

CheesePi: A Raspberry Pi based measurement platform

Liam McNamara*, Ian Marsh*, Sverker Forslin†

*Swedish-ICT, SICS {ljjm, ianm}@sics.se

†PTS {sverker.forslin}@pts.se

I. INTRODUCTION

We aim to to objectively characterise the service users experience from their home Internet connections. The attributes of an Internet connection (e.g., bandwidth and loss rate) dictate the service quality that can be achieved over it. American video use is increasing rapidly with 70% of broadband users under the age of 35 getting some of their TV from online sources. Measurement of such connections is crucial, their characterisation is useful not only for human users (so people know what service they can receive) but also for the users' devices for adaptive behaviour. Furthermore, large-scale characterisation data of individual connections can be collated into a characterisation of the whole network. In this paper we will present a distributed measurement system that we have built and the choices that comprise its design. The IETF are drafting an RFC specifically for the purpose of large-scale broadband measurement called LMAP [1]. We will use the terminology introduced in that document to describe our system.

Our system, CheesePi, is targeted at the Raspberry Pi hardware device to enable always-on simple and reliable monitoring of users' home internet connection. The Raspberry Pi is a low-cost, credit-card sized computer that has recently gained traction with technology enthusiasts. By running CheesePi on a Raspberry Pi (termed a Measurement Agent or MA) connected to their home network a non-expert user can continuously monitor their connection quality. This device can be easily installed and forgotten about, only being accessed when users perceive a possible problem with their connection. A common hardware platform for all MAs gives greater consistency between the collected measurements, it also simplifies the codebase [2]. The result is a common software platform for Measurement Tasks that can host, execute and record arbitrary network behaviour. A different approach would be to install measurement software on a user's gateway, which would also enable passive monitoring of user traffic [3]. However, we are more interested in capturing the experienced network connectivity that devices are able to achieve. This can be heavily dictated by the last hop technology (e.g. Ethernet or WiFi), which would be missed by monitoring at the gateway.

This work is presented in collaboration with the Swedish regulator Post and Telecom Authority (PTS). Their responsibility is to monitor competition and consumer issues in mail, telephony, the Internet and radio. PTS is particularly concerned with expanding connection performance metrics from naive throughput measurements (Mbits/sec) of a particular location and time to something more instructive. Although an easily comparable and widely understood metric, download/upload rates do not necessarily indicate the quality of service that a user will experience (see below).

II. RELATED PROJECTS

For a survey of related research, we refer to the excellent [4]. The Leone EU project, completed in April 2015, assessed 12 new metrics to measure performance factors using the SamKnows infrastructure of 130 "probes" [5]. One metric they identified relevant to CheesePi is the *bitrate reliably streamed* and they trialed it using measurements to Google's YouTube and the BBC's iPlayer VOD platforms. The framework was standardised known as an "Information Model for Large-Scale Measurement Platforms" [1]. The Archipelago Measurement Infrastructure uses 120 Raspberry Pis deployed to perform network monitoring to quantify topology (connectivity) and dynamics (congestion) using active measurements for both IPv4 & v6 [6]. As of August 2015, the developers ARK UCSD, have 120 nodes deployed. Quantifying the congestion is done by a Time-Sequence Ping (TSP), a crafted sequence of pings along

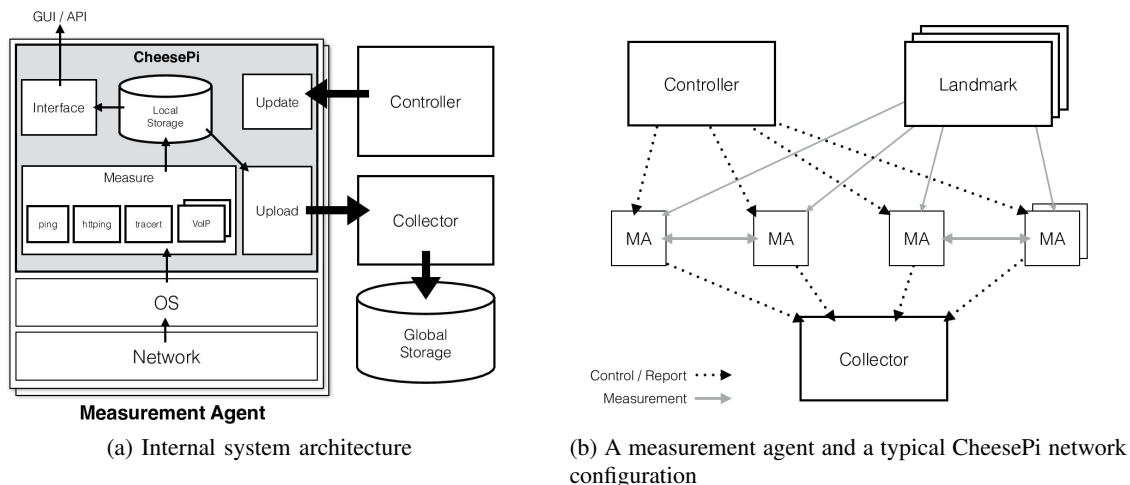


Fig. 1: CheesePi architectural diagrams

a path. Variations in the diurnal variation in delay have been quantified. Fathom has been used in home network environment [7] which has similar goals to ours, but differs in the software/hardware and the external/embedded measurement approaches. Fathom is a Firefox extension that implements the measurement primitives of Netalyzer to measure network characteristics using JavaScript [8]. It adds 3% to the times to load a web page and can timestamp network events to within one millisecond. Measure Broadband America has measurements of 13 ISPs, with 9000 customers, measuring 16 metrics including VoIP and Video streaming measures, which CheesePi also targets. The primarily goal is to show the relative performance of the US's ISPs, more information at www.fcc.gov/measuring-broadband-america.

III. MEASUREMENT TASKS

We aim to create a platform for measurement, whereby a Controller can distribute self-contained Measurement Tasks to be run on MAs. The measurement tools are structured in a modular fashion allowing users to add/submit their own Measurement Tasks. We use a combination of measurements performed from MAs to pre-defined common landmarks (or Measurement Peers) and internally between MAs. The selected external landmarks represent well-connected heavily provisioned domains (e.g., www.facebook.com and www.bbc.com). They are probed with the standard tools: ping, HTTPing and traceroute. Measurements between MAs are conducted in a peer-to-peer fashion for tasks that require cooperation between end-hosts, such as VoIP and VOD streaming. Scheduling cooperation between MAs is non-trivial and can require traversing NAT situations, however the ability for the platform to coordinate collaborative measurements tasks allows more powerful measurements to be easily performed, such as VoIP performance tests. Indeed, we have embedded an automated VoIP call suite that generates quality estimators. Each task exerts a cost on the user's network resources. In future we would like to develop a more adaptive system that aims to optimise its tasks to stay within a data budget while still performing sufficient sampling to capture the network dynamics.

Although not a feature of the access network, a user's local WiFi environment can have a deleterious effect on their apparent network performance. Recording the wireless medium contention can aid analysis of performance and inform users about specific problems with their network. No user traffic is monitored or recorded. To inform the local user we also collect some statistics of the Raspberry Pis internal state, e.g., CPU load and memory. Figure 1 shows how MAs pull measurement schedule and code updates from a remote Controller and finally upload their data to a remote Collector. The process of bootstrapping this architecture onto nodes so they can participate in a distributed measurement infrastructure can be particularly challenging. We simply provide an archive of the distribution on a

secure server and direct the user to run a contained install script. Using Raspberry Pis has the advantage of being able to use the MAC address of the Ethernet interface that is available at a known location on all devices.

IV. STORAGE

Measurements are stored locally on the MAs to allow longitudinal presentation of the connection status without requiring connectivity to the centralised Collector. This is important if the user is experiencing serious network problems. The data is stored in a structured format inside a schemaless database engine, currently InfluxDB but interfaces for other backends are written. The advantage of a schemaless storage engine is that custom built measurement tasks can easily store their output without requiring restructuring of the database. The data Collector accepts periodic uploads from MAs and stores the data in secure and structured storage. This database is processed and collated to provide a queryable information source. Currently this is a monolithic SQL engine but it will be upgraded to cluster storage when we reach the appropriate scale.

V. DATA ANALYSIS INTERFACE

The purpose of operating such a monitoring system is to collect and collate large-scale measurements of peoples Internet connectivity. We believe a user should still gain some utility from running an MA in their home. Therefore we prioritised the development of a visual interface that non-expert users can use to monitor the behaviour of their connection. One of the easiest mediums for this is a website, locally hosted on the Raspberry Pi, that displays graphs and figures accessible via the users web browser on their access device. We plan to extend this interface with simple and expert modes. It is currently user-definable and can be configured through the website to show any combination or processing of stored data. This enables flexible and rapid prototyping of new aspects of interest.




Fig. 2: Local CheesePi Dashboard

As well as a visual interface to the data for human users, we believe it's useful to provide a machine-to-machine interface to detailed characterisation of the connection properties. So that each Internet using application does not need to test and estimate the connection properties themselves. This is unfortunately a common behaviour, where an application will perform a short-lived instantaneous characterisation of the connectivity. Instead, the ability to query a local information source for accurate statistics (e.g., throughput, delay and loss rate etc.) of their Internet connectivity would enable much better configuration for applications behaviour, such as whether to act as a supernode in peer-to-peer networks.

VI. RESEARCH

A. Subjective-objective measurements

One goal in this project is to produce quality ratings from the *network conditions* using both subjective assessment and objective measurements. Defining and refining models, typically by introducing new models, tweaking and/or adding metrics, is very much an active area of research [5]. Our focus is on the conditions that induce buffering/pause events by correlating network conditions to degradations in the perceived quality of media streams. Existing video quality tests do not degrade video sequences by representative network conditions, that is buffering events are not correlated with buffer underrun probabilities. Sessions recorded by the Raspberry Pis will record the media samples, buffer states and subjective ratings by users to provide more representative samples for offline assessment.

B. Shared infrastructure detection

Detecting if separate endsystems share a common path to a server can explain poor correlated spatial and temporal congestion. Where a traceroute cannot reveal topology [9], inference at the end system poses an interesting research challenge [10]. There are implications for regulators, detecting where ISPs collude or could, or for operators, are subsections of our users experiencing poor quality? Our proposal is to use a multi-dimensional mapping and scaling algorithm, t-SNE [11] where a dimension: loss, delay, #hops, technology, represents one measured or labelled factor. The factors will be clustered into a “poor” or “good” quality category. From a machine learning perspective we are actually performing supervised learning on network connections for quality detection. CheesePi will help us in this task by providing a smaller, controllable, observable (to some degree) test network to build an understanding of shared network paths.

VII. PROGRESS AND FUTURE DIRECTION

The project’s current state has a website <http://cheesepi.sics.se>, public code repository and a growing collection of network data metrics. We have measured the delays, both network and web server, and the network path to four major media events¹ via Ethernet and WiFi links, to the streaming server frontends². This SQL collection contains over one million network probes from multiple points.

PTS’s knowledge of the broadband population will be leveraged in the placement of CheesePi nodes. This will ensure sound statistical coverage of the operators access technologies and demographics. Further rollout can recruit Raspberry Pi owners into the CheesePi community to achieve greater scale. We will collect and collate a high quality repository of network data and release it to the public. We believe that network researchers, regulators and designers should collaborate to improve the understanding of digital service dynamics. With some development behind us, we are looking to grow and develop the project given experience from the existing measurement community.

REFERENCES

- [1] IETF, “Large-Scale Measurement Platforms,” <http://datatracker.ietf.org/wg/lmap/charter/>, 2015, [Online; accessed 01-August-2015].
- [2] “EU Regulator: Monitoring quality of Internet access services in the context of net neutrality,” Mar. 2014, [Online; accessed 01-August-2015].
- [3] S. Hätönen, A. Nyrhinen, L. Eggert, S. Strowes, P. Sarolahti, and M. Kojo, “An experimental study of home gateway characteristics,” in *Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement*, ser. IMC ’10. New York, NY, USA: ACM, 2010, pp. 260–266. [Online]. Available: <http://doi.acm.org/10.1145/1879141.1879174>
- [4] V. Bajpai and J. Schonwalder, “A survey on internet performance measurement platforms and related standardization efforts,” *Communications Surveys Tutorials, IEEE*, vol. PP, no. 99, pp. 1–1, 2015.
- [5] “EU 7th Framework program, Large-Scale Measurement Platforms,” <http://www.leone-project.eu>, Apr. 2015, [Online; accessed 01-August-2015].
- [6] D. Clark, S. Bauer, K. Claffy, A. Dhamdhere, B. Huffaker, W. Lehr, and M. Luckie, “Measurement and Analysis of Internet Interconnection and Congestion,” in *Telecommunications Policy Research Conference (TPRC)*, Sep 2014.
- [7] S. Sundaresan, W. De Donato, N. Feamster, R. Teixeira, S. Crawford, and A. Pescapè, “Measuring home broadband performance,” *Communications of the ACM*, vol. 55, no. 11, pp. 100–109, 2012.
- [8] M. Dhawan, J. Samuel, R. Teixeira, C. Kreibich, M. Allman, N. Weaver, and V. Paxson, “Fathom: A browser-based network measurement platform,” in *Proceedings of the 2012 ACM Conference on Internet Measurement Conference*, ser. IMC ’12. New York, NY, USA: ACM, 2012, pp. 73–86. [Online]. Available: <http://doi.acm.org/10.1145/2398776.2398786>
- [9] K. Claffy, D. Clark, and M. Wittie, “The 6th Workshop on Active Internet Measurements (AIMS6) Report,” *ACM SIGCOMM Computer Communication Review (CCR)*, vol. 44, no. 5, pp. 39–44, Oct 2014.
- [10] D. Rubenstein, J. Kurose, and D. Towsley, “Detecting shared congestion of flows via end-to-end measurement,” *Networking, IEEE/ACM Transactions on*, vol. 10, no. 3, pp. 381–395, Jun 2002.
- [11] L. van der Maaten and G. E. Hinton, “Visualizing high-dimensional data using t-sne,” *Journal of Machine Learning Research*, vol. 9, pp. 2579–2605, 2008.

¹Eurovision semifinals and finals, Pacquiao Vs. Mayweather boxing match, the (Internet) leap second and two football cup finals

²including BBC.com, HBO.com, SVT.se