

# RemIX: A Distributed Internet Exchange for Remote and Rural Networks

William Waites<sup>\*†</sup>, James Sweet<sup>\*†</sup>, Roger Baig<sup>‡</sup>,  
Peter Buneman<sup>\*</sup>, Marwan Fayed<sup>‡</sup>, Gordon Hughes<sup>\*</sup>  
<sup>\*</sup>University of Edinburgh, <sup>†</sup>HUBS C.I.C, <sup>‡</sup>guifi.net, <sup>‡</sup>University of Stirling  
Corresponding authors: wwaites@tardis.ed.ac.uk, opb@inf.ed.ac.uk, mmf@cs.stir.ac.uk

## ABSTRACT

The Internet Exchange Point (IXP), an Ethernet fabric central to the structure of the global Internet, is largely absent from community-driven collaborative network infrastructure. IXPs exist in central, typically urban, environments where strong network infrastructure ensures high levels of connectivity. Between rural and remote networks, separated by distance and terrain, no such infrastructure exists. In this paper we present RemIX a distributed IXP architecture designed for the community network environment. We examine this praxis using an implementation in Scotland, with suggestions for future development and research.

## CCS Concepts

•Networks → Network architectures; Wireless access networks;

## Keywords

Internet Exchanges (IXP); Community Broadband

## 1. INTRODUCTION

In remote and rural regions the last-mile problem has been the subject of much focus. Deployments show that high quality access networks can be built in otherwise under-serviced regions [7, 5, 8]. The underlying technologies range in medium (eg. copper or fibre-optic cabling, licensed or unlicensed wireless), energy (eg. solar, wind or mains supplied), and topology. Successful deployments, including our own in Scotland, have two attributes in common: (i) Networks designs are bespoke,

suggesting there is no one-size-fits-all solution; (ii) crucially, communities must be invested and involved [19].

Though remote access network research is far from complete, the next question is increasingly clear: What options do remote, isolated networks have for ‘backhaul’ to interconnect with the rest of the Internet? We define “remote” as far from urban areas where commodified network infrastructure is available. For example long-distance circuits, if and where they are available, are both expensive and difficult to reach. Access networks in remote places serve populations that are dispersed. The lower population density reduces the size of their user-base when compared to their urban cousins. With no options for interconnecting with nearby networks to generate economies of scale, *high-quality* backhaul is prohibitively expensive, if it exists at all.

The absence of resource pooling options for remote networks is the focus of this paper. One such example is operated by the Guifi Foundation [7]. Guifi operates a regional backbone network as a commons. The abstraction that is presented to clients is an exchange point implemented over IP. In this type of network, relationships between end-users are either mediated by Guifi, or implemented as an overlay.

The Internet Exchange Point (IXP) is a long-standing structure that plays a pivotal role in facilitating interconnections between networks [1, 6]. We are motivated by IXPs for two reasons. First, the primary role of an IXP is economic. Member networks can connect  $n$  networks at an IXP with  $n$  circuits, rather than arranging  $O(n^2)$  circuits independently. Second, the IXP model of multilateral public peering leads to high density interconnections, and traffic across the exchange that can be comparable in magnitude to the largest global service providers [1]. Together, they are an indication that such a topology might be used to improve inter-connectivity between networks in under-serviced regions, and to pool otherwise expensive backhaul resources.

In this paper we present RemIX, a *distributed* Internet Exchange for Remote and rural networks. The RemIX architecture is agnostic to underlying technologies, embedding the same principles as the successful re-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

GAIA, August 22-26, 2016, Florianopolis, Brazil

© 2016 ACM. ISBN 978-1-4503-4423-4/16/08...\$15.00

DOI: <http://dx.doi.org/10.1145/2940157.2940162>

remote networks it is designed to serve. It distinguishes itself from IXPs by the vast distances permitted between points of presence, and the lower density of member networks that connect to them. The trade-off between distance and density gives rise to the idea of *lightweight* points of presence. The lightweight nature is advantageous, in that as few as two member networks are sufficient to establish a point of presence.

We describe our RemIX implementation in Scotland. In its current form our deployment services a  $\sim 2000\text{km}^2$  region that spans sea and mountainous mainland. Implementation details are provided, with motivating rationale, so that others may benefit from our efforts. Functionally, our implementation appears to its members as a large Ethernet switching fabric. Crucially, RemIX allows member networks to establish unmediated relationships between themselves.

In the following sections we further motivate IXPs as an ideal model. We then discuss the RemIX architecture in detail. Our deployment is described, along with lessons learned. Finally, a broader context of the local environment is presented before concluding remarks.

## 2. Internet Exchange Points

As part of the decommissioning of the NSFNET, four Network Access Points (NAPs) were created. They were operated by large American telephone companies (MCI, Sprint, PacBell, Ameritech) and designed to prevent partitioning of the commercial Internet [1, 6]. NAPs were prohibitively expensive with technological requirements that created barriers to participation. Soon IXPs emerged as an alternative. IXPs appeared in carrier-neutral facilities allowing dense inter-network connections on a non-discriminatory basis. Presence at an IXP entails freedom to make bilateral arrangements with any other network also present. Worldwide, IXPs now number in the hundreds and are a fundamental feature in the structure of the Internet.

A mirroring of this structure would be useful in joining remote networks. The increases of interconnection density could then be used to pool traffic, and make collective use of expensive resources such as long-distance circuits. However, there are some important differences between the environment of a typical urban IXP and the rural regions, as in the West Coast of Scotland: (i) There are no data centres, carrier-neutral or otherwise; (ii) due to geography there is no single facility where all of the networks could meet.

## 3. REMIX ARCHITECTURE

In this section we present the RemIX architecture. We compare RemIX with IXP architectures, and relate those benefits in the context of remote access networks.

### 3.1 Design Requirements

Our requirements are shaped by three broad goals: (i) establish high-quality backhaul to remote regions;

(ii) ensure backhaul affordability for small access networks; (iii) allow networks to maintain the autonomy that is fundamental to their sustainability. Member networks must be able to connect to one or more transit providers. Members must also be free to arrange and articulate policies among themselves. These requirements imply that a *logical* concentration of inter-network connections is desirable, which suggests a shared switching fabric below the network layer.

Networks that can connect to the same location can do so via an Ethernet switch. This leads to conventional IXP design (Fig. 1a) where member networks connect to a central fabric with their own router that sits inside the IXP facility. Remote networks have no such luxury. In response, we take and distribute the contemporary design of multi-site IXPs (Fig. 1b). Multi-site IXPs present a single logical fabric to members, implemented with switches joined by private circuits.

The RemIX architecture that emerges (Fig. 1c) has no large facility nor physical housing. Instead it is distributed so that *lightweight* points of presence may be established where there are as few as two members. Members either colocate their border routers with the exchange switch, or remotely on the far end of a link, as circumstances dictate.

These circumstances motivate the lightweight nature of points of presence. Since the fabric is distributed, fewer networks that will connect from each site. High port densities are unnecessary. Simultaneously, space and power are both at a premium. For example, a remote port into RemIX could be housed in a small cabinet atop a hill, or in space that is donated by a property owner for this purpose. Equipment is therefore restricted to the small and power-efficient.

## 3.2 RemIX Components

### 3.2.1 Switching Fabric

The exchange itself must mimic an Ethernet switch. Multiple Ethernet-like link options include fibre, 802.11 wireless, licensed wireless, fibre, leased pseudowires. The switching fabric may be implemented on top using BGP-VPLS [11] (as we have in Section 4.2), BATMAN [10], or TRILL [15] protocols. The salient feature between them is MAC address learning to establish an Ethernet switch similar to the Metro Ethernet Forum (MEF) E-LAN interface specification [13].

### 3.2.2 Member Autonomous Systems (ASs)

Among traditional IXPs connected networks are encapsulated into Autonomous Systems (ASs). Among RemIX member networks, policies of the small member networks are different from the Internet’s Default-Free Zone (DFZ). In particular, member networks’ smaller routers will be neither be capable of storing the entire Internet routing table, nor are they likely to announce netblocks large enough to be globally visible. However, AS *encapsulation* enables networks to retain

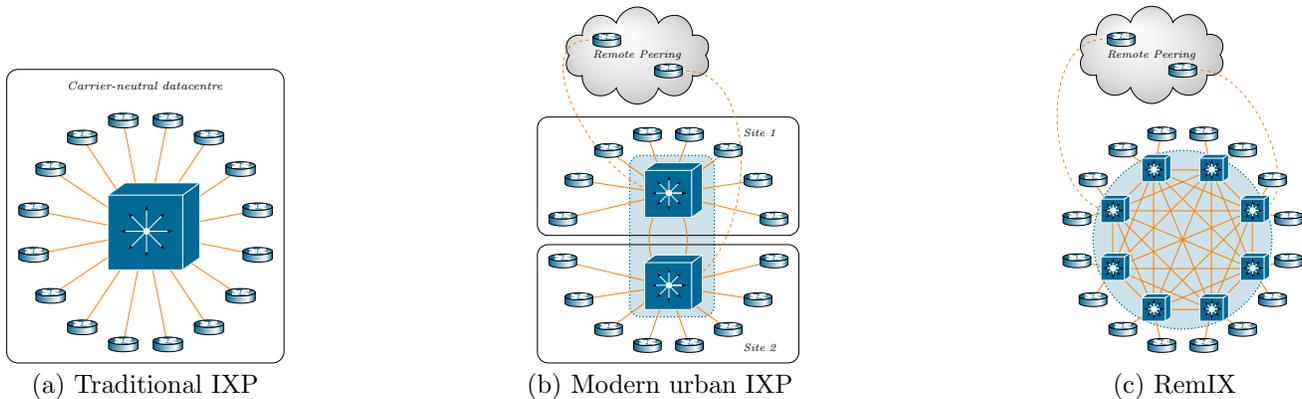


Figure 1: Comparison of exchange point models. Notice density.

their internal structures and methodologies, and to interconnect safely with neighbours. Due to the likelihood of collisions use of private Autonomous System Numbers (ASNs) is inappropriate for this purpose [14], as are private IP addresses for the exchange itself [16].

### 3.2.3 Exchange Transit

RemIX members’ IP address spaces will be small, and need some entity to advertise larger netblocks on their behalf. This suggests a specialized transit provider to mediate between members and the wider Internet. For this reason RemIX members form a confederation with a transit provider that presents them collectively to upstream providers and other exchange points. This is unusual for IXPs: Rarely are transit relationships implemented with exchange points. In RemIX this is normal, and likely necessary to function in the intended environment. We note that transit service should be optional, with no requirement to purchase said provider’s transit as a condition for joining the exchange. Also, nothing prevents other such providers from participating.

### 3.2.4 Auxiliary Services

BGP configuration can be complex. For example, upon connecting to RemIX, member networks need to be configured to peer amongst themselves. The complexity quickly increases as session numbers grow with the square of the number of participants. Instead, IXPs use *route-servers* to repeat member announcements to all others. Route reflectors keep configuration burdens to a minimum. Other useful services such as NTP clocks and looking glasses may be offered in addition.

The overall RemIX architecture is motivated by our own needs in Scotland. In the next section we present our first-phase implementation of RemIX, alongside remarks on usability and directions for the future.

## 4. REMIX DEPLOYED IN SCOTLAND

In this section we describe our first implementation of RemIX in a series of planned deployments across

Scotland. In the West Highlands there is a cluster of 11 small community networks. Their spread across  $\sim 2000km^2$  of sea and mountainous islands makes the construction of an exchange fabric geographically ambitious. Four networks have a history of interconnecting and sharing network resources, pre-established relationships that must be respected in our deployment.

Our deployment’s location is its namesake, the West Highland Internet Exchange (WHIX). Logical and physical layers are described below, with additional lessons and comments drawn from our experience.

### 4.1 West Highland IX at Layer 1

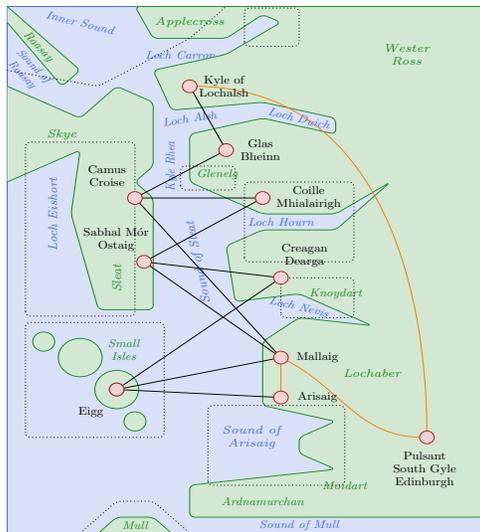
The physical WHIX fabric is overlaid onto a stylized map of the region in Figure 2a. The map itself preserves critical geographical features. Red connected nodes are the connection sites. In a traditional IXP these sites are the Ethernet ports into which subscriber ASs plug-in. WHIX sites are connected by wireless radio links in black, and leased 100Mbps or 1Gbps circuits in orange. The areas enclosed with dotted lines correspond to the service areas reachable from each site.

The WHIX physical topology in Figure 2a appears alongside the member network in Figure 2b. In the latter, unlabeled red nodes are the WHIX points of presence and correspond with the same set of red nodes in Figure 2a. The dashed lines represent the fully connected virtual topology that implements the exchange E-LAN.

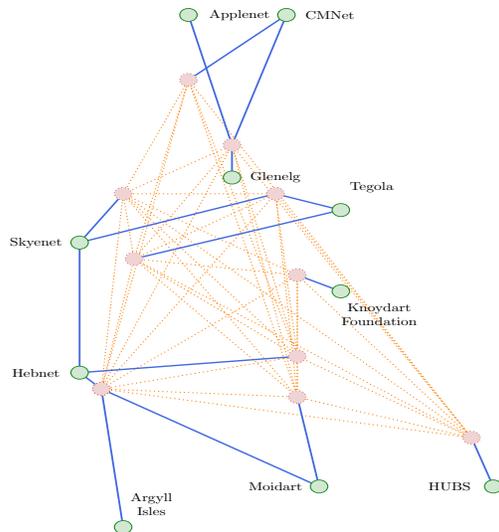
The two places in the region where long-distance ethernet circuits are available on the mainland are the towns of Mallaig and Kyle of Lochalsh. Circuits<sup>1</sup> from these sites connect back to the Pulsant datacentre in Edinburgh to facilitate remote peering — and indeed the provision of Internet access via the exchange point.

The radio links are implemented with Ubiquiti Networks equipment, configured in transparent bridge mode so that they appear as Ethernet from a functional per-

<sup>1</sup>At the time of writing, the Mallaig circuit exists, and that from Kyle is pending.



(a) Physical topology of WHIX.



(b) Member connections to WHIX

**Figure 2: Physical and logical layout of WHIX.** In Figure 2a the dark lines correspond to radio links and the light, curved lines to leased ethernet circuits. In Figure 2b the dashed lines correspond to internal layer-2 circuits forming WHIX switching fabric and the solid lines to member connections.

spective. The switching fabric itself at each of WHIX points of presence is implemented with Mikrotik routers. This choice was made because of their moderate port density, low power consumption, low cost, and adequately featureful MPLS implementation. We revisit this choice in the next section. All equipment is configured to pass Ethernet frames of at least 1600 bytes to provide room for the necessary extra protocol headers for implementing the E-LAN service.

## 4.2 West Highland IX at Layer 2

Layer-2 details are *internal* to WHIX, and invisible to members who only see an Ethernet switch. Also, our implementation decisions are by no means the only possible means of implementation.

In WHIX the requirement for functional equivalence to a MAC address learning Ethernet switch is met using BGP signalled VPLS [11]. This creates a full set of LSP pseudo-wires between every pair of WHIX edge routers. Each WHIX router maintains an OSPF routing protocol adjacency with its neighbours and distributes reachability information for its loopback IP address. All addresses used for this purpose are private IPv4 addresses [16]. This is the basic layer that ensures reachability throughout the distributed fabric. Non-IP traffic is carried via LDP [2] with MPLS labels according to the topology of the underlying OSPF network.

Routers in WHIX establish BGP peering sessions with routers at Mallaig and Sabhal Mòr Ostaig that act as route reflectors [4]. Participating routers use route reflectors to exchange reachability information without requiring a full mesh ( $n^2$ ) of internal peering sessions. The presence of BGP signalling throughout the WHIX

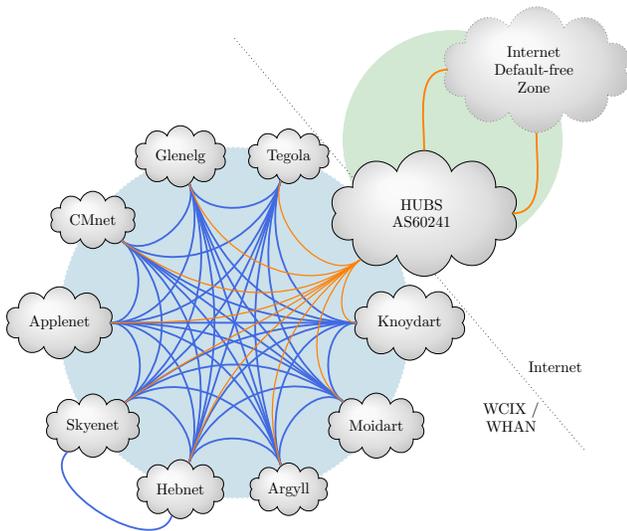
fabric enables the use of multi-protocol extensions [3]. Routers can use extensions to signal a desire to form part of the exchange LAN. The result is a fully meshed VPLS, where each router has a virtual bridge interface that forms part of the exchange LAN.

Interfaces can be added to virtual bridges, as needed, to form part of the exchange. Care must be taken to prevent loops in which members see the traffic that they originate. This is accomplished with a split-horizon method [12]. Equally, members must be prevented from creating bridge loops via their own network by employing MAC address filter on relevant ports.

## 4.3 West Highland IX at Layer 3+

Given logical connectivity between all member ports, it remains to assign IP addresses to their border routers, as well as public infrastructure such as the router server. As mentioned above the use of private IP address space for this purpose is undesirable since it generates risks of collisions with members' own infrastructure. WHIX, and more generally RemIX, is fortunate in this regard: The design meets the definition of an IXP [18, 9], making it possible to acquire IPv4 and IPv6 address allocations from RIPE NCC [17].

The full layer-3 WHIX topology is shown in Figure 3. At this stage member network have everything they need. Members can communicate at layer 2. Each has an IP address at layer 3, an autonomous system number for identification, and their own networks to announce. Bilateral peering arrangements (an otherwise  $n^2$  configuration task) are facilitated by two route-servers, as before at Mallaig and Sabhal Mòr Ostaig. The route servers redistribute reachability information, akin to a



**Figure 3: Autonomous System topology. WHIX members are fully connected; communication over the exchange needs no intermediation.**

route-reflector omits its own ASN from the path.

The transit provider, HUBS (see Section 5.2), is also present at WHIX as a member. In addition to public multilateral peering, it establishes bilateral sessions with members wishing to announce either a default route or full Internet routing tables. HUBS forwards those members' announcements upstream and to their peers. In this way transit, and hence connectivity to the global Internet, is provided over the exchange.

#### 4.4 Deployment Discussion

Our experience motivates higher-level comments to further distinguish RemIX deployments from their larger IXP cousins. Flat networks consisting of a single layer-2 broadcast domain can be plagued by problems that are difficult to troubleshoot. By its very design RemIX requires that members be able to communicate directly without mediation at the IP layer. Like other IXPs RemIX eliminates a large class of potential problems by allowing only unicast and ARP traffic on the exchange. Moreover, members must nominate a specific MAC address for their connections, which reduces the risk of loops and broadcast storms. We also adopt best practices such as quarantines for new connections while they are evaluated for correctness.

IP transit in RemIX also deserves to be addressed. Transit via the exchange, for networks that are not otherwise visible on the Internet, may evoke notions of conflicting interests that beset NAPs. However the similarity is superficial. Here, the transit provider and exchange operator HUBS C.I.C., is a cooperative that exists for the benefit of and is controlled by the members, who are also members of the IXP. As a consequence all parties' economic interests are aligned.

Finally, WHIX' implementation using BGP-VPLS to construct the exchange fabric makes it possible to offer auxiliary point to point pseudo-wire services to its members. This is useful for those members that have need for making connections internal to their networks.

## 5. THE ENVIRONMENT

Scotland holds 1/3 of the area and 10% of the population of the UK. It also has 95 inhabited islands with about 100,000 people. The Scottish Highlands and Islands, where this work is currently focussed, consist of mountainous terrain stretching along a 400km north to south corridor. Islands are scattered on the West, while deep lakes and glens penetrating the mainland to the East. The economy was traditionally maritime, and nearly all habitation is at sea level or in the glens.

### 5.1 Local Infrastructure Development

Fibre in the region has only recently appeared. Much of the telephone network in the region was constructed with microwave links. Infrastructure is improving, but plans terminate at telephone exchanges. Among them, fibre-based services are rare. In the medium term future, local wireless distribution is the only feasible technology for adequate bandwidth and quality of service.

Starting in 2008, the Tegola project [5] started to experiment with technology that would enable communities to build their own wireless networks. The details of Tegola, and its dissemination to nearby communities, are omitted due to space constraints. Relevant to this project is the technology that emerged. Figure 4, for example, features the type of robust, inexpensive relay construction that operates in mountainous region, and that can be constructed by its residents.



**Figure 4: A basic relay**

Many communities have since constructed their own local distribution networks with point-to-point wireless links that can span more than 20km. Expertise is often shared between them, also infrastructure where feasible, yet they operate independently. Constrained by availability, they acquire backhaul via ADSL lines nearby to telephone exchanges. Ethernet services have since emerged in two larger towns, with wholesale pricing that exceeds the budget of any single community. A resolution has two components: An organizational ve-

hicle that combines networks to generate economies of scale, and a supporting network infrastructure.

We have learned that solutions are complicated by both terrain and by culture. In particular we note: (a) Social aspects and organization of communities can fail to align with the ideal “electronic” or networked communities, eg. physical landscape constrains connectivity, while social and economic groups can be determined by vehicles for funding. (b) Local network infrastructure is non-uniform and varies in complexity. (c) Communities that share network resources generally do so in a non-systematic or ad-hoc manner.

## 5.2 HUBS C.I.C.

In response to the local environment and absence of affordable backhaul, the Universities of Edinburgh and Stirling launched HUBS, which is a not-for-profit transit provider whose members are the community networks that it serves. HUBS is also a co-operative where the networks that subscribe also become members. It is the culmination of collaborations between Universities with communities in the West Highlands, and later with community networks in the South Scotland.

The need for RemIX-like functionality arose soon after launch. Two subscriber networks exploited mutual proximity to collaborate on an operational basis. Equipment management and troubleshooting tasks, for example, were shared. Their desire to keep the details internal was complicated because their only interconnection was mediated by HUBS. Circuits were hand-crafted between them, and demonstrated the benefits of bilateral agreements between HUBS members. However, while effective, hand-crafted circuits would fail to scale.

HUBS bridges gaps in backhaul affordability. It has also revealed the benefits emerge when remote and rural networks are able to act collectively in the wholesale telecommunications market, and present a uniform interface to their upstream transit provider. However, a transit-only solution prohibits autonomous arrangements between members unmediated at the IP layer. From this need the RemIX architecture emerges.

## 6. CONCLUDING REMARKS

The features of RemIX described above will be instantly recognizable to anyone who has participated in a regular IXP. This is by design. RemIX is architected to mirror in under-serviced regions, the benefits of IXPs in urban regions. The encapsulation of small community networks in ASs means that they can present a uniform interface to a transit provider, cooperate and share resources. RemIX provides these benefits to members without sacrificing their independence, a necessary attribute for long-term sustainability.

## 7. ACKNOWLEDGEMENTS

This project is supported by the University of Edinburgh, by the Scottish Government and by local indus-

tries including Marine Harvest Scotland and Benchmark Holdings. During an incubation period, WHIX is jointly operated by HUBS and the University of Edinburgh.

## 8. ADDITIONAL AUTHORS

Additional authors: Michael Fourman (University of Edinburgh), Richard Simmons (University of Stirling).

## References

- [1] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger. Anatomy of a large european ixp. In *ACM SIGCOMM*, 2012.
- [2] L. Andersson, I. Minei, and B. Thomas. LDP Specification. RFC 5036 (Draft Standard), Oct. 2007. Updated by RFCs 6720, 6790, 7358, 7552.
- [3] T. Bates, R. Chandra, D. Katz, and Y. Rekhter. Multi-protocol Extensions for BGP-4. RFC 4760 (Draft Standard), Jan. 2007. Updated by RFC 7606.
- [4] T. Bates, E. Chen, and R. Chandra. BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP). RFC 4456 (Draft Standard), Apr. 2006. Updated by RFC 7606.
- [5] G. Bernardi, P. Buneman, and M. K. Marina. Tegola tiered mesh network testbed in rural scotland. In *Proc. of ACM Workshop on Wireless Networks and Systems for Developing Regions*, pages 9–16. ACM, 2008.
- [6] N. Chatzis, G. Smaragdakis, A. Feldmann, and W. Willinger. There is more to ixps than meets the eye. *SIGCOMM Comput. Commun. Rev.*, 43(5):19–28, Nov. 2013.
- [7] Guifi.net. Community network. <https://guifi.net/en/>.
- [8] S. Hasan, Y. Ben-David, M. Bittman, and B. Raghavan. The challenges of scaling wisps. In *Proc. of the Symposium on Computing for Development (DEV)*, 2015.
- [9] HUBS CIC. WHIX Rules. <https://hubs.net.uk/whix>, 2015.
- [10] D. Johnson, N. Ntlatlapa, and C. Aichele. Simple pragmatic approach to mesh routing using BATMAN. In *IFIP Symposium on Wireless Communications and Information Tech. in Developing Countries*, Pretoria, South Africa, 2008.
- [11] K. Kompella and Y. Rekhter. Virtual Private LAN Service Using BGP for Auto-Discovery and Signaling. RFC 4761 (Proposed), 2007. (Updated) RFC 5462.
- [12] M. Lasserre and V. Kompella. Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling. RFC 4762 (Proposed Standard), Jan. 2007.
- [13] MEF. Technical specification 6.2: Metro ethernet services definitions phase 2. <https://www.mef.net/carrier-ethernet/technical-specifications>, 2014.
- [14] J. Mitchell. Autonomous System (AS) Reservation for Private Use. RFC 6996 (Best Current Practice), 2013.
- [15] R. Perlman. Rbridges: transparent routing. In *IEEE INFOCOM*, volume 2, pages 1211–1218. IEEE, 2004.
- [16] Y. Rekhter, B. Moskowitz, D. Karrenberg, G. J. de Groot, and E. Lear. Address Allocation for Private Internets. RFC 1918, Feb. 1996. (Updated) RFC 6761.
- [17] RIPE. IPv4 address allocation and assignment policies for the RIPE NCC service region. RIPE 649, 2015.
- [18] L. Vegoda and T. Lowe. IPv6 address space policy for internet exchange points. RIPE 451, 2009.
- [19] C. Wallace and K. Vincent. Six factors for success in community broadband initiatives. *interactions*, 23(1):36–39, 2015.