Research Statement

Romain Jacob

The energy demand for telecommunication networks keeps rising. As all parts of society need to do their share to meet our planetary sustainability targets, academics are taking a fresh look at how efficiently we are using the energy that powers the core of the Internet. My current research focuses on studying and addressing the limitations and inefficiencies of today's infrastructure to help build a more sustainable Internet for the future generations.

Sustainable networking

I initialized a research line on sustainable networking within Prof. Vanbever's group. My main responsibility is to lead the research activities in that area, which includes

- performing initial exploratory studies, both personally and via undergraduate projects
- applying for funding (from ETH, SNSF, and the EU) and managing the funded projects
- hiring and supervising the doctoral students assigned to these projects
- disseminating the research outcome in talks, tutorials, and conferences
- initiating collaborations with researchers within ETH, Switzerland, and internationally
- coordinating the overall research activity on sustainable networking within the group.

While collaboration with other senior members of the group is welcome, I will maintain scientific ownership of these topics to build my academic reputation in the long term.

My research agenda on sustainable networking is based three complementary pillars: modelling the power demand of today's Internet; reducing the energy cost for serving today's traffic; incentivizing tomorrow's demand towards a more sustainable Internet.

Modeling network energy usage

Figure 1: An Autopower measurement unit.

First, we must understand the existing infrastructure and, more specifically, how much power is needed to forward Internet traffic.

Today, power monitoring in (wired) networks is scarce, high-level, and imprecise. Moreover, while monitoring is important, it is not sufficient to predict the power impact of a change (*e.g.*, increasing traffic volume or running a different protocol). To quantify the energy savings of candidate measures, we need power models for the networking infrastructure. Over the past two years, we developed, implemented, and validated an experimental methodology to derive power models for network routers. This included the design of tools to facilitate remote power data collection (Figure 1), orchestrating and processing power modeling experiments,[1] and publicly serve the collected data online.[2]

This work was supported by an EFCL grant and led to a full paper[3] (under submission).

[1] github.com/nsg-ethz/netpowerbench

[2] networkpowerzoo.ethz.ch

[3] Jacob et al. 2024

Reducing the Internet footprint

Once we know what the current power demand is, we try to drive it down.

The Internet infrastructure is dimensioned for the peak traffic demand, which is growing. However, this growth is not uniform over time and the average network utilization is going down in many networks. This leads to inefficiencies as today's networking hardware has little power proportionality. In this context, one can reduce the power footprint by improving proportionality, as commonly done in embedded systems (duty cycling) or CPUs (C states). In this research, we investigate ways to improve power proportionality in networking hardware today. Initial projects considered turning off unused links[4] or redundant PSUs[5].

[4] Röllin 2023
[5] Béhanzin 2024

This work was supported by an EFCL grant and led to an awarded paper[6].

[6] Röllin et al. 2024

Incentivizing sustainable networking planning and utilization

Looking further, we research ways to improve tomorrow's Internet.

- At the micro level, we consider how we should adapt the hardware and software co-design of the network infrastructure to be more energy efficient. While researching this area, we collected anecdotal evidence of blatant inefficiencies at the firmware and driver levels (*e.g.*, components logically "turned off" but remaining powered on).
- Conversely, at the macro level, we consider incentive mechanisms that could nudge network users (resp. operators) towards a more sustainable demand (resp. capacity planning). This relates to a now-popular concept known as "carbon-aware <something>" whereby one shift workloads in time and/or space to consume electricity when and/or where the energy production has the least carbon footprint.

In addition, we also research ways to rationalize the traffic demand: one of the reasons we consume more and more bandwidth is because it is cheap and easy. Users are lured into the illusion of limitless resources. The name of "cloud" computing perfectly illustrates that idea: we can always get more—or so we are tricked into believing. That is certainly not true, and continuous increase in demand has serious (negative) sustainability impacts. At the same time, we certainly do not need nor benefit from all the traffic that we send and receive over the Internet. Some is useful, some is not; but we consume the network resources anyway because the cost is invisible. To address this, we aim to quantify and expose the footprint of using the network to its user. The hope is, eventually, to foster a more frugal behavior from all the stakeholders, from the end user to the hardware manufacturers.

I currently supervise one doctoral student investigating this research direction.

Other research areas (selected)

Over recent years, I was also involved in research projects in other areas, as co-advisor of doctoral students in Prof. Vanbever's group or during my own academic education.

I briefly describe three of those below.

Replicability in networking research

In many scientific fields, experiments are often not performed, analyzed, and reported in a way that allows to replicate them; this problem is known as the "reproducibility crisis."

I investigated how one could foster and formally specify replicability in networking research, where it is often impossible to replicate the exact same experimental conditions, thus leading to unavoidable performance variability in general. We observed however that, albeit unavoidable, performance variability is not arbitrary. In particular, we proposed that the variability in networking measurements can be classified hierarchically using three different time scales which we call runs, series, and sequels.

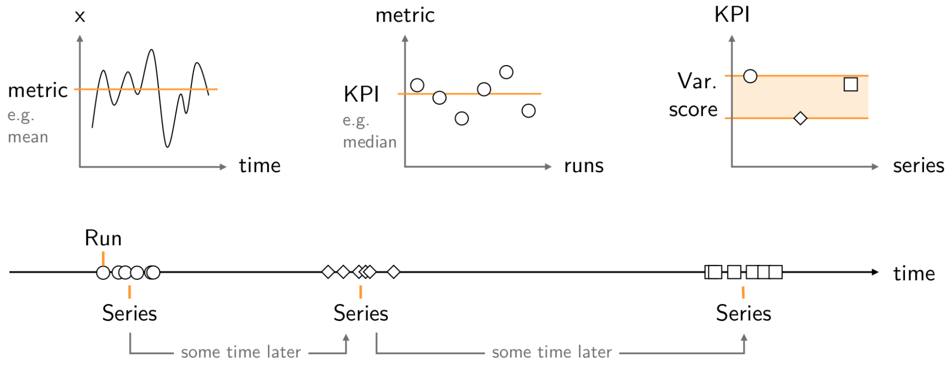

Figure 2: In networking experiments, performance often varies along three different timescales: during the run of the system, across multiple runs, and over longer periods of time. To address the replicability of the results, one must be aware of and account for these different sources of variability.

Based of that observation, we ask two hard questions:

Q1 Given such unavoidable variability, how can we **formalize replicability** and assess whether experimental results are indeed replicable?

Q2 Intuitively, the more repetition one performs, the more reliable the results are. But how can we **quantify** the relation between experimental effort (*i.e.*, the number of repetitions performed) and the confidence in the result? How can we rationally decide how many repetitions are enough?

We proposed answers to these questions with *TriScale*, a framework guiding experiments design and data analysis aiming to foster replicability of performance evaluation in networking. *TriScale* applies appropriate statistical methods to account for the variability of each timescale and provides performance results with quantifiable levels of confidence (*e.g.*, a 95% confidence interval for the median throughput of a flow). The key idea of *TriScale* is to link these methods with the desired levels of confidence to derive the minimal number of repetitions that are required. This allows to answer **Q2**.

Our research illustrates that the formalization of replicability (**Q1**) cannot be set in absolute terms or using binary criteria; some systems naturally exhibit more performance variability than others, thus it is natural to see more variability in experimental evaluations of such systems. Instead, we argue for a quantification of the replicability of experiments; that is, estimating the range of expected variation. It is then up to the different research communities to define acceptable ranges for a specific classes of systems (*e.g.*, congestion-control schemes, multi-hop wireless communication protocols, etc.).

TriScale is implemented, operational, and open source.[7] The framework guides the user through a systematic experiment design and data analysis approach; based on the metrics of interest and desired levels of confidence, *TriScale* computes the minimal numbers of runs, series, and sequels required. Then, given the collected data, the framework computes the appropriate statistics and returns performance results which are guaranteed to hold with the level of confidence specified.

[7] triscale.ethz.ch

Scalable network monitoring

Monitoring the state of a network is challenging, especially at large scale. There is an inherent trade off between the accuracy of the telemetry information collected and the overhead of collecting this data. One interesting approach for scaling network monitoring is to leverage the structure of the telemetry data. For example, the working principles of today's Internet makes it likely that traffic sourced by the same IP address will reach a network via the same entry point; thus, measuring only one packet provides plausible information about many others. One may combine those inferences to significantly expand the monitoring of a network at minimal cost.

I assisted one doctoral student to turn this idea into novel network monitoring systems.[8] The project further developed into a start-up,[9] though I am not personally involved.

[8] Bühler et al. 2023

[9] netfabric.ai

Opacity of discrete event systems

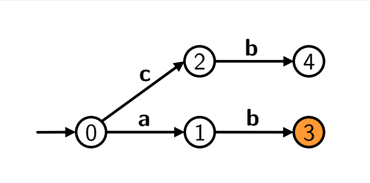

Figure 3: If 3 is the secret state (colored) and **b** is the only observable event, then this automaton is opaque. However, if **a** is also observable, it is not opaque as the observer is sure that the automaton is in state 3 after observing **ab**.

Opacity is an information flow property characterizing whether a given "secret" about a system behavior is hidden from an external observer. It is assumed the observer has full knowledge of the system's structure but only partial observability. Based on its observations, the observer constructs an estimate of the system's behavior. The secret is said to be opaque if the observer's estimate never reveals the system's secret. In other words, the system is opaque if, for any secret behavior, there exists at least one other non-secret behavior that looks the same to the observer.

The notion of opacity comes in many flavors (language- and state-based opacity, probabilistic opacity, etc.) and there has been a lot of work in the past two decades investigating the validation and enforcement of opacity for various classes of discrete event systems. To help further progress, I compiled a survey of the state-of-the-art in that area in 2015; I highlighted some future research directions,[10] some of which have been pursued since.

[10] Jacob et al. 2016

Research philosophy

Some things are important to me regardless of the type of projects I work on. These essentially describe my research philosophy.

Quality over quantity Researchers are too often inclined or pressured to "publish one more paper." I think this is not good for people and not good for science. I much prefer working on fewer projects but dedicating the time necessary to make solid and meaningful contributions to the problem at hand. This tightly relates to my second point.

Research is not about writing papers I found the following quote extremely important for computer science and systems research.

> [A]n article about computational science in a scientific publication is not the scholarship itself, it is merely advertising of the scholarship. The actual scholarship is the complete software development environment and the complete set of instructions which generated the figures—Jon Claerbout, as quoted by Buckheit & Donoho, 1995.

In other words, the main value of the research we produce is not the papers we write; it is the algorithms we develop, the systems we design, the experiments we perform. Consequently, it is essential that we do not limit ourselves to publishing papers, but that we also produce and release the corresponding research artifacts.

Open Science is just science done right I am a strong advocate of the principles of "Open Science." I believe it is crucial that publicly funded research is shared for free and to everyone. We must also strive to make our research replicable by others in order to build confidence in science. I am glad to see these values spreading in the newer generations of researchers and I want to foster this in any way I can. At my own level, I made a public pledge to open science[11] outlining a number of self-imposed rules.

[11] romainjacob.net/pledge-to-open-science

Be proactive One obstacle to the normalization of Open Science in computer science is the lack of open access venues with a so-called diamond publishing model—*i.e.*, free to read and free to publish. Therefore, together with colleagues from several fields, we launched a new diamond open access journal: the Journal of Systems Research (JSys). For 2.5 years I served as one of the Editors-in-Chiefs at JSys, which I believe is a great example of researchers taking the lead to provide themselves with what they need.

jsys.org

Collaborate I very much value collaboration, both within and outside of one's research group. Many of my past projects were collaborations, often with colleagues from other universities. Combining our different expertises has always been very beneficial in my research. As sustainability is a complex and inherently transdisciplinary problem, I expect the need for collaborations in my own research to only increase.

Bibliography

For a complete list of publications, you may refer to my Google Scholar profile.

Romain Jacob, Lukas Röllin, Jackie Lim, Jonathan Chung, Maurice Béhanzin, Weiran Wang, Andreas Hunziker, Theodor Moroianu, Seyedali Tabaeiaghdaei, Adrian Perrig, and Laurent Vanbever. Fantastic Joules and Where to Find Them. Modeling and Optimizing Router Energy Demand. In **IMC (under Submission)**. Zenodo, December 2024. DOI: 10.5281/zenodo.14503450. URL `https://zenodo.org/records/14503450`.

Lukas Röllin. How fast can it wake up? Rethinking network hardware and routing protocols. Master's thesis, October 2023. URL `https://doi.org/10.3929/ethz-b-000637984`.

Maurice Béhanzin. **Reducing Power Conversion Losses in Modern PSUs**. Bachelor Thesis, ETH Zurich, December 2024. URL `https://doi.org/10.3929/ethz-b-000715084`.

Lukas Röllin, Romain Jacob, and Laurent Vanbever. A Sleep Study for ISP Networks: Evaluating Link Sleeping on Real World Data. In **HotCarbon**, 2024. DOI: 10.3929/ethz-b-000680859. URL `https://doi.org/10.3929/ethz-b-000680859`.

Tobias Bühler, Romain Jacob, Ingmar Poese, and Laurent Vanbever. Enhancing Global Network Monitoring with Magnifier. In **NSDI**, pages 1521–1539, 2023. ISBN 978-1-939133-33-5. URL `https://www.usenix.org/conference/nsdi23/presentation/buhler`.

Romain Jacob, Jean-Jacques Lesage, and Jean-Marc Faure. Overview of discrete event systems opacity: Models, validation, and quantification. **Annual Reviews in Control**, 41:135–146, 2016. ISSN 1367-5788. DOI: 10.1016/j.arcontrol.2016.04.015. URL `https://hal.archives-ouvertes.fr/hal-01274956`.