

Statistics-Based Optimal Network Path Exploration and Exploitation

[BSc or MSc]

1 Topic

In conventional computer networks, an entity defines source and destination of a network packet and hands it off to the Internet service provider to deliver the packet. In more modern network architectures, it is possible for a network user to choose not only source and destination of the traffic, but also the complete path (i.e., all intermediate hops) taken by the emitted traffic. This characteristic allows an end user to choose which network paths to take in order to optimize the network connection for their use-case (e.g., picking high throughput paths for bulk data transfer and low latency paths for VoIP or time synchronization). In this project you will work on how to explore the characteristics of paths, as well as how to exploit the gained knowledge in order to maximize the performance of a networked application.

2 Task

We assume that when a networked application starts, it knows which paths are available, but it does not know the characteristics of these paths (e.g., bandwidth, jitter, latency, green-ness). In order to explore the qualities of the paths, one requires to probe them. The overall goal is to find the optimal path (e.g., one that maximizes throughput) as quickly as possible with as few probes as possible. There are two fundamental problems that we want to address in a principled way:

- Measurements are noisy: Due to network congestion, queuing, packet processing, and a myriad of other reasons, the measurements we receive from the probes have an irreducible noisy component with means we never have a perfect picture of the path characteristics at any time.
- Measurements become stale: In practice, if you measure that for example the bandwidth of a path is X at time t , it is not clear for how long of a time this measurement is valuable into the future. In order to extract value from past measurements, we require to model how network characteristics can change over time and how this affects the value of past measurements.

3 Methods

Depending on whether its a Master of Bachelor project/thesis, you will apply methods of varying complexity. In essence however we will focus on Gaussian Processes (GPs) / Stochastic Processes (SPs), Bayesian Optimization (BO), and Optimal Control (OC), in that order. We use GPs in order to model our knowledge about the path characteristics over time. If possible we extend it to more general SPs which work with non-normal distributions. We then use BO to guide exploration and exploitation. Lastly, there is the idea of also applying OC in order to develop more sophisticated exploration strategies by constructing more long-term exploration and exploitation plans.

4 Prerequisites

Required: **Bayesian Probability Theory, Statistics, Linear Algebra**
Nice to have: Gaussian Processes, Bayesian Optimization, Optimal Control

5 Contact

If you are interested, write to francesco.dadalt@inf.ethz.ch. Please attach your CV, transcript of records, and a brief reasoning for why you would like this project.