Passive End-to-End Packet Loss Estimation for Grid Traffic Monitoring

Antonis Papadogiannakis, Alexandros Kapravelos,
Michalis Polychronakis, Evangelos P. Markatos
Institute of Computer Science (ICS)
Foundation for Research and Technology Hellas (FORTH)

Augusto Ciuffoletti
INFN - CNAF

Roadmap

- Introduction
- Passive Packet Loss Measurement Characteristics
- Methodology
- Integration within a Network Monitoring Service
- Experimental Evaluation
- Conclusions

Why to Measure Packet Loss?

- Accurate network monitoring is vital for Grids
  - Resource allocation
  - Scheduling decisions
  - Performance debugging
- Packet loss is an important performance metric
  - Identify poor network conditions
  - Highly affects the TCP throughput and the overall end-to-end data transfer quality

Existing Measurement Tools

- Most existing tools for packet loss estimation use active probes
  - ping, zing, badabing, sting
  - Incur network overhead due to the injected packets
- Existing passive monitoring techniques are based on TCP's loss recovery algorithms
- A passive monitoring tool for packet loss estimation at the IP layer is still missing

Passive Packet Loss Measurement Characteristics (1/2)

- Non-intrusive
  - It does not inject any probe packets
- Real-time measurement of the actual loss ratio
- Scalability
  - Measure end-to-end packet loss between many different domains
- Per-application measurement
  - Differentiated services or rate limiting may result to different loss ratios in the same path
- IP-level measurement

Passive Packet Loss Measurement Characteristics (2/2)

- Limitations:
  - Requires two passive monitors at the ends of the measured path
  - Presence of real traffic in the path
- Can be used as complementary to existing active probing techniques

Our Approach

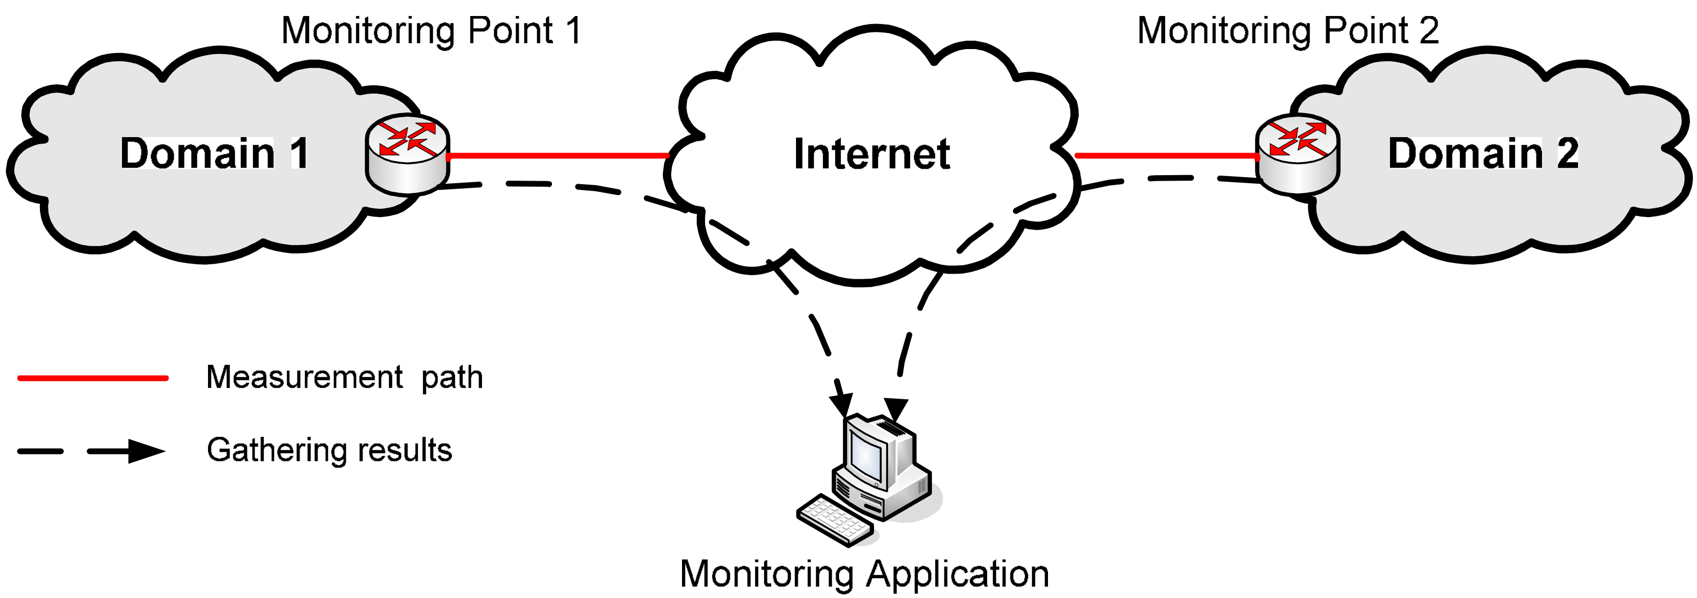

- Based on distributed passive network monitoring
- Two passive monitors at the two ends
- Send periodically information to a central application that computes the loss ratio

A Naive Algorithm

- Count the number of packets at both ends
- The application periodically subtracts the number of packets received from the number of packets that were actually sent

- Major drawback: inaccuracy
  - We cannot accurately synchronize the monitoring points to count the same window of packets
  - Packets in transit are not counted

Our Estimation Methodology (1/2)

- Measure the packet loss in each flow separately
- A flow is defined as a set of IP packets with the same 5-tuple:
  - Protocol
  - Source and destination IP address
  - Source and destination port (for UDP and TCP)
- An expired flow is a flow with no arriving packets for a specified timeout (e.g. 60 seconds)
- Expired flow is well defined: we know the first (e.g. TCP SYN) and the last (e.g. TCP FIN) packet of the flow

Our Estimation Methodology (2/2)

- Each monitoring sensor sends periodically statistical information about the expired flows it has seen
- The monitoring application correlates the statistics regarding the same expired flow
- The difference of the number of packers gives an accurate estimation of the loss ratio for this flow

Passive Monitoring Platform

- In each measurement point we run two daemons:
  - mapid: A passive monitoring daemon that identifies and collect the expired flows
  - mapicommd: A communication daemon that accepts monitoring requests and sends back the results
- Using a distributed monitoring API (DiMAPI) we manipulate multiple monitoring sensors from the same application

Identification of Expired Flows

- Every new packet is associated with exactly one active flow record
- Hashtable for fast lookup
- Linked list with temporal order for immediately identifying the expired flows

A Flow record

Source IP address	Destination IP address
Source port	Destination port
Protocol	
Timestamp of first packet	Timestamp of last packet
Packet counter	Byte counter

Distributed Sensor Management

- The monitoring application communicates with many distributed sensors
- It periodically collects the expired flows from them
- Then, it correlates the pairs of statistics regarding the same flow
- For every matched pair, it computes the packet and byte loss ratios
- It reports both the total loss ratio, and loss ratio per every individual flow

Integration within a Network Monitoring Service (1/2)

- Network Monitoring Element (NME)
  - Offers an interface for measurement requests
  - Plug-in based interface for publishing measurements
  - Access a database that contains information about Grid resources and other NME
- Network Monitoring session for packet loss ratio measurement
  - Identify the source and destination domains
  - Type of service that the measurement corresponds
  - Time period of the measurement (historical, most recent, one-shot, or periodic)

Integration within a Network Monitoring Service (2/2)

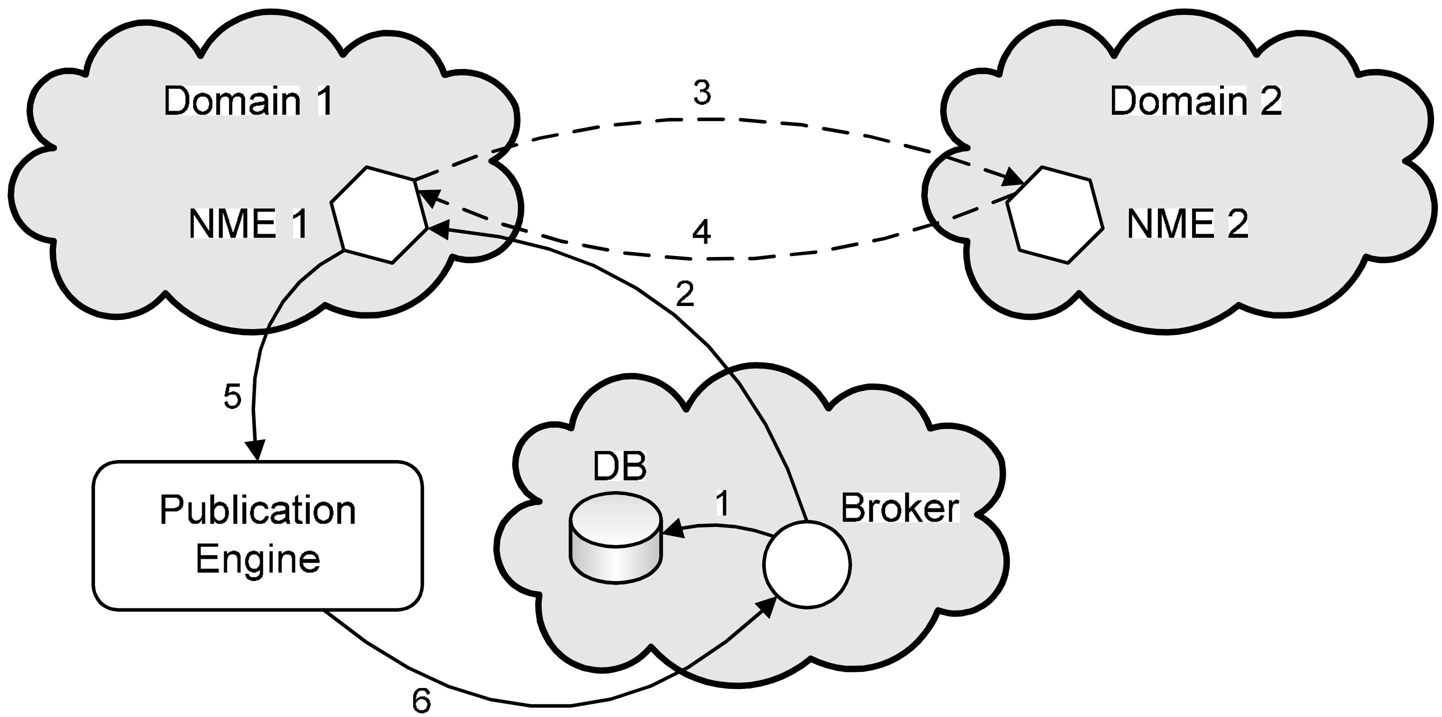

Experimental Evaluation

- Packet Loss Measurement Accuracy
- Experiences with Grid Network Traffic

Artificially generated packet loss

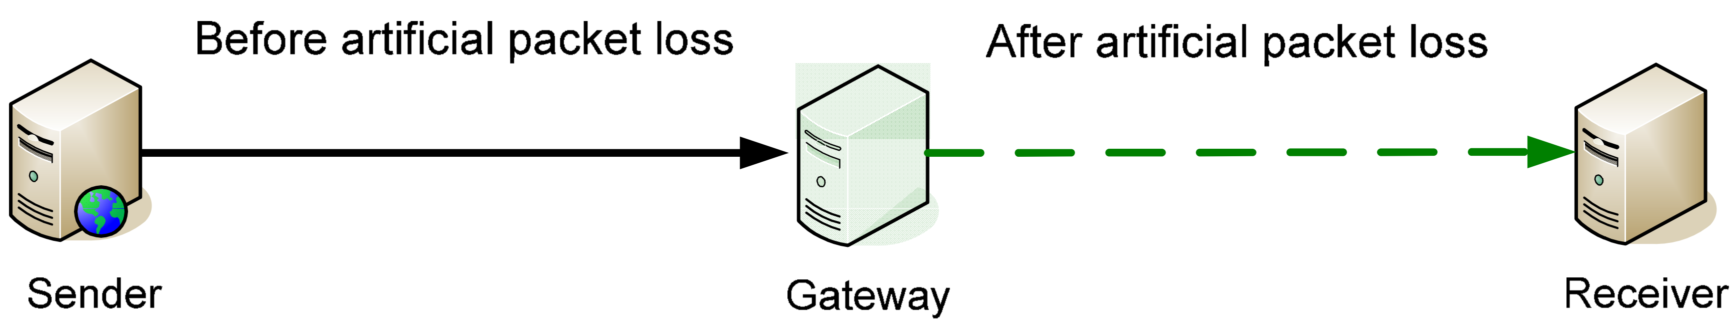

Packet Loss Measurement Accuracy (1/2)

Artificially Generated Loss (%)	Estimated Loss (min/avg/max %)	Measurement Error (%)
0	0.00/0.002/0.01	0.002
1	0.91/0.98/1.06	0.020
5	4.80/5.014/5.13	0.014
10	9.86/10.09/10.18	0.090
25	22.24/24.74/25.32	0.260

- Packet loss measured while generating 10 parallel UDP flows over a 2 hour period
- Accurate results, very close to the generated loss
- Small aberrations due to the probabilistic nature of loss generation

Packet Loss Measurement Accuracy (2/2)

Artificially Generated Loss (%)	Estimated Loss (min/avg/max %)	Measurement Error (%)	Served Requests	Rate (Mbit/s)
0	0.00/0.06/0.17	0.060	2944	38.75
1	1.02/1.078/1.16	0.078	1666	22.23
5	4.92/5.07/5.23	0.070	1058	14.11
10	9.86/10.086/10.12	0.086	290	3.90
25	24.89/25.235/25.50	0.235	0	0.26

- Packet loss measured while performing normal HTTP requests
- Packet loss ratio significantly affects the number of completed requests
- TCP throughput drops dramatically from 38.75 Mbit/s to 0.26 Mbit/s

Experiences with Grid Network Traffic (1/2)

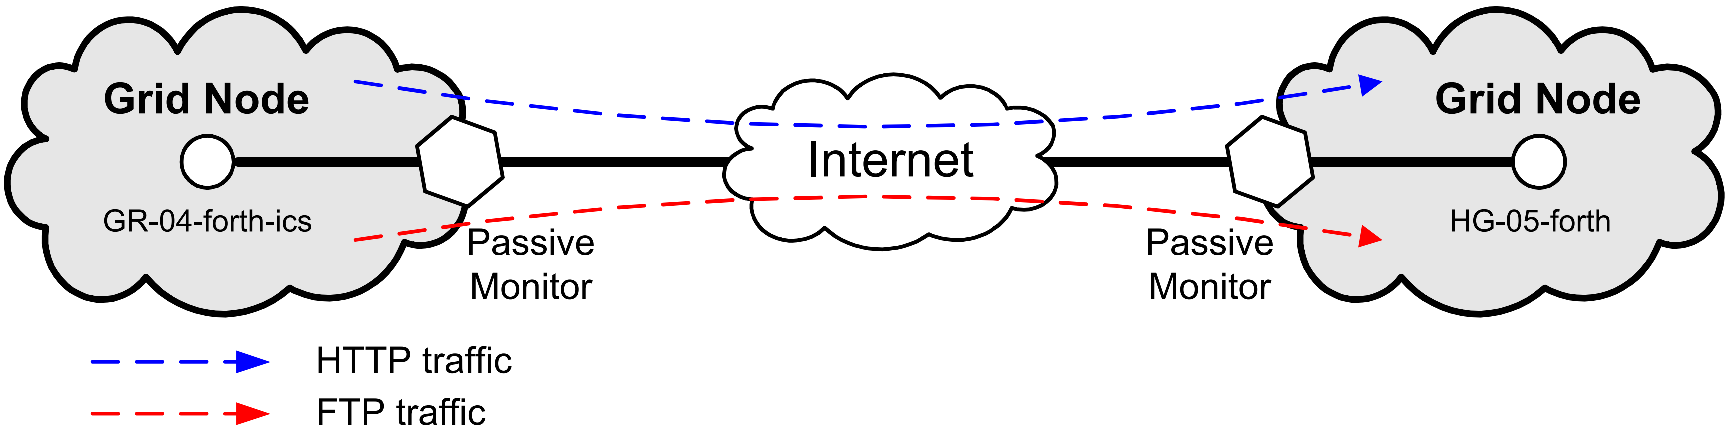

- Deployment of the technique to an operational Grid network path
- Running for a 24-hour period with measurements every 30 seconds
- Generating more traffic using HTTP and GridFTP

Experiences with Grid Network Traffic (2/2)

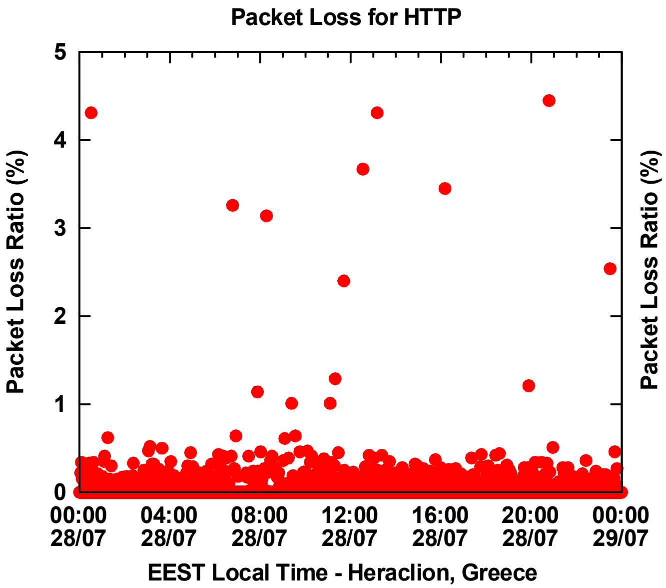

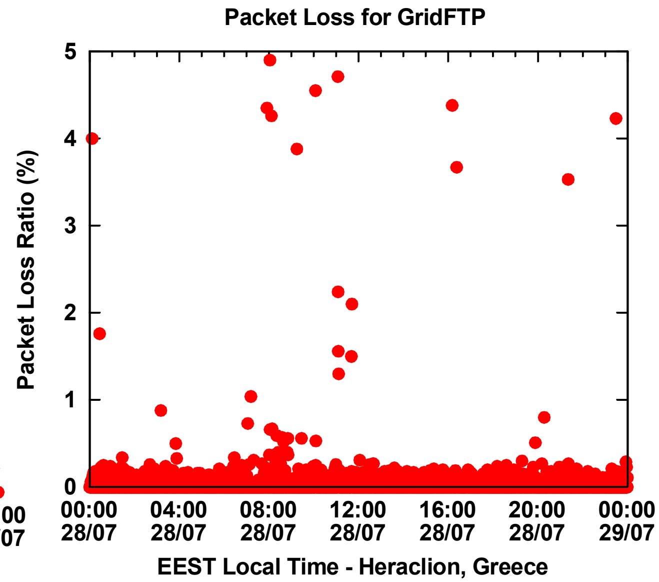

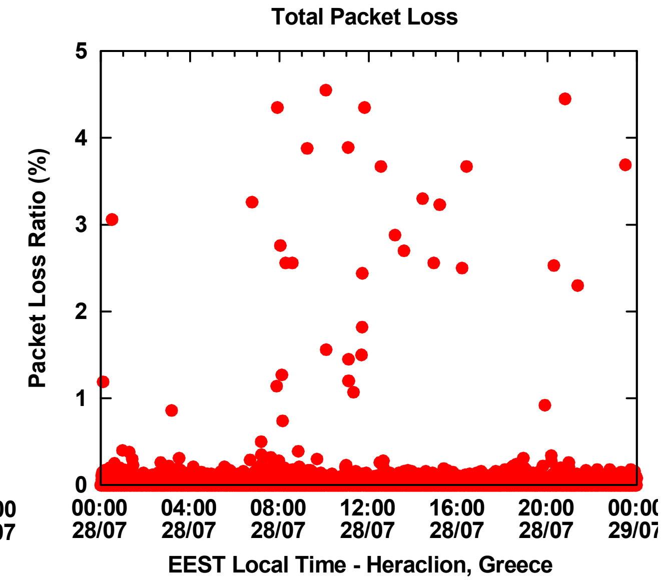

- Bursts of HTTP and FTP transfers result to higher loss rates
- The 83% of the 30-seconds intervals indicate 0% loss ratio (89% and 87% for HTTP and FTP)
- 0.09% overall loss ratio over the 24-hour period for total traffic, 0.13% for HTTP-only and 0.19% for FTP-only traffic

Conclusions

- A novel passive monitoring technique for packet loss estimation between different Grid domains
- Scalable, non-intrusive and real time
- Can be complementary to active monitoring tools
- Based on tracking the expired flows at each monitoring sensor
- Uses a distributed infrastructure for gathering and correlating the results
- Validated using realistic traffic

Thanks!

Any questions?