

Use cases for MPTCP

Costin Raiciu

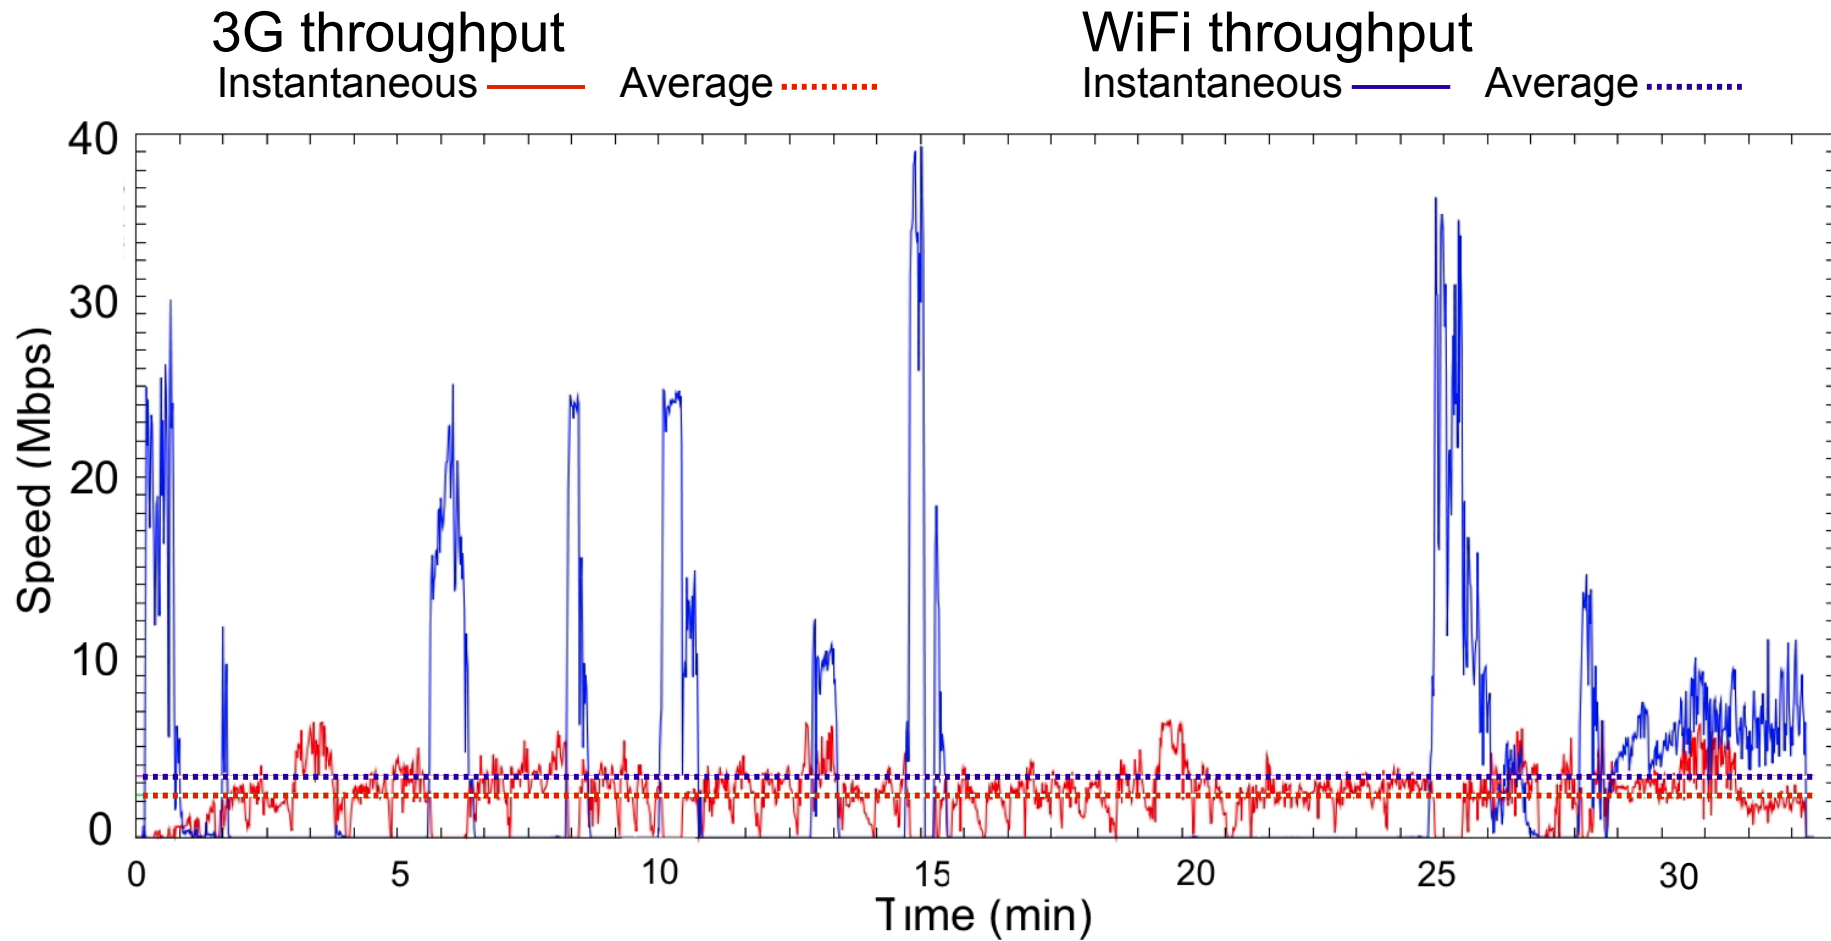
University Politehnica of Bucharest

IETF 87, Berlin

Use-case 1 Cellular to Wifi offload

- Infrastructure:
 - RDS Wifi and 3G deployment in Bucharest underground
- Testing equipment:
 - Laptop running MPTCP
 - Proxy: Squid on Linux(MPTCP) running @cs.pub.ro
- Test: download the Ubuntu distribution via our proxy

Real Mobility in the Bucharest Underground



Use-case 2 Wifi Mobility

- Wifi deployments are ubiquitous
- Important to offload data from cellular nets
- Wifi mobility sucks – lots of downtime
 - AP ranges are very small (many handoffs needed)
 - ***Difficult to know the best AP at any given instant without trying them out***
 - APs belong to different entities, difficult to coordinate

What if we had MPTCP on all clients?


- Multiple NICs
 - Every client has two WiFi NICs (not the case today)
 - Each associates to a different AP
 - Better throughput but worse energy consumption
- Single NIC
 - Typically, APs will be on different channels
 - Client switches dynamically between channels
 - Switch takes 3-10ms depending on NIC
 - Power save mode is used to make APs buffer packets

A Simple Mobile Scenario

Moving from one AP to another

AP1


AP2





Next steps


- How many APs should we connect to?
 - More APs = better robustness but more switches and increased RTT
- Periodic scans to find new APs – don't wait for downtime
- Many heuristics (ongoing work)

Use-case 3 Changing datacenter networks









To utilize the core better we'll use

- Additional server network ports




(datacenter servers nowadays ship with dual or quad port NICs)

- Multipath TCP



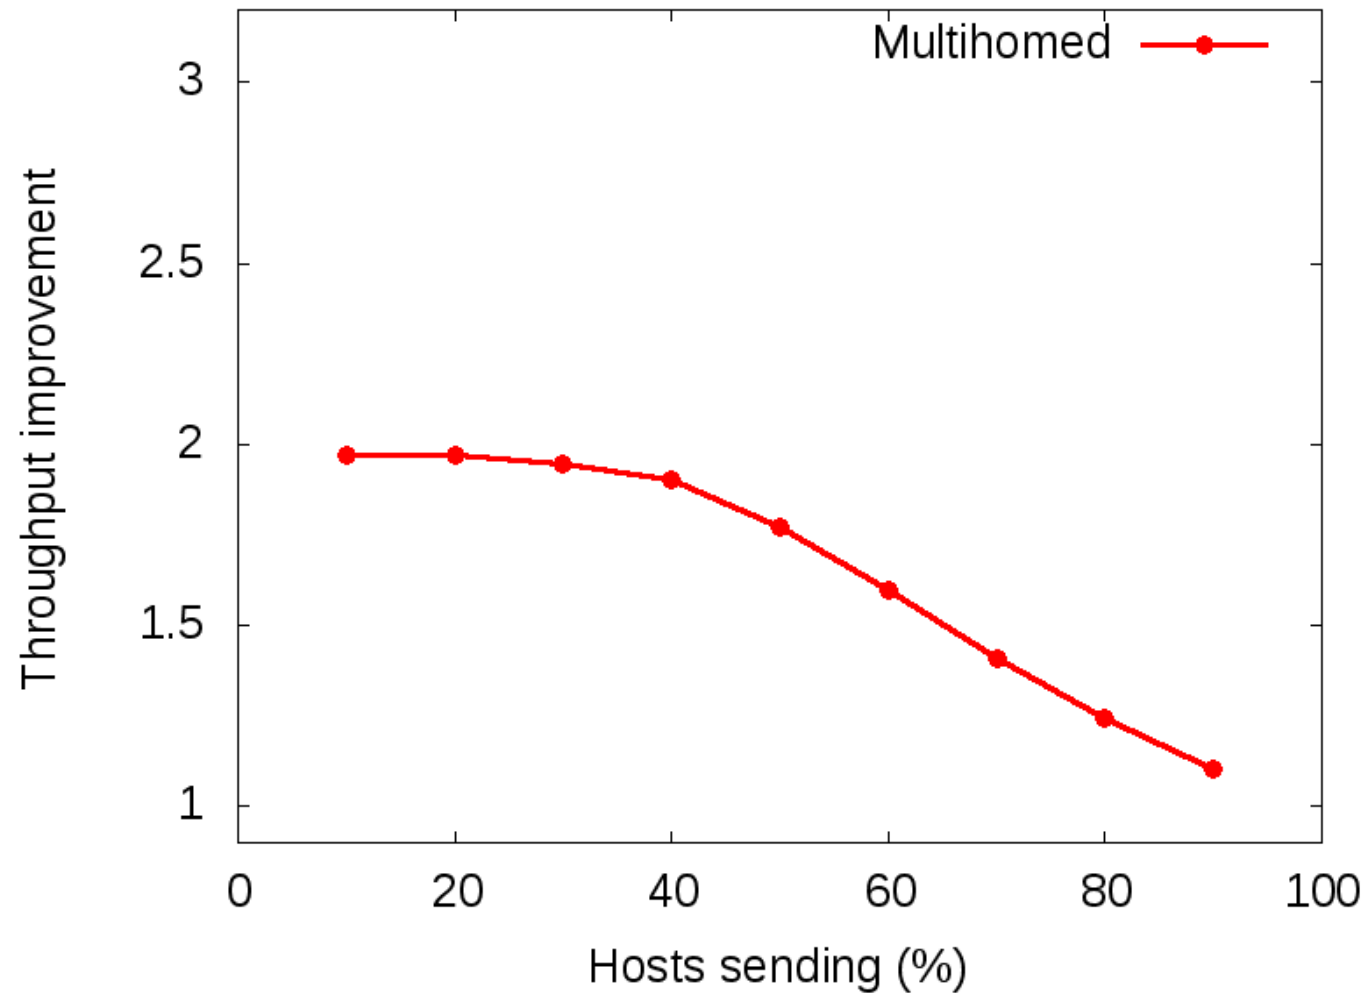
GRIN1




Permutation Traffic Matrix

- Each active server picks a single destination at random and tries to send data at the fastest possible rate
- Each destination may only receive data from a single sender
- Different numbers of active servers for different experiments

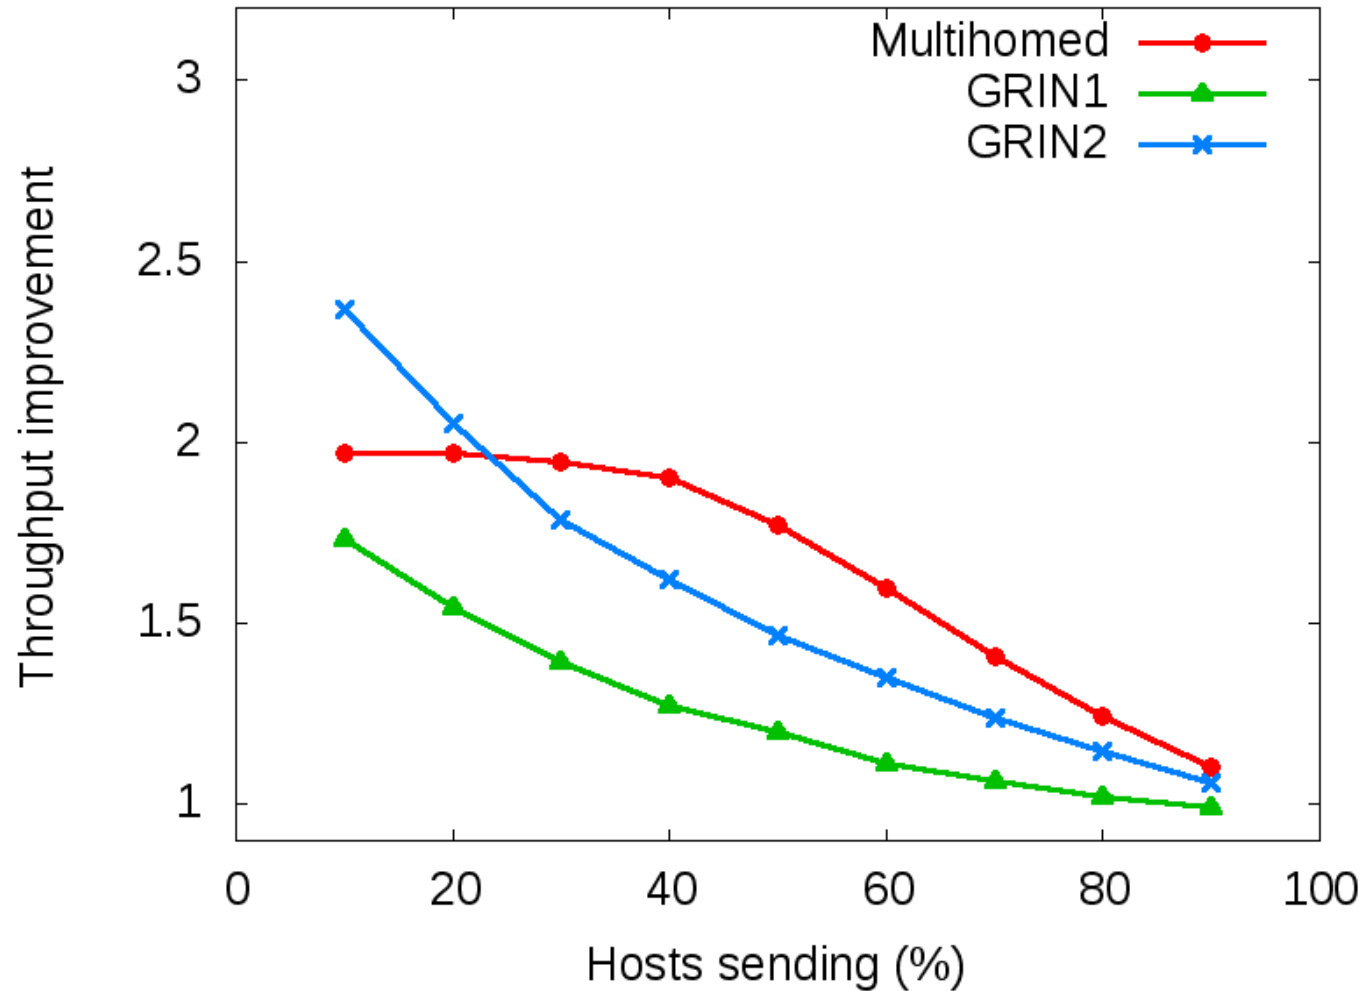
Permutation Traffic Matrix




Permutation Traffic Matrix



Permutation Traffic Matrix



Permutation Traffic Matrix




Use-case 4 Wide-area VM Migration

- Moving VMs across datacenters is useful
 - Unsolved problem: making sure TCP connections survive the migration
- MPTCP's connection identifier enables us to move the endpoint of a connection by just adding a new subflow with the new IP address.
- Xen + VM running Linux with MPTCP
 - Minimal changes needed to hypervisor
 - Works like a charm
- Live connection migration is also possible
 - Except moving processes is tougher...


Use-case 5

Routing a connection via a middlebox



Use-case 5

Routing a connection via a middlebox



MPTCP deployment enables many simplifications/optimizations at other layers in the stack

- L2 mobility
- Better datacenter network topologies
- Migration
- Segment routing
- Responsive TE ?
- Less stress on BGP ?