

Report from the 6th PhD School on Traffic Monitoring and Analysis (TMA)

Idilio Drago
Politecnico di Torino
idilio.drago@polito.it

Fabio Ricciato
University of Ljubljana
fabio.ricciato@fri.uni-lj.si

Ramin Sadre
Université catholique de Louvain
ramin.sadre@uclouvain.be

This article is an editorial note submitted to CCR. It has NOT been peer reviewed.
The authors take full responsibility for this article's technical content. Comments can be posted through CCR Online.

ABSTRACT

This is a summary report by the organizers of the 6th TMA PhD school held in Louvain-la-Neuve on 5-6 April 2016. The insight and feedback received about the event might turn useful for the organization of future editions and similar events targeting students and young researchers.

CCS Concepts

•Networks → Network monitoring; *Public Internet*;

Keywords

PhD School; Organization; Feedback; Lessons Learned

1. BACKGROUND AND ORGANISATION

The series of TMA workshops and TMA PhD schools was initiated under the umbrella of the COST Action IC0703 on “Traffic Monitoring and Analysis” (TMA) that ran from 2008 to 2012 and involved more than 50 research groups across 24 countries [1]. After the closing of the COST Action, the series of yearly TMA events has been continued by a core group of active community members. In the two last editions (2014 at QMU in London, 2015 at UPC in Barcelona) the TMA Workshop and TMA PhD school were co-located in order to foster joint participation to both events by young researchers. The same format was adopted for the 2016 edition hosted by UCL in Louvain-la-Neuve, Belgium.

The 8th edition of the TMA Workshop [2] and the 6th edition of the TMA PhD school [3] took place in April 2016 in Louvain-La-Neuve, Belgium, hosted by the Université Catholique de Louvain (UCL). The PhD school was organized by Idilio Drago from Politecnico di Torino, Italy, Fabio Ricciato from University of Ljubljana, Slovenia, and Ramin Sadre from UCL, Belgium. Ramin was acting also as Local Organizer and co-chair of the co-located TMA workshop. As the previous edition in 2015, the main financial support was provided by a SIGCOMM grant and by a MLAB/Google donation. We also acknowledge the support received by IFIP for the best poster award and website hosting.

The initial program of the PhD school included a total of 5 lectures with on-site speakers from USA, Australia and Europe. As of March 21st 40 students had registered for participation.

Sadly, just two weeks before the event, the terrorist bombing attack took place in Brussels [4]. The Brussels Airport in Zaventem was severely damaged and remained closed until

early April, with hundreds of flights canceled. In the aftermath of the attack, some countries issued security warnings and restrictions for trips to Belgium. As a consequence, just a few days before the event it became clear that some of the planned speakers and a few prospective participants could not be able to reach Louvain-la-Neuve on the planned dates.

After consulting with other people in the community, including representatives from the SIGCOMM committee, it was decided to hold the PhD School in the same location and dates as planned, but with a modified program to adapt to the contingent situation. Specifically, the following changes were implemented:

- Two lectures were rescheduled to be given via video-conferencing, with speakers connecting remotely to the classroom in Louvain-la-Neuve.
- One lecture was canceled and replaced by other interactive sessions moderated by the organizers, as detailed below.

The resulting final program consisted of four lectures, thereof two with on-site speakers and two with remote speakers, and three different kinds of interactive sessions: poster sessions, group brainstorming and toy exercise. Support for remote video lecturing and recording was provided by Meetecho [5]. Additionally, all lectures were web-casted in live streaming through the Meetecho platform, resulting in a few additional remote attendants for each lecture. In total, 37 students participated to the PhD School from 22 institutions and 11 different countries.

2. LECTURES

The PhD School included four lectures covering a broad spectrum of topics in computer networks. In the following, we provide highlights of each lecture. All slides and video recordings of lectures are available online [3].

2.1 Traffic Matrices

Matt Roughan from the *University of Adelaide* opened the PhD School discussing traffic matrices. Traffic matrices are important building blocks, not only for researching purposes, but also for operating the network. They describe the communication volumes between end points in a network. Examples of their practical applications range from operators that need traffic matrices for planing the network or for performing anomaly detection, to researchers that need traffic matrices for designing new protocols.

Traffic matrices are however very hard to obtain. After reviewing basic definitions and the role of traffic matrices in the network field, the talk delved into common practices to measure, manipulate and apply traffic matrices. Matt reinforced that, while Internet traffic can be measured with tools such as Netflow, traffic matrices require further knowledge, such as routing and network topology information. Thus, it is still very uncommon to see real instances of traffic matrices. The lack of data to build real traffic matrices makes it prime to have techniques to generate synthetic traffic matrices. Matt went over many of such techniques, including the Maximum Entropy methods he has been researching for modeling traffic matrices.

Matt concluded with his thoughts about network research in general and his advice for young researchers. He sees that network performance research should not be considered without care about the measurements and the tasks at hand. It is essential to connect the research with real problems: When problems are simultaneously tractable, interesting and useful, odds of success in the research are substantially increased.

2.2 Multipath TCP

Olivier Bonaventure from *Université catholique de Louvain* delivered a lecture on Multipath TCP (MPTCP). First, he presented the motivations and main design principles of MPTCP. Then, he focused on several engineering and implementation aspects of MPTCP, concluding with practical use cases of the protocol.

MPTCP is motivated by the evolution of Internet terminals, which more and more are equipped with multiple interfaces. Examples include mobile devices shipped with both WiFi and 4G, and data-center networks that specify redundant paths between servers. It is common that independent IP addresses are assigned to each interface. TCP is designed to establish a connection between end points identified by unique IP addresses, thus preventing the use of multiple interfaces in a connection.

The main goal in the MPTCP design was to extend TCP — MPTCP should work where regular TCP works. However, middleboxes create several hurdles to TCP extensions. In fact, the Internet is *not transparent*, and several fields of network and transport layer protocols are often touched by middleboxes nowadays. A naive implementation of MPTCP, e.g. based on sending packets on different interfaces without establishing ordinary TCP connections, would fail on most networks. Ingenuity was required to build a deployable protocol. MPTCP establishes ordinary TCP connections on the multiple paths — i.e., from the network point of view, a MPTCP connection looks like multiple regular TCP connections. This design however came with caveats: flow control and congestion control on multiple interfaces are some examples of problems that had to be solved to enable MPTCP.

Olivier showed many success stories of MPTCP. For instance, MPTCP can better utilize the network resources in a data center connected using a FatTree topology. Another success story is the deployment of MPTCP in smartphones, in particular iPhones. Apple seems interested in MPTCP to reduce latency on its voice command application Siri. Once a sub-flow in the MPTCP connection serving Siri is lost, other sub-flows can continue the communication seamlessly, without the need to reestablish connections.

Olivier concluded showing other use cases that highlight how MPTCP is becoming a reality. He called for more use cases, which will bring MPTCP deployment experience. In fact, many of the problems solved in the past 20 years of TCP operations will need to be revisited for MPTCP. It is to be expected that several measurement studies will be required in the coming years to fully understand the implications of MPTCP deployment in the wild.

Finally, the IETF published the Multipath TCP RFC in 2013 (RFC 6824), and an implementation in the Linux kernel is available [6].

2.3 HTTP/2 Adoption and Performance

On the second day, *Matteo Varvello* and *Jeremy Blackburn* from *Telefonica I+D* introduced their work on HTTP/2, covering efforts to unveil the state of HTTP/2 deployment and the performance benefits of adopting the protocol.

They started by reviewing the HTTP evolution, highlighting the distinct features on HTTP/2, e.g., server push and content prioritization. Those features have been motivated by the widely-known limitations of HTTP/1.1, such as head of line blocking and the redundancy on HTTP headers. However, whereas HTTP/2 standardization is ready, it is still unclear whether the protocol will deliver the intended gains, given the lack of broad deployment experience with the new protocol. Matteo and Jeremy are exploring questions around HTTP/2 deployment in two fronts.

First, they presented a measurement platform [7] that monitors HTTP/2 adoption and performance in the wild. Currently, the platform is monitoring the HTTP/2 deployment across more than 1 million websites on a daily basis. It executes several tests, such as to check whether websites fully supports HTTP/2 and to assess the performance of loading websites using both HTTP/1 and HTTP/2. Measurements are freely available to the community on the project website.

Second, Matteo and Jeremy discussed the difficulties for measuring the performance gains brought by HTTP/2. Traditionally, Page Load Time (PLT) is used as a key metric to benchmark web sites. Many approaches to estimate PLT have been proposed and, in particular, the time to observe the `onLoad()` event of the web browser is used as an indication of PLT. While Matteo and Jeremy also rely on the `onLoad()` event to estimate PLT, they acknowledge the lack of connection between the approach and the actual performance perceived by end users. The researchers then described their experience building a crowd-sourced platform [8] to measure PLT.

Overall, HTTP/2 adoption was shown to be increasing at a fast pace, but a large number of websites only advertises support for the protocol, without fully serving content using HTTP/2. The impact of HTTP/2 on PLT seems to be positive, in particular in constrained networks. However, the extent to which end users perceive such improvements is still an open research question.

2.4 Internet-Scale Experimentation

Fabián Bustamante from *Northwestern University* closed the PhD School lectures discussing Internet-scale experimentation. The success of Internet has turned the network into a critical piece of our daily lives. The Internet is now ubiquitous, and many crucial systems rely on it for basic operations. This success however came with a high price for

those experimenting new Internet applications: It is virtually impossible to run a controlled experiment at even fractions of the scale of most systems running over the Internet.

Fabián argued that it is very challenging to design experiments that are representative, reproducible and ethical on Internet scale. Measurement platforms, methodologies, standards, legal and ethical guidelines are missing. He provided several examples of research questions that are faced by means of measurements, and illustrated how inadequate experimentation methodologies can lead to pitfalls and wrong conclusions. As an example, Fabián discussed how one could try to measure the impact of reliability on users' experience. Ideally, such experiment should be designed following classical controlled experiments, where control and treatment groups are randomly selected, to represent users that experience low and high reliability. However, many factors impact such selection, such as the lack of control over people's networks, lack of people's consent to perform measurement, among others.

Fabián presented strategies and good practices for designing sound experiments, including his efforts to create DASU [9] – a measurement platform that works on the edge of the network, close to users' devices. DASU relies on a large number of measurement nodes, which were made possible by aligning end-users' and experimenters' objectives: broadband characterization was used as an incentive to convince people to join the network. Despite the success of DASU, Fabián concluded that Internet-scale experimentation is still in its infancy, and much more work is expected to fill in the gaps for sound Internet-scale measurement platforms.

3. OTHER ACTIVITIES

In addition to classical lectures, the program was enriched with new types of activities specifically designed to stimulate interaction between the students. The format of such activities was somewhat “experimental” and can be certainly refined and improved in future editions (also based on the feedback collected this year). However, the general response by the participants indicates clearly that such kinds of activities were appreciated and should find room also in future editions.

3.1 Poster Sessions

Poster sessions had been introduced in the program already in previous editions as a means to let PhD students learn about other students' work and to foster peer-to-peer interaction among them. Every participant to the school was asked to bring a poster in free format to represent her/his PhD topic, describing the target problem, planned methodology, and possibly showing recent results (if any). Five poster sessions were held during the breaks between the lectures, with coffee and snacks served in parallel in the same room. In each session, 5-6 different students presented their posters to the remaining participants.

This year a small expedient was implemented to catalyze interaction and, at the same time, avoid possible imbalances in the distribution of visitors per poster. Namely, a “mandatory assignment” game was created. Within each session, each student was assigned two posters among the ones being presented for which she/he was required to provide critical but constructive feedback to the respective presenters. The feedback format was very simple, namely 1-2 short sen-

tences handwritten on a post-it, to be attached directly onto the poster. In other words, every student had the freedom to visit any poster in the session, with the constraint of visiting (at least) the two assigned ones. The assignment was generated (semi-)randomly by the organizers and provided to the participants at the beginning of each session. Care was taken to avoid associations between students from the same institution.

According to the feedback received from several participants, this simple expedient was effective to trigger students to visit posters on topics completely different from their own PhD project, that they would have not visited otherwise. Thus, presenters received feedback not only from students with knowledge on the specific topic of the poster, but also from a more general audience with various backgrounds. Furthermore, the explicit task of providing a “critical but constructive feedback” resulted in more focused technical interaction between the (assigned) visitor and the presenter.

Overall, this assignment game has proved very successful. Students appreciated the activity and many of them suggested that next School editions should reserve more time for such activity, beyond the usually short intervals reserved for coffee breaks. All posters presented during the School are available on the PhD School website [3].

3.2 Group Brainstorming Session

A group-based activity was proposed to foster peer-to-peer interaction among the students. The session was divided into two rounds of 30 minutes each. At each round, students were divided in groups of three, according to a pre-defined list prepared by the organizers. We took care of grouping together people (i) from different institutions and (ii) working on diverse topics.

Each group was given the task of formulating a creative research idea based on the fusion of the individual research topics carried out by the group members. The expected outcome of the round was a one-page sketch of the research question(s), expected outcome and a tentative research plan. The students were invited to use their posters as basis for discussion with the other group members.

The brainstorming sessions were in general appreciated, because it offered more opportunities for interaction between students, which is in-line with the feedback received on poster sessions. Thus, the main objective of the activity was achieved. Some groups were able to produce surprisingly rich proposals, despite the limited time for discussions.

Students also highlighted points to be improved in such type of group activity. In particular, some groups had difficulties to spark ideas for collaborations, given the wide-range of topics of group members. The coaching by professors and the organizers helped to speedup the exercise, but it was hard to offer support to all groups simultaneously. Students also suggested the setup of channels for continuing the exercise beyond the PhD School, so that initial ideas produced during the School could be further developed, or potentially extended by others in different groups.

3.3 Toy Exercise

The goal of this session was to raise awareness among the students about some fundamental mistakes and insidious aspects that may be encountered in an apparently simple data analysis task, and trigger them to reflect about the importance of exercising caution and critical sense even in what

may appear at first look a trivial assignment. To this end, Fabio prepared a “toy exercise”, based on synthetic data, designed to emulate a scenario of network data analysis.

The exercise goes like that: Two datasets were provided representing Round-Trip Times (RTT) collected on different networks, the *green* and the *blue* networks, along with a brief scenario description and a summary of the measurement methodology. Based on such input datasets and context information, the students were challenged to indicate which of the two networks yields the “best performance”, and to give quantitative evidence in support of their claim.

The exercise and all support material (datasets and assignment description) were given to the students on the first day of the PhD School, and they were invited to process the data and present findings on the next day. On the second day, Fabio asked some of the students to briefly report their results. Afterwards, he guided the class through a possible resolution path.

Along the way, Fabio highlighted the importance of careful exploration and proper data visualization in order to pinpoint apparent anomalies and inconsistencies in datasets. For instance, proper exploration of the two dataset would reveal subtle but critical differences in the measurement methodology adopted in the two networks (e.g., different timeout values, different definitions of “spacing” between subsequent probes). He drove the students to reconstruct the impact of these and other effects on the data, and in which way completely erroneous conclusions might be drawn by a superficial analysis that fails to recognize these aspects. Fabio highlighted ways to infer possible explanations and countermeasures to typical problems found in real measurement datasets. Based on these “toy” examples, he discussed some general aspects and fundamental issues encountered in the collection, reporting and analysis of measurement data.

This session was conducted in a very tutorial and interactive way. At different stages during the development of the exercise, Fabio “trapped” the students into a wrong answer or conclusion, before driving them to recognize the problem and identify possible corrections. The mismatching between the apparent simplicity of the initial assignment on the one hand, and the richness of the fundamental methodological aspects that emerged during the discussion of the resolution process on the other hand, contributed to capture the attention of students.

The text of the toy exercise as well as slides with the solution can be downloaded from PhD School website [3].

4. FEEDBACK FROM STUDENTS

After the event each student was invited to fill in an *anonymous* online feedback form. The form included several questions and was designed to gauge the appreciation by the students and collect their detailed opinions. Along with numerical scores, students were invited to provide free-text comments on what they liked and disliked about each session and about the School as a whole, and to leave suggestions for future editions. We had 32 respondents out of 37 participants, and most of them provided very detailed textual comments.

Overall, the 2016 TMA PhD School edition was received positively by the participants, with 22 out of 32 respondents rating the whole event “very good” or “excellent” as seen in Fig. 1. Fig. 2 delves into the level of appreciation on individual sessions. Respondents were asked to grade each session




Figure 1: “Overall, how would you rate the event?” (32 respondents).

on a scale from 1 (“Not at all interesting”) to 5 (“Extremely interesting”). Fig. 2 reports the full score distribution (normalized) for each type of session — the four lectures were divided into two groups, depending on whether the speakers were remote or on-site.

It is evident from the figure that remote lectures were less appreciated than lectures with on-site speaker, respectively with 50% and 70% of the answers in the top categories 4 (“Very interesting”) and 5 (“Extremely interesting”). According to the written comments left by several participants, the main problem with remote lectures was the lack of interaction between the speakers and students. This was somehow expected and, in fact, the decision of holding remote lectures was purely due to the contingency situation explained earlier in Section 1.

It appears from Fig. 2 that the two sessions dedicated to “horizontal” interaction between the students were very well-received: the score distribution for poster sessions and group brainstorming are similar to on-site lectures, with 70-80% of top scores. Several students highlighted the importance of meeting and interacting with students from other universities and working on different PhD topics. Another point worth remarking is the importance of catalyzing horizontal interactions that would likely not occur spontaneously. For example, several students reported that they found surprisingly beneficial to interact with other students working on research problems completely away from their target PhD topic. On the other hand, the same students admitted that they would have not paid much attention to those works (e.g., posters) had they not been instigated to do so by the rules of the session – recall from Section 3.1 and 3.2 that poster and group associations were established by the organizers with this purpose.

A bit surprisingly, Fig. 2 shows that the “toy exercise” proposed by Fabio received the highest appreciation, with 28 top scores out of 30 respondents for this session and no negative score. Two ingredients might help to explain the success of this activity: interactivity and concreteness. First, this session was conducted in a highly interactive way between the speaker and students – interaction (or lack thereof) was one of the most cited terms by students in their detailed feedback on all sessions. Second, the exercise was perceived by the students as a hands-on activity, since they had to manipulate actual data and try to provide answers to concrete questions. The exercise was framed into a concrete scenario (realistic, though not real) and illustrated pitfalls in data




Figure 2: “How would you rate each individual session?”. Normalized distribution of scores for each session type.

analysis following a bottom-up approach. Many students recognized that the lessons learned in terms of methodology to analyze measurements can help their daily research.

5. LESSONS LEARNED AND SUGGESTIONS FOR FUTURE EDITIONS

Hereafter we summarise some key “lessons learned” based on the analysis of the student feedback, on informal discussions with some of the participants, and on our own perception of the event. We hope that they will help to further improve the organization of future editions.

- **Interactivity is premium.** School activities should be preferably conducted in two-way interaction, rather than pure one-way communication (as in traditional talks at conferences).
- **Vertical and horizontal interaction.** The main reason for holding a PhD school is to provide room for interaction between the PhD students and ... somebody else than their respective PhD supervisors. While “vertical” interaction between the students and senior lecturers is the primary ingredient, “horizontal” (peer-to-peer) interaction between the students themselves is also an important aspect and should not be overlooked. A good balance between the two should be pursued in the program. Some future activities might well combine both in the same session (e.g., group work coached by senior researchers).

- **Listen. Talk. Do.** In a good interactive session the students *listen* (to the speaker, to other students) and *talk* (e.g., when presenting their work and ideas, or when asking or answering questions). Besides that, it is important to include into the program activities where the students have to *do* things: exercises, assignments, laboratories, hackatons, etc. Whether in group or individually, they should be given the opportunity of trying things (and possibly failing them), and compare / discuss their outcome and experience with the other participants.
- **Proper load.** The program should not be overloaded, and sufficient time should be allowed for informal interaction *outside* the working sessions (e.g., during break and social events).

Acknowledgments

We would like to thank the ACM SIGCOMM, M-Lab and IFIP for sponsoring the 6th TMA PhD School. Thanks to the Meetecho team for the timely and professional support with video conferencing. Thanks to the UCL staff for hosting the PhD School providing local organization support. Thanks to all the speakers who have volunteered to give the lectures. Finally, the organizers would like to thank the students, who were still willing to attend the School despite all unforeseen travel problems.

6. REFERENCES

- [1] TMA COST Action IC0703 website. www.tma-portal.eu/cost-tma-action/general-info.
- [2] TMA 2016 Workshop website. tma.ifip.org/2016.
- [3] TMA 2016 PhD school website. tma.ifip.org/2016/phdSchool/.
- [4] Wikipedia entry for 2016 Brussels bombings. en.wikipedia.org/wiki/2016_Brussels_bombings.
- [5] Meetecho homepage. www.meetecho.com.
- [6] Linux Kernel MultiPath TCP project. <http://www.multipath-tcp.org>.
- [7] Is the Web HTTP/2 yet? <http://isthewebhttp2yet.com>.
- [8] eyeorg – A platform for crowdsourcing Web quality of experience measurements. <http://eyeorg.net>.
- [9] DASU – A platform for measurement experimentation and broadband characterization. <http://www.aqualab.cs.northwestern.edu/projects/115-dasu-isp-characterization-from-the-network-edge>.