# Tunnelling the Internet

Habia Akter [1]

School of Electronic Engineering and Computer Science Queen Mary, University of London, London, United Kingdom

Chris Phillips [2]

School of Electronic Engineering and Computer Science Queen Mary, University of London, London, United Kingdom

**JTDE - Vol 7, No 1 - March 2019** [3]

[4]

☆ 110 [5]

**Abstract**

Despite a considerable increase in Internet capacity, regional congestion is still an issue at certain times of day. Dimensioning the system to provide minimal delay under these transient conditions would be uneconomical. We therefore investigate a scheme that allows end-users to selectively exploit a sequence of mini-tunnels along a path from their origin to a chosen destination. Such tunnels can be advertised centrally through a broker, with the cooperation of the Autonomous System (AS) domain operators, similar to a driver choosing to use a toll road

to avoid potential congestion. It is thus a type of loose source routing. The approach avoids the need for inter-operator cooperation, although such cooperation could enable extending tunnels across AS peers. We explore the benefit in delay reduction for a given concentration of tunnels within a portion of the Internet. We show that a relatively small number of tunnels can provide worthwhile improvements in performance. We consider both when tunnels are randomly distributed and when they are provided close to an AS domain of interest, where traffic congestion is more likely. In this latter case, even a relatively small number of tunnels can benefit a reasonable number of users across a large region.

# Introduction

Although the Internet has proved to be robust and flexible, the delivery of time critical data traversing multiple Autonomous System (AS) domains remains sub-optimal due to the unwillingness of the network operators to support inter-operator signalling coupled with the control of the associated forwarding infrastructure (Schaffrarh, 2009 [6]). Mechanisms for such signalling have been proposed, with functional entities such as the ITU Resource and Admission Control Function (RACF) and the IETF Path Computation Element (PCE) (Chamania, 2012 [7]). Despite the proposal and refinement of these operator-owned control plane entities over many years (Yost, 2015 [8]; Rzym, 2016 [9]; Dasgupta, 2007 [10]), their adoption outside the academic community is no nearer.

The focus of this research is on improving the end-to-end communication performance of the Internet, driven from an end-user perspective. This work is not concerned with establishing end-to-end paths or tunnels: rather, the aim is designing a scheme where short tunnels are made available across individual ASes, and possibly between adjacent ones. These are advertised to the end-users through a âService Brokerâ, providing users with an opportunity to use them, for a small fee, if they wish to do so. To facilitate this, we envision an entity at the usersâ access point that would select the most âappropriateâ path for a given data stream, depending on constraints such as the amount of money the user is ready to pay, the end-to-end delay, and the flow content, for example.

This paper is an extended version of the paper âAre Internet Tunnels Worthwhile?â presented at 28[th] International Telecommunication Networks and Applications Conference (ITNAC), 2018. It includes additional material examining the focused deployment of tunnels.

## Motivation for Tunnels

The main motivation for tunnelling over segments or the entire end-to-end path across the Internet is to overcome limitations inherent in traditional next-hop forwarding. With next-hop forwarding the path taken by the traffic is determined by the router node at each âhopâ using information held in its Forwarding Information Base (FIB). The FIB data is typically constructed based on automatically configured routing information obtained via intra- and inter-gateway routing protocols along with operator policy filtering (Farrel, 2004 [11]). This presents two key issues. First, end-users have no say in how their data is forwarded. Second, lack of traffic differentiation means that information flows along paths based on a simple âleast costâ metric lead to load imbalances and âbest effortâ equal treatment of all traffic, irrespective of its importance to the user.

A tunnelling mechanism, e.g., a classification and label switching mechanism, can be used to address both of these issues. Tunnelling has already been implemented using various technologies (Secci, 2008 [12]). We aim to give end-users some control over choosing the path their data flows through by making the presence of these tunnels visible, advertised centrally using a broker, along with a means of steering traffic between them. Although the broker we are proposing is expected to know where the tunnels are, along with their characteristics, it does not need to know how the tunnels are established or operated. Operator security is not compromised, as the details of the technology used to provide the tunnels is hidden, and their establishment and maintenance remains fully under the control of the operator.

Users can choose to use the tunnels, if they wish, for a nominal fee. The idea of charging customers for better service is not new (Doctorow, 2014 [13]). However, in our case, choosing to use tunnels is optional and it is up to the user which specific flows are directed through them. As such, some customers may be happy to selectively pay to obtain flow transport with a better Quality of Experience (QoE).

In our proposal, the end-user will be the one to decide whether specific tunnels will be used or not, knowing the âfinancial costâ and the expected benefits. Operators are expected to cooperate, as they receive extra revenue by providing the tunnels. However, these tunnels, at least initially, only straddle ingress to egress points of specific AS domains between AS Border Routers (ASBRs). The location, delay, cost and perhaps resilience of these tunnels (comprising an IP address of the ingress ASBR and additional information) are passed to the broker. An entity at the end-userâs access point can see the information advertised by the broker and optionally decide to direct traffic flows via one or more tunnels if the perceived benefits are sufficient relative to the cost involved.

## Net Neutrality

The term ânet neutralityâ was first used in 2003, by Tim Wu, as augmentation of the idea of âcommon carrierâ (which transports data for any person or company with taking the responsibility of any possible loss) for telephone systems (Wu, 2003 [14]). Net neutrality is the idea stated as âall Internet traffic should be treated equallyâ (Honan, 2008 [15]). According to this idea, Internet Service Providers (ISPs) and the governments regulating the Internet treat all the data equally, without making any discrimination or taking different charges by user, content, website, platform, application, type of attached documents (e.g., emails, audio, video), or mode of communication (Rushe, 2017 [16]). Hence, according to this policy, the ISPs cannot prioritise any data over the others while sending it from the source to the expected destination.

Thinking generally, it can be easily stated that a few millisecondsâ delay while sending an email will not bother the sender or receiver much. On the other hand, the same amount of delay in a live video streaming flow can have a noticeable negative impact on the Quality of Experience (QoE) of the user.

A motivation behind our research is that, from the point of view of an end-user, treating all the traffic in the Internet equally creates problems. There is much debate concerning how different traffic should be treated. However, our vision is not about charging users for the services, rather it gives the users the opportunity to choose if they want to pay for getting a better service and also provides some control over how their traffic moves across the Internet. Hence, in a way, we are not against net neutrality: rather, we aim to give more control to the users to decide how they want their traffic to be handled by the Internet.

# Example Framework

## The Tunnel Framework

The basic architecture of the AS-Domain tunnelling framework is shown in Fig. 1. The tunnels shown are assumed to have been set up and maintained by the specific network operators using whatever means they wish. This could involve the use of PCE/RACF signalling; however, this is not essential. The presence of the tunnels is advertised via the Directory Service Broker (DSB). This is explained briefly in the section âBroker Functionâ. The tunnels can be of any technology, though it is expected that many will employ Multi-Protocol Label Switching (MPLS) or be based on optical channels. These tunnels can be both intra- and inter-AS in scope, in the latter case this being achieved through operator peering. Some tunnels may offer 1+1 protection; others may exist between peering operators through Label-Switched Path (LSP) stitching. Unlike the usual traffic sending process, in a 1+1 architecture (spoken of as âOne Plus Oneâ), dual copies are sent through two routes in parallel so that, in case of any network failure, the alternate route can be chosen for receiving the packet flow . However, details of the construction mechanism are considered outside the scope of this research.

Initially, customers for this service are assumed to be Small and Medium Enterprises including financial institutions that wish to transport data quickly without having to incur the costs associated with a leased end-to-end infrastructure. They will have awareness of the sequence of AS domains that their data is passing through and possible alternatives, particularly if Border Gateway Protocol (BGP) reachability information is made available to them via the DSB Internet Map.

Their IT administration, which could be automated software that performs path selection based on cost and other requirements, may wish to choose a preferred path between their own site and a given destination, such as between Customer A and B in Fig. 1. For example, by interrogating the information in the DSB, Customer A wishes to use Tunnel T1 and T3 to hasten the delivery of data between the two sites. Having informed the DSB of this decision, for a small fee Customer A is given tickets for each of the tunnels (i.e. T1 and T3) along with their ingress IP addresses. Tickets are ephemeral so it is unlikely that users can abuse the system extensively.

[17]

Figure 1. User-Selectable AS Domain Tunnelling Framework

## Network Operator Functions

Tunnels traversing multiple domains are hampered by the unwillingness of network operators to support inter-operator signalling, coupled with the control of the associated forwarding infrastructure. However, our system does not depend on any information that the network operators will not share with the DSB due to âtrustâ issues. We assume that cooperative operators will let the broker advertise their available tunnels centrally, with some information such as where the tunnels are, how much the âusage chargeâ is, and what performance they offer to the users, typically in terms of guaranteed traversal delay relative to the no-tunnel alternative. The approach does not require the sharing of any sensitive information such as the mechanism by which the tunnels are established and maintained. However, AT&T and CAIDA already provide some information concerning dynamic network performance (Statista, n.d. [18]; CAIDA, n.d. [19] Huffaker, 2012 [20]).

Our system deals with the tunnels existing between Autonomous System Border Routers (ASBRs) belonging to the same AS. There is no need to know about the internal path of the tunnels. Tunnels straddling AS domains are considered optional, as they would require a peering relationship between operators. However, mechanisms for stitching together LSPs across AS domains could technically be provided if sufficient trust existed between adjacent operators.

## Broker Function

The idea of implementing a service broker has already been proposed in the field of telecommunication to make offers of best service against customer requests (Plummer, 2011 [21]). In (Markakis, 2016) [22], a tunnel broker system is discussed that minimizes the job of the tunnel server by assigning the broker server that handles the user requests and returns the prime configuration to both users and tunnel servers.

We introduce a DSB in order to provide a centralized resource for advertising the AS tunnels to the end-users, giving them the opportunity to choose to some extent their desired path across the inter-network. We assume it will have the map view of the ASes that the broker can show the end users, indicating which ASes are adjacent to each other and, in the case of cooperative ASes, information concerning their tunnels will be included. The broker presents the location of tunnels to the users superimposed on an AS view of the Internet (or a portion of it) and the users have the opportunity to choose whether their traffic is directed through one or more tunnels in a particular sequence. This provides a form of loose source routing. Furthermore, certain ASes may show information concerning their degree of congestion. This allows the end users to selectively choose to use a tunnel to detour traffic away from the congestion, or to provide preferential treatment across the congested AS.

To clarify more, the DSB does not retrieve topology maps. It generates a map view from information that is either passed to it from the ISPs or which can be obtained using âtracerouteâ and/or BGP update messages. We naturally assume that the operators that are willing to cooperate will also pass some information saying whether the tunnels or their default forwarding environment are busy at a particular time and this information can be made available in the proposed brokerâs map view. In short, the DSB provides an Internet Map showing the tunnel locations, their usage charge and some statistics regarding the performance they offer.

However, the broker does not tell the user how to get across the network. It provides a view of the topology, with the cost. The map can also include ASes present in the network that are not cooperating with the broker. In this instance, only their ASBR interconnection with other ASes will be available.

The DSB also provides a single brokerage point whereby the user can request a sequence of tunnel permits (tickets) so that traffic can use a tandem arrangement of multiple tunnels between a source and destination. The DSB is effectively the customer-facing entity where operators advertise their tunnels and the transactions that can be made.

## End-User Function

The end-user would be expected to install software in his/her network. The software would obtain the visualization part of the Internet map from the DSB. It also needs to know where the tunnels are available for the users and what are the tunnelsâ starting and end points. The software will get some information from the user, e.g.:

- The source and destination ASes for the data of the user to be sent.

- Expected service of the users, where delay and other constraints will be the measure.
- Amount of money the user wants to pay.

Knowing the preferences of the user, the software will be able to:

- Tell the least cost path using Dijkstraâs Algorithm.
- Find the path with tunnels.
- Compare the constraints and the financial cost.
- Suggest the better route for the traffic.

Initially, the decision will be made depending on two factors: the benefit and the (financial) cost.

The user software will get the same visualization of the Internet map from the broker, including the location and expected service provided by the tunnels and the usage cost. However, details of this mechanism and how the financial model operates are considered beyond the scope of this paper. This paper focuses on assessing the magnitude of the benefit in terms of delay performance, if such tunnels existed.

# Evaluation

## Implementation

A framework has been constructed to investigate the benefits of using different percentages of tunnels present in a part of the Internet for sending data from one AS to another.

Some regional internet topologies at the AS level are generated using the topology generator PFP (Positive Feedback Preference) developed by Mondragon and Zhou in 2004 (Zhou, 2006 [23]) and used as input for the tool we have developed. The main reason for choosing PFP is that it is a phenomenological model for AS-level internet topology which can precisely reproduce a number of topological characteristics, e.g., degree distribution, rich club connectivity, maximum degree, shortest path length, short cycles, disassortative mixing and betweenness centrality (Zhou & Mondragon, 2004 [24]). The PFP model starts from a small random AS-graph and keeps growing where, at each step, new nodes are attached to old nodes and old nodes also peer with other old nodes (Zhou, 2006 [23]). The probability of a node gaining a new link, which is a function of the node degree, is calculated as 0.048 (Clegg, 2010 [25]). The more links a node has, the more is its chance to obtain further links. The developers of PFP have explained the consequence as âthe rich not only get richer, but they get proportionately richerâ (Zhou, 2006 [23]; Clegg, 2010 [25]).

The AS topology developed from the PFP is fed into the framework developed for this research, which then produces another topology at the level of ASBRs, assuming there is a peering of border routers (formed by one from each of the connecting ASes) at the point where the two AS domains are connected to each other. We are aware that the route between the adjacent border routers of two connecting ASes does not necessarily have to be one-to-one; rather, there can be one-to-several connections. However, we currently confine ourselves to one-to-one ASBR peering.

Moreover, within a single AS, the border routers are inter-connected into a full mesh, but the connections need not necessarily be direct; rather, more than one internal hop may exist between a pair of border routers. Our system does not require this knowledge, nor do operators need to share this information. Therefore, the topology view of the broker is not necessarily a complete one. We can call it a âsanitizedâ or an âartificialâ view of the Internet map. It just shows how the various ASes are inter-connected at the AS level. Depending on this AS view, the ASBR topology is produced.

We then use Dijkstraâs Algorithm to calculate the least cost routes for traffic to be sent from any source AS to any destination. The paths include the ASBRs that the traffic needs to traverse to reach the destination. This is the no-tunnel least cost path.

For now, the cost of the routes is considered using the metric of âdelayâ in milliseconds. A data packet typically needs to go through 4 to 6 hops within a given AS while traversing across a number of ASes to reach the destination (BegtaÅ¡eviÄ, n.d. [26]). Hence, an intra-AS tunnel having the ingress and egress points in the same AS can reduce the delay that is experienced relative to the normal no-tunnel intra-AS links. This is particularly true if the normal pathways are congested and some form of priority is given to the tunnels, be that through the use of separate optical channels or queueing priority along shared links.

Fig. 2 illustrates a simple example of alternate no-tunnel and tunnel paths within an AS.

[27]

Figure 2. Example Intra-AS path with and without Tunnels

In Fig. 2, the source and destination ASes are S and D and the traffic is assumed to traverse through another AS to reach the destination, which has a tunnel T with ingress point A and egress point B. The dotted lines represent a normal intra-AS pathway including routers inside the AS.

For now, along each of the links, the associated cost is the (mean) delay in milliseconds. The four types of delay contributing to the total end-to-end delay are: transmission (Tx) delay, propagation delay, processing delay and queueing delay. The propagation delay between the ASBRs A and B will be same for the no-tunnel normal path and the tunnel. Ramaswamy (2004 [28]) shows that the processing delay matters although both processing and transmission delays are proportionately small. Hence, queueing delay is the one that typically contributes most to the delay experienced. It also confirms that processing and queueing delays are the ones that are usually considered in terms of measurements and simulations.

The amount of delay experienced via tunnels versus no-tunnel intra-AS paths and the corresponding cost ratio have been chosen carefully after doing some research on Internet delay measurements (Ramaswamy, 2004 [28]; Zeitoun, 2004 [29]; Choi, 2004 [30]; Carlsson, 2004 [31]). Keeping the hop count in mind, our experiments have been run considering the normal intra-AS path cost as *3x* milliseconds and *4x* milliseconds, where the cost for using a tunnel is *x* milliseconds. Then, for a congested traffic situation, where the queuing delay must be high, the cost for normal intra-AS path is set as *15x* milliseconds.

Our tool uses Dijkstraâs Algorithm to calculate the no-tunnel least cost path depending on these allocated costs. After that, the tool generates a given percentage of tunnels in the produced AS topology. Taking the expected number of tunnels as user input, the tool places different percentages of tunnels in randomly chosen ASes and calculates the least cost path again, considering the tunnels in the chosen ASes. The least cost path includes the tunnels if and only if the delay cost of the tunnels is less than that of the no-tunnel paths. For now, we assume an AS that is selected for hosting tunnels has them arranged in a full mesh between the ASBRs of the AS.

## Results

We have performed a number of simulations in order to access the benefits of using tunnels in a regional network topology. To start with, a small topology of 7 ASes is fed into the PFP model to grow it to a larger topology of 30 ASes for two different node degrees â 3 and 4 â with the probability of a node obtaining a new inter-AS link/adjacency of 0.04, to investigate the benefit for both cases.

Taking the PFP-generated AS-level topology as an input, the framework produces a topology at the ASBR level. Next, Dijkstaâs Algorithm calculates the least cost path from every AS to all the remaining ASes. Then, the presence of 5%, 10%, 15%, 20%, 25% and 30% tunnels is added to the topology and least cost paths are again calculated for every tunnel percentage.

For now, no inter-domain tunnels have been considered and the cost of a link between the peering border routers of two adjacent ASes is set to 1 ms.

The benefit of the tunnels being present is calculated as follows:

*Benefit from AS âAâ to AS âBâ for x% tunnels = [cost from A to B using no tunnels minus the cost from A to B when x% tunnels are present] ms*

The costs are automatically calculated using Dijkstraâs algorithm for each least cost path and then the average and standard deviation of these differences is calculated. It should be noted that in many cases there will be no cost benefit of going via one or more tunnels when they are remote from the original no-tunnel pathway. This tunnel-placement process is repeated 10 times for a given overall AS topology and the average and standard deviation of the benefit are calculated and the results plotted.

Result for Different Topologies

Setting the average node degree to 3, five topologies each having 30 ASes and similar properties are generated by the PFP topology generator. They are then fed into our tool and tunnel placement is repeated 10 times. In each case, the ratio of the cost of a tunnel in an AS to that of a normal no-tunnel path is set at 1:3 (delays are considered in milliseconds) and the average is calculated for the average and standard deviation of the benefit for different percentages of tunnels. Table 1 summarises the results.

Table 1. Average and standard deviation of the benefit for using tunnel(s) (in milliseconds)

| % of tunnels | Topology 1 | Topology 2 | Topology 3 | Topology 4 | Topology 5 |
|---|---|---|---|---|---|
| | Average/ Standard Deviation | Average/ Standard Deviation | Average/ Standard Deviation | Average/ Standard Deviation | Average/ Standard Deviation |
| 5% | 0.325057/ 0.492979 | 0.10698/ 0.425952 | 0.309425/ 0.601801 | 0.191724/ 0.461871 | 0.053334/ 0.28789 |
| 10% | 0.430345/ 0.647781 | 0.438125/ 0.736941 | 0.46006/ 0.790095 | 0.381609/ 0.74438 | 0.218391/ 0.574282 |

| | | | | | |
|---|---|---|---|---|---|
| 15% | 0.822418/ <br> 0.893881 | 0.59418/ <br> 0.899783 | 0.657011/ <br> 0.953498 | 0.498851/ <br> 0.834226 | 0.558391/ <br> 0.786062 |
| 20% | 0.98318/ <br> 1.044247 | 0.847816/ <br> 1.057714 | 0.818391/ <br> 1.049085 | 0.774713/ <br> 1.040576 | 0.703908/ <br> 0.921514 |
| 25% | 1.217595/ <br> 1.120945 | 0.945149/ <br> 1.138616 | 0.872184/ <br> 1.092653 | 0.950345/ <br> 1.144512 | 0.777931/ <br> 0.977332 |
| 30% | 1.272562/ <br> 1.134269 | 1.104503/ <br> 1.230331 | 0.998161/ <br> 1.173999 | 1.071724/ <br> 1.144512 | 0.914943/ <br> 1.028713 |

As expected, as the proportion of tunnels increases so does the average benefit. When the percentage of tunnels is small, the average benefit is marginal. However, from the standard deviation, we can see that even when the percentage of tunnels is low, for some users located close to the tunnels, considerable benefit is still achievable.

Of the five generated topologies, the results of one of the topologies (Topology 4 in Table 1) are now considered in detail.

Different Cost Ratio

For the same topology, different cost ratios are considered for 10 runs. While choosing the cost ratios, at first we have been conservative and considered the average delay cost for no-tunnel paths as *3x* milliseconds and *4x*, where the average delay for a tunnel is *x* milliseconds (as explained in the Implementation section). Then we have considered a situation representing traffic congestion where the average no-tunnel linkâs cost is *15x*. At certain times, the Internet can be busy, impacting on the end-to-end delay. Usually, queueing delay makes a greater contribution in such cases.

Fig. 3 presents a graph plotting the average benefit for different percentages of tunnels for Topology 4 from Table 1.

[32]

Figure 3. Average of Cost Benefit for different cost ratios

It is clear from the graph that, for all cost ratios, the benefit increases as there is an increase in the percentage of tunnels present in the Internet. With a ratio of 1:3, the average delay for sending data in topology 4 is 4.97 ms which is reduced by a minimum of 0.19 ms when 5% of ASes have tunnels in them. The average benefit gradually reaches almost 1.08 ms for 30% tunnels. For a tunnel/no-tunnel ratio set to 1:4, the average end-to-end delay without the use of tunnels is 5.97 ms. With 30% tunnels, this end-to-end delay goes down by 2.03 ms.

It can be seen that the average improvement is relatively small when the tunnelâs average delay cost is one-third or one-quarter of the normal no-tunnel average delay. This is not surprising, as many paths would incur a costly diversion to reach tunnel(s), particularly when they are few in number. Even so, a decrease of almost 2 ms compared with almost 4 ms to 6 ms could still be of attraction to at least some end-users for specific application services.

Conversely, for the âbusyâ period, the average benefits associated with greater cost ratios are noticeably high. When we consider the cost associated with a no-tunnel link in a congested AS as 15 ms, the average end-to-end delay for the same topology is calculated as 16.97 ms. As expected, exploiting tunnels within this AS lowers the delay to a great extent, resulting in more average benefit. For 10% tunnels the average benefit is more than 5 ms and for 30% it reaches almost 9.6 ms.

The standard deviation of the benefit is plotted in Figure 4.

Figure 4 Standard Deviation of Cost Benefit for different cost ratios

The increasing standard deviation shows that, between a smaller number of source-destination pairs, the cost benefit can be substantial. Indeed, it is worth noting that, during peak hours or when specific high-demand events occur, the intra-AS queueing delay can be tens of milliseconds if not more. If tunnels bypass such âhot spotsâ, the delay cost benefit could be orders of magnitude, providing end-users considerable benefit in terms of delay.

Altering the Node Degree

The PFP generator is used again to generate an AS-topology from the same initial 7-node seed graph that has been used to generate the topology used in Fig. 3 and 4. However, this time the graph evolution is altered by setting the average node degree to 4. As with the previous simulations, we have considered the use of tunnels under normal traffic conditions and during a period of localised congestion, where the ratio of the cost for tunnel to that of no-tunnel path is 1:4 and 1:15 within the specified ASes.

Then, for both cases the average of the delay cost benefit for the presence of 5%, 10%, 15%, 20%, 25% and 30% tunnels was calculated and is shown in Fig. 5.

Figure 5. Average of Cost Benefit for different cost ratios

For the AS level topology, the average end-to-end delay without any tunnel is 5.97 ms, which decreases when tunnels are available to end users. If the tunnel has an average delay cost of $1/4^{th}$ of the normal intra-AS link path, then it gives an average end-to-end delay benefit of 0.32 ms, which increases with the number of tunnels and, for 30% tunnels, reaches 1.62 ms.

For the busy period conditions, we assume that the tunnel will have an average delay of $1/15^{th}$ of the average normal intra-AS link delay. For the no-tunnel topology, the average end-to-end delay is 14.94 ms. Clearly, the graph shows that the availability of different percentages of tunnels adds benefit by improving the average delay cost. For 15% of tunnels the average reduction in delay is 4.73 ms and for 30% it is almost double, 8.95 ms; and it is approximately 6 times more than the benefit we observed for the ratio of tunnel/no-tunnel cost of 1:4.

Hence, it is clear from the graphs that, during peak times, even the presence of a small percentage of tunnels can provide noticeable benefit to many users by decreasing the delay cost.

Considering âHotspotâ Area

In the Introduction, we mention that the lack of traffic differentiation can lead to load imbalances and âbest effortâ equal treatment of all traffic irrespective of its importance to the user. Keeping this situation in mind, we conducted a series of simulations for a situation where all the source ASes want to send their data traffic to a particular destination AS over the Internet. This approximates the situation where the destination AS hosts a popular server farm or data centre.

In this case, the framework is changed in such a way that, upon calculating the expected number of tunnels for a specified tunnel-percentage, it generates the tunnels in ASes adjacent to the destination AS first. If the number of expected tunnels is more than the number of adjacent ASes, then the rest of the tunnels are generated to the ASes that are one hop away, and so on. Thus, the tunnels are organised into approximately concentric rings around the destination AS.

Noting the benefits of tunnel-usage are more pronounced and meaningful for âpeak timeâ situations, we have again run 10 simulations for the same topology with 30 ASes, as used in Figure 5 (with a node degree of 4), where the average delay cost for a tunnel is 1 ms and that of a normal path is 15 ms. Taking AS2 as the destination AS, and assuming each of the remaining 29 ASes act as the source domains, Figure 6 presents the graph of average and standard deviation of the benefit for using tunnels around a âhotspotâ destination.

[35]

Figure 6. Average and Standard Deviation of Cost Benefit for Tunnel-No Tunnel = 1:15

The baseline average end-to-end delay for sending data to AS2 from all the other domains is calculated to be 9.5 ms. It is clearly observed that both the average and standard deviation of the delay cost benefit of employing a given percentage of tunnels is relatively high compared the ones we have observed in Figures 3 to 5, even for only 5% of tunnels near a hotspot destination in the network topology.  Hence, in the case of known âhot spotsâ in terms of desirability and possible congestion, access to low delay tunnels becomes particularly attractive.

# Discussion

Using the developed framework, we can examine the delay benefits that intra-AS tunnels might bring to the Internet. It shows that there is a benefit for even 5% tunnels in the network for some users, though this is dependent on how close the tunnel alternatives are to the default traditional pathway.

To show the variation in benefit between source-destination pairs, we provide the standard deviation. This is because the spread of values provides an indication of the percentage of customers that can obtain a certain level of benefit. For different percentages of tunnels, the average benefit is not going to be great for all the users and the standard deviation shows that at least for some customers there is substantial benefit if their desired path lies close to one or more tunnels. If they are further away, the benefit of the tunnel(s) is offset by the additional number of hops to reach them.

However, as one standard deviation only encompasses about 68% of a Normally distributed population, it is worth noting that for some users the cost benefit would be appreciable. Indeed, if the standard path experiences delays brought about by âhot-spotâ congestion, then tunnel alternatives become much more attractive. Furthermore, our current investigations avoid the use of more ambitious inter-AS tunnels, and the broker function need only have access to limited information, ameliorating any security risks. The key benefit, from our perspective at least, is that it gives end-users some choice over how their data is treated by the network and it is up to the user which traffic should be directed via these tunnels.

A further aspect of user-selectable tunnels not explored in detail in this paper is their ability to partially pin down a route to circumvent âless desirableâ areas of the Internet. This allows users to use one or more tunnels to ânail-upâ a loose source-routed path between the source and destination. Using âheartbeatâ probe messages the regular path could be monitored to a limited degree and, if an outage is detected, the traffic flows could be directed via a suitable tunnel to the destination to avoid the anticipated area of concern. This has the advantage that the action can be implemented rapidly, without waiting for BGP to announce an alternative pathway around the outage location.

## Conclusions

This paper introduces a tunnelling framework allowing cooperation between end-users and transport service providers via a simple brokerage mechanism. This is done in such a way that trust issues to do with the tunnel details and AS domain internal architecture are not compromised. The paper at first takes a conservative approach to the introduction of low delay-cost tunnels in an Internet region, typically comprising about 30 AS domains. We avoid tunnels spanning ASes as this would typically require cooperation between service providers. Instead we focus on intra-AS tunnels that are added in a relatively low concentration. We show that some benefit is available, but its magnitude is dependent on the proximity of users to suitable tunnels and the delay performance ratio between the tunnel/no-tunnel intra-AS path alternatives. Not surprisingly, when the tunnels allow a user to circumvent hot spots, their benefit can be appreciable.

We have also considered an AS topology having a particular domain as a hotspot destination, representing an AS hosting a datacentre etc. As the traffic tends to concentrate as it moves towards this âhot spotâ, locating tunnels in the region close to this focal point provides more benefit to more users, should they choose to avail themselves of the tunnelled path alternatives.

In summary, we believe that end-user selectable access to tunnels provides a suitable degree of choice whilst avoiding the issues of ânet neutralityâ and would allow better management of the Internet as demands on its resources continue to grow.

## References

Akter, H., & Phillips, C. (2018). Are Internet Tunnels Worthwhile? *28<sup>th</sup> International Telecommunication Networks and Applications Conference (ITNAC)*, 21-23 November.

BegtaĂ¡eviĂ, F., & Van Mieghem, P. (n.d.). Measurements of the Hopcount in Internet. Available at http://circuit.ucsd.edu/~massimo/ECE158A/Handouts_files/hop-count.pdf [36].

CAIDA. (n.d.). Topology Research. Retrieved from: http://www.caida.org/researchâ/topology/ [37]. Last accessed September 26, 2017.

Carlsson, P., Constantinescu, D., Popescu, A., Fiedler, M., & Nilsson, A. (2004). Delay Performance in IP Routers, *2nd International Working Conference (HET-NETs â04)*.

Chamania, M., Drogon, M., & Jukan, A. (2012). An Open-Source Path Computation Emulator: Design, Implementation, and Performance. *Journal of Lightwave Technology, 30*(4).

Choi, B-K., Moon, S., Zhang, Z-L., Papagiannaki, K., & Diot, C. (2004). Analysis of point-to-point packet delay in an operational network. *IEEE Infocom 2004*, 7-11 March.

Clegg, R.G., Di Cairano-Gilfedder, C., & Zhou, S. (2010). A critical look at power law modelling of the Internet,â *Computer Communications*, *33*(3), 259-268, February.

Dasgupta, S., De Oliveira, J.C., & Vasseur, J-P. (2007). Path-Computation-Element-Based Architecture for Interdomain MPLS/GMPLS Traffic Engineering: Overview and Performance. *IEEE Network Magazine*, *21*(4), 38-45, July-August.

Doctorow, C. (2014). Internet service providers charging for premium access hold us all to ransom. *The Guardian*, 28 April. Retrieved from: https://www.theguardian.com/âtechnology/2014/apr/28/internet-service-providers-charging-premium-access. [38]

Farrel, A. (2004). *The Internet and Its Protocols: A Comparative Approach*. USA: Morgan Kaufmann.

Honan, M. (2008). Inside Net Neutrality: Is your ISP filtering content?*Macworld*, 12 February. Retrieved from: https://www.macworld.com/article/1132075/web-apps/ânetneutrality1.html. [39] Last accessed August 2018.

Huffaker, B., Fomenkov, M., & Claffy, C. (2012). Internet Topology Data Comparison.*Technical Report*, Cooperative Association for Internet Data Analysis (CAIDA), May.

Markakis, E., Sideris, A., Alexiou, G., Bourdena, A., Pallis, E., Mastorakis, G., & Macromoustakis, X. (2016). A virtual network functions brokering mechanism, *International Conference on Telecommunications and Multimedia (TEMU)*, IEEE, 25-27 July.

Plummer, D., Lheureux, B., Cantara, M., & Bova, T. (2011). Cloud Services Brokerage Is Dominated by Three Primary Roles. *Gartner Research Note* G00226509, 23 November.

Ramaswamy, R., Weng, N., & Wolf, T. (2004). Characterising the Network Processing Delay*IEEE Global Telecommunications Conference (GLOBECOM '04)*, 29 November-3 December.

Rushe, D. (2017). Net neutrality: Amazon among top internet firms planning day of action.*The Guardian*, 6 June. Retrieved from:https://www.theguardian.com/technologyâ/2017/jun/06/net-neutrality-amazon-etsy-kickstarter-protest [40]. Last accessed August 2018.

Rzym, G., Wajda, K., & Rzym, K. (2016). Analysis of PCE-based path optimization in multi-domain SDB/MPLS/BGP-LS network. *18th International Conference on Transparent Optical Networks (ICTON)*, 10-14 July.

Schaffrath, G., Werle, C., Papadimitriou, P., Feldmann, A., Bless, R., Greenhalgh, A., Wundsam, A., Kind, M., Maennel, O., & L. Mathy, L. (2009). Network virtualization architecture: proposal and initial prototype. Applications, Technologies, Architectures, and Protocols for Computer Communication, pp. 63-72 in *Proceedings of the 1st ACM workshop on Virtualized infrastructure systemsand architectures*, Spain: Barcelona, August.

Secci, S., Rougier, J-L., & Pattavina, A. (2008). On the selection of optimal diverse AS-paths for inter-domain IP/(G)MPLS tunnel provisioningâ, *IEEE 4th International TeleÂcommunication Networking Workshop on QoS in Multiservice IP Networks*, 13-15 February.

Statista. (n.d.). AT&T - Statistics & Facts. Retrieved from: https://www.statista.comâ/topics/1252/atundt/ [41]. Last accessed September 30, 2017.

W. Tim, 2003. âNet Neutrality, Broadband Discriminationâ, *Journal on Telecommunications And High Technology Law, Volume 2, pp 141-176.*

Yost, J.R. (2015). The Origin and Early History of the Computer Security Software Products Industry. *IEEE Annals of the History of Computing*, *37*(2), 46-58.

Zeitoun, A., Chuah, C-N., Bhattacharyya, S., & Diot, C. (2004). An AS-level study of Internet path delay characteristics. *IEEE Global Telecommunications Conference (GLOBECOM '04)*, 29 November-3 December.

Zhou, S. (2006). Characterising and modelling the Internet topology- The rich-club phenomenon and the PFP model. *BT Technology Journal*, *24*(3), July.

Zhou, S., & Mondragon, R. J. (2004). Accurately modelling the Internet topology.*Physical Review E*, *70*(6), No. 066108, December.
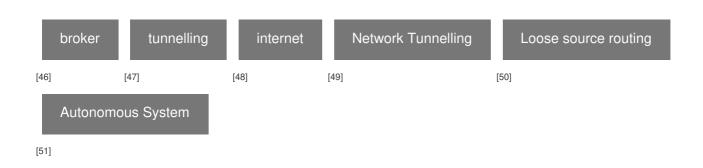
**Article PDF:**

175-article_text-1904-1-10-20190419.pdf [42]

**Copyright notice:**

**Cite this article as:**

| broker | tunnelling | internet | Network Tunnelling | Loose source routing |
|---|---|---|---|---|
| [46] | [47] | [48] | [49] | [50] |

| Autonomous System |
|---|
| [51] |

**Links**
[1] https://telsoc.org/journal/author/habia-akter [2] https://telsoc.org/journal/author/chris-phillips [3] https://telsoc.org/journal/jtde-v7-n1 [4] https://www.addtoany.com/share#url=https%3A%2F%2Ftelsoc.org%2Fjournal%2Fjtde-v7-n1%2Fa175&amp;title=Tunnelling%20the%20Internet [5] https://telsoc.org/printpdf/2416?rate=TNbzHcc2JZWfW8WalPx4kEd_5mCf2O1IlDNlMccnXHl [6] https://telsoc.org/journal/jtde-v7-n1/a175#Schaffrath_2009 [7] https://telsoc.org/journal/jtde-v7-n1/a175#Chamania_2012 [8] https://telsoc.org/journal/jtde-v7-n1/a175#Yost_2015 [9] https://telsoc.org/journal/jtde-v7-n1/a175#Rzym_2016 [10] https://telsoc.org/journal/jtde-v7-n1/a175#Dasgupta_2007 [11]

https://telsoc.org/journal/jtde-v7-n1/a175#Farrel_2004 [12] https://telsoc.org/journal/jtde-v7-n1/a175#Secci_2008 [13] https://telsoc.org/journal/jtde-v7-n1/a175#Doctorow_2014 [14] https://telsoc.org/journal/jtde-v7-n1/a175#Wu_2003 [15] https://telsoc.org/journal/jtde-v7-n1/a175#Honan_2008 [16] https://telsoc.org/journal/jtde-v7-n1/a175#Rushe_2017 [17] https://telsoc.org/sites/default/files/images/tja/175_fig1b.jpg [18] https://telsoc.org/journal/jtde-v7-n1/a175#Statista_nd [19] https://telsoc.org/journal/jtde-v7-n1/a175#CAIDA_nd [20] https://telsoc.org/journal/jtde-v7-n1/a175#Huffaker_2012 [21] https://telsoc.org/journal/jtde-v7-n1/a175#Plummer_2011 [22] https://telsoc.org/journal/jtde-v7-n1/a175#Markakis_2016 [23] https://telsoc.org/journal/jtde-v7-n1/a175#Zhou_2006 [24] https://telsoc.org/journal/jtde-v7-n1/a175#Zhou_Mondragon_2004 [25] https://telsoc.org/journal/jtde-v7-n1/a175#Clegg_2010 [26] https://telsoc.org/journal/jtde-v7-n1/a175#Begta%C5%A1evi%C4%87_nd [27] https://telsoc.org/sites/default/files/images/tja/175_fig2.jpg [28] https://telsoc.org/journal/jtde-v7-n1/a175#Ramaswamy_2004 [29] https://telsoc.org/journal/jtde-v7-n1/a175#Zeitoun_2004 [30] https://telsoc.org/journal/jtde-v7-n1/a175#Choi_2004 [31] https://telsoc.org/journal/jtde-v7-n1/a175#Carlsson_2004 [32] https://telsoc.org/sites/default/files/images/tja/175_fig3.jpg [33] https://telsoc.org/sites/default/files/images/tja/175_fig4.jpg [34] https://telsoc.org/sites/default/files/images/tja/175_fig5.jpg [35] https://telsoc.org/sites/default/files/images/tja/175_fig6.jpg [36] http://circuit.ucsd.edu/~massimo/ECE158A/Handouts_files/hop-count.pdf [37] http://www.caida.org/research%E2%80%8C/topology/ [38] https://www.theguardian.com/%E2%80%8Ctechnology/2014/apr/28/internet-service-providers-charging-premium-access. [39] https://www.macworld.com/article/1132075/web-apps/netneutrality1.html.%20 [40] https://www.theguardian.com/technology%E2%80%8C/2017/jun/06/net-neutrality-amazon-etsy-kickstarter-protest [41] https://www.statista.com/topics/1252/atundt/ [42] https://telsoc.org/sites/default/files/tja/pdf/175-article_text-1904-1-10-20190419.pdf [43] https://telsoc.org/copyright [44] http://doi.org/10.18080/ajtde.v7n1.175 [45] https://telsoc.org [46] https://telsoc.org/topics/broker [47] https://telsoc.org/topics/tunnelling [48] https://telsoc.org/topics/internet [49] https://telsoc.org/topics/network-tunnelling [50] https://telsoc.org/topics/loose-source-routing [51] https://telsoc.org/topics/autonomous-system