The Coalition Peering Domain: A New Entity In The Routing Landscape*

Manish Lad, Saleem Bhatti, Peter Kirstein, Stephen Hailes
Department of Computer Science,
University College London,
Gower Street, London, WC1E 6BT
{M.Lad | S.Bhatti | P.Kirstein | S.Hailes}@cs.ucl.ac.uk

15 March 2005

Abstract

Current business models for the service provision of domestic Internet connectivity focus on individual users or parties. In the local-area, the IEEE 802-standards technologies (wired and wireless) promote easy interconnectivity and resource sharing between local users. This has stimulated the appearance of ***community networks*** — ad hoc networks residing at the edge of, but still connected to, the Internet. Currently, such activities are seen as both disruptive and difficult to sustain, breaking traditional network-service business models and causing discontinuity of the network architecture. We introduce a new architectural entity, the ***coalition peering domain***, that would allow structure and control to be added to such ad hoc edge networks. We highlight addressing and routeing challenges to be resolved for the inclusion of such network usage within the existing network architecture.

1 Introduction

The discussion and position we present in this paper looks at a particular usage of community-oriented network connectivity at the edges of the Internet; that of aggregating network capacity across the community network members. We take the position that this desire for aggregate capacity is likely to grow. We propose an outline of how this usage may evolve and be included in the general Internet architecture by highlighting various relevant challenges. By its nature, the viewpoint presented here is forward-looking but takes examples from network usage scenarios that are in existence today. The work is ongoing.

1.1 Background

Local Area Networking capabilities allow users to interconnect easily multiple machines or devices to utilise more efficiently and flexibly both their local resources and their access to the wide-area. Use of both wired Ethernet and wireless IEEE 802.11 standards have increased, while costs of both types have fallen significantly in recent years; manufacturers now integrate them into their equipment (e.g. laptops, desktops and ADSL gateways). Many consumer operating system platforms (such as Windows and MacOS X) also provide improved networking support. They enable very simple local network set-up in a plug-and-play manner by configuring 'connection sharing' automatically through a combination of Network Address Translation (NAT) and automatic address allocation using the Dynamic Host Configuration Protocol (DHCP).

Advances in Internet and wide-area access technologies also have had a great impact on connectivity both for the home user and for the mobile user. The data rates of wired Internet access technologies (e.g. ADSL) have increased substantially, improving the speed at which users are able to download and access content, and fueling the demand for various types of (multimedia) content that often involve large downloads of data. Additionally, peer-to-peer applications, especially peer-to-peer file-sharing applications, have become popular and so users may now be interested in high upstream capacity on access links as well as high downstream capacity.

*Department of Computer Science, University College London, Technical Research Note Number: RN/05/03

However, it is only very recently that levels of subscription to technologies like ADSL have increased to their current levels in some parts of the world. Meanwhile, in many parts of the world, there still remains a large proportion of home users for whom primary connection to the Internet is through older, slower technologies: mainly analogue modems but some ISDN. There still remains also a large number of users for whom connectivity in a mobile environment is through a GSM connection. Thus we have a growing situation where there is very good local-area connectivity supported by wired and wireless networking technologies, but where wide-area data rates (wired and wireless) are improving relatively slowly and vary greatly leaving a large disparity.

1.2 Coalition-based peering required

We take the position that to exploit fully their available wide-area connectivity, users should collaborate, exploiting the statistical multiplexing gain in scenarios where wide-area links are shared. We illustrate this with two examples: a local community, within a geographically localised neighbourhood; and a group of mobile users, within proximity of each other. In both examples, the users have local wireless connectivity (e.g. using 802.11 standards), and have Internet connectivity; although local community users have relatively higher data rate connectivity (e.g. using wired DSL or cable modems) than mobile users (e.g. GSM, GPRS abd 3G).

To introduce structure and control to this environment we propose a new architectural entity within the routeing landscape, the ***coalition*** peering domain, in which collaborating users within a community together form a larger inter-network that allows them to aggregate their individual network connectivity.

Our work in this area is just beginning, and we seek to properly define the problem space. We present here a discussion as a series of questions and suggest some possible directions to follow to find solutions.

In Section 2 of this paper, we discuss the ideas of coalition-based networking: what coalitions are and how they might operate. We discuss in Section 3 why coalitions are needed and we present in Section 4 the architectural challenges they may introduce. We present relevant work in Section 5 and conclude in Section 6 with a summary.

2 What is a Coalition Peering Domain?

2.1 What leads us to propose the use of coalitions?

Exploiting the local-loop to avoid the costs of laying new data network cabling has driven the evolution of analogue modem technologies and then lead to the development of domestic ISDN and xDSL services specifically designed for digital connectivity. With the ever increasing desire for faster connectivity from users, xDSL services now offer multi-megabit data rates using the same physical infrastructure that at one time offered only a few 10s of Kb/s. Research continues into pushing further the limits of the existing local-loop. However, the legacy of the installed-base — features such as poor cabling installations, distance from exchanges and sometimes the slow enabling of exchange equipment — still restrict service provisioning for some users. Additionally, although the data rates offered by such wide-area connectivity are approaching the lower end of those possible in the local-area (a few Mb/s), they still fall well short of those possible at the higher end (up to 100Mb/s with proprietary extensions to 802.11g, and up to 1Gb/s with Ethernet).[1]

The wide-area wireless data market is less mature than the wide-area wired market thus, disparities between local-area and wide-area data rates are more pronounced. Wireless wide-area connectivity currently offers data rates of a few 10s of Kb/s with plans for 3G systems to offer a few 100s of Kb/s (or perhaps a few Mb/s at best). Yet many advanced services readily could be made available if these data rates were to improve. Therefore, even as connectivity technology advances in the mobile device arena, services shall also advance and so users will still desire higher data rates.

So, for those with poorer domestic connectivity and for those in mobile environments, the coalition-based approach to connectivity proposed here provides a possible solution for exploiting wide-area connectivity by aggregating individual links, thus bringing greater wide-area capacity to the coalition over any existing infrastructure, wired or wireless. It may be useful also more generally even where connectivity is richer. However, the approach raises important architectural challenges related to routeing.

[1] However, some countries including South Korea and Japan, have more advanced wide-area DSL connectivity of 100Mb/s and higher.

2.2 What is meant by a Coalition Peering Domain?

A coalition is defined as [2]:

1. a : the act of coalescing : UNION b : a body formed by the coalescing of orig. distinct elements : COMBINATION
2. : a temporary alliance of distinct parties, persons, or states for joint action

In static home and local community environments, relatively cheap 802.11-based wireless access routers can be used with extension antennae: omni-directional antennae extending the range to potentially a few hundred metres, and directional antennae allowing connectivity up to several kilometres. These have allowed some users to connect together directly their home networks, creating small inter-networks on an ad hoc basis to form a local neighbourhood *mesh* or *community network*. As the number of such initiatives grows, affiliations are formed to promote their use and growth [1]. This connectivity relies on individual, and usually informal, ad hoc peering agreements between those within radio frequency range of each other's wireless base-stations. However, as the numbers of such peering arrangements increase and begin to overlap, we can talk of a *coalition* within the community and the formation of a *Coalition Peering Domain (CPD)*.

2.3 How might a Coalition Peering Domain operate?

Figure 1 illustrates a number of local peering agreements between individual Coalition Members (CM). Some of these members are single nodes (individual users with a single machine), while some represent local networks (users with multiple machines networked together, for example a home or personal-area network). Multiple sets of peering agreements join together to form a single overall community or *coalition*. Such coalitions may be formed on specific premises agreed between individual peers or across the coalition (e.g. basic peering, traffic forwarding, resource provision, resource pooling etc.).

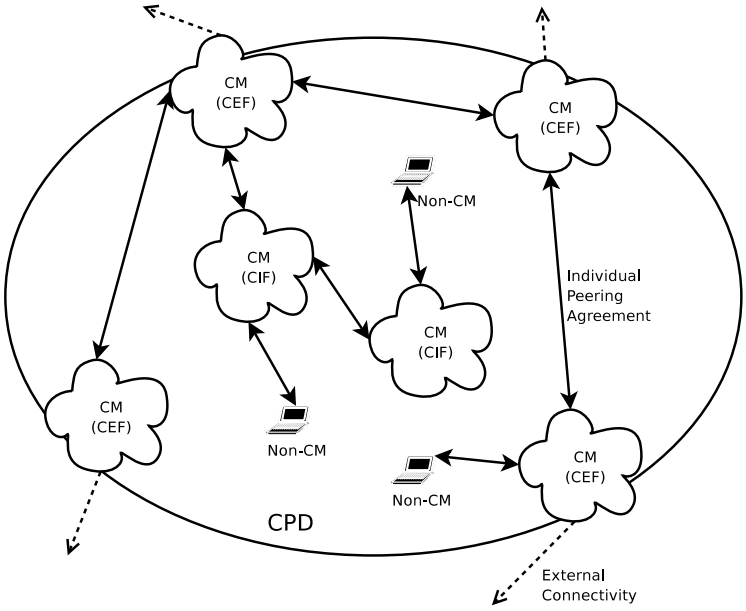

Figure 1: Coalition-Based Connectivity Architecture

This architecture allows a local neighbourhood community or a group of mobile users to pool their connectivity resources, through the combined (wired or wireless) wide-area connections of a subset of the CMs.

Coalition members who have wide-area connectivity (or more generically, connectivity outside the coalition) are said to reside at 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

[2] Merriam-Webster Dictionary Online http://www.m-w.com/

may forward outgoing packets on their coalition-egress link, but in a more interesting case they may forward some of these outgoing packets by 'spraying' (distributing) them via their coalition-internal interfaces to other member CEFs within range, who then forward the packets outside the CPD. Thus outgoing traffic is distributed across many CEFs, potentially enabling a higher upstream data rate by aggregating all CM-uplinks. This approach is especially useful when the local capacity between a number of CMs is greater than or equal to their individual uplink capacity to a common remote entity (e.g. the Internet), and their individual connectivity is not fully loaded.

Coalition members who do not have connectivity outside the coalition act as *Coalition-Internal Forwarders (CIFs)*. The forwarding of coalition-internal traffic (that sourced and sinked within the CPD) might be performed using modified forms of standard inter-domain or ad hoc routeing protocols as agreed within the coalition. However CIFs forward coalition-outbound traffic by directing it towards their 'nearest' CEF for coalition egress. This traffic can be sprayed across peer CEFs by the receiving CEF as described above. 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.

In this context, coalition members represent a reasonably static group of nodes or local networks that form peering agreements between each other. It is also possible for such connectivity to be extended to non-coalition members, for example mobile/roaming nodes travelling through the CPD. These may peer dynamically with a CIF or directly with a CEF as they pass within radio range.

2.4 How might a community or an individual benefit?

We present here a very basic set of performance metrics to try and highlight the benefits and the implications of this approach with respect to the amount of aggregate traffic that can be handled by the coalition. For the time being, we ignore traffic models and effects of media access control.

Let us assume the simplest case where all CMs are 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). 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:

$$C_S = C_T - C_C \qquad (1)$$

We assume that the CPD edge consists of M CEFs, each CEF, m, providing some ingress capacity, C_{I_m}, and some egress capacity, C_{E_m}. The total CPD ingress capacity, C_I, and the total CPD egress capacity, C_E, are respectively:

$$C_I = \sum_{m=1}^{M} C_{I_m} \quad (2) \qquad\qquad C_E = \sum_{m=1}^{M} C_{E_m} \quad (3)$$

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

$$G_I = \frac{C_S}{C_I} \quad (4) \qquad\qquad G_E = \frac{C_S}{C_E} \quad (5)$$

These gain factors may be used by prospective 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 were to increase.

For example, in 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 1, $C_C = 0$, 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 5 that the gain for using the packet spraying mechanism solely to increase coalition-egress traffic is 27.5. This means that the coalition-egress capacity may be increased by a factor of 27.5 before the CPD ceases to benefit further.

By increasing all member CEFs' upstream ADSL connection speeds to 512Kb/s, we find from eqn 5 that the corresponding gain is 13.75. This means that the coalition-egress capacity may be increased by a factor of 13.75 before the CPD ceases to benefit further.

We may take another example of a CPD that consists of four PDAs acting as CEFS. They share an 802.11b channel for intra-CPD data transmission, and each have 34Kb/s GPRS wide-area connectivity. In this case, we find from eqn 5 that there is a greater corresponding gain factor of approximately 80.

Of course, as well as gaining capacity, the CPD gains robustness from being multi-homed.

3 Why do we need a Coalition Peering Domain?

Why do we need a *coalition?*. Why can we not use existing routeing architectures and protocols?

The formation of local peering agreements and their grouping into a CPD may initially be thought of as special cases of Administrative Domains (AD), ad hoc networks, or even conventional IP edge networks. Externally, a CPD may appear as a single multi-homed entity resembling an AD. Internally, the topology of a CPD may resemble that of an ad hoc network. The question thus arises as to why the coalition-based approach being proposed here is necessary because extra-CPD or inter-CPD communication may be possible through the use of existing EBGP protocol mechanisms, and intra-CPD communication may be possible through the use of intra-AD IBGP, existing ad hoc protocols or inter-domain routeing protocols. However, although characteristics of existing routeing mechanisms exist within the community network scenario, there are significant differences in the coalition-based approach proposed here. The implementation of CPDs present challenges for routeing, policy, trust establishment, addressing and management of connectivity, that are not tackled directly by either existing BGP, existing ad hoc routeing, or existing intra-domain IP routeing mechanisms.

A CPD is a *collection* of collaborating, small ADs: Although it may appear so externally, a CPD does *not* represent an AD under the control of a single organisation or entity, but rather a collaborative group of such entities. This is because administrative responsibility is distributed across the CPD with each CM maintaining a degree of autonomy that provides local control over its own resources. Thus a CPD represents a multi-homed *virtual* entity residing at the edge of, and connected to, the Internet (or possibly another CPD). To implement a BGP-based peering system, it would be necessary for each CM to be allocated an Autonomous System (AS) number and to run EBGP.

However, as each CM only represents a small local network or a single node, implementing this would raise significantly the level of complexity by increasing exponentially the number of routeing table entries across the CMs. Although financial costs may not be a problem as BGP implementations are available for PC-based platforms[3], there are some pracitical issues. The CMs are likely to be transient (users turning their equipment on and off as required), this could cause stability problems with using BGP. Also as coalitions are formed between domestic users, they are unlikely to have the expertise (or the inclination) to be able to configure, maintain and administer a network using BGP. Therefore, applying the BGP approach to the CPD is neither feasible nor desirable.

CPDs require some sharing of policy between CMs: To be able to form a CPD, CMs must share some aspects of common policy including the criteria by which they may form the coalition or by which they may join an existing coalition. This is in contrast to the usage of BGP where there is no requirement to share policy.

Policy mechanisms and peering agreements also need to be examined carefully. For example, reaching a particular local peering agreement may depend on certain criteria that dictate which services a given CM *must* or *must not* provide, both within that local peering agreement and within any local peering agreements that it forms with other CMs.

Some more mature groups are already beginning to establish simple policy-based approaches [1].[4]

Routing may be distributed across a CPD edge: Appearing externally as a multi-homed virtual entity, a CPD is able to distribute routeing functions across its edge through the use of the packet spraying mechanism described earlier. Traditional intra-domain routeing and ad hoc network approaches have focussed 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 domains and ad hoc networks as a direct extension of a larger infrastructure, thus requiring them to either discover efficient routes to a very wide set of destinations, or to route towards specific designated domain or network gateways. However, the coalition-based approach

[3] For example: the Quagga Routing Suite (http://www.quagga.net/) or the XORP eXtensible Open Router Platform (http://xorp.org/).

[4] As use of the coalition-based approach matures, successfully developed 'standard' policies with known semantics could be made available openly in a similar model to software licencing repositories (e.g. http://opensource.org/)

focuses on finding a route to the *edge* of the CPD. From there, packets are distributed across the edge of the CPD to take advantage of the aggregate uplink afforded by the CPD. This means that intra-CPD routeing need only discover efficient routes to a small set of destinations (i.e. one or more of the CPD's CEFs) thus easing the burden (perhaps through long-lived caching) on potentially resource-poor CIFs. This also provides a robustness as multiple coalition-egress points are utilised simultaneously.

Optimisation in the CPD: The use of localised metrics may be required to optimise routeing within the CPD, through the selection of CPD-internal paths and paths to the CPD-edge. Such optimisation is not possible with a path vector approach as used in BGP. Modified forms of intra-domain protocols or ad hoc approaches would be required. Additionally, with CMs likely to be transient, any routeing protocol would also need to be able to detect topology changes quickly, adapt to those changes and stabilise in a timely manner.

Mechanisms for load-balancing also may need to be employed between CEFs distributing egress traffic amongst each other, as well as by CIFs taking responsibility to distribute traffic evenly across CEFs.

CPD formation involves establishing trust relationships: In the past, ad hoc and opportunistic networking approaches have tended to focus on automated discovery, negotiation and routeing between neighbouring nodes that are assumed to fully-trust each other. However, as with current community area networks [1, 2], coalitions are likely to be organised at the human-level. This may involve either personal meetings or other forms of out-of-band, high-level interaction to establish trust. Although the subject of trust establishment is a complex problem, the human-level interaction implies that there must be a basic level of trust before local peering agreements can be reached, so a trust mechanism is implicit to a certain degree. This may evolve over time, based on a number of factors including direct past experience, and also recommendations or references dependant on the level of trust-weighting placed by a CM on the recommender.

Although CPD formation may not be possible in a totally automated fashion, some levels of automation could be achieved through the application of policy over existing automation and discovery mechanisms.

4 What architectural challenges may coalitions introduce?

We see from the discussion in the previous section that the implementation of CPDs present challenges for routeing, addressing and management of the connectivity that are not tackled directly by either existing ad hoc routeing mechanisms or existing IP routeing mechanisms.

4.1 Routing

The coalition-based approach to connectivity raises a number of issues for the operation of transport layer protocols that rely on the network layer IP address as part of the transport protocol state (such as TCP). By spraying coalition-egress packets for onward routeing across a multi-homed CPD's edge, a receiver shall receive packets that have the same source IP address but that may have taken multiple paths to their destination. Response packets from the receiver shall however follow normal routeing back to the source address and so enter the CPD through a single CEF link.

Attempting to increase the coalition-ingress data rate to exploit the multi-homed CPD would require either the remote party to have sufficient knowledge to spray CPD-ingress traffic across the destination CPD's edge (across multiple CEFs), or the placement of a localised (coalition-area) 'middlebox' at a provider's premises handling reverse spraying across the CPD edge. Both these methods are problematic as they tie mechanism to policy and to provisioning within the network. The former burdens the remote party with the storage of extra state and policy while the latter increases the number of points of failure within the network and may require all CMs to be subscribed to one provider. It is essential that mechanism is separated from policy and that neither policy nor mechanism is placed 'within' the network path for maximum scalability and flexibility of deployment.

Spraying packets solely to increase the coalition-egress data rate, and accepting return packets through a single CEF, still provides a gain for CPD members as demonstrated in the earlier examples.

4.2 Addressing

Each peering agreement is uniquely formed between peering CMs. How then may CMs make multi-hop forwarding decisions to determine both the destination of intra-CPD traffic and the shortest path to a CEF for coalition-egress traffic? Unique addressing between peering CMs must be propagated across the CPD to allow intra-CPD

routeing. Additional issues arise with the use of functions such as firewalls and NATs.

To resolve this, the formation of a CPD requires some management of addressing that spans all local peering agreements that form the CPD. This is non-trivial because it involves the operation of a centralised function (address allocation) within an environment where administrative responsibility is fully-distributed. A 'distributed NAT' function may be required across the CPD edge so that CEFs can co-ordinate address allocation and packet forwarding within the CPD. Ingress and egress filtering on firewalls would need to be aware of the 'distributed NAT' function and the address allocation.

4.3 Inter-CPD communication

It may be possible for a CM from one CPD to form a local peering agreement with a CM from another CPD. It is also quite likely that within one CPD, there may be CMs who are members simultaneously of multiple other CPDs. This may be because their membership of each CPD serves a distinct set of their requirements, it may be a result of geographic separation, or it may be simply a result of the human-level interaction (or lack there of) and negotiation undertaken to form a CPD. These CMs provide a bridge between the CPDs.

Losing a CEF's egress capability within a multi-homed CPD that has multiple uplinks to the wide-area does not affect the CPD's reachability to the wide-area, only its aggregate capacity. In contrast, a bridging-CM is far more critical because it is likely to be the only (or one of only very few) CEFs within a CPD that provides connectivity to another CPD. Should this fail, traffic destined for the remote CPD must traverse the wide-area.

4.4 Other challenges

Considering our discussion above, there are other relevant challenges, which we do not discuss here due to lack of space. These include the effects of multipath routeing such as effects on packet delay characteristics and reordering of packets within flows. There is also the problem of managing the CPD: as well as the addressing and routeing issues already noted, there is the problem of managing CPD peering agreements and trust relationships. Some level of automation for such mechanisms will be required so that the systems are usable by domestic users.

Also, we have considered that the CPD appears as a single multi-homed entity and there is no visibility of the CPD semantics outside the CPD. However, there may be a business opportunity for providers in that they may choose to support directly connectivity to CPDs and so this would change the external interaction with other routeing entities such as ISPs.

There is also the discussion of how service provision business models would need to evolve to support the notion of a coalition and, related to this, the secondary benefits and opportunities that may result, such as localised service provisioning — there are many types of localised service [3] for which traffic need only traverse the local CPD or traverse directly between multiple CPDs via bridging-CEFs.

5 Surely someone has looked at this already?

The 'MAR commuter mobile access router' [4] provides an architecture that is somewhat close, in terms of connectivity aggregation, to the approach proposed in this paper. However, a key difference is that MAR focusses on a multi-homed *hotspot* model, rather than a collaborative distributed effort. The MAR device connects 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; 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).

The 7DS Peer-to-Peer Information Dissemination and Resource Sharing system [5] 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. The coalition-based approach we propose provides a greater degree of aggregation by distributing to multiple CEFs. 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.

Both the MAR and the 7DS systems may provide a number of valuable lessons for the development of the coalition-based approach proposed here.

The DIRAC software-based wireless router [6] provides a distributed router architecture composed of a Router

Core (RC) and a Router Agent (RA). This may be useful inside a CPD boundary where routeing functions can be shared and distributed, especially in scenarios involving inter-CPD communication. We plan to investigate further the merits of this within the CPD context.

6 Summary And Conclusions

We have presented an architectural outline that would enable communities or groups of users to better utilise their wide-area connectivity and resources through interconnection and collaboration via their local-area capabilities. This is achieved by adding structure, the **coalition peering domain**, to the otherwise ad hoc community networks residing at the edge of the Internet. The idea proposed may be attractive to both fixed local communities, and groups of users willing to collaborate in a long-lived mobile environment (e.g. a meeting room, a train journey, etc.)

The proposed approach requires investigations into a number of existing research areas including peering, addressing, routeing, trust and security policy within fixed and mobile wireless environments.

Although we have not dedicated a separate discussion on the subject of security, we have highlighted that without a basic level of trust between peers, local peering agreements and thus coalition peering domains cannot be formed. Varying degrees of trust and security would need to be established using a combination of human-level involvement, and distribution of administrative responsibility allowing individual coalition members to maintain local control of their own resources.

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

7 Acknowledgements

The initial ideas on the use of a coalition have arisen from the work of Defence Research & Development Canada (DRDC) on coalition-based dynamic VPN infrastructures, that focus on providing absolute local control for each coalition member site. Various discussions and deployment activities for their system were undertaken at UCL during 2003-2004.

References

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

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

[3] Robert X. Cringely. The Limits of SpongeBob SquarePants. One Canadian’s Wireless Neighborhood Network Could Someday Serve Us All, 30 September 2004. http://www.pbs.org/cringely/pulpit/pulpit20040930.html.

[4] 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.

[5] 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.

[6] Petros Zerfos, Gary Zhong, Jerry Cheng, Haiyun Luo, Songwu Lu, and Jefferey Jia-Ru Li. DIRAC: a software-based wireless router system. In *MobiCom ’03: Proceedings of the 9th annual international conference on Mobile computing and networking*, pages 230–244. ACM Press, 2003.