

Saratoga

fast data transfer... from space



draft-wood-tsvwg-saratoga

Wood, Eddy, Smith, Ivancic, Jackson, McKim

NASA Glenn Research Center,
Surrey Satellite Technology Ltd (SSTL)

Lloyd Wood

June 2013

Saratoga is a transfer protocol that is...

- ***scalable***

Can transfer kilobytes, megabytes, gigabytes, petabytes, exabytes... suitable for exascale computing.

- ***fast***

intended to run as quickly as possible across dedicated links. (Can support congestion control for shared use.)

- ***reliable***


while providing high-speed delivery of content.

- ***robust***

Originally developed for use from space.

Short summary of *Saratoga*

- A simple, fast data transfer protocol ideal for file transfers across private links or for delay/disruption-tolerant networks (DTNs).
- In use by Surrey Satellite Technology Ltd (SSTL) to transfer remote-sensing imagery from its IP-based DMC satellites.
- NASA Glenn improved base *Saratoga* design to create a new version of *Saratoga* described to the Internet Engineering Task Force (IETF): **draft-wood-tsvwg-saratoga**.
- We have developed and tested *Saratoga* with RTEMS-based computers on SSTL's UK-DMC satellite and in a ground-based testbed.
- There are multiple protocol implementations.




Disaster Monitoring Constellation (DMC)

www.dmcii.com

SSTL build and help operate an international constellation of small sensor satellites.

These satellites orbit in a sun-synchronous orbital plane for rapid daily large-area imaging (640km swath width with 32m resolution). Can observe effects of natural disasters. Imaged the effects of Hurricane Katrina and the Indian Ocean Tsunami.

Government co-operation. Each government finances a ground station in its country and satellites. Ground stations are networked together. Newer satellites are added to the constellation over time. Now on second generation of satellites, with added commercial investment.



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

DMC in use: after Hurricane Katrina, 2005



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


Small part of an image from the NigeriaSat-1 satellite, taken on Friday 2 September 2005, for US Geological Survey.

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

All imagery delivered by using *Saratoga* over UDP/IP.

***Saratoga* is in daily operational use.**

Saratoga use in satellite 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 using S-band. Newer satellites also have 20/40/80 Mbps at X-band for added hi-res cameras; faster downlinks (100+ Mbps) are in use on second generation. Uplink is a minimum of 9600bps for command and control. Uplink speeds have also increased... 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 from first generation, that's about 650MB of useful data that can be transferred in a high pass – if the downlink is filled at line rate with back-to-back packets. Link utilization and efficiency *really matter*. SSDRs take scheduled turns filling downlink.

Simple delivery mechanism



STATUS acknowledgement pacing is set by request pacing



-  DATA packet
-  DATA packet requesting STATUS ack update
-  STATUS packet sent as requested

Basic *Saratoga* design

- Flood data packets out *as fast as you can*. No specified congestion control, since usual use has data only going one hop over a private link. (No specified timers means no *timeouts*, so *Saratoga* will also be good for the very long propagation delays found in deep space.)
- Every so often *while flooding*, ask for acknowledgement STATUS from file receiver. Receiver can also send acks if it thinks it needs to, or to start/restart/finish transfer.
- STATUSes are Selective Negative Acknowledgements (SNACKs) indicating received packets and any gaps to fill with resent data (and with enough information so that intelligent sender rate control or congestion control can be optionally added if needed locally).
- Other telemetry can be multiplexed in as IP traffic.
- That's it. But just how big are the files being transferred?

File sizes can be *large*, streams can be *fast*

- For the DMC satellites, imaging files are big – typically up to a few gigabytes at 32m resolution; larger for newer cameras. So we believe there is a need to transfer *large* amounts of sensor data – gigabytes and up. This will only increase over time as sensors and network links improve.
- But ad-hoc/sensor nets also need to transfer small files; guessing a range limits use.
- So we allow a range of file-descriptor pointers to be advertised: 16/32/64/128-bit file descriptors.
- If file is less than 64KiB, use 16-bit offsets. If file is larger but less than 4GiB, use 32-bit offsets...
- 16-bit offsets are *always* supported. Others are optional. Draft diagrams are all 32-bit, because that fits 80 columns.
- 128-bit descriptors are useful for petabyte-sized file transfers or for fast high-speed streams. Future-ready.

Saratoga packets

BEACON

Sent periodically. Describes the *Saratoga* peer:

- Identity (e.g. EID)
- capability/desire to send/receive packets.
- max. file descriptor handled: 16/32/64/128-bit.

REQUEST

Asks for a file via 'get', directory listings, deletes.

METADATA

Sent at start of transaction.

Describes the file/bundle:

identity for transaction

file name/details, including size.

descriptor size offsets to be used for this file
(one of 16/32/64/128-bit pointer sizes.)

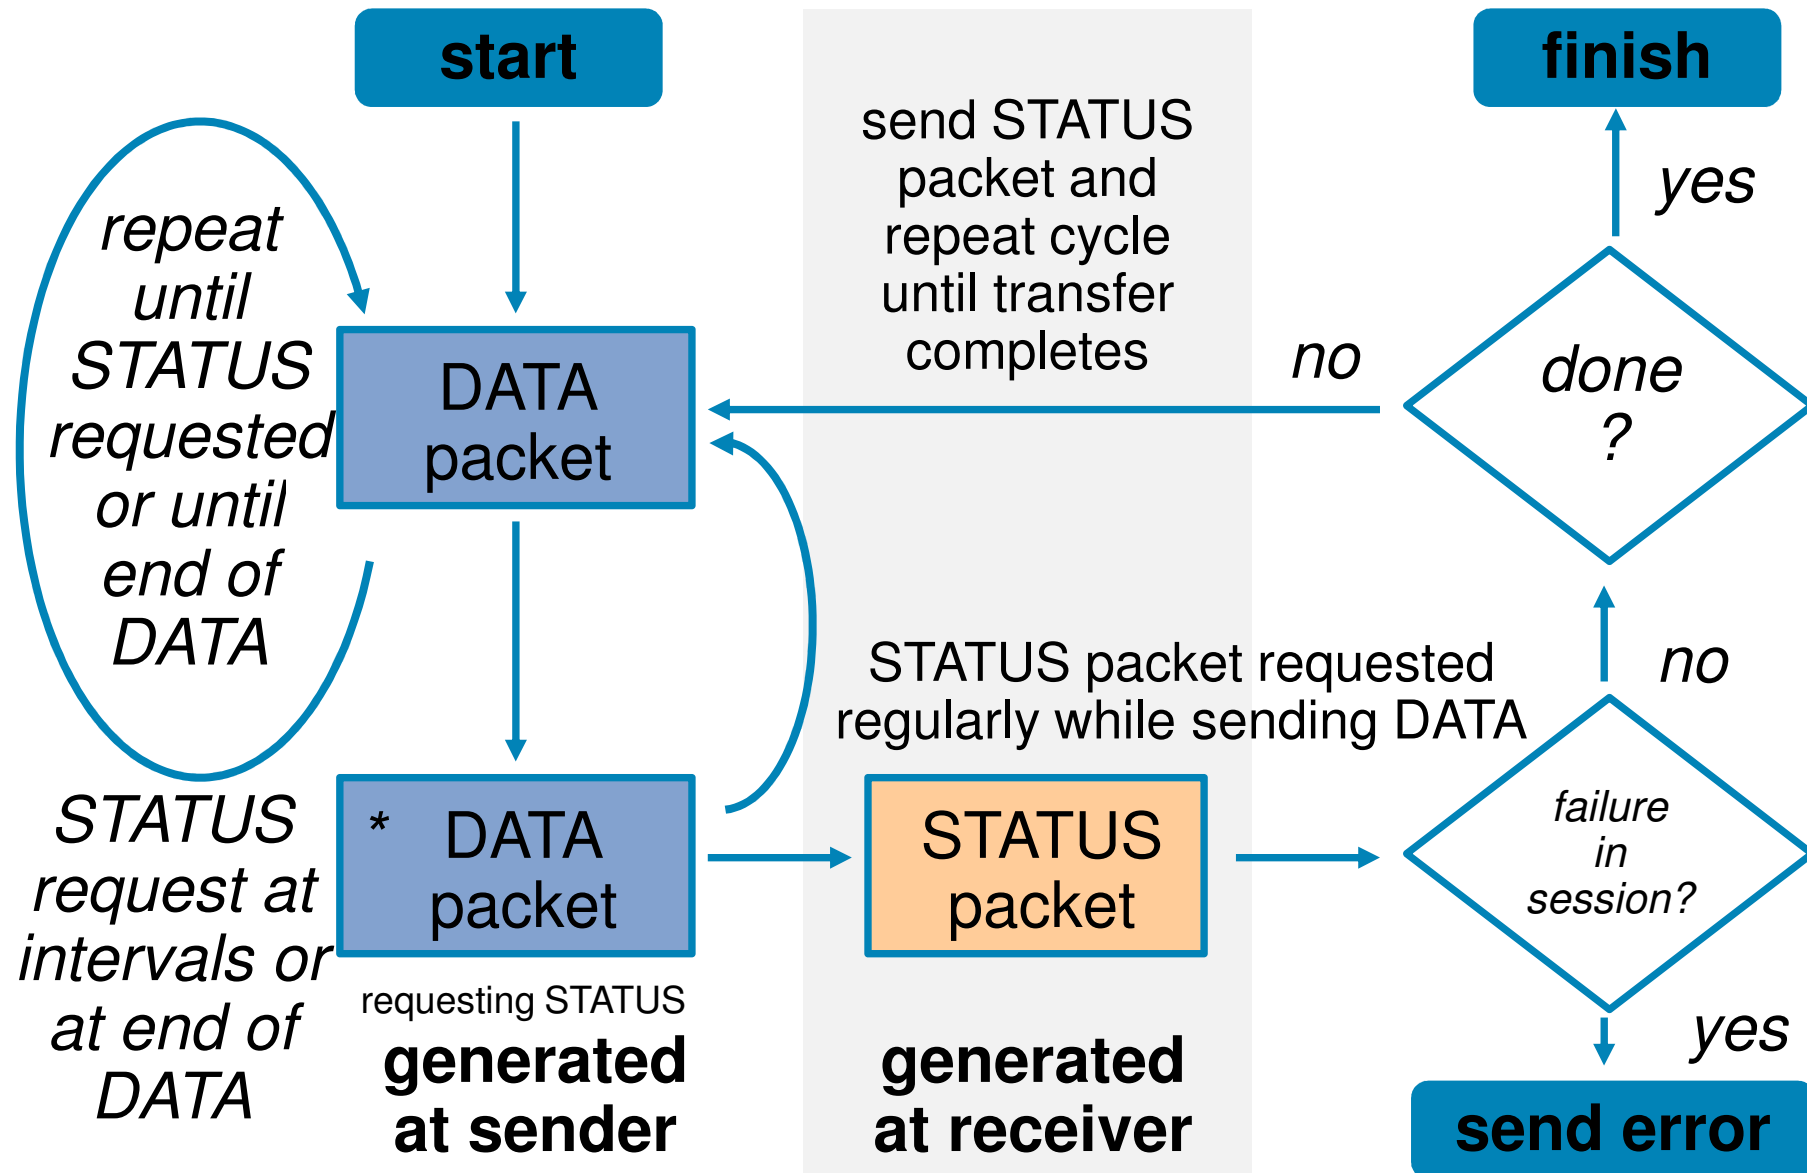
DATA

Uses descriptor of chosen size to indicate offset for data segment. May request an ack.

STATUS

Selective negative ack (SNACK). 'Holestofill'
hole edge offsets indicate missing 'holes' in data.

Heart of a *Saratoga* transfer session



Congestion control can be added

- Enough information is sent and received between peers to support congestion control, if needed.
- Congestion control is not mandated; primary use case across dedicated private link filled by single flow doesn't need it.
- But this could be implemented if required for traffic across the public, shared, terrestrial Internet.
- Sender-based 'TCP friendly' congestion control mechanisms have been studied.
- Described in **draft-wood-tsvwg-saratoga-congestion-control** and **draft-eddy-tsvwg-saratoga-tfrc**.
- Simply rate limiting *Saratoga* to a set speed, rather than the full output rate of an interface, is of more immediate interest to private users for predictable performance.

Other optional *Saratoga* features

- Streaming of data is supported. This allows *Saratoga* to be used for real-time delivery outside the file-based paradigm.
- Uses link-local multicast to advertise presence and discover peers. Data can be sent to local multicast addresses for multiple peers to receive.
- Can support checksums for reliable file delivery.
- Also supports UDP-Lite for when errors in data received can be tolerated. Header content is always checked so that the information *about* the data is error-free. UDP-Lite use is expected to be rare; reliable transfers of data are the norm.

Optional DTN bundling

- *Saratoga* can also support Bundle Protocol delivery for Delay-Tolerant-Networking (DTN), as a reliable “bundle convergence layer.”
- This is described in **draft-wood-dtnrg-saratoga**.
- NASA Glenn conducted the first tests in space of the ‘Interplanetary Internet,’ from SSTL’s UK-DMC satellite, by sending DTN bundles of imaging data from orbit using *Saratoga* (2008).
- Bundle support is now seen as of little value to *Saratoga* users. Not a key feature for these DTNs.

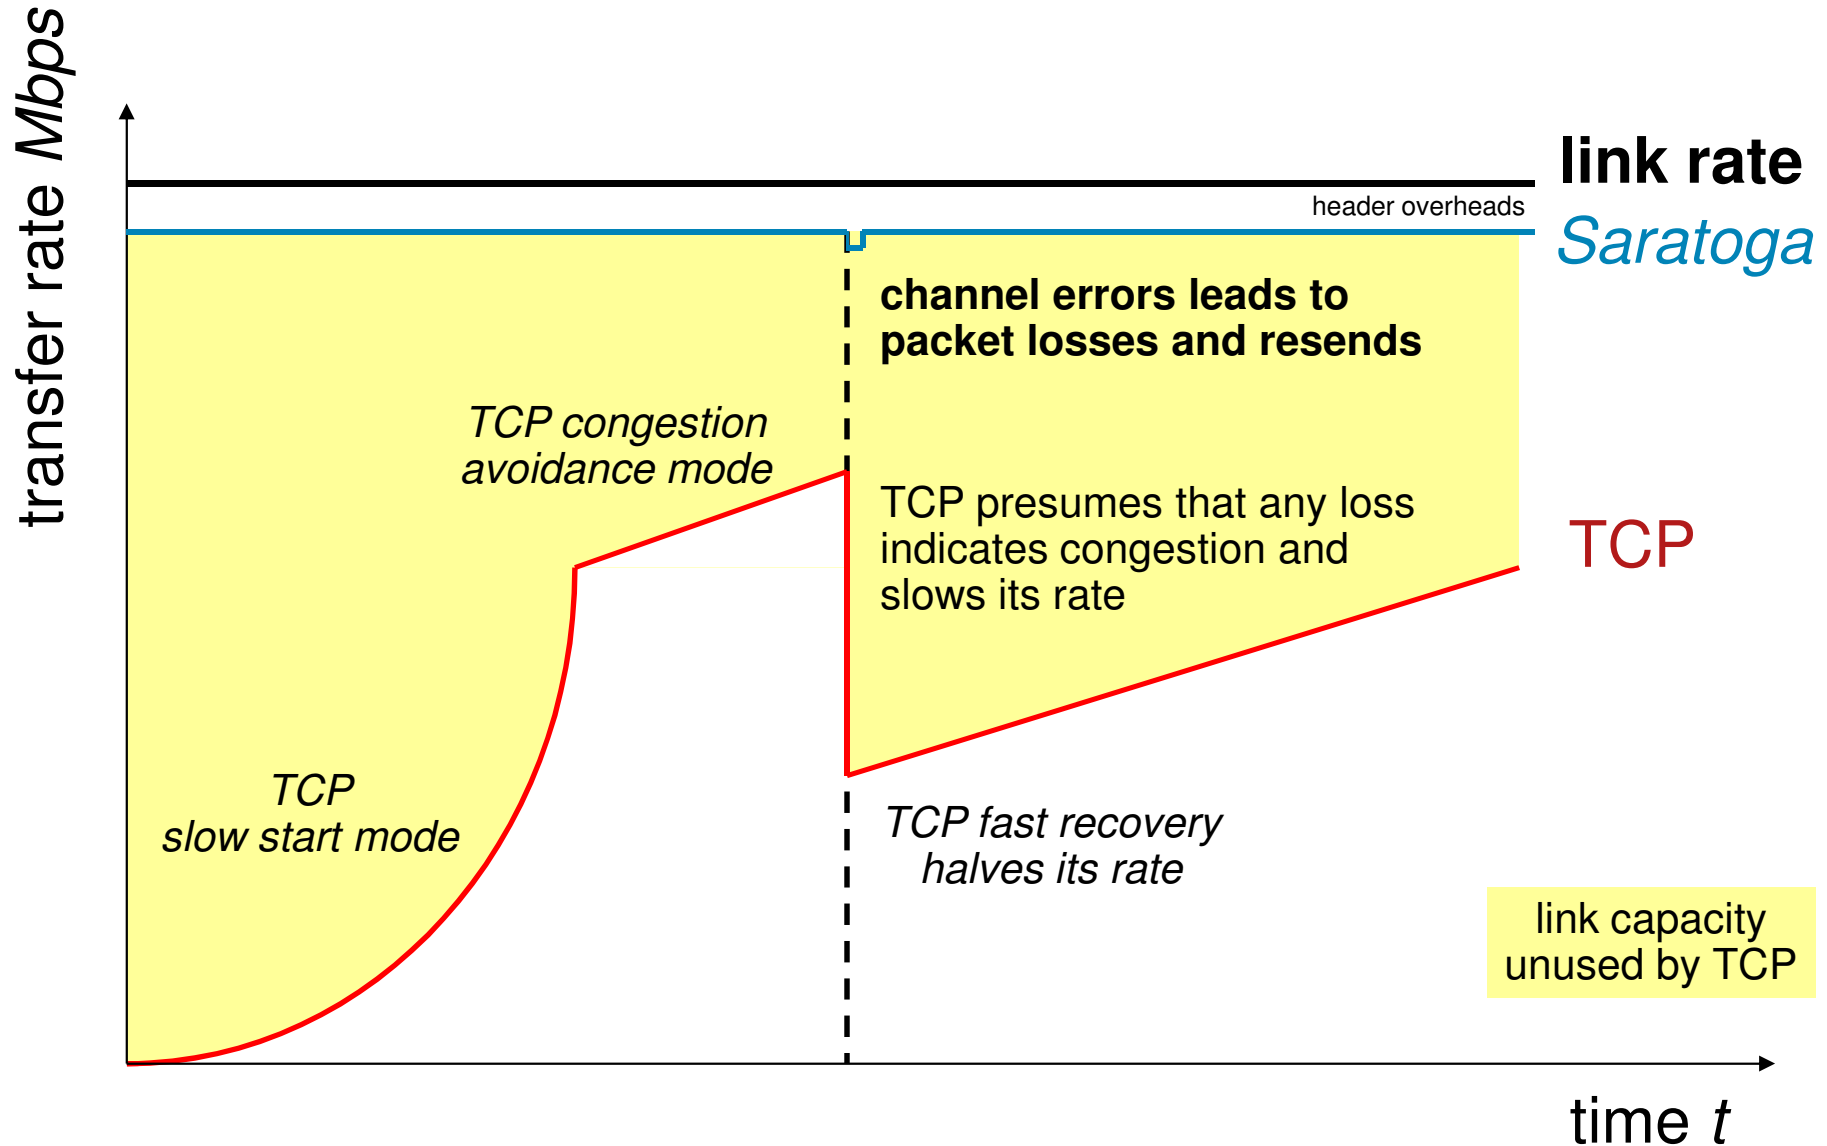
Saratoga can provide reliable transfers

- *Saratoga* 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 optional end-to-end checksum (CRC32c, MD5 or SHA-1) over an entire file being transferred can increase confidence that a reliable copy has been made, or that file fragments have been reassembled correctly. Strong link-layer checksums are optional.
- The Delay-Tolerant Networking research group “Bundle Protocol” and its convergence layers lack any reliability checks. We have spent time examining that shortfall and proposing ways of adding reliability into 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 vs TCP – a single flow



Our approach to DTN networking

- **We believe that the Internet Protocol (IP) is useful for operational use in delay or disruption-tolerant and in sensor networks.** IP runs across many links already; its engineering is well-understood.
- How IP is used differs between the public Internet and private space links (shared contention *vs* dedicated scheduling models – this discourages TCP reuse) but the base IP protocol is the same.
- DMC shows use of IP both on the ground and in space, with the ground station acting as a gateway between different kinds of link behaviour and link use. DMC satellites provide a real DTN scenario, with long disruptions between contacts.

Possible applications of *Saratoga*

- In *any* private IP-based network where data must be moved as rapidly as possible over dedicated links.
- In private sensor networks generating large quantities of data and moving it from sensors to central processing, such as in remote-sensing satellites and for radio astronomy.
- In delay-tolerant networks (DTNs).

Current status of *Saratoga*

- SSTL's *Saratoga* remains in daily use on DMC satellites for mission-critical downloads of Earth imagery. Ground and space implementations.
- Other implementations underway.
- Public *Saratoga* perl implementation and Wireshark plugin for interoperability testing of private implementations.

What prompted the name *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.**



<http://tools.ietf.org/html/draft-wood-tsvwg-saratoga>

<http://saratoga.sf.net/>

with thanks to
Will Ivancic, Wes Eddy, Jim McKim
Chris Jackson and Charles Smith



NigeriaSat-2 (launched 2011)
DMC second generation