Simplifying Manageability, Scalability and Host Mobility in Large-Scale Enterprise Networks using VEIL-click

Sourabh Jain and Zhi-Li Zhang
University of Minnesota-Twin Cities
{sourj, zhzhang}@cs.umn.edu

Current Trends and Future Networks

- ☑ Large number of mobile users and systems
- ☑ Large number of smart appliances
- ☑ High bandwidth core and edges
- ☑ Heterogeneous technologies
- ☑ More and more Virtual appliances:
  - ⊙ Virtual Servers
  - ⊙ Virtual PCs

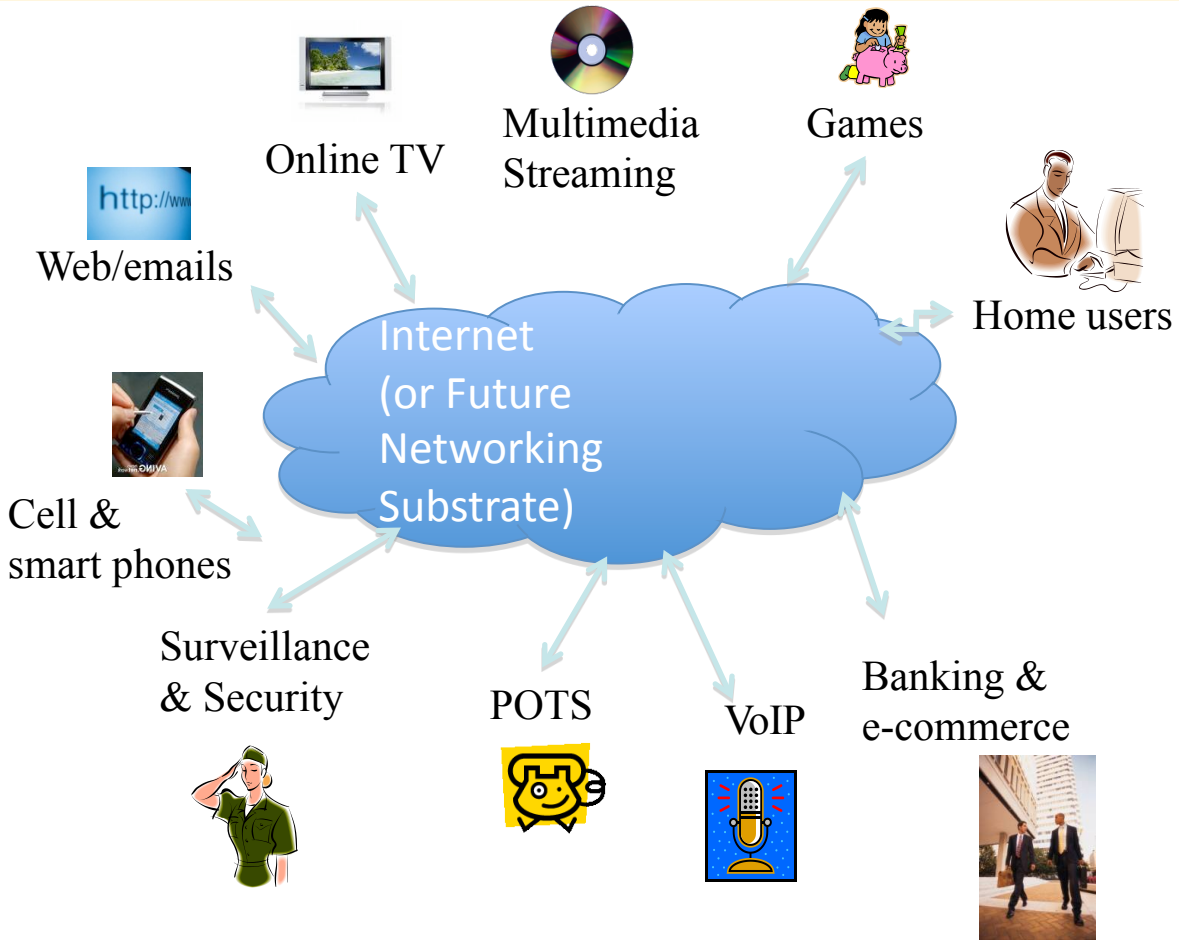

Challenges posed by These Trends

- **Scalability:** capability to connect tens of thousands or more users and devices
  - routing table size, constrained by router memory, lookup speed
- **Mobility:** hosts are more mobile, "virtual servers" are also mobile
  - need to separate location ("addressing") and identity ("naming")
- **Availability & Reliability:** must be resilient to failures
  - need to be "proactive" instead of reactive
  - need to localize effect of failures
- **Manageability:** ease of deployment, "plug-&-play"
  - need to minimize manual configuration
  - self-configure, self-organize, while ensuring security and trust
- .....

How Existing Technologies Meet these Challenges?

Ethernet/Wireless LANs (L2)

- ☑ Pluses:
  - ⊙ plug-&-play, minimal configuration, better mobility
- ☑ Minuses:
  - ⊙ (occasional) data plane flooding, sub-optimal routing (using spanning tree), not robust to failures
  - ⊙ Not scalable to large (wide-area) networks

IPv4/IPv6 (L3)

- ☑ Pluses:
  - ⊙ better data plane scalability, more "optimal" routing, ...
- ☑ Minuses:
  - ⊙ control plane flooding, global effect of network failures
  - ⊙ poor support for mobility
  - ⊙ difficulty/complexity in "network renaming"
  - ⊙ Esp., changing addressing schemes (IPv4 -> IPv6 transition) requires modifications in routing and other network protocols

IP address Management & Mobility

IP address (re)assignment creates management overhead:

- *Careful IP configurations*
  - DHCP servers need to maintain state
  - Static assignment requires manual effort
- *Breaks the mobility*
- *Firewall re-configurations*

How to enable live migrations of "virtual machines"?

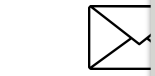

Recent Proposals

- SEATTLE [SIGCOMM'08] , VL2 [SIGCOMM'09], TRILL, LISP
  - Shortest path routing using link state routing protocol on Ethernet switches
  - ID Location separation for better mobility
  - Seattle uses DHT style lookup, VL2 uses a directory service for flooding free lookup
  - No flooding on data plane
  - However, control plane still uses flooding!
- ROFL [SIGCOMM'06], UIP [HotNets'03]
  - DHT style routing for scalability
  - Uses flat labels for mobility
  - However, these may incur significant routing stretch due to no topology awareness
- No fundamental support for advanced features such as:
  - Multipath routing
  - Fast Failure Rerouting

Overview of VEIL-Click

- Prototype implementation of recently proposed **VIRO routing** framework
  - Enables benefits of VIRO for existing Ethernet networks
    - Creates **Scalable, robust and efficient Ethernet networks**
    - Built-in support for multipath routing & fast failure rerouting
  - Virtual Ethernet ID Layer
    - Re-use MAC addresses for **topology-aware structured** *vid*s
    - *vid*s act as the location for the hosts, while IP addresses are used as persistent identifiers
    - Enables a fully backward compatible design, with no modifications to existing host-devices

VEIL Design Elements

- VEIL Switches
- VEIL Centralized Controller (VCC)
- [unchanged] Host Devices, and Ethernet/Wireless Routers.

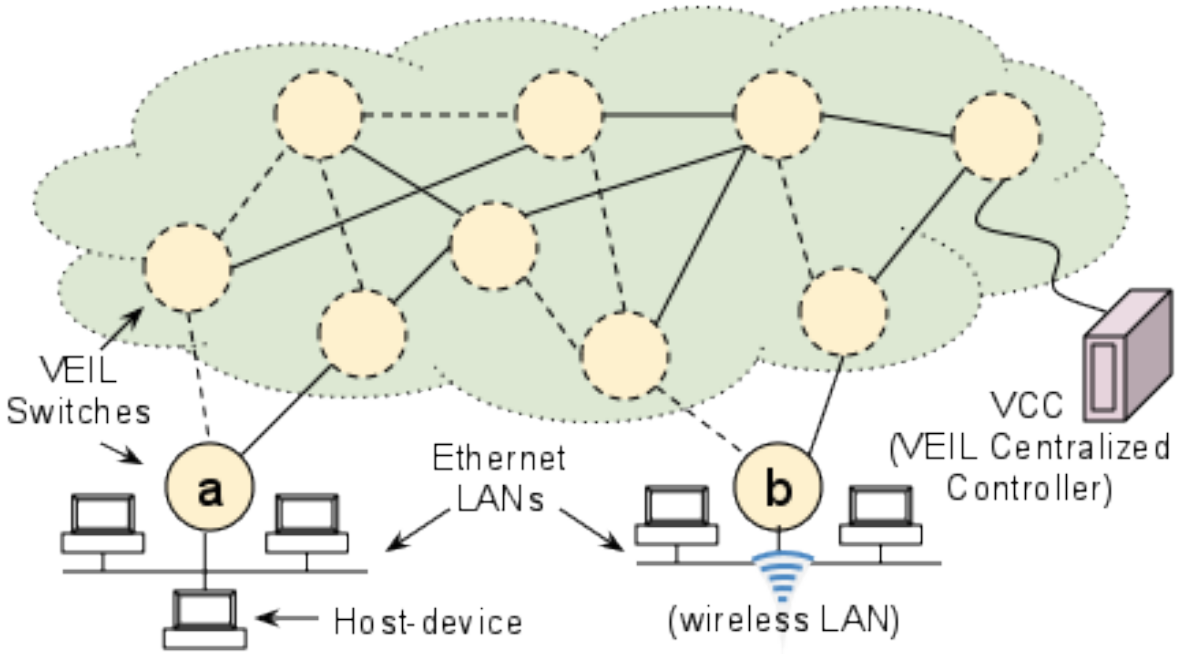

IP address range for all the hosts: 192.168.0.0/16 Mask: 255.255.0.0

VEIL: Virtual Ethernet ID Layer

- Re-use 48-bit MAC addresses for vid
  - For backward compatibility with existing Ethernet protocols and devices
- vid structure:
  - **switch vid** (32 bits)
    - assigned to switches using the vcc
    - Host-device inherit the switch vid from the switch they directly connect to
  - **host id** (16 bits)
    - assigned by "host-switches"
    - uniquely identify hosts directly connected to a switch.
  - End hosts agnostic of their vids
  - Host switch performs vid/MAC address translation
  - Backward compatible w/ Ethernet, 802.11, etc.

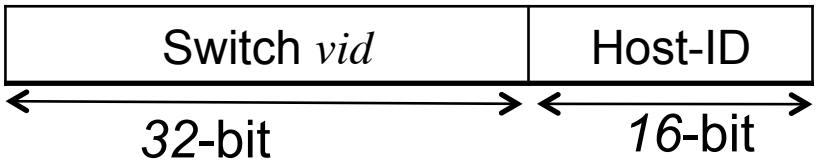

VEIL: <IP/MAC, vid> Mapping

- Host-switch:
  - a switch directly connected to the host
  - discover host MAC/IP through ARP, and assign vid to host
  - host-switch publishes IP → vid mappings at an "access-switch"
- Access-switch:
  - a switch whose *vid* is closest to hash (*IP address* of the host)

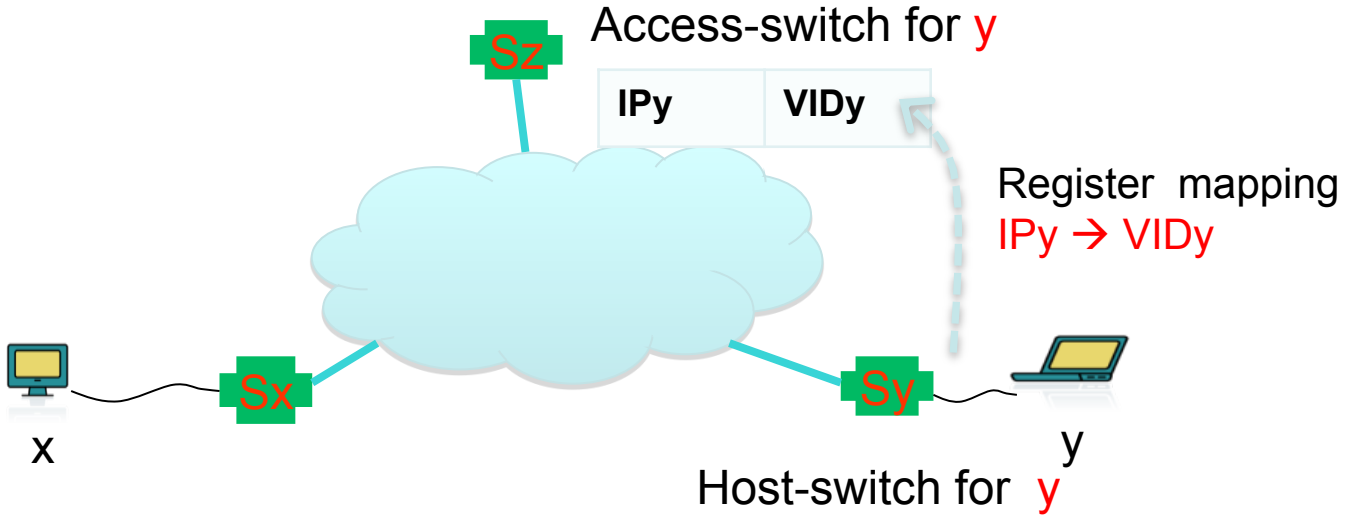

Address/vid Lookup & Data Forwarding

- ☑ Use DHT look-up for address/vid resolution with local cache
- ☑ vid to MAC address translation at last-hop

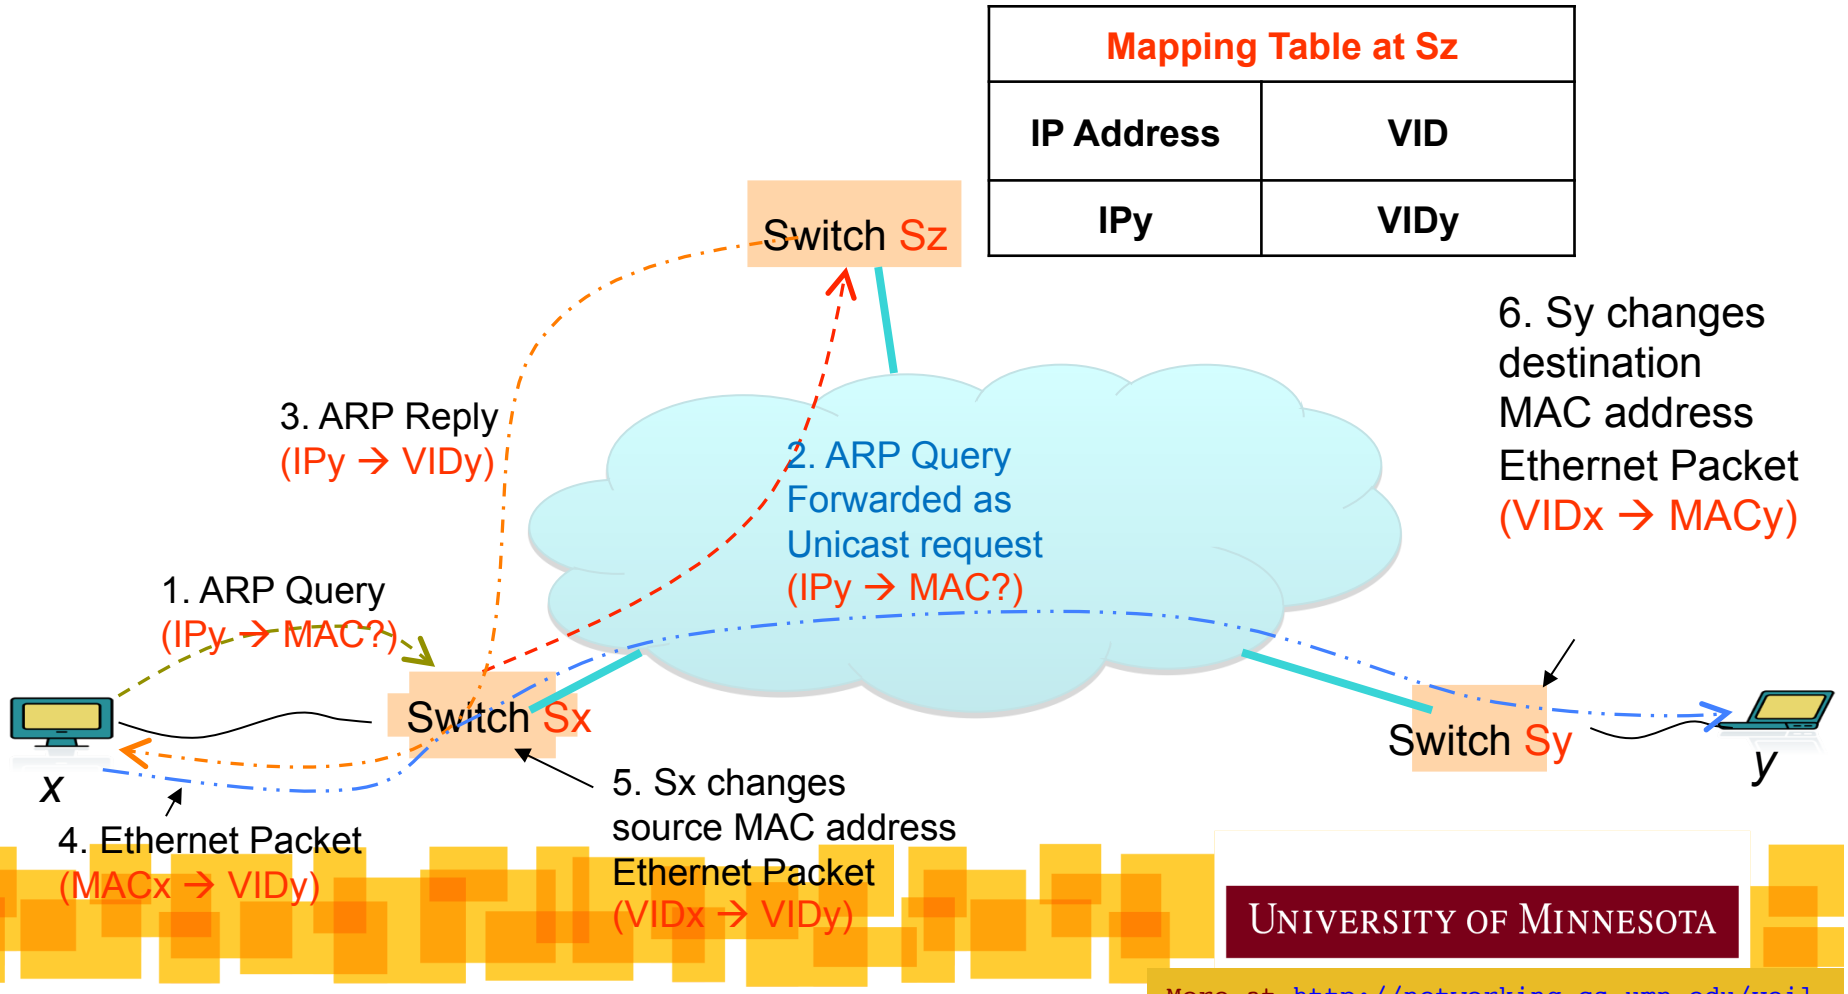

Seamless Host Mobility

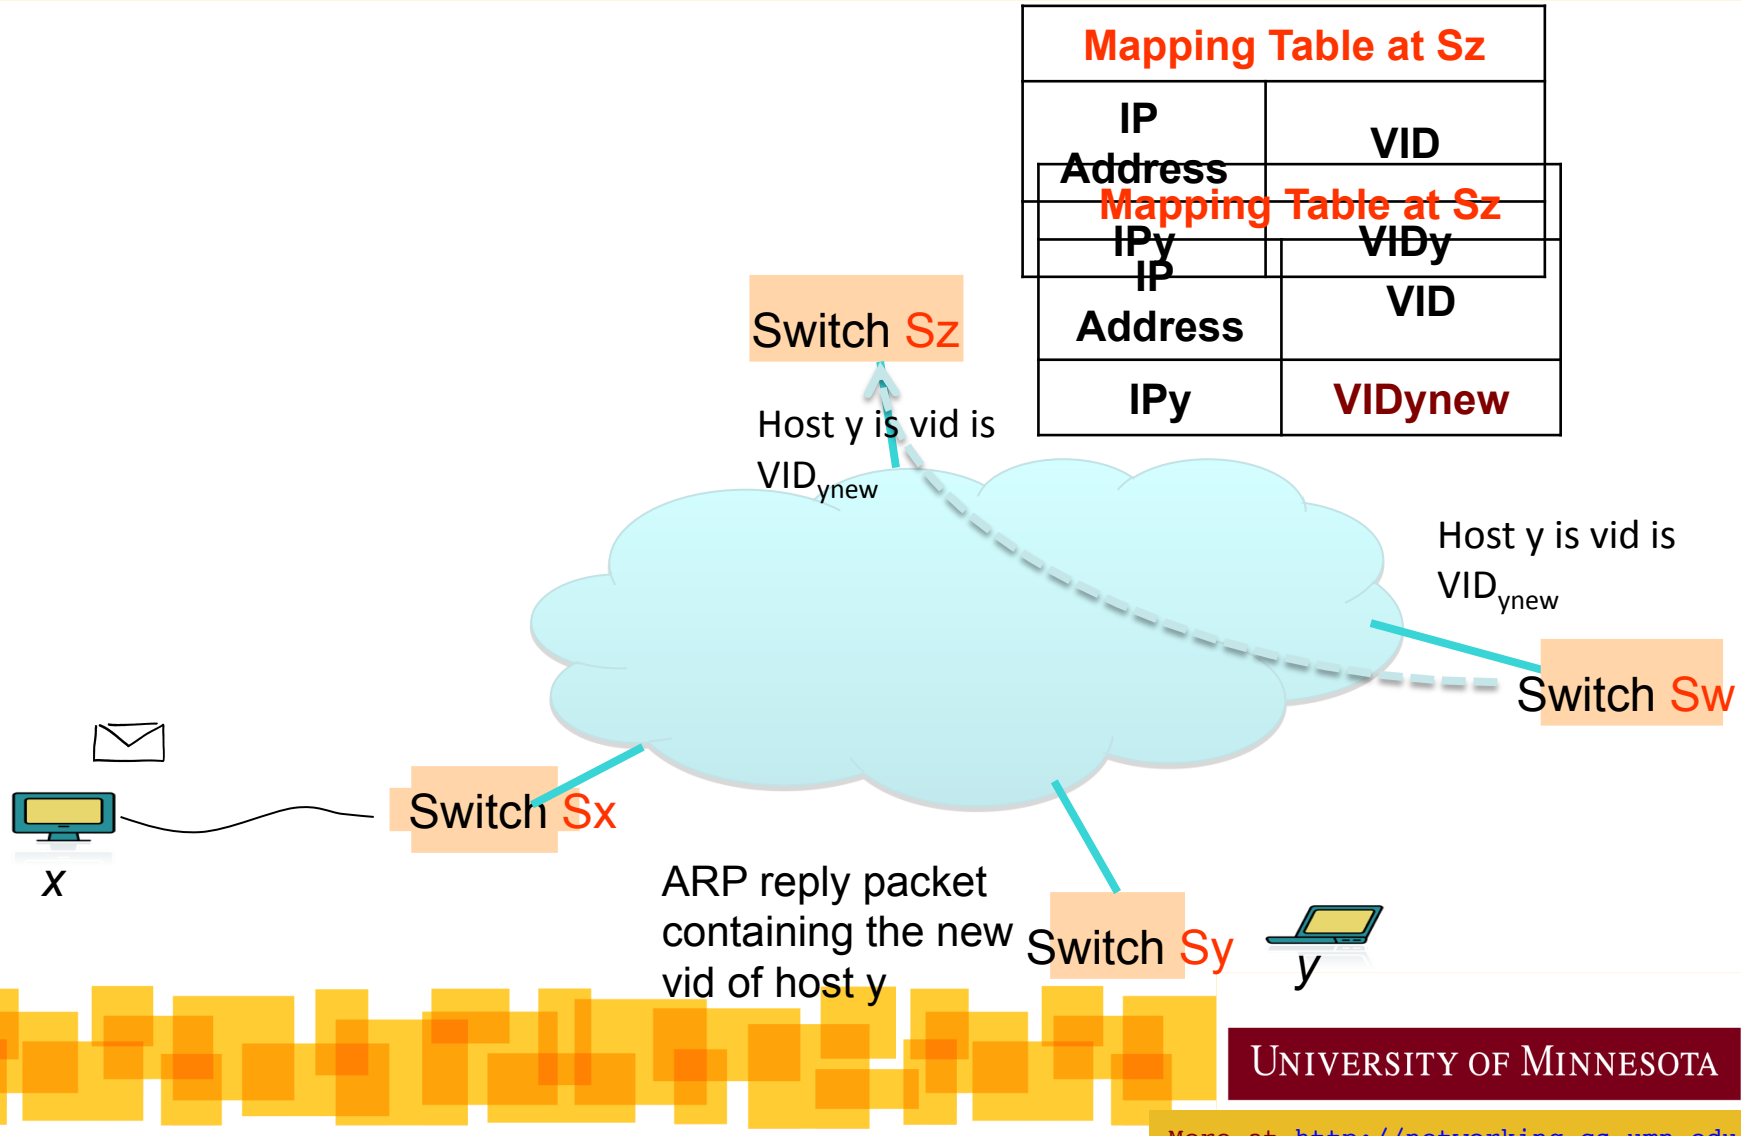

VEIL-Click: An initial prototype

- ☑ Implementation of VEIL architecture using Click Modular Router framework
- ☑ VEIL-Click enabled switch consists of:
  - ⊙ A linux machine
  - ⊙ Multiple network interfaces
  - ⊙ Click Modular Router
  - ⊙ VEIL as Click elements

VEIL-Click: Modules and Interaction

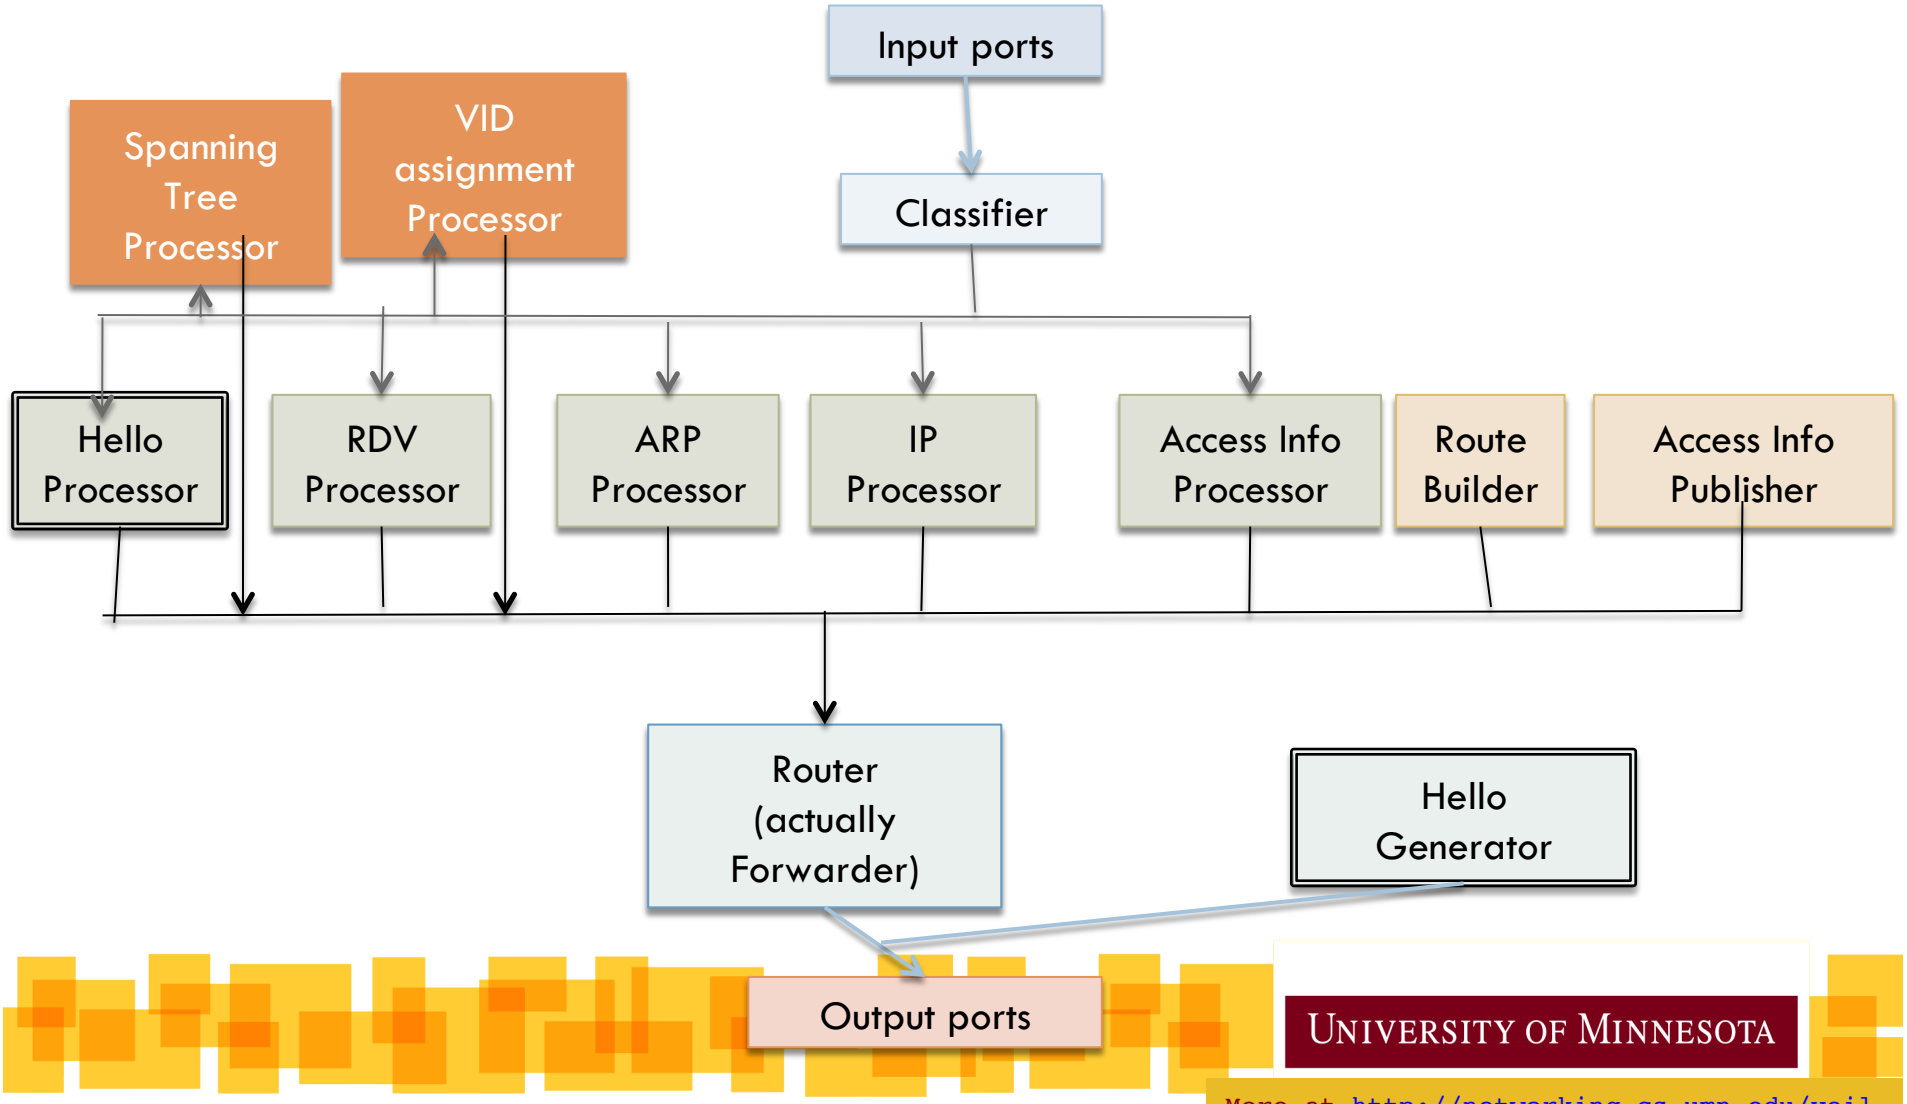

A prototype switch

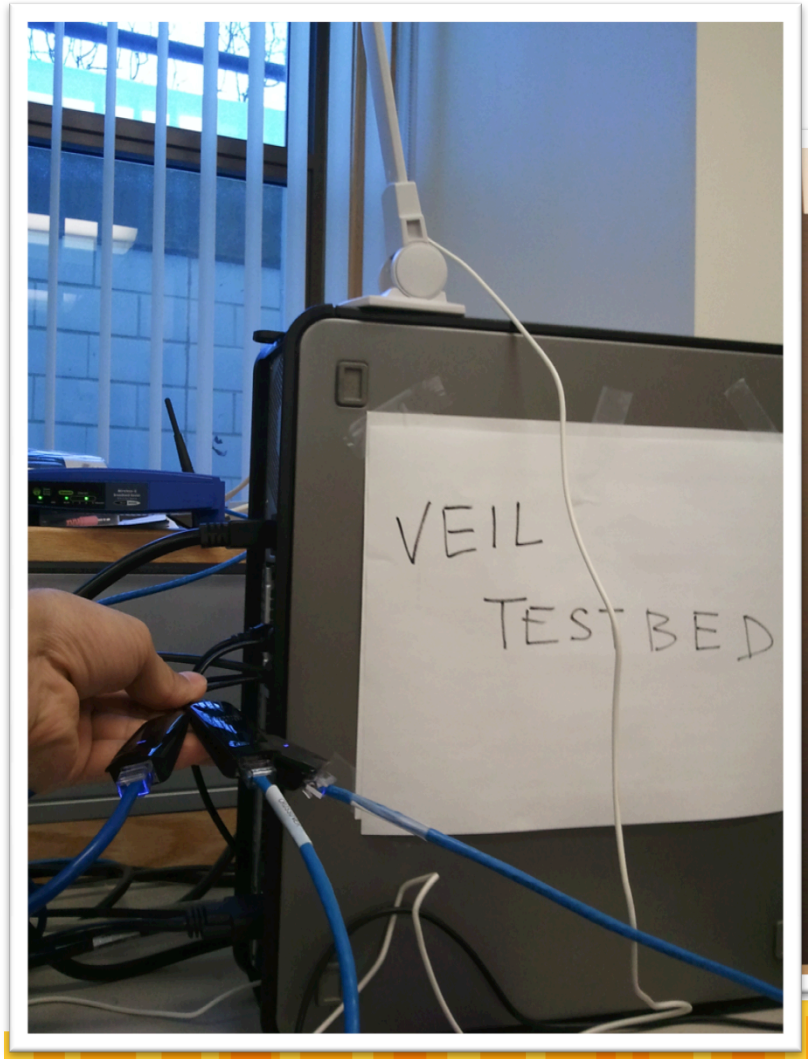

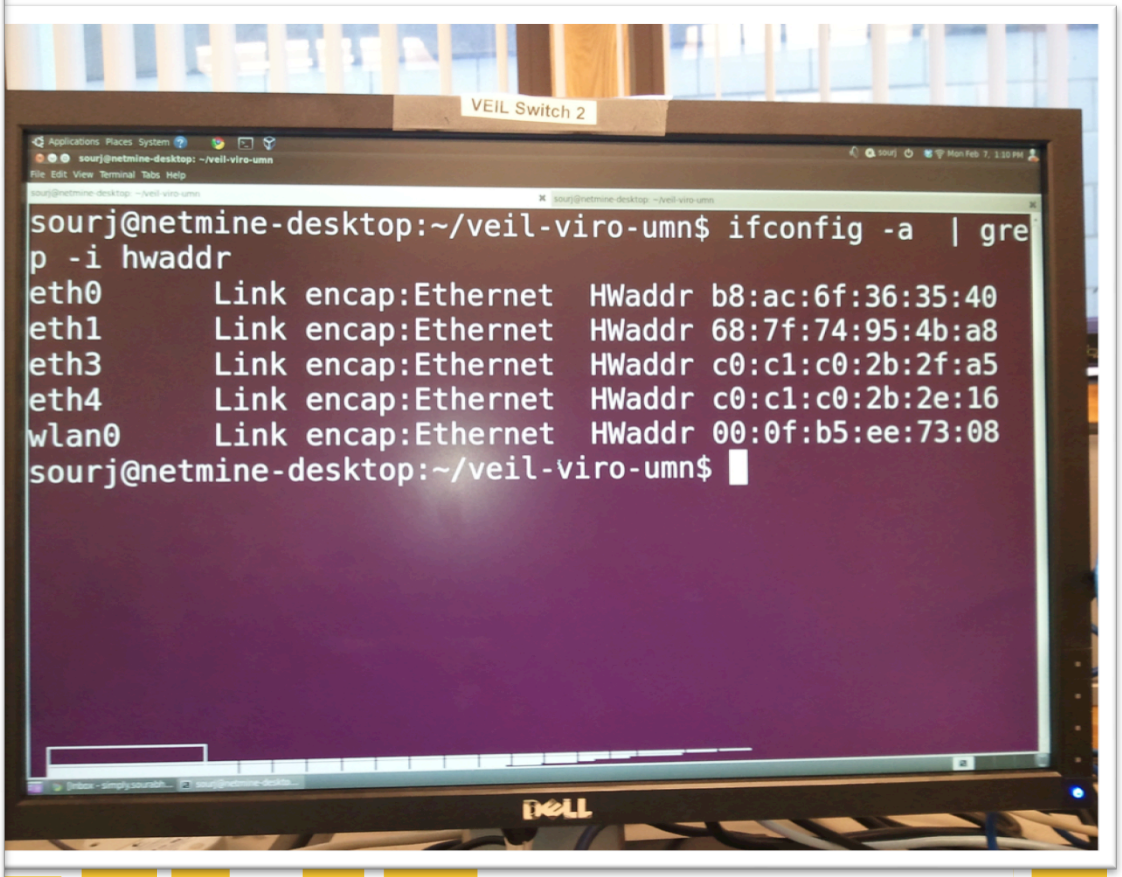

UNIVERSITY OF MINNESOTA

More at http://networking.cs.umn.edu/veil

VEIL-Click: Host Mobility Demo

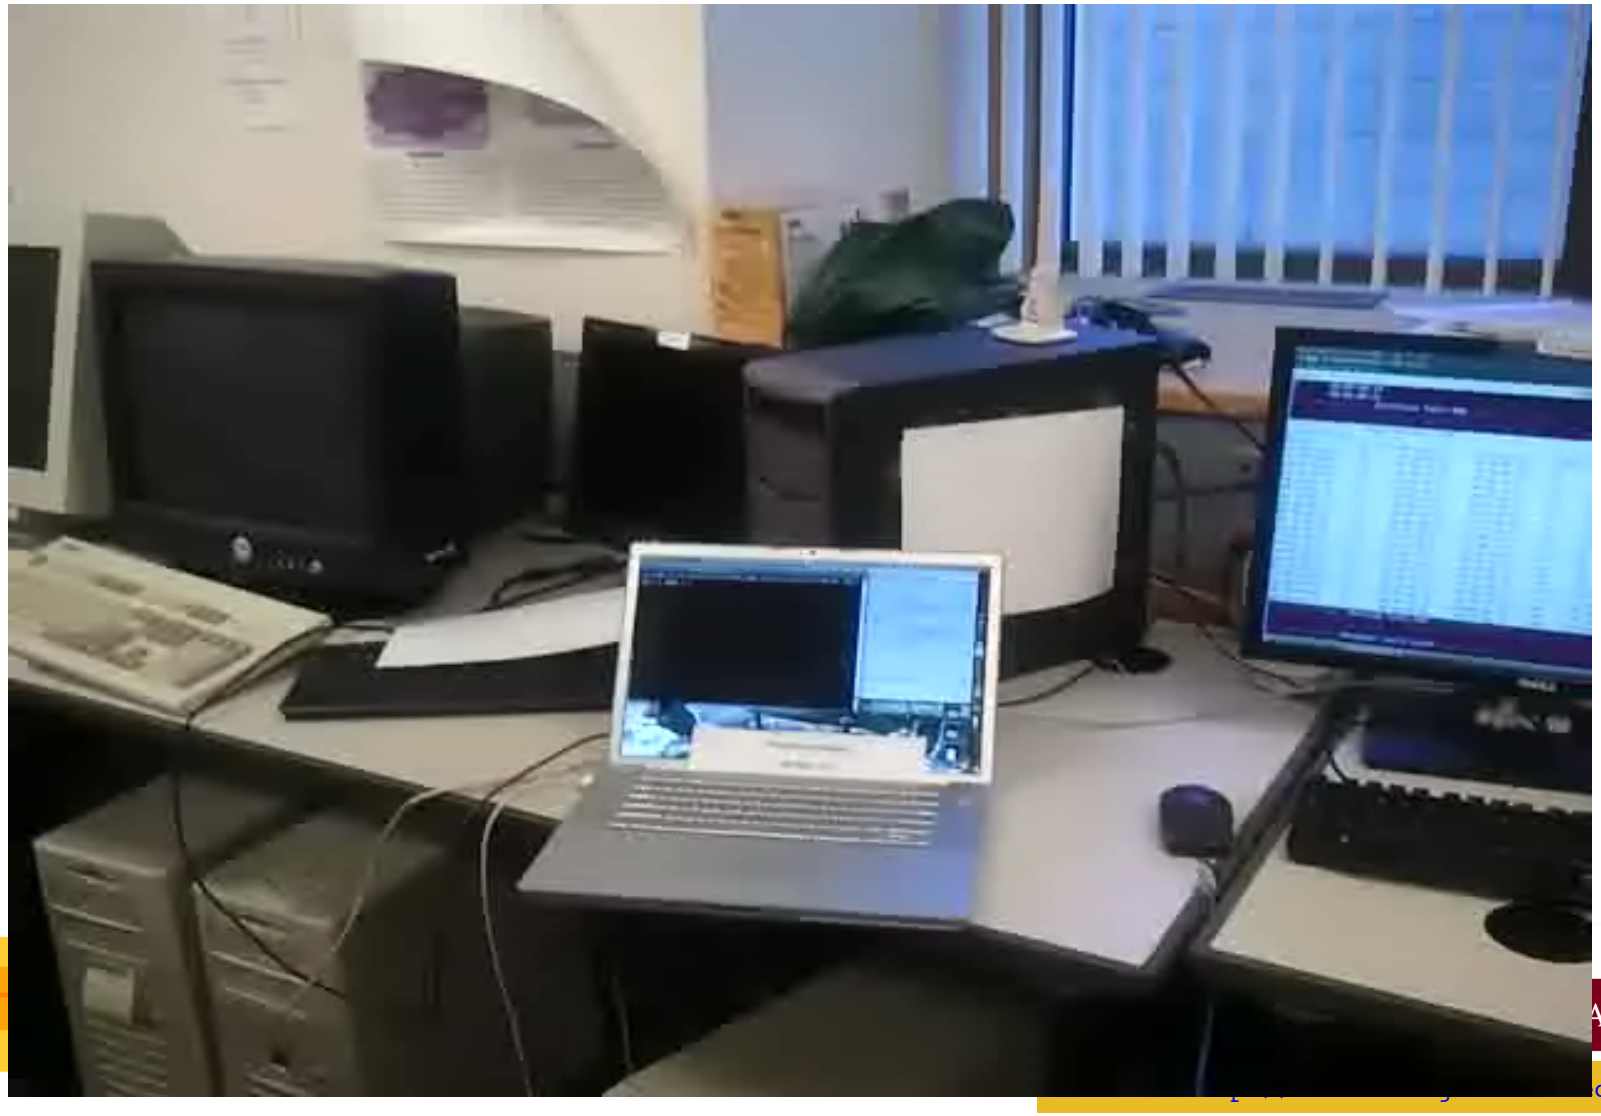

Throughput during the mobility

Throughput remains more or less stable, with minimal disruptions during the transitions from one switch to another!

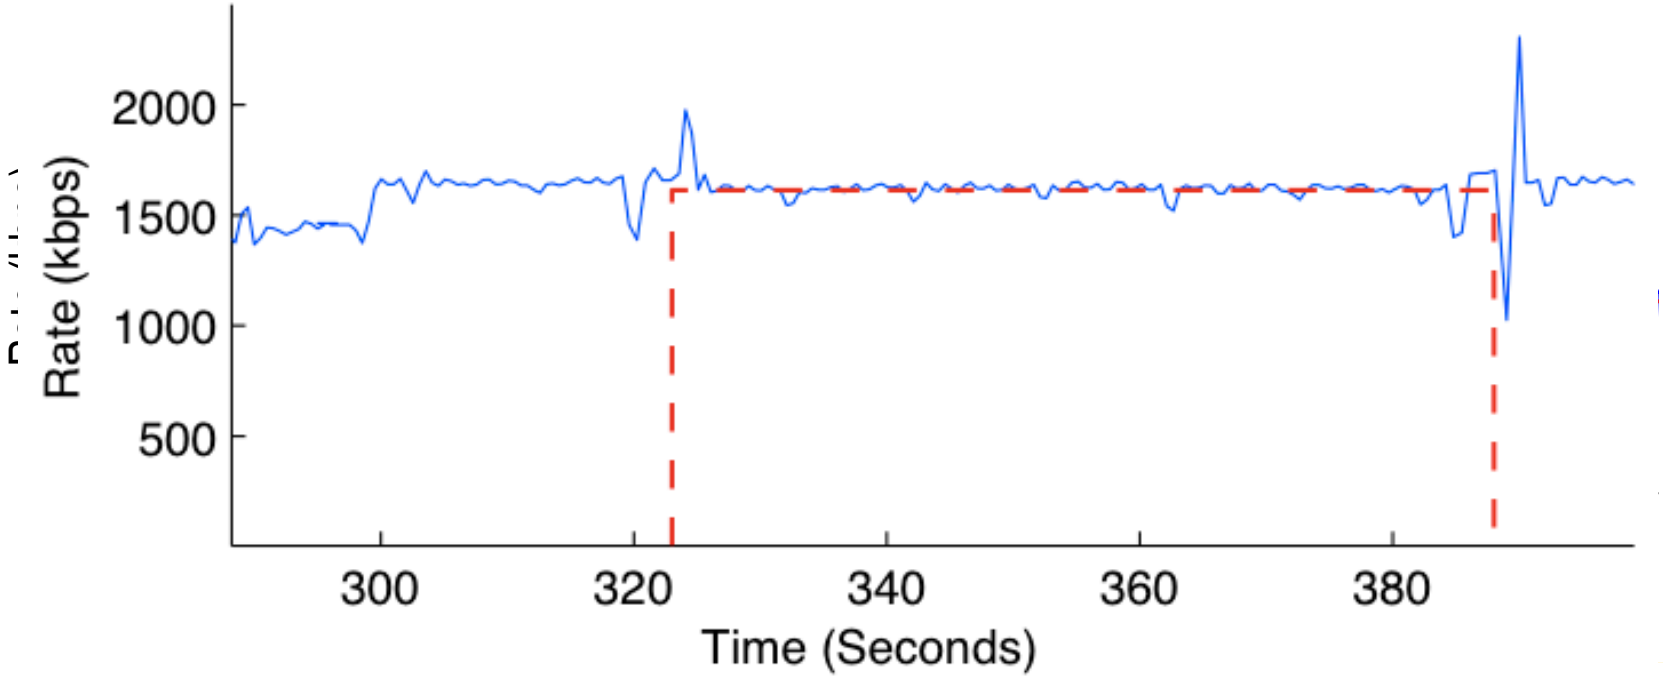

Conclusion & on-going work

- VEIL: Enables large-scale, efficient & robust Ethernet networks
  - Practical realization of VIRO
    - VIRO provides a scalable & robust substrate for future networks
    - No flooding in both data and control planes
  - Backward compatible
    - compatible with current host protocols (such as ARP etc)
  - Enables (nearly) configuration-free networks
  - Essential for seamless mobility
  - Built-in support for Multi-path routing
- Ongoing work:
  - Prototype using OpenFlow based switches
  - Inter-domain routing issues

Please visit http://networking.cs.umn.edu/veil for:

- Demo videos,
- List of related publications,
- Source code!

Or simply search online for "VIRO VEIL"

Thanks!