Motivation for Coalition-Based Community Networking [*]

Manish Lad, Saleem Bhatti, Steve Hailes,
Peter Kirstein, Adam Greenhalgh
University College London
Gower Street, London. WC1E 6BT
{M.Lad, S.Bhatti, S.Hailes, P.Kirstein, A.Greenhalgh}@cs.ucl.ac.uk

Abstract

Advances in local-area networking have enabled individuals to interconnect their computers or local networks, to form neighbourhood meshes. However, such interconnection, or peering, is usually undertaken by the individuals on an ad hoc basis. The mechanisms that are employed currently; such as tunnelling, private addressing, and network address translation; are inefficient and require a level of expertise to configure and maintain. This excludes many from participating and benefiting from such collaborations. We examine why formalising the individual peering relationships using existing routing mechanisms is not a realistic approach. We present a new architectural element, the *Coalition Peering Domain*, to add structure to these ad hoc peerings, while allowing individuals to maintain local control, and we discuss the challenges that must be faced.

1 Introduction

A new class of community-area networks [1, 2] are emerging. They involve individuals connecting together directly their home and personal-area networks, on an ad hoc basis, forming a local neighbourhood *mesh* or *community network*. However, there are no defined protocols to facilitate the formation of these community-area networks. The peering agreements between pairs of community members are unique, requiring manually customised static configurations.

We propose that some structure should be added to the ad hoc peering agreements, in a way that encourages collaboration while maintaining local control in a dynamic manner. This is achieved through the formation of a ***Coalition Peering Domain*** wherein there is collaboration between individuals for mutual benefit. The viewpoint presented here is forward-looking and provides some initial analysis on possible future solutions for the emerging networking relationships within the community-area context.

In Section 2 of this paper, we discuss the reasons why existing mechanisms do not provide optimal solutions when applied to this class of community-area networks. In section 3 we discuss some existing work relevant to community-area networking. We present the idea of coalition-based community peering in Section 4, and discuss the challenges faced by such a scheme. Finally, we conclude with a summary in Section 5.

2 Motivation

Community-area networks have begun to emerge as a result of a growing disparity between the data rates available in the wide-area and in the local-area. Although advances in Internet and wide-area connectivity technologies have had a great impact on connectivity both for the home user and for the mobile user, their data rates still fall well short of the ever increasing data rates possible in the local-area. Advancing user applications and increasing user expectations have aided to further accentuate this disparity. Meanwhile the increasing availability and the decreasing cost of network-capable consumer equipment has fuelled an increase in device and network interconnectivity. So, not only do such community-area networking initiatives already exist, but they are likely to stay and continue to grow. Such initiatives exhibit some key characteristics that distinguish them from existing classes of networks.

***Collections* of collaborating, autonomous systems:** The resulting community networks are essentially multihomed, with many ingress and egress connections to the wide-area. Administrative responsibility is *distributed* across the community. Thus, they do not represent a single Administrative Domain (AD) that is under the control of a single organisation or entity, but rather a *collaborative* group of such entities.

Although existing mechanisms for addressing and routing

[*] Department of Computer Science, University College London, Technical Research Note Number: RN/05/16

are employed by individuals to administer their peering agreements, such mechanisms do not provide an optimal solution. These mechanisms are designed to operate in network environments where administrative responsibility is not distributed. Formalising the relationships between community members is therefore also not feasible using existing techniques alone.

The need to limit complexity: Proposing to employ existing BGP mechanisms [3] between community members is an extremely complex and heavy-weight solution. It requires first that each community member be allocated an Autonomous System (AS) number. This poses great problems for the existing Internet infrastructure, because, even the emergence of a small number of such communities would lead to a big explosion in the lengths of autonomous system paths and the numbers of routing table entries. Such back-pressure onto the core network is made even worse by the transient nature of community members' equipment: community members may switch on and off arbitrarily their equipment. Magnified by the explosion in the number of ASs and route entries, the resulting flapping of routes would lead to severe stability problems throughout the infrastructure.

Secondly, the configuration and maintenance of BGP requires a level of knowledge and expertise that is very unlikely to be available to most community members. Misconfiguration is already a source of problems within the existing Internet infrastructure [4], leading to unnecessary routing load and instability in the core routing tables. Poorly configured systems within community-area contexts would serve only to add to load, and propagate further any stability problems throughout the network.

Finally, employing existing BGP mechanisms requires that the equipment used is capable of supporting the relevant protocol and policy systems; this is particularly infeasible within resource-poor mobile or personal-area network environments, and would effectively exclude them from participating in such community network initiatives.

***Multi-homed virtual* edge entities:** Traditional intra-domain ad hoc routing mechanisms [5] have focused on finding the single most efficient route on a source–to–destination basis, where the destination may be either inside or outside the local domain or ad hoc network. This models the domain or ad hoc network as a single AD that is, either disconnected, or a direct extensions of a larger infrastructure. This in turn requires them either to discover efficient routes to a very wide set of destinations, or to route towards a specific designated domain or network gateway (representing a single point of failure).

However, the multi-homed nature of the newly emerging class of community networks allows them to be seen as composite, virtual entities. They reside at the edge, but remain connected to, the Internet.

The establishment of trust relationships: In the past, ad hoc and opportunistic networking approaches have focused on the automated discovery, negotiation and routing between neighbouring nodes that are all assumed to trust each other. However, this cannot be applied to the community network environment because it is very unlikely that members would be willing to trust all others unconditionally.

The newly emerging class community networks are organised instead at the human level. This may be through either personal meetings or other forms of out-of-band interaction. This implies a basic, local, level of trust before any peering agreements between them can be reached, so, a level of trust is implied. However, this type of coordination is unsustainable in the long term and inhibits the future evolution of such community networks. These out-of-band interactions limit the efficiency with which community networks can be formed and may expand, requiring human-level intervention at every stage.

3 Existing Work

The CUWiN project [6] has developed and released open source software to enable wireless community-area networking. The aim of the project is to allow “users to buy bandwidth in bulk and benefit from the cost savings”. The project offers an installation CD that simplifies and automates all technical set-up and configuration, including “loading the networking operating system and software, sending out beacons to nearby nodes, negotiating network connectivity, and assimilating into the network”. However, this level of automation means that there is a lack of emphasis being placed on the level of local control available to individual users. This is unlikely to be acceptable to all individuals wishing to participate in the newly emerging class of community-area network initiatives discussed here. The CUWiN approach relates closely to existing ad hoc networking approaches where connectivity between nodes is open and fully-automated, through the transmission of beacons; new nodes transmitting beacons are incorporated automatically into the mesh network. Thus the resulting mesh architecture appears to become a single edge network that extends the larger Internet, providing similar single-path source–to–destination routing mechanisms for traffic. Additionally, the CUWiN software appears to require higher-powered machines rather than potentially low-powered mobile devices. This limits the range of possible deployment scenarios, excluding any involving lower powered and mobile devices.

The ‘MAR commuter mobile access router’ [7] provides an architecture for aggregating multiple heterogeneous types of wide-area connectivity. However, it focuses on a multi-homed *hotspot* model of access with the placement of a 'MAR' device in moving vehicles. The device provides a range of local connectivity access (wired and wireless) for commuters. It is connected to the wide-area via multiple wireless interfaces, which it uses "simultaneously, to build a better combined wireless communication channel" and to provide bandwidth aggregation; externally it appears as a NAT box. However, this relies on all local users gaining wide-area access via a single provider (i.e. the MAR device) and thus represents a single point of failure. It does not take into account the possibility that individuals may have some wide-area connectivity that could be better utilised.

The 7DS Peer-to-Peer Information Dissemination and Resource Sharing system [8] provides a mechanism for self-organised connection sharing. However this focuses on a more traditional model of sharing individual wide-area connections among multiple devices, specifically when such connections are temporarily idle, by treating the mobile device as a temporary gateway. However, there is an opportunity to better utilise all available wide-area links to take full advantage of a higher aggregate wide-area capacity available to all users. Load balancing mechanisms are also provided in 7DS, but again these are based on the selection of single (least loaded) gateways rather than distribution across multiple CEFs as undertaken by the coalition-based approach.

Thus, there are a number of existing projects that enable community-area networking. However, by focusing on solutions that rely on the use of existing mechanisms for administration and operation, they suffer the same shortfalls that the underlying mechanisms have when applied to the newly emerging class of community-area networks.

4 The Coalition Peering Domain

Our solution aims to formalise the ad hoc nature of relationships within existing community network initiatives, while allowing administrative responsibility to remain distributed, allowing community members to retain local control.

4.1 Principles

Figure 1 illustrates a number of collaborative efforts or 'local peering agreements' between pairs of community members. These peerings may be either as simple as links interconnecting different pairs of community members, or more complicated associations controlled through policy defined locally by the community members. As the numbers of such local peering agreements begin to increase and to intersect between community members, we refer to the creation of a *coalition* within the community and the formation of a *Coalition Peering Domain (CPD)*.

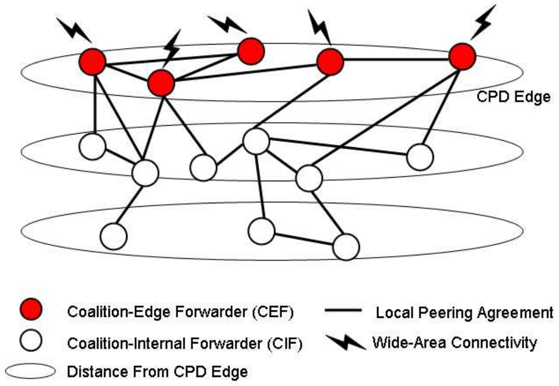

Figure 1: CPD Architecture

Each Coalition Member (CM) may represent an individual with either a single node, or a local network. Coalition members who have wide-area connectivity (or more generically, connectivity outside the CPD) form together the edge of the CPD and act as *Coalition-Edge Forwarders (CEFs)*; they are the CPD ingress–egress points, allocating some proportion of their external connectivity for this purpose. In the simplest case they may forward outgoing packets on their CPD-egress link. However, in a more interesting case they may forward some of these outgoing packets by 'spraying' (distributing) them, across the CPD edge, via their CPD-internal interfaces to other member CEFs within range, who then forward the packets outside the CPD. Thus outgoing traffic is distributed across multiple CEFs, so enabling a higher upstream data rate by aggregating multiple CM egress links. This type of wide-area connectivity aggregation is an example of collaboration between individuals for mutual benefit. This approach is useful when the local capacity between a number of CMs is greater than or equal to their individual egress capacity to a common remote entity.

Coalition members who do not have connectivity outside the CPD, or who choose not to make available their wide-area capability to other CMs, act as *Coalition-Internal Forwarders (CIFs)*. The forwarding of CPD-internal traffic (the traffic traversing between CMs) may be performed using modified forms of standard inter-domain or ad hoc routing protocols. In the example of collaboration for the purpose of wide-area connectivity aggregation, CIFs forward CPD-outbound traffic by directing it towards their 'nearest' CEF for CPD egress. This traffic can be sprayed across the CPD edge by the receiving CEF as described above, thus, CIFs may also benefit. Of course, CIFs may also use mechanisms for load balancing and take responsibility for spraying directly to multiple CEFs, depending on the physical connectivity of the CPD.

4.2 Basic Performance Metrics

We present here a very basic set of performance metrics to try and highlight the benefits and the implications of this approach. For the time being, we will ignore traffic models and effects of media access control, but including such analyses would affect only the detail of the analyses and not the general principle.

Let us assume the simplest case where there are N CMs, all within radio range, and so can communicate with each other; in this case the entire CPD makes use of a single shared media with total capacity C_T (e.g. either 11Mb/s for 802.11b or 55Mb/s for 802.11g). Let us assume also that all N CMs are CEFs peering together in a circular topology, and that they treat all packets with equal priority, attempting first to egress traffic directly, then to spray it to a neighbouring CEF if the buffer is full. The wide-area egress of packets at each CEF is represented by an M/M/1/b process with a packet arrival rate of λ and a packet egress rate of μ. As each CEF's buffer size is limited to b, $\lambda\pi_b$ arriving packets find the buffer full and are thus either sprayed to the next neighbouring CEF, or dropped. Therefore, at each CEF, λ is composed of packets generated locally, and sprayed packets arriving from the previous neighbouring CEF.

The probability that a packet is egressed from the CPD can be expressed as a geometric series:

$$\begin{aligned} P_{EGRESS} = (1-\lambda\pi_b)+ \\ (\lambda\pi_b(1-\lambda\pi_b))+ \\ (\lambda{\pi_b}^2(1-\lambda\pi_b))+\cdots \end{aligned} \quad (1)$$

We assume that a packet may be sprayed repeatedly by consecutive CEFs until it cycles back to the originating CEF in the circular topology, after which it may be dropped if the buffer is full. Therefore, the probability that a packet is egressed can be expressed as the sum to N of the geometric series in eqn 1:

$$\begin{aligned} P_{EGRESS} &= \frac{(1-\lambda\pi_b)(1-\lambda{\pi_b}^N)}{1-\lambda\pi_b} \\ &= 1-\lambda{\pi_b}^N \end{aligned} \quad (2)$$

However, in a traditional single path routing M/M/1/b process with no packet spraying, we find that the probability that a packet is egressed is reduced by order N to:

$$1-\lambda\pi_b$$

Assuming a general intra-CPD communication traffic level of C_C (i.e. non-coalition-egress traffic), the total CPD capacity available for packet spraying C_S is:

$$C_S = C_T - C_C \quad (3)$$

Of the N CEFs that form the CPD edge, each one, n, provides some ingress capacity, C_{I_n}, and some egress capacity, C_{E_n}. The total CPD ingress capacity, C_I, and the total CPD egress capacity, C_E, are respectively:

$$C_I = \sum_{n=1}^{N} C_{I_n}$$

$$C_E = \sum_{n=1}^{N} C_{E_n}$$

We assume that the total CPD capacity available for packet spraying is divided equally between all CMs:

$$C_{S_n} = \frac{C_S}{N} \quad (4)$$

We define the ingress and egress *CPD gain factors*:

$$G_I = \frac{C_{S_n}}{C_I} \quad (5)$$

$$G_E = \frac{C_{S_n}}{C_E} \quad (6)$$

These gain factors may be used by CMs to assess the potential benefit of joining a CPD. They allow also some assessment of the overall performance gain for the CPD. When the gain factor is greater than 1, it indicates that there is still a benefit from adding further CEF ingress or egress capacity respectively. When the gain factor is less than 1, it indicates that there is no benefit to the CPD. The optimum value for the gain factor is 1 and the CPD may only gain further if C_S (and so C_{S_n}) were to increase.

If we take an example of a CPD that uses an 802.11g channel for intra-CPD data transmission, and that consists of 8 CEFs each with an ADSL connection to the wide-area of 1Mb/s downstream and 256Kb/s upstream, assuming in the best case that there is no general intra-CPD communication traffic (i.e. in eqn 3, $C_C = 0$, so $C_S = C_T = 55Mb/s$), and that all CEFs choose to allocate all available external capacity for edge forwarding, we find from eqn 4 that a given CM m's available capacity for packet spraying $C_{S_n} = 6.875Mb/s$ and thus from eqn 6 that the gain for using the packet spraying mechanism solely to increase coalition-egress traffic is approximately 3.44. This means that the coalition-egress capacity may be increased by a factor of 3.44 before the CPD ceases to benefit further.

By increasing all member CEFs' ADSL connection upstream speed to 512Kb/s, we find from eqn 6 that the corresponding gain is approximately 1.72. This means that the coalition-egress capacity may be increased by a factor of 1.72 before the CPD ceases to benefit further.

We may take another example of a CPD that uses an 802.11b channel for intra-CPD data transmission, and that consists of 4 PDAs acting as CEFs, each with 3G wide-area connectivity of 384Kb/s downstream and 128Kb/s upstream. In this case, we find from eqn 6 that there is a greater corresponding gain factor of approximately 5.5, even though the intra-CPD channel has a lower data transmission capability.

4.3 Challenges

Although the CPD architecture provides a means to enable easier collaboration between individuals, while maintaining local control, it is still disruptive to the existing architecture and service provision models. There are a number of challenges that must be overcome.

CMs need a mechanism to communicate with each other once they have formed local peering agreements. This implies that some form of addressing scheme should be employed within a CPD. However, addressing is a centralised function that would, in this context, need to be applied to a distributed system. This is a non-trivial task and needs careful consideration.

Although the use of Network Address Translation (NAT) devices may provide an obvious solution for community-area networks, they pose a number of problems that may limit the overall usefulness of the CPD. There is an increased chance that the arbitrary use of private addressing may lead to clashes between potentially peering CMs. They limit the operation of some types of applications as well as the abilities to apply security at the IP layer, and ultimately introduce unnecessary complexity in configuration and maintenance.

If we assume that all CEFs have a globally reachable address (allocated to them by their wide-area connectivity provider), then CEFs may communicate with each other by using these global addresses but routing them locally via the local peering agreement links. This can be enabled through the operation of an existing (or modified form of) routing protocol between the CEFs across the CPD edge. If we assume also that CEFs have additionally a block of addresses that they may sub-allocate (e.g. an IPv6 /64 block allocation), then CEFs may sub-allocate portions of this address space to any CIFs with which they peer. Routing of traffic to CIFs can thus take place through the address allocation hierarchy. This would also behave well with reverse path or CPD-ingress traffic because receiving CEFs can forward it to the allocating CEF for onward CPD-internal routing.

Even though traffic destined for a specific remote destination may be sprayed across the CPD edge, the reverse path still relies on standard routing. This means that individual CEFs may be burdened with a greater volume of return path CPD-ingress traffic. However, the asymmetry of most wide-area connectivity technologies may be sufficient to offset this inequality. Alternatively, the burden for reverse spraying may be placed on either the remote party, or a provider-controlled device located beyond the CPD edge. Such multi-path routing would however, have implications for higher layer protocols that rely on the underlying routing infrastructure. Traffic may arrive at its destination with some delay or in an unordered fashion thus causing problems for delay-intolerant applications.

De Couto et al [9] have shown through a number of experiments that, the forwarding of packets on a shortest path basis within multihop wireless networks would be unlikely to result in a choice of paths with the best throughput. Thus, our earlier assumption that CIFs should forward CPD-outbound traffic by directing it towards their 'nearest' CEF for CPD egress may not be an efficient mechanism for the routing of CPD-outbound traffic within a CPD formed with wireless links. The forwarding of packets to the CPD edge may need to be re-evaluated by each CIF along each hop, depending on whether the default path of local peering agreements to the CPD edge contains poorer quality links than alternative routes.

5 Summary and Conclusions

The formation of community networks is a growing trend. It has been particularly aided by recent advances in local-area network technologies, making them much more affordable and easily available. We have presented an architectural element that would enable groups or communities of individuals to better utilise their wide-area connectivity resources, through collaboration, by using the local-area connectivity between them. This is achieved by adding structure, the **Coalition Peering Domain**, to the otherwise ad hoc community inter-networks residing at the edge of the Internet.

Although we have focused on community-area networks, there are a number of applications and scenarios that may benefit from the CPD architecture. The multi-homed nature of the CPD architecture may be especially useful for networks that require a high degree of robustness or survivability in connectivity to the wide-area. Assuming that a CPD is formed with multiple CEFs, should one CEF fail or become disconnected from the wide-area, the CPD as a whole still maintains some wide-area connectivity.

Scenarios involving the co-ordination of evacuations in an emergency, and scenarios involving the co-ordination of aid and relief efforts following a natural disaster are both examples of situations in which a diverse set of devices and communications technologies may be used. The available device resources and connectivity capabilities may be quite limited in some cases. Such scenarios may benefit greatly from the formation of a CPD for the purpose of wide-area connectivity aggregation, allowing better utilisation of the available connectivity to the wide-area.

In fact the formation of a CPD for the purpose of connectivity aggregation can be useful in any scenario where there is some degree of heterogeneity in the available wide-area connectivity, and where the data rates of local-area connectivity exceed individual wide-area data rates.

In conclusion, we take the position that a coalition-based approach would enable individuals to share connectivity resources in a controlled manner, but there are a number of technical challenges that should be researched further.

References

[1] Consume “trip the loop, make your switch, consume the net”. http://consume.net/.

[2] FreeNetworks.org volunteer co-operative association. http://scoop.freenetworks.org/.

[3] Y. Rekhter and T. Li. A Border Gateway Protocol 4 (BGP-4). RFC 1771 (Draft Standard), March 1995.

[4] Ratul Mahajan, David Wetherall, and Tom Anderson. Understanding BGP misconfiguration. In *SIGCOMM ’02: Proceedings of the 2002 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 3–16, New York, NY, USA, 2002. ACM Press.

[5] IRTF RRG Ad hoc Network Systems Research Subgroup. http://www.flarion.com/ans-research/.

[6] Champaign-Urbana Community Wireless Network. http://www.cuwireless.net/.

[7] Pablo Rodriguez, Rajiv Chakravorty, Julian Chesterfield, Ian Pratt, and Suman Banerjee. MAR: a commuter router infrastructure for the mobile internet. In *MobiSYS ’04: Proceedings of the 2nd international conference on Mobile systems, applications, and services*, pages 217–230. ACM Press, 2004.

[8] M. Papadopouli and H. Schulzrinne. Connection Sharing in an Ad Hoc Wireless Network among Collaborating Hosts. In *Proc. International Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV)*, pages 169–185, June. 1999.

[9] Douglas S. J. De Couto, Daniel Aguayo, Benjamin A. Chambers, and Robert Morris. Performance of multi-hop wireless networks: shortest path is not enough. *SIGCOMM Comput. Commun. Rev.*, 33(1):83–88, 2003.