

Resource Pooling for the Future Internet

NOKIA

**Lars Eggert
Nokia Research Center &
Helsinki University of Technology**

**2nd Japan-EU Symposium on the
New-Generation Network and Future Internet
Tokyo, Japan, October 13-14, 2009**

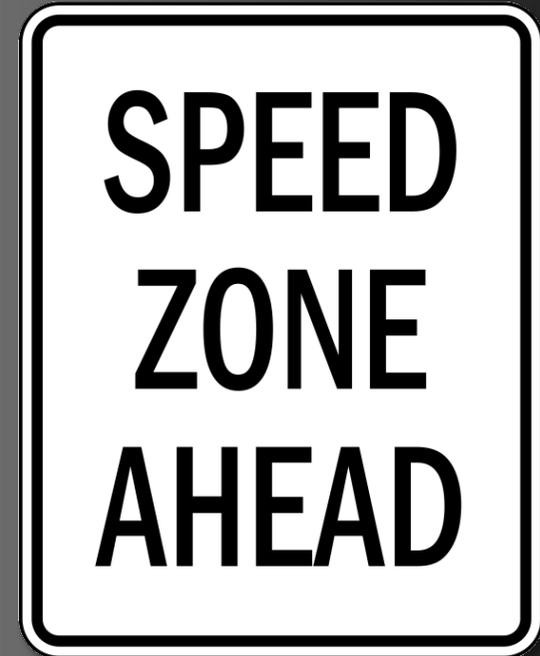
Premise

we need a **more robust Internet** than we can get from simply making better individual components

most hosts will be **mobile**, with **multiple radios** in use concurrently
applications are becoming more demanding & important

we don't have a clean slate =
incentives are more critical than ever!

incentive misalignment kills many technically viable proposals
*partial benefit from partial deployment
change is easiest where pain is felt*



Resource Pooling for the Future Internet

make the network resources behave like a **single, pooled** resource

increase reliability, flexibility and performance

by exploiting **parallelism** and **redundancy** in the network

through an **evolution** of the current Internet protocols



Components of a Resource Pooling Architecture

Multipath Transport

transmit the data of a single transport connection along multiple paths

coupled congestion control loops are critically important (Kelly/Voice, etc.)

Resource Accountability

expose the impact of its resource usage to an end system (and “charge” for it)

creates the correct incentives to behave smartly

Multipath Routing

make multiple, disjoint paths available between two endpoints

the only real way to robustness is via diversity

Components of a Resource Pooling Architecture

Multipath Transport

transmit the data of a single transport connection along multiple paths

coupled congestion control loops are critically important (Kelly/Voice, etc.)

Resource Accountability

expose the impact of its resource usage to an end system (and “charge” for it)

creates the correct incentives to behave smartly

Multipath Routing

make multiple, disjoint paths available between two endpoints

the only real way to robustness is via diversity

Multipath Transport Overview

multipath transport treats multiple end-to-end paths as a single pooled resource

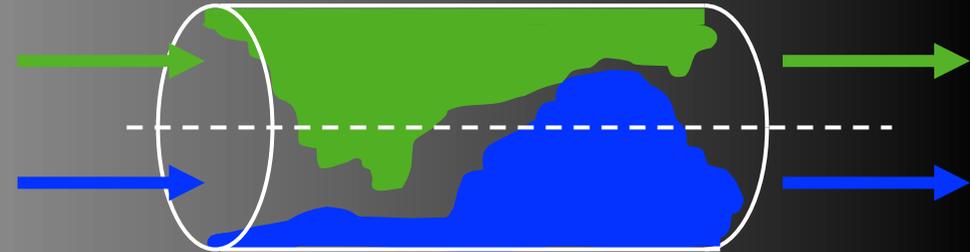
both in terms of robustness and bandwidth

linked congestion control for bottleneck fairness

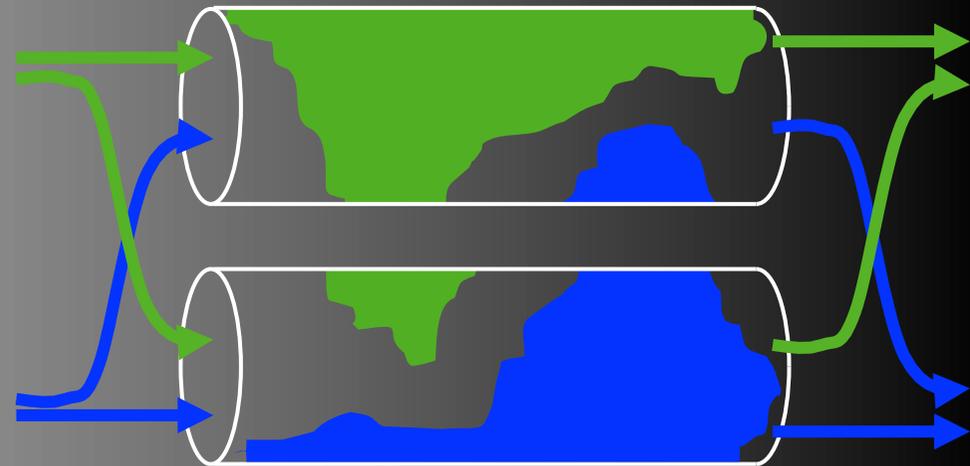
traffic moves away from congested paths naturally

larger bursts can be accommodated

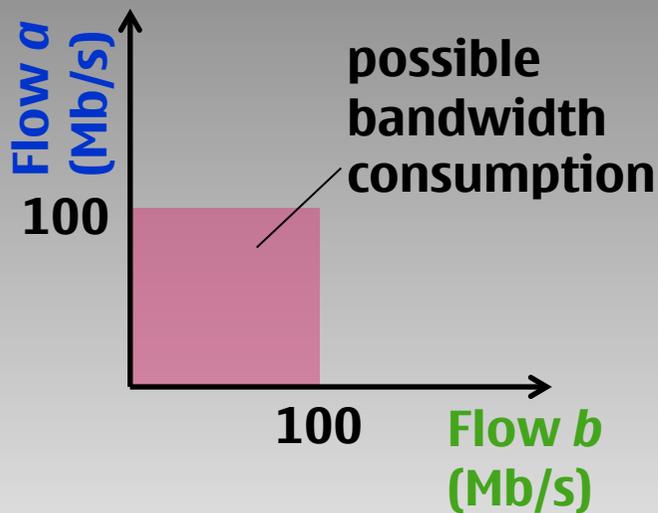
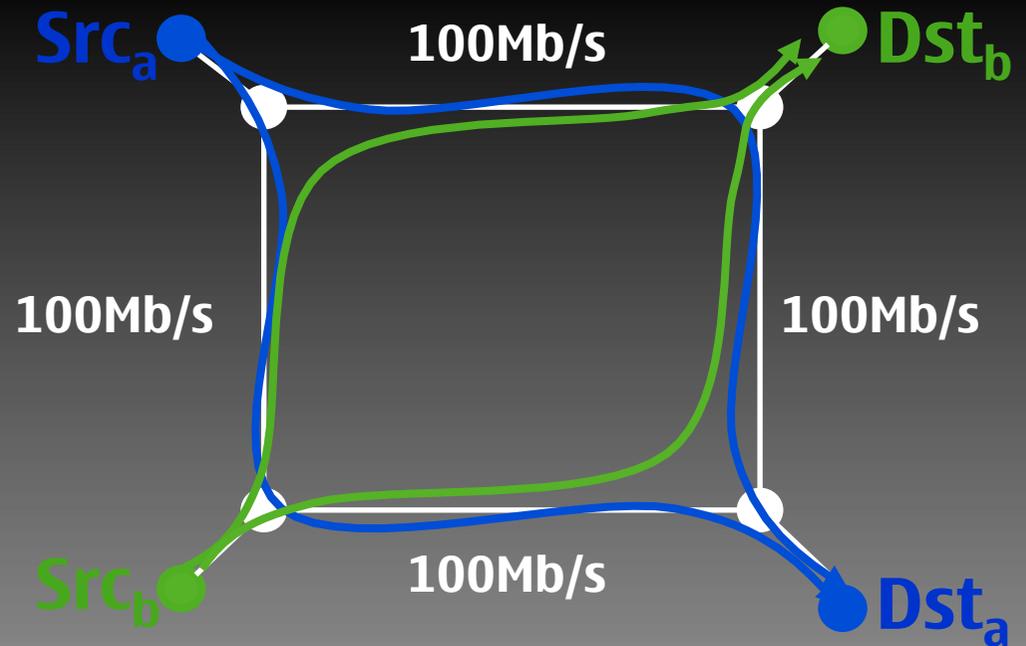
traditional Internet:



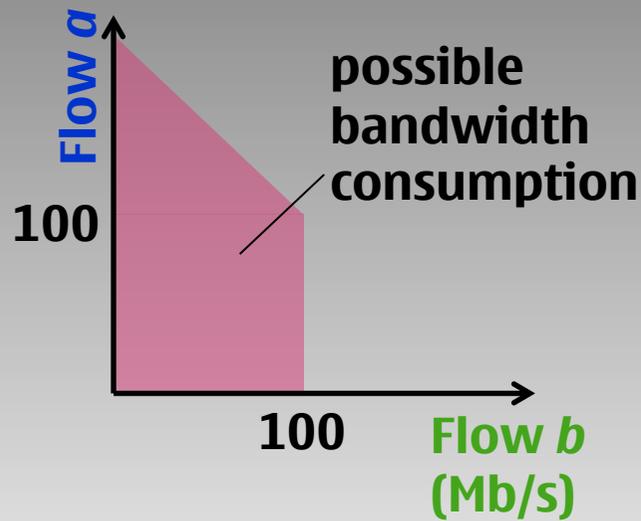
multipath resource pooling:



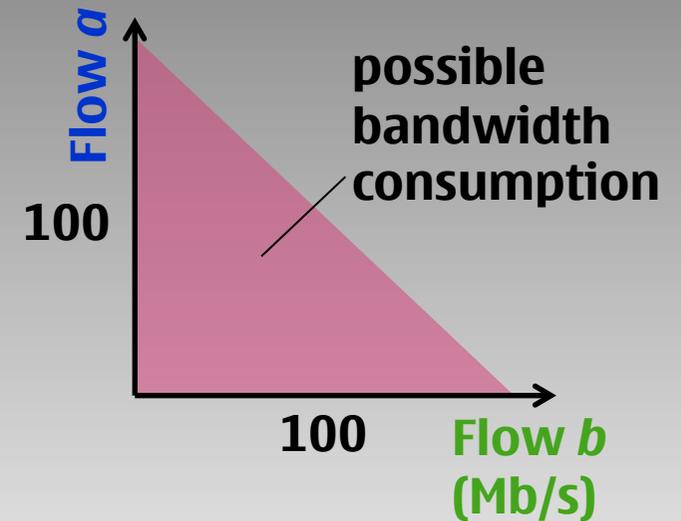
Multipath Transport Allows a Wider Range of Traffic Matrices



no multi-path flows



only flow a is multi-path



both flows are multi-path

Internet Transport means TCP

applications and higher-layer protocols use TCP to transmit data across the Internet

this is unlikely to change anytime soon (quite the opposite)

so any other transport protocol is a non-starter (incentives!)

but standard TCP only transmits along a single path...



Resource Pooling + TCP = Multipath TCP

extend TCP to allow one connection to send data along multiple paths between the same two end systems

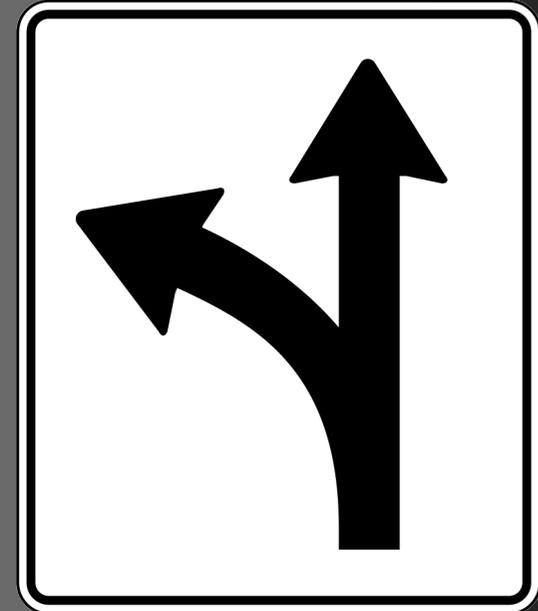
concurrently, or changing over the lifetime of the TCP connection

flexibility, performance, robustness

must be fully backwards compatible

- same socket API for apps

- looks like standard TCP (+ options) to network



Current Status

initial research is complete

protocol engineering is underway

IETF standardization starts 2H2009

“Multipath TCP” Working Group

strong interest from Apple, Google,
Microsoft

Windows 8 roadmap?

prototype code available



Components of a Resource Pooling Architecture

Multipath Transport

transmit the data of a single transport connection along multiple paths

coupled congestion control loops are critically important (Kelly/Voice, etc.)

Resource Accountability

expose the impact of its resource usage to an end system (and “charge” for it)

creates the correct incentives to behave smartly

Multipath Routing

make multiple, disjoint paths available between two endpoints

the only real way to robustness is via diversity

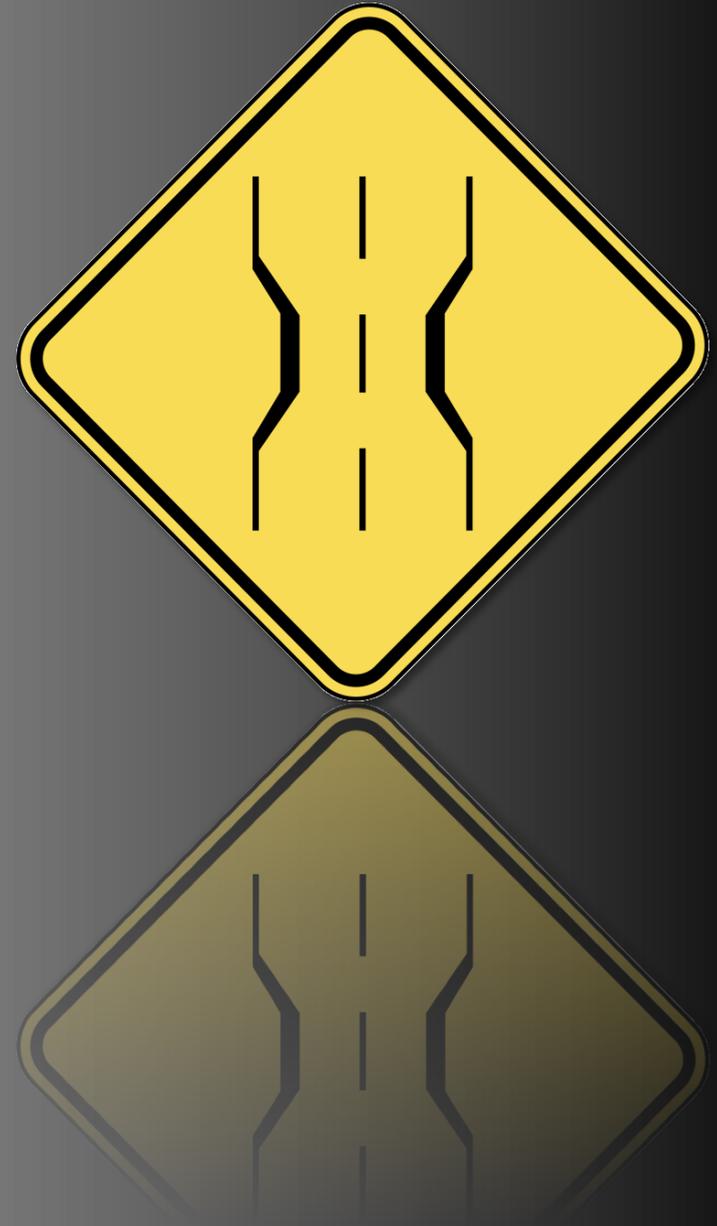
Resource Accountability

the current Internet does not adequately account for whose sessions use which fraction of its communication resources

example: ISPs account for traffic volume when they should be accounting for “congestion caused”

a Future Internet must have more useful accounting & reporting

so users of the network have the information & incentives to behave smartly



Congestion Exposure

a mechanism for exposing how much a user's traffic is contributing to current path congestion

to the user himself and the ISPs

ISPs charge for "congestion caused"

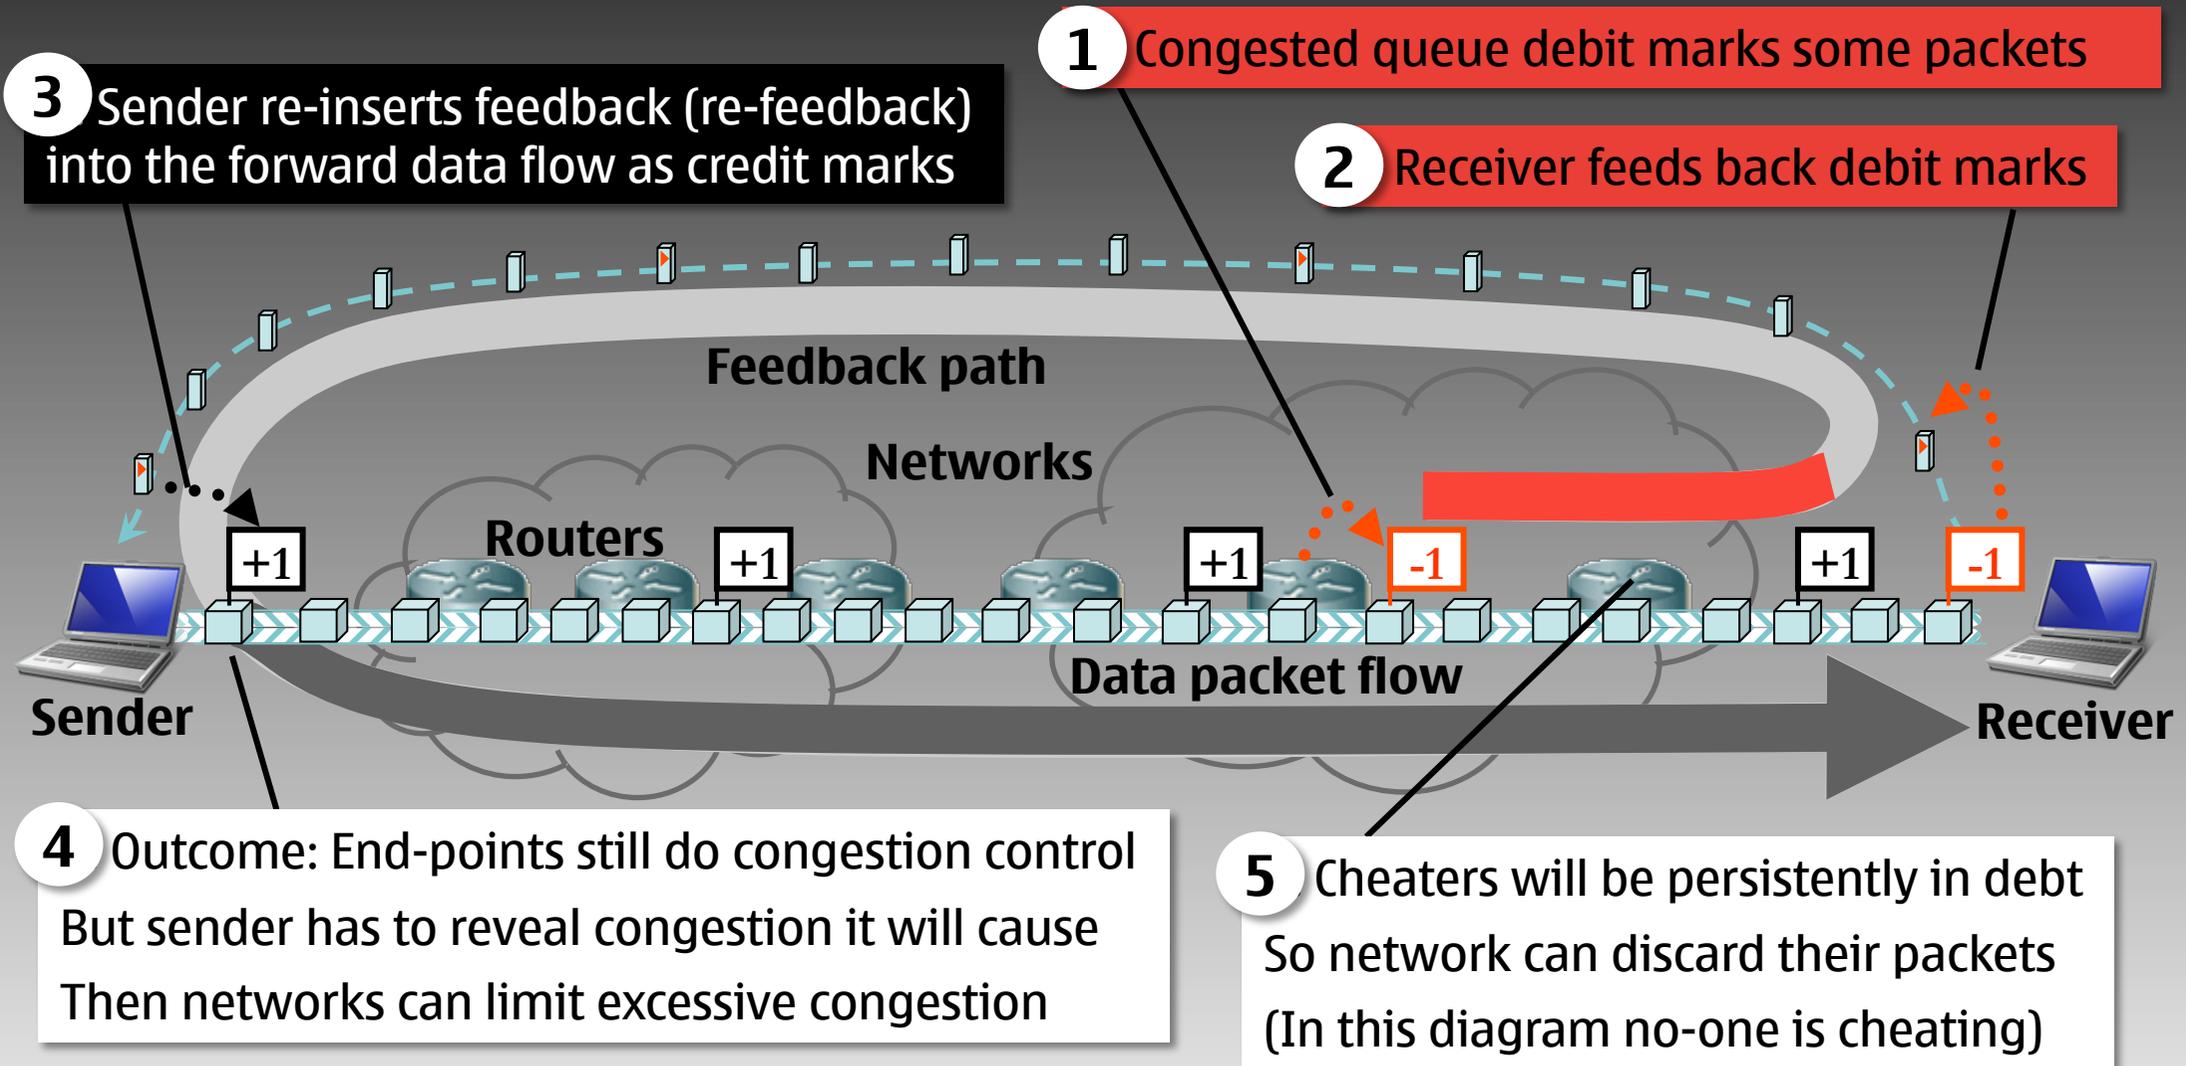
users have an incentive to use applications & protocols that are smart about avoiding congestion

or they can decide to pay a premium
creates flexibility, choice & incentives

proposed protocol: Re-Feedback



Re-Feedback



Current Status

research/engineering on congestion exposure (“re-feedback”) has progressed farthest

“Congestion Exposure” BOF in Nov 2009 at IETF-76 in Hiroshima, Japan

resource accounting & control ideas for other communication resources are being investigated



Components of a Resource Pooling Architecture

Multipath Transport

transmit the data of a single transport connection along multiple paths

coupled congestion control loops are critically important (Kelly/Voice, etc.)

Resource Accountability

expose the impact of its resource usage to an end system (and “charge” for it)

creates the correct incentives to behave smartly

Multipath Routing

make multiple, disjoint paths available between two endpoints

the only real way to robustness is via diversity

Multipath Routing

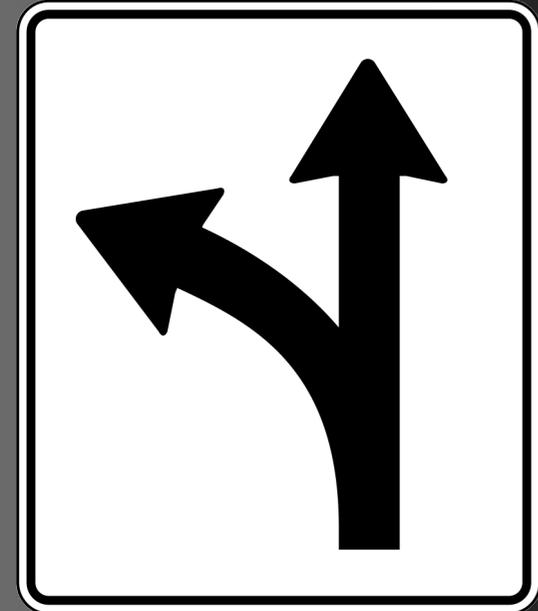
how can we bring the benefits of multipath transports to hosts that only have one **single** network interface?

when the ISP they connect to is multihomed

or when multiple paths exist “deeper” in the network

(multipath theory says that complete path disjointness is not required, as long as there is no shared path bottleneck)

Internet routing = BGP



Resource Pooling + BGP = Multipath BGP

computing & using multiple BGP paths
at the same time

potentially better traffic engineering
better resilience + convergence
increase utilization of the network

challenges

incremental deployment
loop-freeness
end-to-end consistency



Work in Progress

loop-freeness & incremental deployment

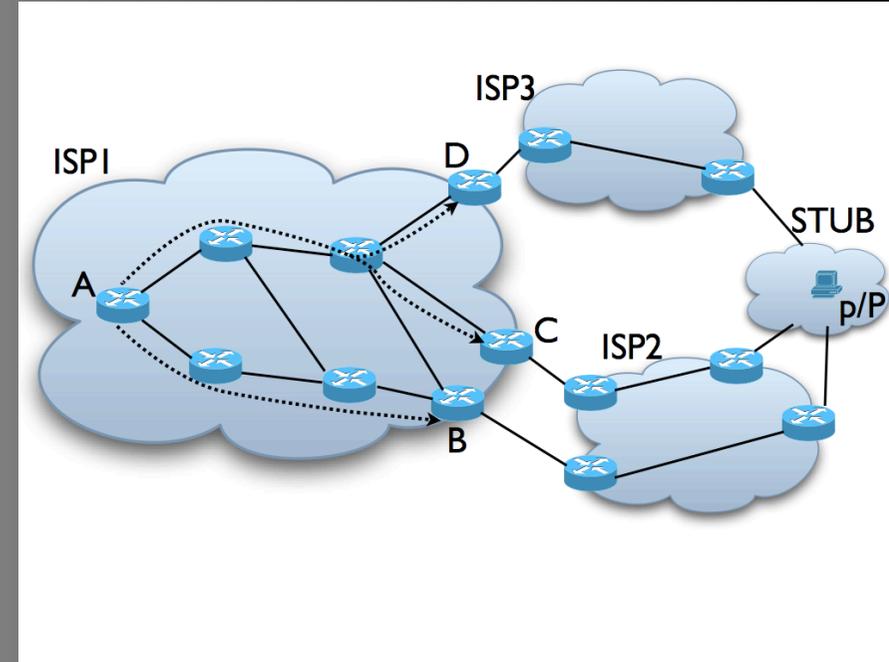
only propagate a single path to avoid loops

longest used (not longest known)

first result: understand the dynamics

slightly faster convergence in some chosen topology

many more results needed to understand resulting path diversity



The **trilogy** Project



Develop a unified control architecture for the Future Internet that adapts in a scalable, dynamic & robust manner to local operational and business requirements.

Develop and evaluate new technical solutions for key Internet control elements: reachability & resource control.

Assess commercial and social control aspects of the architecture & technical solutions, including internal & external strategic evaluation

Funded by the EU under FP7 for 3 years (2008-10)
<http://www.trilogy-project.eu/>

