Master Course
Computer Networks
IN2097

Prof. Dr.-Ing. Georg Carle
Christian Grothoff, Ph.D.
Stephan Günther

Chair for Network Architectures and Services
Department of Computer Science
Technische Universität München
http://www.net.in.tum.de

Chapter:
Network Measurements

Acknowledgements:
Anja Feldmann, Constantine Dovrolis

Chapter: Network Measurements

- Goals
- Architecture & Mechanisms
- Protocols
  - IPFIX (Netflow Accounting)
  - PSAMP (Packet Sampling)
- Anomaly detection

Why do we measure the network?

- Network Provider View
  - Manage traffic
    - Predict future, model reality, plan network
    - Avoid bottlenecks in advance
  - Reduce cost
  - Accounting
- Client View
  - Get the best possible service
  - Check the service („Do I get what I've paid for?)
- Service Provider View
  - Get information about the client
  - Adjust service to demands
  - Reduce load on service
  - Accounting
- Security view
  - Detect malicious traffic, malicious hosts, malicious networks, …
- Researcher View
  - Performance evaluation (e.g., "could our new algorithm handle all this real-world traffic?")

But why should we do it at all?

- Do we really have to?
  - The network is well engineered
  - Well documented protocols, mechanisms, …
  - Everything built by humans ⇨ no unknowns
    (compare this to, e.g., physics: does higgs boson exist?)
  - In theory, we can know everything that is going on
    ⇨no need for measurements?

- But:
  - Moving target
    - Requirements change
    - Growth, usage, structure changes
  - Highly interactive system
  - Heterogeneity in all directions
  - The total is more than the sum of its pieces

- And: The network is built, driven and used by humans
  - Detection of errors, misconfigurations, flaws, failures, misuse, …

Network Measurements

- Active measurements
  - “intrusive”
  - Measurement traffic is generated and sent via the operational network.
    (Examples: ping, traceroute)
  - Advantages
    - Straightforward
    - Does not depend on existing traffic by active applications
    - Allows measurement of specific parts of the network
  - Disadvantages
    - Additional load
    - Network traffic is affected by the measurement
    - Measurements are influenced by (possibly varying) network load

Example: Packet pair probing

- Packet Pair (P-P) technique
  - c.f. van Jacobson, Keshav
- Send two equal-sized packets back-to-back
  - Packet size: L
  - Packet TX time at link i: L/C_i
- P-P dispersion = time interval between last bit of two packets
- Without any cross traffic, the dispersion at receiver is determined by bottleneck links (i.e., slowest link):

$$\Delta_{out} = \max\left(\Delta_{in}, \frac{L}{C_i}\right)$$

$$\Delta_R = \max_{i=1,\ldots,H}\left(\frac{L}{C_i}\right) = \frac{L}{C}$$

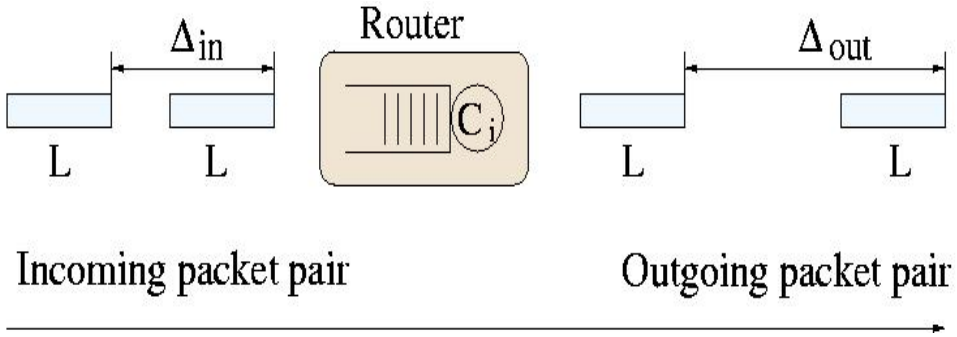

Network Measurements II

- Passive measurements (or **Network Monitoring**)
  - "non-intrusive"
  - Monitoring of existing traffic
  - Establishing of packet traces at different locations
  - Identification of packets, e.g. using hash values

  - Advantages
    - Does not affect applications
    - Does not modify the network behavior

  - Disadvantages
    - Requires suitable active network traffic
    - Limited to analysis of existing / current network behavior, situations of high load, etc. cannot be simulated/enforced
    - Does not allow the transport of additional information (time stamps, etc.) within measured traffic

Network Monitoring

- Applications of network monitoring
  - Traffic analysis
    - Traffic engineering
    - Anomaly detection
  - Accounting
    - Resource utilization
    - Accounting and charging
  - Security
    - Intrusion detection
    - Detection of prohibited data transfers (e.g., P2P applications)
  - Research

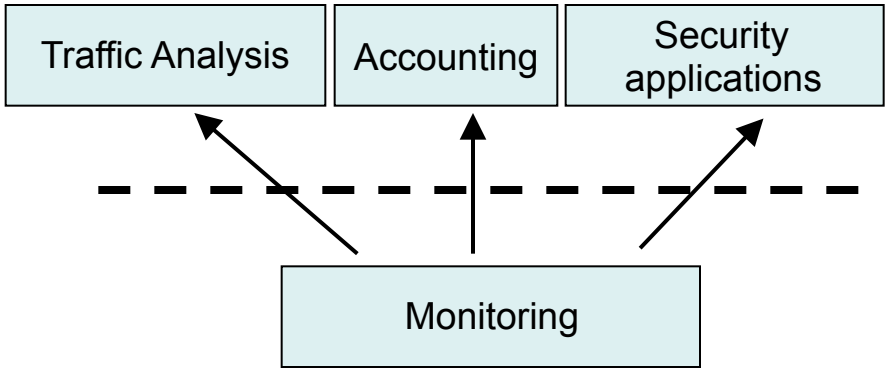

- Issues
  - Protection of measurement data against illegitimate use
  - Applicable law (“lawful interception”, privacy laws, …)

Network Measurements III

- Hybrid measurements
  - Modification of packet flows
    - Piggybacking
    - Header modification
  - Advantages
    - Same as for “passive”
    - additional information can be included (time-stamps, etc.)
  - Disadvantages
    - Modifying of data packets may cause problems if not used carefully

Measurement Types (summary)

- Active Measurements
  - Intrusive
  - Find out what the network is capable of
  - Changes the network state
- Passive Measurements (or network monitoring)
  - Non-intrusive
  - Find out what the current situation is
  - Does not influence the network state (more or less)
- Hybrid: combination of active and passive methods
  - Possible approach: alter actual traffic
  - Reduce the impact of active measurements
  - Might introduce new bias for applications

Flow-based Traffic Measurements

- **Export timeouts** (*flow expiration*)
  - *inactive timeout* ➔ export at the end of flow
  - *active timeout* ➔ export periodically for long-lived flows
  - timeouts can be configured (e.g., granularity ~60 sec)

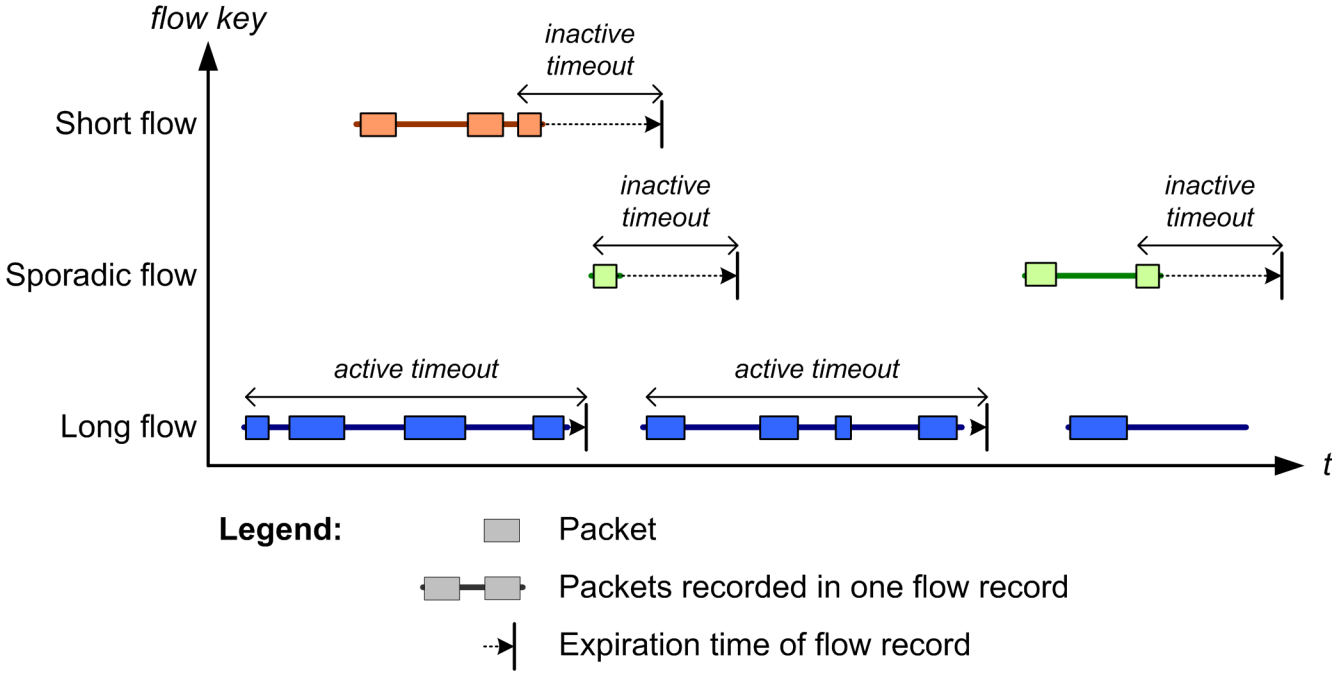

IPFIX: IP Flow Information Export

- IPFIX (IP Flow Information eXport) IETF Working Group
  - Standard track protocol based on Cisco Netflow v5…v9
- Goals
  - Collect usage information of individual data flows
  - Accumulate packet and byte counter to reduce the size of the monitored data
- Approach
  - Flows are represented by IP 5-tuple (protocol, srcIP, dstIP, srcPort, dstPort)
  - For each arriving packet, statistic counters of appropriate flow are modified
  - Whenever a flow is terminated (TCP FIN, TCP RST, timeout), its record is exported
  - Sampling algorithms can reduce the # of flows to be analyzed

IPFIX – Work Principles

- Identification of individual traffic flows
  - 5-tuple: Protocol, Source IP, Destination IP, Source Port, Destination-Port
  - Example: TCP, 134.2.11.157, 134.2.11.159, 2711, 22
- Collection of statistics for each traffic flow
  - # bytes
  - # packets
- Periodical statistic export for further analysis

Flow	Packets	Bytes
TCP, 134.2.11.157,134.2.11.159, 4711, 22	10	5888
TCP, 134.2.11.157,134.2.11.159, 4712, 25	7899	520.202

IPFIX - IP Flow Information Export Protocol

- RFCs
  - Requirements for IP Flow Information Export (RFC 3917)
  - Evaluation of Candidate Protocols for IP Flow Information Export (RFC3955)
  - Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information (RFC 5101)
  - Information Model for IP Flow Information Export (RFC 5102)
  - Bidirectional Flow Export using IP Flow Information Export (IPFIX) (RFC 5103)
  - IPFIX Implementation Guidelines (RFC 5153)
- Information records
  - **Template Record** defines structure of fields in **Flow Data Record**
  - Flow Data Record is a data record that contains values of the Flow Parameters
- Transport protocol: transport of information records
  - SCTP must be implemented, TCP and UDP may be implemented
  - SCTP should be used
  - TCP may be used
  - UDP may be used (with restrictions – congestion control!)

IPFIX – Terminology

- IP Traffic Flow
  - A flow is defined as a set of IP packets passing an observation point in the network during a certain time interval. All packets belonging to a particular flow have a set of common properties.
- Observation Point
  - The observation point is a location in the network where IP packets can be observed. One observation point can be a superset of several other observation points.
- Metering Process
  - The metering process generates flow records. It consists of a set of functions that includes packet header capturing, timestamping, sampling, classifying, and maintaining flow records.

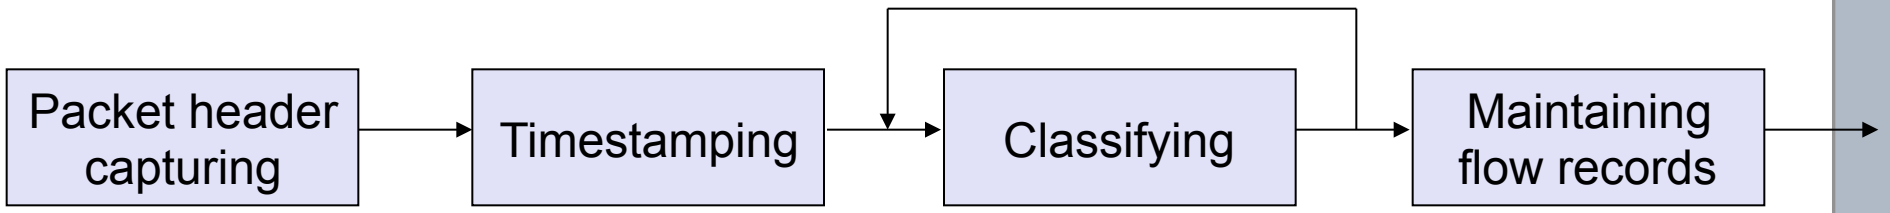

IPFIX – Terminology II

- Flow Record
  - A flow record contains information about a specific flow that was metered at an observation point. A flow record contains measured properties of the flow (e.g. the total number of bytes of all packets of the flow) and usually also characteristic properties of the flow (e.g. the source IP address).
- Exporting Process
  - The exporting process sends flow records to one or more collecting processes. The flow records are generated by one or more metering processes.
- Collecting Process
  - The collecting process receives flow records from one or more exporting processes for further processing.

IPFIX – Devices

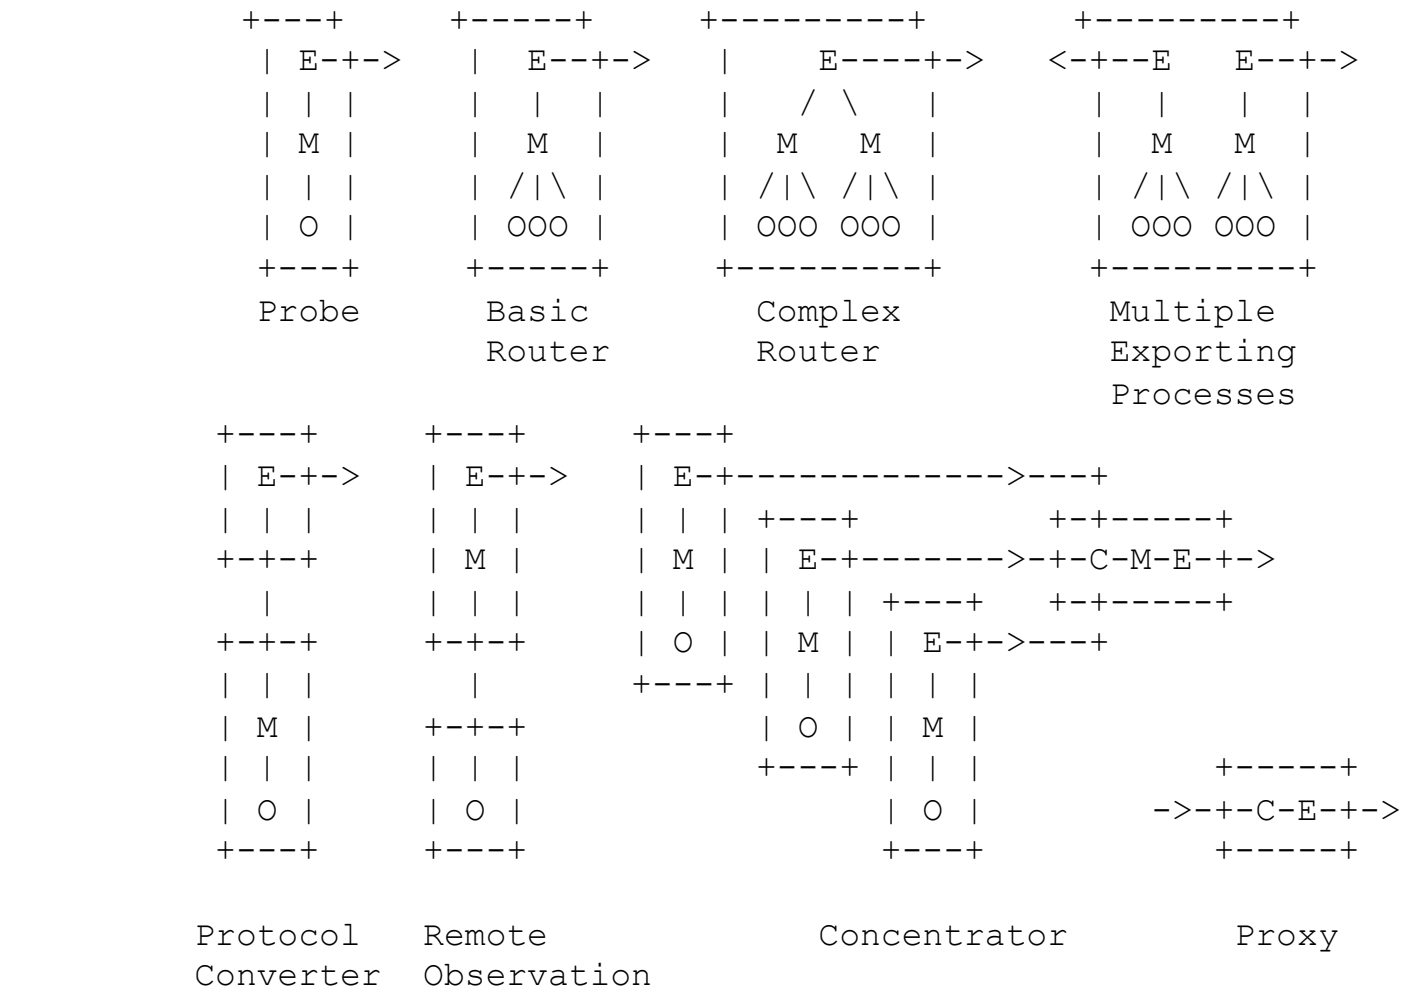

Generation of Time Series from Flow Data (2)

- Aggregation according to start times

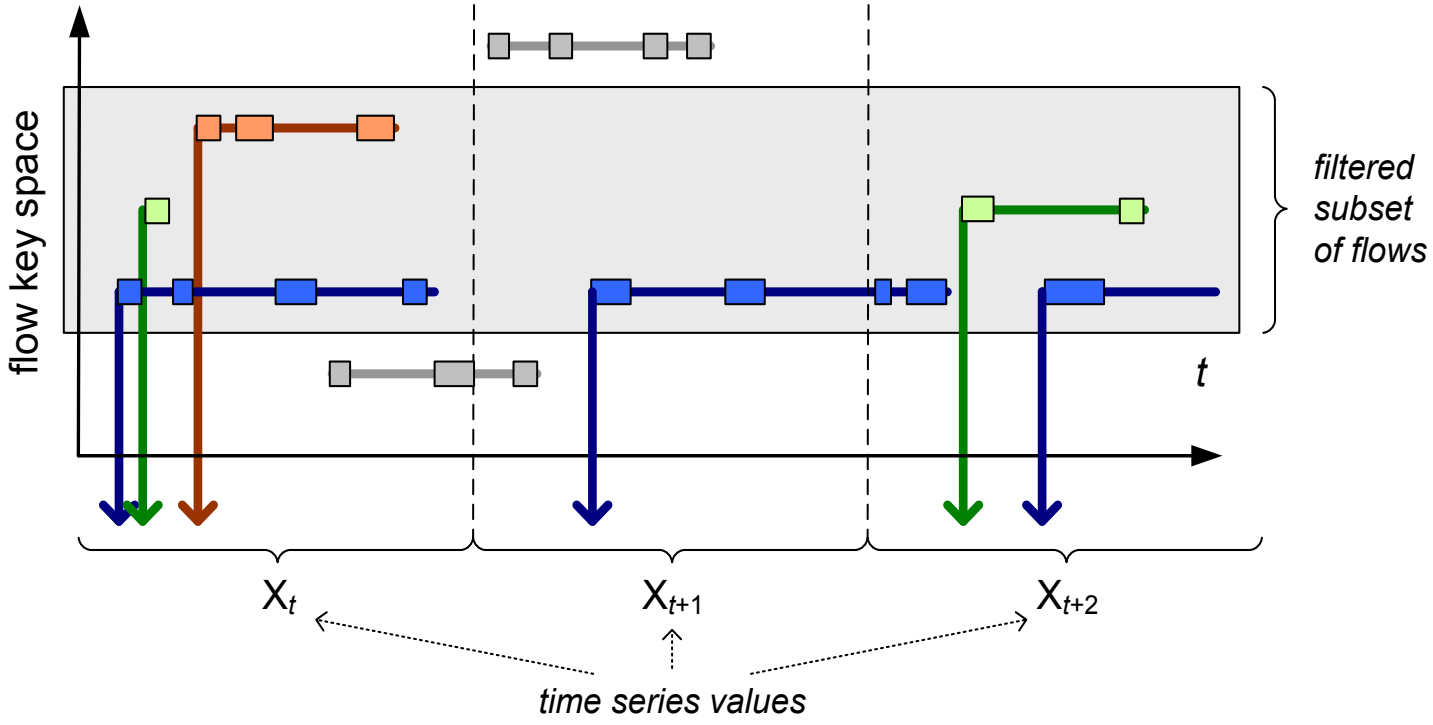

- Metrics to describe traffic volume
  - # Bytes, # packets, # flows

- Description of Flow-Key distributions

Generation of Time Series from Flow Data (2)

- Aggregation according to flow start time stamps

Start	Ende	Quell-Adr.	Ziel-Adr.	Prot.	Quell-Port	Ziel-Port	Byte-Anz.	Paketanz.
...	...	...	...	...	...	...	...	...
34:44	34:59	10.0.1.5	10.0.1.7	6	3458	80	11043	22
34:45	35:04	10.0.1.7	10.0.1.5	6	80	3458	83921	35
35:07	35:15	10.0.1.4	10.0.1.7	6	58312	25	73849	44
...	...	...	...	...	...	...	...	...
39:58	39:10							
40:00	40:13	10.0.1.5	10.0.1.7	6	3458	80	11043	22
...	...	...	...	...	...	...	...	...

5 minutes interval

- Advantage
  - Amount of data independent of traffic volume, and traffic characteristics

Packet Filtering – Algorithms

- Mask/match filtering
  - Based on a given mask and value
  - In the simplest case, the selection range can be a single value in the packet header (e.g., mask out the least significant 6 bits of source IP address, match against 192.0.2.0)
  - In general, it can be a sequence of non-overlapping intervals of the packet
- Router state filtering
  - Selection based on one or more of the following conditions
    - Ingress/egress interface is of a specific value
    - Packet violated ACL on the router
    - Failed RPF (Reverse Path Forwarding)
    - Failed RSVP
    - No route found for the packet
    - Origin/destination AS equals a specific value or is element of a given list

Packet Filtering – Algorithms II

- Hash-based filtering
  - Hash function h maps the packet content c, or some portion of it, to a range R
  - The packet is selected if h(c) is an element of S, which is a subset of R called the selection range
  - Required statistical properties of the hash function h
    - h must have good mixing properties
      - Small changes in the input cause large changes in the output
      - Any local clump of values of c is spread widely over R by h
      - Distribution of h(c) is fairly uniform even if the distribution of c is not

Packet Filtering – Algorithms III

- Hash-based filtering (cont.)
  - Usage
    - Random sampling emulation
      - Hash function (normalized) is a pseudo-random variable in the interval [0,1]
    - Consistent packet selection and its application
      - If packets are selected quasi-randomly using identical hash function and identical selection range at different points in the network, and are exported to a collector, the latter can reconstruct the trajectories of the selected packets

        ⇨ Technique also known as *trajectory sampling*
      - Applications: network path matrix, detection of routing loops, passive performance measurement, network attack tracing

IPFIX – Applications

- Traffic engineering
  - Comprises methods for measurement, modeling, characterization, and control of a network
  - The goal is the optimization of network resource utilization
- Traffic profiling
  - Process of characterizing IP flows by using a model that represents key parameters such as flow duration, volume, time, and burstiness
  - Prerequisite for network planning, network dimensioning, etc.
  - Requires high flexibility of the measurement infrastructure
- Usage based accounting
  - For non-flat-rate services
  - Accounting as input for billing
  - Time or volume based tariffs
  - For future services, accounting per class of service, per time of day, etc.

IPFIX – Applications II

- QoS monitoring
  - Useful for passive measurement of quality parameters for IP flows
  - Validation of QoS parameters negotiated in a service level specification
  - Often, correlation of data from multiple observation points is required
  - This required clock synchronization of the involved monitoring probes
- Attack/intrusion detection
  - Capturing flow information plays an important role for network security
  - Detection of security violation

    1) detection of unusual situations or suspicious flows

    2) flow analysis in order to get information about the attacking flows

Packet Sampling

- PSAMP (Packet SAMPling) WG (IETF)
- Goals
  - Network monitoring of ultra-high-speed networks
  - Sampling of single packets including the header and parts of the payload for post-analysis of the data packets
  - Allowing various sampling and filtering algorithms
    - Algorithms can be combined in any order
    - Dramatically reducing the packet rate
- Benefits
  - Allows very high-speed operation depending on the sampling algorithm and the sampling rate
  - Post-analysis for statistical accounting and intrusion detection mechanisms