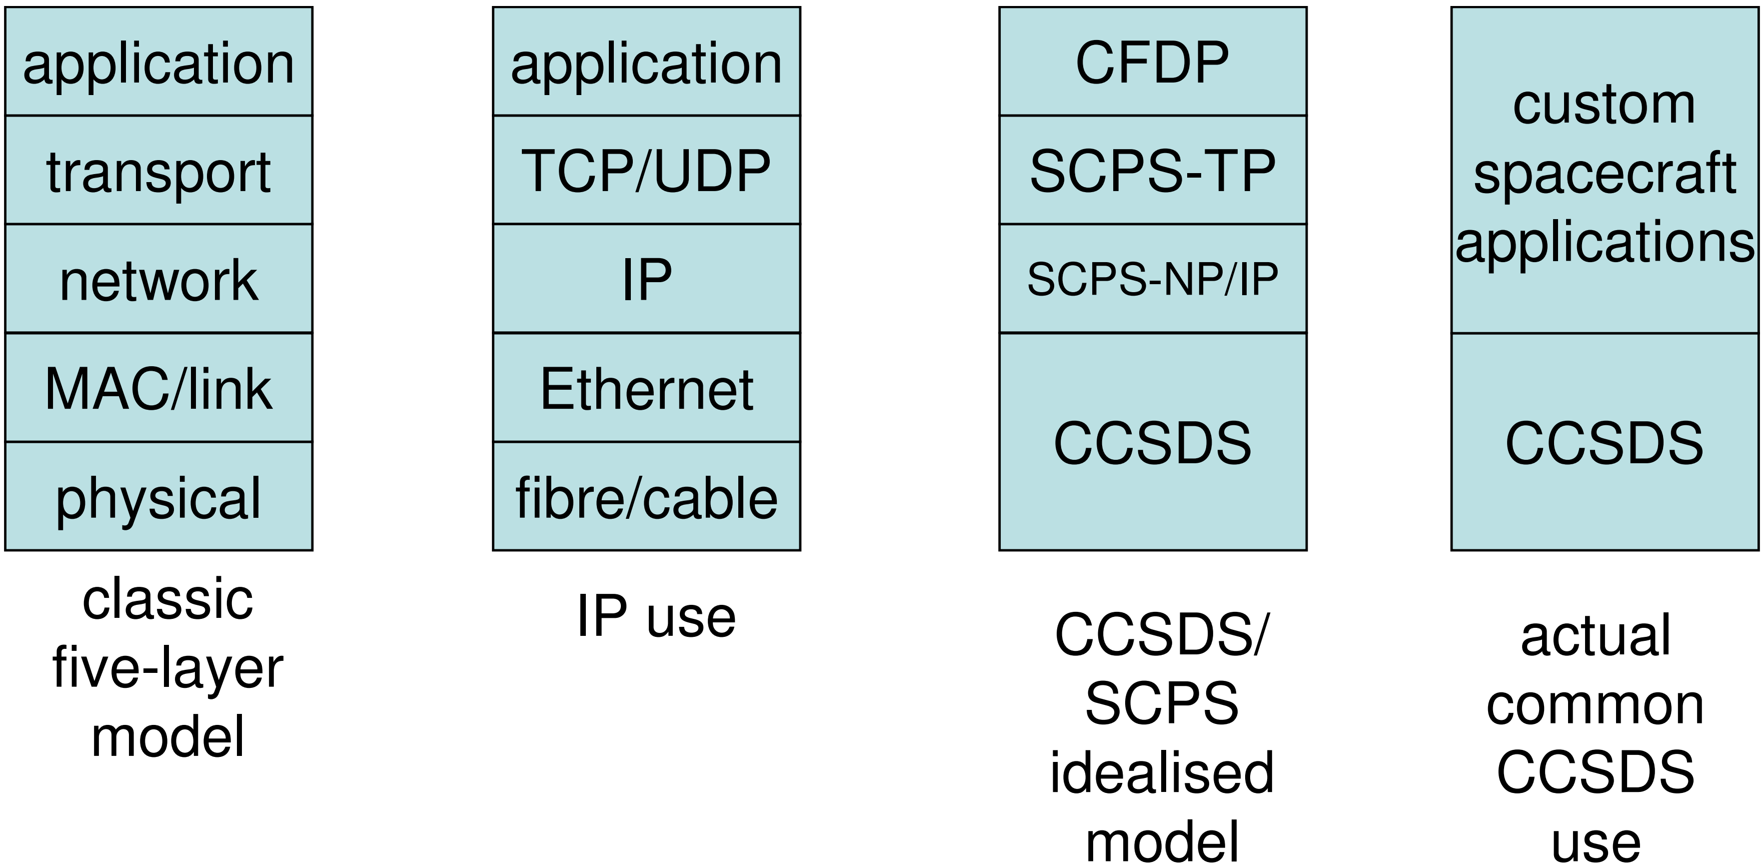

Existing convergence layers for the Bundle Protocol

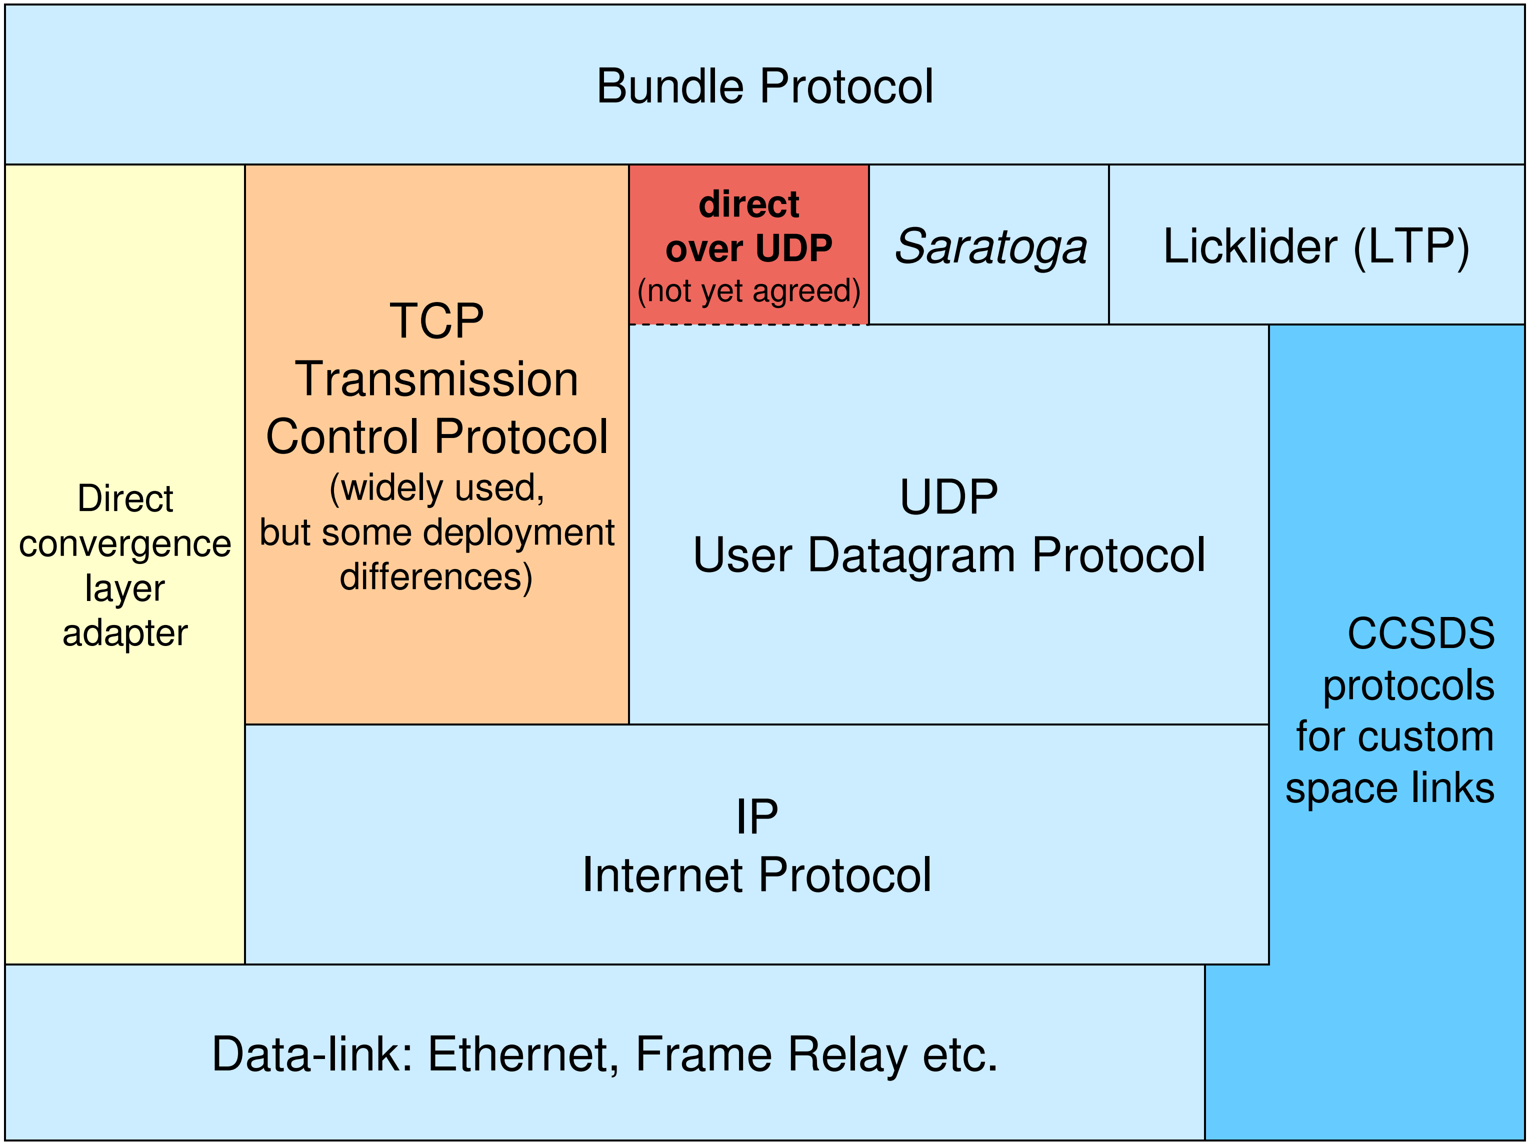

Most Bundle Protocol use is over IP. Except for the CCSDS world, of course.

Scott Burleigh's approach to space networking

- Careful consideration of legacy CCSDS base and slow migration.
- Gave us CFDP, Licklider LTP, and Bundle Protocol.
- CCSDS File Delivery Protocol (CFDP) a bit like Bundle Protocol – layers over everything: TCP, UDP CCSDS, you name it.
- CFDP (lite) in use by Messenger Mercury probe and Deep Impact mission. Licklider LTP a 'stripped down' CFDP in some ways.

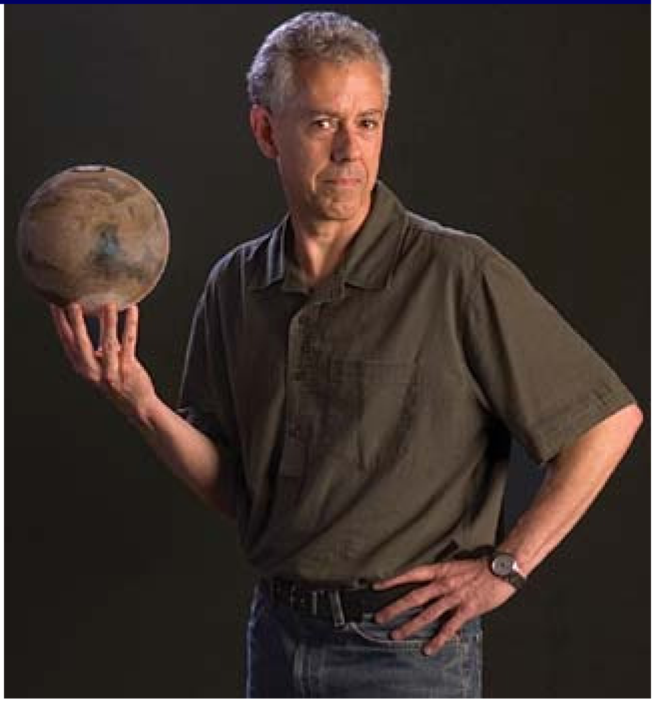

IEEE Spectrum, Aug 2005.

at NASA JPL

CFDP
LTP
Bundle Protocol

CCSDS

Extending the Internet into space

- NASA JPL gives DERA's STRV-1b an IP address (1996).
- NASA Goddard flies IP stack on SSTL's UoSAT-12* (2000). This encourages SSTL to adopt IP.
- Cabletron router on Russian module of ISS. NASA uses IP in shuttle experiments, e.g. VoIP with Cisco SoftPhone tested from *Atlantis* (Feb 2001). These culminated in CANDOS,* tested onboard *Columbia* (Jan 2003).
- NASA gets SpaceDev to launch CHIPSat (Jan 2003).
- SSTL adopts IP with DMC (AlSAT-1 launched in Nov 2002, UK-DMC *et al*. Sep 2003, Beijing-1 Oct 2005). Cisco and SSTL fit CLEO mobile access router on UK-DMC satellite, alongside imaging payloads.
- MidSTAR-1* and SSTL's CFESat launch (March 2007).

*Keith Hogie's team at NASA Goddard was instrumental in use of IP in these projects.

Keith Hogie's approach to space networking

- Grew tired of reinventing the wheel and reimplementing CCSDS protocols with tweaks.
- Proposed basic standards use – IP in standard Frame Relay over ISO standard HDLC. Showed that this worked with Surrey and on final *Columbia* mission.
- HDLC can also be carried as bitstream over CCSDS, as CCSDS supports 'tape recorder bitstream' – easier than CCSDS ways to carry IPv4 (and other ways to carry IPv6). CCSDS community really not keen on that.

IEEE Spectrum, Aug 2005.

Contracts to NASA Goddard.

Effectively designed SSTL's approach to networking with UoSAT-12.

HDLC/FR/IP

Disaster Monitoring Constellation (DMC)

Surrey Satellite Technology Ltd (SSTL) build and help operate an international constellation of small sensor satellites.

The satellites share a sun-synchronous orbital plane for rapid daily large-area imaging (640km swath width with 32m resolution). Can observe effects of natural disasters.

Government co-operation: Algeria, Nigeria, United Kingdom, Turkey and China.

Each government finances a ground station in its country and a satellite. Ground stations are networked together. Further satellites planned.

www.dmcii.com

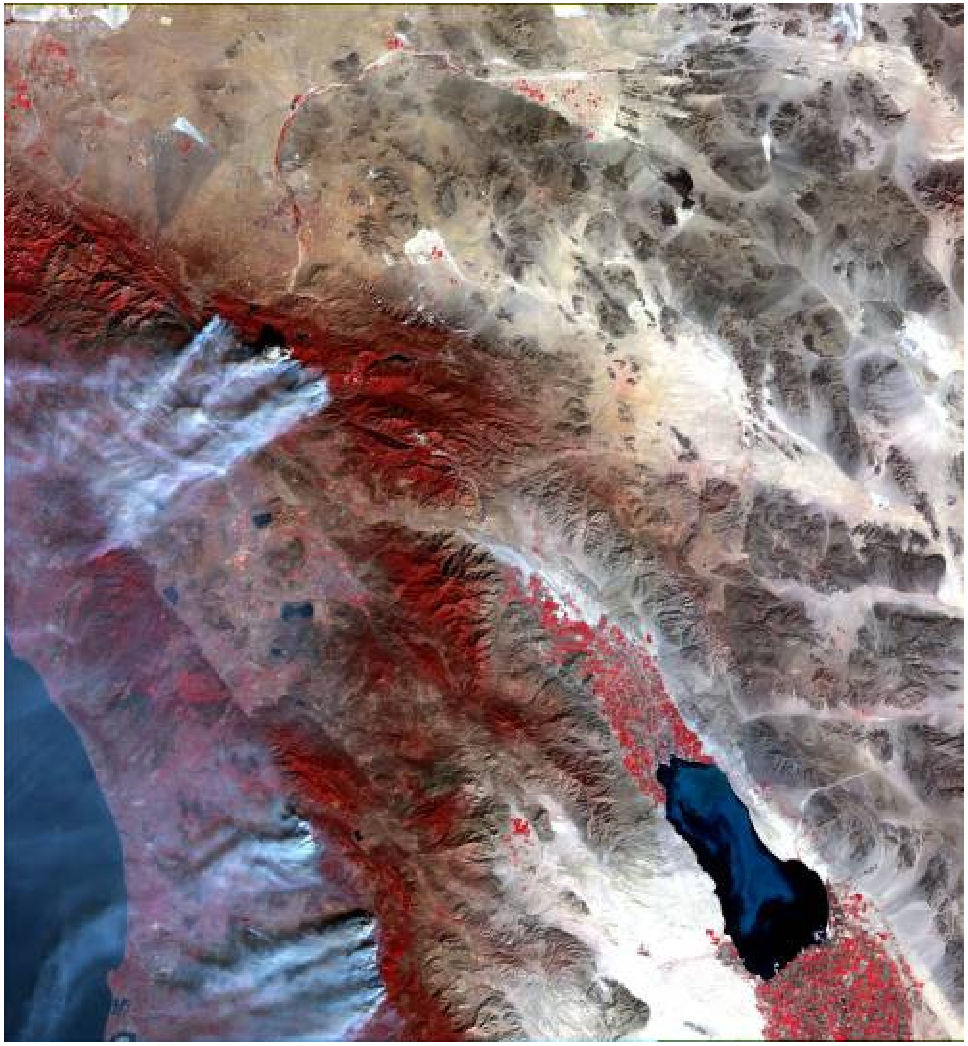

fires in California, 28 October 2003 (UK-DMC)

DMC in use: after Hurricane Katrina, 2005

www.dmcii.com

In this false-color image, dry land is red. Flooded and damaged land is shown as brown.

Small part of an image taken by the Nigerian DMC satellite on Friday 2 September, for the US Geological Survey.

DMC is working as part of the United Nations International Charter for Space and Major Disasters.

Imagery delivered by using Internet Protocol.

Our previous work in space

- UK-DMC satellite, with Cisco router onboard, launched with other satellites into low Earth orbit, September 2003.
- UK-DMC and sister satellites are based around use of Internet Protocol (IP). IP works for satellite and payload communication and control.
- IP internetworking of satellite and router tested and validated by international collaboration and demonstration of Virtual Mission Operations Center (VMOC) at Vandenberg in June 2004.
- IPv6 and IPsec tested in orbit, March 2007.
- Cisco router still works after five years in space.
- Router for geostationary satellite being reading for launch (IRIS on Intelsat-14, late 2009).

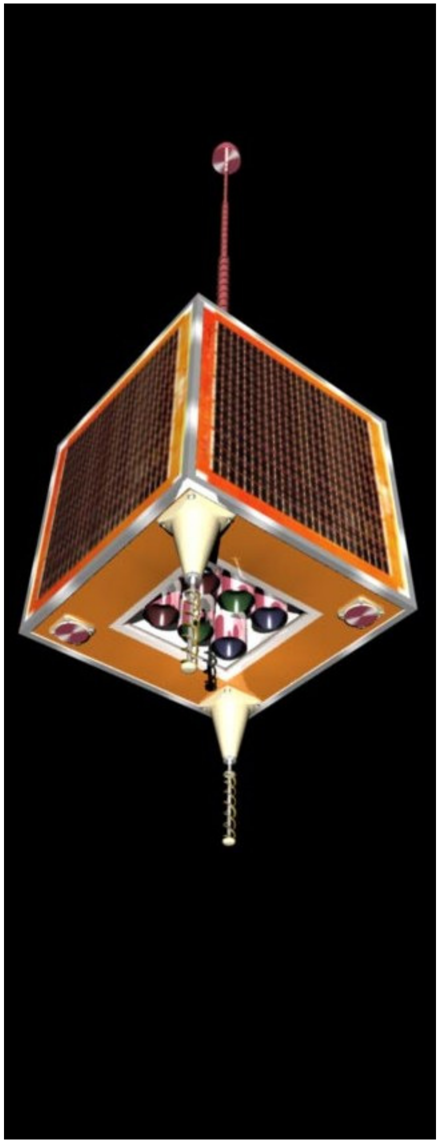

Existing network environment for the DMC

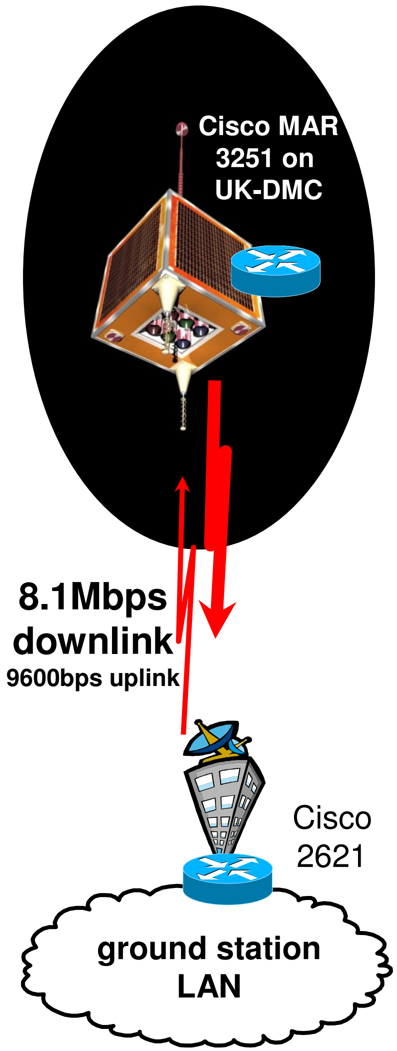

Satellite: each DMC satellite has multiple onboard computers. For housekeeping (the On Board Computer, OBC), for image capture and packetised transmission (the Solid State Data Recorders, SSDRs), for redundancy and survival. Interconnected by IP over 8.1Mbps serial links for data and slower CANbus for backup control; really a custom-built LAN.

CLEO: Cisco router was able to fit into UK-DMC satellite's onboard network by connecting to OBC and SSDRs using common serial interfaces.

Ground: SSTL's design for its ground station LANs uses IP. Satellites communicate with PCs on LAN via S-band radio space-ground link. IP over 8.1 Mbps serial stream from downlink commercial modem goes into a rack-mounted Cisco 2621 router, which forwards IP packets onto the LAN. SSTL's ground station LAN is connected to and an integral part of SSTL's corporate IP network.

CLEO router in payload tray

CLEO mobile router and serial card talk to other computers onboard via simple serial links.

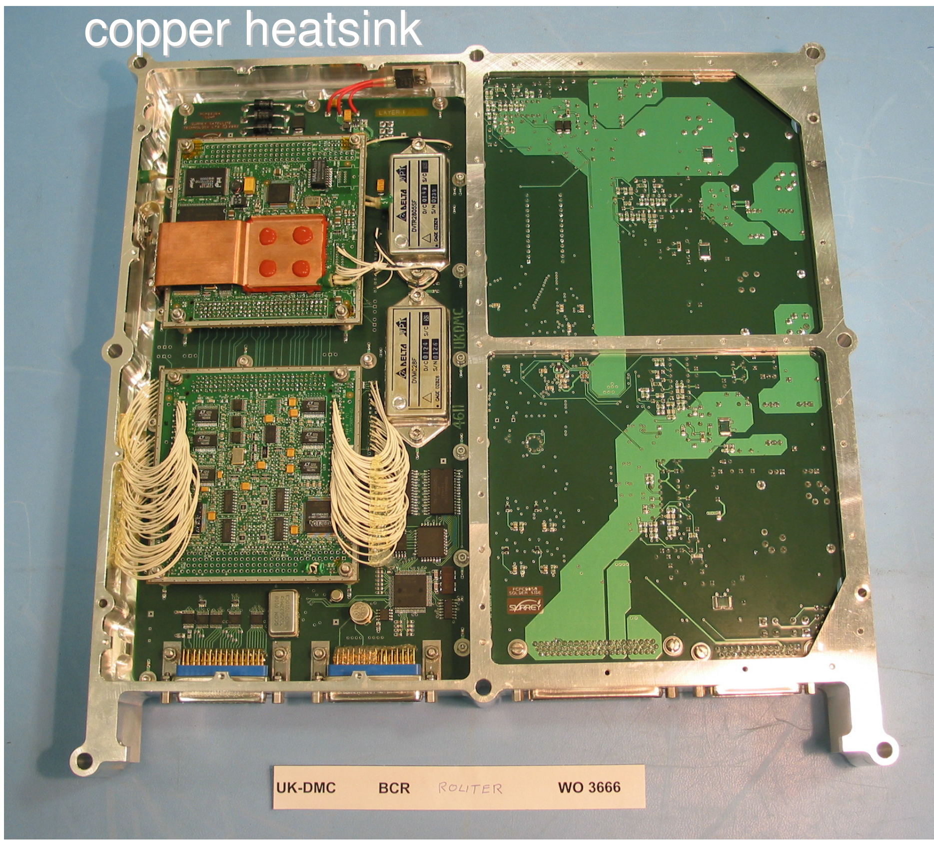

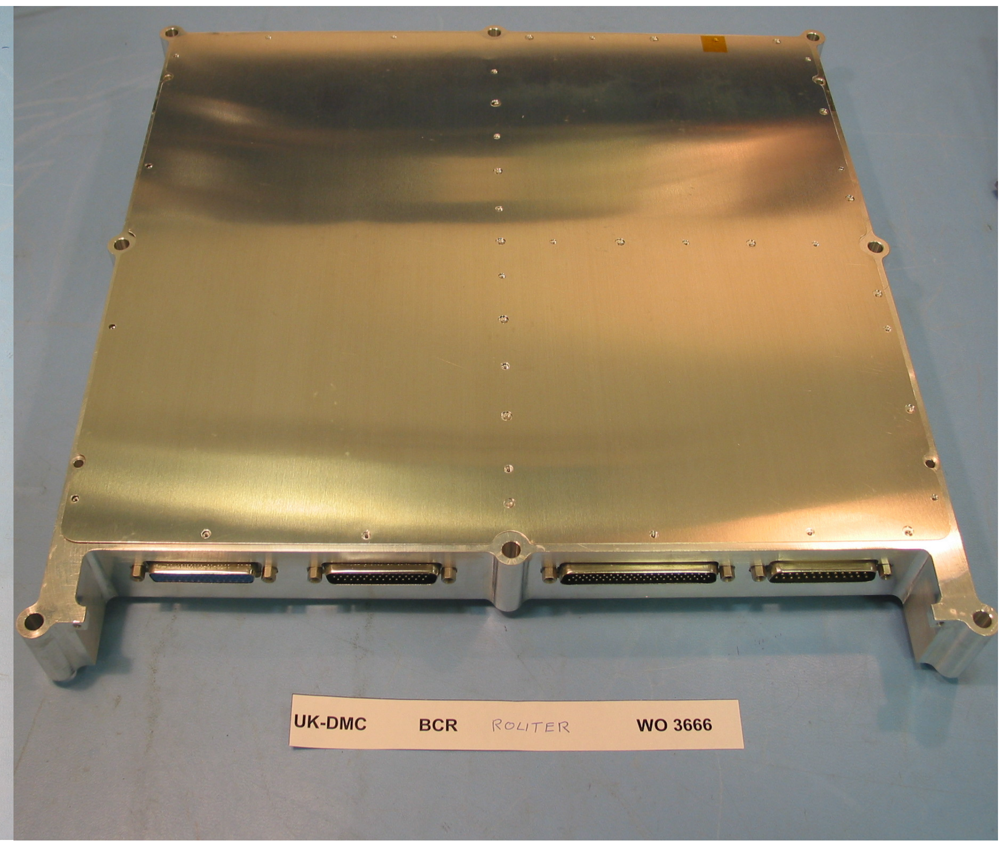

Ground testbed enables new development

- Ground-based testbed loaned to **NASA Glenn** was key to success of VMOC testing to deadline.
- IPv6 configs prepared and tested by NASA Glenn, before being copied to CLEO on orbit.
- Testbed now being used for software development.
- Delay-Tolerant Networking Research Group's 'Bundle Protocol' put on *Saratoga.*

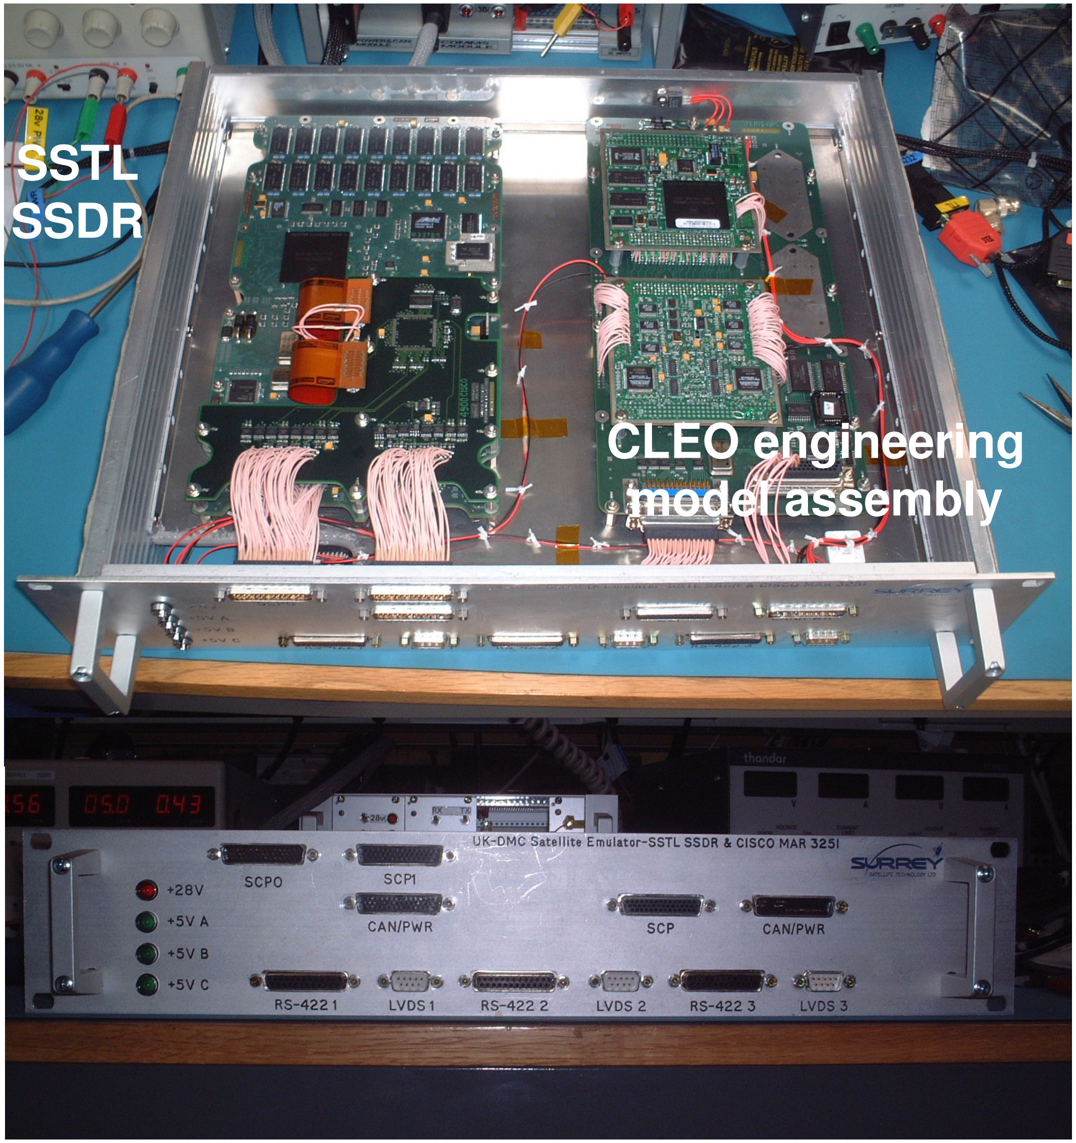

Short summary of *Saratoga*

- *Saratoga* is a simple and fast file transfer protocol ideal for file transfers over private links or for delay/disruption-tolerant networks.
- Developed and in use by Surrey Satellite Technology Ltd (SSTL) to transfer remote-sensing imagery from IP-based LEO satellites.
- NASA Glenn has improvements to the base *Saratoga* design to create a new version of *Saratoga:* **draft-wood-tsvwg-saratoga-02.txt**.
- We already have multiple implementations (in Perl, Python, and C, on Linux and RTEMS).
- We are developing and testing *Saratoga* with RTEMS-based computers on SSTL's UK-DMC satellite and in a ground-based testbed.

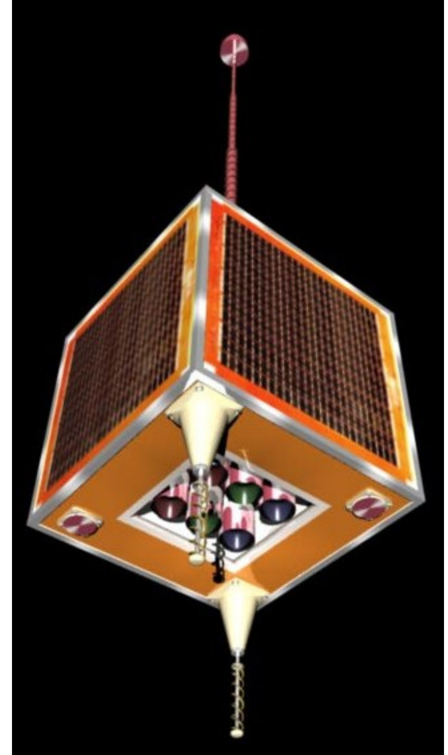

UK-DMC
satellite

How is *Saratoga* used in DMC operations?

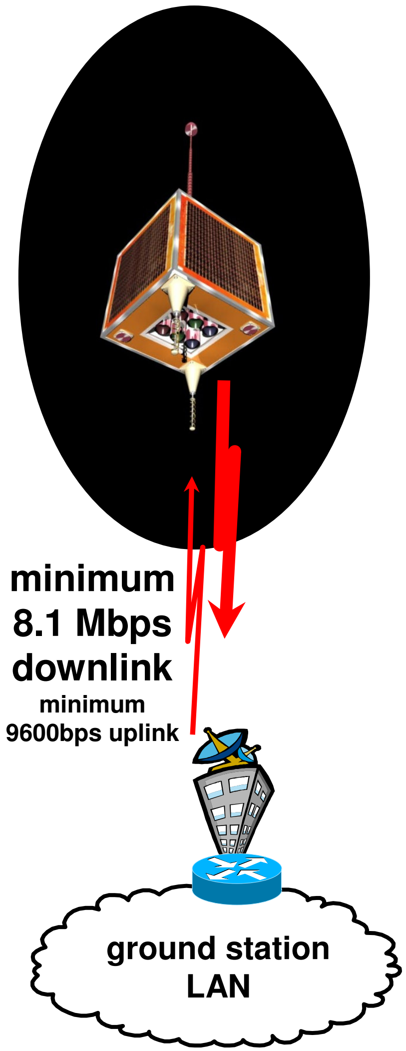

Each DMC satellite has multiple onboard computers. The Solid State Data Recorders (SSDRs) control cameras and store and download images using *Saratoga* over UDP/IP.

DMC downlink for image files is a minimum of 8.1Mbps. Newer satellites also have 20/40 Mbps X-band downlinks for added hi-res cameras; faster downlinks (100+ Mbps) are planned for future missions. Uplink is only 9600bps for command and control. Uplink speeds are also likely to increase… but only to 38400 bps.
Very asymmetric; 850:1 or worse downlink/uplink ratio.

As much data as possible must be transferred during a pass over a ground station. Passes may be up to fourteen minutes, depending on elevation. At 8Mbps, that's approximately 650MB of useful data (about a CD-ROM-full) that can be transferred in a high pass – if the downlink is filled with back-to-back packets at line rate. Link utilization and efficiency *really matter*. SSDRs take scheduled turns filling link.

Our approach to DTN networking

- **We believe that the Internet Protocol (IP) is useful for operational use in delay or disruption-tolerant networks.** Being convenient and cheap are compelling reasons to use IP for DTN. IP runs over many links already. Implementing support for custom "DTN bundle" convergence layers directly over all these links simply isn't scalable or cost-effective. Many IP-based protocols can be reused for DTN.
- The DMC is an example of using IP both on the ground and in space, with the ground station acting as a gateway between different types of network links.
- How IP is used differs between ground and space (link use, shared contention *vs* dedicated scheduling models – this discourages TCP reuse) but the base IP protocol remains the same. DMC satellites provide a real DTN scenario, with long disruptions between passes over ground stations.

Saratoga can provide reliable transfers.

- *Saratoga* always uses the UDP checksum to cover header and payload. This is consistent but not that strong (one's-complement), and not end-to-end.
- An end-to-end MD5 checksum over the entire file being transferred increases confidence that a reliable copy has been made, or that fragments have been reassembled correctly. Strong link-layer checksums are optional.
- The IRTF DTN research group Bundle Protocol and its convergence layers lack reliability checks.
- We have spent time examining that shortfall and proposing ways of adding reliability back to the DTNRG bundle protocol.

Why *Saratoga* instead of TCP?

- For high throughput and link utilization on dedicated links, where a single TCP flow cannot fill the link to capacity.
- For links and link use where TCP's assumptions about loss/congestion/competition simply don't hold.
- Able to cope with high forward/back asymmetry (>850:1).
- Long delay use – eventually TCP will fail to open a connection because its SYN/ACK exchange won't complete. TCP has many unwanted timers.
- Simplicity. TCP is really for a conversation between two hosts; needs a lot of code on top to make it transfer files. A focus on just moving files makes *e.g.* sequence nos. simpler. Having SNACKs means that handling sequence number wraparound when streaming becomes easy.

Saratoga is not content-based

- Saratoga transfers files or streams; its limited metadata support does not identify content with MIME. (DTN “bundles” don’t support MIME. MIME helped email and the web succeed.)
- We have proposed a way of reusing HTTP for delay-tolerant networks – HTTP-*DTN* **draft-wood-dtnrg-http-dtn-delivery-01**.
- That gives one way to identify content with MIME for applications to pick up on, using the well-understood HTTP interface. HTTP*-DTN* can be carried over *Saratoga* or TCP streams, depending on the link environment being used.

Why is it called *Saratoga*?

Photo # 19-N-84312 USS Saratoga underway in Puget Sound, 15 May 1945

USS Saratoga **(CV-3) is sunk off Bikini atoll.**
Chris Jackson of SSTL has dived there.

Bundle Protocol tests in space

- On Surrey Satellite Technology's UK-DMC satellite, in January and September 2008. Used Bundle Protocol over *Saratoga*.

  downloaded real operational sensor data, transferred fragments across Internet from Surrey to NASA Glenn.

- On NASA JPL EPOXI (Extrasolar Planet Observation and *Deep Impact* Extended Investigation) comet probe, October 2008. Used Bundle Protocol over LTP over CFDP over lots of stuff. DINET – Deep Impact DTN Experiment

  uploaded pictures to probe, got them back again. Implemented ground network simulating other probes.

Sensor data we downloaded using Bundle Protocol

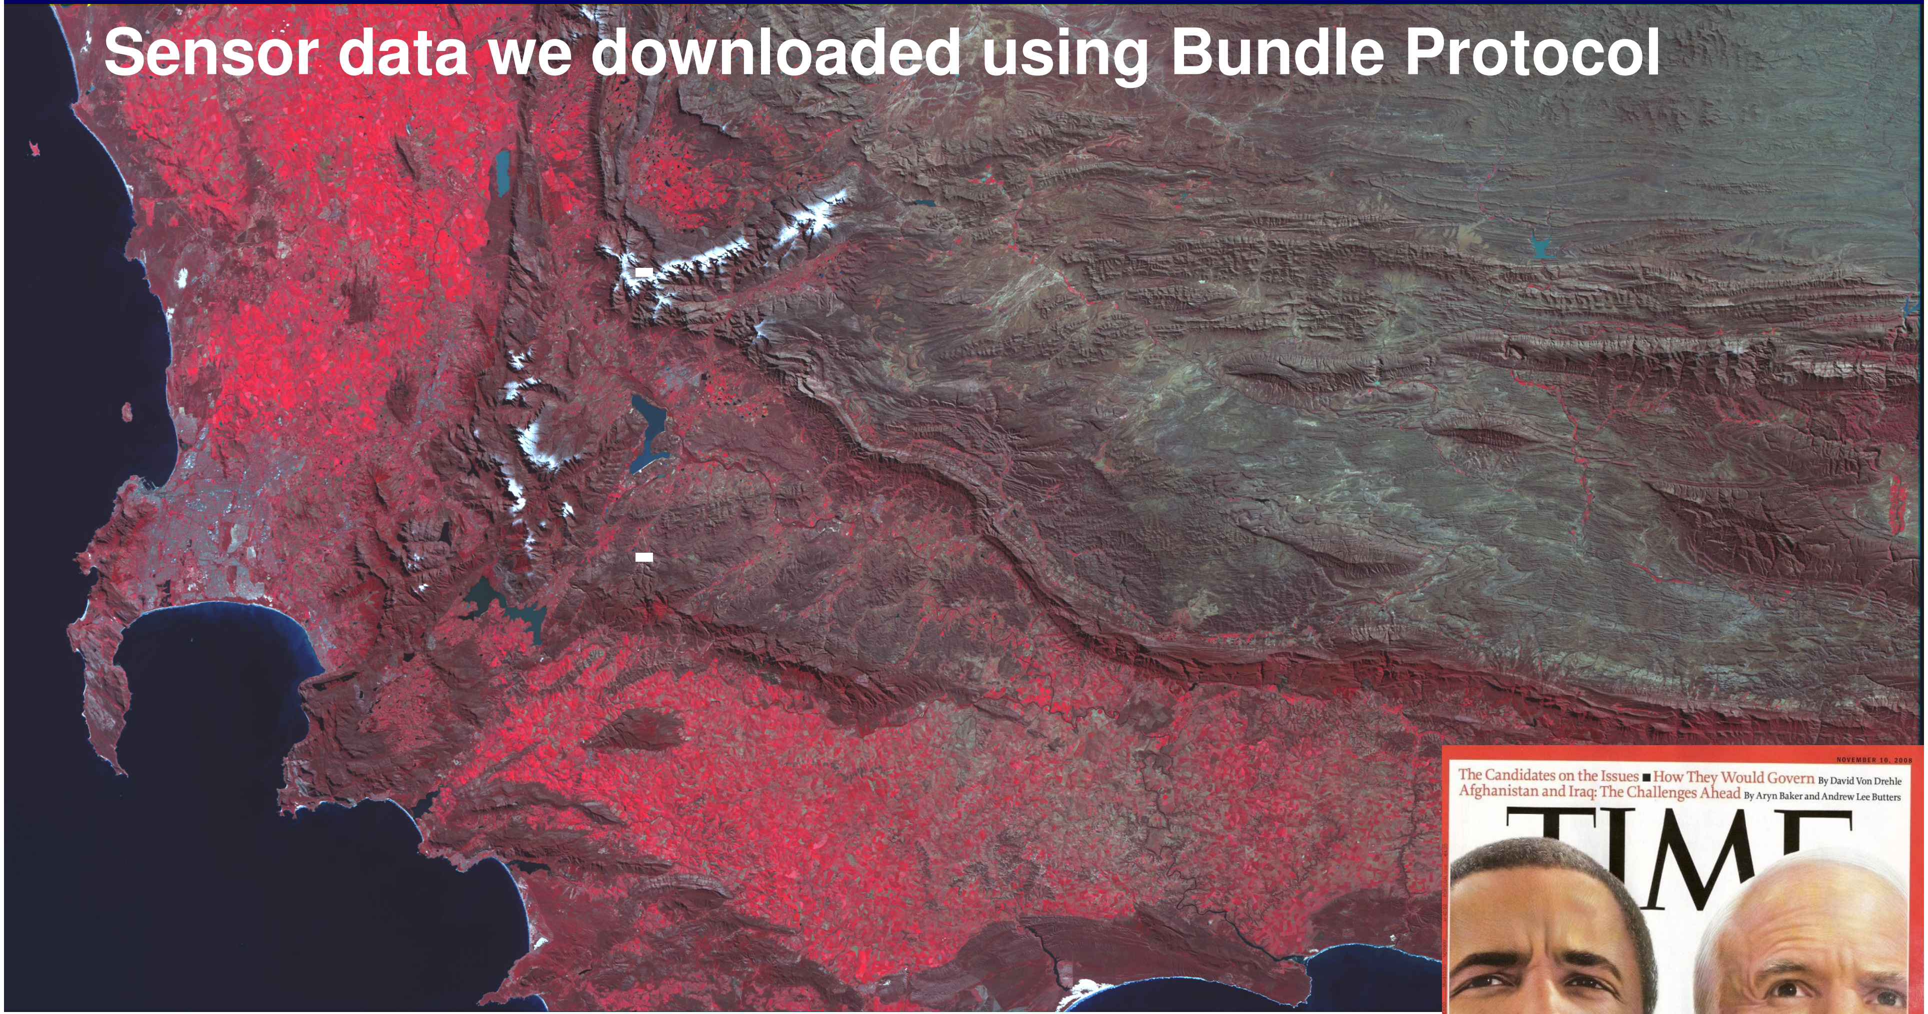

150MB image transferred from UK-DMC satellite using Bundle Protocol over *Saratoga* with proactive fragmentation, 25 August 2008.

TIME Magazine best inventions of the year **#9 Orbital Internet**, 10 November 2008 issue – before EPOXI tests announced.

Example picture of Licklider sent to EPOXI and back

- Uplink at 250 bytes/sec, downlink at either 110 or 20,000 bytes/sec.
- One-way propagation delay of 81s, down to 49s four weeks later.
- About 300 images relayed through Deep Impact and back as a ‘DTN router’

Networking stacks used in these experiments

UK-DMC tests, Jan/Sep 2008

after Hogie. Max possible bundle size: 4GB

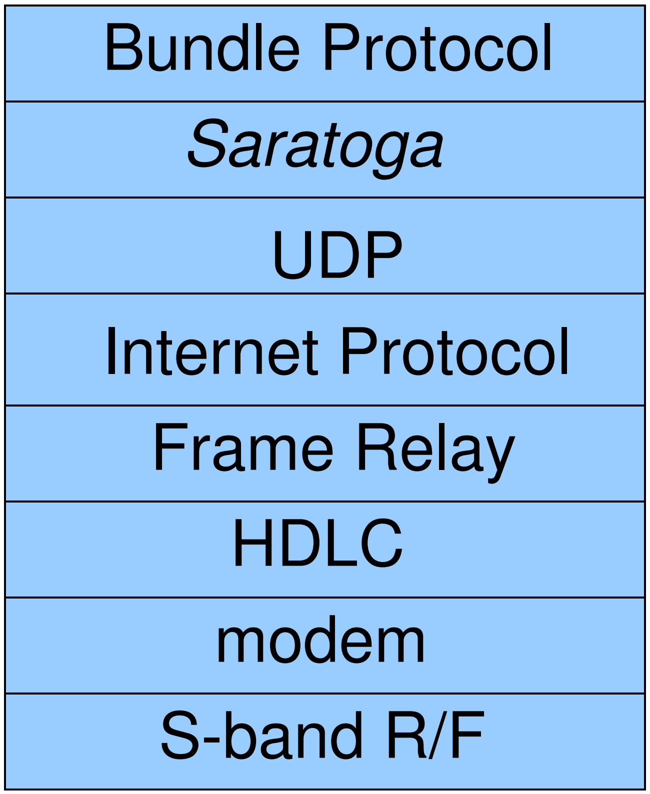

Bundle security not implemented onboard either spacecraft.

Deep Impact, Oct 2008

after Burleigh. Max possible bundle: 64K

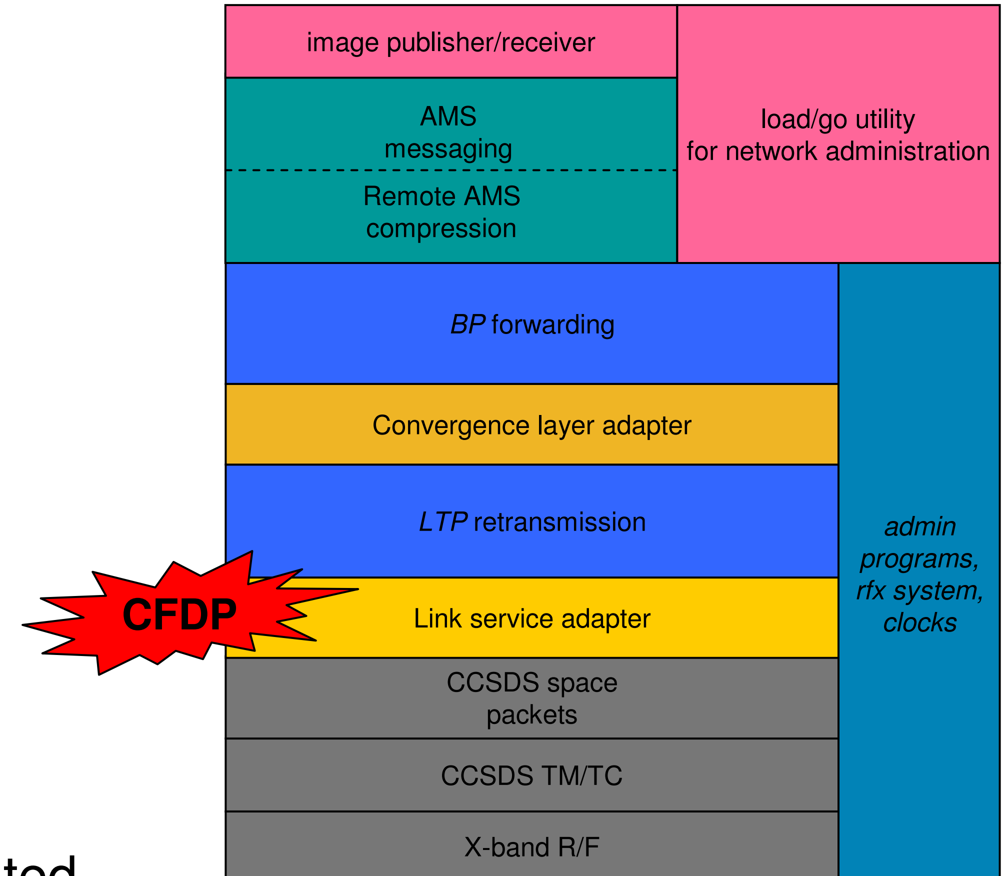

Scott Burleigh, IETF 73 DTNRG meet, 20 Nov 2008

We have discovered problems with bundling

- **Reliability**. No error detection, and reusing security to give reliability is not ideal. Errors seen Jan 2008.
- **Timing**. Every bundle agent is expected to know current UTC time. This has limits in space (relativity, eventually). Leap seconds must be communicated. Synchronization is a problem; bundles can get dropped as expired (Jan 2008).
- **Convergence layer adapters.** Pretty much all use and deployment is over IP – *except* for CCSDS.
- naming schemes/routing/QoS/management.
- No content identification *a la* MIME and HTTP…

An alternative to bundling: HTTP*-DTN*

- MIME describes the things we move around the network. The most successful protocols support MIME.
- HTTP is the simplest MIME wrapper.
- HTTP provides infinitely-flexible text metadata.
- Use HTTP hop-by-hop between neighbouring DTN nodes.

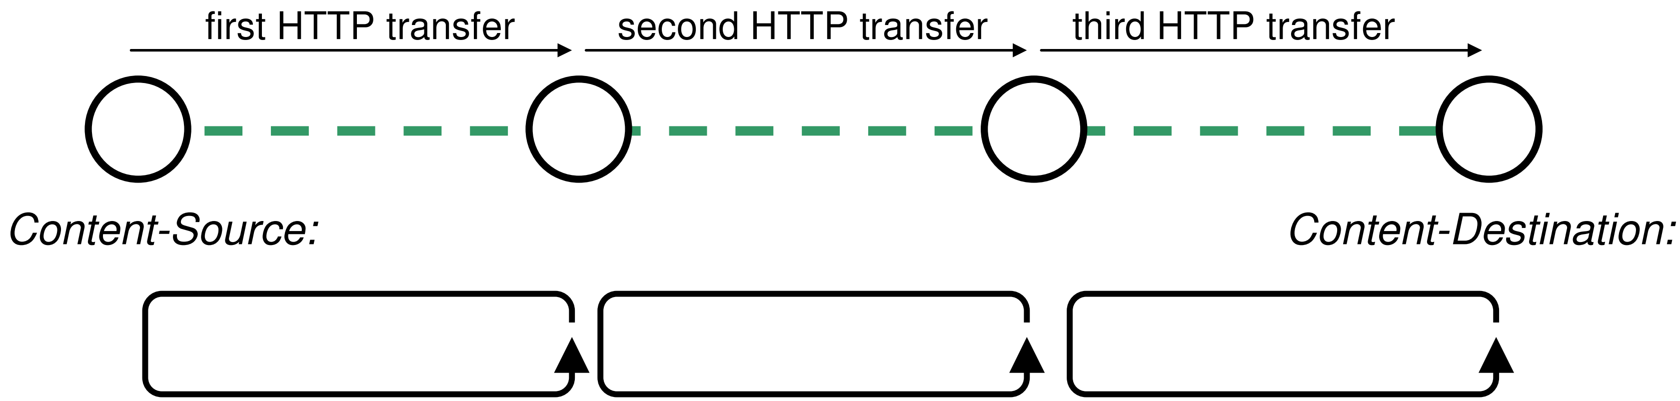

- Allow HTTP to be run over different transports: TCP, SCTP, *Saratoga*… HTTP can be separated from TCP's limitations. Divide HTTP from transport to make a true session layer. What HTTP requires from transport isn't that onerous.

What makes HTTP*-DTN* special?

- Two new **Content-*** headers:

  Content-Source: where the object is originally from

  Content-Destination: final destination
- Basic HTTP rule: **Content-*** headers are special . If Content-blah is unfamiliar, reject the transfer.
- This makes HTTP*-DTN* separate from, and not polluting, existing web. Unlikely to alarm W3C.
- Optional e2e reliability over payloads by reusing existing **Content-MD5:** header or similar.
- Header/metadata reliability a bit trickier – may need new headers. HTTP already supports ‘per hop’ limited-scope headers.
- New Package- headers can *package* related objects together, track if they’ve all arrived or not.

HTTP-DTN is the waist in *this* hourglass

HTTP is the universal session glue.
choose the transport to suit the conditions; TCP in traditional Internet, *Saratoga* for high performance on dedicated links. Separate the session control from transport, link and traffic conditions.

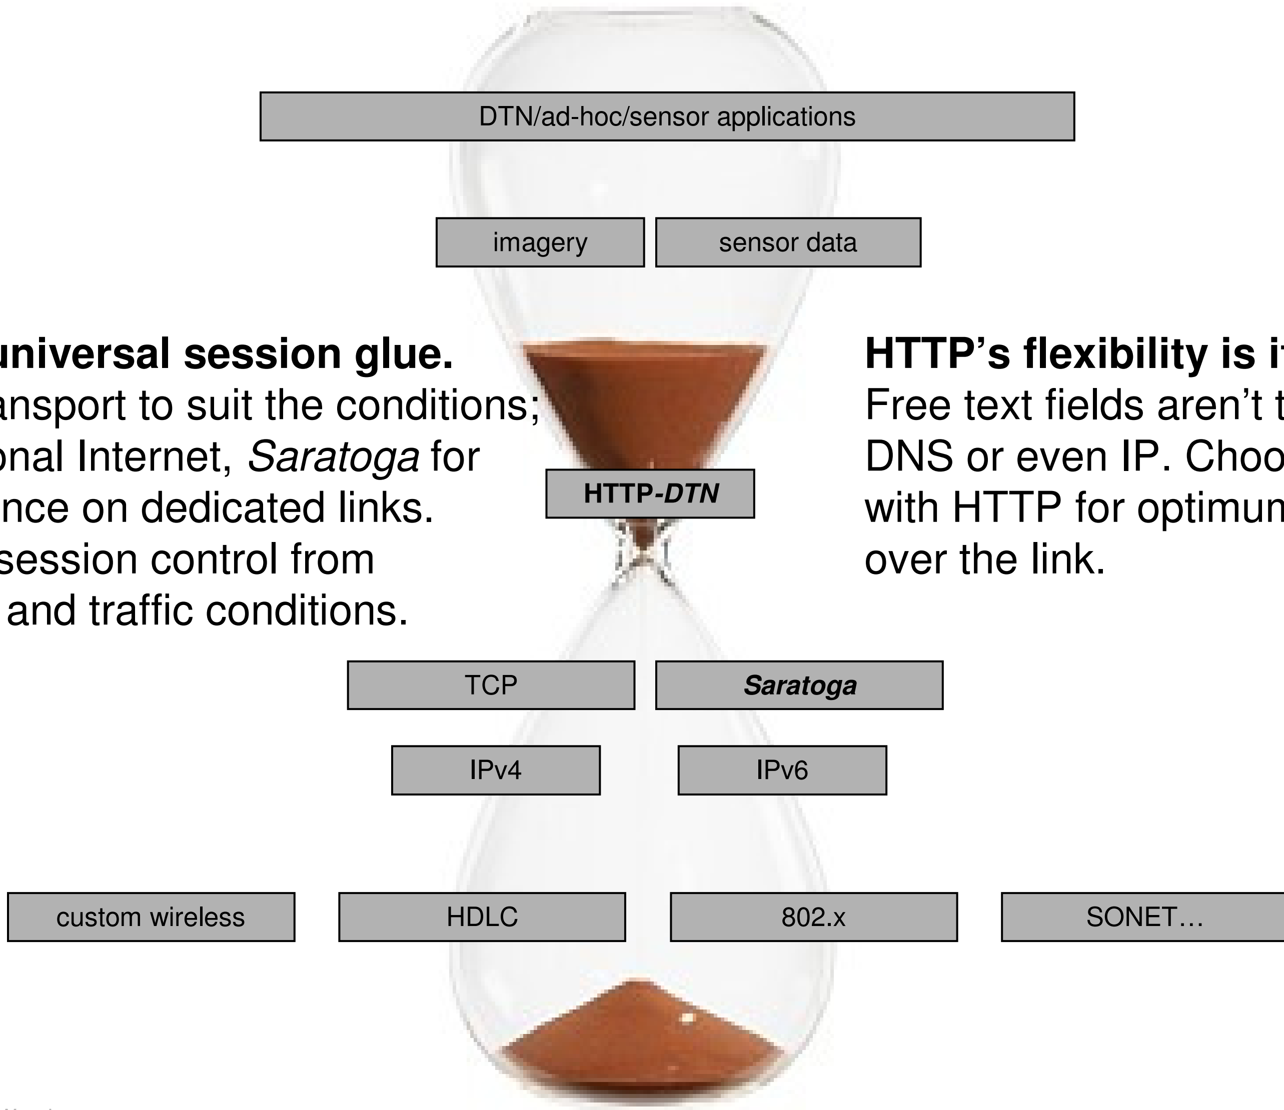

HTTP's flexibility is its strength
Free text fields aren't tied to TCP, DNS or even IP. Choose what to use with HTTP for optimum performance over the link.

HTTP-*DTN* advantages

- Text fields aren't tied to IP, TCP or to DNS. Could implement HTTP over own stack, with own routing namespace, etc.
- Doesn't require a two-way session; HTTP PUT can be entirely unidirectional.
- Reuses large body of existing code and well-understood functionality. Only minor changes.
- Possible to build on top of HTTP-*DTN* base to reuse pieces of web infrastructure, e.g. SOAP.
- Conceptually very *very* simple.

CCSDS vs IP/FR/HDLC for space – brief history in a slide

time

CCSDS stack world (NASA JPL) '70s-

- Custom protocols over CCSDS links.
- Standards by fiat – adopted by very small community, but are ISO standards as CCSDS is ISO subgroup.
- No clear layering, for optimisation of link budgets via high integration. Not modular, must be re-engineered when lower layers change waveform: costly.
- Claimed use on hundreds of (incompatible) missions, over decades.
- Legacy streaming, rather than packet/files. Missions do not use files.
- CCSDS links do not carry IP well.
- Expensive to (re)develop whenever link coding changes the waveform.

IP stack world (NASA Goddard/Glenn) '90s-

- UDP/IP over frame relay/HDLC.
- IETF and other commercial standards are widely used and implemented.
- *Made CLEO router in orbit and first DTN 'bundle' tests in space possible.*
- Fully layered and modular, for plug and play and simple integration. Just change modem to support a new physical coding.
- Used on over twenty missions since 1999, *e.g.* DMC satellites.
- Full support of packet/file transmission and networking concepts.
- Carries IP very cleanly and well.
- Cheap to develop; tools and knowledge widely available.

CFDP – CCSDS File Delivery Protocol. Like FTP with many added features. Intended to be used over both IP and over CCSDS – but is not in use over IP.

DTN Delay-Tolerant Networking bundles. Replaces CFDP. Intended to run over CCSDS and over IP via adaptation layers. Developed cheaply over IP. But CCSDS wants DTN over legacy CCSDS links and waveforms to migrate their userbase.

Some thoughts

- Standards for space are a good thing – but why does space need its own standards?
- Why doesn't CCSDS let ISO-standard HDLC run over ISO-standard CCSDS protocols?
- Is the Bundle Protocol a failure for not meeting the needs of its various problem spaces?
- Is such a complex and fragile bundle format suited to harsh errored ad-hoc conditions?
- What about the end-to-end principle?

For more information
google UK-DMC bundle

Lloyd's papers on this are at:
http://info.ee.surrey.ac.uk/Personal/L.Wood/publications/