Summary Report on the IPNSIG Architecture Working Group Workshop on DTN Routing

2 September 2022

Objectives

- Perform a “Consumer Reports”-like apples-to-apples review of all well-defined technologies that have been proposed for managing the forwarding of bundles through the Solar System Internet.
- Rule out those technologies that have fundamental problems.
- Examine in more detail those technologies that look practical.
- From this analysis, identify the relevant identifiers that need to be managed in IANA registries (or equivalent).

Considerations (1 of 2)

- By what delay-tolerant mechanism(s) does this technology obtain the information on which it decides which next-hop node to forward a given bundle to? (E.g., route computation, scoring/ranking of neighboring nodes, other?)
- If the technology entails computation of end-to-end routes through the network, by what delay-tolerant mechanism(s) does it obtain the time-varying topological information on which routes are computed? (E.g., does it expect this information to be managed? Does it expect to discover this information in some other way?)

Considerations (2 of 2)

- How does this technology:
  - Cause high-value (nominally, high-priority) data to be delivered before lower-value data?
  - Maximize the utilization of transmission opportunities?
  - Maximize throughput in the network?
  - Scale up to a network of 100,000 nodes?

ESA BP routing

- All possible destinations are listed in Next-Hop Table.
- Each destination in NHT is mapped to one or more Convergence-Layer Adapters, which are individually activated by external signals.
- Each CLA is a **stack** of Convergence-Layer Elements, each of which defines a connection using a specified protocol.
- Design enables bundle to be forwarded either to a specific node or to whichever node receives it.
- All information populated by management.
- Not intended for use in spacecraft or for large networks.

SABR /CGR

- Definitions of forwarding, route, and routing
- Analogy between routing and travel planning: roadmap (IP) vs flight schedules (BP)
- Contact plan is like a flight schedule, a time-varying topology.
- For each destination, construct a contact graph from the contact plan and use Dijkstra search to find best path through the graph.
- For each bundle, use the best path for the bundle's destination.
  - If best path doesn't work for a given bundle, find more (per Yen's algorithm).
- Must deal with contact failures, overbooking, other exceptions.

REDMARS

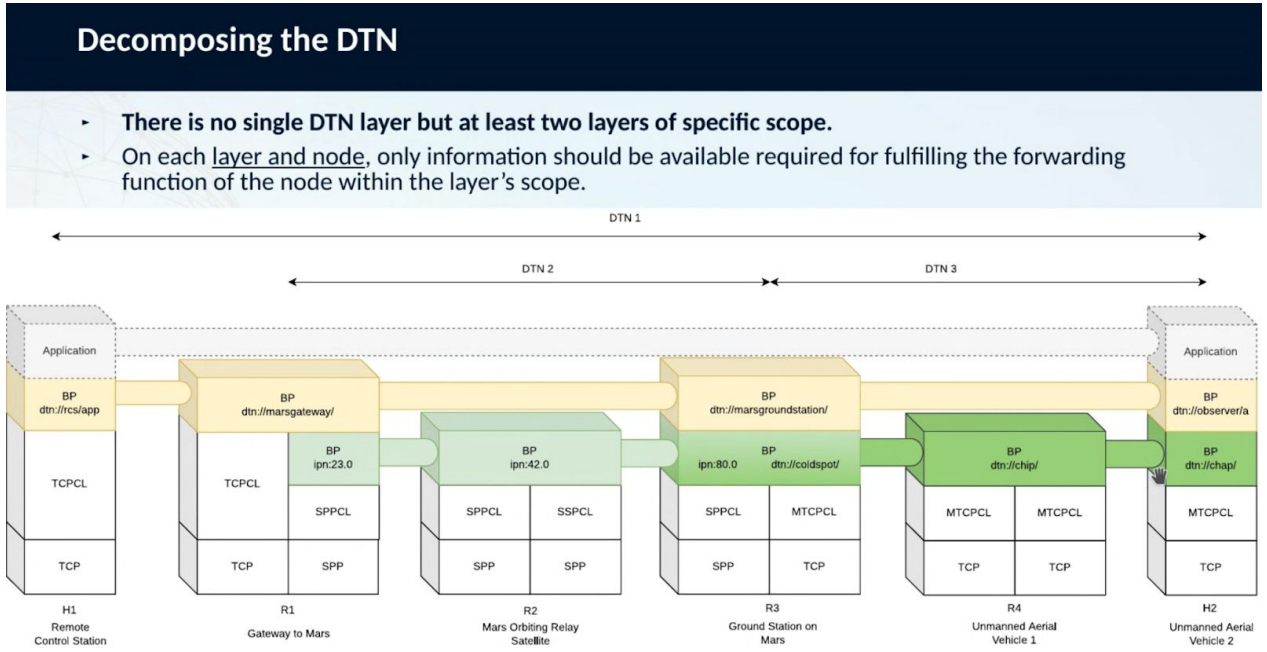

Leverages bundle-in-bundle encapsulation (BIBE).

- Generic bundle handling interface: Bundle Dispatcher Module receives bundle parameters, returns forwarding instructions.
- For Ring Road, a list of neighboring nodes identifying endpoints that are reachable via each neighbor; bundles are queued for specific contacts, may be reordered per priority.

Spacetime

- Large compute resource:
  - Models the network
  - Anticipates state changes
  - Automatically re-tasks assets accordingly, pre-emptively
  - Re-tasks in real time when unanticipated state changes are detected
- No integration with Bundle Protocol at this time; all network activity is IP, including all routing and forwarding.

SPSN

- Replacement for SABR/CGR.
  - Based on a node multigraph rather than contact graphs.
    - Much smaller computation problem.
  - Computes path during per-bundle path selection, rather than in advance.
    - Enables bundle size to be considered during computation of best path.
  - When computing paths, computes paths for all destination nodes rather than just one.
  - Improved mechanism for managing allocation of contact capacity to bundles (volume management).

PRoPHET

- Based on exchanged history of encounters and transitivity.
- Table asserts delivery predictability for each node, aging over time.
- A copy of a given bundle is forwarded, according to a selected strategy, to each node for which the probability of delivery to the bundle’s destination exceeds the local node’s own probability of delivery to the destination.

Spray & Wait

- No knowledge of the network is assumed.
- For each bundle, the optimal number of copies that may be in transit in the network (N) is computed at the source node.
- Upon contact with a node:
  - Authorization to forward some number of copies of the bundle (Q, where Q < N) is conveyed to that node.
  - The forwarding node's authorization to forward additional copies of the bundle in the future is reduced by Q.
- Eventually no further copies may be forwarded by any node, except to the final destination. Now we wait for contact with the final destination node.

OCGR

- Extends CGR with contacts in which we have imperfect confidence.
- Relies on contact discovery as per Neighbor Discovery Internet Draft.
- History of discovered contacts is propagated via Saga protocol.
- Predicted contacts, in which we have less than 100% confidence, are computed from aggregate discovered contact history.
- Limited contact confidence results in limited route confidence; bundle delivery confidence gradually grows as copies of the bundle are forwarded via these routes, until a threshold is reached.
- NOTE: RUCoP (Routing under Uncertain Contact Plans) may be an alternative; need more information.

IRF (inter-regional forwarding)

- In-situ computation of routes through a contact plan citing a billion nodes is infeasible. Instead, divide the network into *regions*, each of which comprises all nodes cited in the contacts of a single contact plan of manageable size.
- To send a bundle to a node that is not in the source node's region, forward it through a sequence of passageway nodes that are members of topologically adjacent regions. Passageway connections coerce the region topology into a tree structure.
- Region topology is discovered by transmitting copies of a probe bundle to all locally reachable passageways, recursively, and noting ultimate delivery results returned from passageways at leaf regions.

Discussion – general model

- Similar to the D3TN "bundle dispatcher module" concept:

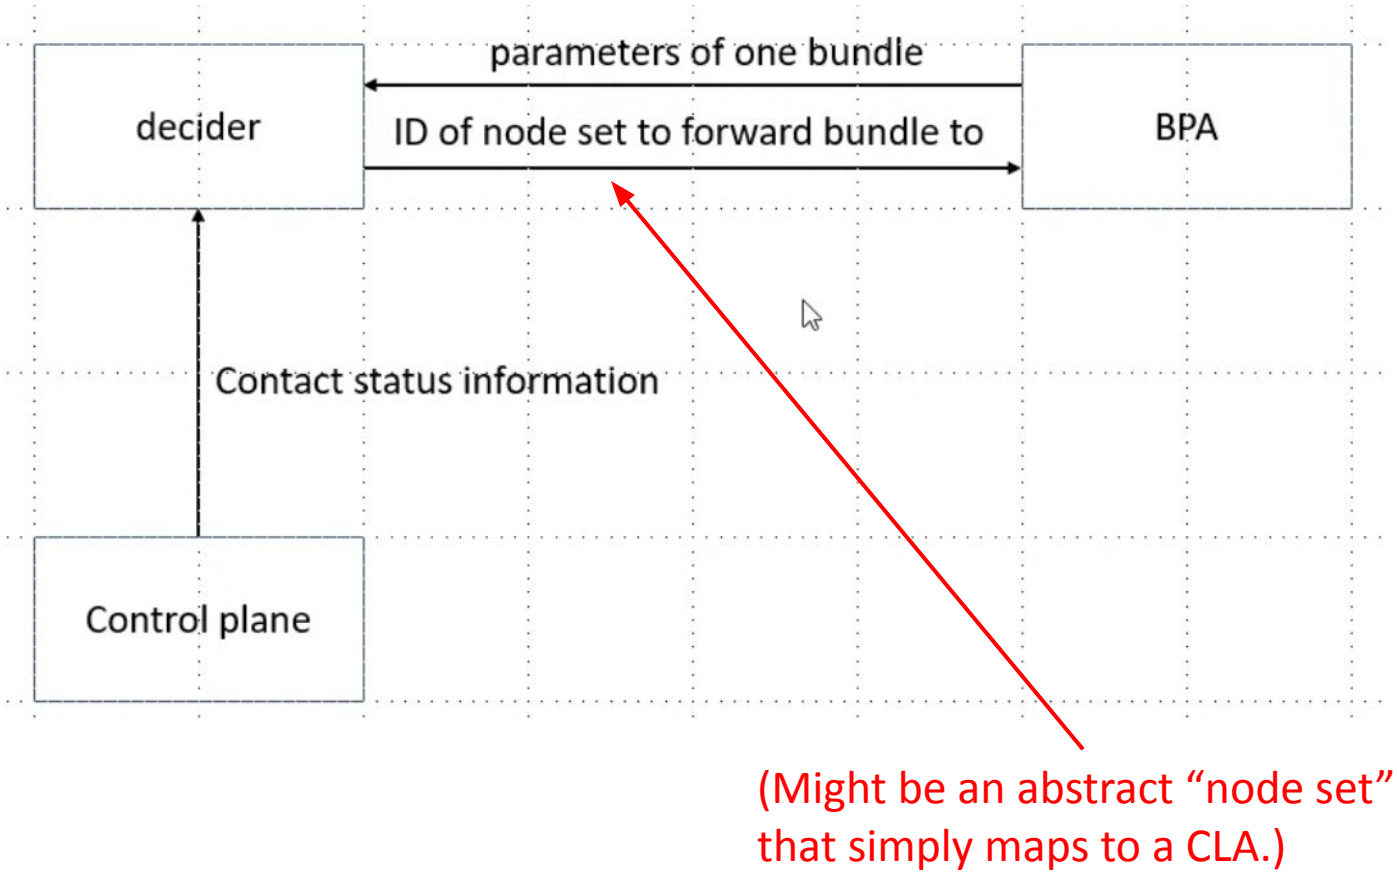

	ESA BP routing	SABR/CGR	REDMARS	Spacetime	SPSN	PRoPHET	Spray & Wait	OCGR	IRF
Source of forwarding information	Management	Routes	Any; for Ring Road, info is managed	Routes	Routes	Statistics	Random	Routes	Passageway discovery
Source of topology information	n/a	Contact plans	Any; for Ring Road, info is managed	Models	Contact plans	n/a	n/a	Predicted contacts, per SAGA protocol	n/a
Priority support	Not identified	Overbooking support, etc.	Only per intra-regional subnet; Ring Road – yes	Not identified (IP routing)	Not identified	Not identified	Not identified	Per CGR	Only per intra-regional subnet
Maximizes utilization	By management	Effective Volume Limit	Only per intra-regional subnet; Ring Road – mgt	Not identified (IP routing)	Contact Partitioning	No	No	Per CGR	Only per intra-regional subnet
Maximizes throughput	Single copy is forwarded end-to-end	Single copy is forwarded end-to-end	Single copy is forwarded end-to-end	Single copy is forwarded end-to-end	Single copy is forwarded end-to-end	Selective forwarding	Limited flooding	Selective forwarding	Only per intra-regional subnet
Scales to 100,000 nodes	No	Only as supported by IRF	Yes; stacked networks	Only as supported by BP (future)	Yes; rapid computation	Yes	Yes	Only as supported by IRF	Yes