

# Improving the Interactions Between Transport & Network Mechanisms

Lars Eggert, Magnus Westerlund and Jari Arkko

IETF-66, Montréal, Canada

# Background

- disconnect between the INT and TSV areas on how transports and mobility mechanisms interact
- there's been a number of BOFs and BOF proposals in this problem space
- sometimes with problematic scopes, sometimes got formed and ran out of steam
- anyway, these are indications that there is energy in the IETF/IRTF to do something in this area
- this talk attempts to encourage such work and provide a possible scope

# Layering Is Good

- layers in the network stack provide communication abstractions
- they expose well-defined sets of operations & information
- hide layer-internal intricacies from their users (and details of layers further “down”)
- this is good! modular design at work

# Example: Network Layer

- abstraction is something like
  - “will deliver your packets in some order”
  - “may deliver multiple copies of some packets”
  - “may not deliver some others”
- hides other network-layer functionality, such as
  - packet fragmentation/reassembly
  - route computation and forwarding

# But in Practice...

- users of the network-layer abstraction, i.e., mostly the transport protocols, have made additional assumptions about it
- these assumptions are the basis of many key transport-layer mechanisms, such as
  - congestion control
  - flow control
  - reliability mechanisms

# Example Assumptions

- hosts remain at the network port identified by an IP address for long times
- packets between the same src/dst addresses mostly follow the same path
- paths change on time scales that are orders of magnitude greater than the RTT
- path characteristics change on similarly large time scales
- connectivity along a path is very rarely disrupted

# Reality Check

- many of these assumptions are no longer generally true throughout the whole network
- especially with recent/proposed network layer extensions, such as MIP, HIP, SHIM6, NEMO, etc.
- but also simply because recent link technologies are different
  - network-based mobility
  - link-layer retransmissions
  - non-congestion-related packet loss

# Consequences

- traditional transport mechanisms are performing less well than in the past
- not news: resulted in a gazillion of “optimize transport protocol  $X$  for scenario  $Y$ ” proposals
  - where  $X$  is mostly TCP
  - and  $Y$  = satellites, 802.11, GSM, 3G, ad hoc networks, high bit-error links, etc.
- but vast majority of these proposals are band aids
  - specific fixes for limited scenarios
  - not appropriate for a general-purpose Internet



# What Could Be Appropriate?

- idea: extend the communication abstraction that the network layer provides to its users
- but keep it independent of specific...
  - network-layer extensions
  - link technologies
  - application or deployment scenarios
- result should be universally and incrementally deployable

# Not a New Idea

- existing mechanisms already enhance the network-layer communication abstraction in this way
  - ECN: “I’m about to start dropping these packets”
  - Quickstart: “you may send me packets at rate  $n$ ”
  - XCP: evolved congestion control framework
- and don’t forget about ancient stuff like ICMP
  - unreachable: “this host/network is not here”
  - source quench: “stop sending so fast”
- all these define new pieces of network-layer information that new transport mechanisms act on

# General Principle?

- examples on the previous slide follow same general principle
  - provide additional information about network-layer events to transport protocols
  - information should be advisory & optional: transports shouldn't depend on them
  - new transport mechanisms act on this to improve operation and performance
- seems like a useful general principle for approaches in this space

# What Next?

- TSV and INT ADs started a discussion list for people interested in this problem space
  - [ternli@ietf.org](mailto:ternli@ietf.org)
  - <https://www.ietf.org/mailman/listinfo/ternli>
- goal is to identify interested parties and discuss what work could be done where in the IETF/IRTF (and then start doing it...)