Increasing Application Performance In Virtual Environments Through Run-time Inference and Adaptation

Ananth I. Sundararaj

Ashish Gupta

Peter A. Dinda

Prescience Lab

Department of Computer Science

Northwestern University

http://virtuoso.cs.northwestern.edu

Summary

- Dynamically adapt unmodified applications on unmodified operating systems in virtual environments to available resources
- The adaptation mechanisms are application independent and controlled automatically without user or developer help
- Demonstrate the feasibility of adaptation at the level of collection of VMs connected by Virtual Networks
- Show that its benefits can be significant for two classes of applications

Outline

- Virtual machine grid computing
- Virtuoso system
- Networking challenges in Virtuoso
- Enter VNET
- VNET, VTTIF ⇨ Adaptive virtual network
- Evaluation
- Summary

Virtual Machine Grid Computing

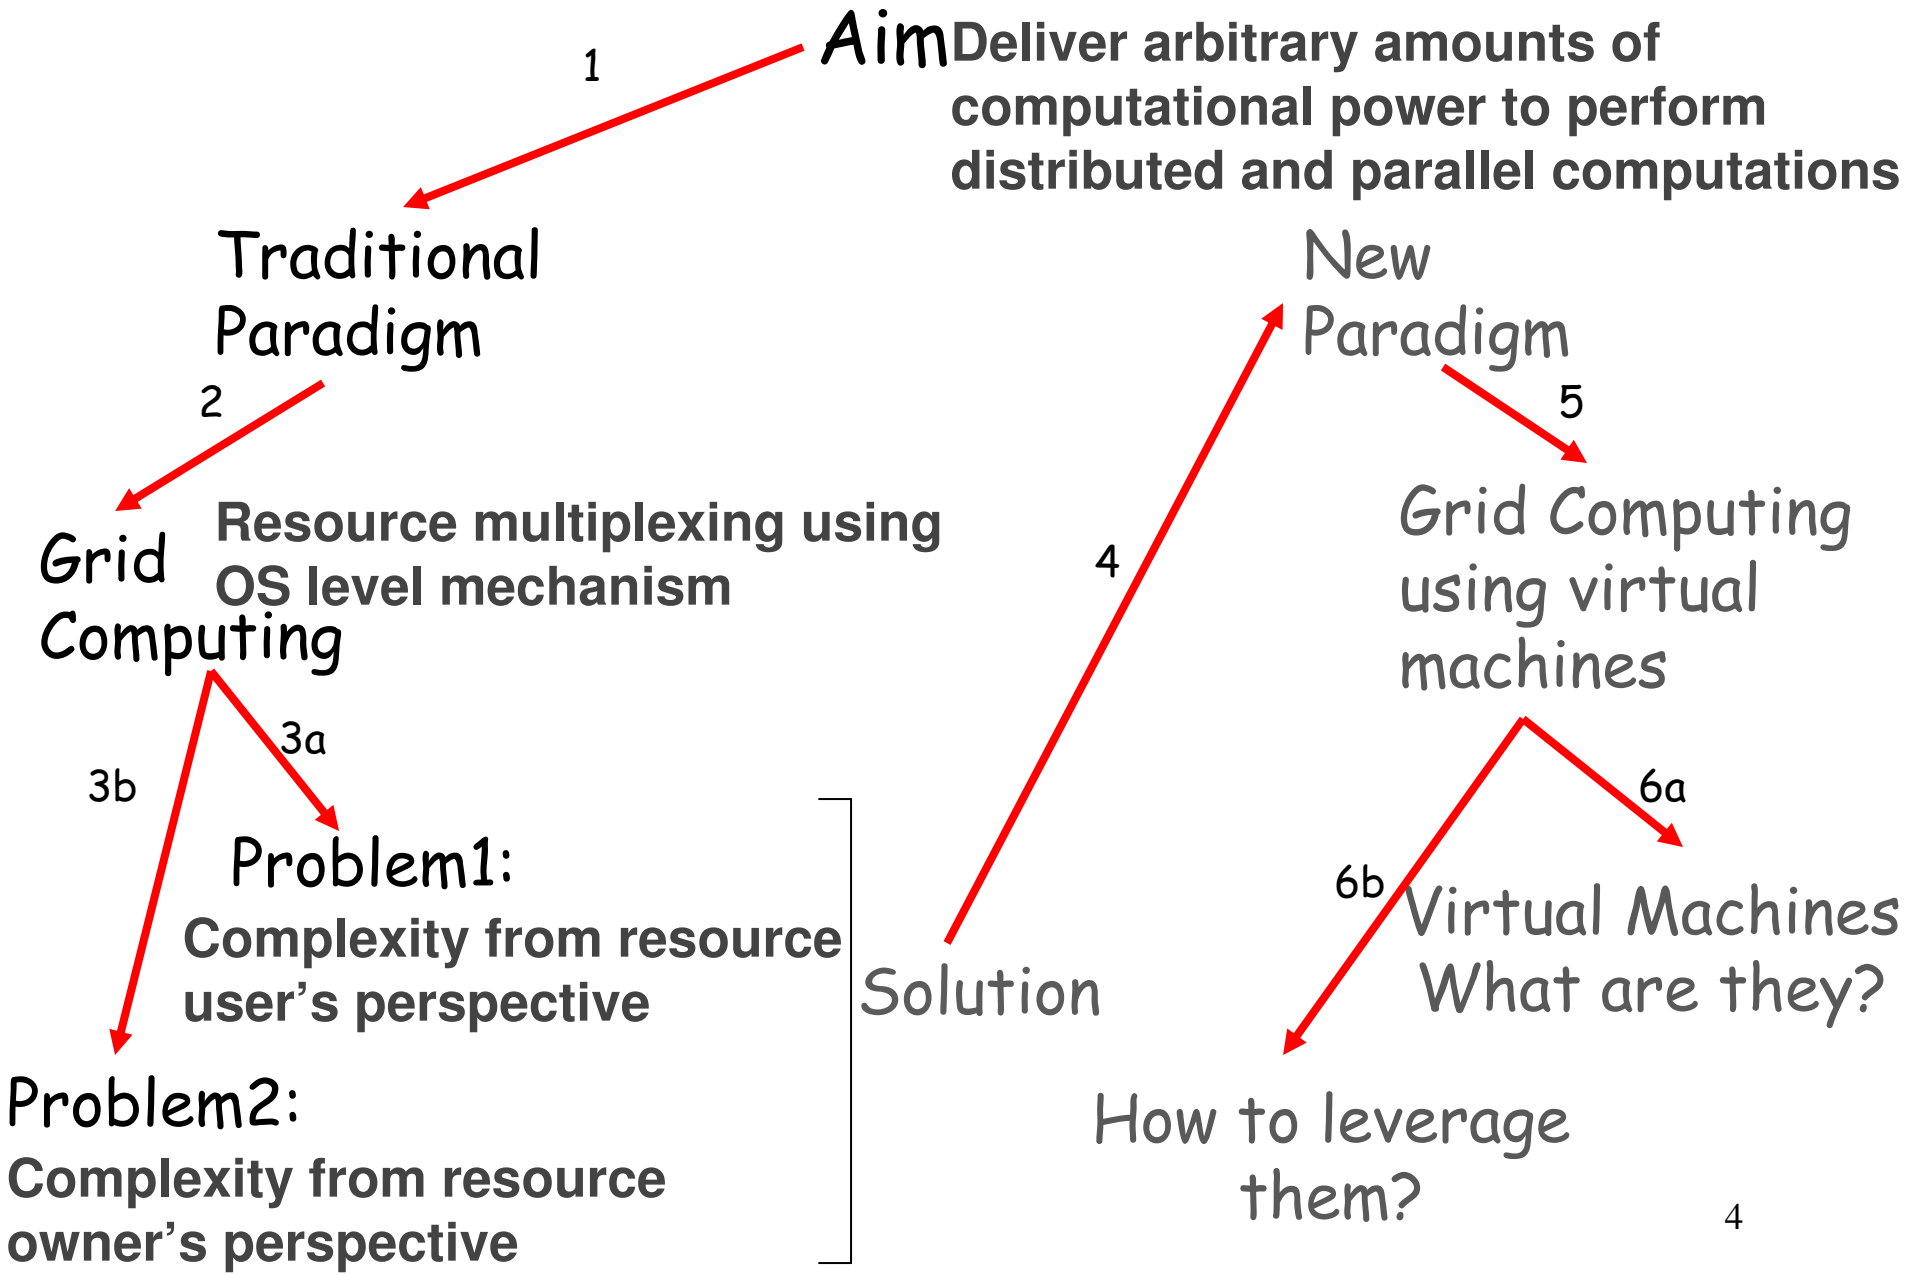

Virtual Machines

Virtual machine monitors (VMMs)

- Raw machine is the abstraction
- VM represented by a single image
- VMware GSX Server

The Simplified Virtuoso Model

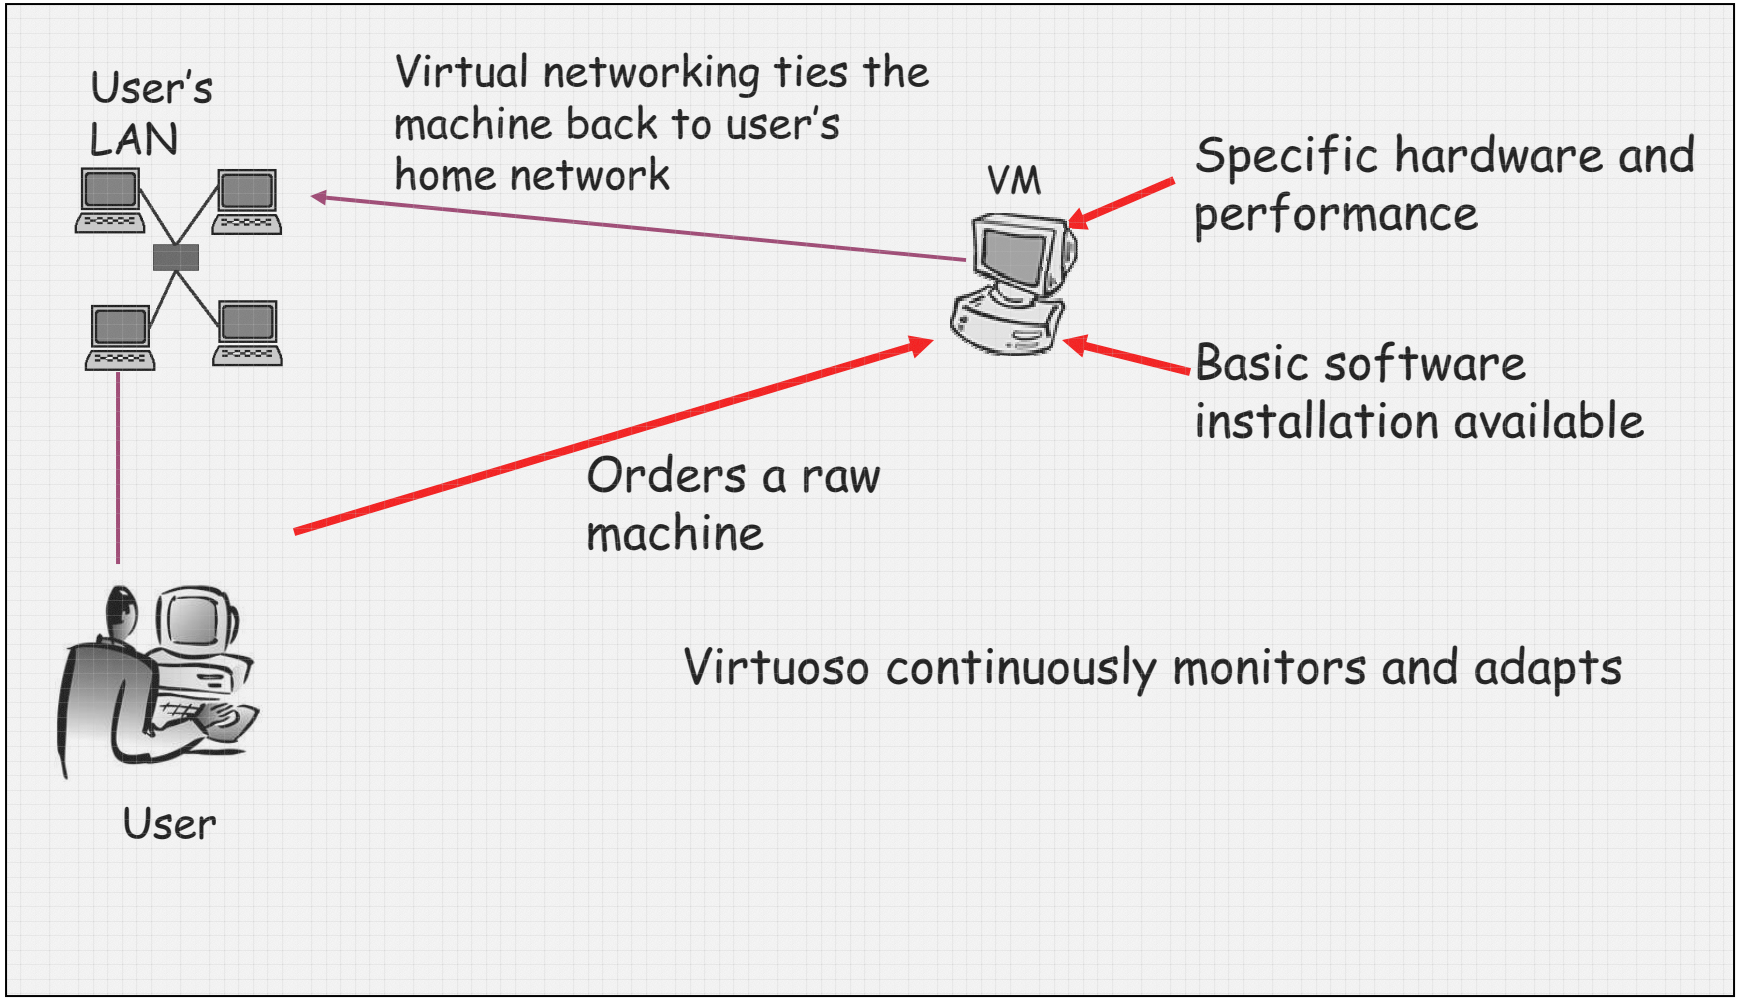

User's View in Virtuoso Model

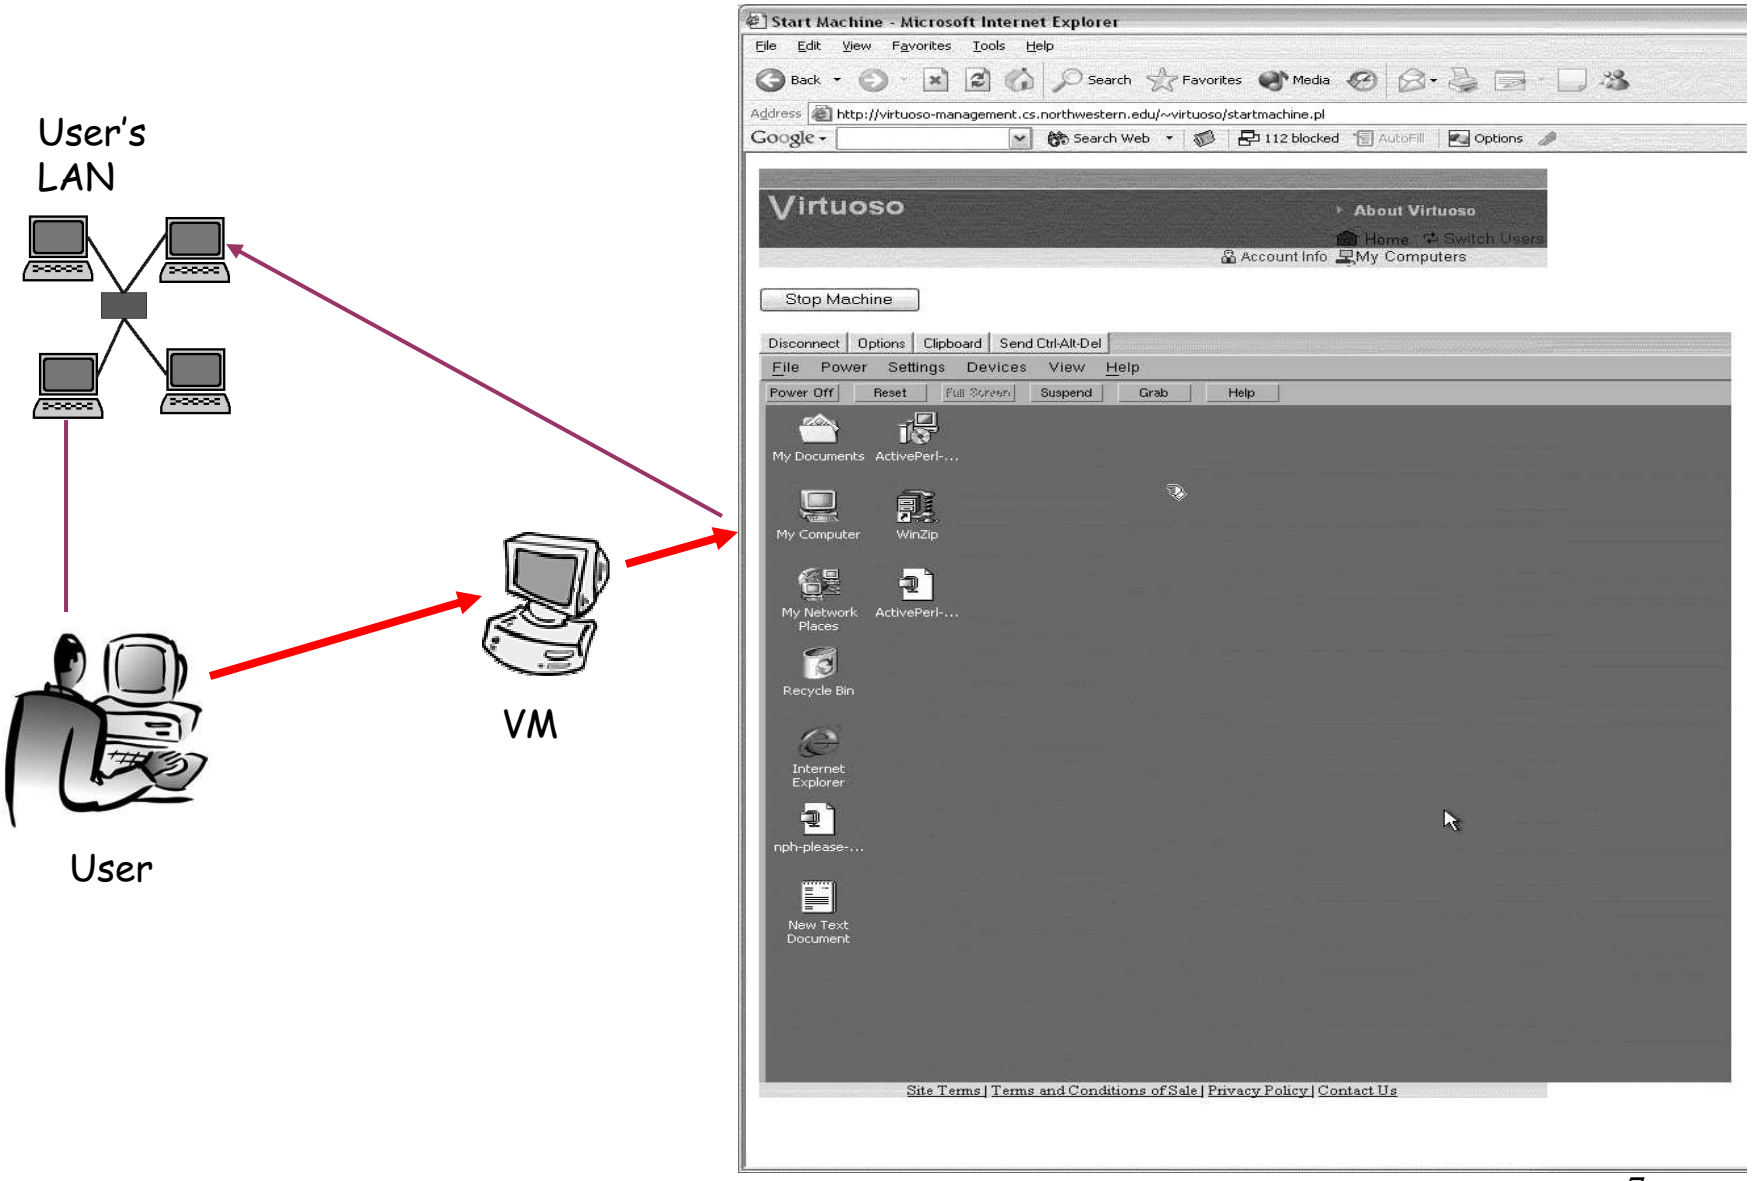

Outline

- Virtual machine grid computing
- Virtuoso system
- Networking challenges in Virtuoso
- Enter VNET
- VNET, VTTIF ⇨ Adaptive virtual network
- Evaluation
- Summary

Virtual Networks

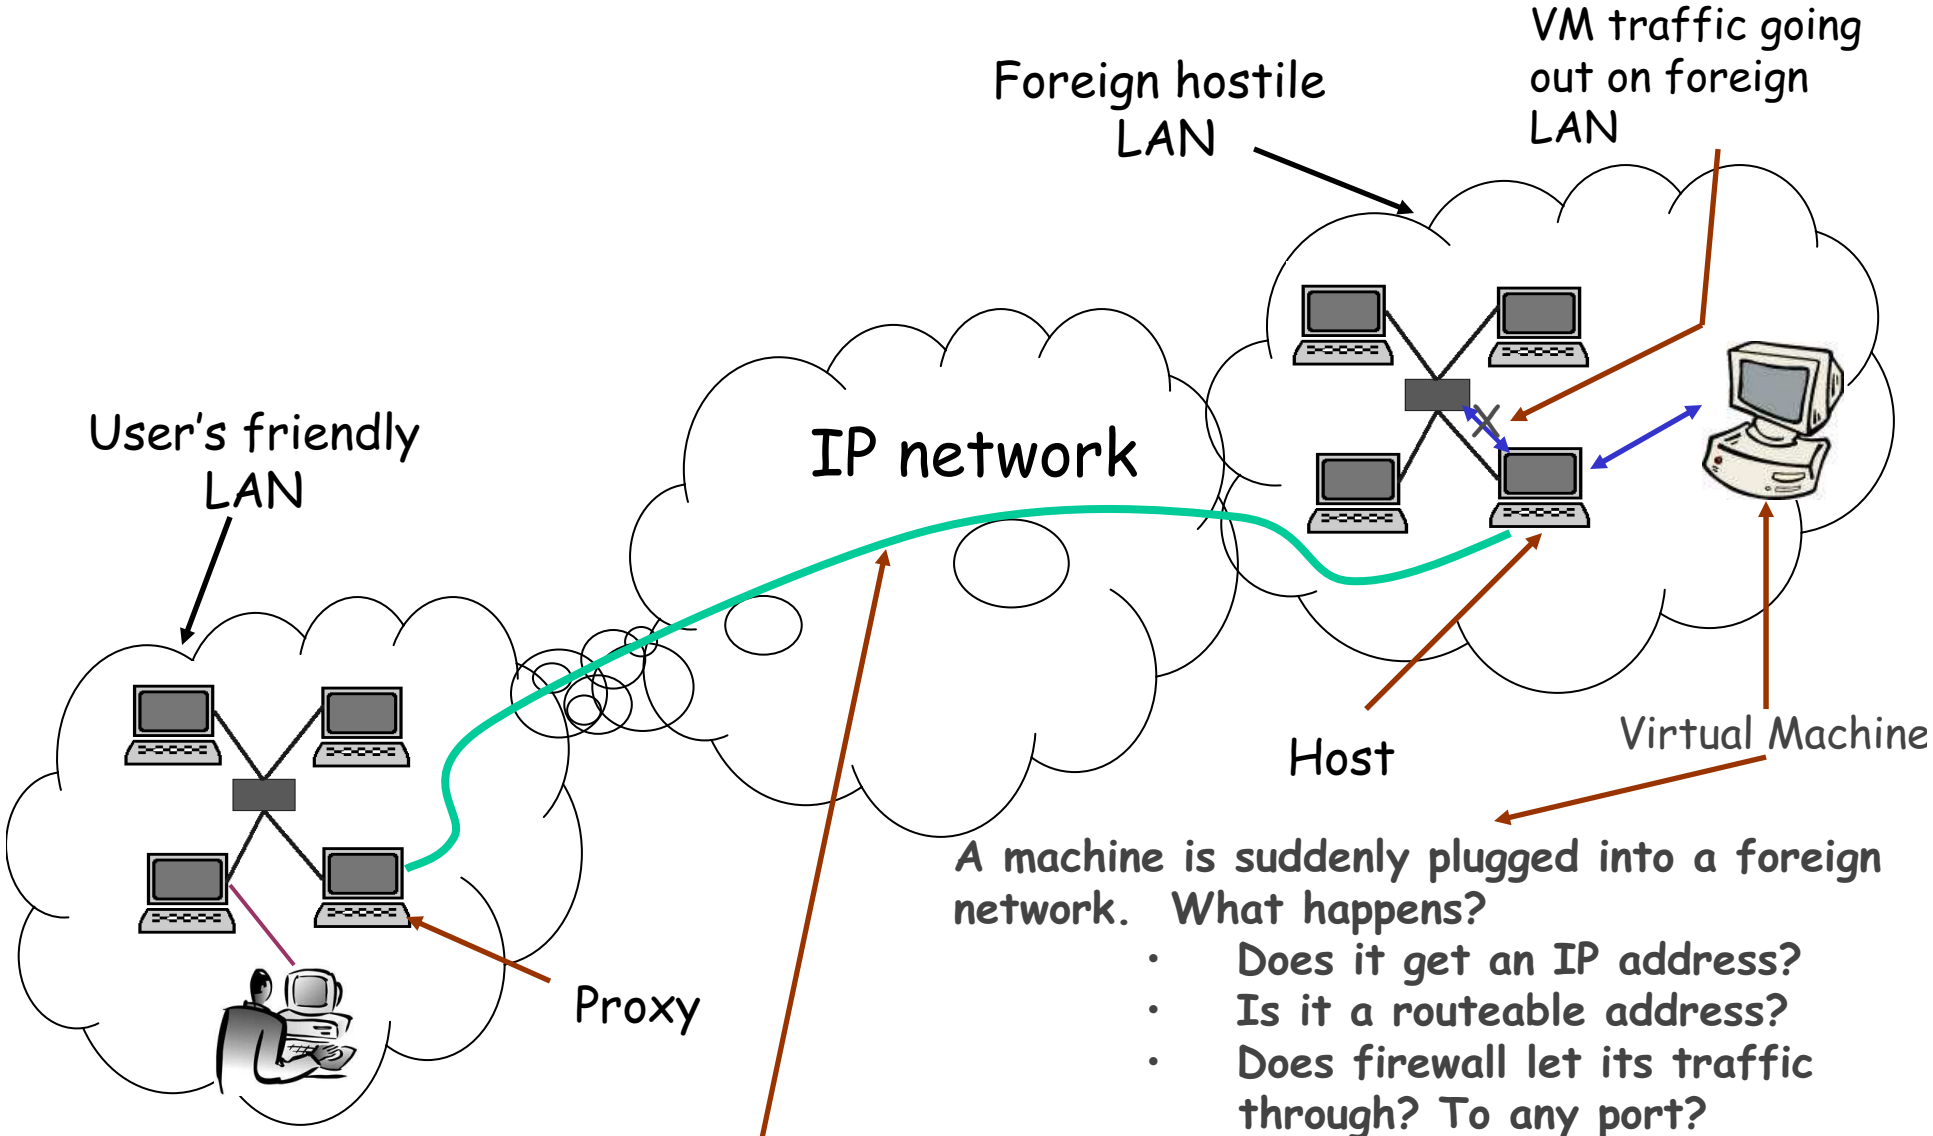

A machine is suddenly plugged into a foreign network. What happens?

- Does it get an IP address?
- Is it a routeable address?
- Does firewall let its traffic through? To any port?

VNET: A bridge with long wires

A VNET Link

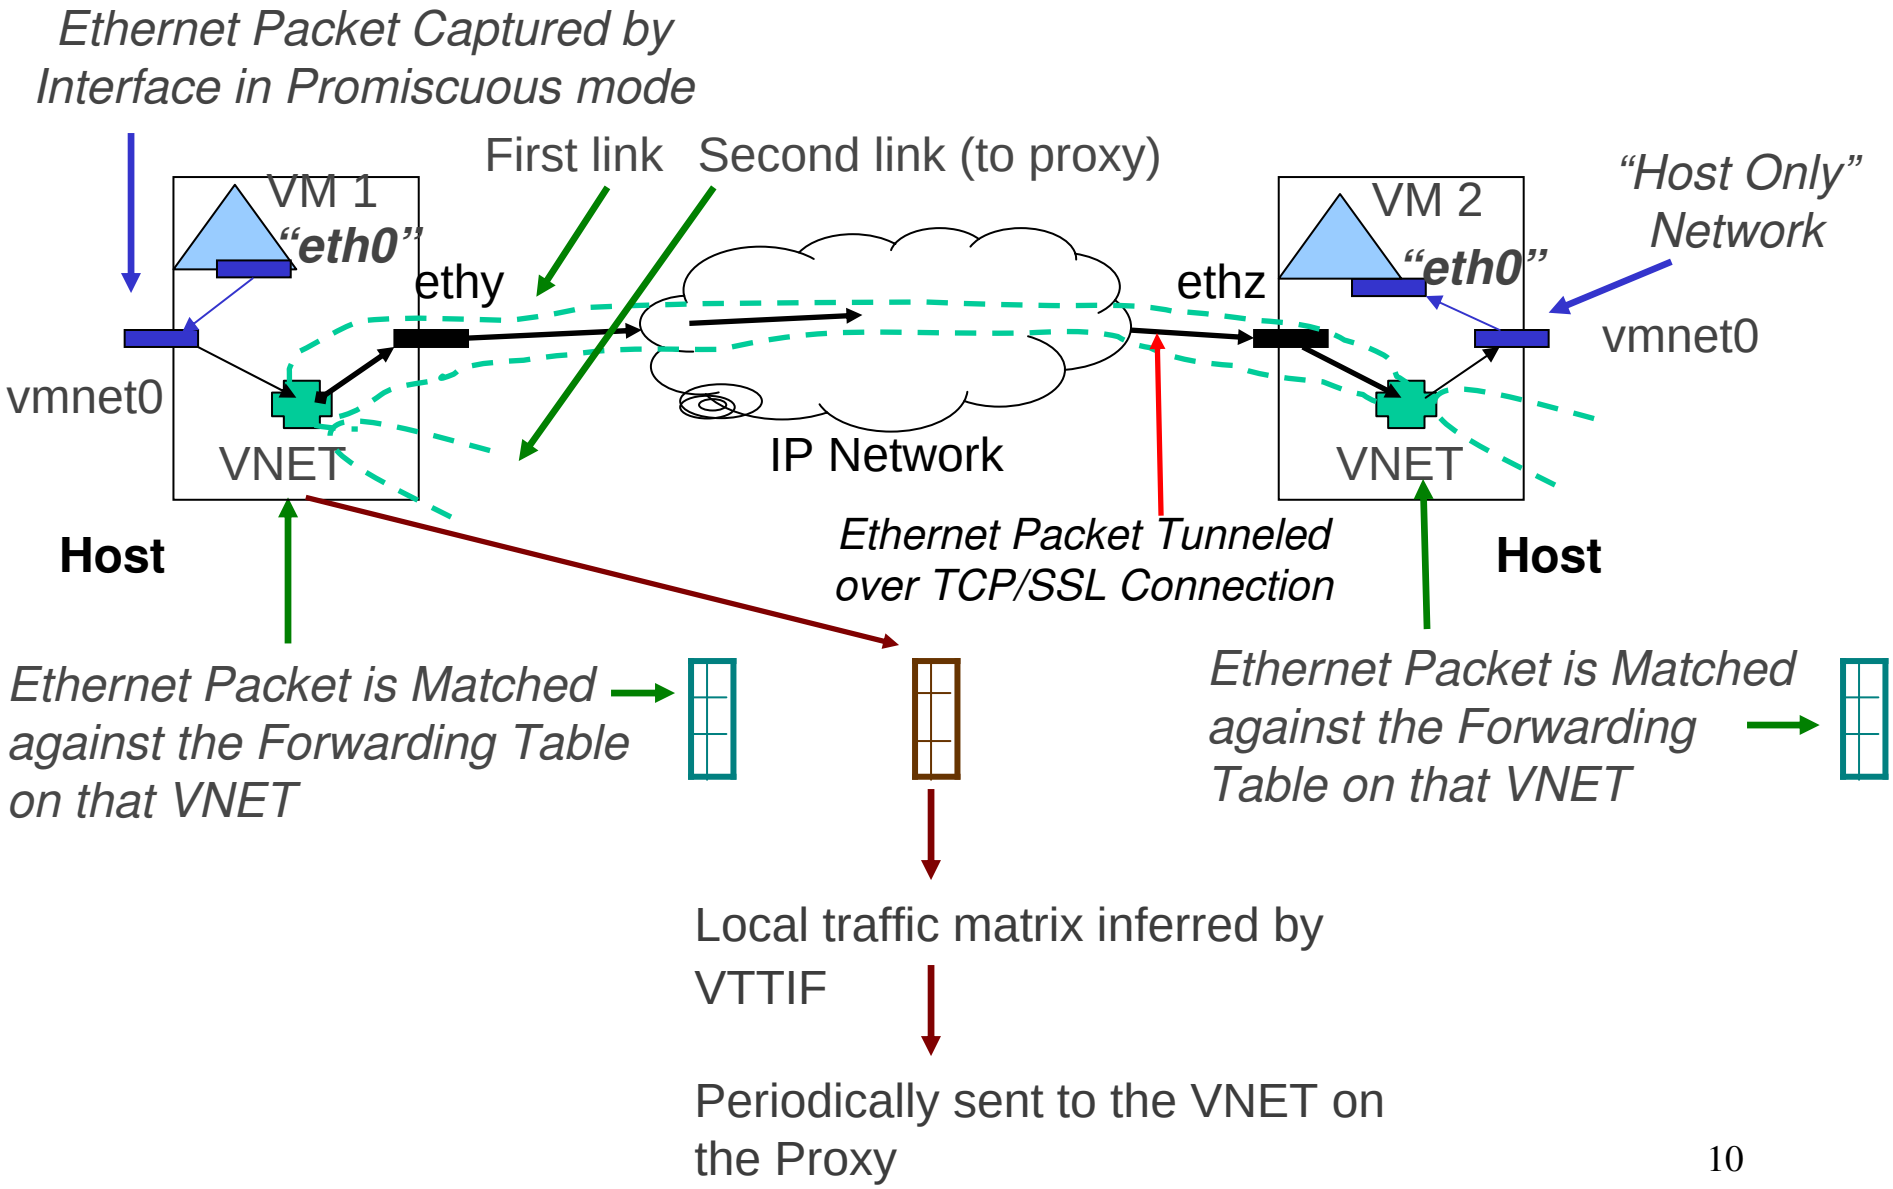
Ethernet Packet Captured by Interface in Promiscuous mode
First link
Second link (to proxy)
VM 1
"eth0"
ethy
ethz
VM 2
"eth0"
"Host Only" Network
vmnet0
vmnet0
IP Network
VNET
VNET
Host
Host
Ethernet Packet Tunneled over TCP/SSL Connection
Ethernet Packet is Matched against the Forwarding Table on that VNET
Ethernet Packet is Matched against the Forwarding Table on that VNET
Local traffic matrix inferred by VTTIF
Periodically sent to the VNET on the Proxy

Virtual Topology and Traffic Inference Framework (VTTIF) Operation

Ethernet-level traffic monitoring

```
-y.loc-serv > nera-x.1035:  udp 172
1.224.18.www-http > 137.133.57.68.1255: tcp 664
1.224.18.www-http > 137.133.57.68.1255: tcp 834
133.57.68.1255 > 209.1.224.18.www-http: tcp 0 (DF)
1.224.18.www-http > 137.133.57.68.1255: tcp 326
1.224.18.www-http > 137.133.57.68.1255: tcp 1160
133.57.68.1255 > 209.1.224.18.www-http: tcp 0 (DF)
1.224.18.www-http > 137.133.57.68.1255: tcp 1106
```

VNET daemons collectively aggregate a global traffic matrix for all VMs

	h1	h2	h3	h4
h1		7.7	7.6	7.8
h2	13.1		6.6	6.5
h3	13.5	6.4		6.6
h4	13.2	6.5	6.5	

*numbers indicate MB of data transferred.

Application topology is recovered using normalization and pruning algorithms

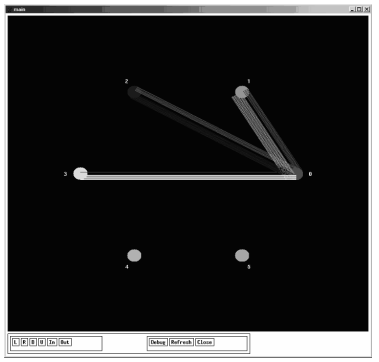

Dynamic Topology Inference by VTTIF

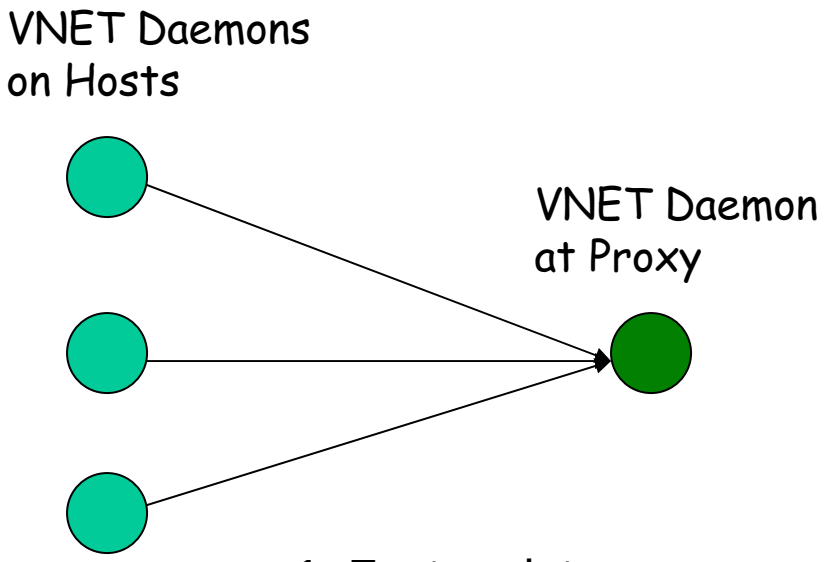

1. Fast updates

0.7	0.6	0.5	0.2
1.0	0.2	0.3	0.0
0.4	0.3	0.2	0.1
0.2	0.3	0.5	0.3

Aggregated
Traffic Matrix

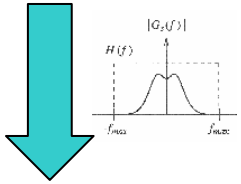

0.5	0.6	0.5	0.3
0.9	0.3	0.3	0.2
0.4	0.3	0.3	0.4
0.3	0.3	0.5	0.3

Smoothed
Traffic Matrix

2. Low Pass Filter
Aggregation

Topology change output

3. Threshold
change detection

Outline

- Virtual machine grid computing
- Virtuoso system
- Networking challenges in Virtuoso
- Enter VNET
- VNET, VTTIF ⇒ Adaptive virtual network
- Evaluation
- Summary

Adaptation

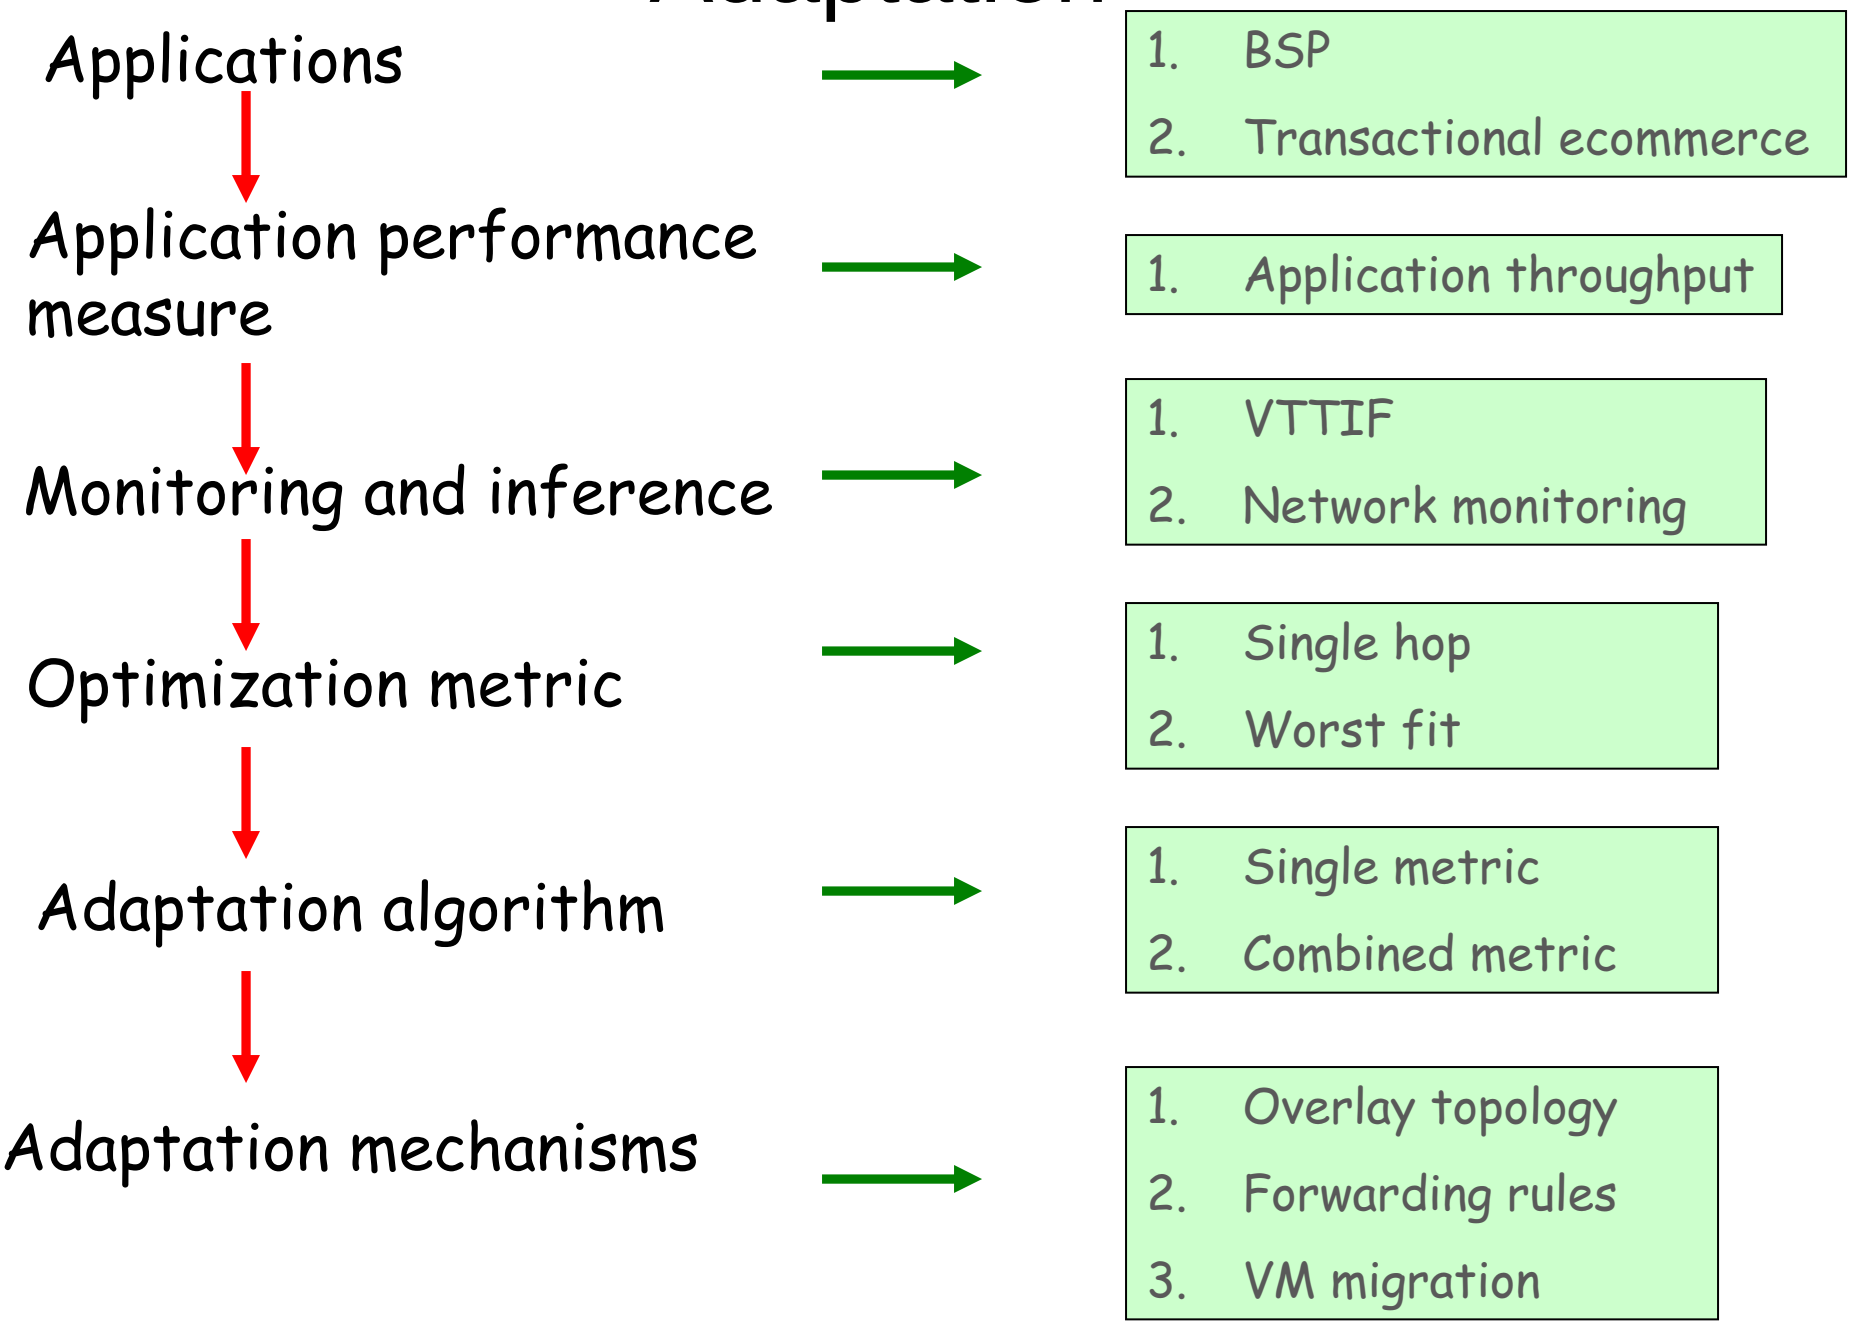
Applications
1. BSP
2. Transactional ecommerce
Application performance measure
1. Application throughput
Monitoring and inference
1. VTTIF
2. Network monitoring
Optimization metric
1. Single hop
2. Worst fit
Adaptation algorithm
1. Single metric
2. Combined metric
Adaptation mechanisms
1. Overlay topology
2. Forwarding rules
3. VM migration

Optimization Problem (1/2) Topology Only

Informally stated:

- Input
  - Network traffic load matrix of application
- Output
  - Overlay topology connecting hosts
  - Forwarding rules on the topology
  - q Such that the application throughput is maximized

The algorithm is described in detail in the paper

Illustration of Topology Adaptation in Virtuoso

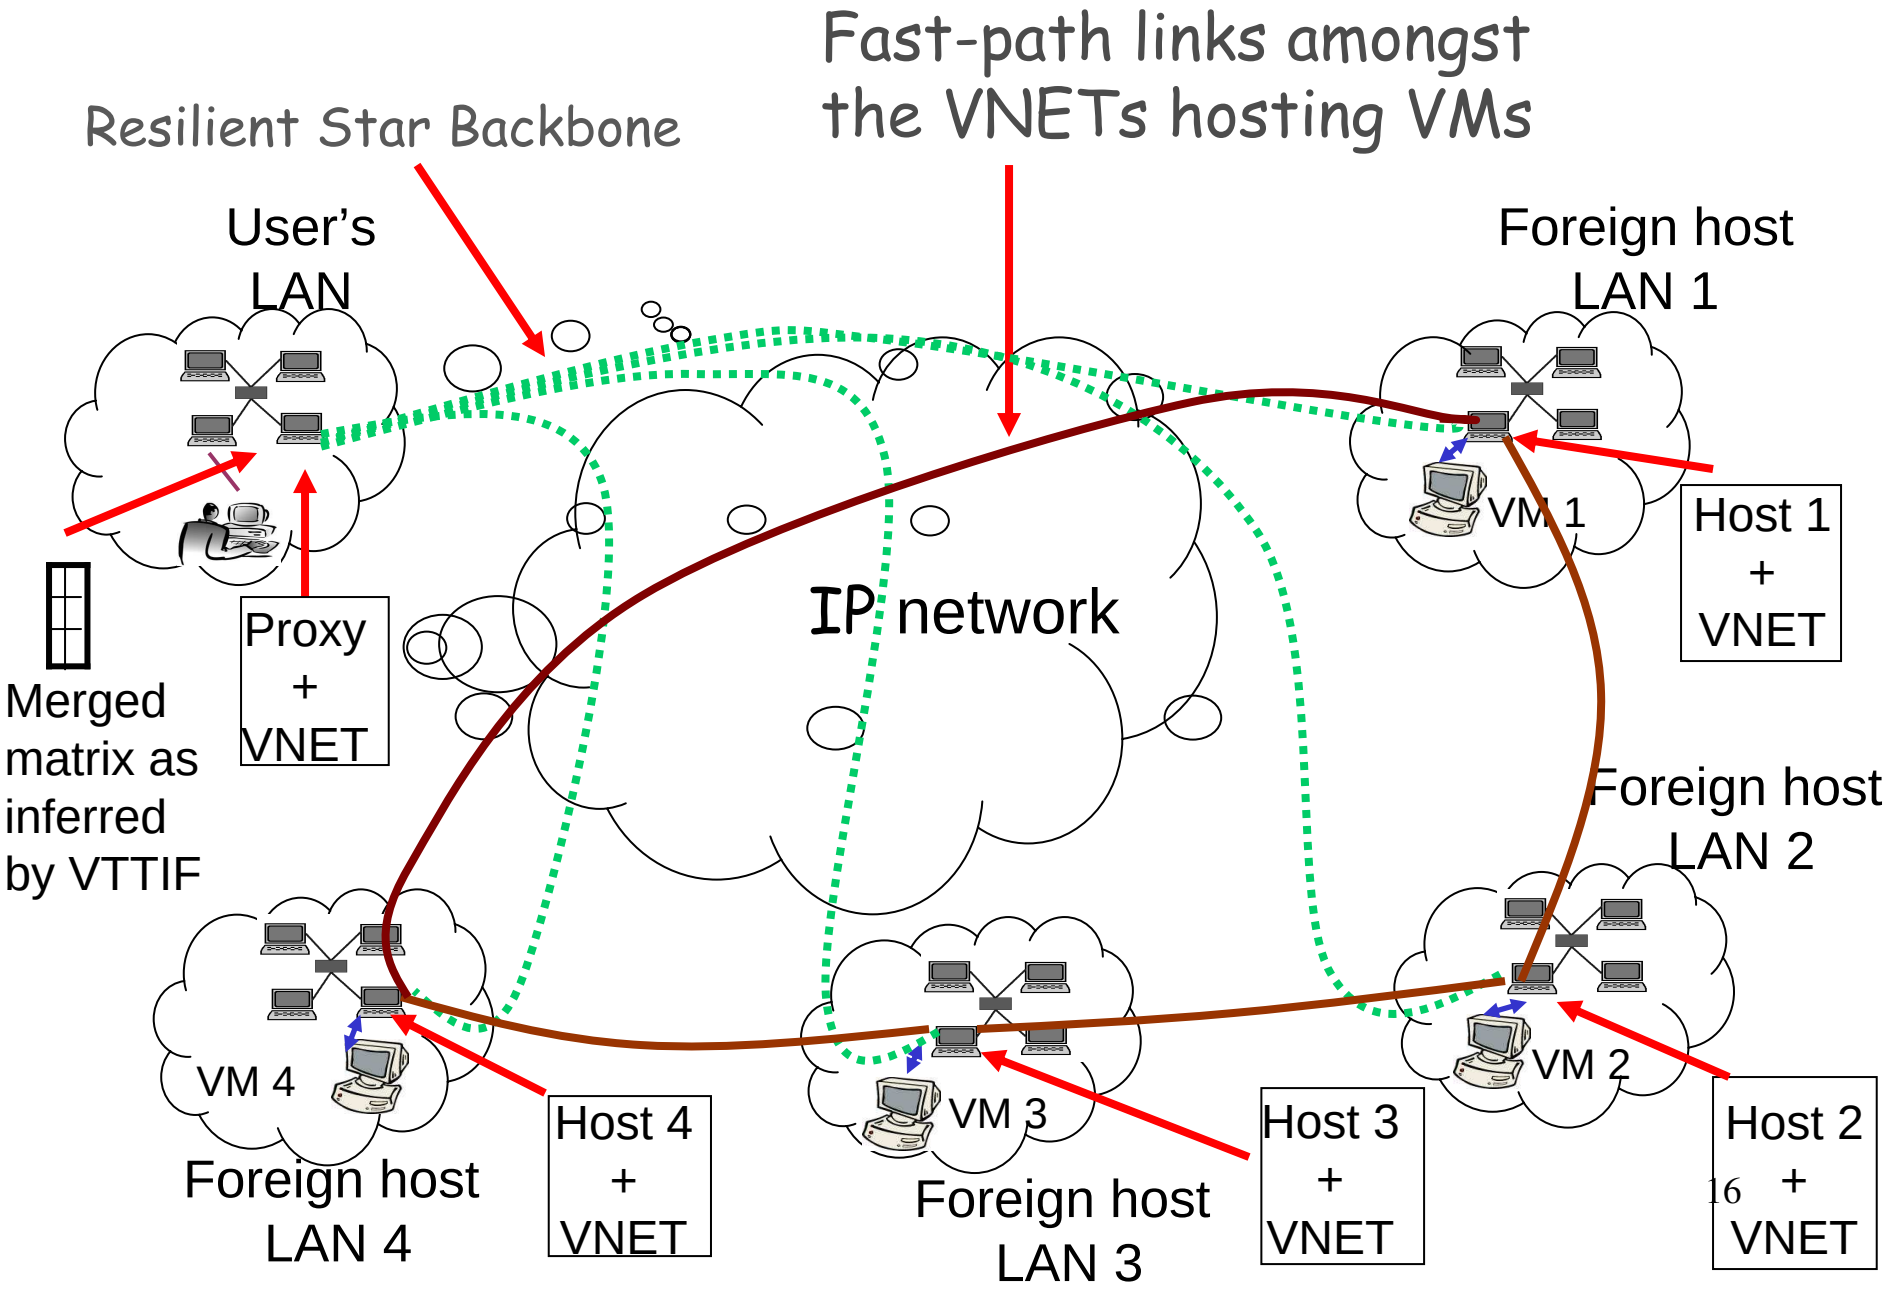

Evaluation

- Reaction time of VNET
- Patterns: A synthetic BSP benchmark
- Benefits of adaptation (performance speedup)
  - Eight VMs on a single cluster, all-all topology
  - Eight VMs spread over WAN, all-all topology

Wide-Area testbed

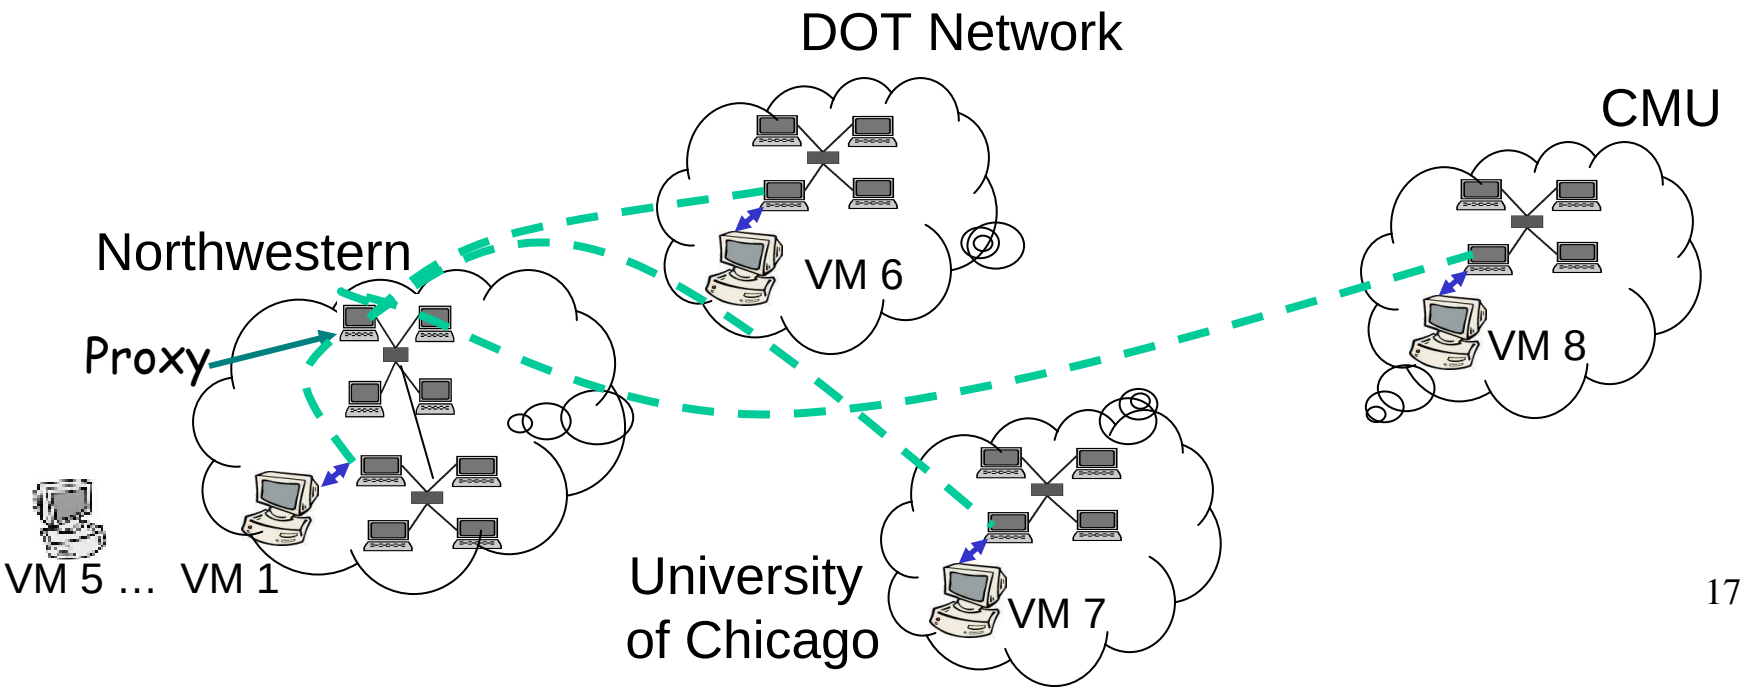

Reaction Time

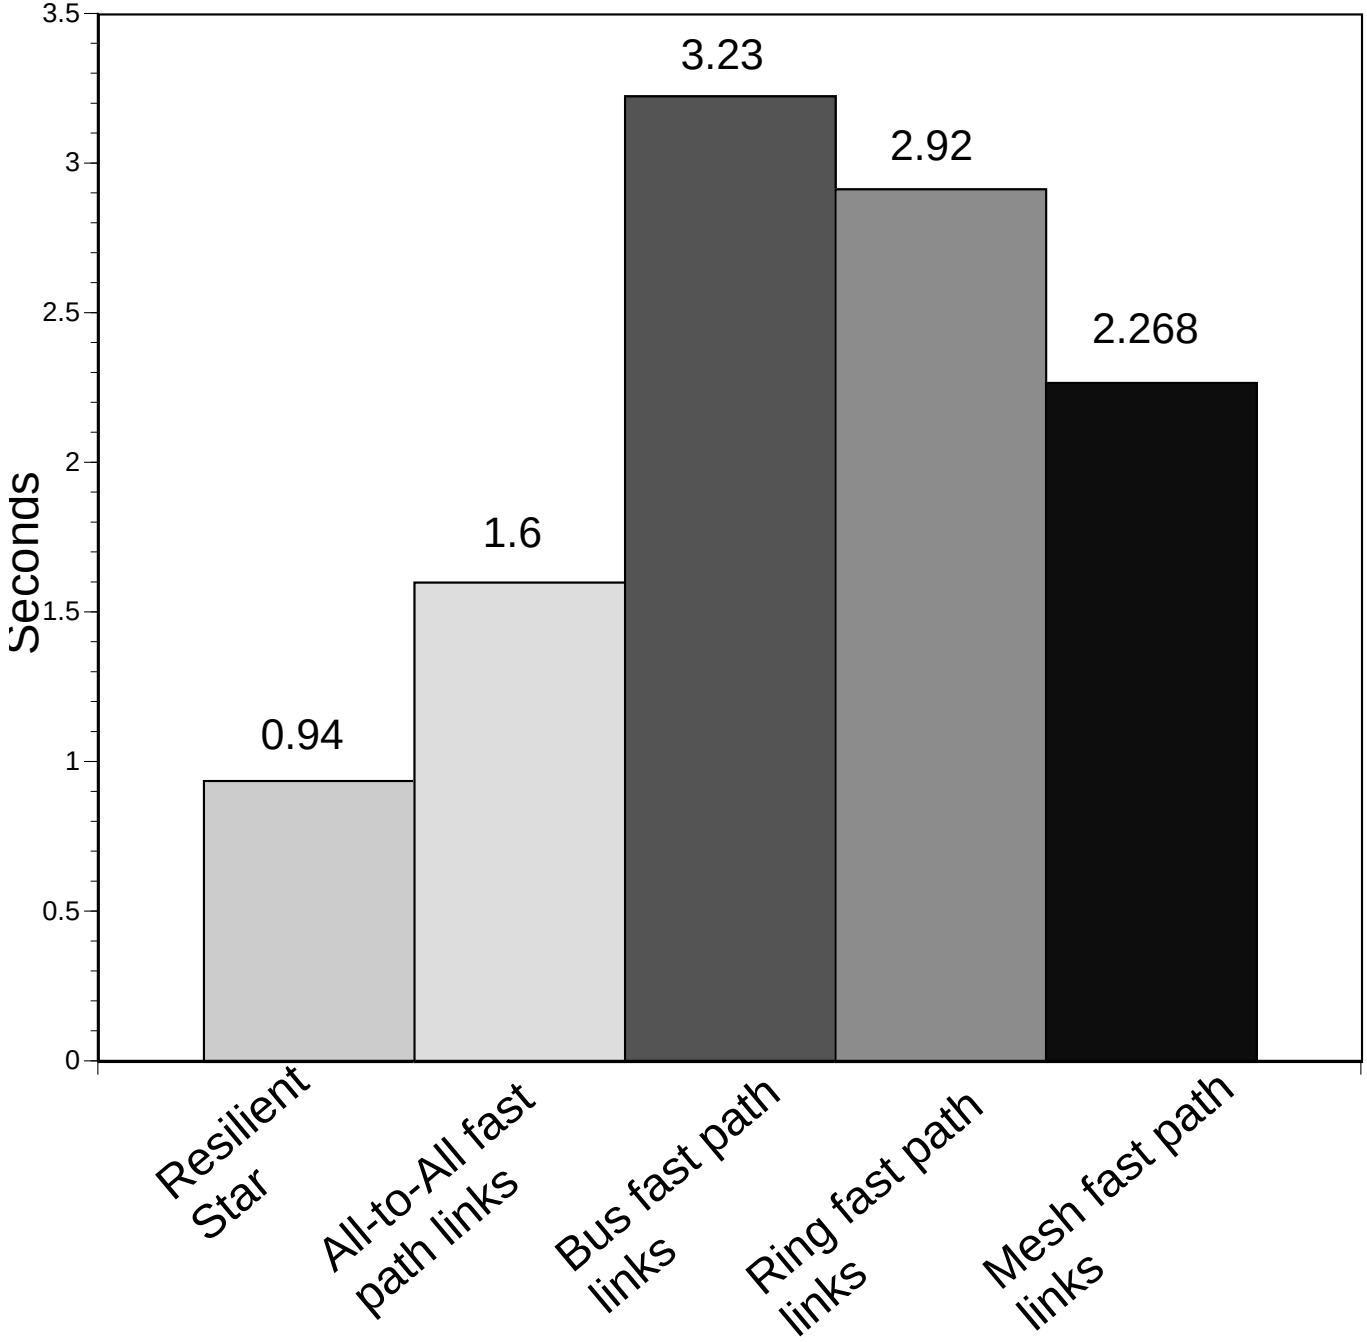

Benefits of Adaptation

Benefits accrued as a function of the number of fast-path links added

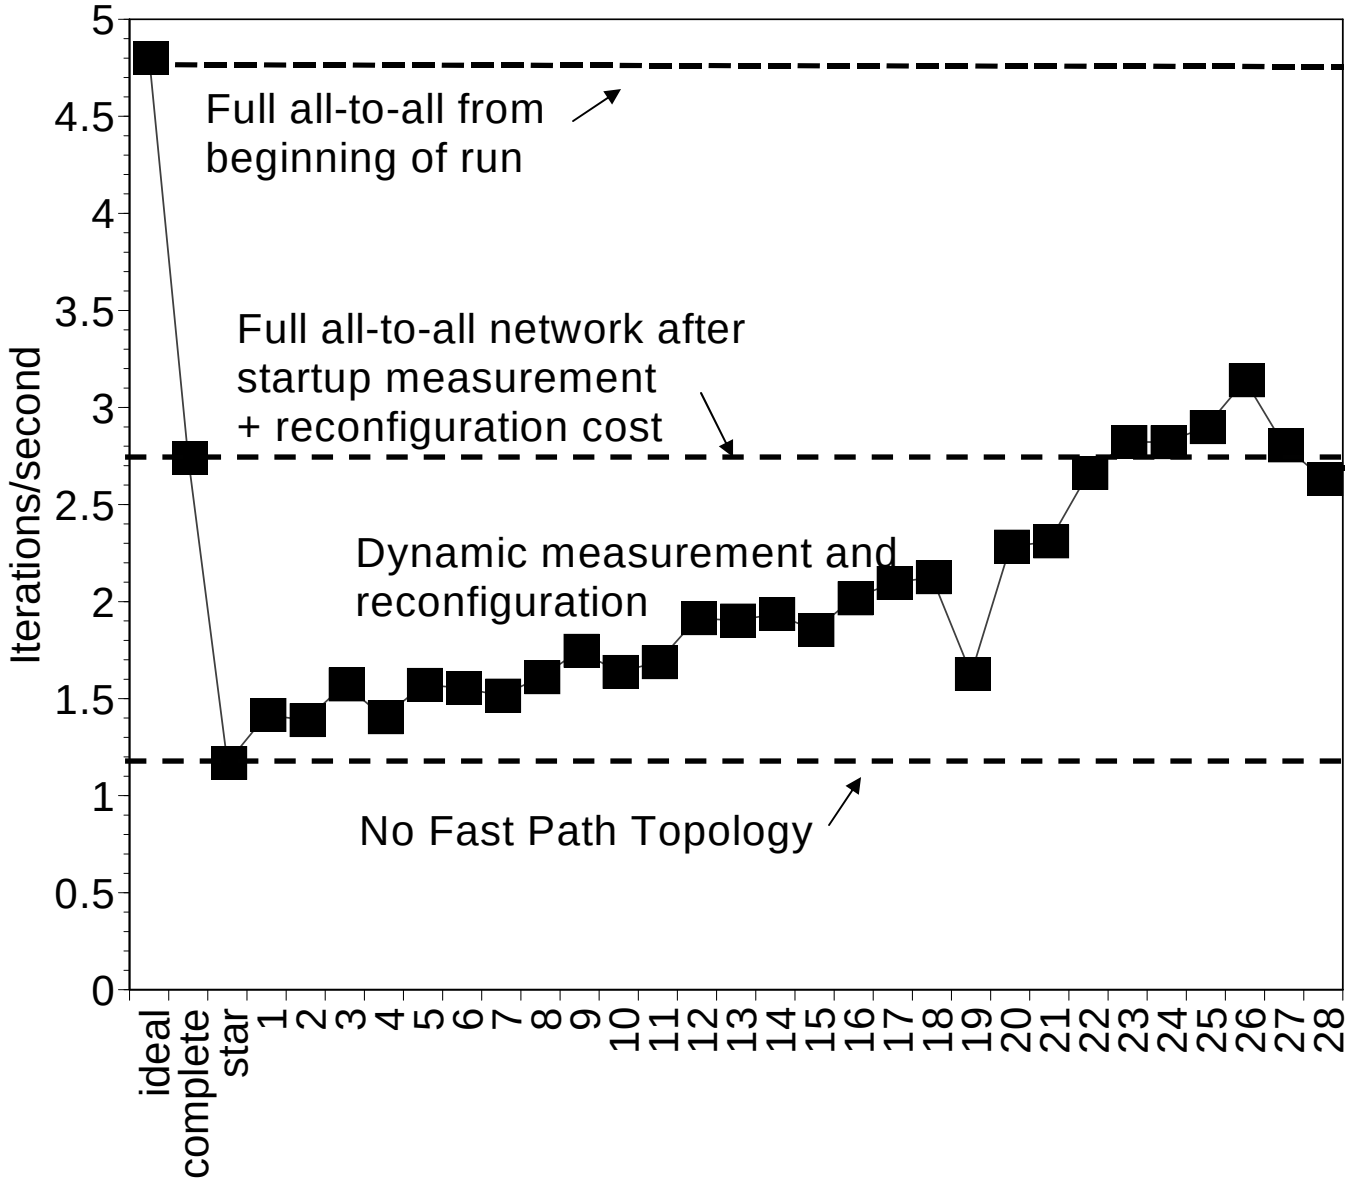

- Patterns has an all-all topology
- Eight VMs are used
- All VMs are hosted on the same cluster

Benefits of Adaptation

Benefits accrued as a function of the number of fast-path links added

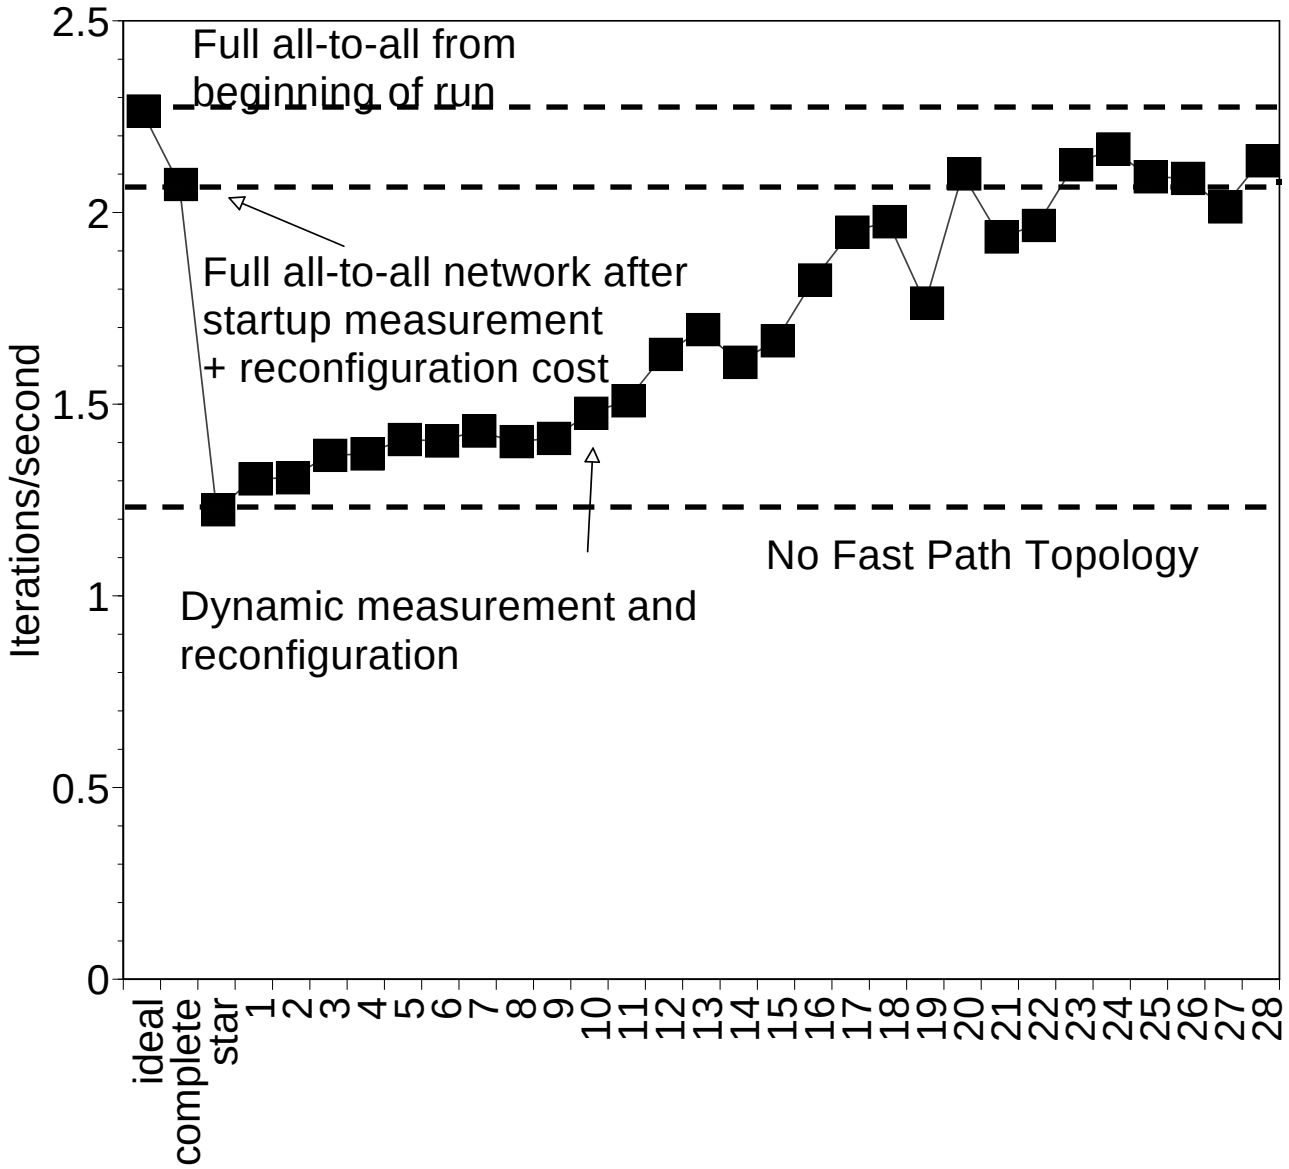

- Patterns has an all-all topology
- Eight VMs are used
- VMs are spread over WAN

Optimization Problem (2/2)
Topology + Migration

Informally stated:

- Input
  - Network traffic load matrix of application
  - Topology of the network
- Output
  - Mapping of VMs to hosts
  - Overlay topology connecting hosts
  - Forwarding rules on the topology
  - Such that the application throughput is maximized

The algorithm is described in detail in the paper

Evaluation

- Applications
  - Patterns: A synthetic BSP benchmark
  - TPC-W: Transactional web ecommerce benchmark
- Benefits of adaptation (performance speedup)
  - Adapting to compute/communicate ratio
  - Adapting to external load imbalance

Effect on BSP Application Throughput of Adapting to Compute/Communicate Ratio

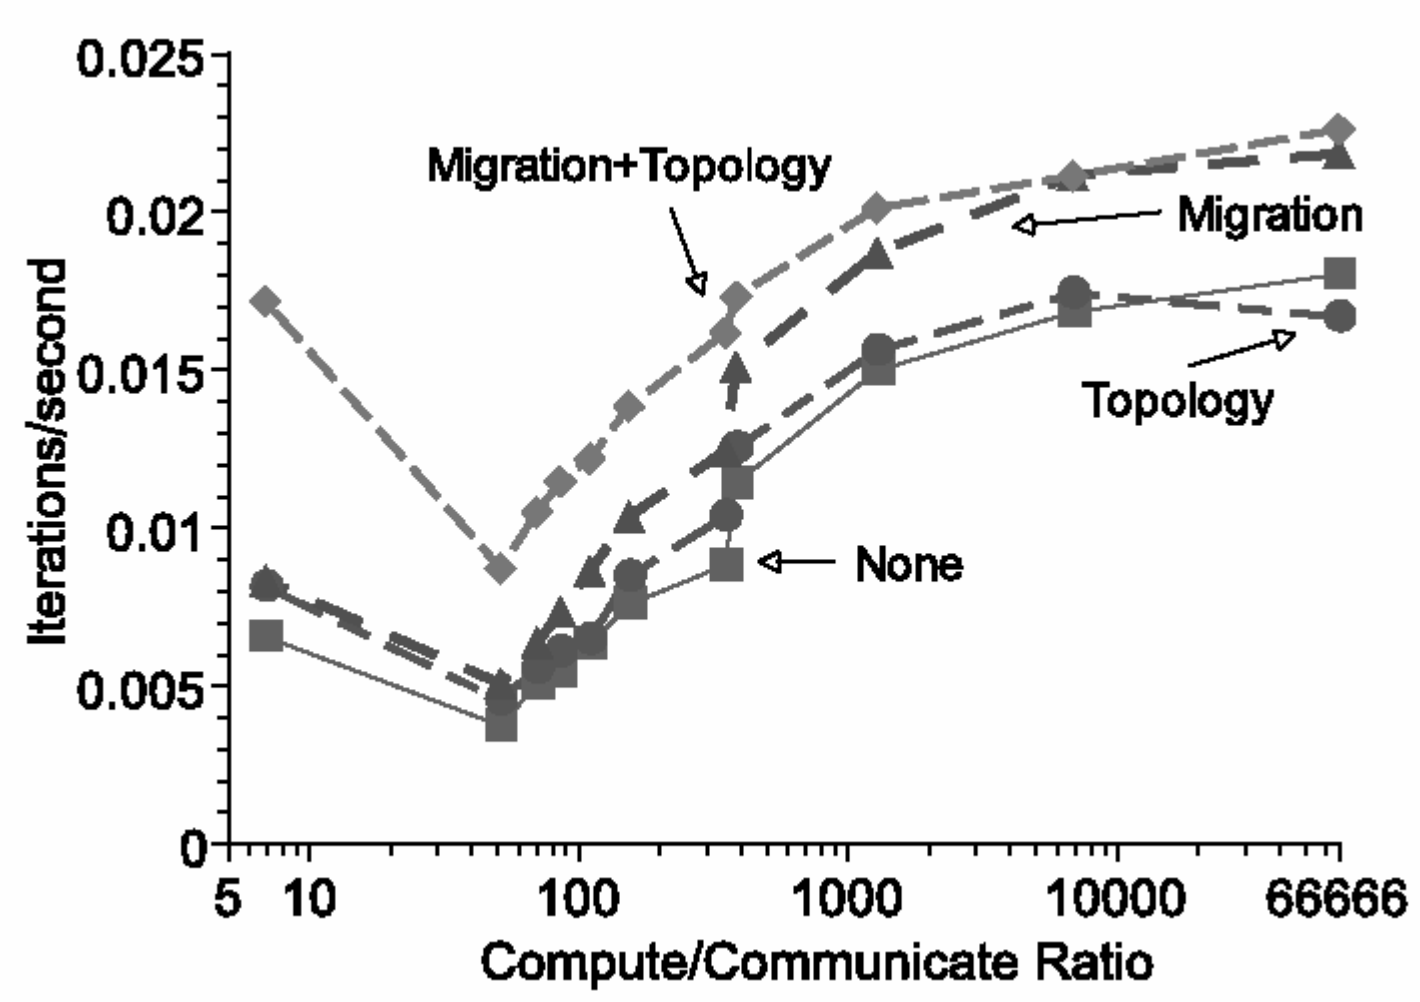

Effect on BSP Application Throughput of Adapting to External Load Imbalance

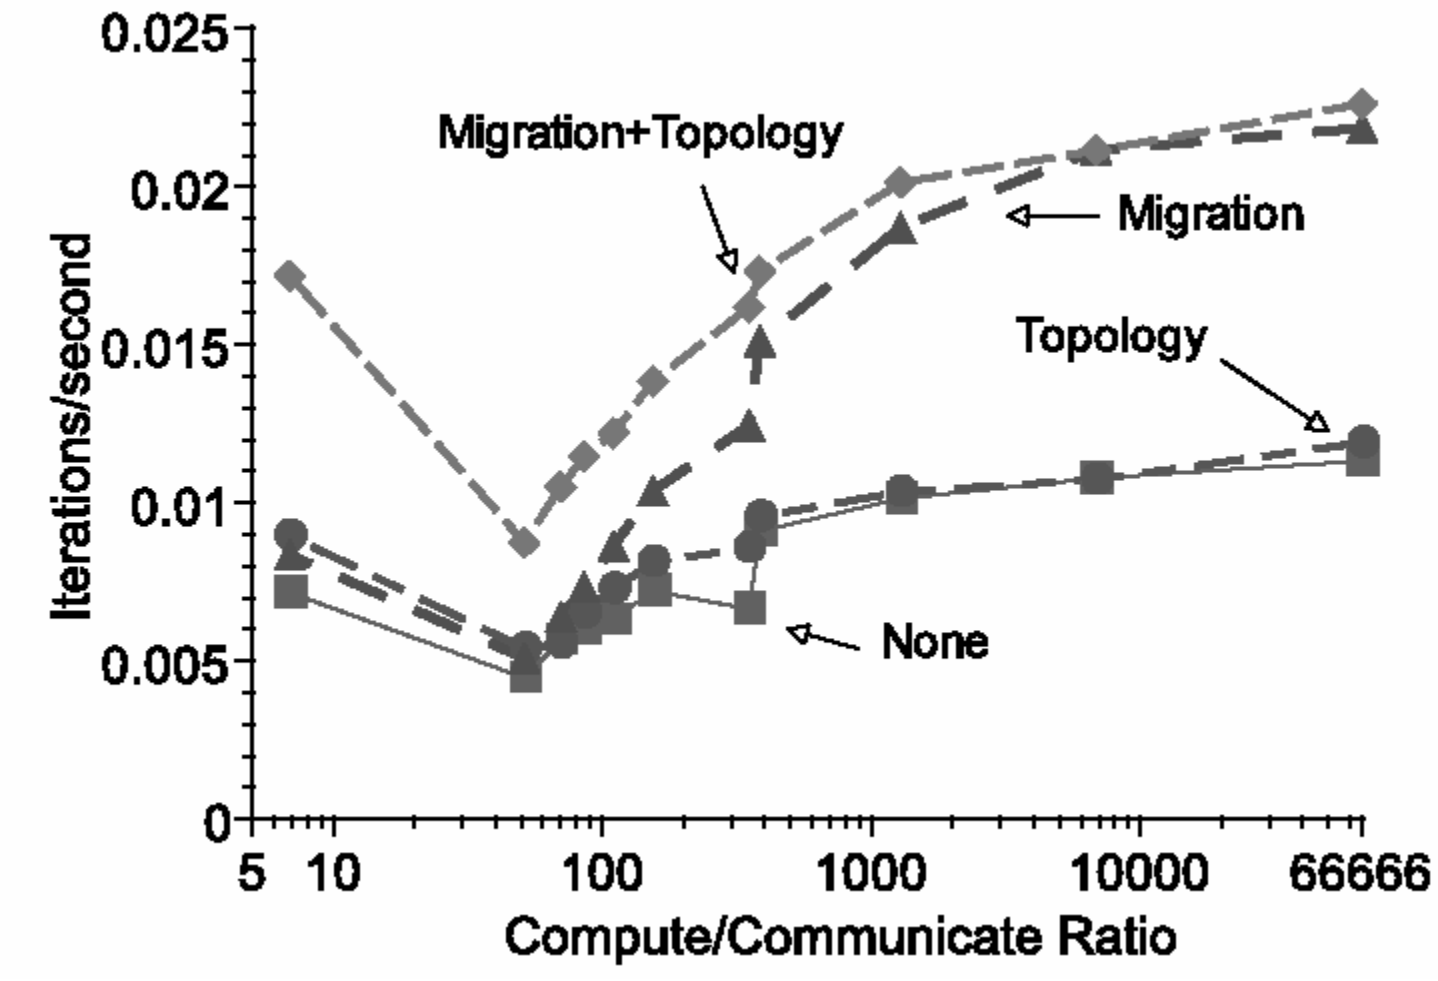

TPCW Throughput (WIPS) With Image Server Facing External Load

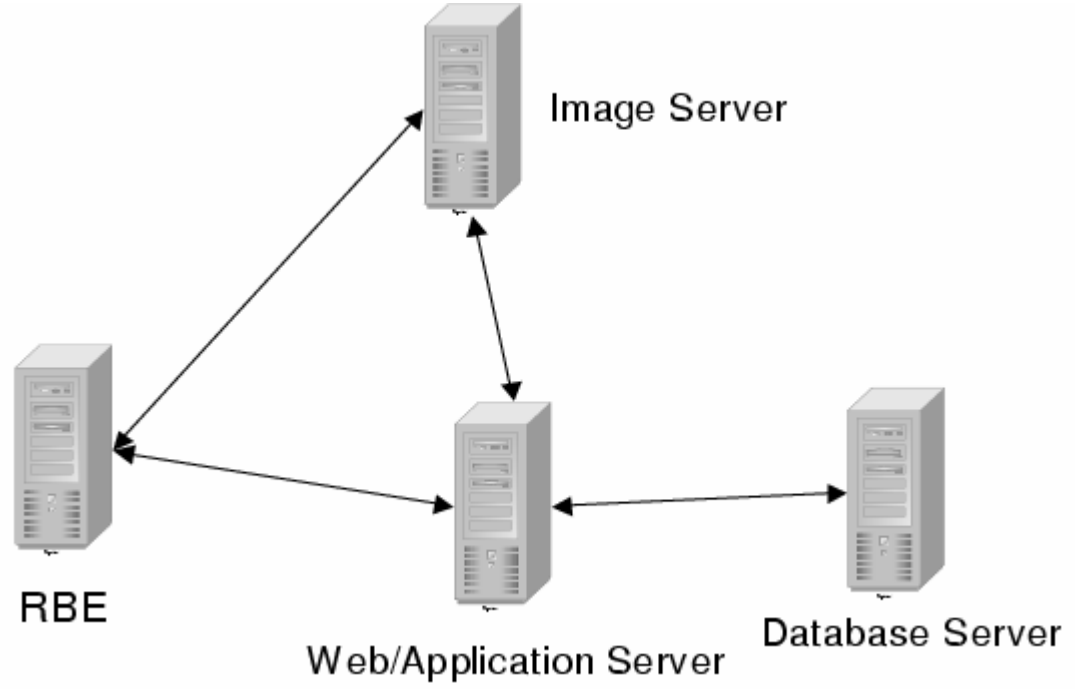

	No Topology	Topology
No Migration	1.216	1.76
Migration	1.4	2.52

Outline

- Virtual machine grid computing
- Virtuoso system
- Networking challenges in Virtuoso
- Enter VNET
- VNET, VTTIF ⇨ Adaptive virtual network
- Evaluation
- Summary

Summary

- Dynamically adapt unmodified applications on unmodified operating systems in virtual environments to available resources
- The adaptation mechanisms are application independent and controlled automatically without user or developer help
- Demonstrate the feasibility of adaptation at the level of collection of VMs connected by Virtual Networks
- Show that its benefits can be significant for two classes of applications

For More Information

- Future Work
  - Free network measurement (Wren) – Collaboration with CS, W&M
  - Applicability of a single optimization scheme

- Related Talk at HPDC 2005
  - J. Lange, A. Sundararaj, P. Dinda, “Automatic Dynamic Run-time Optical Network Reservations”
  - Wednesday, July 27, 2:00 P.M.

- Please visit
  - Prescience Lab (Northwestern University)
    - http://plab.cs.northwestern.edu
  - Virtuoso: Resource Management and Prediction for Distributed Computing using Virtual Machines
    - http://virtuoso.cs.northwestern.edu
    - VNET is publicly available from above URL