Delay Tolerant Networking Research Group (DTNRG)

Kevin Fall & Stephen Farrell
Intel Labs / WHOI
and
Trinity College Dublin

kfall@intel.com
stephen.farrell@cs.tcd.ie

Challenged Networks

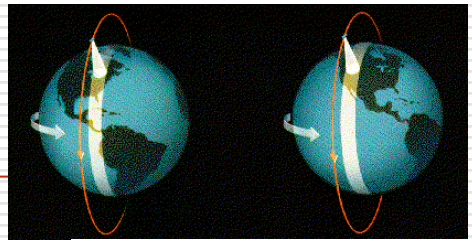

- Random/predictable connectivity
- Big delays, low bandwidth
  - satellites (GEO, LEO / polar)
  - exotic links
    - deep space comms
    - underwater acoustic comms
- Big delays, high bandwidth
  - Buses, mail trucks, patrol vehicles, zebras, etc.

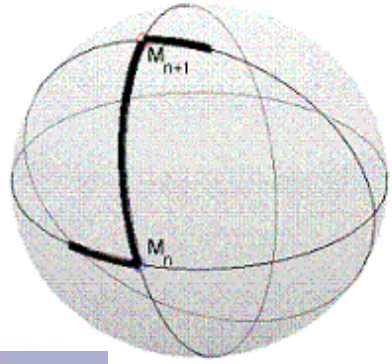

Internet Assumptions (in practice)

- Topology graph may change a bit, but remains connected [even in MANETs]
- Fixed-size node labels remain topologically-related
- E2E path has modest delay at most
  - Control loops on O(one RTT)
- E2E path doesn't have really big, small, or asymmetric bandwidth
- Not much re-ordering
- Routers are trusted to forward packets
- Paths not very lossy (< 1%)
- In-network storage is limited / short-term
- End stations are more reliable than routers

Internet for Challenged Networks?

- Is the topology/routing approach ok?
- Is the data plane model still good?
- What happens when one or more of the Internet assumptions don't hold (strongly)?
- Do:
  - Applications break or have intolerable performance?
  - Communications become impossible?
  - Elements of the system become less secure?

Performance Enhancing Proxies

- Perhaps the bad links can be 'patched up'
  - If so, then TCP/IP might run ok
  - Use a specialized middle-box (PEP)
- Types of PEPs [RFC3135]
  - Layers: mostly transport or application
  - Distribution
  - Symmetry
  - Transparency

TCP PEPs

- Modify the ACK stream
  - Smooth/pace ACKS -> avoids TCP bursts
  - Drop ACKs -> avoids congesting return channel
  - Local ACKs -> go faster, goodbye e2e reliability
  - Local retransmission (snoop)
  - Fabricate zero-window during short-term disruption
- Manipulate the data stream
  - Compression, tunneling, prioritization

Evolving Topology & Addressing

- IP uses fixed 32(128) bit addresses assigned based on topology location
  - couples location with identification
  - not inherently secured
  - aggregable ~ "scalable" [KK77]
- Name-based and flat routing
  - helps separate ID from topology
  - can be linked with application uses
  - non-aggregable (but maybe "scalable")
    - see results in DHT schemes + compact routing

Evolving the Data Plane

- Datagrams are a poor fit with
  - connection-oriented / cloud network
  - small frames (sensornets, atm)
  - poor links and unusable network storage
- Application Data Units (ADU's)
  - tailored to the application's desires
  - might be stored / retransmitted by network
  - convenient security unit

Evolving End-to-End

- Recall ‘fate sharing’ (Clark):
  - *it is acceptable to lose the state information associated with an entity if, at the same time, the entity itself is lost*
- But state (e.g. for reliability & security) doesn’t need to be in the endpoint for the duration of a dialog
- The network can participate (and hold state), but it’s unwise to distribute critical state

What to Do?

- Some problems surmountable using existing Internet/TCP/IP model
  - 'cover up' the link problems using performance enhancing proxies (PEPs)
  - Mostly used near "edges"
  - Brittle wrt asymmetric routing, security
- But some environments never have an e2e path (or a low-loss e2e path)
- Yet we want our applications to work

Delay-Tolerant Networking

- Major Goals
  - Support interoperability across 'radically heterogeneous' networks
  - Tolerate large delays and major disruptions
- While maintaining
  - Flexibility and extensibility in support of innovation
  - Decent performance for networks with low loss/delay/errors

DTN Architecture Components

- Naming
  - generalized URI (many address families)
  - late binding (mapping) to location
- Application Data Units
  - variable-sized messages (with options)
  - can be signed, fragmented, timestamped
- Store and Forward Operation
  - 'plug-in' routing algorithm framework
  - persistent storage for store-and-forward
- Per-(overlay)-hop & E2E security

DTN Application Model

- DTN API for sending/receiving ADUs
  - agent handles bundle processing
  - asynchronous sends
  - asynchronous receipts with callbacks
- Callbacks
  - persistent registrations (~ socket bindings that span reboots)
  - can re-invoke original program or do something else
- Options for: error/ACK reporting

DTN Research

- Selected Research papers ('top tier')
  - SIGCOMM 2003– the architecture
  - SIGCOMM 2004– routing in DTN
  - SIGCOMM 2005– use of erasure coding
  - Infocom 2005/6– vehicle routing
  - Infocom 2006 (x2): sensor nets, MaxProp
  - Mobicom 2006 – Rural Internet Kiosks
  - ToN 2007 (x2) – multi and single copy routing
  - Infocom 2007 – DTN throwboxes
  - SIGCOMM 2007 – DTN as a resource allocation problem
  - Mobicom 2007 – DieselNet
  - IMC 2007 – Forwarding diversity in PSNs
  - IEEE JSAC 2008 – architecture retrospective
- Book: Delay and Disruption Tolerant Networking
- At least 6 PhD theses (inc Berkeley, Umass, UW, TCD)

State of the DTNRG

- Very active / prolific RG
  - 5 RFCs published: 4838, 5050, 5325-7
  - 14 active RG IDs + 4 author IDs
  - 3rd IETF Interop event (6+ impls)
- Diversity
  - Academic, govt, commercial (somewhat)
  - US, Canada, Ireland, Finland, Japan, Germany, UK, Sweden, India, …
- Originally meeting every other IETF
  - Now more often (outside IETF, etc)

Deployments / Demos

- SNC; N4C EU/FP7 Project
- DARPA DTN, Dielselnet, WNaN
- NASA CCSDS, etc.
- Canada – India Kiosks
- NICT - several

Issue: Checksums

- At present, no checksums
  - "checksum ciphersuite" defined
  - But security for DTN is optional
- Some controversy here:
  - Meta-data not protected
  - But, right to drop errant pkts or not?

Issue: Naming

- Naming fields are variable-length
- No real address concept in DTN
- Naming properties (today)
  - Multiple name spaces (via URI schemes)
  - 'dtn:' URI scheme reserved
  - Considering 'dtn: <URI>' format
- Late binding converts names to underlying addresses when needed
  - Some exception for network mgmt

Issue: clock synchronization

- DTN requires sync'd clocks
  - For scheduled routes
  - For expiring bundles
- Some umbridge by implementors
  - Impoverished nodes lose the time
  - Consequence: lost bundles
- How "loose" to sync clocks

Issue: routing

- Routing in DTN is both path selection and protocol selection
- Lots of research papers, (5+ yrs!)
- Not much implemented:
  - Static routing (used in tests most often)
  - DTLSR (DT link state)
  - Contact Graph Rtg (scheduled)
  - Prophet (opportunistic)
- Little real-world experience

Future: content & security

- DTN is a ‘store-carry-forward’ model
- Stored data is similar to caching
- Names can name objects
- Multiple ‘views’ on an object
  - Multi-level security on data
  - Multi-encondings

Future: a pretty big question

- Some folks want to go standards track
- Potential Issues:
  - Commercial interest
  - Energy and leadership
  - Specifies not-necessarily-IP protocols
  - Clash with other work (apps area)

Conclusions

- DTNRG is active and healthy
- DTN research = “slow success”
  - Top tier research pubs years running
  - RFCs and interoperability testing
  - Influenced govt funding / activities
- DTN as basis for ‘next-gen’ Internet
  - Reasonably well-specified and tested
  - Seems well-matched for ‘content’

Thanks

- Delay Tolerant Networking Research Group (DTNRG)
  - http://www.dtnrg.org
  - dtn-interest@mailman.dtnrg.org