Towards MultiPath TCP Adoption: Challenges and Opportunities

A. Kostopoulos, H. Harma, T. Levä, B. Heinrich, A. Ford and L. Eggert

NGI 2010 – 6th Euro-NF Conference on Next Generation Internet
June, 2010

Athens University of Economics and Business – RC

Introduction

- Internet → *reach fundamental capability limits*
- Increasing resilience and performance requirements
- FI research → *develop new architectures and protocols that address emerging technical deficiencies.*
- *Design solutions that deliver effective and efficient control of resource sharing.*

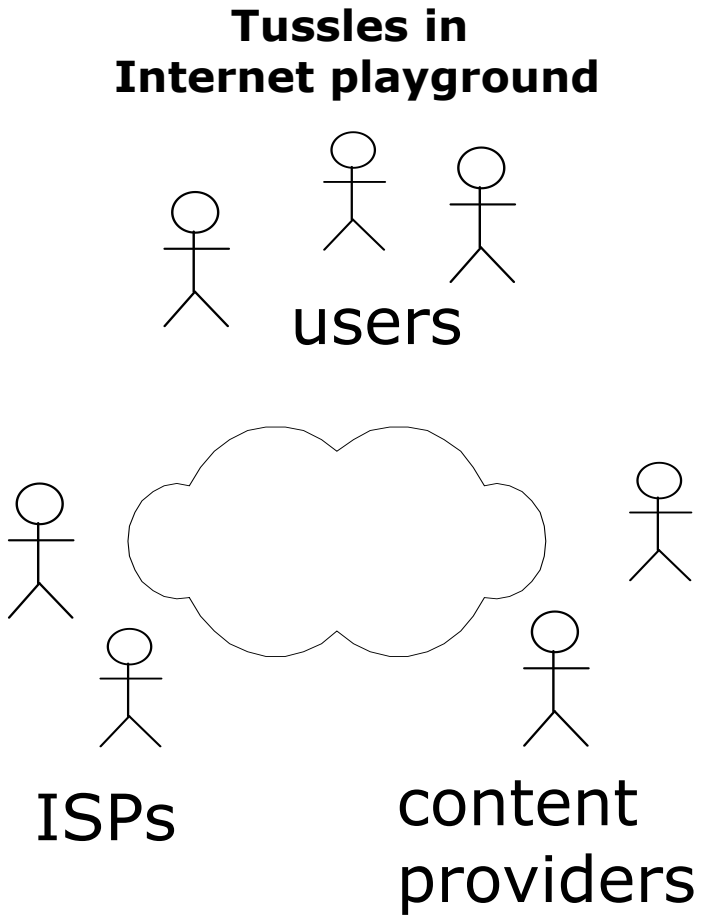

Motivation

- Transport protocols → *only a single path between a source and a destination*
  - limits the achievable throughput.
- Firewalls / middleboxes reject packets which are not using TCP or UDP
  - → affected the deployment of other transport layer protocols
  - → MPTCP has to overcome this challenge as well.

- **Multipath TCP** uses multiple paths at the same time to transmit the data belonging to a single TCP connection.
- In the case of congestion along a path, or even a complete path failure, MPTCP shifts traffic onto other available paths that have available capacity.
- Reliability, flexibility and throughput.

Outline

- The network protocol deployment differs from the diffusion of end user centered innovations (consumer products).
  - Operating system vendors play a major role.
  - Users cannot directly select network stacks for their end systems.

- New challenges for the involved stakeholders.
  - *A framework for analyzing MPTCP deployment*
  - *Key factors that make MPTCP deployable*
  - *Identification of the involved stakeholders*
  - *Deployment process*
  - *Possible scenarios that facilitate the required steps to support MPTCP adoption.*

Proposed Adoption Framework

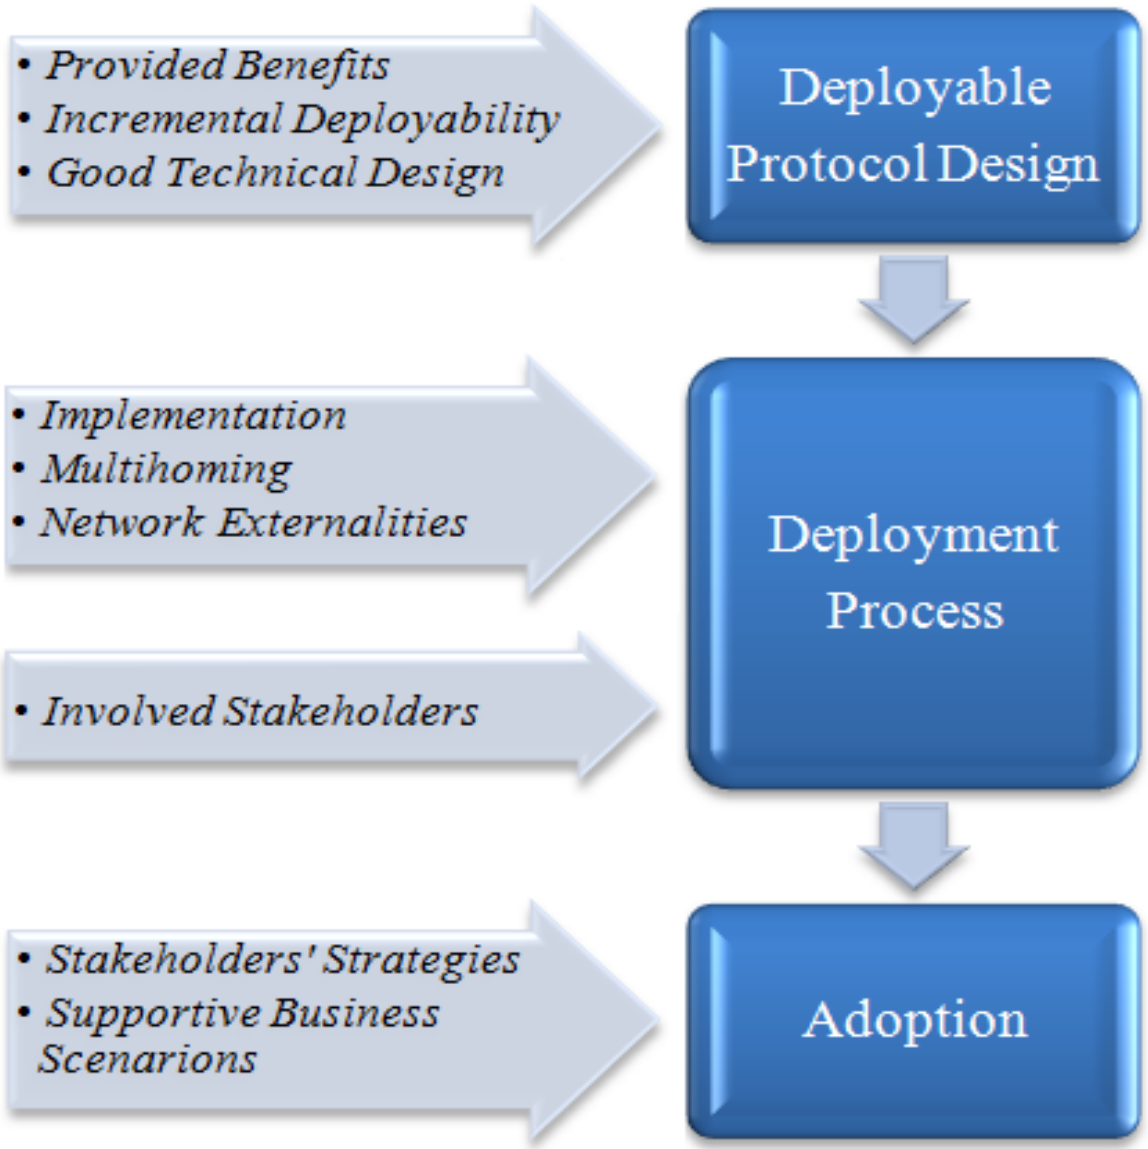

Provided Benefits

- **A real need is met** → An identified problem solved better than other approaches.
- An MPTCP connection uses several paths for a single connection at the same time:
  - *In case of congestion or a failure along one path, MPTCP can make greater use of less congested alternate paths.*
  - *MPTCP pools the available capacity along all paths for a single connection → faster transfers than traditional TCP.*
  - *Coupled congestion control*

- Mobile (battery-powered devices): Sending and receiving data across multiple radio interfaces increases the energy consumption of network communication
- Interesting feature → switch an established connection between different paths or
  - aggressively switch an MPTCP connection to the most energy-efficient path

Uncoupled vs. Coupled CC

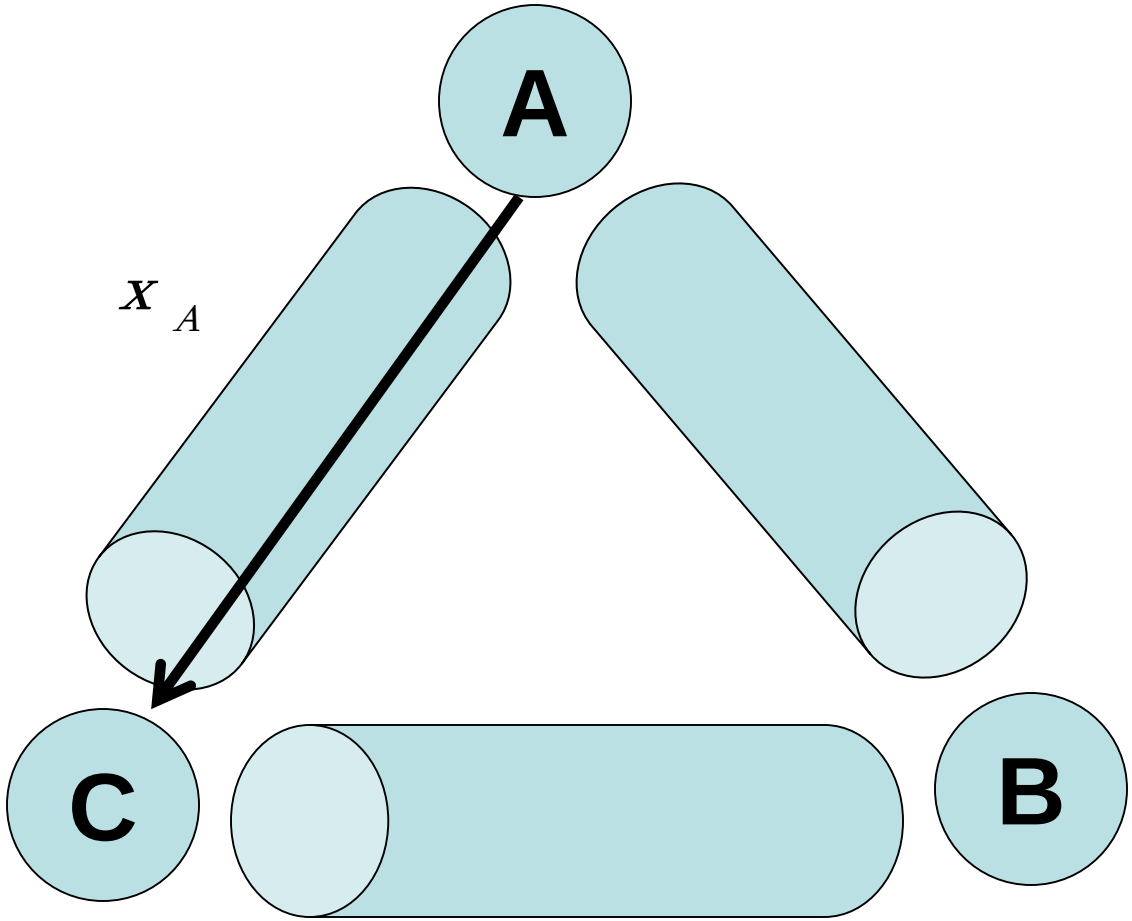

Uncoupled vs. Coupled CC

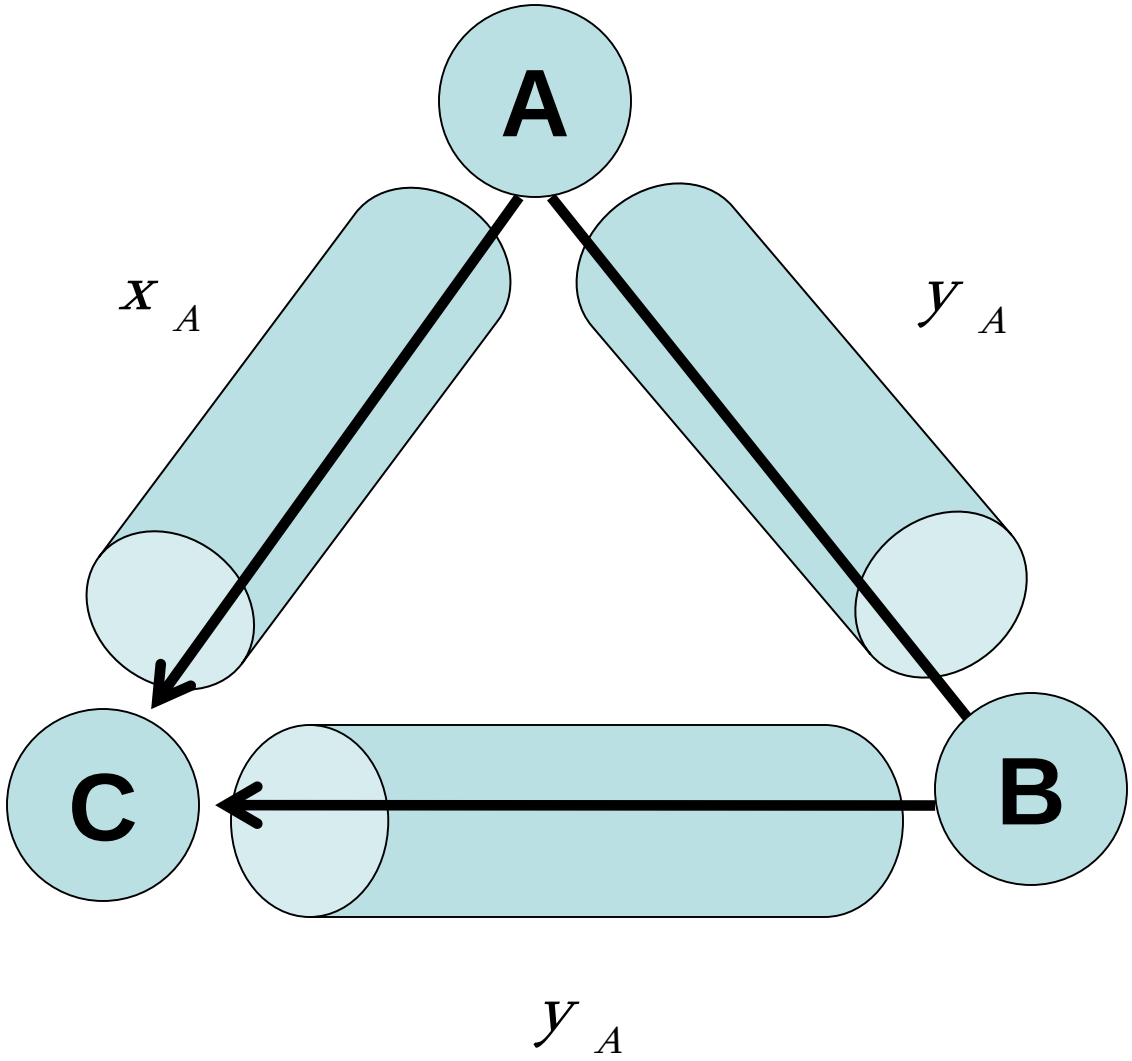

Uncoupled vs. Coupled CC

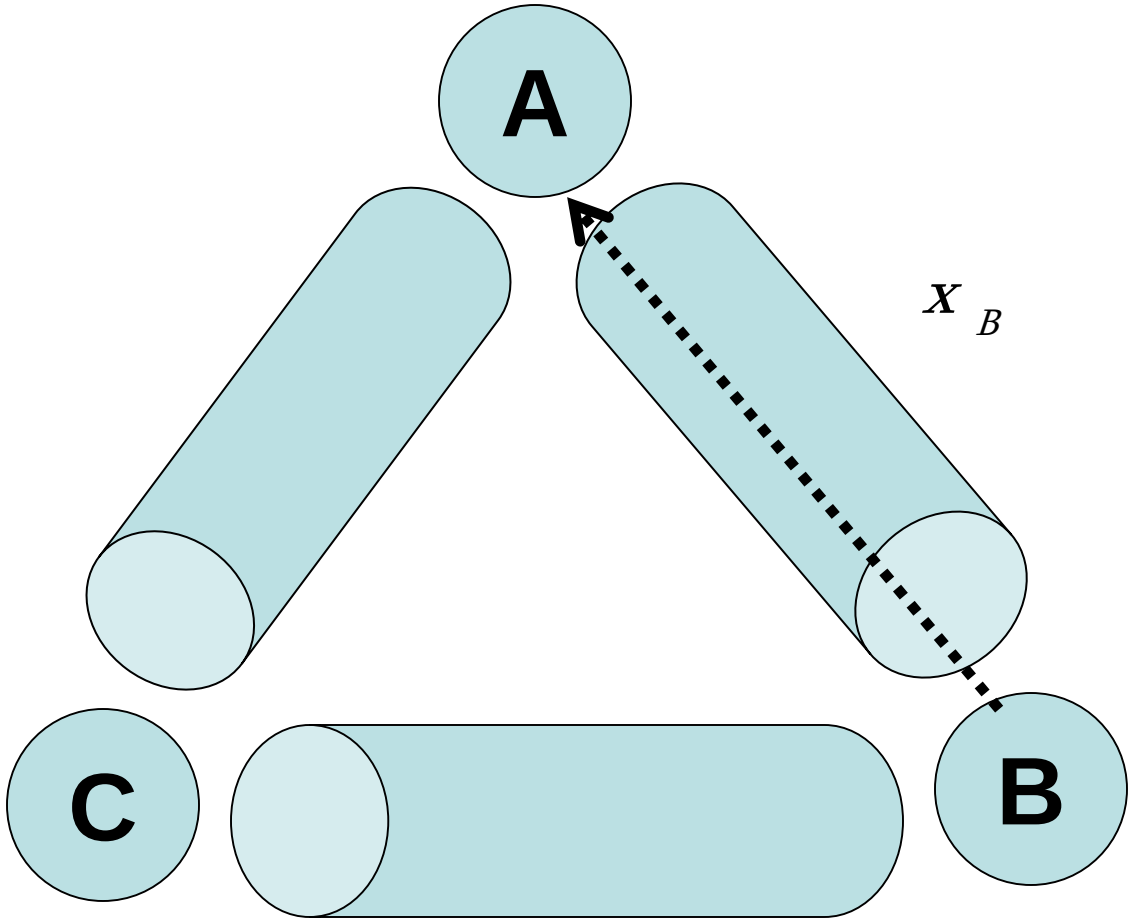

Uncoupled vs. Coupled CC

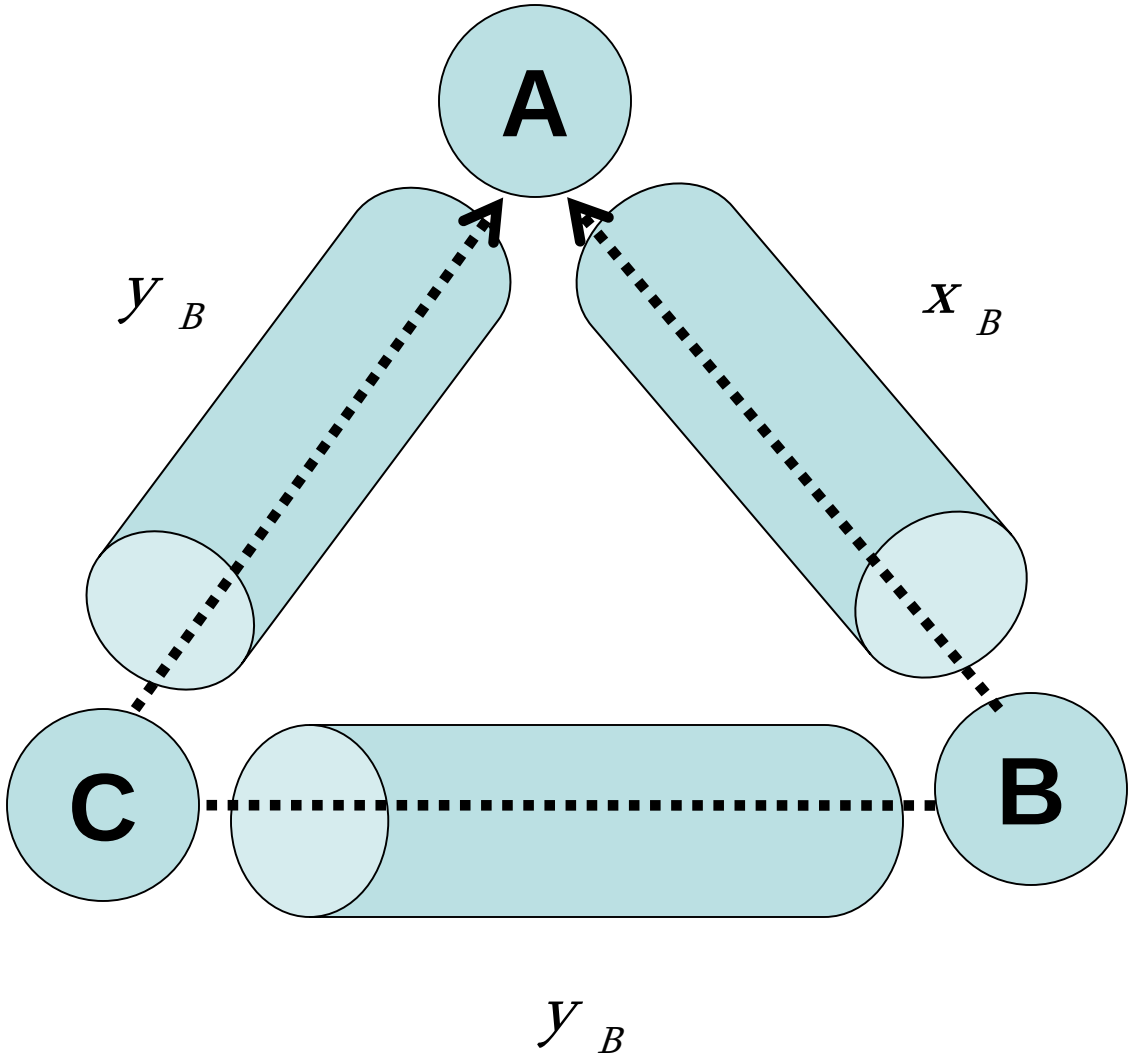

Uncoupled vs. Coupled CC

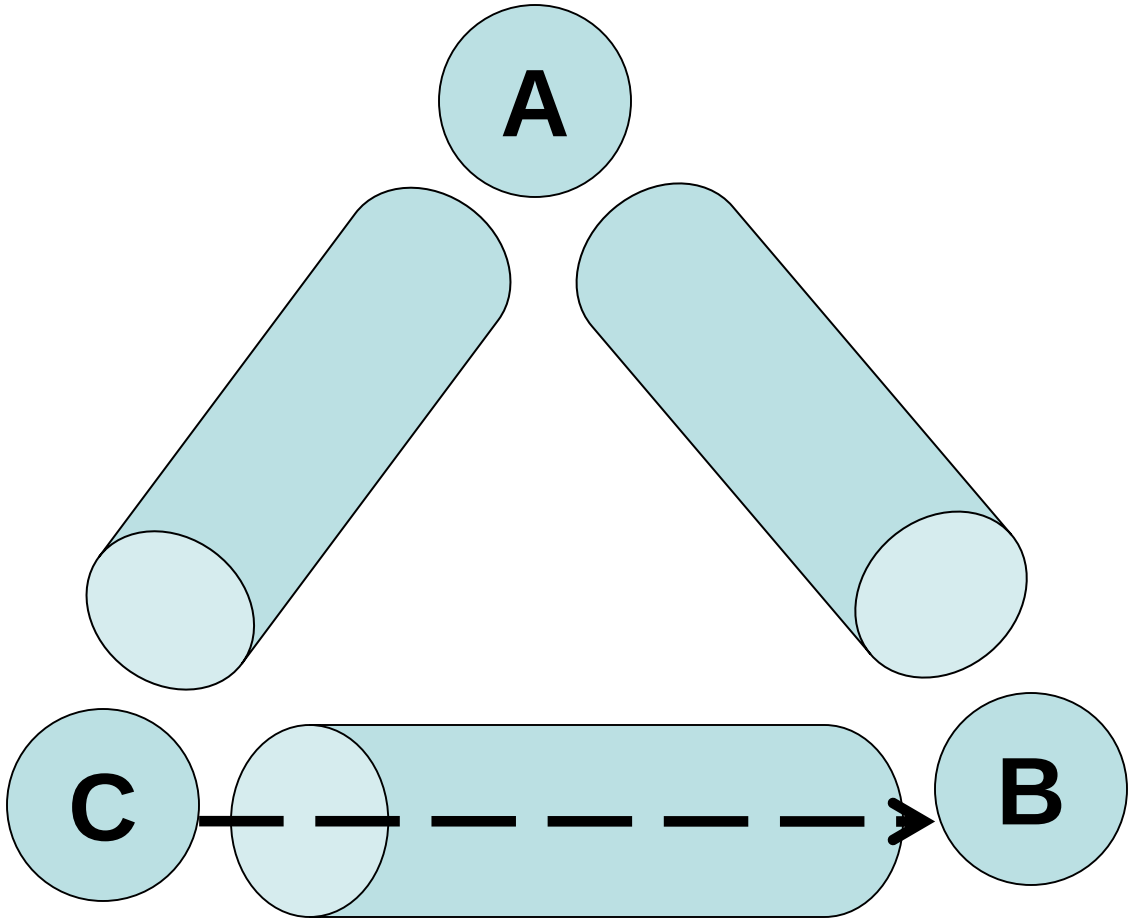

x_C

Uncoupled vs. Coupled CC

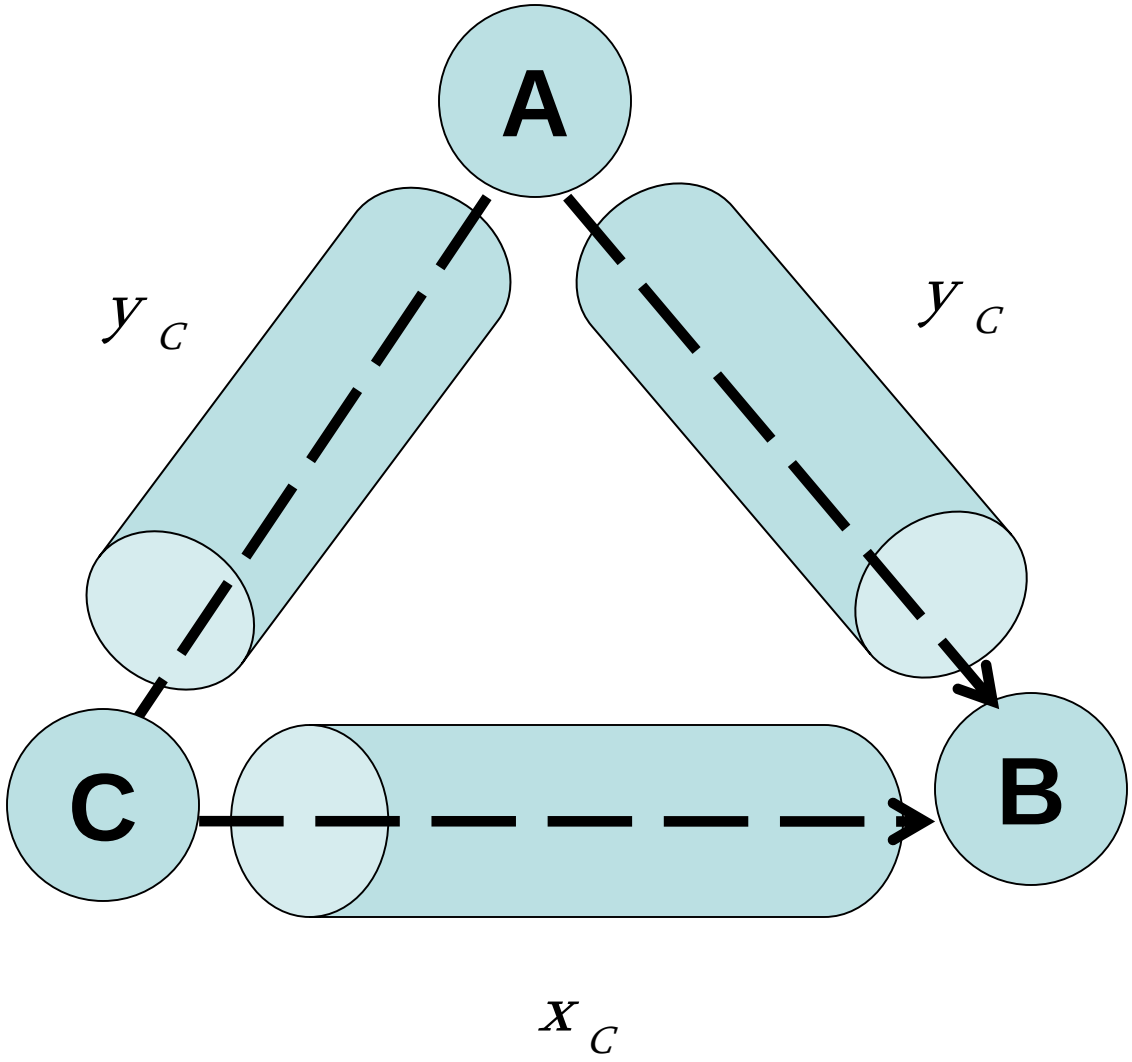

Uncoupled vs. Coupled CC

- **Uncoupled CC**

$$\max\left[U(x)+U(y)\right]$$

$$s.t.\ \ x+2\cdot y\le C$$

- Total throughput is not maximized:

$$x+y=\frac{3+2\sqrt{2}}{4+3\sqrt{2}}\cdot C\ <\ C$$

- **Coupled CC**

$$\max\left[U(x+y)\right]$$

$$s.t.\ \ x+2\cdot y\le C$$

- Total throughput is **maximized**

$$x+y=C$$

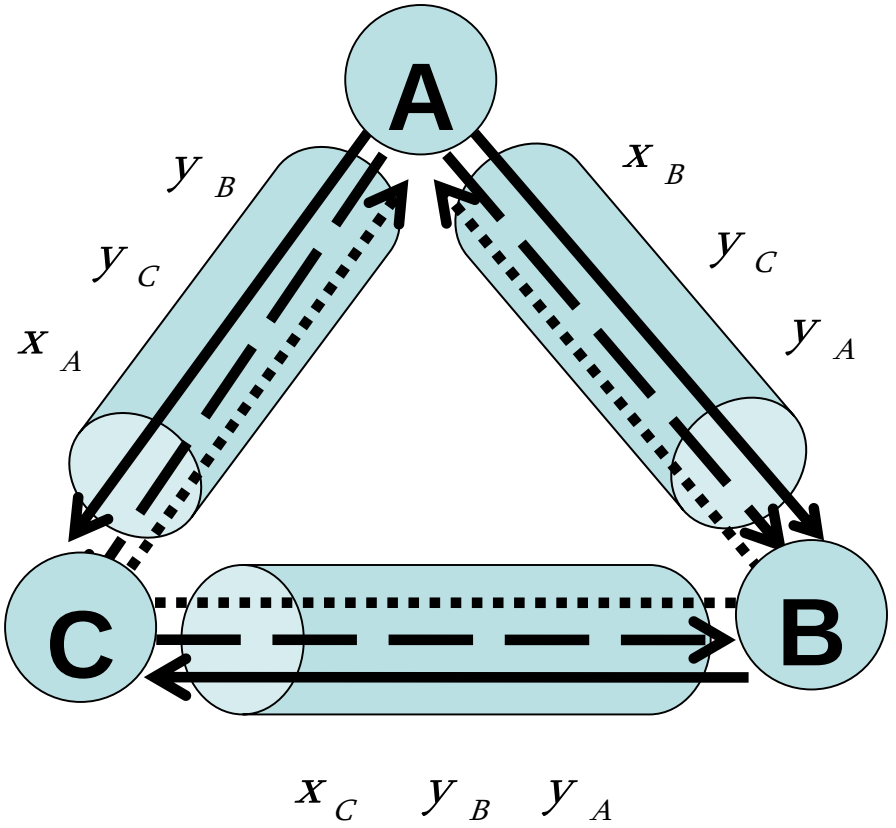

Incremental Deployability

- The deployment of a new technology is encouraged when related technologies already exist.
- **Applications:** backward-compatible extension of standard TCP. Offers an unmodified *sockets* API → not need to modify / recompiled applications
- **Network:** each MPTCP flow → like a single standard TCP connection with some new option headers. The connection starts as a normal TCP.

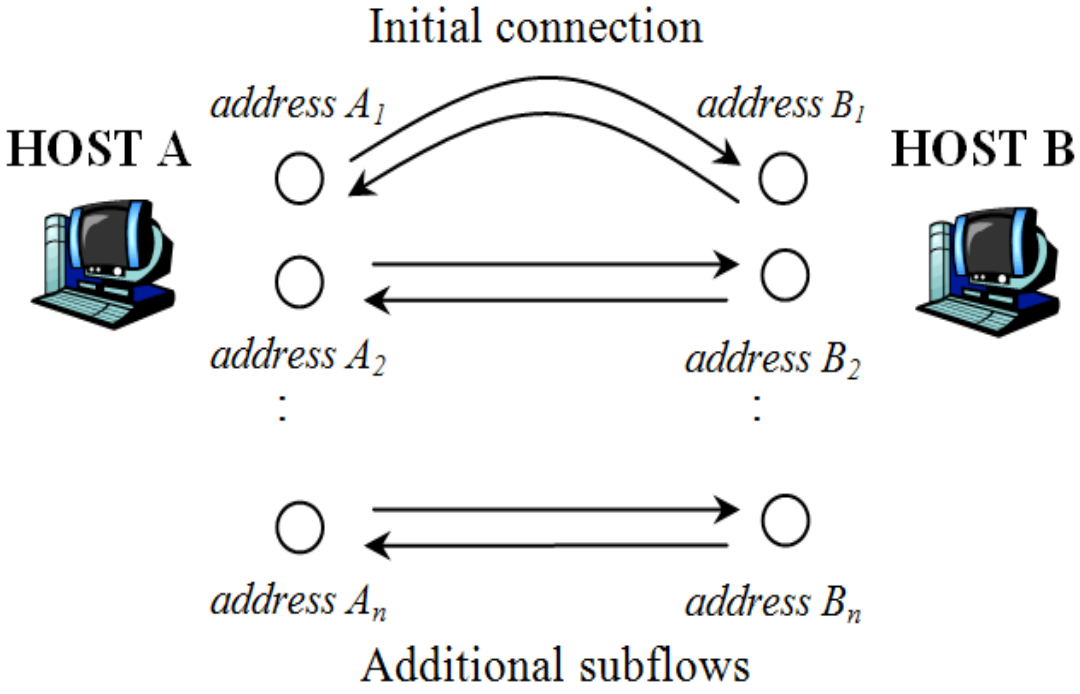

Good Technical Design

- Designing a protocol that follows “good principles” enhance deployment and interoperability.
- **Design for Tussle:** multiple stakeholders with conflict interests interact
  - Tussle-aware protocol designs have better chances at deployment in the long-term

❑ ***Resource Pooling***

- *when resources in a network can be pooled, effectiveness of the network will be improved.*
  - *pooling* → a set of resources appear as a single resource of aggregate capacity.
  - ➢ MPTCP → rp mechanism, sends data along multiple paths, uses ccc algorithm, allows the traffic load to be relocated to /spread over several paths.

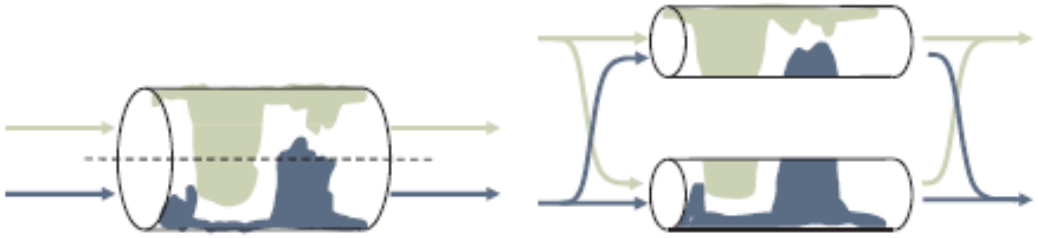

Good Technical Design

Information Exposure

- *sufficient information about resource usage should be exposed to support an efficient allocation*
  - MPTCP monitors the congestion signals on each individual subflow, to respond appropriately to resource usage and congestion by shifting load between the subflows.

Fuzzy Ends

- *end points allowed to delegate functions to the network*
  - MPTCP → end-to-end, but
  - proposed architecture → sufficiently extensible to allow the development of MPTCP proxies, placed within the network without the need of endpoints to be multi-homed themselves.

Deployment Process

- **Key Stakeholders**
  - OS vendors → *implement MPTCP in OS for use on end systems*
  - End users (i.e., individual users, service providers, CDNs) → *own end systems*
  - ISPs → *provide connectivity for multi-homing*

- **Fundamental requirements for MPTCP Deployment**
  - ***Availability of OS Implementation***
  - ***Installation of MPTCP-capable OS to end systems***
  - ***Multi-homing***

  - **Key role of other end points and network externalities**

Deployment Process

- ***Availability of OS Implementation***
  - Changes required only to the TCP/IP stack of end systems

  → *an OS update that adds MPTCP support needs to be available*

  ❖ Key Stakeholder: OS vendors
- Motives for implementing MPTCP in OSs
  - **Pressure from end-users** →
    - *if they are MPTCP aware / demand problem solution that MPTCP alleviates*
  - **Pressure from (high) application developers** → *their products could be enhanced*
  - **Own business interest** → *direct business benefits (NOKIA – Ovi)*
  - **Competitive environment** → *"leader" role of an open source OS – incentives for commercial OS vendors*

❑ Actual usage – **enabled by default** in shipping configuration

Deployment Process

- ***Installation of MPTCP-Capable OS to End Systems***
  - ❖ Key stakeholder: <u>end-users</u> → ultimate control over their devices
  - conscious decision to deploy MPTCP
    - **content providers** → increase QoS
    - **"heavy users"** → large traffic volumes
  - get the MPTCP unbeknownst to them
    - **domestic users** → purchase a new device / automatic OS updates
- ISPs may foster MPTCP use → *providing an MPTCP proxy service*
  - intercepts standard TCP traffic generated by end systems and translates it to MPTCP.

Deployment Process

- ***Multihoming***
  - Key stakeholders: end-users → *acquire additional Internet access connections*

    ISPs → *business interests (more access links, MPTCP proxies)*
  - User's desire for multihoming
    - Run MPTCP, but most probably:
      - need for ubiquitous access for a mobile user
      - back-up connections for content providers.
  - Many end-users may already have multi-homing capability available (enterprises, academic networks) → *no required hardware updates*
  - ISPs have monetary motives to improve support for it:
    - *possibility to sell more access connections*
    - *MPTCP can also help ISPs to balance the load in their networks*

Deployment Process

- ***Other End-points and Network Externalities***
  - large number of users that adopt MPTCP
  - ✓ the probability of a successful MPTCP connection establishment is increased
  - ➢ a **client** and a **server**
    - i.e. Google deploy MPTCP → *update a significant fraction of its servers at the same time (externalities in jumps)*
  - ➢ between **two clients**
    - Specific interest for peers he mostly connects to → *if he often accesses a specific service, it is important to him that the particular service is MPTCP-capable*

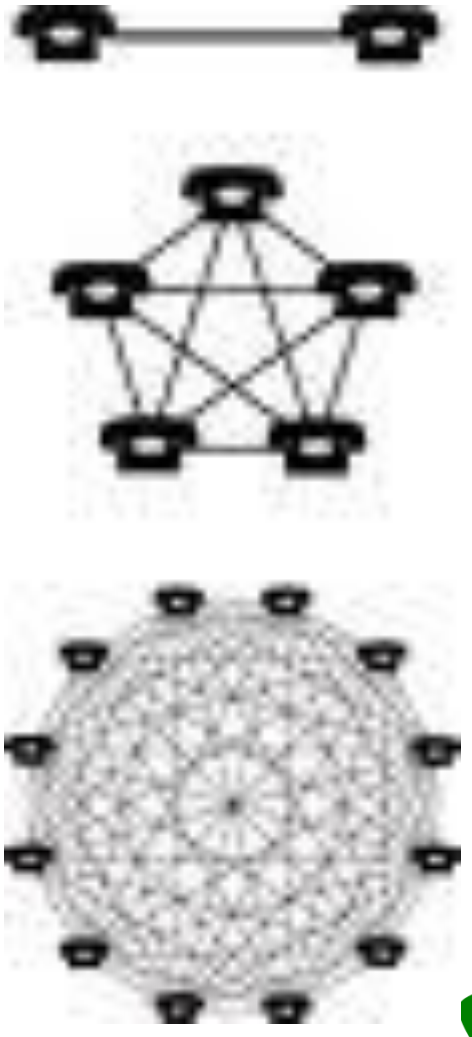

Scenarios Supporting Adoption

– ***Both Ends in one Hand***

- Multi-homed devices and content/application servers are under the control of one stakeholder, i.e. companies that provide a mobile device for their employees to use company applications remotely over WLAN/3G could significantly benefit from MPTCP.
- An end-user accessing content using access from a provider which controls both end-user devices and content servers, (Ovi / iTunes service, both delivering devices and services/content).
- Consumers probably MPTCP-unaware → *perhaps opportunistic adopters of MPTCP when it is implemented by device manufacturers*.

✓ The deployment in the client devices (OS vendor's enabling MPTCP by default)
→ *key driver to the adoption on the client-side if the end-user is multi-homed already*.

Scenarios Supporting Adoption

– ***Lobbying***

- Lobbying towards OS vendors who have to implement the new protocol in their network stacks.
- Key players (i.e. Microsoft – 85% of PCs OS worldwide)
  - → *convinced of the merits of MPTCP*
- Organizations that represent end-users with a vital interest for MPTCP
  - → *take on the initiative*

– ***End-user decision***

- "Heavy" users and operators of large content sites
  - → lots of data – direct interest in the increased resilience and throughput.
- Once the protocol has been made available by OS vendors → decision for adoption depending on:
  - *the involved cost for OS upgrade installation*
  - *additional cost physical access lines for multi-homing*
  - *availability of MPTCP enabled clients or peers.*

Scenarios Supporting Adoption

– ***ISP Support***

- Multi-homing → considerable barrier to MPTCP adoption due to costs for additional connections.
- Offer cheaper access bundles → incentives for ISPs (*lock-in, improved traffic engineering control*)
- Virtual Multipath Operators could offer such bundles by buying / leasing access lines, potentially of different kinds, from other ISPs → *increase in competition – driving factor for ISPs to offer their own price-reduced bundles*

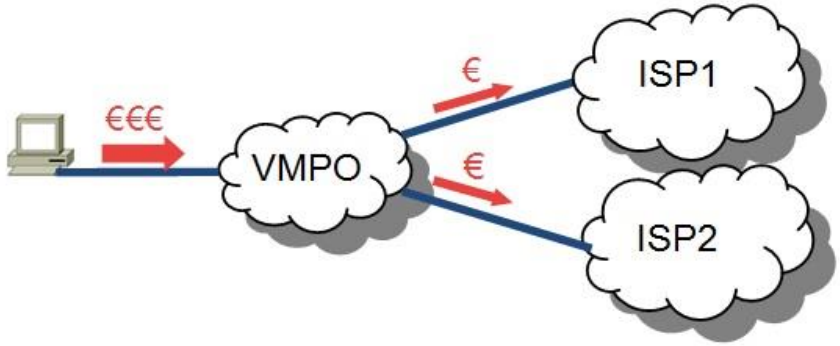

- MPTCP-enabled access from ISPs → a value-added service by providing a MPTCP proxy service to end-users → Cost-effective solution – no requirement for extra access link /MPTCP-enabled OS.
  - increased resilience and throughput, but limited to one access connection / not realize the full potential benefits of MPTCP

Conclusions and Future Work

- The adoption of transport protocols differ from diffusion of end user-centered innovations
- Performance, reliability, flexibility beneficial, **but** not the main drivers for adoption.
- Role of end users in not of primary importance → *not necessarily conscious adoption decision*
- Mainly in hands of operating system vendors → *deployment decision to enable by default*.
- The deployment of MPTCP-enabled OS will take different channels:
  - roll out on new devices delivered with new operating systems
  - automatic software updates to the deployed base (often without awareness of end users)
  - intentional installation by operators of large sites (e.g., content providers)

- **Future Work**
  - Compare the MPTCP-like solutions in other layers → is transport the proper layer?
  - How efficient is MPTCP for short flows?
  - Applying different pricing schemes

THANK YOU FOR YOUR ATTENTION !!

Alexandros Kostopoulos (alexkosto@aueb.gr)

Athens University of Economics and Business

Network Economics and Services Group (http://nes.aueb.gr)