Saratoga: Efficient Transport over Short-Lived Links

Wesley M. Eddy
Verizon / NASA
weddy@grc.nasa.gov

with:
Lloyd Wood (Cisco),
Jim McKim (RSIS / NASA),
Will Ivancic (NASA),
Chris Jackson (SSTL)

http://www.ietf.org/internet-drafts/draft-wood-dtnrg-saratoga-01.txt

Original Use Case

- Large file downloads of images from SSTL's Disaster Monitoring Constellation (DMC) satellites
- DMC satellites use IP for communications
  - in daily operation since 2002
  - 5 satellites currently in orbit (at least 3 more planned)

Beiing-1

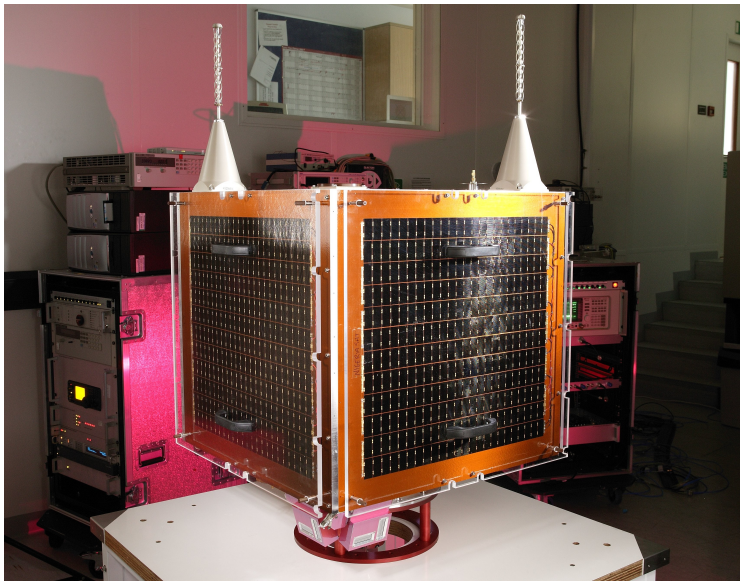

UK-DMC / NigeriaSat

USGS
science for a changing world

dmc
International Imaging

Disaster Monitoring Constellation (DMC)

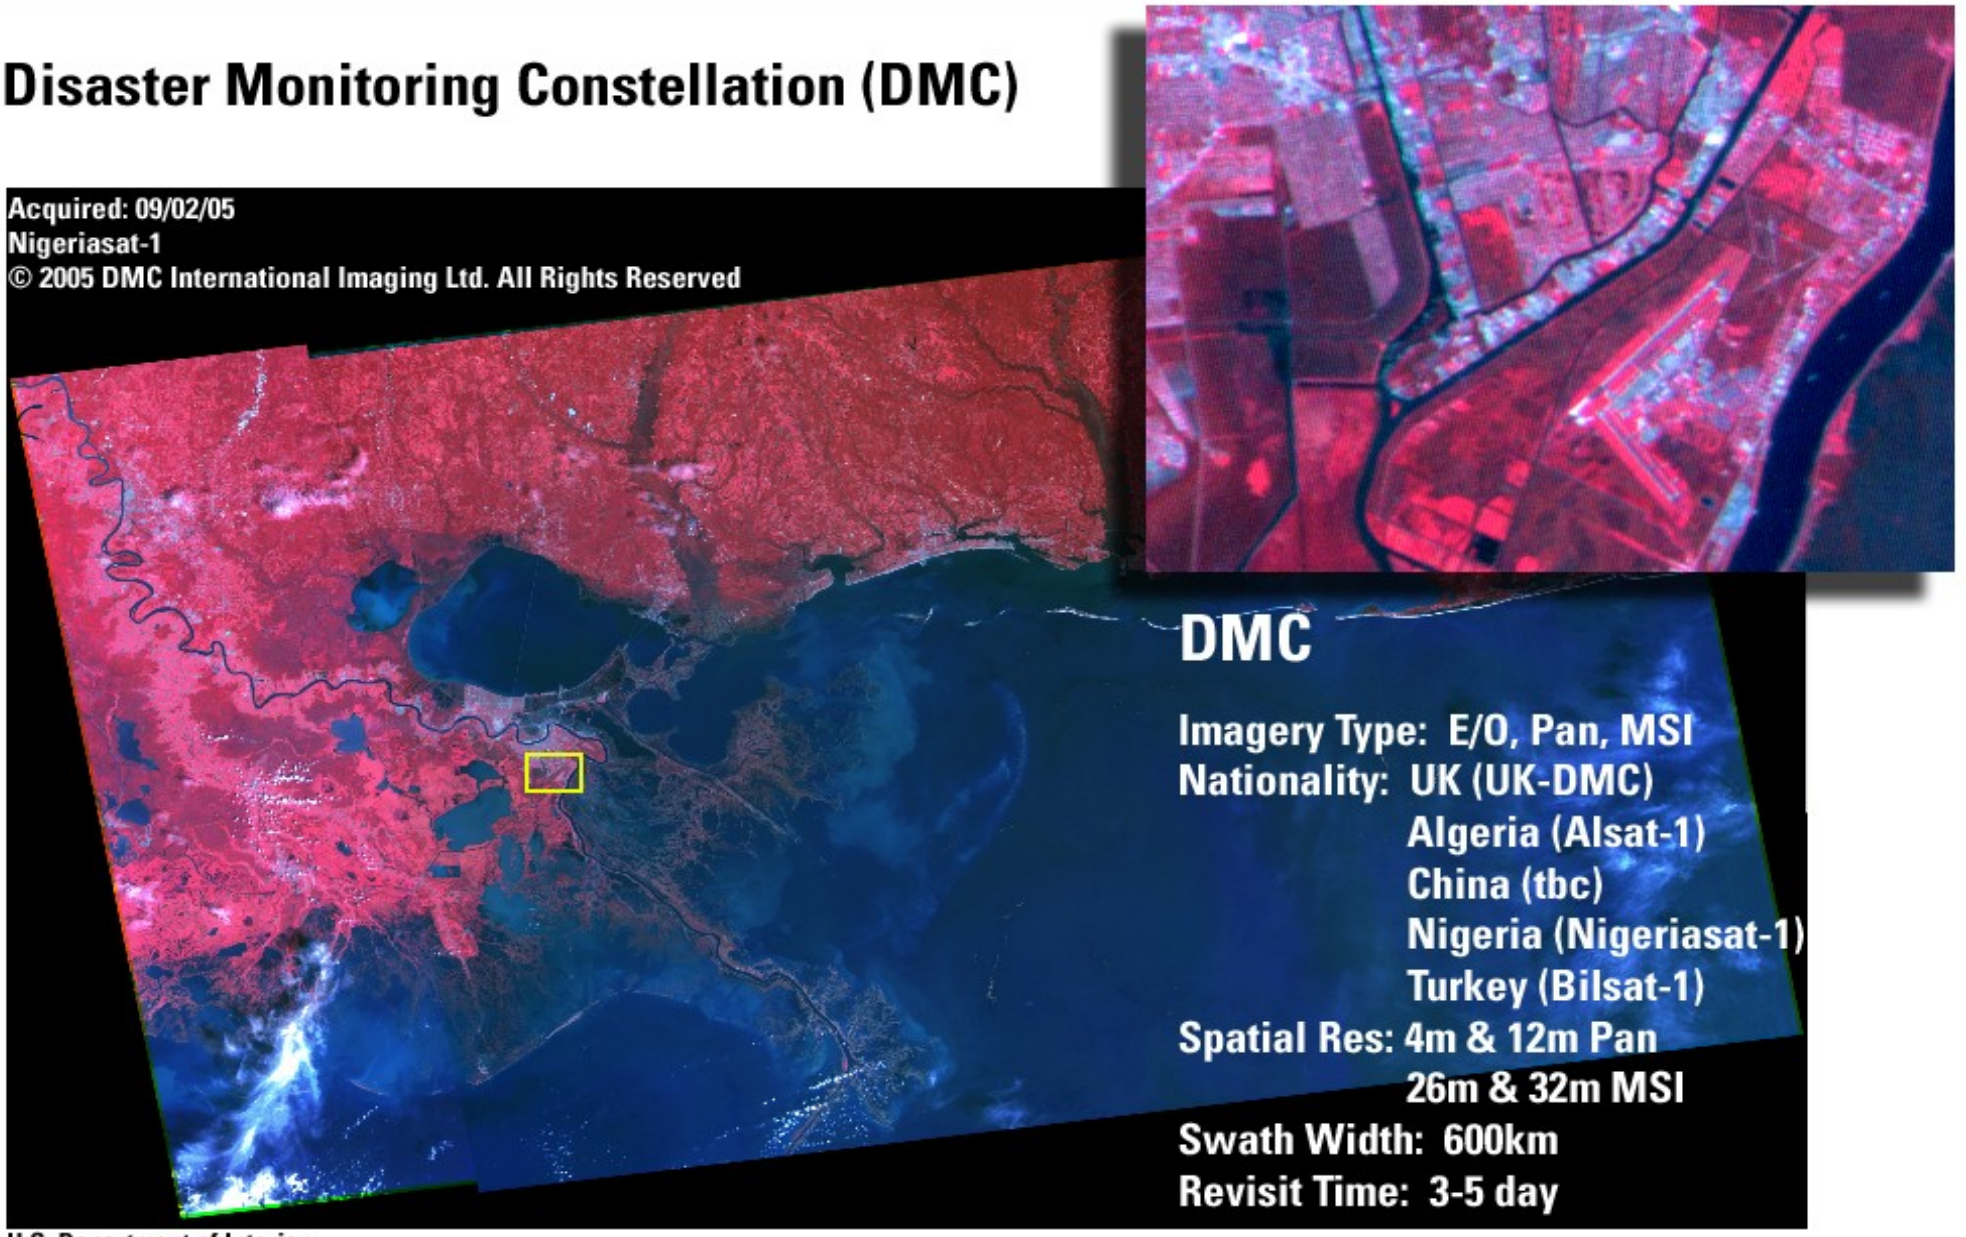

U.S. Department of Interior
U.S. Geological Survey

http://gisdata.usgs.net/hazards/katrina/satelliteinfo/DMC.jpg

Environment

- Asymmetric Links
  - 40 Mbps downlinks
  - 9.6 kbps uplinks
- Point-to-Point Links
  - link is the e2e path
  - no (little) other traffic -- explicitly scheduled
- “Low-end” CPUs
  - 200 Mhz PowerPC
- Potential for corruption-based losses
- Short-Lived Links
  - 8-12 minutes of connectivity
  - Image files up to several GB
    - need to fill link and do so quickly!

Assumptions

- The unshared point-to-point link could lead to bogus assumptions about acceptable sending behavior
  - Saratoga congestion control options are flexible:
    - may be externally configured with a rate
    - may use built-in acknowledgements and timestamps with some other scheme
  - Implementations do/will fully conform to TSVWG's “UDP Guidelines”

Basic Operation (“Put”)

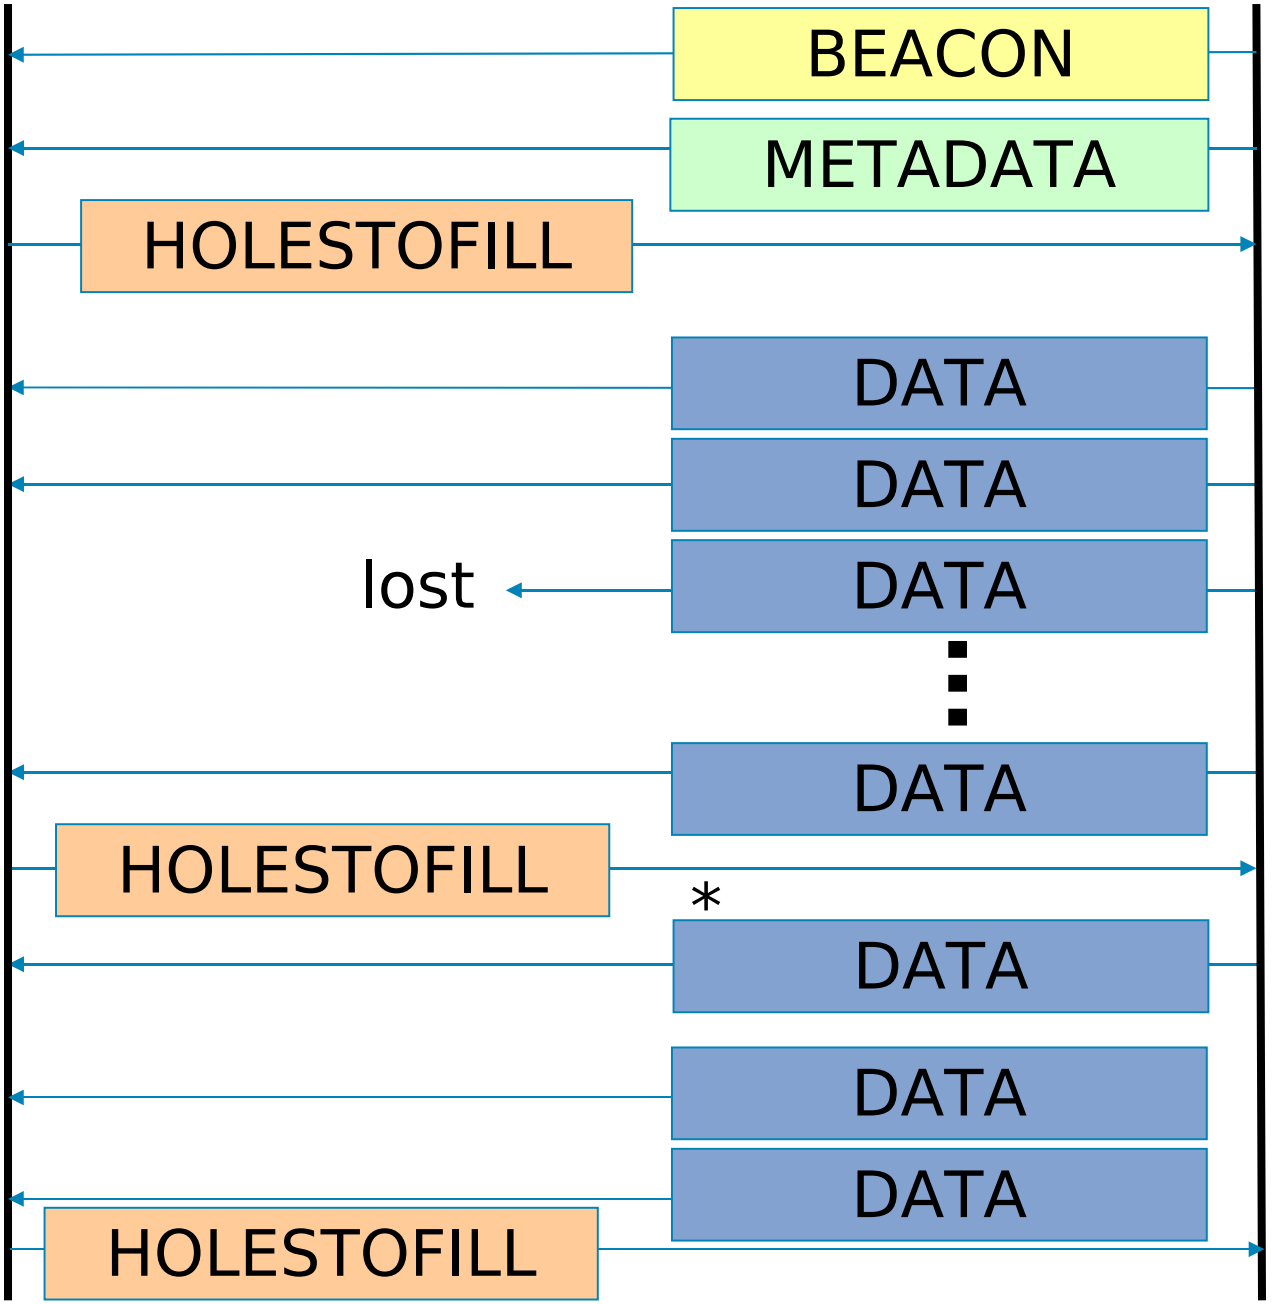

Describes object

Lost data creates hole in copy at receiver.

Ack requested and sent describing hole to be filled.

Data Packet Format

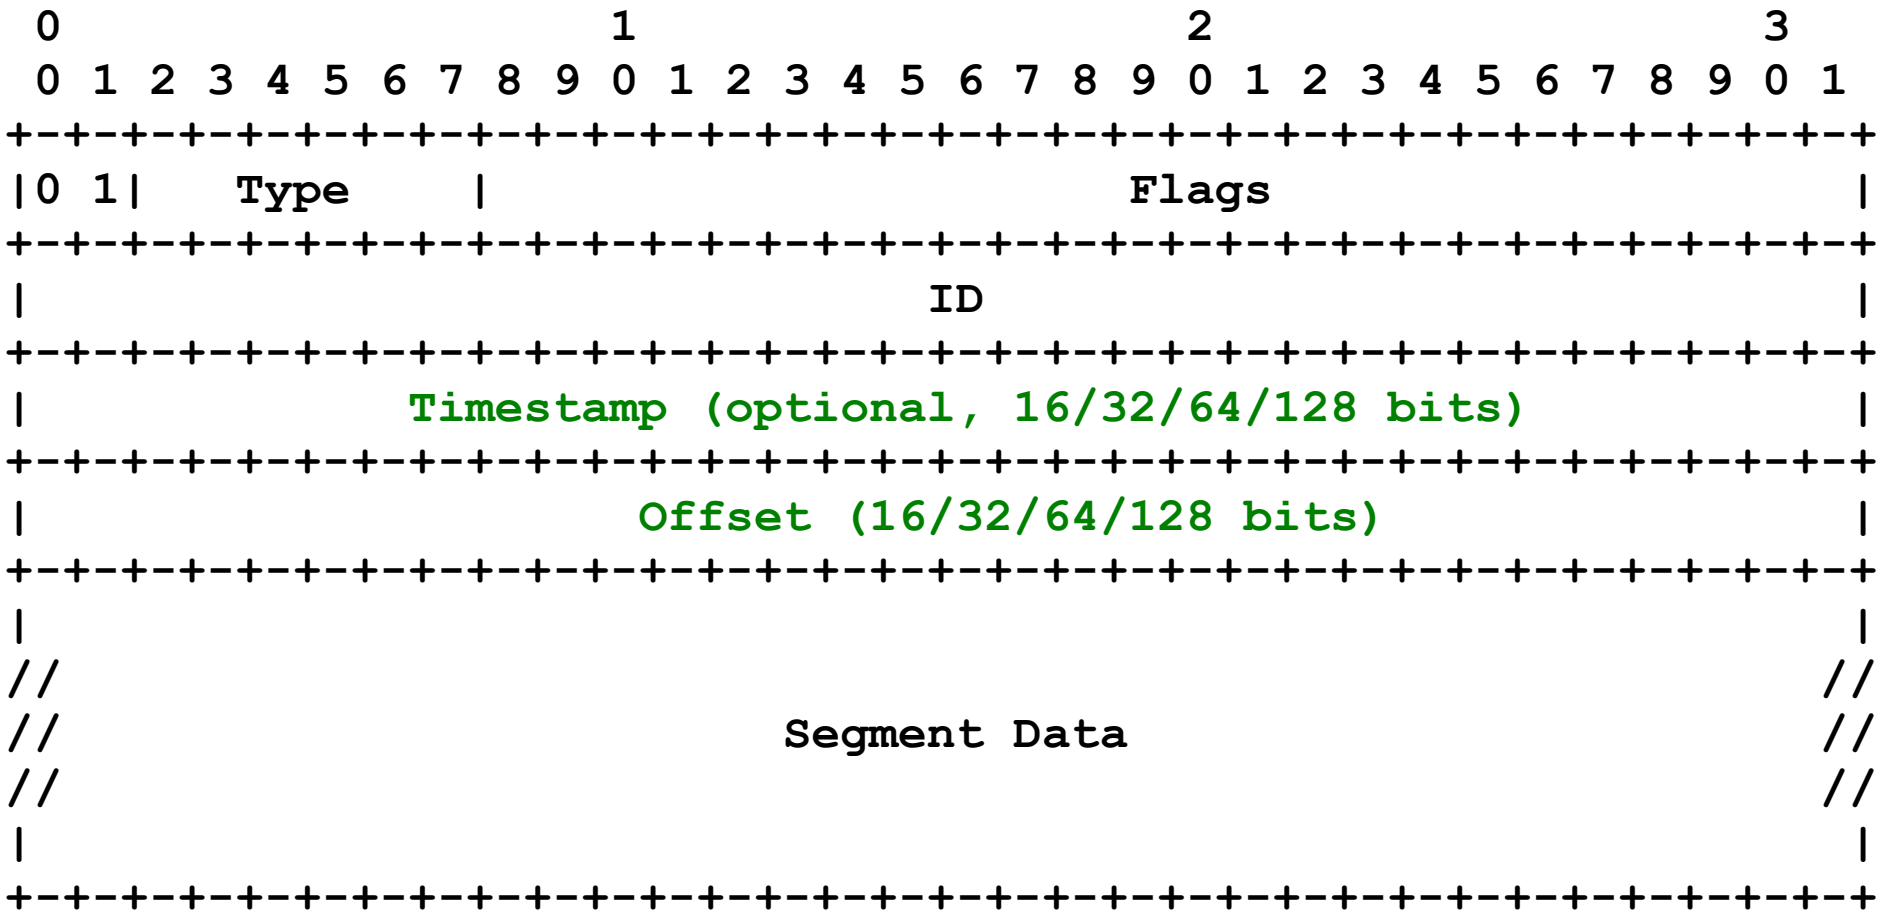

Feedback Packet Format (HOLESTOFILL)

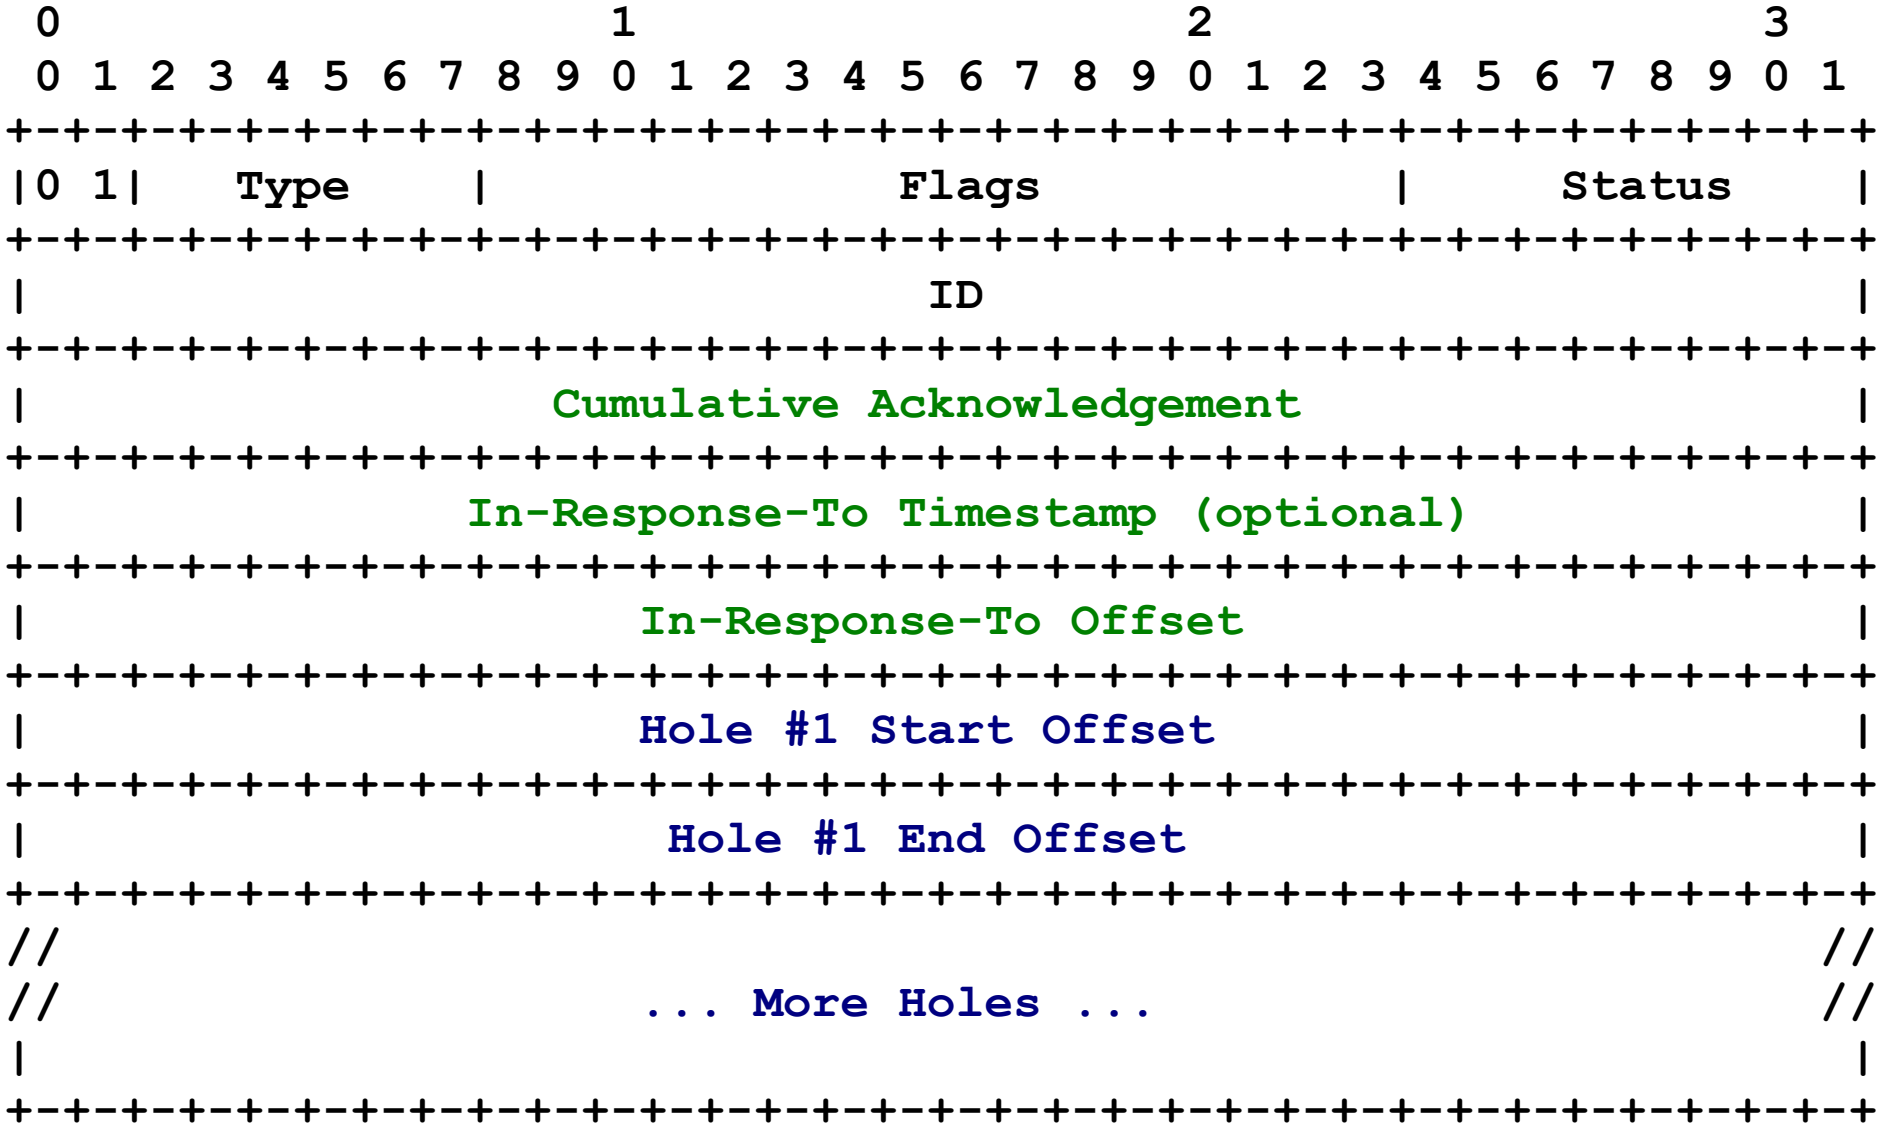

Unique Properties

- BEACON packets can (optionally) be sent even when no transfers are in-progress
  - advertise presence / capabilities
- Sends named content, not just bit blocks/streams
  - allows “directory” requests and “resuming”
- Supports file transfer, streaming, and DTNRG Bundle exchange
  - uses either UDP or UDP-Lite
- Allows both “push” and “pull” for content dissemination
- Field sizes can vary between 16/32/64/128 bits depending on content needs and platform constraints

Reliability

- (1) Protocol header sanity; (2) retransmission of segments; (3) delivery of errored content
- Per-segment
  - Relies on UDP checksum for error detection in both headers and payload
    - Can use UDP-Lite for header-only coverage if an application desires delivery of errored-content
  - SNACK and/or timer-based retransmission is completely optional for data
- Per-object
  - ACK-based retransmission for metadata
  - MD5 checksum over complete object (optionally computed / verified)
    - Stronger than UDP checksum, and guards against errors in reassembly or processing

Demonstrated Utility

- Measurements show Solid-State Data Recorder with 200 Mhz CPU fills entire 8 Mbps downlink with 9.6 kbps uplink and startup time of 1 RTT
  - also holds for 40 Mbps downlink
- TCP's main problems are the stream-based ordering and constrained ACK path
  - Accepted bound is roughly 50:1 ACK ratio

Downlink Rate	Uplink Rate	Raw Asymmetry
8 Mbps	9.6 kbps	833
40 Mbps	9.6 kbps	4167
210 Mbps	38.4 kbps	5469

Field Processing Efficiency

Saratoga fields are faster to encode and decode than the SDNVs used in LTP

Fixed length packet formats also allow faster parsing of the overall packet structure

C Encoding / Decoding Times

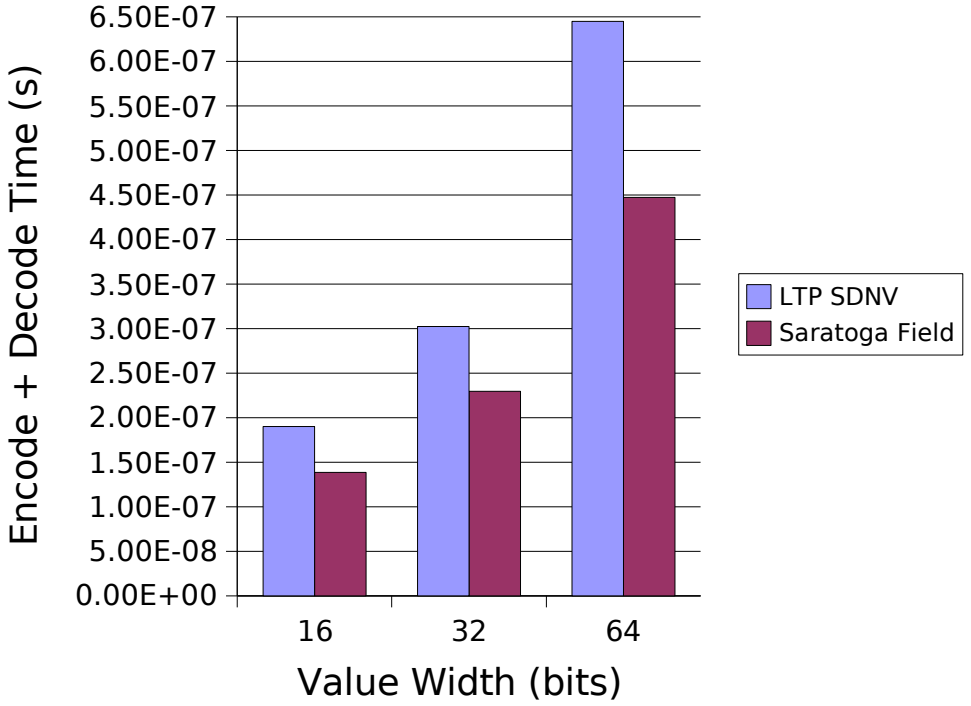

Python Encoding / Decoding Times

Goal: Experimental RFC

- SSTL & DMCII have depended on the “running code” for several years in day-to-day business operations.
- Likely applicable to similar scenarios for content distribution in PANs, free-space optical, proximity networking, Square Kilometer Array, grid computing
- If other uses are found outside space / satellite community, it can then move to Proposed Standard

Other Uses

- Any situation where:
  - HTTP or FTP are unusable because of asymmetry
  - connectivity may be frequently disrupted, or even unidirectional
  - large amount / rate of data and need for full link utilization is infeasible for traditional protocols
- Please talk to us if this could help you!

http://www.ietf.org/internet-drafts/draft-wood-dtnrg-saratoga-01.txt