

Progressive DTN adoption for CubeSat operators

Sarah Theouille

Olivier De Jonckère

LIRMM · CSUM






Montpellier, France

Montpellier University Space Centre: designs and operates its own CubeSats
— a 12U mission in preparation.


Delay/Disruption-Tolerant Networking (DTN)



Store, carry, forward: a node keeps the *bundle* until the next contact.

DTN architecture: RFC 4838 · Bundle Protocol v7: RFC 9171


Delay/Disruption-Tolerant Networking (DTN)




Store, carry, forward: a node keeps the *bundle* until the next contact.

DTN architecture: RFC 4838 · Bundle Protocol v7: RFC 9171

Delay/Disruption-Tolerant Networking (DTN)



CLA Convergence Layer Adapter BP Bundle Protocol
- - High delay transmission - - Message retention




Ground station Orbiter Orbiter trajectory
Stored bundle Bundle transmission

Long delays, long disruptions: no need for a permanent end-to-end path.

BP is an **overlay protocol**: it runs over any transport, through a CLA.

Delay/Disruption-Tolerant Networking (DTN)



Long delays, long disruptions: no need for a permanent end-to-end path.

BP is an **overlay protocol**: it runs over any transport, through a CLA.

Why BP?

A protocol for the future.

IETF

standardisation: RFC
4838 (DTN), RFC
9171 (BPv7)

CCSDS

space standards,
including
*Schedule-Aware
Bundle Routing*

IOAG

inter-agency
recommendations,
lunar and solar
networks

IPNSIG

carries the vision of
the *Solar System
Internet*

Industrial and academic interest

NASA LCRNS

ESA Moonlight

NASA-ESA LunaNet

D3TN

Aalyria

Spatiam

Tolero

Areas of interest

TOOLING

- Protocol — BP + CSPCL
- Routing — A-SABR
- Interoperability — Charon

MODELS

- Routing optimisation
- Internetworking

PARTNERS AND CONTACTS

- CSUM
- Students from Montpellier and elsewhere
- INRIA, D3TN, University of Bologna ...

AMBITIONS

...

The challenge

How do we introduce DTN without being intrusive?

CSP, tested and flight-proven

the satellite's internal network,
widespread among CubeSats
— but not DTN

BP is appealing...


... but changing the stack
is a risk, and not every
application speaks BP

Small players

are not ready to make
the leap before
assessing the gains

- **Our lever:** BP is an overlay — it can run *on top of* CSP, and delay tolerance can be tested **without changing the stack.**


CSPCL



A small adapter between BP and CSP.

CSPCL is the missing CLA — and CSP remains the internal network, on the CAN bus.

Level 1: talking to CSP




- ▶ **RDP**: acknowledged CSP connections.
- ▶ **SFP**: fragmentation — the maximum bundle size is no longer limited by the size of a CSP packet.
- ▶ **SocketCAN**: access to the CAN bus through Berkeley sockets.


What TCPCL does for TCP, CSPCL does for CSP.

Since BP is an overlay protocol, we might as well use Internet technologies (TCPCL) for the part of the DTN network that overlaps the Internet.

Level 2: bundle reliability over CSP




An end-to-end ACK, specific to CSPCL.



3 different Bundle Protocol
implementations, one shared core.


Challenge 2



How can existing applications become delay-tolerant?

Goal: **service continuity**. They already fly, are qualified, and speak CSP: rewriting them is not an option.

Charon



A proxy: CSP packets, wrapped in bundles.

On arrival, Charon unwraps them back into plain CSP.

Neither Alice nor Bob changed a single line of code.

We have all the bricks. How do we validate them?

A full chain

heterogeneous, end to end


Realistic routing

with the A-SABR library

Real hardware

and induced disruptions

What the full chain looks like



4 nodes, 3 BP implementations, 3 transports: validated by an interoperability test.

Adaptive library for Schedule-Aware Bundle Routing (A-SABR)

A framework, developed in Montpellier, to instantiate different routing strategies.

New concepts, proposed in response to the difficulties encountered with CGR (*Contact Graph Routing*).

- ▶ Volume management
- ▶ Node resources
- ▶ Internetworking support
- ▶ Microcontroller support
- ▶ Multiple backends, for many different algorithms

No reactive quality of service in DTN: resource management goes through routing optimisation.

The testbed



- **Disruption test:** a node killed mid-transfer — bundles are kept, then delivered on restart.

Why PolarFire SoC: RISC-V running Linux, representative of the avionics of CSUM's 12U mission.

Future work

Towards microcontrollers

The target CubeSat is a fairly large one, running Linux. One goal is to extend the approach to smaller ones: A-SABR is already compatible, CSPCL and Charon still need to be ported, and a microcontroller BPA is in sight with our German partners (picoD3TN).

Optimisation model

to derive the theoretical limits of routing, and serve as a reference for benchmarks

Cloud-native D3TN

in space: a dedicated PCB to run containerised delay-tolerant experiments

A flying laboratory

for researchers' experiments, DTN or not

Acknowledgements



Thanks to everyone who contributed to the project — including many students.

Open source, contributions welcome: github.com/DTN-MTP

THANK YOU

Questions?

Progressive DTN adoption
for CubeSat operators

Sarah Theouille

Olivier De Jonckère

LIRMM · CSUM

github.com/DTN-MTP