

"Free" Public Access to Internet Services via satellite for rural UK

AIC. Mohideen^{#1}, G. Fairhurst^{#2}, KGD. Tharangie^{#3}, S.T. Nandasara^{##4}

Abstract— Internet Access is regarded by many as a key public service in modern societies. The UK Government is pursuing a Digital by Default agenda as part of its digital communications infrastructure strategy to make Internet connectivity available to all. While it has seen success in enabling connectivity to the majority of the population, the geographical location of the final 5% of hard-to-reach premises pose technical challenges, preventing traditional broadband services from being cost effective. In most remote and hard-to-reach areas, satellite broadband can provide connectivity, however, the present design of service offerings is unable to provide affordable or “Free” Internet service models combined with acceptable performance. One alternative is to offer lower-effort (LE) access to Internet services to provide “Free Internet”. In contrast to terrestrial networks, providing LE services over a satellite network requires optimal cross-layer traffic engineering (TE) to offer acceptable performance. This paper presents an experimental service platform called Rural-PAWS (Rural- Public Access Wi-Fi Services). This was designed to support free high-speed access to government services, with basic access to traditional Internet services over satellite. The Rural-PAWS model uses a prototype multilevel service model that distinguishes Government Digital Services from LE traffic. The platform was deployed in a 12 months study over satellite across 8 pilot sites. This study sought to understand actual service requirements, user perceived performance of web access over satellite-based LE service, while investigating technical challenges to enable deeper understanding of whether it is technically viable to offer such services to hard-to-reach communities in the UK. The experimental results show that an LE service over satellite network can be provisioned, provided that (i) the web browser parameters are optimised for satellite networks and (ii) suitable application layer protocols are enabled over satellite networks.

Keywords— Free Internet, Satellite, Lower Effort, Diffserv, Active Queue Management, HTTP/2.0.

I. INTRODUCTION

Access to the Internet has become a basic need and a legislated right in most developed countries. In April 2002, Vint Cerf stated, “The internet is for everyone” [1]. Governments around the globe have adopted progressive measures to provide Internet access to their citizens. In the UK, the government is pursuing a ‘Digital by Default’ policy under the digital communications infrastructure strategy [2], with the delivery platform for public services moving to online provision [3]. Projected efficiency savings are highlighted as major benefits [4].

Manuscript received on 01st June, 2018. Recommended by Dr. Peter Mozelius on 23rd Oct, 2018.

A.I.C. Mohideen, G. Fairhurst and K.G.D. Tharangie are from University of Aberdeen, School of Engineering. (*alhaff_mohideen@yahoo.com*, *gorry@erg.abdn.ac.uk*, *tharangie@yahoo.com*).

S.T. Nandasara is from Colombo School of Computing. (*stm@ucsc.cmb.ac.lk*).

Concurrently, the UK government's GBP 530 million Broadband UK (BDUK) programme (aligned with the European Union's Digital Agenda [5]) is supporting industry-led deployment of both ‘superfast’ and ‘standard’ broadband across the UK [6]. To date, the government has invested GBP 1.7 billion in its rural broadband programme, that aimed to provide superfast broadband of 24Mbit/s or above to at least 95 % of premises in the UK by 2017 and to ensure that ‘virtually all households’ benefit from a speed of at least 2Mbit/s, again by 2017 – as stated in the ‘Universal Service Commitment’ (USC) [7].

The UK government committed to ensuring the provision of access to basic broadband for all premises, seeking to raise the legal entitlement (Universal Service Obligation) to a service of 5Mbps by 2017 [7]. It is claimed over two million households (11% of users) receive less than 2Mbps [8], such locations being termed “slow spots”. An estimated 160000 households cannot receive broadband at all, or at a reasonable cost, and are in what are termed “not spots” [8].

The provision of universal broadband access to achieve digital inclusion for the “final-few” has proved to be a challenge not only because of the socio-economic barriers associated with non-use of the Internet, but also from technology perspective, because of the difficulties in providing traditional (copper or fibre) digital infrastructure to many currently unconnected or under-served premises [9]. Legacy infrastructure and terrain constraints make next-generation-access (NGA) rollout difficult and/or costlier.

A series of recent papers elaborate on published Ofcom data and outline the geography of broadband availability, from a UK perspective. Accurate reporting of the availability of digital infrastructure at a specifically local level is challenging. Ofcom maps are useful to highlight variation at a regional level – 3G mobile phone coverage and fast, reliable broadband coverage remains poor across large swathes of northern and southern Scotland, northern England, East Anglia, south-west England and Wales, although such maps mask large variation in connectivity within regions.

BDUK efforts have focused on roll-out fixed-line broadband infrastructure and improve mobile phone coverage, making for a dynamic and changing landscape in terms of mapping the availability of digital infrastructure. However, the more remote and sparsely populated an area the more likely it is to experience slow or no broadband connectivity [10]. Nearly one-third of those living in ‘deep rural’ areas of England, Wales and Scotland say that their Internet speed is always too slow for what they want to do [11]; a situation described as ‘Two-Speed Britain’ as described by John Ferrington et al in [11]. This shows an internal digital divide in terms of Internet accessibility.

Satellite broadband offers suitable technology to provide Internet access to remote and hard-to-reach areas constrained by infrastructure and terrain [12]. Further, studies have introduced the notion of Free Internet over satellite or LE services in urban and rural UK [13, 14, 15]. In this study, we

sought to understand actual service requirements, user perceived performance in terms of Webpage load-time over LE or limited basic service (2Mbps) over satellite networks, and to inform whether it is technically viable to offer such services to people in the UK.

Our work defined a prototype multi-level service model that could be adopted by commercial satellite Internet providers, to provide the Rural-Public Access WiFi Service (Rural-PAWS). Implementing this access model showed how satellite networks could offer a service model that combines high-speed “Free” access to government services (4Mbps-up/8Mbps-down) together with basic access (lower effort service) (265Kbps-up/2Mbps-down) for the general Internet service. Our system was built using open tools and an iPOS [18] satellite terminal.

We then explore the implications of these services by studying the web access performance using application layer protocols such as Hyper Text Transfer Protocol HTTP/1.1, HTTP/1.1+TLS (HTTP/TLS) and HTTP/2 (h2 mode). We show how optimising web browsers and suitable application protocols we can maximise web performance using both the BE and LE services by measuring the page load time (PLT) for a range of webpages.

II. FREE INTERNET

The notion of “Free Internet” has existed for over a decade, studied and often associated with digital divide or inclusion from social science perspective. A variety of service models have been studied to enable access to the Internet, over years. Low Cost Denominator Networking (LCD-Net) models [10] and User Provided Network (UPN) [19] models have matured to support the concept of Internet sharing.

In UPN, a micro-provider or an owner of a broadband Internet connection, provides connectivity to unknown users who are within the connectivity radius of the Access Point based on a return-for-sharing incentive [20]. This broadband sharing could use a fixed LE service or on-demand sharing. UPNs such as FON [21], OpenSpark [22], Wifi.com [23] have produced millions of micro Free Internet providers; BT FON provides a fixed 512Kbps [24] over terrestrial networks. The successful growth of UPNs is attributed to the urban nature of dense population where there is availability of super-fast broadband infrastructure.

However, the sparse distribution of premises in rural environments makes the UPN and micro-provider concept impractical and other approaches need to be found to connect these people. A real issue is that the cost of providing service in rural locations is often higher, because the technology required (e.g. satellite has a higher cost per unit of data sent, and also a greater cost of installation). Together this therefore suggests that the offered service may be less, or need to be subsidized. At the same time, the cost of providing government services (education, advice, healthcare, etc.) is also higher for a remote location. This naturally leads to a question about whether the cost-saving from offering free high-speed access from government services could be used to co-fund a basic (lower effort) broadband service. If such a scheme were to be successful, there also could be opportunities for subscribers to transition to high-speed paid services, providing new market opportunities to ISPs.

III. A LOWER EFFORT SERVICE (LE)

IP based networks are ubiquitous and used to deliver a range of services from web access to demanding real-time services. Most common Internet providers offer a consumer service where all traffic is assigned to the Best-Effort (BE) service model.

Commercial customers are often encouraged to subscribe to a service that provides higher assurance of delivery and higher capacity. These services normally are charged with at a higher price, and can preferentially use capacity impacting the performance of traffic in the BE class. A set of services can be realised using the Differentiated Services (DiffServ) model. In this model, packets entering a network domain may already be marked with a DSCP in the IP header, or one is using a traffic classifier. Packets with a specific DSCP are mapped to one of a set of Behaviour Aggregates (BAs). All packets assigned to a BA receive the same treatment, but packets assigned to another BA may preferentially gain access to capacity. The way that a BA is processed is known as a Per Hop behavior (PHB). The IETF has defined a set of standards PHBs [25] e.g.:

- The expedited forwarding (EF) PHB [25] is used where it is desired to reduce as much as possible the queuing delay for a flow in a network node. EF traffic may pre-empt other traffic. The EF PHB is typically used in conjunction with traffic conditioning at the domain edge.

- The best effort (BE) PHB [25] is the default behaviour where traffic is forwarded only if capacity remains after satisfying higher PHBs.

- The assured forwarding (AF) PHB [25] group tries to reduce the probability of packet loss in the case of congestion. Three levels of drop precedence are defined. In congestion, packets with a lower drop precedence are dropped first.

- The Lower Effort (LE) PHB [26] is currently being specified by the IETF. This traffic is known also known “scavenger” and is only sent if there is no other traffic awaiting transmission at the network interface.

Traffic sent using the LE service seeks to not impact BE traffic, and hence the LE PHB will relinquish use of the Internet to other classes of traffic. This attempts to make available capacity that would otherwise be unutilised. Numerous uses have been suggested for the LE PHB, e.g., for background traffic of low precedence, such as bulk data transfers with low priority in time, none time-critical backups, larger software updates, web search engines while gathering information from web servers, etc. We propose using a LE service model for the delivery of Free Internet web access, so that it protects the paid BE and AF traffic.

IV. ACTIVE QUEUE MANAGEMENT OF DELAY

In the Rural-PAWS service model, the upstream capacity for LE traffic is intentionally limited to 256 Kbps. This could become a severe bottleneck to performance for applications using this service, unless advanced traffic management is introduced.

Current web access utilized the Transmission Control Protocol (TCP) [25] to move the data over the network service. Any extra delay introduced by the network could significantly degrade the user experience for a latency-sensitive TCP application, such as a web browser.

Delay can result from a common design used in network devices, in which all traffic shares a simple FIFO buffer in front of a bottleneck router. If this buffer is small, it tends to fill faster and excess packets need to be dropped. A packet drops serves as a signal of congestion, slowing the transmission rate of the sender. On the other hand, if the buffer is large, the traversal time of a packet (“sojourn time”) also becomes large and adds to the RTT, which impacts all flows that share the common buffer (this is sometimes known as “bufferbloat”). The problem is to know the correct size of buffer to use to optimise performance.

A first stage to protecting delay-sensitive traffic is to ensure “free” traffic is assigned to a LE services. This protection the BE and higher assurance services using a Hierarchical Token Bucket (HTB). Explicit bandwidth allocation for BE and LE services can be achieved for each service (BE and LE). However, it does not really control the delay resulting from traffic sent within a class.

Active queue management (AQM) [26] offers a solution to this problem of sharing the bottleneck buffer without incurring unnecessary delay. New AQM methods such as Flow-Queue-Controlled-Delay (FQ-Codel) [28] support flow differentiation. The Rural-PAWS LE service model implements FQ-Codel. This comprises two different queue management algorithms that use two different processes. One algorithm makes a decision whether to packet drop (at the head of the queue) based on sojourn time and a target delay (Controlled-Delay or “CoDel”), while the other enqueues incoming packets onto different sub-flows based on a hashing function (“Flow-Queue”). The hashing function is based on up to 5 elements from the packet header. This combined algorithm helps separate latency-sensitive flows from long lasting bulk TCP flows. Sub-flows on the egress are serviced using byte based Deficit-Round Robin mechanism. The latency sensitive application performance over satellite could be further improved by parameterising the default FQ-CoDel configuration [29].

The implementation in Rural-PAWS model is described further in the section on the Rural-PAWS concept.

V. INTERNET VIA SATELLITE

A satellite network consists of a satellite terminal (ST), one or more (traffic) gateways, the satellite, and the Network Control Centre (NCC). STs are outdoor dish terminals generally mounted at locations with line of site to the satellite. The ST offers an IP local area network (LAN) interface. Cost-effective Internet via satellite has become available to consumers through a satellite system (e.g. [30]). Fig. 1 shows typical components in a satellite network.




Fig 1. Setup of the rural PAWS testbed

Satellite systems are usually configured with a higher capacity available in the forward link (FL), carrying Internet traffic from the satellite gateway to the STs. In contrast, the Return Link (RL) has less capacity and use of this capacity is controlled by the NCC, which allocates capacity by scheduling

ST transmissions using Multi-Frequency Time-Division-Multiple-Access. Typical commercial services offer tens of Mbit/s for the FL, and up to 10 Mbit/s for the RL [32].

Rural PAWS selected a service that used Internet Protocol over Satellite (IPoS) [31]. This is the most widely deployed satellite Internet system based on a proprietary standard developed by Hughes Network Systems.

VI. WEB ACCESS OVER SATELLITE

HTTP/1.1 and HTTP/TLS have been the predominant application layer protocols used for web access until recently. An estimated 83% of web access is performed over HTTP [33]. HTTP/1.1 operational design is based on a request/response model. This poses a major performance hurdle, and constrains the performance of web access due to the growing characteristics and dynamics of webpages and web applications.

One important issue is that request/response design of HTTP/1.1 introduces Head of Line (HoL) blocking. Over a TCP connection, a request for an object can be made only after the response for the previous request is received in full by the application layer.

To alleviate this performance problem, most web client (browser) designs adopt TCP connection parallelism. This allows the download of web objects from a web resource server to be multiplexed. Web browsers each define a limitation on the number of parallel TCP connections that a browser can open towards a server (6 by default in Google Chrome and 8 in Mozilla Firefox). The limitation on the TCP parallelism is necessary to reduce the collateral damaging to other flows sharing an Internet bottleneck.

Over the satellite RL, capacity request and allocation mechanisms can cause additional access delays, further increasing the total round trip time (RTT) above the general 0.5S resulting from the satellite path of geostationary satellite networks. The additional delay can impact Internet services that are based on TCP. The impact of delay is particularly significant for short-lived flows (e.g. web access). The performance of TCP based short-lived flows becomes limited by the window growth of TCP. The TCP window growth is impacted by the increase in path delay.

To counter the performance limits, most commercial satellite networks have introduced Performance Enhancement Proxies (PEP) [36] at the edge of satellite networks as a solution to mitigate the performance obstacle caused by long-path delay. A common PEP design splits a TCP connection. This violates the end-to-end semantics of TCP, but improves performance by allowing TCP enhancements to be tailored for the satellite networks [34], especially to improve web performance.

Browsers could also be tuned. Some browsers define a default static threshold of 250ms before pronouncing a TCP connection idle and starting a new connection (kMaxConnectRetryIntervalMs static parameter in Google Chrome [35] and network.http.connection-retry-timeout parameter in Mozilla Firefox [36]). This default value has been optimised for terrestrial networks, and is not suitable for satellite networks.

More recently, HTTP/2 [37] has emerged. This new design is inspired by the experimental protocol SPDY developed by Google [38]. This removes the design defect of HoL blocking in HTTP/1.1 [37] by introducing a framing layer that offers

bidirectional multiplexing with interleaved requests and responses over a single persistent TCP connection. In addition, the framing layer dependent HTTP/2 binds to TCP's end-to-end semantics when implemented over Transport Layer Security (TLS) [39].

The use of HTTP/2 could change the need for PEPs on a satellite network. HTTP/2 uses a client/server index table to eliminate the re-transmission of objects. The header data compression (using HPACK) further reduces the amount of data transferred. These mechanisms were a common method used in satellite PEPs.

Implementing HTTP/2 over TLS provides secure connection through intermediaries. However, the TLS handshake adds an additional 2 RTTs delay to the TCP three-way handshake (3WHS), on a satellite network any additional delays could significantly degrade the performance when using a satellite network.

VII. THE RURAL-PAWS PILOT SERVICE

A Rural-PAWS pilot was deployed ‘in the wild’ in April 2014. This provided Internet connectivity to four households of ten individuals, with two additional engineering terminals in Aberdeen. In January 2015, this connectivity was extended to another four households located in rural south west Shropshire on the English/Welsh border. Pilot sites focus on households who either had not used an Internet home service or without acceptable infrastructure in the area.

A. Rural-PAWS Design

The Rural-PAWS service was implemented using a wireless router (Netgear WNDR-3800). This ran a custom built OpenWrt [40] operating system (OS), a Linux-like OS for router devices. An integrated application package, called “Whitelister”, Fig 2., was developed for OpenWrt OS using cURL [14] and Dig [42] libraries. The whitelister resolves destination IP addresses of new connections to its domain name indicated in HTTP header.

The Whitelister holds the most recent 15 IP addresses in a cache, when a new IP is detected it is added to the cache head. The name resolving occurs only when a new IP address is found. A fast text fitting algorithm traverses through a list of predefined government services related or “whitelisted” Top-Level Domains (TLD) to match the resolved names from HTTP header. The domains used in this study were: .ac.uk, .gov.uk, .judiciary.uk, .mod.uk, .nhs.uk, .parliament.uk, .police.uk, .sch.uk, .bl.uk, .british-library.uk, .jpsc.uk, .nls.uk, .supremecourt.uk.

The whitelisted traffic is marked with the assured forwarding DSCP (AF42) using a netfilter [43] mangling rule while all remaining traffic was marked with a LE DSCP. Network signaling packets were marked with the expedited forwarding DSCP (EF), ensuring these continued to be passed even when the network was congested, and capacity was exhausted. Fig:2 shows the Whitelister architecture.

A QoS module in the whitelister utilised the Traffic Control (TC) command to implement static bandwidth allocation with queue discipline (qdisc) and classes using Hierarchical Token Bucket (HTB) algorithm [44]. Rural-PAWS model defines two additional diffserv classes, one for network signaling and high-speed service to government services (4Mbps-up/8Mbps-down, for traffic marked with a AF42 and EF DSCP), and one for basic service (265Kbps-up/2Mbps-down, for traffic marked as LE) to traditional Internet services.

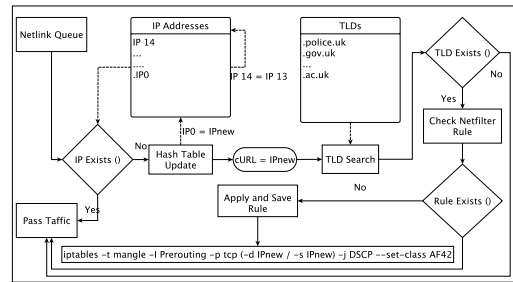


Fig 2. PAWS – Whitelister architecture

FQ-CoDel active queue management algorithm is attached to HTB qdiscs to minimise the queuing delay of competing flows or traffic. FQ-CoDel is a byte based deficit-round-robin algorithm, designed to drop packets from the largest sub-flow upon queue memory exhaustion. Netfilter CLASSIFY target is used for classifying the traffic into qdiscs. Netfilter CLASSIFY target allows classification with more match options than using U32 filters.

Fig 3, shows the traffic classification. The TC rule is applied on ingress and egress. Since the incoming traffic cannot be conventionally controlled, TC is applied on the ingress with the support of the Intermediate Functional Block to avoid packets being dropped under congestion.

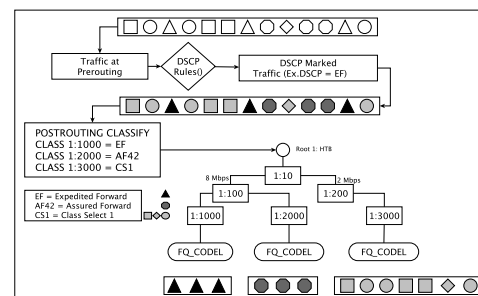


Fig 3. QoS Module in Whitelister

VIII. TRAFFIC MEASUREMENTS AND STATISTICS

A. Methodology

A replica of the Rural-PAWS system was used for our laboratory experiments. Performance was assessed by examining the web page load time (PLT) of synthetic webpages using HTTP/1.1, HTTP/TLS and HTTP/2. HTTP/1.1 has been the de-facto application layer protocol for web access. HTTP/TLS (HTTPS) is used for web access over a secure channel. The new HTTP/2 is a byte stream based successor of HTTP/1.1. All experiments and performance measurements were obtained over an IPOS satellite network. We chose to run our experiments on an operational IPOS system to as our results would reflect a real-world scenario.

Apache 2.2 webserver with self-signed SSL certificate was implemented on an Ubuntu 12.04 LTS platform with kernel 3.17. The HTTP/2 compatibility on the web server was enabled using the SPDY/3 module (mod_spdy) provided by Google. HTTP/2 was not integrated into any open-source web servers at the time of study hence, we utilised mod_spdy. Web-sharding was not considered in the experiments because we expect the next generation of web content to eliminate the advantage of sharding. All experiments were performed using Google Chrome (Vers. 36) as the web browser, and utilized

the benchmarker tool to capture the PLT measurement metrics. The selection of our testbed tools were purely based on availability and compatibility with our purpose of study.

Our results explored a range of webpage size and compositions, by using three cases of webpage lengths (of 500 KB, 1500 KB, and 2500 KB) and three cases of object sizes (5 KB, 20 KB and 100 KB). This produced a total of nine webpages with a homogeneous object size. The number of objects in the tests varied between 5 (for a page size of 500 KB and the object size is 100 KB) and 500 (for a page size of 2500 KB and object size is 5 KB). Each experiment was repeated 10 times to mitigate the variability of download duration on the average PLT. We inherited this dataset from an interdisciplinary study reported in [46].




Fig 4. HTTP/1.1 traffic dynamics for 2500KB webpage over 8Mbps satellite network.




Fig 5. HTTP/2 traffic dynamics for 2500KB webpage over 8Mbps satellite network.

Experiments were executed to assess the impact of the static nature of web browsers on application layer protocols (HTTP/TLS and HTTP/2). The Chrome default static threshold (250ms) was used and then manually changed to 1500ms. The new threshold was selected to be larger two full RTTs over the satellite network.

IX. RESULTS AND DISCUSSION

A. Impact of Protocols on Web Access Performance

The Fig 4 and 5 respectively illustrate the TCP connection dynamics for web access (using a webpage with the length of 2500KB - 500 web objects of size around 5 KB) using HTTP/1.1 and HTTP/2 (rate samples were taken at every 100ms with default chrome settings).

Fig 4, evaluates the Chrome web browser with HTTP/1.1. This initiated a predefined number (6) of parallel TCP connections.

There may be benefit in limiting the number of parallel TCP connections for a typical terrestrial path where a higher aggregate and bursty flow (due to short path) may cause high congestion, packet loss and degrade the overall network performance. However, for a satellite path this approach may limit the performance (even with the use of application layer PEPs). The HTTP/1.1 design requires each object is handled by a separate request/response transaction, as a result each

connection carries around one-sixth of the total number of objects.

The bandwidth-delay product is much higher (about 650 KB) for a satellite path than that for a typical terrestrial path. Since the web objects used in our experiments were of around 5KB in size, each TCP connection could carry a maximum of one object per RTT, and hence could retrieve six concurrent objects per RTT. Therefore, six connections were insufficient to effectively utilize the satellite path, resulting in low throughput (230 KB/s) while accessing the web over HTTP/1.1. Much less than the available 2 Mb/s capacity provided by the Rural-PAWS LE service model.

The HTTP/1.1 performance can be improved by a satellite transport PEP. However, this cannot be used when a web page is accessed over HTTPS (HTTP/TLS), and the long path accessing a webpage with numerous small objects over HTTPS could result in serious performance deterioration.

The TCP window size limits performance when transferring a large object, resulting in large transfer times, since the browser is unable to utilize the available satellite capacity.

The use of TLS with HTTP/1.1 requires two RTTs in addition to the TCP three-way handshake (one RTT). However, a three RTT connection opening is strictly required only the first time the site is accessed. In subsequent accesses the web client could indicate the session-ID of a previous TLS session to resume the session (with session caching, RFC 5246, is supported by the server) or use a session ticket, RFC 5077, which was previously attached by the server for a certain TLS session. Both mechanisms reduce the time to complete the TLS handshake to just one RTT (the abbreviated TLS handshake). The technical barriers of HTTPS could attribute to poor web access performance over saturated satellite networks. Therefore, a cross layer mechanism is needed in web clients to dynamically adjust to the path characteristics.

Fig 5, shows the TCP connection persistency for HTTP/2. Using HTTP/2 the browser initiates a single TCP persistent connection per web resource server. The figure also shows the received/transmitted times.

HTTP/2 opens one persistent TCP connection to multiplex the web objects in frames over multiple virtual bidirectional byte streams. Fig 4 shows the PLT of a webpage of size 2500KB, composed of 500 objects each of size 5KB. This was loaded in around 12 seconds with a throughput of 1.7Mbps.

The improvements when using HTTP/2 are mainly attributed to the bidirectional binary framing, multiplexing and header compression. Binary framing and multiplexing helps avoiding the HoL blocking that existed in HTTP/1.1. Further, HTTP/2 is less verbose than its predecessor HTTP/1.1. HTTP/2 performed well for the satellite network, partly because the satellite networks experience much lower levels of packet than a typical terrestrial network. However, loss of a TCP segment over a satellite network can severely impact the HTTP/2 performance, because the traffic utilizes a single TCP connection.

Figures 4 and 5 represent the HTTP/1.1 and HTTP/2 dynamics over an 8Mbps satellite network. HTTP/1.1 suffers from the design obstacles, whereas HTTP/2 benefits from multiplexing and header compression.

The throughput in HTTP/2 can be limited by the maximum number of concurrent virtual streams allowed by a web server. The default HTTP/2 configuration in the SPDY/3 module allowed only up to 100 concurrent streams per session. A non-optimal value for the connection retry timeout, caused multiple

TCP connections over the satellite network to be opened and connection errors for HTTP/2 at start, in contrast to the use of 1 persistent TCP connection for the entirety of the HTTP/2 session.

B. Impact of Page Composition on Web Access Performance

Fig 6 and 7 show the PLT of webpages with respect to the number of objects retrieved using HTTP/1.1, HTTP/2 and HTTP/TLS and the default setup for a satellite network. The results from the experiments suggest that the number of objects is an important metric in determining the PLT when using HTTP/1.1 and HTTP/TLS. These results were obtained with a FL capacity of 2Mbps and 8Mbps and the default configuration for the Linux TCP/IP stack.

The PLT was higher as the number of objects increased.. The number of objects influences the PLT in parallel to the total page length in HTTP and HTTP/TLS. For example, the PLT for a web page consisting of 500 objects of size 5 KB (the 2500 KB, small objects label in the figure) was approximately 100 s, whereas the PLT for a page of 25 objects of 100 KB (2500 KB, large objects label) took only around 10 s. The dependency of PLT on the number of objects with HTTP/1.1 and HTTP/TLS was attributed to the request/response design pattern of the HTTP/1.1 and partly because of repeated header information.




Fig 6. Impact of webpage composition over 2mbps




Fig 7. Impact of webpage composition over 8mbps

HTTP/1.1 re-transmits header data, an estimated 500-800 bytes of header data is re-transmitted on the path for each request and response, a request of 100 objects with an initial TCP window of 3 would cost an additional 6-8 RTTs to transfer the data. This rate limitation imposed by the HTTP/1.1 request/response model is particularly significant for long delay paths, and is one of the main motivations for using PEPs. The capacity becomes less relevant compared to the number of objects with no contention.

The PLT using HTTP/1.1 is significantly lower than with HTTPS when a page is accessed directly. In reality, this performance is only obtained because the HTTP/1.1 connection is not persistent and represents the worst case for HTTP/1.1. A non-persistent HTTP connection is needed for each object. However, when persistence is enabled, the performance of HTTP/1.1 and HTTPS are similar.

When HTTP/2 is used, the PLT is around 10-12 seconds with any object size (small object, medium object and large) for a 2500 KB web page. Performance was not strongly dependent on the number of objects, because objects could be transmitted concurrently. Further, the HTTP/2 performance did not strongly dependent on page size. Most web pages (500 KB, 1500 KB and 2500 KB) were downloaded in a period between 10-12 seconds. Therefore, HTTP/2 could be suitable for web access over low capacity satellite networks or the Rural Paws LE service.

C. Browser Limitations

In satellite networks, a PEP may break the end to end TCP connection. The PEP can perform the 3WHS with clients and then fetch web object for a client by opening ten connections towards the server. This allowed a PEP to boost performance for HTTP/1.1. The use of TLS with HTTP/1.1 and HTTP/2 over satellite prevents the PEP accelerating the transfer.

The longer delay presented by a satellite path can impact the operation of protocols. In our experiments, the chrome client initiated six TCP connections even for a HTTP/2 web access. After opening the connections, one was chosen for data transfer. The other five connections were later dropped. When the Chrome client was re-compiled with 1500ms kMaxConnectRetryIntervalMs, the number connection drops reduced. This behaviour impacts web access performance.

A browser could be adapted to the satellite path characteristics based on metrics such as initial RTT (iRTT) and historical knowledge on PLT for the betterment of web access performance. Fig 8 shows the behaviour before the change, while Fig 9 shows the behaviour after the change. This pathology was also evident using a Firefox client.




Fig 8. Browser configuration impact on web access over HTTP/2 performance: Before

D. Web Protocol Improvements and Web Access Performance over LE Service

The results show that when a large number of objects are multiplexed, HTTP/2 largely outperforms HTTP/1.1 and HTTP/TLS. In particular, when the web page has more




Fig 9. Browser configuration impact on web access over HTTP/2 performance: After

than hundred objects, HTTP/2 completes the web page transfer in around 10-12 seconds, while a client using HTTP/1.1 or HTTP/TLS spends several tens of seconds using the default settings.

When a web page has fewer than ten objects, HTTP/1.1 performed better than HTTP/2 (HTTP/2 over TLS); this was partly because HTTP/1.1 could move 10 objects in two RTTs after the TCP connection was established, whereas using HTTP/2 with TLS would have required an additional 2 RTTs for TLS handshake before any application data could be transferred.

HTTP/2 performance could have been improved if protocol parameters at the client had been optimally configured for the satellite path. For example, the default TCP send and receive buffer for the HTTP/2 connection was around 130 KB. This did not have benefit when the end-to-end connection was split by a PEP. This was insufficient when the client and server were connected directly without a PEP. Also, the default initial connection size used in the Flow Control mechanism and the maximum number of permitted parallel streams was less than required for optimal use of the satellite network. Finally, the HTTP/2 client connection-retry-timeout was set 250 ms, too small for the satellite case.

HTTP/2 allows servers to send an entire web page over a single TCP connection by introducing methods to multiplex the request/response. HTTP/2 also introduces a mechanism to push data to the client without an explicit request from the client, thus eliminating precious RTTs. The HTTP/2 improvements are complemented by a series of TCP transport enhancements. Examples include TCP Fast Open [97] and the larger Initial Window [98]. Both of these enhancements can significantly reduce the delay associated with completing a short transfer using TCP.

Latency can have a very significant import on web performance. An array of latency reduction mechanisms have been identified that can benefit web sessions [45]. The combination of recommendations removes inessential RTTs. Combined with the new application layer protocol; these modifications could significantly enhance the web access performance over satellite networks.

Proportionate Rate Reduction (PRR) [99] algorithm is an alternative to the widely deployed Fast Recovery and AIMD. PRR seeks to minimise excess window adjustments by growing the actual window size at the end of recovery as close as possible to the ssthresh.

Our results show, the reduction in web access performance is not primarily a result of the limited capacity. The poor performance is largely a result of the design of the web protocol (HTTP/1.1).

The performance over a satellite path can be limited by the use of default network parameters in browsers. These parameters are by default fine-tuned for the typical terrestrial paths. Web browser configurations such as the number of parallel TCP connections, the default connection retry timeout values could be optimized to achieve better performance over satellite paths.

X. CONCLUSIONS

The Rural-PAWS study sought to understand the technical issues and inform the feasibility of providing an LE service using satellite terminals to provide capacity for hard-to-reach communities in the UK, by developing new affordable or free service models combined with acceptable performance. The Rural-PAWS platform was deployed for a 12 month pilot. This was instrumented to collect usage data from eight pilot sites to shed light on actual service requirements, user perceived performance of a limited basic service (2Mbps) or the LE service to enable deeper understanding of whether it is technically feasible to offer a free service to people in remote locations across the UK. The experimental results show that a satellite network based LE service is feasible. User perceived performance is acceptable using current web browsers, and could be optimised for satellite networks.

ACKNOWLEDGMENT

The research reported here is supported by the award made by the RCUK Digital Economy programme to the dot.rural Digital Economy Hub [award reference: EP/G066051/1]. The authors would also like to thanks John Farrington and Fiona Williams for their work developing the Rural PAWS concept and the help they provided in analysis of the pilot study.

REFERENCES

- [1] V. Cerf. "The Internet is for Everyone", IETF RFC 3271, April 2002.
- [2] The digital communications infrastructure strategy, 18 March 2015 [Online] Available at: <https://www.gov.uk/government/publications/the-digital-communications-infrastructure-strategy/the-digitalcommunications-infrastructure-strategy> [Accessed: 5 June 2015]
- [3] Policy paper: UK Government Digital Strategy, December 2013, [online] Available at: <https://www.gov.uk/government/publications/government-digital-strategy/government-digital-strategy> [Accessed: 5 June 2015]
- [4] Research and analysis, Digital Efficiency Report, 2012 Cabinet Office, UK Government, [Online] Available at: <https://www.gov.uk/government/publications/digital-efficiency-report/digital-efficiency-report#annex-1-list-of-case-studies> [Accessed: 10 June 2015]
- [5] Digital Agenda for Europe 2020, 2012, European Commission, [Online] Available at: <http://ec.europa.eu/digital-agenda/en/our-goals/pillar-iv-fast-and-ultra-fast-internet-access> [Accessed: 13 June 20 15]
- [6] Guidance, Broadband Delivery UK, updated September 2013, Department for Culture, Media and Sport [Online] Available at: <https://www.gov.uk/broadband-delivery-uk> [Accessed: 15 June 2015]
- [7] Availability of communications services in UK cities, ofcom, June 2014 [Online] Available at: http://stakeholders.ofcom.org.uk/binaries/research/infrastructure/UK_cities.pdf [Accessed: 20 June 2015]
- [8] Alternative technologies for broadband in rural areas [Online] Available at: <http://www.cswbroadband.org.uk/about-broadband/self-help-guides/alternative-technologies-for-broadband-in-rural-areas/> [Accessed: 20 June 2015]
- [9] Infrastructure Report 2014 Ofcom's second full analysis of the UK's communications infrastructure [Online] Available at: <http://stakeholders.ofcom.org.uk/binaries/research/infrastructure/2014/infrastructure-14.pdf> [Accessed: 25 June 2015]
- [10] L. Townsend, A. Sathiaseelan, G. Fairhurst, C. Wallace, "Enhanced broadband access as a solution to the social and economic problems of the rural digital divide", *Local Economy*, 2013, 28(6) 580-595

- [8] John Farrington, Lorna Jennifer Philip, Caitlin Doyle Cottrill, Pamela Abbott, Grant Blank, William Dutton, “Two-Speed Britain: rural Internet Use”, 2015, ISBN 978-1-85752-034-7
- [9] Lutz, Erich, Markus Werner, and Axel Jahn., “Satellite systems for personal and broadband communications”. Springer Science & Business Media, 2012.
- [10] Sotirios-Angelos Lenas, Vassilis Tsaoussidis, “Enabling Free Internet Access at the Edges of Broadband Connections: A Hybrid Packet Scheduling Approach”, *Mobile Computing and Communications Review*, Volume 18, Number 1, January 2014
- [11] Arjuna Sathiaseelan, Jon Crowcroft, LCD-Net: Lowest Cost Denominator Networking, *ACM SIGCOMM Computer Communication Review*, Volume 43, Number 2, April 2013
- [12] David Ros, Michael Welzl, ‘Less-than-Best-Effort Service: A Survey of End-to-End Approaches’, *IEEE Communication Surveys & Tutorials*, vol. 15, no. 2, second quarter 2013
- [13] Hughes Networks, “IP over Satellite (IPoS) - The Standard for Broadband over Satellite”, 2007, [Online] Available at: <http://defense.hughes.com/resources/ip-over-satellite-ipos> [Accessed: 1 July 2015]
- [14] Secchi, R., Mohideen, A. & Fairhurst, G. "Evaluating the Performance of Next Generation Web Access via Satellite" Wireless and Satellite Systems. Pillai, P., Hu, Y. F., Otung, I. & Giambene, G. (eds.). Springer , pp. 163-176 14, 2015
- [15] R. Sofia, P. Mendes, ‘User-provided networks: consumer as provider’, *IEEE Commun. Mag.*, v .46, pp. 86–91, 2008
- [16] FON Community, [Online] Available at: <http://www.fon.com> [Accessed: 2 July 2015]
- [17] OpenSpark Community, [Online] Available at: <http://open.sparknet.fi> [Accessed: 2 July 2015].
- [18] Wifi.com, [Online] Available at: <http://www.wifi.com> [Accessed: 2 July 2015]
- [19] BT FON Community, [Online] Available at: <http://www.btfon.com> [Accessed: 2 July 2015]
- [20] Babiarz J, Chan K, Baker F, “Configuration Guidelines for DiffServ Service Classes”. IETF RFC 4594, Aug. 2006.
- [21] A Lower Effort Per-Hop Behavior (LE PHB) [Online] Available at: <https://tools.ietf.org/html/draft-ietf-tsvwg-le-phb-04> [Accessed: 5 July 2015]
- [22] TRANSMISSION CONTROL PROTOCOL, IETF RFC 793, September 1981
- [23] IETF Recommendations Regarding Active Queue Management, IETF RFC 4594, July. 2015.
- [24] Kokku, R., Bohra, A., Ganguly, S., and A. Venkataramani, "A Multipath Background Network Architecture", *Proceedings of IEEE INFOCOM 2007*, May 2007.
- [25] T. Høiland-Jørgensen, “Technical description of FQ CoDel,” February 2014. [Online]. Available: https://www.bufferbloat.net/projects/codel/wiki/Technical_description_of_FQ_CoDel [Accessed: 15 July 2015]
- [26] Kulatunga, C., Kuhnz, N., Fairhurst, G. and Ros, D., “Tackling Bufferbloat in Capacity-limited Networks”, *Networks and Communications (EuCNC)*, 2015 European Conference, doi:10.1109/EuCNC.2015.7194103
- [27] ETSI EN 301 790. "Digital Video Broadcasting (DVB); Interaction Channel for Satellite Distribution Systems", 2003.
- [28] ETSITS102354, “BSM; Transparent Satellite Star-B(TSS-B); IPover Satellite(IPoS)Air Interface Specification”, 2005.
- [29] DVB-RCS2 Return Channel Satellite DVB Fact Sheet - August 2012
- [30] HTTP Archive, [Online] Available at: <http://httparchive.org> [Accessed: 15 July 2015]
- [31] Luglio, M., Roseti, C. and Zampognaro, F., “Performance evaluation of TCP-based applications over DVB-RCS DAMA schemes”, *Int. J. Satell. Commun. Network.*, 27: 163–191. doi: 10.1002/sat.930, 2009
Google Chromium, [Online] Available at: https://code.google.com/p/chromium/codesearch#chromium/src/net/socket/client_socket_pool.h&q=kMaxConnect%20file:%5Esrc/net/socket/&sq=package:chromium&l=170 [Accessed: 20 July 2015]
Mozilla Firefox Bug, [Online] Available at: https://bugzilla.mozilla.org/show_bug.cgi?id=592284 [Accessed: 25 July 2015]
- [32] M. Belshe, R. Twist, M. Peon, and A. Melnikov, “Hypertext Transfer Protocol version 2”, IETF Internet Draft, draft-ietf-httpbis-http2-17, Feb 2015.
- [33] M.Belshe, M. Peon, R. Twist: SPDY protocol - draft 3, 2012 [Online] Available at: <http://www.chromium.org/spdy/spdy-protocol/spdy-protocol-draft3> [Accessed: 25 July 2015]
- [34] The Transport Layer Security (TLS) Protocol, IETF RC 5246, August 2008
- [35] OpenWrt, [Online] Available at: <https://openwrt.org/> [Accessed: 8 August 2014]
- [36] cURL, [Online] Available at: <https://curl.haxx.se/> [Accessed: 10 August 2014]
- [37] dig, [Online] Available at: <http://linux.die.net/man/1/dig> [Accessed: 14 August 2014]
- [38] Netfilter/iptables, [Online] Available at: <http://www.netfilter.org/> [Accessed: 15 August 2014]
- [39] TC, [Online] Available at: <http://lartc.org/manpages/tc.txt> [Accessed: 20 August 2014]
- [40] Briscoe, B., Brunstrom, A., Petlund, P., Hayes, D., Ros, D., Tsang, I., Gjessing, S., Fairhurst, G., Griwodz, C. and Welzl, M., “Reducing Internet Latency: A Survey of Techniques and their Merits”, *IEEE Communications Surveys & Tutorials*, [Online] Available at: <http://dx.doi.org/10.1109/COMST.2014.2375213> [Accessed: 25 July 2015]