

QoS Discussion

Informal meeting Sunday November 17
Continuation Wednesday November 20

Existing Drafts of interest

- **Quality of Service for ICN in the IoT -**
[draft-gundogan-icnrg-iotqos-01](#)
- **QoS Treatments in ICN using Disaggregated Name Components -** [draft-anilj-icnrg-dnc-qos-icn-01](#)
- **Flow Classification in Information Centric Networking -**
[draft-moiseenko-icnrg-flowclass-04](#)
- **Considerations in the development of a QoS Architecture for CCNx-like ICN protocols -** [draft-oran-icnrg-qosarch-02](#)

Topics- what do people want to talk about

- What are the interesting research this can support?
- Architecture? Do we want to work on this?
- Experiences from IoT-land - how much should this drive our designs?
- Who cares, anyway?
- Do we only want Diffserv-like features, or do we want Intserv-like features too?
- Who will implement and measure what we do?

Notes

- Icn has some intserv-like capabilities - maybe this is a good thing to work on
 - Already per- interest-data state to build on
 - Maybe the effort is less than what we saw for IP
- Can we cleverly manage caches and other resources - this might be a strong argument
- If we do this, how to encode in packets
 - Doing in the name may encode routing info in the name.
 - What about multiple names?
 - This might be the strongest argument for separating flow classification from QoS treatment.
- What about putting QoS information in the routing protocols -
 - Use matching rules in the routing distribution/calculation
 - Two things - using routing to express classification versus doing QoS based routing - these are probably separate
- Need to clean up our terminology and be consistent on what we mean by classification and treatment
 - Specific: clear up meaning of flow control versus congestion control versus QoS treatment

More notes 2

- Started from scratch (at least inside Cisco) - want to move forward.
 - Need a common draft showing how things fit together
 - Make progress on the individual pieces in parallel
- What do we do with Dave's QoS architecture draft?
 - Tentative - let Dave publish as individual submission
 - Then start over as RG activity
- Write down on pros/cons of each of the proposed encoding mechanisms and the consequences.

More notes 3

- Resource coordination across nodes can improve service for everybody.
 - This may be easier with ICN than IP because of stateful forwarding and caching.
 - Also have more resource tradeoffs to exploit (e.g. bandwidth versus cache space)
- Non-orthogonality of resources versus QoS treatment:
 - Low latency interaction with high reliability - does high reliability allow unbounded impact on low latency traffic
 - One thought - latency has higher precedence to keep PIT state, while high reliability has higher precedence on cache space.
 - How much of this is specific to low-end (IoT) - Carofiglio et.al have argued that it never makes sense to have less than a BDP of PIT resource.
- Where to take the flow classification draft?
- If you distribute classification via routing, who is authoritative for what?
 - What gets precomputed via routing and what gets computed on demand when Interests arrive?

More notes 4

- Is there any difference in requirements for queues between IP and ICN?
 - Types of queues? Russian-doll, WFQ, SFQ, PQ, etc.?
- Running on top of IP versus an L2? Any differences?
 - How do we leverage underlying machinery?
 - What about tunnels?
 - Variable bandwidth - can scheduling be done accurately based on Interest/Data exchanges?
- Inter AS QoS - can ICN do better than Diffserv?
 - Diffserv was totally inadequate in these cases
 - Forced to ignore or remark.

Notes 5 (Wednesday)

- Asymmetry (problem or opportunity?)
- Relevance to Time-sensitive Networking (network layer support?)
 - Can ICN QoS address some deterministic networking objectives?
- Multicast & QoS (multi-destination)
- Network coding relationship?
 - NWCRG will discuss NW coding and congestion control (and QoS?) this week
- CCNInfo (Multipath communication with producers)
- IoT use case: resource-constrained
 - fine granular resource control, more resources, more options to implement QoS
- Inter-domain (multi-stakeholder in general): can ICN-QoS facilitate/enable QoS across domain/network boundaries
- Stateless forwarding model in ICN makes it difficult to coordinate resource usage
 - Seems good research topic
- Generalize concepts and some mechanisms?
- Terminology document does not address QoS -- could be useful exercise to develop good common understanding
 - Congestion control vs. flow control
 - Queuing semantics
- Content delivery, multimedia streaming use cases -- how would ICN fit into deployment?
- deadline-driven transport & interaction with queuing: how would these concepts relate to ICN QoS?
-

Notes 6 (Wednesday)

- Receiver-driven resource allocation
 - Capacity sharing and optimizations
 - Accounting for resource usage
 - Receiver-driven congestion control did not work well -- unclear we have insights into why this was the case
 - Object fairness vs. transmission fairness (flow fairness not meaningful in ICN context) -- cf. Edmund Yeh, <https://conferences.sigcomm.org/acm-icn/2014/papers/p117.pdf>
 - “Data-driven QoS”
- ICN-QoS: “Cache Pollution” with selected content...

Notes 7 (Wednesday)

- QoS Arch Considerations Draft
 - Do a few more edits and then try to get it published through ICNRG soonish
 - so that it would not be part of a QoS architecture work item in ICNRG
- Scope of potential QoS work item in ICNRG
 - Terminology!
 - Document research challenges -- and “what not to do”
 - QoS-based forwarding and QoS-based routing (preferential treatment for selected users/flows/objects)
 - Design team? Curator? Facilitator?
 - Would be good to have implementations and experiments soon: opportunity for hackathon work...

Notes 8 (Wednesday) -- Discussion of additional Ideas

- Low-latency, deadline-driven QoS
 - Earliest-deadline first (EDF)-like schemes
 - EDF on per-link basis probably not a good idea -- e2e EDF scheduler for INTEREST/DATA exchanges (will affect protocols)
- Insights from INTSERV
 - 2 services:
 - Controlled load service
 - Guaranteed service
 - How far do you allow the actual delivered service to diverge from specified service description
 - Explicit measure for applications to allow some divergence (could also be interesting in TSN world...)
 - Poster-child: Real-time multimedia
- Delay vs jitter