DS-TP: Deep-Space Transport Protocol

Ioannis Psaras[‡], Giorgos Papastergiou[‡], Vassilis Tsaoussidis[‡] and Nestor Peccia[⋆]

[‡](ipsaras, gpapaste, vtsaousi)@ee.duth.gr

[⋆]nestor.peccia@esa.int

[‡]COMputer NETworks (COMNET) Group

Democritus University of Thrace, Xanthi, Greece

http://comnet.ee.duth.gr/

[⋆]Mission Operations Center

European Space Agency (ESA)

Darmstadt, Germany

Introductory Notes (1/4)

- The communication infrastructure is (perhaps) the most important part of a space mission.
- Scientific data needs to travel through deep-space links to reach institutes and research labs on Earth.
- Therefore, communication or connectivity time becomes a crucial point for the success of a space mission.

Introductory Notes (2/4)

- Reduced communication/connectivity time leads to extraordinary buffering requirements.
- There exist two optimization approaches:
  1. Extend infrastructure (not considered here).
  2. Increase the amount of data transferred within the given timeframes (our focus).
- Therefore, our focus is to optimize transport protocols.

Introductory Notes (3/4)

Deep-space links are characterized by:

- Extremely High Propagation Delays (e.g., 8 to 40 minutes from Earth to Mars)
- High Link Bit Error rates (e.g., up to 10^{-3})
- Intermittent Connectivity
- Bandwidth Asymmetry (e.g., up to 1000:1)

1. Clearly, TCP cannot operate under such conditions.
2. Dynamic, IP or IP-like routing will not work on the deep-space.

Introductory Notes (4/4)

The *Deep-Space Transport Protocol* introduced here,

- is an open-loop, rate-based protocol,
- transmits data on a fixed, pre-determined rate, according to the line rate,
- uses both positive and negative ACKs as the feedback scheme of choice,
- incorporates a novel, efficient retransmission mechanism, called *Double Automatic Retransmission* (DAR).

Related Studies

- TP-Planet: over-qualified (includes congestion control).
- RCP-Planet: un-reliable.
- SCPS-TP: operates mainly according to TCP.
- Saratoga: pretty simple and efficient but slow compared to DS-TP.
- CFDP: application layer protocol, with transport layer functionalities. Similar to Saratoga, but less efficient.
- LTP: includes a unique mechanism to differentiate between blocks of data that need 100% reliability and blocks of data that do not.

Deep-Space Transport Protocol: Basic Components

1. **Rate-based Transmission**: Transmit data on the line rate (i.e., data transfers are pre-scheduled and therefore congestion is not an issue).
2. **Mixed ACK-SNACK Strategy**: ACKs are used for buffer space release at the sender side. SNACKs are used for retransmission signalling.
3. **Double Automatic Retransmission**: DAR transmits, intentionally, each packet twice. The "gap" between the two transmissions is determined by the channel error rate. For example, if PER=33% the transmission sequence is: 1-2-**1**-3-4-**2**-5-6-**3**

Deep-Space Transport Protocol: DAR (1/5)

- *Actual Rate* and `c_seqno`
- *Retransmission Rate* and `r_seqno`
- *Line Rate = Actual Rate + Retransmission Rate*
- *Retransmission Rate = Packet Error Rate*
- *Actual Rate = (1 - PER%)·Link Rate*

Deep-Space Transport Protocol: DAR (2/5)

One redundant packet is sent every $\frac{1}{PER} - 1$ original packets.

For example, if $PER = 20\%$, the transmission sequence is:

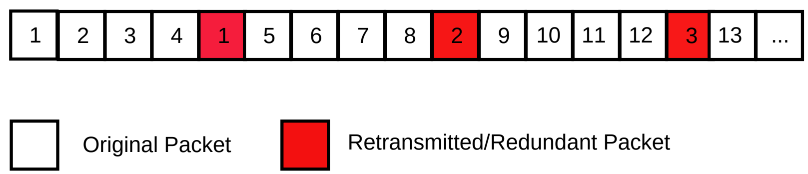

Figure 1: Transmission Sequence when $PER = 20\%$

Therefore:

(1) $$r\_seqno = \frac{c\_seqno - 1}{\frac{1}{PER} - 1}$$

Deep-Space Transport Protocol: DAR (3/5)

According to Equation 1, the packet with sequence number `c_seqno` will be retransmitted after `diff_pkts`:

$$diff\_pkts = [(\frac{1}{error\_rate} - 1) \cdot c\_seqno] - r\_seqno. \quad (2)$$

Deep-Space Transport Protocol: DAR (4/5)

ACK/SNACK types and their Functionality

- DS-TP uses positive ACKs for retransmission buffer space release and SNACKs to trigger retransmissions.
- DAR uses two types of SNACKs (i.e., SNACK$_1$ and SNACK$_2$.
- SNACK$_1$s are used for link error rate measurement at the sender.
- SNACK$_2$s are used to trigger retransmissions.

Deep-Space Transport Protocol: DAR (5/5)

- SNACK_1 is sent when there is a hole at the receiver's buffer space
- The sender uses SNACK_1s to measure the link error rate. SNACK_1 does not trigger retransmission.
- According to Equations 1 and 2 the receiver knows the arrival time of the retransmitted packet.
- In case the packet does not arrive, the receiver sends SNACK_2.
- SNACK_2 triggers packet retransmission.

DS-TP Theoretical Evaluation (1/5)

We grapgh $r\_seqno = \frac{c\_seqno - 1}{\frac{1}{PER} - 1}$ and $diff\_pkts = [(\frac{1}{error\_rate} - 1) \cdot c\_seqno] - r\_seqno$ in Fig. 13.

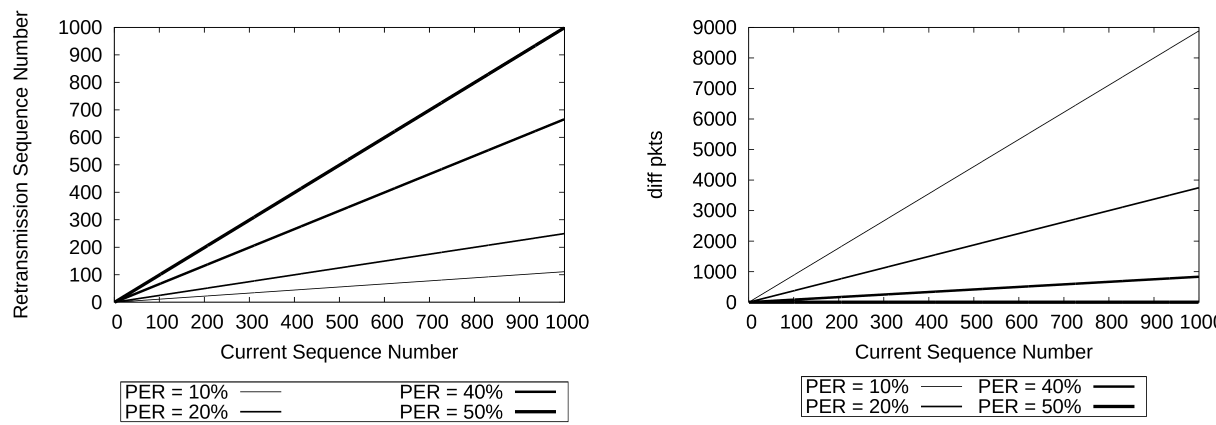

(a) Current vs. Retransmission Sequence Number

(b) Number of Packets between Original Transmission and Retransmission

DS-TP Theoretical Evaluation (2/5)

- A x Mbps link can transfer $\frac{x}{8}$MBps or $\frac{1024 \cdot x}{8}$KB/s.
- Therefore `diff_pkts` require: $diff\_time = \frac{8 \cdot diff\_pkts}{1024 \cdot x}$ seconds.

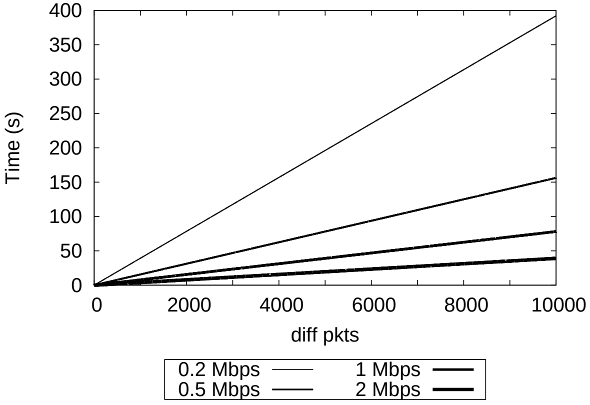

Figure 2: Time Interval between Original Transmission and Retransmission

DS-TP Theoretical Evaluation (3/5)

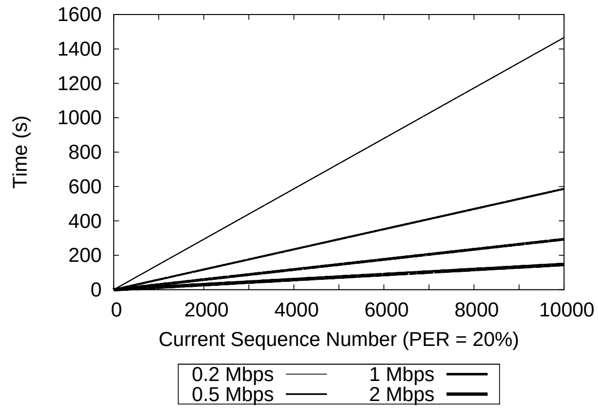

(a) PER=20%

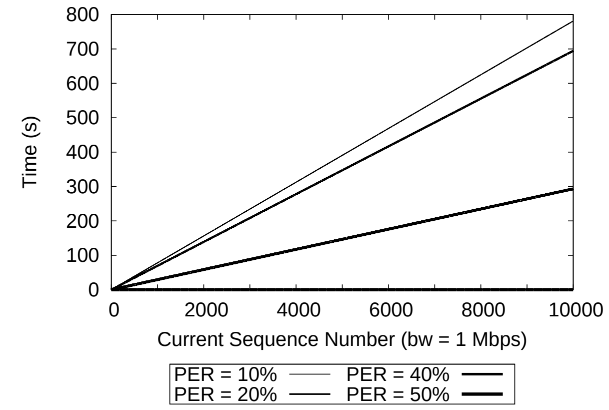

(b) bw=1Mbps

- We see that the DAR retransmission interval may be up to 25mins for the 10,000th packet.
- If the reverse link propagation delay is smaller than the retransmission interval, then the DAR's functionality is cancelled.

DS-TP Theoretical Evaluation (4/5)

- In particular, if $diff\_time > Pr.\ D.\ Rev.$ or $diff\_time > \frac{RTT}{2}$, then retransmission should be triggered by SNACK$_1$.
- The sender calculates the sequence number boundary that cancels DAR.
- When the above sequence number becomes the *current* sequence number (i.e., c_seqno), then the sender sets $r\_seqno \leftarrow c\_seqno$.

DS-TP Theoretical Evaluation (5/5)

Overhead

The overhead introduced by DS-TP is mainly determined by the (intentional) DAR retransmission mechanism. In order to calculate DS-TP's retransmission overhead, we modify Equation 1 (i.e., $r\_seqno = \frac{c\_seqno-1}{\frac{1}{PER}-1}$), as follows:

(3)
$$retr = \frac{fs-1}{\frac{1}{y}-1} \simeq \frac{fs}{\frac{1}{y}-1} = \frac{fs \cdot y}{1-y},$$

where fs is the size of the file to be transmitted and y the PER.

Protocol Evaluation Framework (1/2)

- We compare (theoreticaly) the performance of DS-TP with the *Fixed-Rate Transport Protocol* (FR-TP).
- FR-TP is similar to Saratoga and CFDP.
- FR-TP transmits data on a fixed, predetermined rate, equal to the line rate.
- SNACKs are sent to the sender only after the file transfer is complete (i.e., the sender has transmitted all data into the transmission link).

Protocol Evaluation Framework (2/2)

- DS-TP's functionality is modified accordingly, in order to compete fairly with FR-TP.
- According to this protocol setup, the file transfer is divided into rounds:
  1. During the 1^{st} round the sender transmits the file.
  2. During the 2^{nd} round the sender (re-)transmits the packets lost during the 1^{st} round.
  3. The file transfer is complete once all packets are successfully delivered to their destination.
- Therefore, we attempt to evaluate the performance of the protocols in terms of *rounds*.

DS-TP vs FR-TP (1/11)
FR-TP (1/2)

- During the 1^{st} round FR-TP transmits fs MBs; $fs \cdot y$ MBs are lost and need to be retransmitted.
- During the 2^{nd} round FR-TP transmits $fs \cdot y$ MBs; $(fs \cdot y) \cdot y$ MBs are lost and need to be retransmitted.
- During the n^{th} round FR-TP transmits $fs \cdot y^{n-1}$; the sender will need to retransmit $fs \cdot y^n$ MBs.
- We assume that once $fs \cdot y^n < 1$ (packet) the file transfer is complete.

DS-TP vs FR-TP (2/11)

FR-TP (2/2)

Therefore, FR-TP needs n_{frtp} rounds in order to complete the file transfer:

(4) $$n_{frtp} = \log_y(y^n) = \log_y(\frac{1}{fs}) = \frac{\log \frac{1}{fs}}{\log y}$$

DS-TP vs FR-TP (3/11)
DS-TP (1/9)

- During the 1^{st} round, DS-TP transmits $fs + r_1$ MBs, in total, where r_1 are the DAR retransmissions.
- During the 1^{st} round, $fs - r_1$ KBs are sent once and r_1 KBs are sent twice.
- Provided that the channel PER applies uniformly for the total number of packets:
  - $fs - r_1$ are lost with probability y, and
  - r_1 are lost with probability y^2, where $r_1 = fs \cdot \frac{y}{1-y}$.

DS-TP vs FR-TP (4/11)

DS-TP (2/9)

- Therefore, during the 1^{st} round, a_1 packets are lost:

(5) $$a_1 = (fs - r_1) \cdot y + r_1 \cdot y^2.$$

- Substituting r_1 into Equation 5, we have:

(6) $$a_1 = fs \cdot y \cdot (1 - y).$$

DS-TP (3/9)

Similarly,

- During the 2^{nd} round, DS-TP transmits $a_1 + r_2$ MBs, in total, where r_2 are the DAR retransmissions.
- During the 2^{nd} round, $a_1 - r_2$ KBs are sent once and r_2 KBs are sent twice, where
  - $a_1 - r_2$ are lost with probability y, and
  - r_2 are lost with probability y^2, where $r_1 = fs \cdot \frac{y}{1-y}$.

DS-TP vs FR-TP (6/11)
DS-TP (4/9)

- Therefore, during the 2nd round, a_2 packets are lost:

(7) $$a_2 = (a_1 - r_2) \cdot y + r_2 \cdot y^2.$$

- Substituting r_2 into Equation 5, we have:

(8) $$a_2 = fs \cdot y^2 \cdot (1 - y)^2$$

DS-TP vs FR-TP (7/11)
DS-TP (5/9)

- DS-TP will complete the file transfer, when:

(9) $$fs \cdot y^n \cdot (1-y)^n < 1\,packet$$

- Hence, DS-TP needs n_{dstp} rounds to transfer a fs MBs file: $n_{dstp} = \log_{[y\cdot(1-y)]}[y\cdot(1-y)]^n = \log_{[y\cdot(1-y)]}(\frac{1}{fs}) \Rightarrow$

(10) $$n_{dstp} = \frac{\log \frac{1}{fs}}{\log(y\cdot(1-y))}$$

DS-TP vs FR-TP (8/11)
DS-TP (6/9)

- We divide Equations 4 and 10 by parts, in order to obtain DS-TP's gain against FR-TP, due to DAR:

(11) $$n_{ratio} = \frac{n_{frtp}}{n_{dstp}} = \frac{\frac{\log \frac{1}{fs}}{\log y}}{\frac{\log \frac{1}{fs}}{\log(y \cdot (1-y))}} = 1 + \frac{\log(1-y)}{\log y}.$$

- We see that the performance difference depends on the PER, only.
- The performance difference increases with the error rate.

DS-TP (7/9)

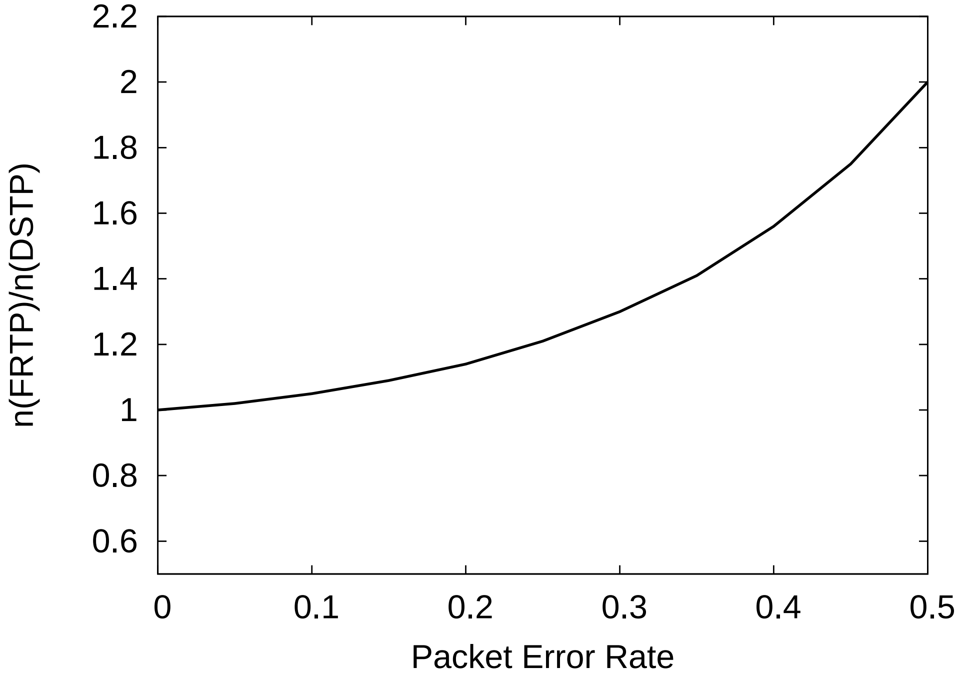

Figure 3: Performance Increase due to DAR in terms of Rounds

DS-TP vs FR-TP (10/11)

DS-TP (8/9)

- In absolute numbers, the difference in rounds between DS-TP and FR-TP is given by:

$$n_{diff} = n_{frtp} - n_{dstp} = \frac{\log(1-y) \cdot \log \frac{1}{fs}}{\log y \cdot \log(y \cdot (1-y))} \tag{12}$$

- In contrast to n_{ratio}, we see that n_{diff} depends, apart from the link error rate, on the file size as well.

DS-TP (9/9)

- DS-TP is up to 8 rounds faster than FR-TP.
- The difference increases even more for larger file sizes.
- Note that the performance difference *does not* depend neither on the link speed, nor on the RTT.

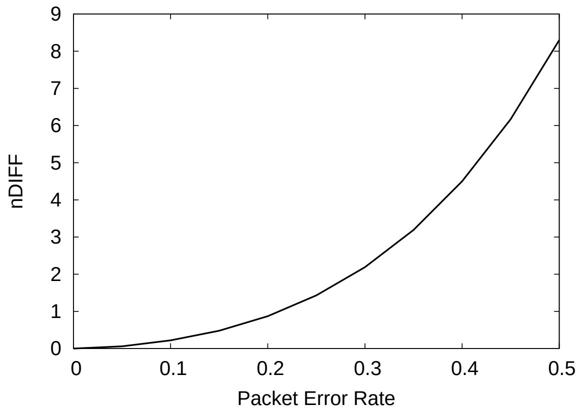

(a) File Size = 100MB

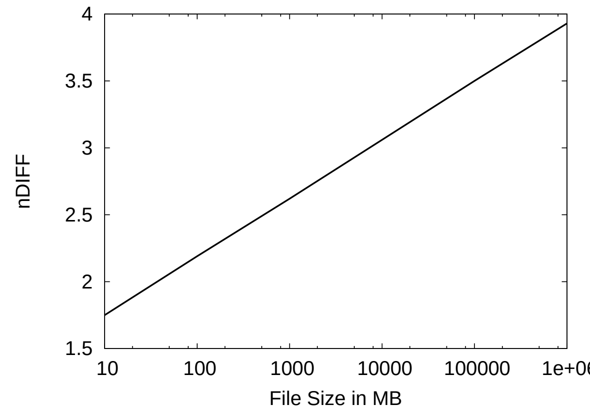

(b) PER = 0.3%

Conclusions

- DS-TP is a novel approach to deep-space communications.
- DS-TP’s radical design guidelines reveal that decoupling of conventional ideas and protocols can improve the performance of deep-space communications.
- DS-TP outperforms conventional transport protocols by a factor of 2, under specific network conditions.
- Further investigation is needed in order to determine implementation details, which where left open in the current study.