Comparing Information Centric and Delay Tolerant Networking

Kevin Fall

Qualcomm Research, Berkeley

Oct 23, 2012 / LCN

Clearwater, FL

The views expressed are solely those of the author and do not necessarily represent the views of, and should not be attributed to, Qualcomm or its subsidiaries.

Networking Today

- Existing TCP/IP-based Internet
  - Every interface has 32 or 128-bit IP address
  - Identity and location tied together
  - Routing establishes single end-to-end path to host
  - Most traffic uses virtual connection (with TCP) using small best effort transfer unit (datagram)
  - Security of channels between hosts; firewalls
  - Early binding of host name to address
    - Binding security (DNSSEC) just starting

Historical Motivators for Change

- DNS – host table files too big and hard to manage
- CIDR – routing scalability concerns
- IPV6 – running out of IPv4 addresses
- TLS, IPsec, DNSSEC – security
- Delay/Disruption Tolerant Networking
  - Not always connected
  - Not always using same networking stack
- Information/Content Oriented Networking
  - Connections to hosts not of paramount importance
  - Content caching and security are; make them better

Goal

- In Delay Tolerant Networking (DTN), primary areas of focus have on
  - Supporting heterogeneous networks & naming
  - Handling disruption / disconnection
- Information Centric Networking (ICN) has been introduced as a new architectural area
  - What is it all about?
  - Does it relate to, borrow from, or replace DTN?

Major Tenets of DTN

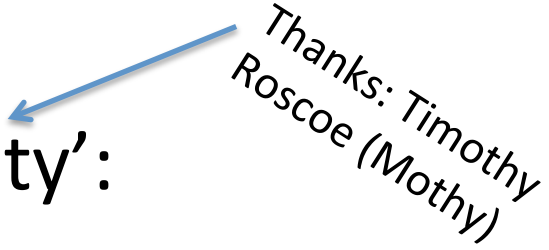

- To support 'radical heterogeneity':
  - 'Entities' can be named
  - (hosts, groups, predicates, anything else)
  - Adaptation procedures ('convergence layer')
    - For data delivery and routing in regions
- To support delay/disruption tolerance:
  - Per-node storage holds data for (long) time
  - Store-carry-forward / *n*-copy forwarding
  - Late binding (of identifiers to locators)

Information-Centric Networking

- Motivator:
  - *Data is much more important than conversations*
  - So change focus to the content / information
- A few ICN related projects:
  - NDN (CCN) – Named Data Networking (USA/NSF)
  - PSIRP->PURSUIT – Pub/Sub Internet (EU/FP7)
  - 4WARD/Netinf – Networking Information (EU/FP7)
  - DONA – Data-Oriented Network Architecture (UCB)
  - PODNET – Mobile Distribution of UG content (CH/ETHZ)
  - CONVERGENCE – Pub/Sub based on VDI (EU/FP7)
  - COMET – Content Mediator/Content Aware Nets (EU/FP7)
  - CBMEN – Content Based Mobile Edge (USA/DARPA)

ICN Tenets

- Primary unit of interest is (named data) object
- Names (IDs) are location independent
- Pub/sub [FIND/REGISTER] as basic net operations
  - Dates back to Stanford TRIAD Project (2000/2001)
- Routing function supports a pub/sub model
  - Routing on IDs (name based) natively -or-
  - Mapping between locators and IDs
  - Caching is a by-product of in-network storage
- Security primarily on content, not channel

Nuances of Naming

- What needs to be named?
  - Interface, stacks, data, groups, users
  - Brittle if names change with topology or org structure
- Structure of the namespace(s)?
  - Scope, uniqueness and authority (e.g., hierarchy or flat)
  - Impact on search/routing scale and management
- Is resolution required?
  - Early may fail faster but be wrong in high delay settings
  - Impact on latency and robustness
- Coupling and persistence?
  - Coupling can be algorithmic (e.g., "ni" names)
  - Or may require persistent storage & management for lookup

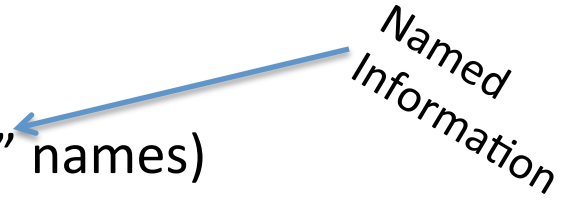

DTN and ICN Naming

- DTN names endpoints/groups or predicates
  - Using a URI-based format with scheme and SSP
    - Supports multiple simultaneous name formats
  - Designed originally as a push model to destination
  - Evolution since then using queries extension (BPQ)
- ICN names data and matches data to interests
  - Receivers express ‘interests’ in data
  - Which may be delivered from any node (cache)
  - Not naming nodes eases mobility, but raises issues

Naming Syntax and Scope

- DTN leaves exact naming open for definition
  - But starts with URI and regex match baseline
- ICN has not come to agreement (yet?)
  - Flat – couples to content, hard to aggregate
  - Hierarchical – easy to aggregate, hard to change
  - (intermediate options: scope ID)
- Resolution is required if ID->Loc is an overlay
  - Stretch versus scalability implication for routing

Routing and Forwarding

- DTN and ICN route on IDs (not nec. addresses)
  - No fixed limit on size or # of identifiers
  - No issue (or need for, hopefully) NAT
  - No major issue of re-binding or multihoming
- DTN routing has a concept of contacts
  - And times/durations they become active
  - Lots of schemes in literature (some have time loops)
- ICN (NDN) routing can be naturally multipath
  - With loops avoided by recognizing interests/data
- NDN does not address particular hosts
  - But instead content objects

More on Routing

- We understand unicast routing on a graph
  - Paths, stretch, etc. expressed as f(V, E)
  - Most often on a static (often regular) graph
- We understand something about multicast
  - Minimum Steiner tree computation is NP-hard
  - We've tried lots of others (PIM, MOSPF, RPM…)
- So, we are faced with:
  - (Distributed), dynamic, (multi-copy) routing that is not bounded in ICN by |E| or |V|… rich area?

Storage

- Storage used to be comparatively expensive
  - But thanks to Moore's Law isn't now
- Packet switching uses short-term storage
  - Store and forward queues (volatile)
  - Statistical multiplexing; short-term rate adaptation
- Just putting more into routers isn't always good
  - Buffer bloat – more latency for no benefit
- And persistent storage is now cheap too
  - Kryder's Law – 2x magnetic density increase / 18 mos

Storage

- DTN uses storage primarily for persistence and disruption tolerance
  - N-copy routing, restart, custody transfer
- ICN uses it for latency reduction and DT
  - Moderate-term storage acts as a cache
  - Also tolerates modest link outages/disruptions
- Some projects make it visible to applications
  - Lockers (SAILOR), DECADE (RFC6392)
  - Might change how media applications operate

Security Models

- Today: mostly *channel* security
  - 802.1X, EAP, IPsec, TLS [endpoint is a host/port]
- A few exceptions that focus on *content*
  - DNSSEC [observes sender maybe != author]
- Basic needs for content security
  - Integrity
  - Authentication
  - Confidentiality
  - Provenance

Data Access Control

- Many people want controlled sharing of their private information
- Establishing the threat and trust model
  - Does Bob trust his private content on Alice's box?
    - The DRM problem
  - Does Alice trust (anybody's) content on her box?
    - Isolation, sandboxing, taint tracking, IFC
  - Does an ICN network stack require attestation?
    - Does a user know server is checking signatures?

Availability

- New potential areas for DoS in ICN
  - REGISTERs and FINDs both generate traffic
  - Anonymity already built in if no source ID concept
  - Long ID parsing
  - Fragmentation interaction (e.g. with signatures)
  - Crypto processing
  - Traffic against mapping system (if one exists)
- Concern regarding diags in DTN Bundle Proto

Forwarding Performance

- DTN has not focused on this
  - And one might consequently expect slow results
  - Yet VoIP-over-DTN has been demonstrated, and…
- ICN projects have focused on this
  - In particular, variable length names (which are applicable to DTN as well) can be forwarded at reasonably high speed

Chunking and Transport

- DTN Bundle Protocol – High-Effort Delivery
  - Custody transfer between custodians
  - Preference for reliable delivery transports
  - But does not guarantee e2e in-tact delivery
- Content Chunking
  - DTN has proactive and reactive fragmentation
  - ICN approaches differ-
    - Packet size stays below path MTU (NDN/CCN)
    - Multi-Level Encodings for Objects

Table of Issues/Features

Feature	DTN	ICN
Push model	yes	No but “preplacing” content ~ similar
Interest	recent	yes
Storage	Persistent	Transient (persistent is add-on)
Channel sec	Option	No (may be in transport protocol)
Content sec	Option	yes
1-way links	Yes	No
Custodian	Integral	Separate
Node IDs	Yes	Varies [No (NDN) / Yes (Netinf)]
Conv. Layer	Yes	Yes (Netinf – explicit) / Yes (CCN – effectively)
Lifetimes	Yes	Yes (on data and on interests)
Fragments	At any node	Source only
Multicast	Nacent	Recv-driven (implicit)
Names	Regex on strings (URI)	Prefix-based names (CCN); flat names (Netinf)

Deployment Considerations

- ICN-like capabilities at the application layer continue to evolve elsewhere
  - DECoupled Application Data Enroute (IETF)
  - CDN Interconnection (IETF CDNI – RFC 6707)
  - HTTP 2.0 (IETF HTTPBIS)
- So will ICN (or DTN) be widely deployed?
  - Install base would be difficult to overturn
  - But multiple niches will continue to exist
    - Long latency, disruption prone, data center

Some Remaining Challenges

- Privacy as balanced with 'caching everywhere'
  - Users may wish to control distribution of their interests or visibility of what's cached near them
- Scalability balanced with 'name everything'
  - Routing to many more objects than nodes in a topology; indirection adds latency and stretch
- System optimality versus business rules/policy
  - Algorithms (e.g., BGP) not so elegant when forced to be tweaked according to externalities
    - Like network traffic engineering, media policies

Common Research Themes

- Routing / forwarding scalability
  - Objects not constrained by physical topology size
  - Long, variable-length names not like fixed 32 bits
  - Discovery of local nodes/objects/attributes
- In-network storage management
  - Cache eviction, custody, DoS resistance, priority
  - Multicast operations over time
- Security and privacy
  - Scalability, revocation, resource exhaustion
  - Content/policy-enforcing gateways
  - Threat model

Conclusions I

- DTN has focused on architectural components
  - Storage, custody transfer, timing, security
  - Framework for naming, routing [pluggable]
- Also, operations on some unusual networks
  - SCF ad-hoc networks, high delay, one-way
- ICN has come to some similar conclusions
  - But has focused on naming and content
  - Mostly for moderately-well-performing links

Conclusions II

- So, has DTN influenced ICN?
  - Design similarities would suggest ‘yes’
- Can DTN do what ICN wishes to do?
  - With framework components, probably so
  - But ICN does hold latency to be important
- Can ICN do what DTN does?
  - Probably so for fairly well-performing networks
  - Others are harder (1-way links, persistence, delay)

Some References

- A Survey of Information-Centric Networking, IEEE Communications, July 2012
- Information Centric Networking: Seeing the Forest for the Trees, HotNets 2011
- http://www.irtf.org/{icn,dtn}rg (2 URLs)

Thanks

kfall@qti.qualcomm.com