

Resource Pooling and the Trilogy Project



Lars Eggert

(with thanks to Mark Handley, Damon Wischik and Marcelo Bagnulo)

The IEEE 22nd Annual Computer Communications Workshop (CCW)

Steamboat Springs, CO, USA

October 22-24, 2008



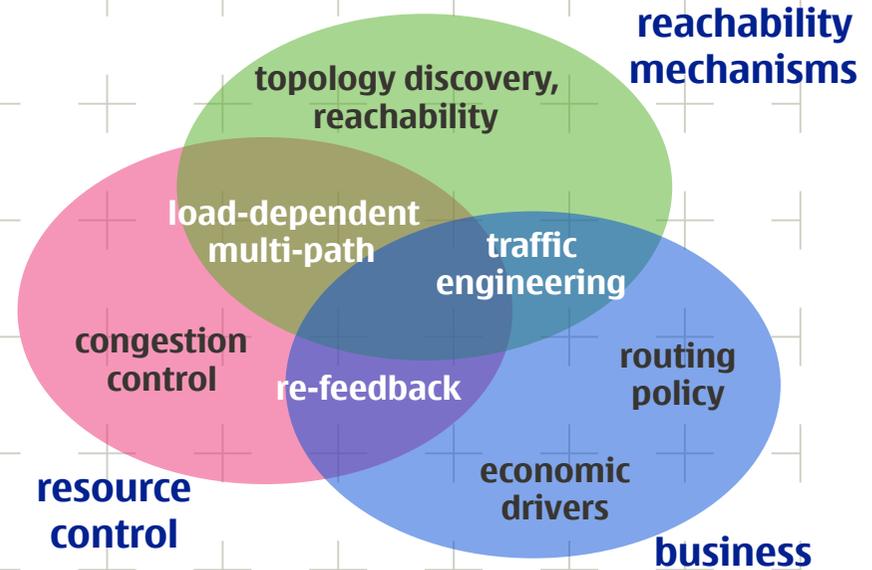
NOKIA



Develop a **unified control architecture for the Future Internet** that can adapt in a scalable, dynamic and 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 our architecture & technical solutions, including internal & external strategic evaluation



Funded by the EU under FP7 for 3 years (2008-10)

Total volume: 9.15M€
EU: 5.82M€

~60 person-years total

<http://www.trilogy-project.eu/>



The architectural requirements have changed

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

traditional routing can't solve this in a scalable way

applications are becoming more demanding (VoIP, TV, Games)

most of the end-systems will be mobile, with multiple radios that can be used simultaneously

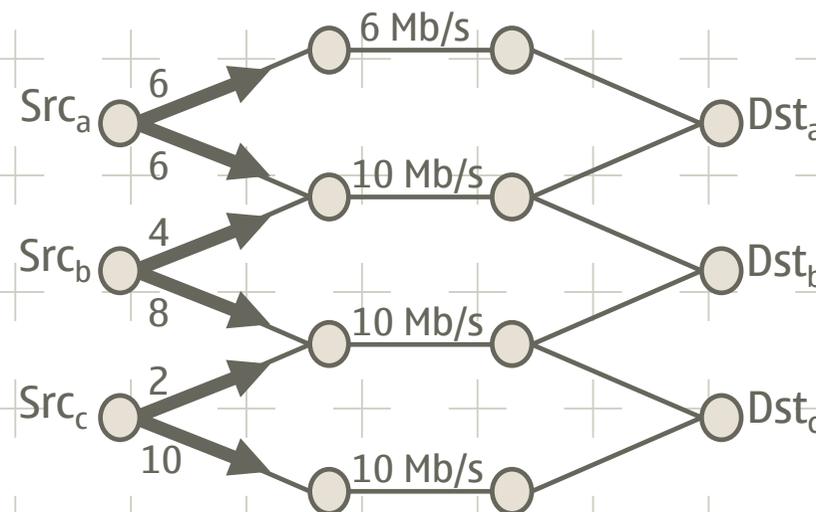


Resource pooling

make a network's resources behave like a single pooled resource

aim is to increase reliability, flexibility and efficiency

method is to build mechanisms for shifting load between the various parts of the network



Resource pooling is not new...

Routing

BGP traffic engineering

slow, manual process to pool
resources across peering links

OSPF/MPLS traffic engineering

slow, mostly manual process
to pool resources across
internal ISP links

BT, AT&T and others

dynamic alternative routing

Elsewhere

multi-homing

pool reliability & capacity

Google, Akamai, CDNs

pool reliability & bandwidth

BitTorrent

pool capacity & reliability

Theoretical foundations

Kelly and Voice

Key, Massoulié and Towsley



Resource pooling for the Internet

multipath

only real way to get robustness is redundancy

multi-homing, via multiple addresses

can aggregate routing information

mobility, via adding and removing addresses

no need to involve the routing system

or use non-aggregatable addresses



Approach

multipath-capable transport layers

- use multiple subflows within transport connections
- congestion control subflows independently
- traffic moves to the less congested paths

note: the involvement of congestion control is crucial

- link the back-off parameters for stability and fairness (Kelly+Voice)
- you can't solve this problem at the IP layer

moves some of the stresses out of the routing system

- might be able to converge slowly and no-one cares

eventually, routing system should expose in-network multipath availability, so single-homed end systems benefit



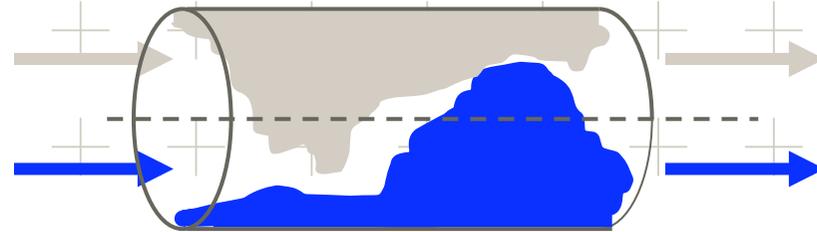
Multipath transport

multipath transport allows multiple links to be treated as a single pooled resource

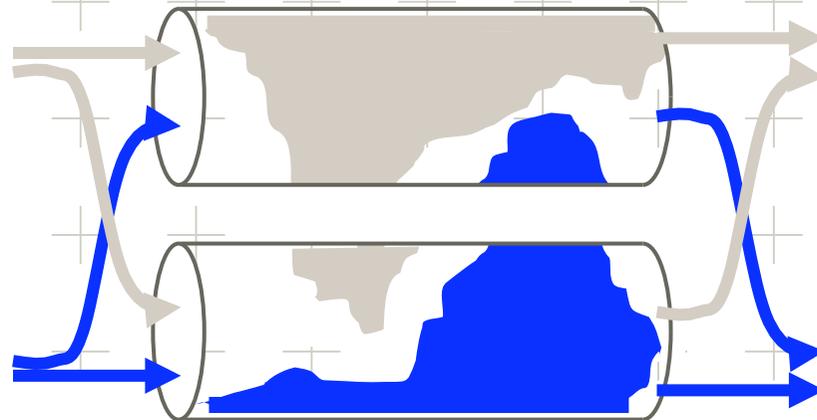
traffic moves away from congested links

larger bursts can be accommodated

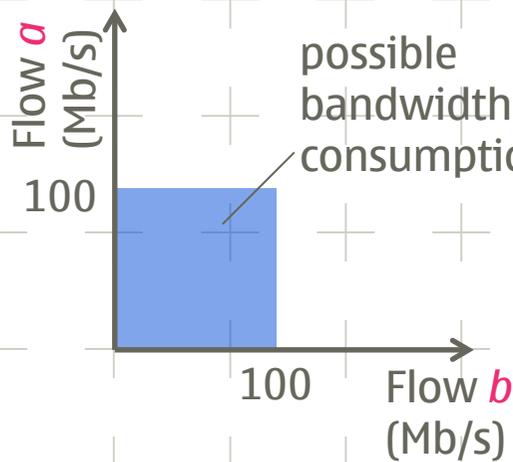
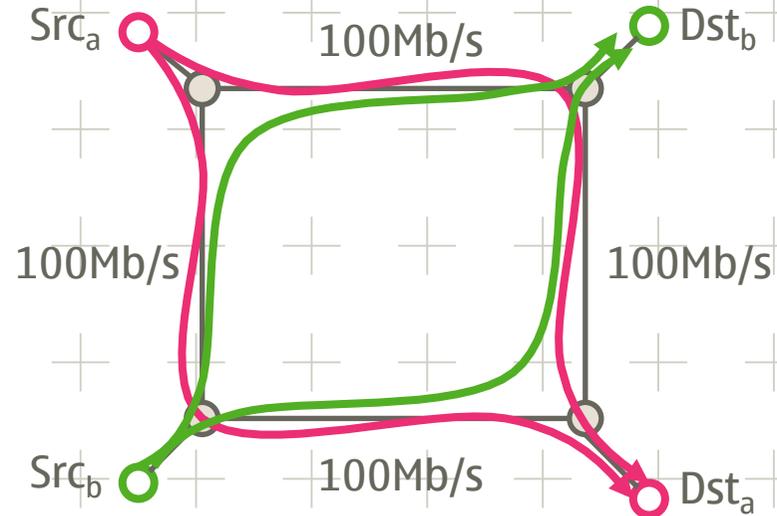
ARPAnet resource pooling:



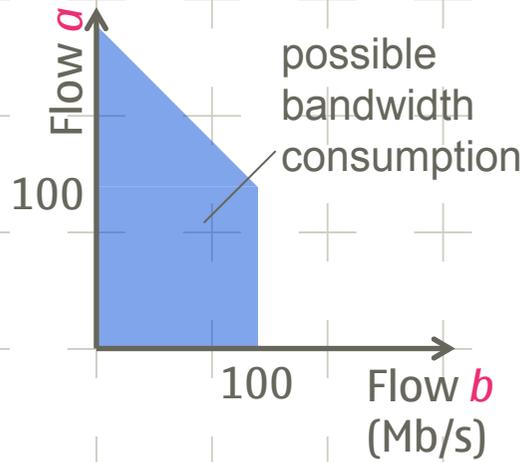
multipath resource pooling:



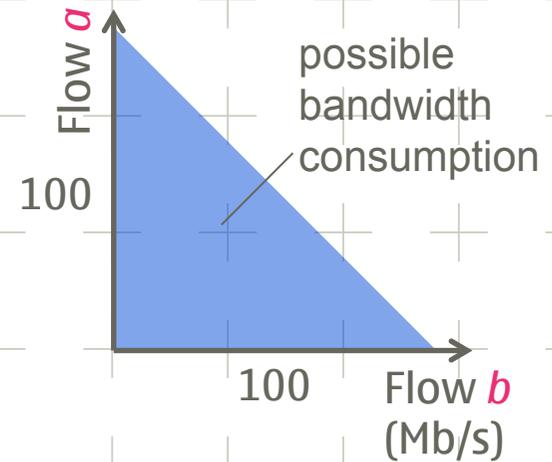
Resource pooling allows a wider range of traffic matrices



no multi-path flows



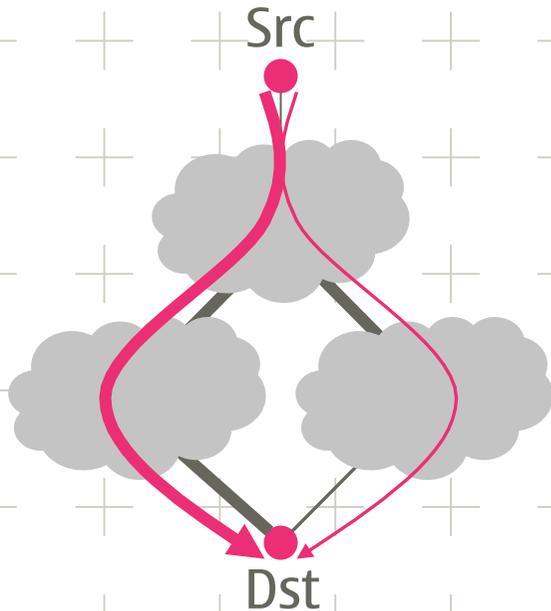
only flow a is multi-path



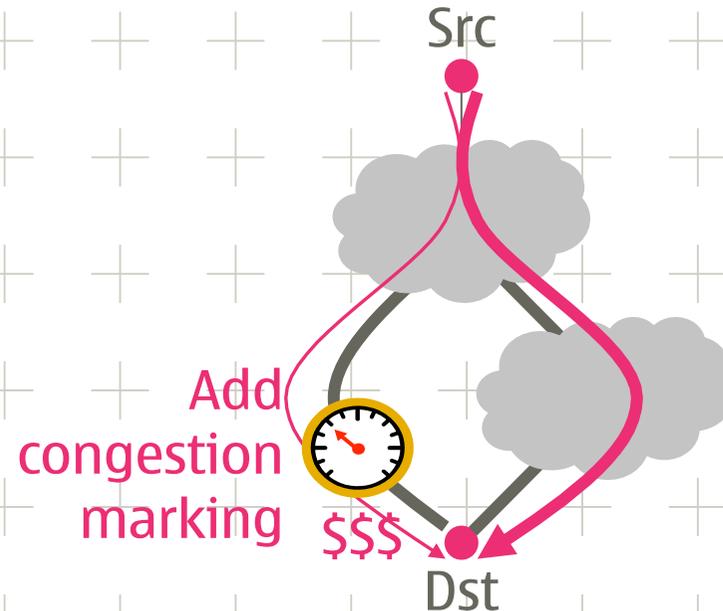
both flows are multi-path

Multipath traffic engineering

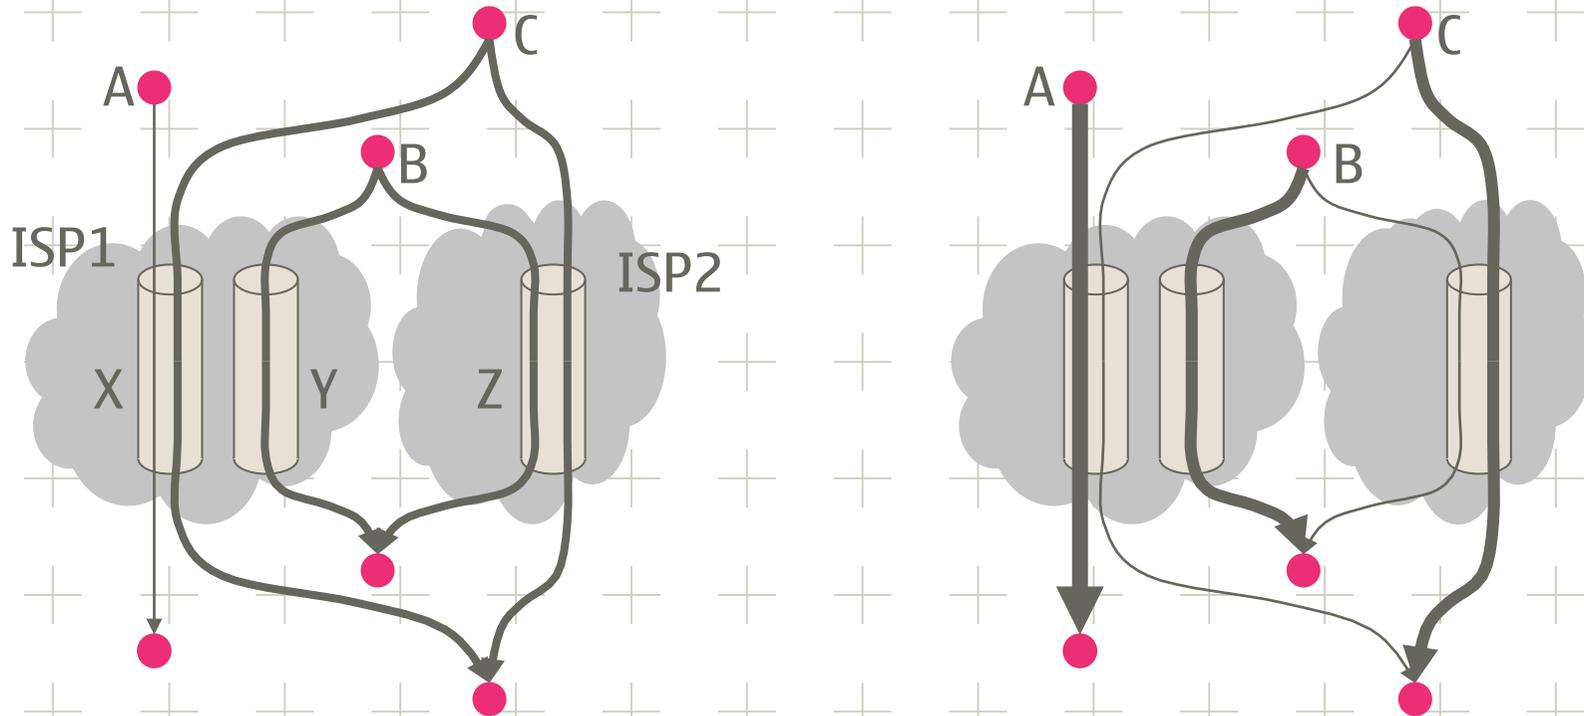
Balancing across
dissimilar speed links



Balancing across
dissimilar cost links



End-systems can optimize globally (often ISPs cannot)



Where are we today?

good theoretical understanding of the issues (past work by others)

Kelly and Voice; Key, Massoulié and Towsley

Trilogy is working on the details for TCP & BGP

how well does this work in practice?

are there cases where multipath does worse?

how much of the traffic engineering problems does this solve?

how much remains to be done in routing?

how to manage such dynamic networks?

(Trilogy is also investigating other topics)

<http://www.trilogy-project.eu/>

