

ITU-T Workshop on IP Traffic Flow Measurement

(Geneva, Switzerland, 24 March 2011)

Introduction to IP Traffic Flow Measurements and Packet Sampling

Dr.-Ing. Tanja Zseby

Fraunhofer Institute for Open Communication
Systems (Fraunhofer FOKUS)
Berlin, Germany

Geneva, 24 March 2011



Overview

- Motivation
- IP Packet Observation
- Limitations
- Aggregation and Selection
 - Aggregation: IP Flow Measurements
 - Selection: Filtering and Sampling
- Multipoint Measurements
- IP Flow Information Export
- Summary

Motivation

- Accounting
- Security
- SLA Validation
- Fault Detection
- Traffic Engineering
- Traffic Profiling
- Research (Experiment Supervision)

Packet and Flow Measurements

Observation of IP Packets

Packet Attributes at Observation point

O_A – Observation point

c_i – Packet content (header and payload)

t_i – Arrival time

P_i – IP Packet Record

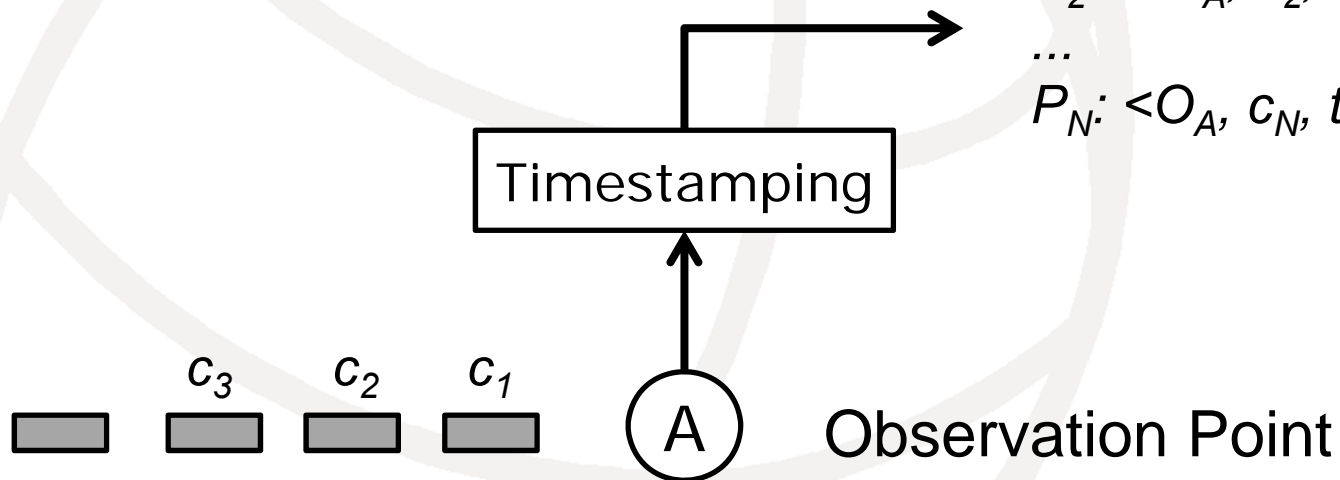
IP Packet Records

$P_1: \langle O_A, c_1, t_1 \rangle$

$P_2: \langle O_A, c_2, t_2 \rangle$

...

$P_N: \langle O_A, c_N, t_N \rangle$



Calculation of Metrics

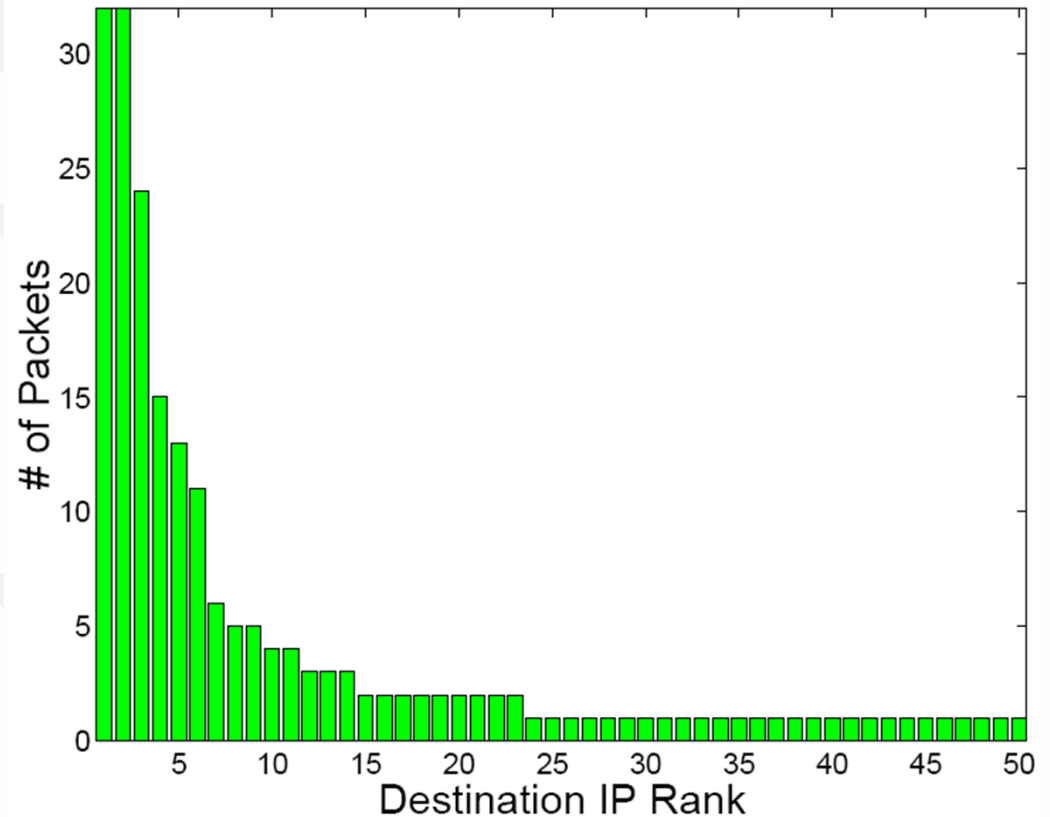

IP Packet Records

$P_1: \langle O_A, c_1, t_1 \rangle$

$P_2: \langle O_A, c_2, t_2 \rangle$

...

$P_N: \langle O_A, c_N, t_N \rangle$



Quelle: [LaCD05] Lakhina, Crovella, Diot: Mining Anomalies Using Traffic Feature Distributions, *SIGCOMM* 2005.

Problem: Limited Resources

Specialized
Hardware

Aggregation

Processing

Limited
Resources

Storage

Transport

Improved
Algorithms

Selection

Capturing/processing all packets → too high effort

Limitations

■ Capturing

- On router → capturing competes with routing tasks
- Specialized hardware → expensive, multiple devices needed

■ Storage

- E.g. flow cache on routers

■ Data export

- Transmission capacity
- Effort to support reliable transport

Specialized Hardware

- Specialized capture cards
 - 10 Gbit Ethernet, full line rate capturing
 - High precision time stamping
- Examples
 - Endace DAG 9.2X2 (March 2010)
 - Napatec NT20E Capture
 - NTT Advanced Technology PRESTA 10G
- But: Expensive

Aggregation: IP Flows

IP Flow: set of IP packets with common properties [RFC5101]

IP Flow Records

Flow 1: $\langle N_1, \mu_1, \sigma_1, \dots \rangle$

Flow 2: $\langle N_2, \mu_2, \sigma_2, \dots \rangle$

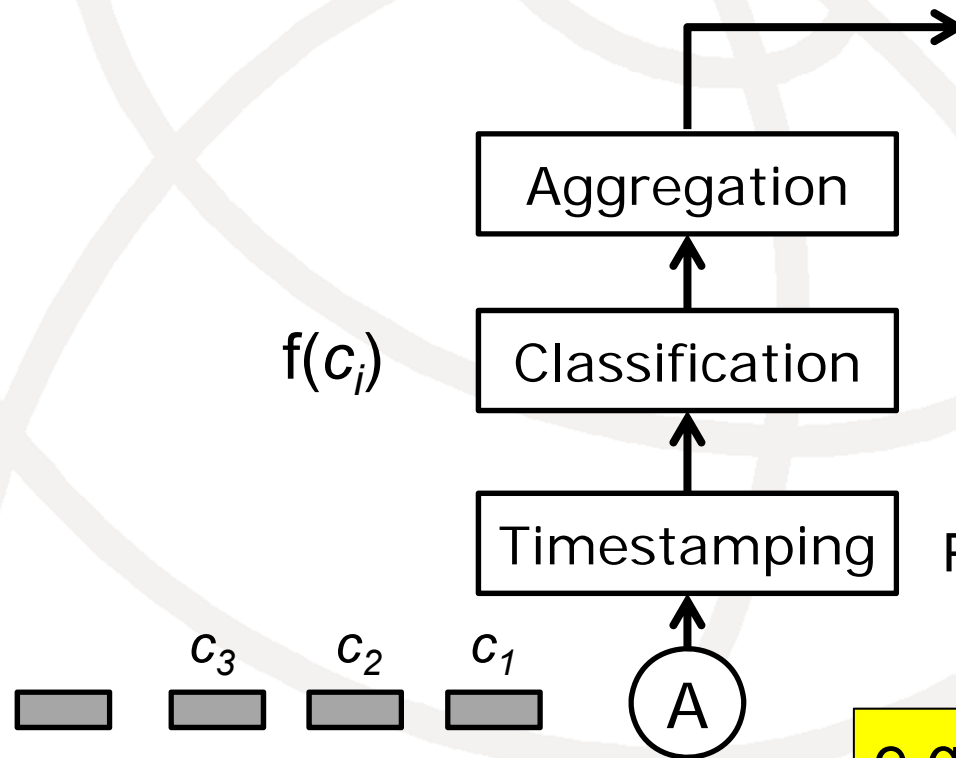
Flow 3: $\langle N_3, \mu_3, \sigma_3, \dots \rangle$

Flow 1: P_1, P_4, P_6, \dots

Flow 2: P_2, P_3, P_8, \dots

Flow 3: P_5, P_7, \dots

Packet Records: P_1, P_2, P_3, \dots



e.g. Cisco NetFlow, IPFIX

Calculation of Metrics




e.g. classification according to IP addresses

IP Flow Records

Flow 1: $\langle N_1, \mu_1, \sigma_1, \dots \rangle$

Flow 2: $\langle N_2, \mu_2, \sigma_2, \dots \rangle$

Flow 3: $\langle N_2, \mu_2, \sigma_2, \dots \rangle$




Entropy
 $H(X)$

Calculation of metrics possible
if suitable classification (flow characteristics)

Example: Usage-based Accounting

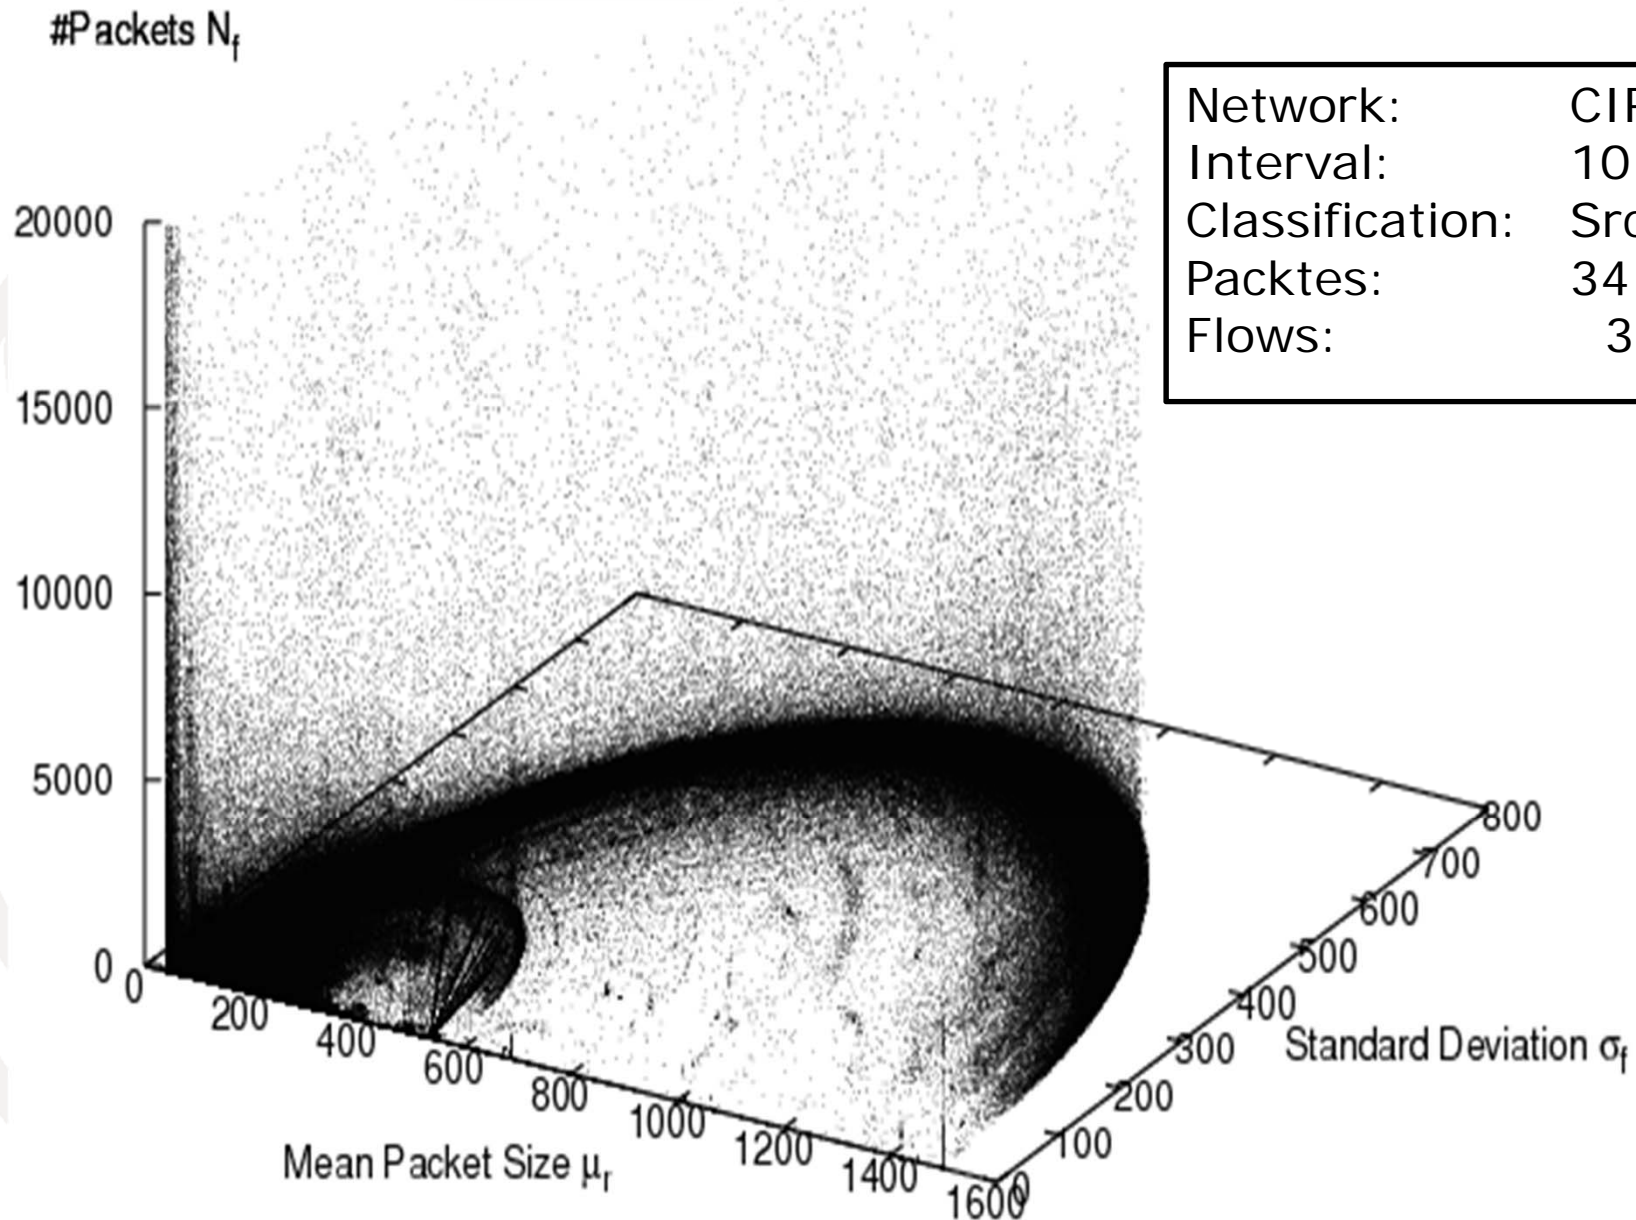
Flow Measurements



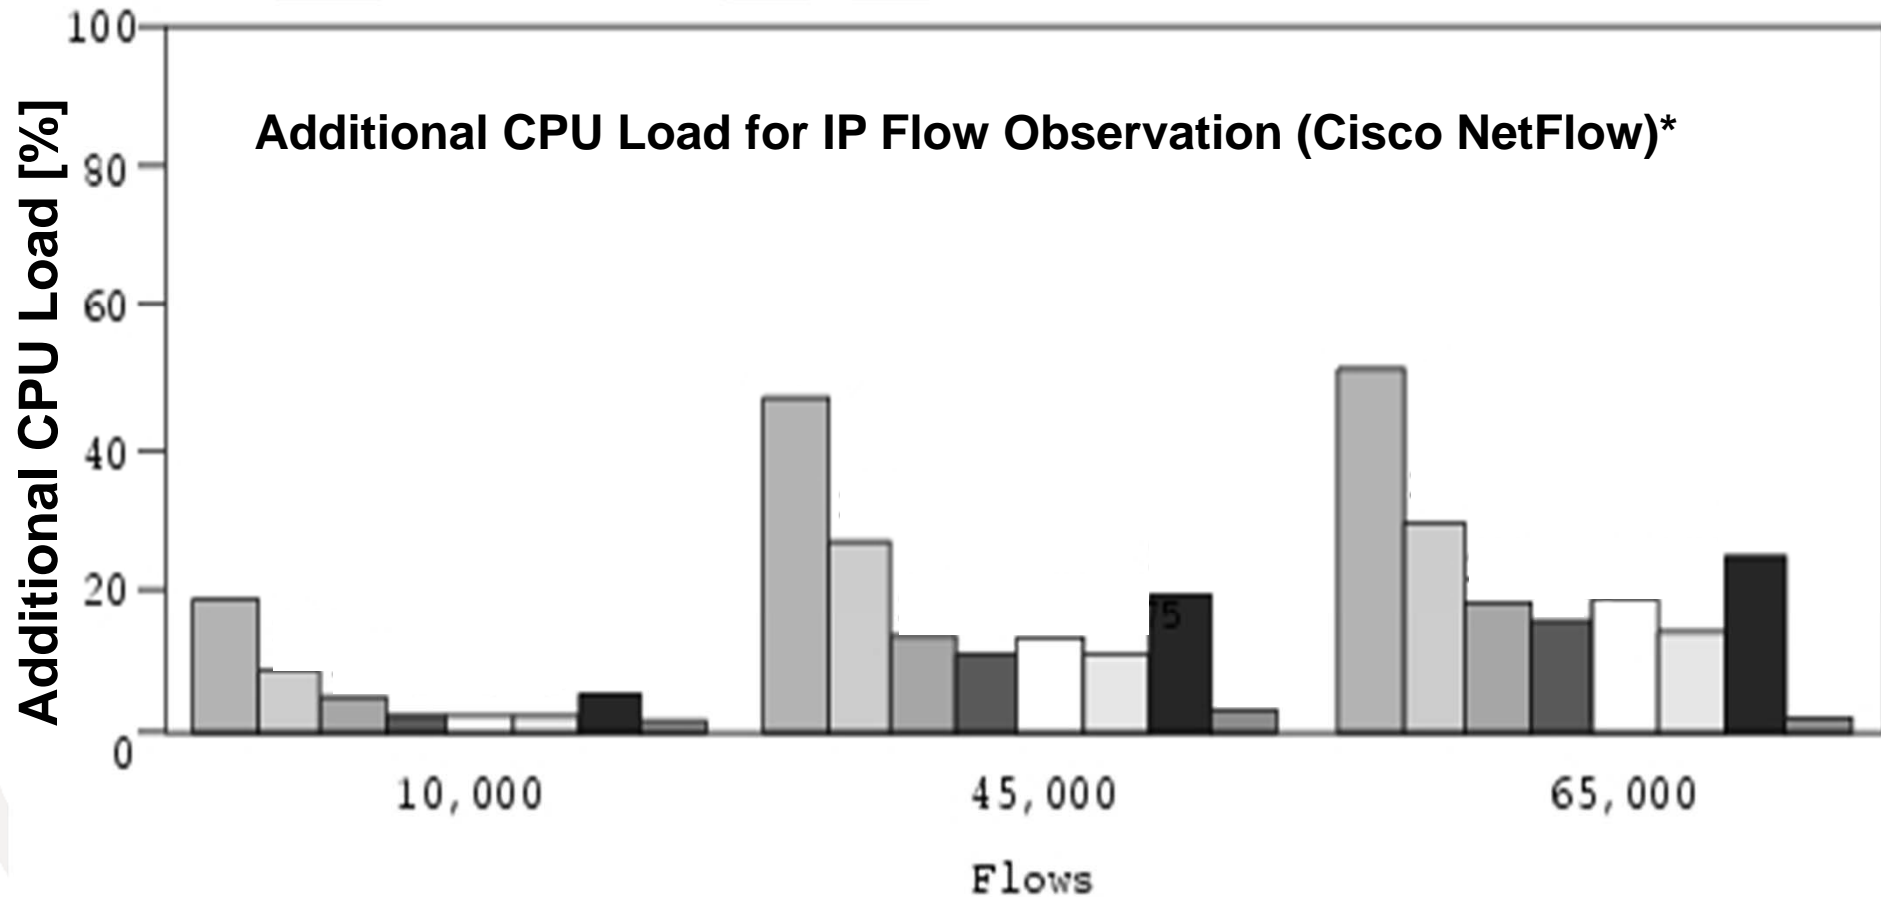
Example: Usage-based accounting

- Flow classification
 - ➔ Based on source network addresses
 - ➔ Dependent on tariff model
- Generation of Accounting records
 - ➔ Capture packet sizes
 - ➔ Calculate number of bytes per flow
- Today: Cisco NetFlow
 - ➔ Flow classification based on flow keys

Example: IP Flow Observation



IP Flow Observation on Routers



Still high resource consumption

*Quelle: NetFlow Performance Analysis, Cisco Whitepaper, May 2007

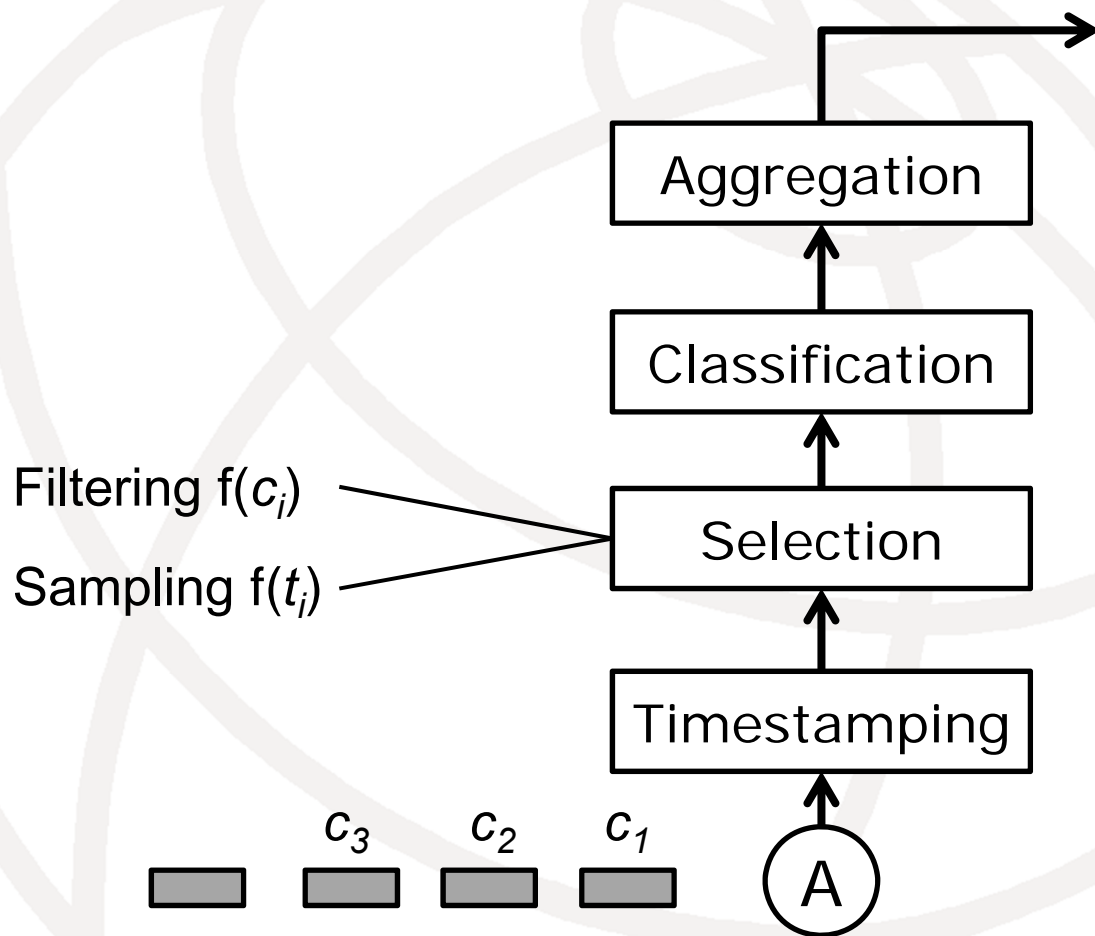
Packet Selection

IP Flow Records

Flow 1: $\langle \hat{N}_1, \hat{\mu}_1, \hat{\sigma}_1, \dots \rangle$

Flow 2: $\langle \hat{N}_2, \hat{\mu}_2, \hat{\sigma}_2, \dots \rangle$

Flow 3: $\langle \hat{N}_3, \hat{\mu}_3, \hat{\sigma}_3, \dots \rangle$



$P_1, \cancel{P_2}, P_3, \cancel{P_4}, \cancel{P_5}, P_6, P_7, \dots$

Exact value substituted by estimate

Estimation Accuracy

Selections- and
estimation method

Selection parameter
(Interval, size, etc.)

Characteristics of
parent population

Estimation accuracy

highly dynamic
hard to predict

Calculation and transmission of estimation accuracy

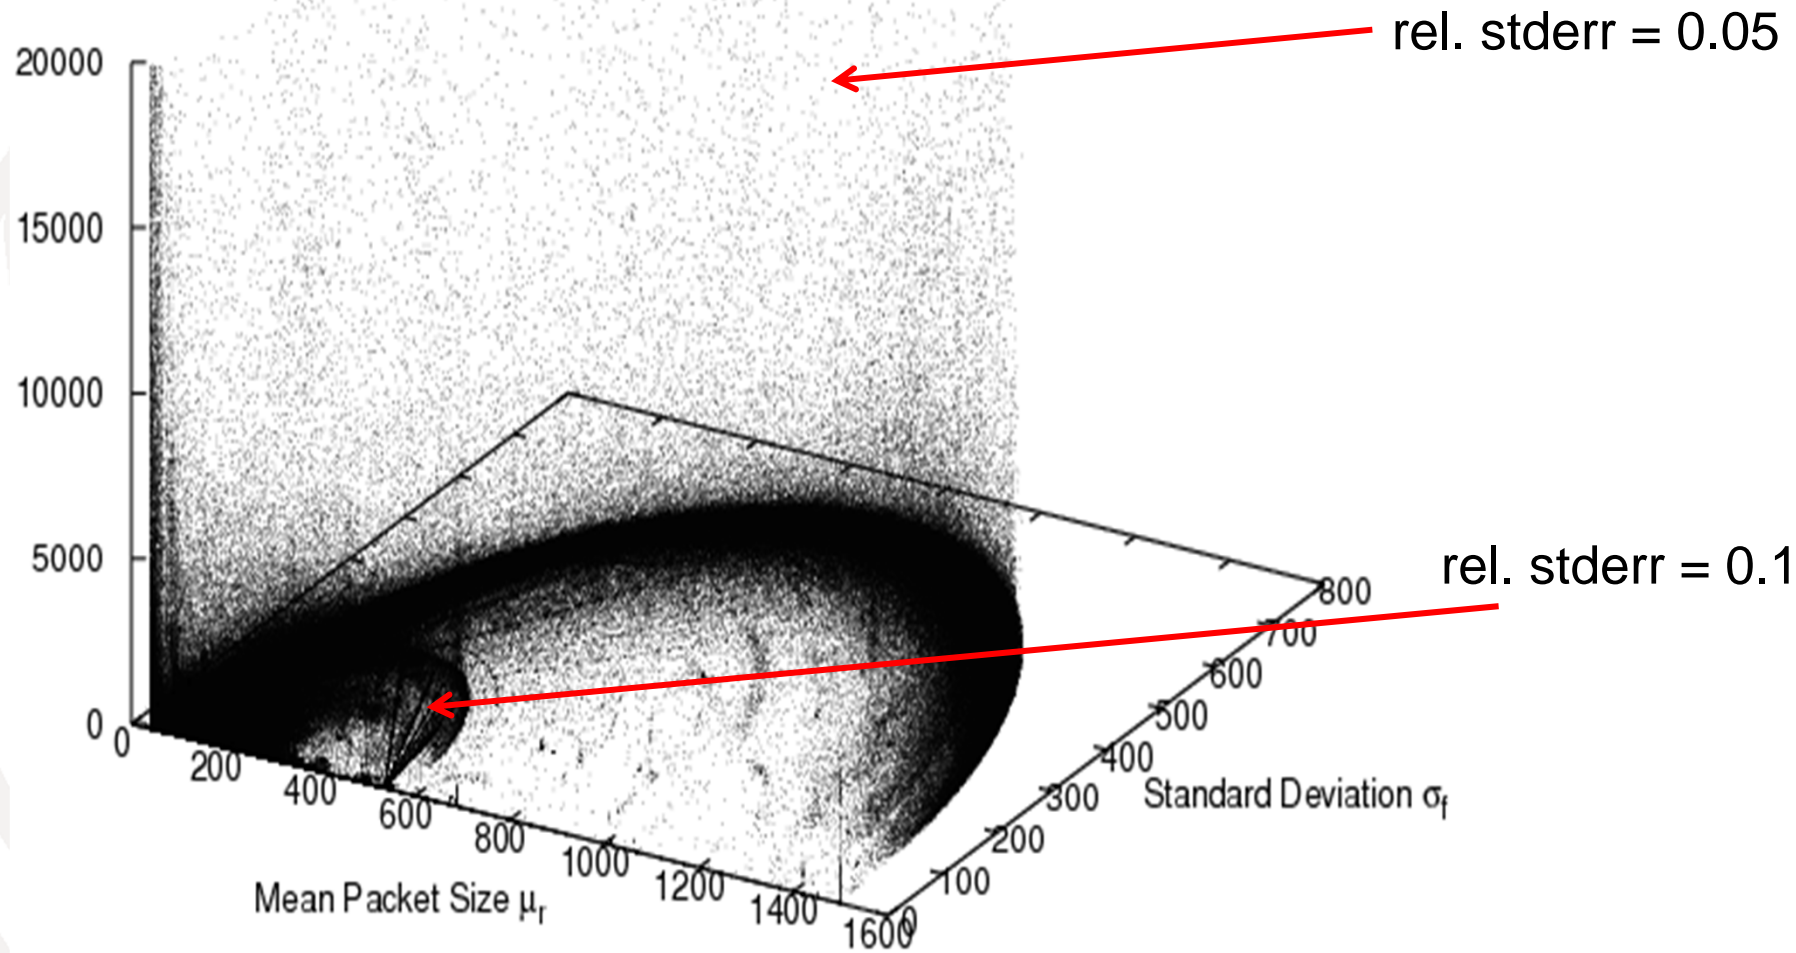
During the
measurement

Per interval
Per IP Flow

Based on
sample

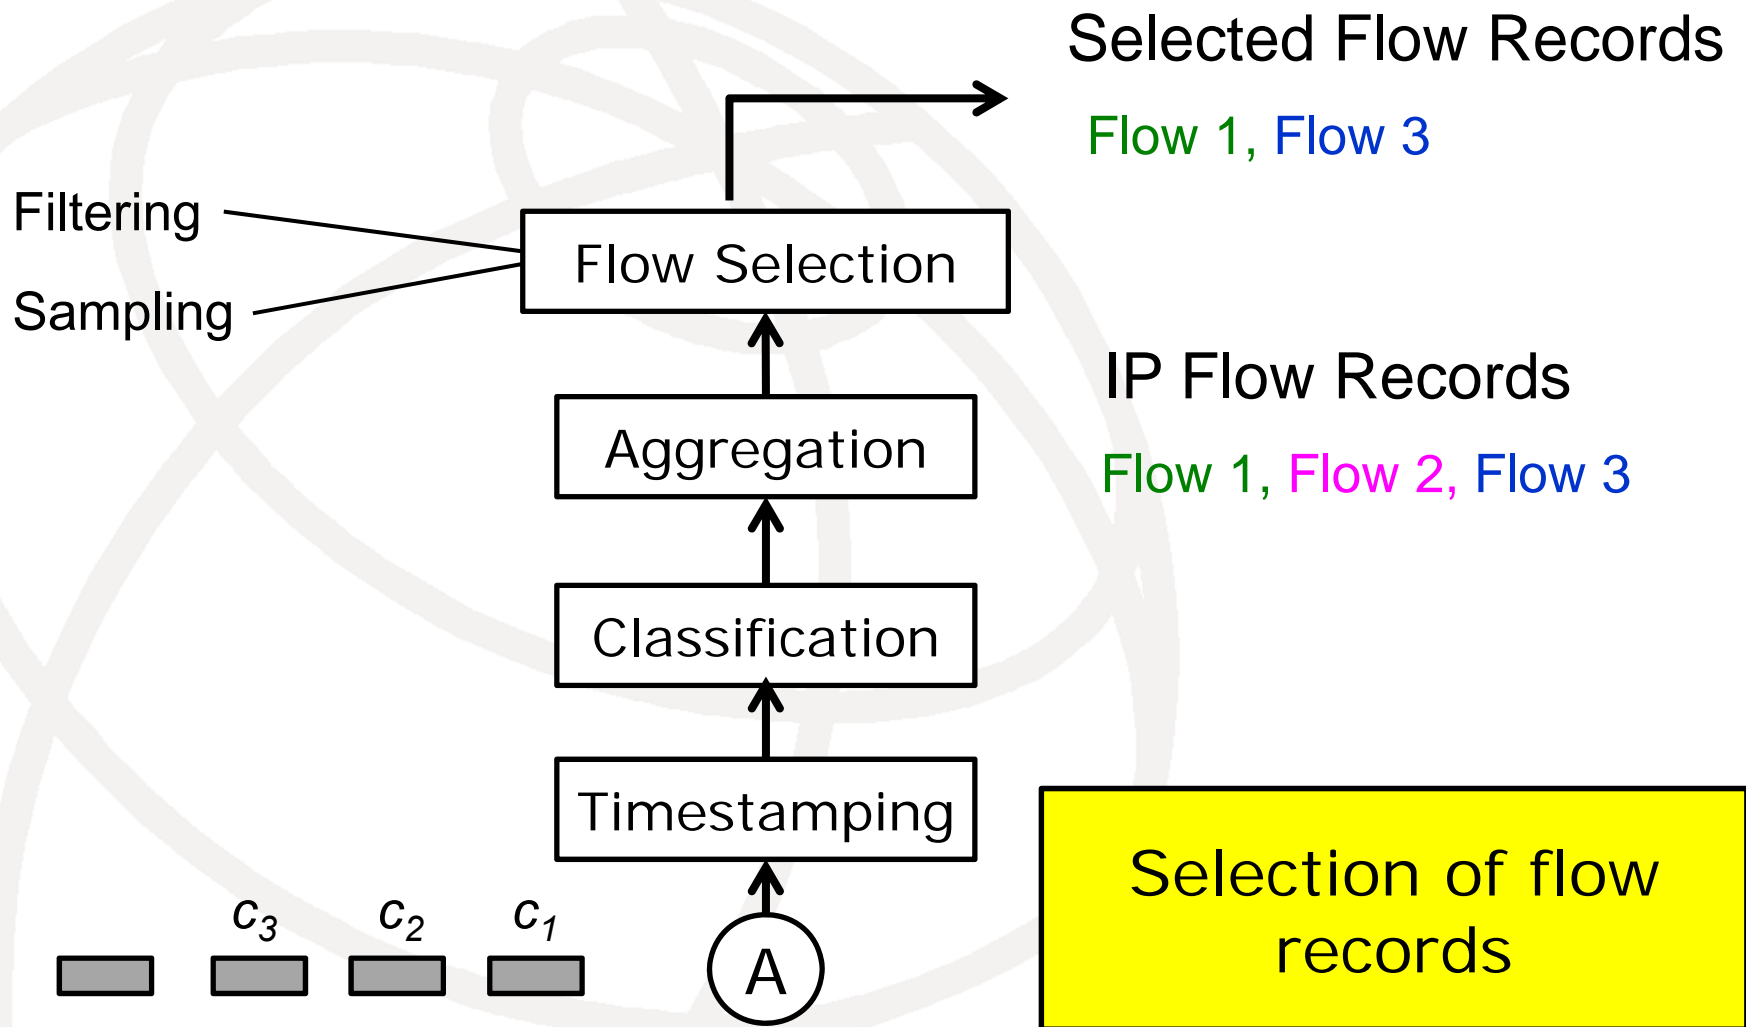
Example: Estimation of Flow Volume

#Packets N_f



Additional Value per Flow

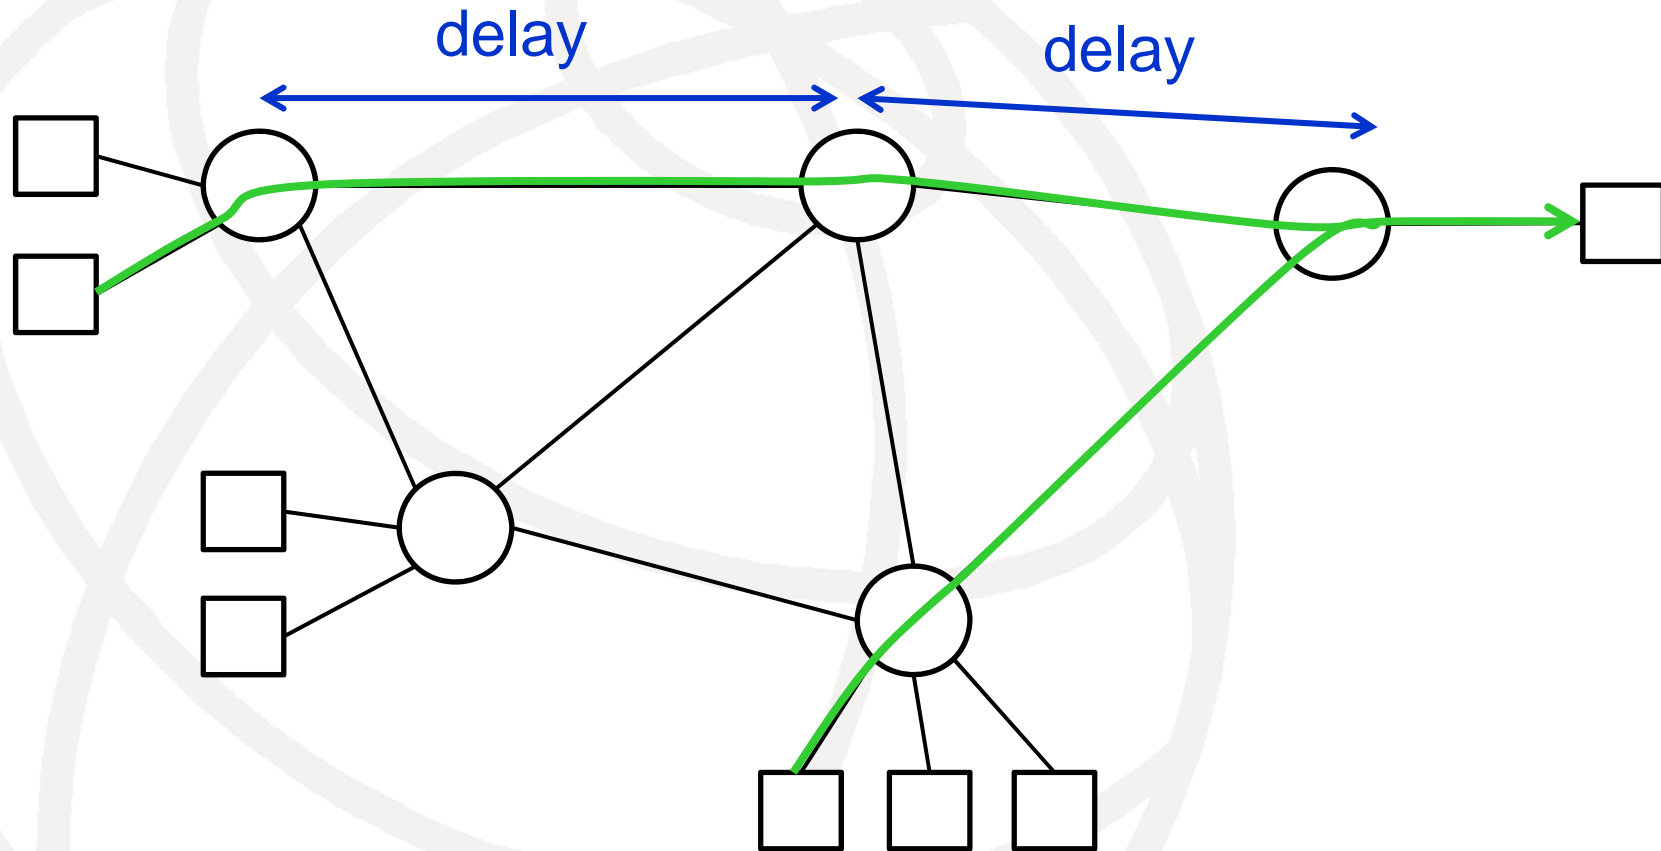
Flow Selection



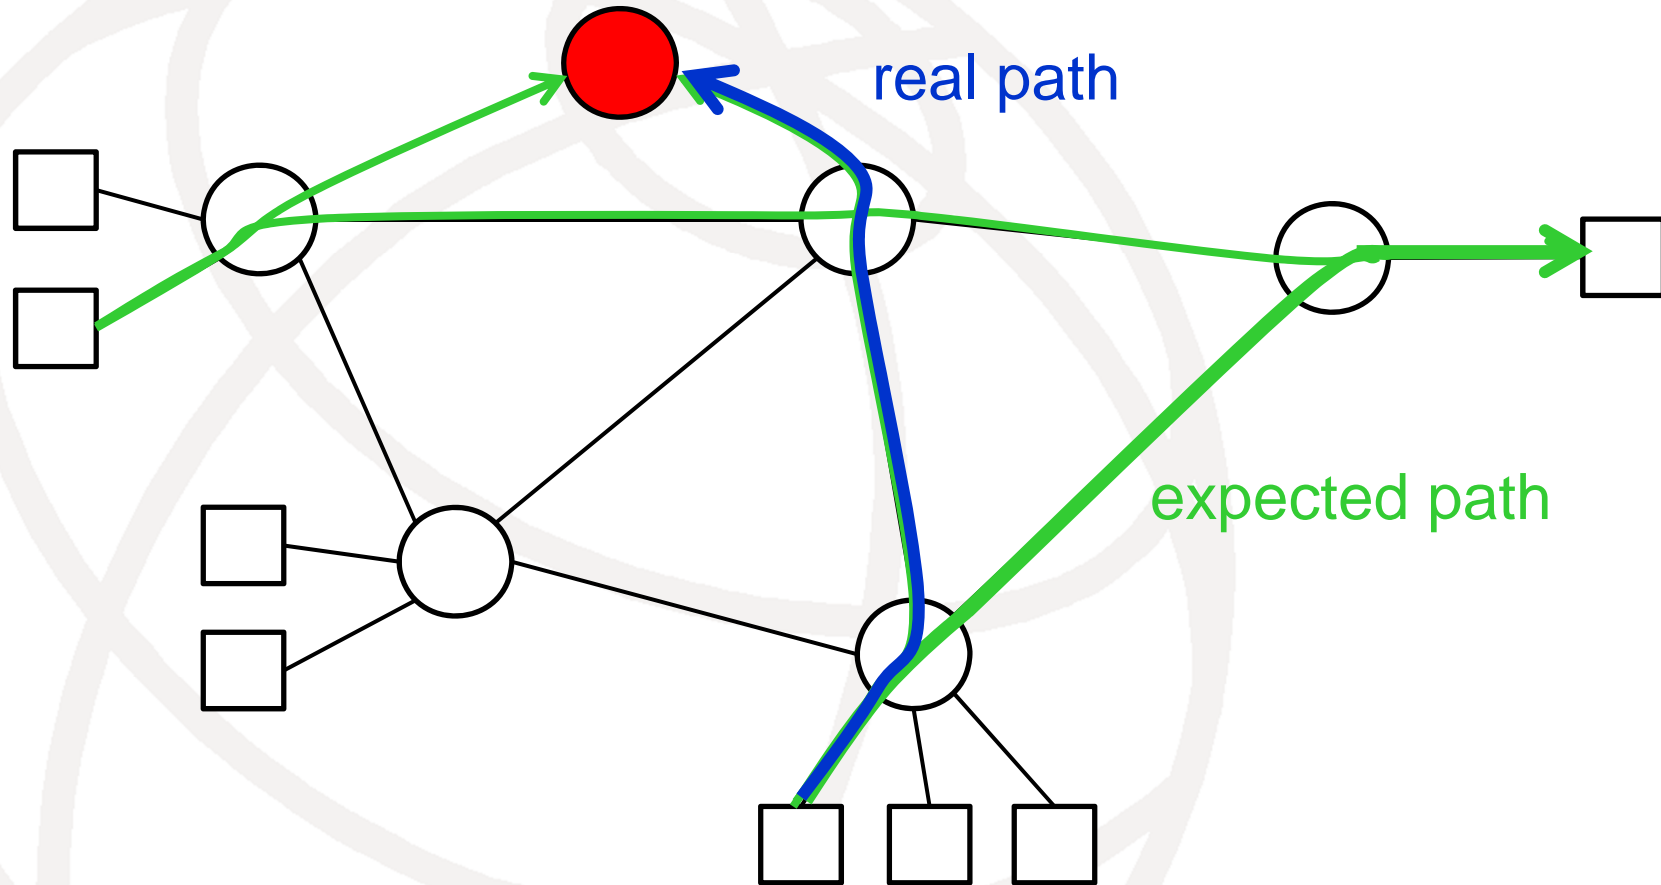
Multipoint Measurements

- Metrics: Path, one-way delay, loss
- But Challenges:
 - Positioning of Measurement Points
 - Inter-domain Observations
 - Cooperation of network operators
 - Exchange of data, data protection
 - Synchronization of measurement processes
 - Time synchronization → NTP, GPS,...
 - Synchronization of data selection

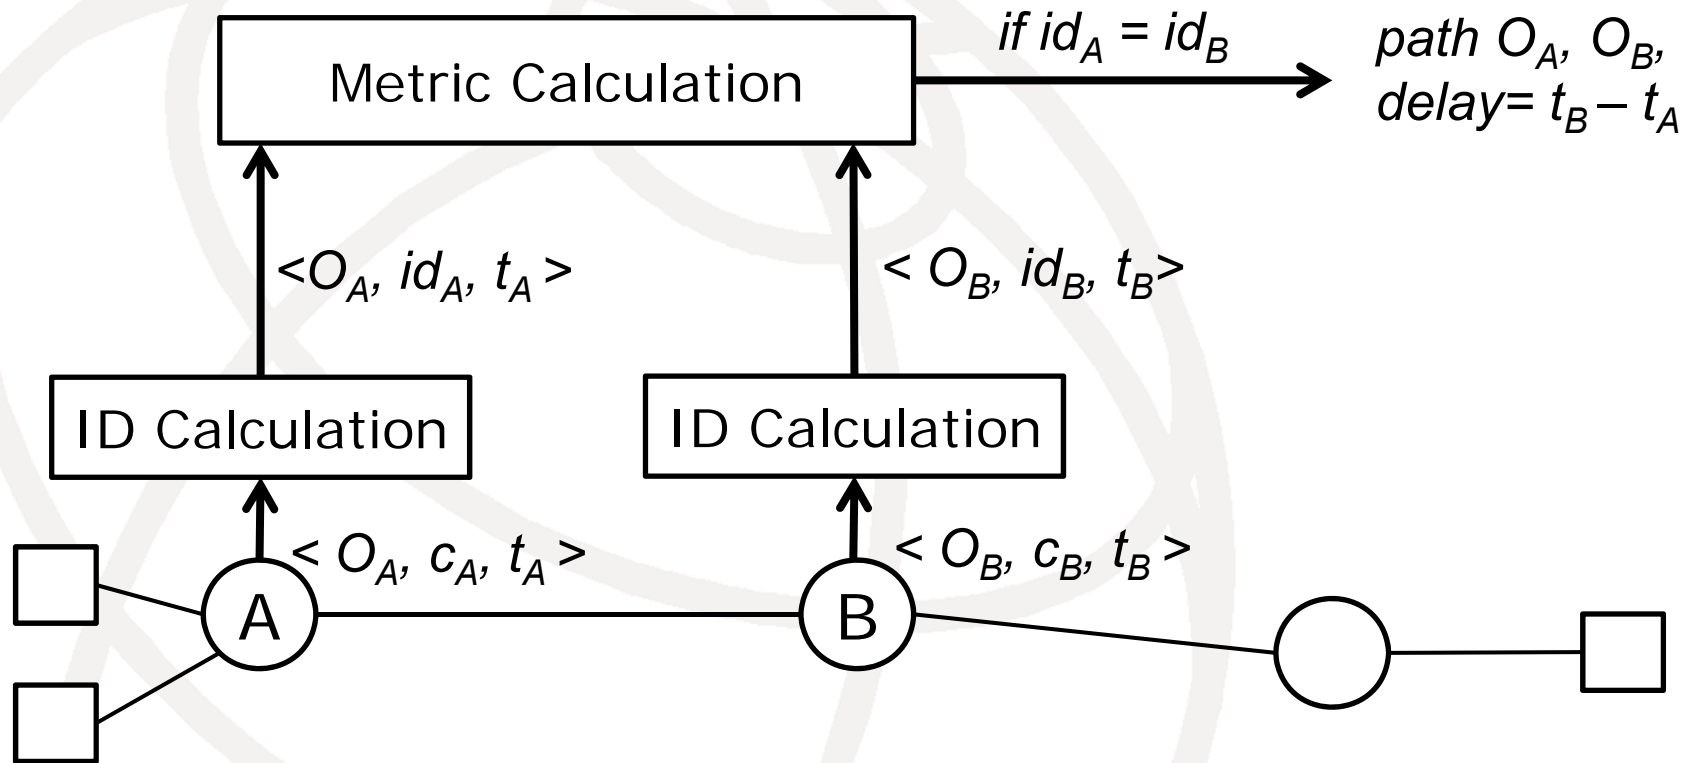
Example: SLA Validation



Example: Attack Detection

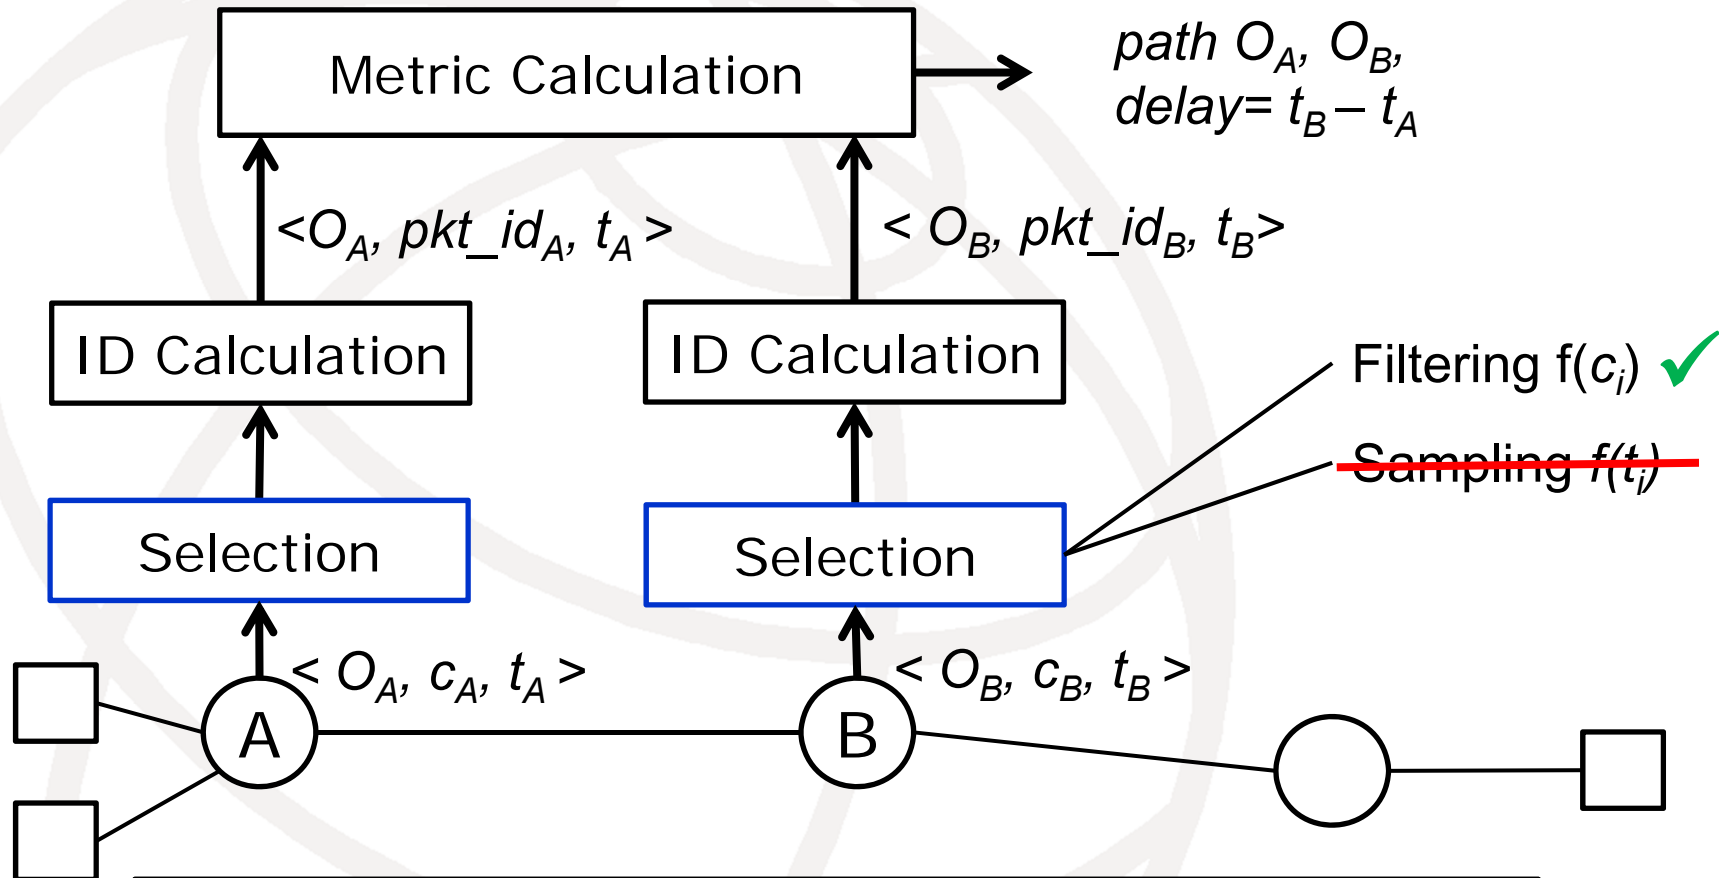


Multipoint Measurements



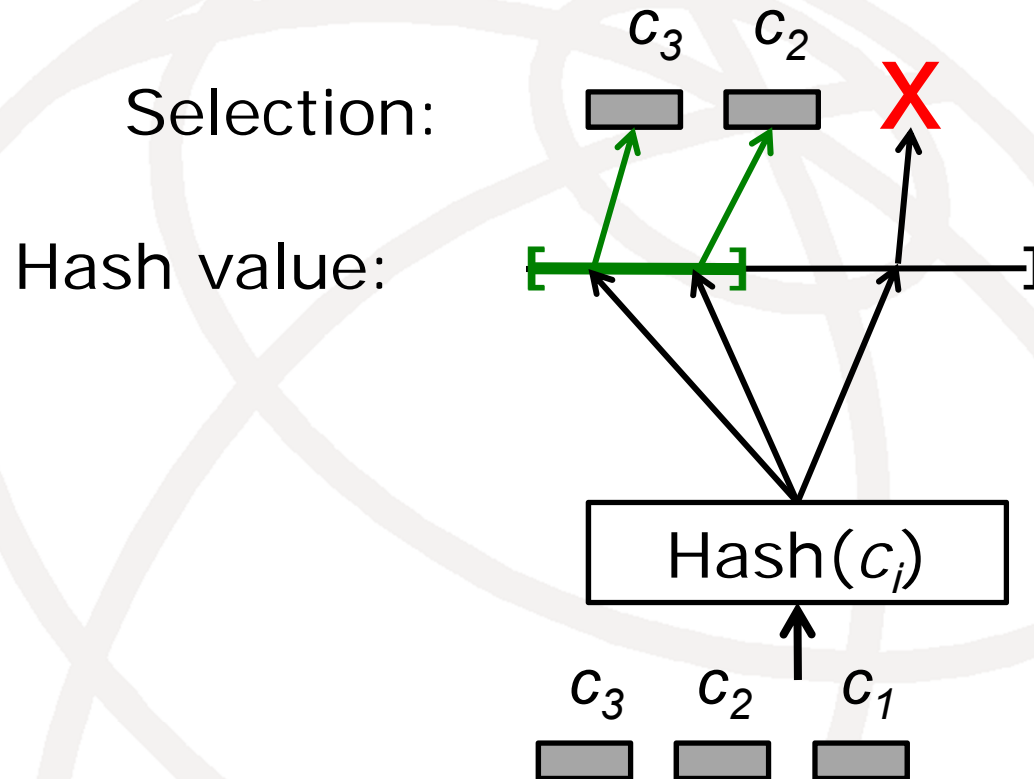
Packet ID: recognize packet at different observation points

Multipoint Packet Selection



Goal: Ensure selection of same packet at all observation points

Hash-based Selection [RFC5475]



Duffield, Grossglauser: "Trajectory Sampling for Direct Traffic Observation", IEEE/ACM Transactions on Networking, vol. 9, 2001

[RFC 5475] Zseby, Molina, Duffield, Niccolini, Raspall. Sampling and Filtering Techniques for IP Packet Selection, RFC 5475, Standards Track, March 2009.

Challenges

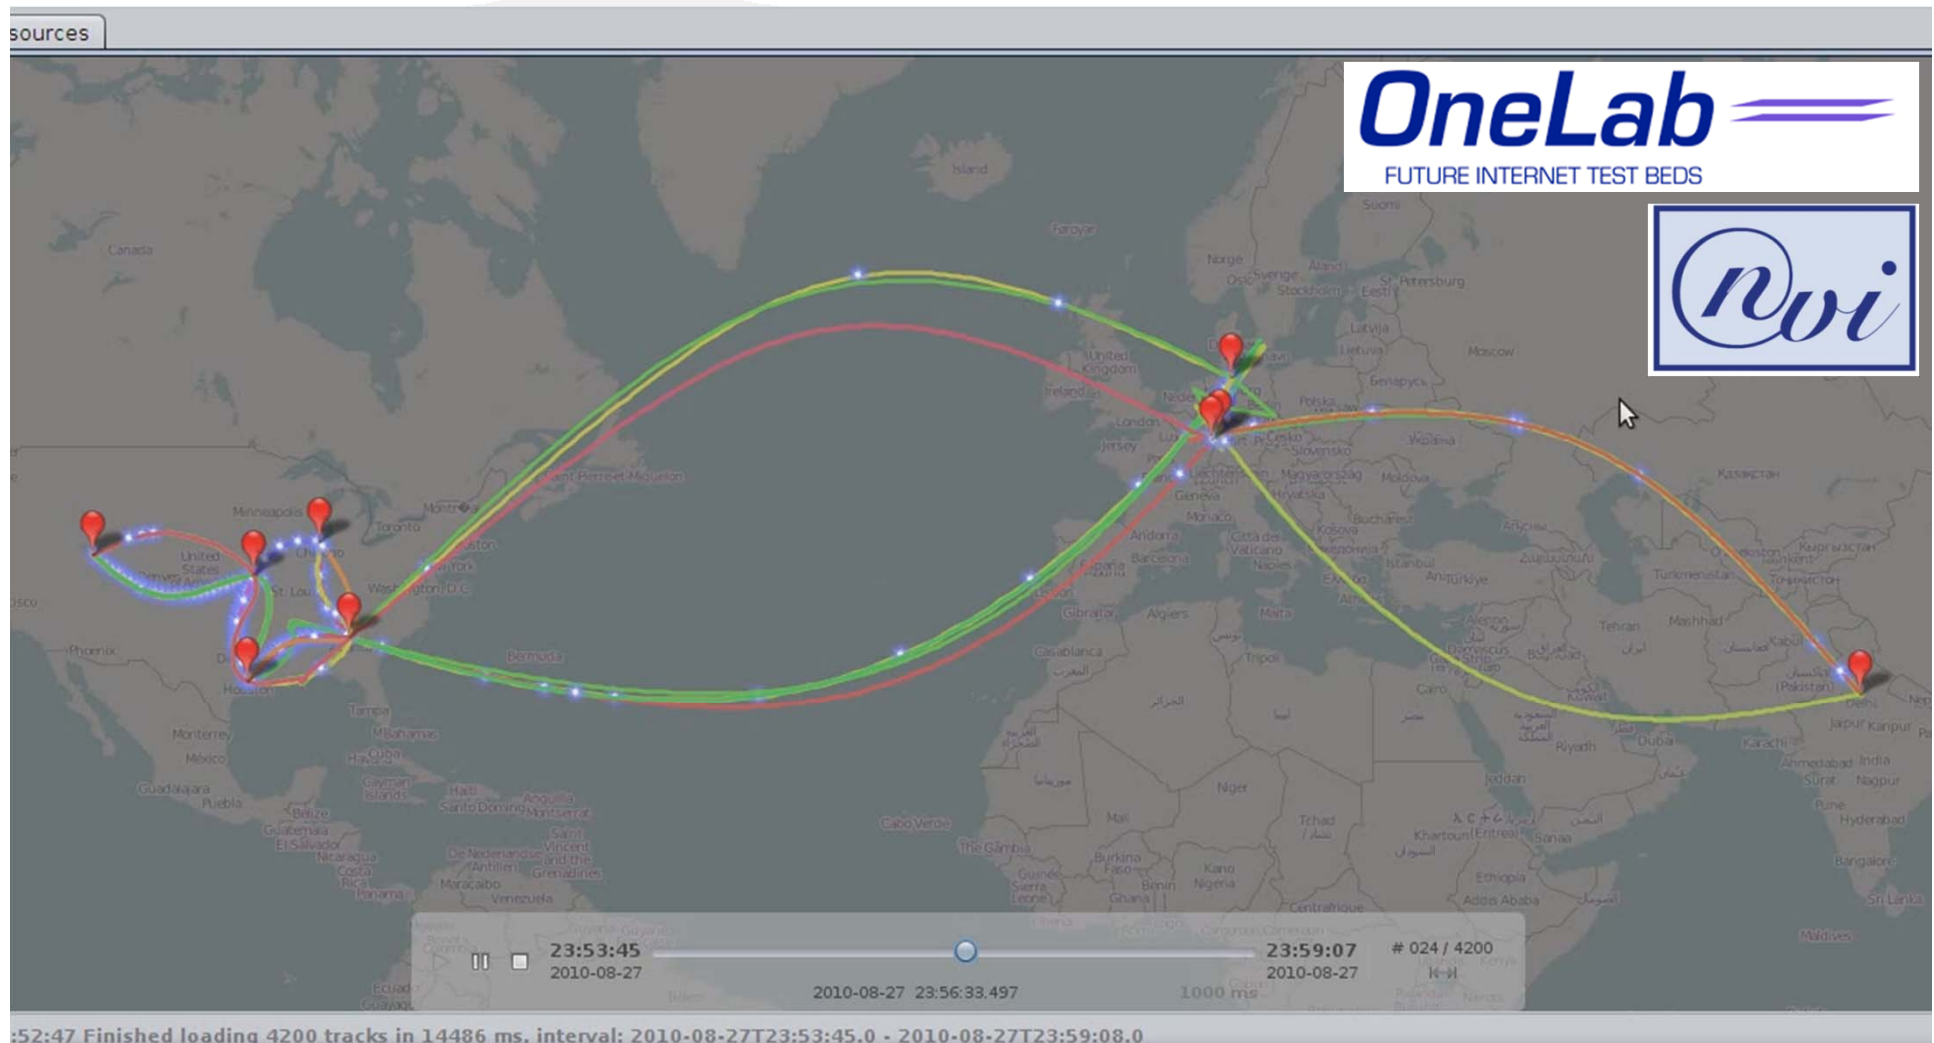
- Hash input: suitable header fields
 - Invariant on the path
 - Variable between packets
- Suitable hash function*
 - Performance
 - Representativeness of selection
- Dynamic adaptation of selection rates
 - Configured vs. attained selection rate
 - Coordinated adaptation to available resources

* [HeSZ08] Henke, Schmoll, Zseby: Empirical Evaluation of Hash Functions for Multipoint Measurements, ACM Comput. Commun. Rev. CCR 38, 3, July 2008.

Packet Tracking

- Packet Tracking Software
 - ➔ Follow the path that a packet takes
 - ➔ Based on hash-based selection
- Deployed in PlanetLab Europe
 - ➔ Service for PlanetLab users (researcher)
 - ➔ Capture path of packets in network experimenters
 - ➔ Used for multipath routing, overlay experiments

Packet Tracking in Federated Environment



Demonstration at SIGCOMM 2010

IP Flow Information Export (IPFIX)

- Future protocol for data export
 - „Successor“ of Cisco NetFlow
 - Supports packet and flow measurements
 - Allows flexible flow definitions
 - Integration of data selection methods possible
- Information Elements (IE)
 - Flow information in IEs (e.g. #packets, bytes)
 - Many information elements exists (see www.iana.org/assignments/ipfix/ipfix.xml)
 - Vendor-specific IEs possible

IPFIX Information Elements

- Example: Usage based accounting
 - IE: sourceIPv4Address (source IP)
 - IE: destinationIPv4Address (destination IP)
 - IE: ipDiffServCodePoint (DiffServ class)
 - IE: octetDeltaCount (#octets in the Flow)

- Further Examples
 - IPFIX Applicability Statement (RFC 5472)

Summary

- Many applications require measurements
- Problem: Limited resources
- Solution: Aggregation and Data Selection
 - Aggregation: IP Flow Measurements
 - Selection: Filtering and Sampling
- Multipoint Measurements
 - Measurement of path, quality (loss, delay)
 - Synchronization needed
- IPFIX Standard
 - flexible flow and packet reporting



Thank You!

tanja.zseby@fokus.fraunhofer.de