

Proposal for a new COST Action in Traffic Monitoring and Analysis

Author: Fabio Ricciato
Forschungszentrum Telekommunikation Wien (FTW), Vienna, Austria
ricciato@ftw.at

Revised: 6th May 2006

Proposed Title:

Traffic Monitoring and Analysis Methods in Support of Network Operation and Engineering

Motivations:

In the recent years studies and research activities on Internet Traffic Monitoring have surged, enabled by the availability of capture hardware (DAG cards, wiretaps) and large-storage solutions (e.g. RAID) at accessible cost. There are now several research groups -mainly in EU and US- involved in developing new tools and methods to acquire, analyze and interpret traffic data from the live network in non-intrusive way, and they already address a range of different network environment (e.g. wired access, broadband backbone, campus WLAN, 3G WAN). We will use the term Traffic Monitoring and Analysis (TMA for short) to collectively refer to such activities.

Current TMA research activities include the following applications:

- Anomaly Detection and Intrusion Detection: to detect security threats, attacks, infections and any other macroscopic change in the traffic pattern, deliberate or not, that require attention by the network staff.
- Performance Monitoring: to quantify the current network performances at a large-scale; they are the basis to detect global drifts and/or local degradation that trigger intervention and revision by the network staff.
- Network Tomography and Topology Discovery: to infer the status of internal network elements from a parsimonious number of measurement points; they can be used to detect network problems (e.g. capacity bottlenecks, failures) in a cost-effective way.
- [... *anything else to be added ???*]

The types of monitored data (e.g. NetFlow data, tcpdump traces, BGP records) and acquisition methods vary depending on the particular application. Beyond the differences, there are some underlying commonalities across TMA research activities: adoption of non-invasive measurement infrastructure, scalable storage solutions, methods for representation, indexing and retrieval of large heterogeneous datasets, methods for visualization and exploration of multidimensional datasets, adoption of advanced techniques from different areas (statistics, signal-processing, data-mining, etc.). In addition, the strongest unifying factor is perhaps the fact that virtually all TMA activities can be applied in real networks to improve the effectiveness of the network operation and engineering processes.

The scale and the functional complexity of the modern packet networks make it difficult and costly to rely on the classical methods for network monitoring and

management which were based solely on the few data available directly from the network equipments (counters, logs, SNMP). New techniques based on TMA can be successful and cost-effective in solving several practical problems found in the operation of a real network. When integrated and consistently coordinated (e.g. using common tools), individual TMA techniques can collectively evolve towards a new paradigm for operating and engineering modern telecommunication networks in a cost-effective way.

This vision requires a stronger coordination -ideally a collective collaboration- between the various research groups active in the TMA area. Several groups have access to operational networks, academic or commercial (in force of private collaborations with operators). There are already initiatives aimed at promoting the sharing of traffic traces and pieces of monitoring tools (see [CCR1,CCR2] and references therein). However much more can be done for improving the sharing of know-how, lessons-learned, problems found during practical deployment, ideas for real-world exploitation of known techniques etc.

It would be desirable to set a platform for improving the coordination, networking and information sharing between the EU research groups in this field. A new COST Action might be a first step in this direction. This would facilitate early dissemination of preliminary results and ongoing work, promote the setup of ad-hoc bilateral collaborations, and improve networking between the community members. As a side effect, those few groups that have access to commercial networks can share their operational experience with other groups, and together improve the level of practical applicability of the proposed TMA solutions. This is essential to help TMA research to move from research papers into real-world network operation and engineering practice.

References:

[CCR1] Shannon et al, The Internet Measurement Catalog, ACM Computer Communication Review, vol. 35, n. 5, October 2005.

[CCR2] M. Faloutsos, Public Real Data Repositories and Measurement Tools, ACM Computer Communication Review, vol. 36, n. 2, April 2006.