



BSense: A Flexible and Open-Source Broadband Mapping Framework

Giacomo Bernardi, Damon Fenacci, Mahesh K. Marina, Dimitrios P. Pezaros

► To cite this version:

Giacomo Bernardi, Damon Fenacci, Mahesh K. Marina, Dimitrios P. Pezaros. BSense: A Flexible and Open-Source Broadband Mapping Framework. 11th International Networking Conference (NETWORKING), May 2012, Prague, Czech Republic. pp.344-357, 10.1007/978-3-642-30045-5_26 . hal-01531128

HAL Id: hal-01531128

<https://inria.hal.science/hal-01531128>

Submitted on 1 Jun 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

BSense: A Flexible and Open-Source Broadband Mapping Framework

Giacomo Bernardi¹, Damon Fenacci¹,
Mahesh K. Marina¹, and Dimitrios P. Pezaros²

¹ School of Informatics, The University of Edinburgh, UK

² School of Computing Science, University of Glasgow, UK

Abstract. We present, **BSense**, a flexible broadband mapping system for broadband coverage and quality assessment of broadband connections in a given geographic region. For coverage related analysis, it relies on data that is either obtained from ISPs or generated based on technology models and information about infrastructure sites. Broadband quality assessment in **BSense** is via host-based measurements using our multi-platform and technology-adaptive software client which periodically runs as a background process on users' computers. The host-based software measurement approach employed in **BSense** is not only cost-effective but is also flexible and reduces measurement bias. **BSense** also incorporates a flexible broadband quality index for summarizing the collective effect of various underlying attributes such as download/upload speeds and latency. **BSense** system has been implemented based on open-source software components. The usefulness of the **BSense** system is demonstrated using two real world case studies, one on identifying notspots in Scotland and the other on broadband quality assessment in a rural part of Scotland through pilot deployment.

Keywords: Broadband mapping, performance measurement, broadband quality index, broadband access technologies

1 Introduction

Broadband mapping is the process of assessing broadband coverage, quality and market for a given geographical region (e.g., country, province, city). Broadband *coverage* assessment is aimed at identifying “*notspots*”, i.e., locations not serviced by any broadband access technology. For areas that are covered, assessing broadband *quality* in those areas is more interesting. Quality is measured using a set of performance metrics such as download/upload speeds, latency, jitter and packet loss rate. Several technology-specific and network provisioning factors affect quality in practice (e.g., line length, number of concurrent users, contention ratio, backhaul capacity). *Choice* and *cost* associated with broadband subscription are additional aspects that are of interest for broadband mapping, especially to consumers and regulators. In order to determine the amount of choice that a consumer has, one needs to find out various access technologies and Internet Service Providers (ISPs) available at the consumer's location. Greater choice usually

also implies lower cost (per Mbps) for the consumer. Moreover, choice and cost both tend to depend on the coverage and quality aspects — poor broadband coverage or quality in a region correlates well with lack of choice and/or higher costs for consumers in that region. Note that all these four aspects – coverage, quality, choice and cost – vary with time though timescales of change may differ widely between them.

Interest in broadband mapping has been growing recently in tune with increase in consumer awareness and recognition by governments on the importance of high-speed Internet access for all citizens. Different countries have launched national broadband mapping programs (e.g., [1–4]) to quantify the existing state of broadband delivery and to track the progress towards achieving targets set forth in national broadband plans, especially in view of the on-going debate on the role of public funding and regulation in enabling fast and universal broadband access. Beyond these government-initiated efforts, other broadband mapping examples include [5, 6]. *Despite these various efforts, we identify the lack of an open and flexible broadband mapping framework, key to effective and consistent mapping exercises.* Open specification of assessment methodology and metrics is important for audit, whereas the use of open-source software results in lower cost implementations which in turn enable broadband mapping efforts in developing regions. Flexibility is also important to accommodate the wide range of broadband access technologies and to factor in latest advances and best practice in terms of measurement techniques. Moreover, these efforts take one of two different approaches we refer to as *model based* (e.g., [7–9]) and *measurement based* (e.g., [10, 11, 5, 12, 6, 13, 14]) even though *neither is enough* — model based approaches fail to capture the discrepancy between expected and actual broadband quality experienced by consumers, whereas measurement based approaches are clearly not useful for identifying notspots. We further discuss these different broadband mapping approaches in §2.

In this paper, we present an open and flexible broadband mapping framework called **BSense** (Fig. 1(a)) along with its implementation based on open-source software components (§3).

- **BSense** incorporates both model based and measurement based approaches keeping in mind the observation that each is useful for a different purpose. For coverage related analysis, it relies on data that is either obtained from ISPs or generated based on technology models and information about infrastructure sites. Broadband quality assessment in **BSense** is based on host-based continuous measurement stream obtained using our multi-platform and technology-adaptive software client, which periodically runs as a background process on users’ computers. As elaborated in §2, *the measurement paradigm employed in BSense results in a lower-cost and flexible alternative to [10, 11], and reduces measurement bias compared to [5, 12, 14].*
- *Unique to BSense is a flexible broadband quality index* for summarizing the collective effect of various underlying attributes such as download/upload speeds and latency. Specifically, we propose to separately model user preference concerning various performance attributes through a specific instan-

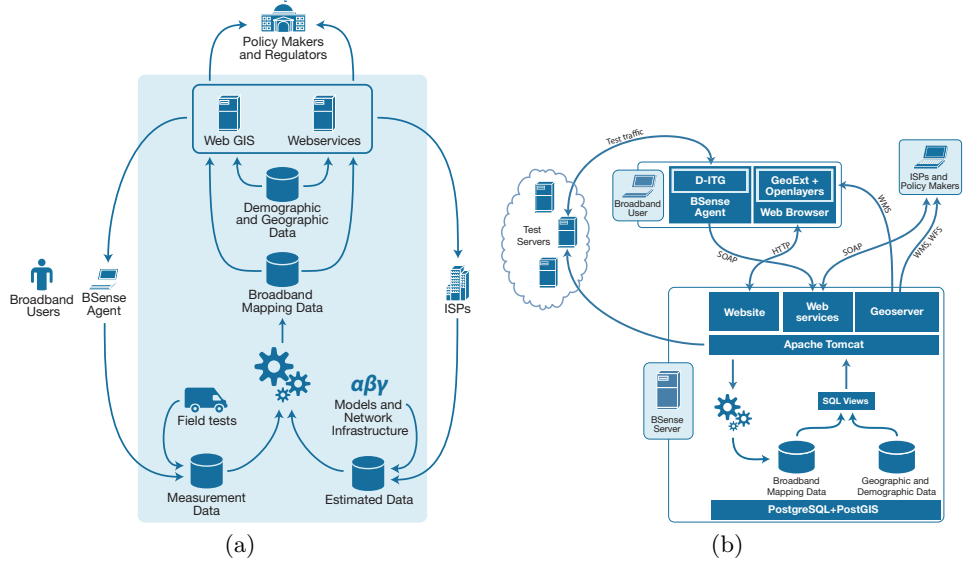


Fig. 1: (a) *BSense* broadband mapping framework illustrated. (b) *BSense* software architecture.

tiation from a flexible family of utility functions and then combine them to produce an overall index by leveraging the general framework of multi-attribute utility theory [15].

To demonstrate the usefulness of the *BSense* system, we use two real world case studies — one on identifying notspots in Scotland and the other on broadband quality assessment in a rural part of Scotland through pilot deployment involving 60 real users over a three month period (§4). Also, evaluation of our measurement methodology shows that it compares favorably to the hardware based measurement approach [10] as well as with a recently proposed sophisticated measurement technique [16]; see [17] for these results. We should, however, clarify that our focus is not on measurement techniques. Though we use D-ITG [18], an open-source measurement technique we have access to, in our current *BSense* implementation, our design is flexible enough to allow the use of other measurement techniques such as [16, 19]. *We regard the BSense framework, its realization using open-source components and the two real-world case studies demonstrating its utility as our key contributions in this paper.*

2 Related Work

Broadband mapping approaches can be broadly classified into two categories: (1) *model based* and (2) *measurement based*. With the model based approach, broadband coverage and speeds are estimated based on theoretical or empirically derived models of access technologies, knowledge of network infrastructure (e.g.,

locations of phone exchanges, mobile network base stations) and configurations (e.g., contention ratio, radio parameters). For ADSL, see [7, 8] for examples of such models and their use in estimating broadband coverage. Similar approaches can be followed to estimate 3G mobile broadband coverage [9]. Such data is inherently optimistic as it does not consider various practical impediments (e.g., line quality, contention). Nevertheless, it is useful for coverage analysis in the absence of any measurement data. Note that this is the approach followed in [1–3] using the data obtained from ISPs.

Measurement based approaches involve actual measurement of broadband connections, a necessity for assessing broadband quality in a region. None of the existing measurement approaches we are aware of rely on measurement data from ISPs. Measurement based approaches can be further divided into *hardware-based* and *software-based* approaches.

The hardware-based measurement approach (also called gateway-based approach) involves deploying a customized hardware box that directly connects to home broadband router for a representative *sample* of users, and using the gathered statistics across all such boxes to estimate statistics for the whole *population*. It is pioneered by SamKnows [10] for UK Ofcom and US FCC sponsored broadband speed studies, and is also considered recently in the academic community [11]. This approach can be expensive than an equivalent software based measurement approach. Moreover, it is also limited in terms of flexibility — unviable for mobile/wireless broadband quality assessment and for analysis at a fine-grained geographic granularity (e.g., city) when planned at national level.

Software-based measurement approaches broadly come in three varieties:

- *Web based*: Using web-based speed tests such as [5, 12] can only gather sporadic and geographically non-uniform measurement data and also suffer from measurement biases (e.g., users taking speed tests may have poor broadband connections or may not belong to a representative sample).
- *Host daemon based*: This approach relies on a measurement agent running in the background on user computers and can overcome the limitations of the web-based approach. However, existing systems following this approach [13, 6] lack openness and flexibility expected from a broadband mapping system.
- *Consumer and ISP independent*: Dischinger et al. [14] present an interesting approach that does not require cooperation from either the consumers or the ISPs. This approach relies instead on certain specific but standard functionality from routers (e.g., responding with TCP RST packets upon receiving unsolicited ACKs). Such functionality may be disabled due to security concerns. If a particular ISP does not support this functionality on all its broadband routers, then that ISP is effectively ignored by this approach, introducing a measurement bias and thus undesirable from a broadband mapping perspective.

BSense uses a model based approach as well as a host daemon based software measurement approach, each for a different use case.

3 BSense Broadband Mapping Framework

As shown in Fig. 1(a), **BSense** framework views broadband mapping as a cooperative exercise involving various stakeholders — consumers, ISPs, policy makers and regulators. In [17], we elaborate on the incentives that can be provided to different stakeholders to continually contribute to the exercise. **BSense** brings together different types of data. Estimated broadband data from ISPs is fed into the **BSense** database via webservice API calls. Broadband users (consumers) are the key source of continuous measurement data for the mapping system. This is enabled by a lightweight software agent termed *BSense Agent* that runs in the background on a user computer and periodically communicates with **BSense Test Servers** to measure technical attributes of user’s broadband characteristics such as download speed, upload speed and latency. Digital geographic data from country-specific sources and demographic data from population census are additionally used as layers underneath estimated or measured broadband statistics to generate broadband coverage or quality maps, respectively.

3.1 Design and Implementation

Fig. 1(b) depicts the **BSense** software architecture, which is described in the rest of this section.

BSense Server The **BSense** server side includes a web server for hosting a public website that users can access for registering and downloading the **BSense** agent software and subsequently to retrieve their broadband connection statistics. The web server also supports a set of web service API calls over SOAP for interaction between the **BSense** system and various stakeholders. The current API consists of the following calls: `BroadbandTestRecord`, `AddPackage`, `EditPackage`, `DeletePackage`, `AddEstimatedData`, `EditEstimatedData`, `DeleteEstimatedData` and `LookUpMappingData`. The purpose of these calls is evident from their names; for further elaboration, see [17]. These API calls are handled by a server side component that enforces security and access control, validating the input and checking whether an API call is made by a party with the required permissions.

In addition to SOAP based web services, **BSense** provides external access via the Open Geospatial Consortium’s standard WMS (Web Map Service) and WFS (Web Feature Service) to obtain raster and vector geo-referenced images, respectively, of a geographical area of interest; most open-source and commercial GIS software products can directly use WMS and WFS services. **BSense** also provides an in-built web application based on WMS, developed using open-source GeoExt (<http://www.geoext.org/>) and OpenLayers (<http://openlayers.org/>) frameworks, to further ease access to broadband maps and their visualization.

Coming to data management, our implementation uses the open source PostgreSQL (<http://www.postgresql.org/>) database management system augmented with the PostGIS (<http://postgis.refractory.net/>) extensions to handle

spatial data. See [17] for further details (database schema, etc.). Access to the **BSense** database via WMS and WFS is enabled by the well known open-source GeoServer (<http://www.geoserver.org>), a Java software that allows users to view and edit geospatial data. Besides broadband mapping data, **BSense** allows the storage of additional geospatial layers with geographic and demographic data that are useful when generating maps.

BSense Agent Each broadband user participating in the **BSense** based mapping exercise runs a software agent (**BSense Agent**) that facilitates continuous and cost-effective measurement of the user’s broadband connection. As such, the agent is a key element of the **BSense** framework for gleaning the quality of broadband provisioning in a given region. Given the diversity of operating system (OS) platforms used by consumers in the real world, the agent should function on different commonly used platforms to avoid measurement bias. The **BSense** agent was designed explicitly keeping in mind this requirement for multi-platform support. Specifically, it has been developed using Qt (<http://qt.nokia.com>), an open-source and cross-platform application/UI framework; Qt is available for Windows, Linux and Mac OS X.

We first give a high-level overview of the measurement process. A participating broadband user would download the agent from a public website like the one we developed (<http://broadbandforall.net>) and install it on the user’s home computer. The agent runs in the background and periodically wakes up to perform a measurement test of the user’s broadband connection³. Each measurement test consists of the following sequence of steps (Fig. 2(a)):

1. The agent queries **BSense** server to get the details of the measurement test to be performed.
2. **BSense** server replies with an “experiment definition” (elaborated below) as well as details for a test server to be used (e.g., IP address, port number).
3. **BSense** server also simultaneously notifies the test server about the impending measurement test from the user’s agent.
4. The agent interprets and follows the experiment definition received, generating the traffic flow requested and/or receiving the incoming traffic to/from the specified test server.
5. Upon test completion, the agent summarizes the traffic traces from the test and uploads it to the **BSense** server.

Since the effectiveness of the mapping framework improves with a larger number of participating users, the mapping system should be scalable and robust to server failures. In the case of **BSense** server, this can be achieved through the use of a server farm (the current approach) or by installing the **BSense** server on a cloud-based hosting. The use of multiple test servers as in our current design also contributes towards scalability and fault tolerance. As regards to the location of

³ The time interval between measurements (i.e., measurement frequency) is a customizable parameter whose setting is a tradeoff between gathering fine-grained measurement samples over time and measurement overhead. The value is set to 15 minutes in the current implementation.

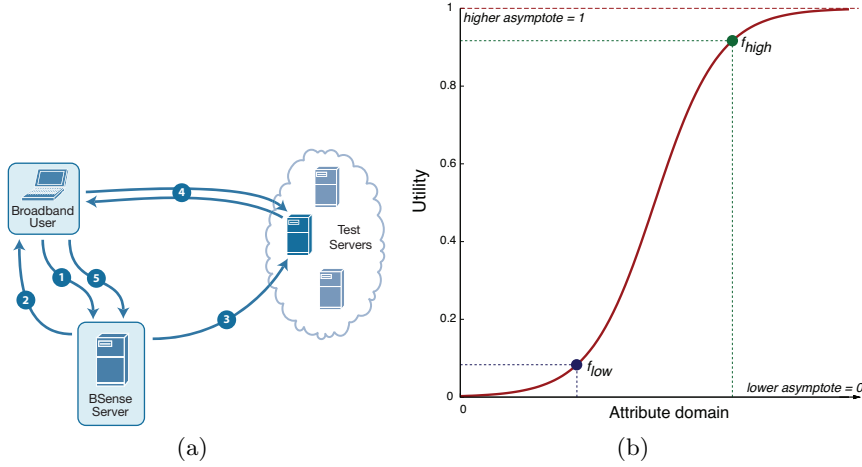


Fig. 2: (a) The sequence of steps making up a measurement test initiated by an instance of the *BSense* agent. (b) The sigmoid utility function can be seen as a transfer function between a given attribute f and the perceived utility associated with specific values of f . In fact, the function shown is a modified sigmoid function given in equation (1) to realize zero utility when the value of f is zero.

test servers, we advocate their deployment at neutral Internet exchange points (IXPs) (e.g., the ones listed at <http://www.euro-ix.net/>) to avoid introducing bias against users of some ISPs. In the current implementation, however, test servers are co-located with the *BSense* server farm.

BSense agent could in principle use any multi-platform network performance measurement tool. In our implementation, we choose a widely used traffic generator called D-ITG [18] with *BSense* Agent acting as a wrapper application. While D-ITG can be seen just as a placeholder in the current implementation, it also has several attractive features that have influenced our choice such as the following: open source; can be made to work on different platforms and behind most common types of NATs with minimal effort; and provides a high degree of flexibility when it comes to traffic generation.

We now briefly describe the experiment specified to the agent every time it is about to do a measurement test. An experiment is defined as a set of traffic session specifications with each session potentially consisting of multiple concurrent or partially overlapping flows. Specifically, each experiment in our context is a sequence of three traffic sessions: initial ping-like UDP session with short packets to measure latency, jitter and packet loss rate, followed by an upstream traffic session and then a downstream traffic session. See [17] for detailed parameter settings. Note that in a host-based software measurement approach like ours, it is possible to have in-home wireless network use and other active use of the broadband connection from other devices within the home concurrently with broadband connection measurement traffic. However, the problem of cross-traffic

on the network path is unavoidable with end-to-end performance measurement even with hardware-based measurement approach. We believe it is best dealt with via statistical filtering.

3.2 Broadband Quality Index

As noted at the outset, coverage and quality are the two key aspects of interest for broadband mapping. While broadband coverage in a particular location can be quantified as a binary variable, the same is not true for broadband quality as the latter is dependent on several underlying technical attributes such as download speed, upload speed and round-trip latency. Due to the lack of standard ways to summarize the collective impact of those several attributes, the focus is often solely on download speeds even though it is widely recognized that upload speeds and latency also have to be considered at the very least.

Defining an *index* is a common approach to deal with problems of the above nature. Only work we are aware of that tries to address the issue of developing a broadband quality index is surprisingly a sociological study [20] relying on expert surveys to determine the relative importance of various technical attributes. More importantly, the work in [20] only provides a very specific approach to defining the broadband quality index while we are interested in defining a more flexible and general framework.

Our main idea in addressing this issue is to model each attribute impacting broadband quality as a *utility function* and then draw upon the *multi-attribute utility theory (MAUT)* [15] to define the *broadband quality index (BQI)* as a composite function of utility function values for the individual attributes.

In our model, denote $F = [f_1, \dots, f_n]$ for the set of network attributes (features) to be included in the BQI, covering important attributes characterizing a broadband connection. For the sake of concreteness, we focus on three key performance attributes in this paper: download speed (d , in Mbps), upload speed (u , in Mbps) and round-trip latency (l , in milliseconds).

We first aim to identify a suitable family of *realistic* single-attribute utility (SAU) functions for modeling user preferences about individual attributes, and then consider their composition into a multi-attribute function. For the attributes under consideration (i.e., download speed, upload speed and latency), we observe that sigmoid functions, whose graphs are “S-shaped” curves (see Fig. 2(b)), better reflect user satisfaction. This is because improvement in utility from improving any of these attributes beyond a point is marginal. Equally, when these attributes are below a certain threshold (for speeds) and above a certain threshold (for latency), the change in utility is again marginal. In between these extremes, the improvement in utility with improvement for any of these attributes is noticeable and substantial.

Thus, we define a set of utility functions $u_f(f)$, each defined on a given attribute $f \in F$ as:

$$u_f(f) = \frac{e^{a+bf} - e^a}{1 + e^{a+bf}} \quad (1)$$

where parameters a and b determine the nature of each utility function curve.

A pragmatic approach to specify a utility function for each attribute would be to have the **BSense** administrator pick two strategic values of each attribute f and provide their corresponding utility u values:

$$(f^o, u^o) \text{ and } (f^*, u^*) \quad \text{s. t.} \quad u^o = u_f(f^o) < u^* = u_f(f^*) \quad (2)$$

These two points can be carefully picked so that they represent the utility of low-end and high-end broadband connections (e.g., with $u^o = 0.2$ and $u^* = 0.8$ or $u^o = 0.1$ and $u^* = 0.9$). Intuitively, the lower knee in the curve represent the value of the attribute which is deemed as insufficient, and the upper knee describes the attribute value which is good enough for the service. As a consequence, “poor” broadband connections that are only able to offer attribute values (e.g., download speeds) below the lower point offer only marginal utility to the users. Similarly, the incremental utility above the upper threshold is only marginal. The low and high values can, for example, be based on current policies and regulations (e.g., the “Universal Service Obligation” to set the bottom bar that ISPs have to provide and the policy maker must enforce) or current state-of-the-art (e.g., the fastest commercially available service to set the top bar).

The parameters a and b for each of our SAU functions can then be derived from their two specified corresponding ‘knobs’ by solving the following pair of equations obtained by substituting values from (2) in equation (1).

To provide a quality index for each broadband connection, we first need to generate a summary statistic from all the measurements that have been gathered by **BSense** for that connection. In the current implementation, we use the median values of download speed, upload speed and latency. These values along with their respective a and b parameter values as input to equation (1) determine the utilities of the broadband connection in question with respect to each of those three attributes.

Multi-attribute utility theory assists us in combining the various SAU functions in a single equation, whose form depends upon the particular independence conditions fulfilled by the different SAU functions. For simplicity, we assume mutual additive independence in this paper. Then the resulting multi-attribute utility function can be represented as:

$$u(f) = \sum_{f \in F} k_f u_f \quad \text{s.t.} \quad \sum_{f \in F} k_f = 1 \quad \text{and} \quad k_f > 0, \quad \forall f \in F \quad (3)$$

A further simplifying assumption would be to have all scaling constants (weights) k_f to be equal. This is reasonable given that our main purpose is to demonstrate the value of multi-attribute utility theory and utility functions in providing a flexible framework for defining broadband quality index.

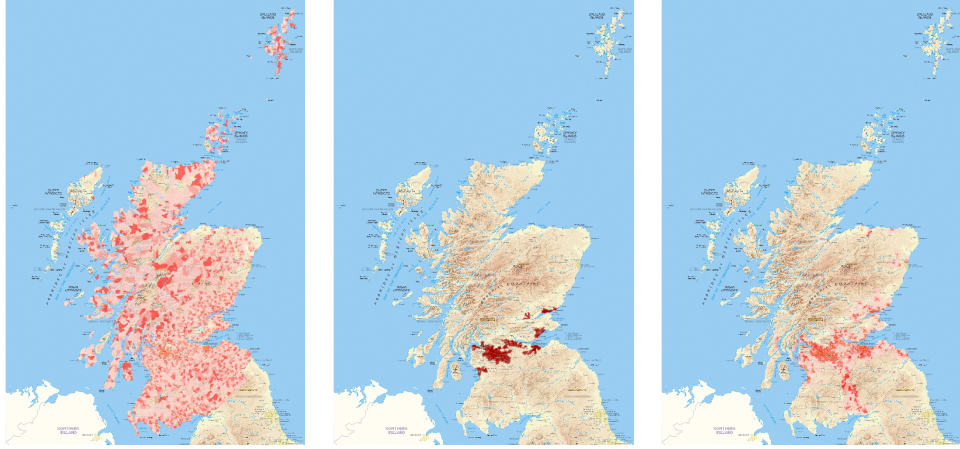


Fig. 3: *BSense* generated broadband map for various access technologies and ISPs based on their estimated data: (a) ADSL — BT Wholesale; (b) Cable — Virgin and Smallworld; (c) 3G mobile broadband — Orange. In each of these maps, postcode areas with the corresponding service are colored with darker colors indicating faster expected speeds or better mobile coverage.

4 Case Studies

4.1 Broadband Coverage Analysis for Scotland

In this case study we show the benefit of **BSense** for understanding broadband coverage, using Scotland as the setting. Such studies would rely upon estimated coverage and speed data from ISPs whenever available. For this study, we mimicked the way ISPs would contribute to the **BSense** mapping system by trawling through the public websites of different ISPs to determine whether an ISP covers a particular postcode and if so, the estimated download speeds from the ISP’s viewpoint, for each of the 152,000 postcodes in Scotland. This information is then fed into the **BSense** estimated database via the web service API calls.

Fig. 3 shows the broadband coverage in Scotland for different access technologies based on the estimated data from ISPs collected as described above. For the 3G mobile broadband case, we show data for only one network operator for clarity but the coverage for other mobile network operators is similar. From these maps, we observe that ADSL is the dominant access technology with cable and mobile confined mainly to population centers in the central belt and north east.

Now focusing on ADSL alone, we estimate the notspots in Scotland with respect to a threshold download speed. Specifically, a postcode area is considered to be a notspot if ADSL service with estimated download speed above the specified threshold cannot be supported within that area. This may be because residences in the postcode area are too far away from their nearest phone ex-

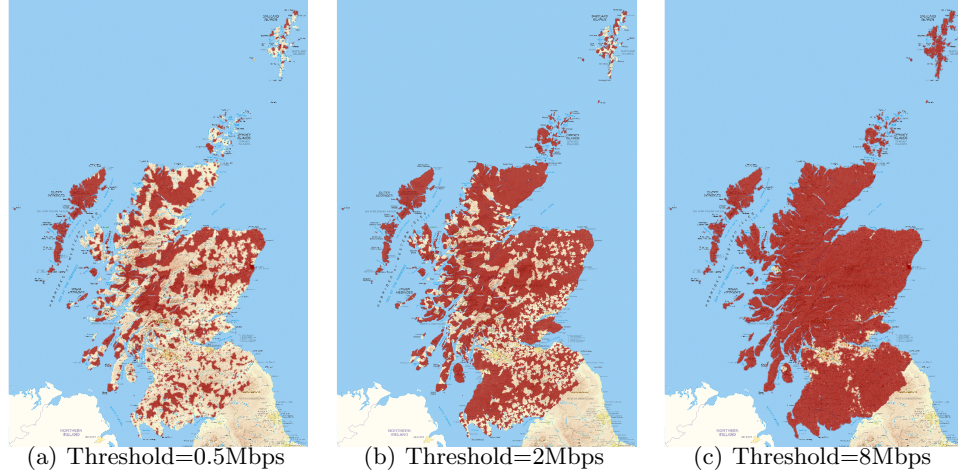


Fig. 4: *BSense* generated map of notspots in Scotland that lack an ADSL broadband service supporting download speed greater than the indicated threshold. Notspot postcode areas are shaded in red.

changes, for example. We consider three different threshold values (512Kbps, 2Mbps, 8Mbps). Resulting notspot maps produced using **BSense** are shown in Fig. 4. It can be clearly seen that most postcode areas outside of the central belt of Scotland (with the two main cities of Edinburgh and Glasgow and having the largest population concentration) become notspots as the threshold is increased. While it is true that satellite based broadband covers virtually the whole of Scotland, the large round-trip latencies associated with the satellite technology (as shown using measurements in the next case study) makes it less attractive.

4.2 Broadband Quality Measurement: A Pilot Study

In this case study, we assess the broadband quality in a rural part of Scotland. Specifically, we focus on the area around the Isle of Skye located in the north-west of Scotland. We also consider the neighboring archipelago of the ‘Small Isles’ and the mainland rural areas of Glenelg and Knoydart peninsulas. This region is quite diverse in terms of demographics, terrain and broadband service provisioning, making it a well suited region for our broadband performance measurement study. It has a population of around 10 thousand people spread across a handful of small towns, several small villages and scores of isolated dwellers in the farming lands. Several different access technologies used for broadband provisioning. In total, 15 phone exchanges are located in the area. Although every resident has access to a landline, broadband connection types vary. A few phone exchanges are enabled for ADSL2/ADSL2+, which is available only from the telecom incumbent (BT). Other exchanges offer ADSL (8Mbps download speed) service, and a few are enabled only for “Exchange Activate” (512Kbps)

ADSL service. There are no FTTH deployments in the area, and cable and 3G coverage are non-existent. Due to a recent broadband reach initiative from the Scottish government, some of the users in rural and remote areas in previously notspot areas in Scotland, including those in our study area, now connect via subsidized yet relatively expensive satellite connections. In addition, residents in a small part of this area connect via Tegola, an experimental/community long distance WiFi network we have deployed four years ago [21].

Through publicity of our pilot broadband quality assessment initiative via email, local press and word of mouth, we managed to find 60 volunteers in the area who were willing to install and run our **BSense** agent software. Half of these users are connected to the Internet via ADSL lines to different exchanges and differing line lengths, whereas 18 users connected via our Tegola network; remaining volunteers used satellite connections. Over a 3-month period, we measured the broadband connections of each of the volunteer users, keeping track of median values of download/upload speeds and latency measurements for each user. We collected around 40,000 measurements in total. To study the broadband quality index across users and access technologies, we used the following parameter settings for the individual utility functions (see Section 3.2), all reasonable given the type of broadband connections in the study area:

- *Download speeds*. Low-end: 2Mbps with a utility of 0.1. High-end: 24Mbps with a utility of 0.9.
- *Upload speeds*. Low-end: 1Mbps with a utility of 0.1. High-end: 5Mbps with a utility of 0.9.
- *Latency speeds*. Low-end: 200ms with a utility of 0.1. High-end: 20ms with a utility of 0.9.

Fig. 5(a) shows the results. The top graphs show the utility function values for each of the three performance attributes. Each data point in the plots corresponds to a user with the color of the data points indicating the access technology used. Clearly and as expected, satellite users have poor utility values and are clustered together at the worst extreme. Wireless users on the Tegola network, on the other hand, not only experience high speeds exceeding 20Mbps but also are subject to greater variability in speeds because of the shared nature of access. ADSL users also exhibit greater variability in speeds like wireless users but because of different reasons — due to differences in broadband capabilities of the associated phone exchanges and differences in line lengths; most ADSL users fall in between satellite users and wireless users in this area. The bottom graph shows the combined effect of the three attributes. For ease of interpretation, we scale up each user’s index value to a percentage value between 0% and 100%. Results for different access technologies shown as CDFs follow directly from the top graphs given our choice for the multi-attribute utility function (an equal weighted sum of individual attribute utilities). When index values for different users are geographically rendered on a map, however, the result is quite revealing (see Fig. 5(b)). Here coloring is done at the ward level — all users belong to a ward (an area with about 50 residents and 20 households) are aggregated together. We observe that remote parts of Knoydart and the Small Isles (colored

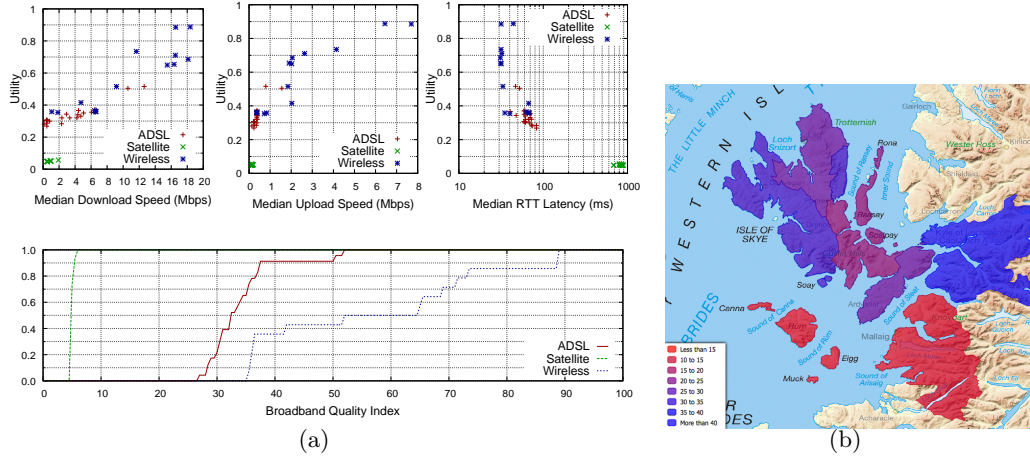


Fig. 5: (a) Pilot study results. Top graphs showing the utilities for different attributes, different users and different access technologies. Bottom graph summarizes the broadband quality index values for users of different access technologies as CDFs. (b) Map generated using the results from the pilot study aggregating users belonging to a ward together and colors representing different ranges for broadband quality index values.

in red) fare poorly, whereas adjacent ward above Knoydart has the best index as a result of high-speed wireless connections from the Tegola network. Wards on the Isle of Skye have intermediate index values as it mainly consists of ADSL users.

5 Conclusions

In this paper, we have developed a flexible framework for broadband mapping called **BSense** that incorporates both model based broadband coverage data and broadband performance measurement data from users. **BSense** framework also incorporates a flexible specification of broadband quality index based on utility functions and multi-attribute utility theory. We implemented **BSense** using open-source tools and use it demonstrate the value of **BSense** approach for broadband coverage and quality assessment with two real-world case studies. Our future work will focus on enhanced measurement techniques that are robust to various sources of variability. We would like to also enhance the broadband quality index by considering additional attributes, variability of each attribute and relationships among various attributes. Finally, we would like to extend **BSense** for mobile broadband mapping.

References

1. US National Broadband Map. <http://broadbandmap.gov/>.
2. UK Fixed Broadband Map. <http://maps.ofcom.org.uk/broadband/index.html>.
3. Broadband Map for Germany. <http://www.breitbandatlas.de>.
4. Ofcom. UK Fixed Broadband Speeds. <http://stakeholders.ofcom.org.uk/market-data-research/telecoms-research/broadband-speeds/speeds-nov-dec-2010/>, Mar 2011.
5. Ookla Speedtest. <http://www.speedtest.net/>.
6. Grenouille. <http://www.grenouille.com/>.
7. T. H. Grubescic. Spatial Data Constraints: Implications for Measuring Broadband. *Telecommunications Policy*, 32(7), Aug 2008.
8. Ofcom. Assessment of the Theoretical Limits of Copper in the Last Mile. http://www.ofcom.org.uk/research/technology/research/emer_tech/copper/, Jul 2008.
9. Ofcom. 3G Coverage Maps. <http://www.ofcom.org.uk/radiocomms/ifi/licensing/classes/broadband/cellular/3g/maps/3gmaps/>, Jul 2009.
10. SamKnows. <http://www.samknows.com/broadband/>.
11. S. Sundaresan et al. Broadband Internet Performance: A View From the Gateway. In *Proc. ACM SIGCOMM*, 2011.
12. Network Diagnostic Tool (NDT). <http://www.measurementlab.net/measurement-lab-tools>.
13. Isposure. <http://www.isposure.com/>.
14. M. Dischinger, A. Haeberlen, K. Gummadi, and S. Saroiu. Characterizing Residential Broadband Networks. In *Proc. Internet Measurement Conference (IMC)*, 2007.
15. S. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*, chapter 16. Pearson, 2010.
16. P. Kanuparth and C. Dovrolis. ShaperProbe: End-to-end Detection of ISP Traffic Shaping using Active Methods. In *Proc. Internet Measurement Conference (IMC)*, 2011.
17. G. Bernardi, D. Fenacci, M. K. Marina, and D. P. Pezaros. BSense: A Flexible and Open-Source Broadband Mapping Framework. <http://homepages.inf.ed.ac.uk/mmarina/papers/EDI-INF-RR-1408.pdf>, Dec 2011. The University of Edinburgh, Informatics Research Report EDI-INF-RR-1408.
18. A. Botta, A. Dainotti, and A. Pescapè. Multi-Protocol and Multi-Platform Traffic Generation and Measurement. In *Proc. IEEE INFOCOM*, 2007. Demo.
19. C. Kreibich, N. Weaver, B. Nechaev, and V. Paxson. Netalyzr: Illuminating The Edge Network. In *Proc. Internet Measurement Conference (IMC)*, 2010.
20. M. Vicente and F. Gil de Bernab. Assessing the Broadband Gap: From the Penetration Divide to the Quality Divide. *Technological Forecasting & Social Change*, 77(5), Jun 2010.
21. G. Bernardi, P. Buneman, and M. K. Marina. Tegola Tiered Mesh Network Testbed in Rural Scotland. In *Proc. ACM MobiCom Workshop on Wireless Networks and Systems for Developing Regions (WiNS-DR'08)*, Sep 2008.