Challenges: Opportunities For Coalition-Based Community Networking[*]

Manish Lad, Saleem Bhatti, Peter Kirstein
{M.Lad | S.Bhatti | P.Kirstein}@cs.ucl.ac.uk

Department of Computer Science
University College London
Gower Street, London. WC1E 6BT

01 March 2005

Abstract

Current business models for the service provision of Internet connectivity focus on individual users or parties. In the local-area, the use of wireless technologies promotes easy interconnectivity and resource sharing between local users, leading to 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 a discontinuity of the network architecture. We introduce a new architectural entity, the ***coalition domain***, that allows structure and control to be added to such ad hoc edge networks. By examining a number of tussles that arise between parties with the adoption of such an approach, we show that it is feasible to include such network usage within the existing network architecture, and we discuss the challenges and the new opportunities that it brings with it.

1 Introduction

The discussion we present in this paper looks at a particular usage of wireless network connectivity at the edges of the Internet. We propose how this may evolve and be included in the general Internet architecture. 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.

Local Area Networking capabilities have improved greatly in recent years, allowing 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. A uniform set of hardware technologies enable such interconnectivity, through both wired Ethernet and wireless IEEE 802.11 standards. The costs of both types have fallen dramatically in recent years and 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-access technologies have increased substantially, improving the speeds at which users are able to download and access content — many have moved from analogue modems to digital subscriber lines (DSL). Multi-megabit wireless connectivity (approximately 10-100 Mb/s) also may be cheaply and easily configured within either the home or the local-area. Local-area connectivity based on IEEE 802.11 standards also offers configurable mechansims that provide basic security thus further improving users' confidence in those systems [1]. However 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

---

[*] This paper has been submitted to the 11th Annual International Conference on Mobile Computing and Networking. Paper: 1568957932, Status: Pending review. The authors' names and affiliations, and also the Acknowledgement section were removed before submission in accordance with the format guidelines for the double-blind review process.

[1] However some users do not actaully enable these security mechanisms and maintain only the basic factory configuration

environment is through a GSM connection. Thus we have a growing situation where there is very good local-area connectivity supported by wireless networking technologies, but where wide-area data rates (wired and wireless) are improving relatively slowly and vary greatly.

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 take as our main example in this paper a community network: a neighbourhood of users who are close to each other geographically, have wireless connectivity (e.g. using 802.11 standards), and have connectivity to the Internet (e.g. using wired DSL, cable or analogue modem; or wireless satellite, GPRS or 3G). There is a win–win situation: while users share connectivity to improve their data rates to the wide-area, they also create new opportunities for both service providers and equipment manufacturers.

In essence, we show a new architectural entity evolving in which groups or communities may use their individual local networks together to form a ***coalition*** domain.

In Section 2 of this paper, we discuss the idea of coalition-based connectivity. Section 3 explores the incentives for the various actors to adopt the coaltion-based approach. In Section 4 we comment on relevant work and then we conclude with a summary in Section 5.

2 Coalition-Based Connectivity

2.1 Motivation

There is a tension between the increasing demands of users and the capability of existing network technologies that provide access to Internet connetcivity. Users want to maximise the value for money that they receive from any product or service that they purchase. This leads to an increased demand to push the existing capabilities of connectivity and associated hardware to their maximum.

The evolving use of the copper local-loop infrastructure is a good example of this tension. 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 ISDN and xDSL services specficially designed for digital connectivity. With the ever increasing need and desire for faster connectivity from users, xDSL services now offer multi-megabit data rates using the same physical infrastrcuture that at one time only offered a few 10s of Kb/s. Research continues into pushing further the limits of the existing local-loop. The data rates offered by such wide area connectivity are approaching the lower end of the wireless data rates possible in the local-area (a few Mb/s), but they still fall well short of the higher end (up to 100Mb/s with proprietary extensions to 802.11g). However, the legacy of the installed-base (features such as poor cabling installations, distance from the exchange and the (sometimes) slow enabling of exchange equipment) restricts service provisioning for some users.

The wide-area wireless data market is newer than the wide-area wired market, at least on the consumer side. Wireless consumer data services accessible through relatively cheap, small, mobile networked devices with multiple interfaces (such as mobile phones and ‘super PDAs’), mean that use of these devices will increase. Such devices may have, for example, integrated Bluetooth (a few 10s of Kb/s to a few 100s of Kb/s) and 802.11b (at 11Mb/s). So, the disparity between the local area data rates and the wide-area data rates are more pronounced here; the wireless wide-area connectivity is currently offering data rates at 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). Therefore sharing wireless wide-area communication channels is even more attractive than sharing wired wide-area connections because of the potential relative gain.

The coalition-based approach to connectivity proposed here provides a supplimentary solution for exploiting wide area connectivity by aggregating individual links and so bringing greater wide-area capacity to the coalition from any existing infrastructure, wired or wireless.

2.2 Principles

The advancing markets for wireless communication technologies have enabled a potential shift in the model of access connectivity, taking some control away from connectivity service providers and placing it in the hands of consumers. Relatively cheap 802.11-based wireless access routers can be used with extension antennae: omnidirectional antennae extending the range to potentially a few hundered 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 local neighbourhood *community networks*. 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)*.

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 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 peers or acrosss the coalition (e.g. basic peering, traffic forwarding, resource provision, resource pooling etc.).

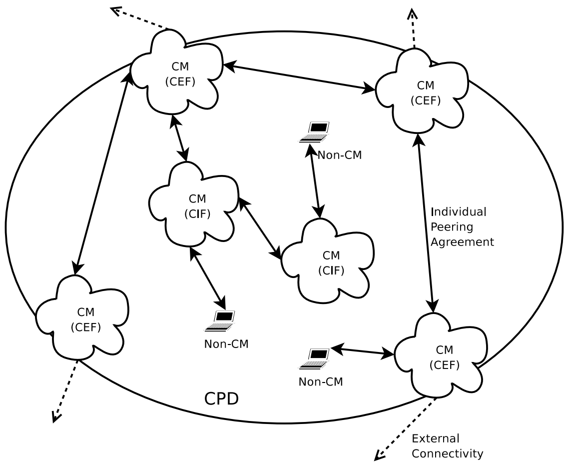

Figure 1: Coalition-Based Connectivity Architecture

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

Community 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 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 outisde the CPD. Thus outgoing traffic is distributed across many CEFs, potentially enabling a higher upstream data rate by pooling all the community uplinks. This approach is especially useful when the available capacity between CMs is greater than their individual uplink capacity.

Community 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) may be performed by using standard local-area or ad hoc routing 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 exmaple 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.

2.3 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 with respect to the amount of aggregate traffic that can be handled by the coalition. 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:

$$C_S = C_T - C_C \tag{1}$$

We assume that the CPD edge consists of M CEFs, each node, 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:

$$C_I = \sum_{m=1}^{M} C_{I_m} \tag{2}$$

$$C_E = \sum_{m=1}^{M} C_{E_m} \tag{3}$$

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

$$G_I = \frac{C_S}{C_I} \tag{4}$$

$$G_E = \frac{C_S}{C_E} \tag{5}$$

These gain factors may be used by CMs to assess the potential beneift of joining a CPD. They allow also some assement 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.

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 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' ADSL connection upstream speed 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 uses an 802.11b channel for intra-CPD data transmission, and that consists of 4 PDAs acting as CEFs, each with 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.

2.4 Why A Coalition?

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

The formation of local peering agreements and their grouping into a CPD may initially be thought of as a special case of an ad hoc network or even simply just a special case of a conventional IP edge network. Although elements of ad hoc architecture and conventional IP networking exist, there are significant differences with the coalition-based approach proposed here.

CPDs are not equivalent to autonomous systems: The IRTF RRG Ad hoc Network Systems Research Subgroup [2] describes an ad hoc network as "... an autonomous system of routers (and associated hosts) connected by wireless links — the union of which form an arbitrary graph". The key term 'autonomous system' (AS) implies a network under the control of a single administrative authority. However, a CPD does *not* represent an AS 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 complete local control over its own resources, whilst co-operating with other CMs. However, the CMs share some aspects of common policy which include some criteria by which they may form the coalition.

CPD formation involves trust establishment: 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, coalitions are organised at the human level. This may be through either personal meetings or other forms of out-of-band interaction. This implies a basic level of trust before local peering agreements can be reached, so a level of trust is in-built. Thus although formation may not be possible in a totally automated fashion, some levels of automation could be achieved through the application of policy on existing automation and discovery mechanisms.

CPDs are multi-homed edge entities: Although the topology within a CPD may resemble that of an ad hoc network, a CPD represents an individual *multi-homed* entity sitting at the edge of, and connected to, the Internet (or possibly another CPD). Traditional ad hoc network approaches have focussed on finding the most efficient route on a source-to-destination basis (where the destination may be either inside or outside the ad hoc network). This models ad hoc networks as an extension of a larger infrastructure, thus requiring them to either discover efficient routes to a very wide set of destinations, or to route towards a single gateway for the entire network (which then represents a single point of failure). However, the coalition-based approach 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 routing 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 on potentially resource-poor CIFs.

Thus the implementation of CPDs present challenges for routing, addressing and management of the connectivity within the mobile and wireless environment that are not tackled directly by either exisiting ad hoc routing mechanisms or exisiting IP routing mechanisms.

2.5 Architectural Issues

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). With a combination of the multi-homed CPDs and the spraying of

---

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

coalition-egress packets across the CPD's edge for onward routing, 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 routing 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 receiver to have sufficient knowledge to spray return traffic across the receiving 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 receiver 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.

As well as CPDs being multi-homed, there are additional issues to consider when functions such as firewalls and NATs are used. A 'distributed NAT' function may be required 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.

3 Tussle Spaces And Emerging Opportunities

The coalition-based approach is potentially highly disruptive in nature. Although it may be implemented at the edge by the end users, its effects are wider-reaching. We can identify a number of actors upon whom the adoption of such coalition-based connectivity would impact:

- End Users / Coalition Members
- Internet and Network Service Providers
- Equipment Manufacturers and Software Vendors

Initially it may appear that coalition members have the most to gain while others gain little or nothing. This leads to a number of tussles among the various actors involved, who have divergent interests [3]. However on further examination we see that these tussles may catalyse a number of new opportunities and new models for service provision. We present the various tussles that exist and outline the challenges and new opportunities that they may be transformed into.

3.1 Economics

3.1.1 Models Of Pricing

The current models of Internet and wide area connectivity require end users to subscribe to a specific provider and involve direct payment for connectivity. There are two models of subscription: metered — where there is a charge per unit time of connectivity or per unit of data transfered, and unmetered — where there is a monthly or annual flat-rate charge perhaps with some traffic capping.

The coalition-based connectivity approach breaks this model of connectivity access on three counts:

1. Coalition members' egress traffic is distributed via multiple CEFs, each of which may receive wide area connectivity through a different provider. For any given CEF, not only does some of its traffic bypass its own provider by traversing neighbouring CEFs' respective providers, but by doing so, the CEF potentially also gains a greater uplink capability than it has subscribed for with its own provider.
2. As coalition-egress traffic is distributed via multiple CEFs, providers find themselves in the situation of forwarding traffic that does not all originate from only CEFs subscribing to them, but originates instead from other non-subscribers.
3. CIFs gain wide area connectivity by sending their traffic via CEFs, and thus may benefit from wide-area connectivity via multiple providers without directly subscribing to any of them.

So, some CMs provide transit for other CMs on services that are sold for individual use, and all CMs benefit from higher capacity levels without individually subscribing for them. Thus with the current models of connectivity provision, service providers lose out instead of maximising sales by either attracting more customers from a particular community, or inducing existing ones to pay for higher capacity connectivity.

By maintaining a policy of disallowing onward connectivity sharing, providers risk losing customers to competitors who are willing to permit such practice. Yet although competitive fear may force providers to permit such practices, their comparitive overall service offerings remain similar and entering into direct price competition can prove to be very expensive for all parties.

Thus there appear to be few incentives for providers to encourage or support the formation of such community-oriented coalition-based networks. However, the existing tussles between the coalition-based connectivity approach and providers' pricing models could catalyse opportunities for new models of service provision with collateral benefits for providers. A number of non-price competitive strategies could be used by providers to influence multiple groups of customers and so increase market share.

For example, providers could offer connectivity services specifically targetted towards communities, encouraging coalition members to purchase particular quantities of products, and offering incentives in return guaranteeing specific levels of service for coalition traffic. This may create the added incentive for all coalition members to purchase from the same provider and/or upgrade in parallel their service provision from that provider, as it increases the capacity available for the coalition overall.

The coalition-based approach to connectivity may have also a detrimental financial effect on any CMs paying for wide area connectivity on a metered basis. By acting as a CEF, they would incur the added cost of forwarding peer CM traffic. The implications of this are complex because on the one hand their added costs may be offset by the benefit of receiving tit-for-tat forwarding by other CMs, but on the other hand such forwarding may not be balanced between CMs, leading some CMs to pay more than they otherwise would individually. To overcome this tussle between peering CMs, a model of payment could be applied locally allowing CMs to receive financial remuneration for any forwarding they provide (and pay for any forwarding they use from others). The deployment of such a model is a non-trivial task as it would require detailed logging, auditing and feedback mechanisms across the CPD placing additional burden on potentially resource-poor devices. This type of onward-selling may also conflict with the terms and conditions of some service providers. Another solution may be for providers to offer special coalition-oriented tariffs.

3.1.2 Customisation

The recent advances in mass production of wireless technologies have made wireless access points much more affordable for home users. With little effort, users may customise off-the-shelf equipment to extend its range of capabilities. The motivation for users to co-operate opens newly emerging equipment and its software to being reverse engineered to increase its flexibility and feature set (e.g. by attaching extension antennae or by loading custom firmware). In most cases modifications to hardware or software voids waranty. Wishing to minimise costs, users are likley to buy cheaper equipment and pay for fewer features as they will install new software to override factory defaults. Such practices reduce original equipment manufacturers' level of control and potential revenue from exisiting streams.

However, consumers tend to spend a minimal amount of time modifying equipment if a specific need does not arise. For example, a study into the development of wireless networking in London [4] ran an 'Air Stumbling' (as opposed to 'war driving') experiment from a light aircraft with "… a directional antenna, a GPS and a laptop running network discovery program Netstumbler". It showed that out of 1525 nodes seen, 50% were 'open' and "… approximately 40% of access points are running with the manufacturers factory default SSID settings". While not a definitive measure, the figures seem to indicate that a significant portion of node owners may be non-technical and have found it sufficient to leave factory settings unchanged. This shows evidence of a potentially expanding market for out-of-the-box products aimed at allowing non-technical customers to participate in community-oriented networking activities without needing to customise heavily their equipment.

By designing and manufacturing equipment that is flexible and simple to configure and to modify, manufacturers increase the likelyhood of product success and benefit from the greater revenue that that success brings with it.[3]

This principle applies equally for software vendors. By designing and engineering software that specifically allows non-technical consumers to benefit easily from customised usage within a coalition network scenario, the software product is likely to attract greater demand and produce greater revenue.

3.1.3 Changes In Traffic Patterns

As backbone operators sell capacity, it is in their interests to encourage the generation of more traffic to increase revenues. However, the coalition-based approach to connectivity at the edge shall draw some traffic (that which is localised to the coalition) away from the aggregated-level networks, confining it to the edge. This would affect backbone operators on two counts:

1. As less traffic is generated for aggregate-level network traversal, revenues may fall from reduced demand for capacity.

2. With less traffic traversing aggregate-level network links, under-utilisation may mean that previously incurred over provisioning costs take longer to recover.

---

[3] An example is the Linksys WRT54g series wireless router (http://www.linksys.com/products/product.asp?prid=508&scid=35), which quickly became very popular on its release. Not only was it easy to re-flash the firmware on it, but the procedure remained an open option without any attempts from the manufacturer to prevent it.

The benefits of coalition-local traffic remaining local means that ISPs and other access network operators may see less of the disruptive traffic that they ‘dislike’ (e.g. peer-to-peer). The use of a coalition-based connectivity approach at the edge of the Internet may thus provide new opportunities for methods of traffic control, of traffic shaping and better utilisation of available capacity within the backbone. These may catalyse new models of service provision and open up new revenue streams.

Additionally the benefits from use of coalition-based connectivity may result in increased sharing of resources in a locally distributed environment, something that has not been used widely in a community area environment. As applications advance over time to take advantage of such environments, users at the edge become accustomed to higher resourcing within their respective CPDs. This may fuel the development of more demanding types of bandwidth-intensive services extending outside of individual CPDs, leading to higher demands being placed on access networks and ultimately increasing backbone traffic again.

3.2 Trust

While discussing the changes in the Internet since its inception, Clark et al [3] state that “… users don’t trust each other. The users of the Internet no longer represent a single community with common motivation and shared trust.” The coalition-based approach reintroduces communities on a local scale and within them, members *must* trust each other to some degree, for without this a coalition cannot be formed. Coalition members have to open themselves up, trusting each other with their traffic and in the worst case potentially leaving themselves open to attack or abuse of resources.

The human-level aspects of the nature of local peering agreements between coalition members must be emphasised here. Local peering agreements are ultimately formed between *owners* of the nodes or the individual local networks that form a CPD. We propose that a non-automated, or not totally automated, process is involved during the formation of local peering agreements and their aggregation into CPDs thus reinforcing the cohesiveness of a CPD. The fundamental goal is to provide an implicit level of trust and security tailored to the needs and requirements of individual coalition members such that each coalition member is able to maintain complete local control of its own resources.

The formation of local peering agreements between two parties therefore implies a sufficient level of trust between them to reach an agreement in the first place. This level of trust may vary, directly related to the number of services provided across the peering agreement (a greater number of services implying greater trust). This provides a degree of implicit, basic trust and security throughout a CPD. Control is exercised through policy, defining conditions and levels of access to specific resources for peers; this may be propagated also transitively between peering agreements throughout the CPD. This is however non-trivial and requires dependance on an independant trust mechanism being in place, capable of Authentication, Authorisation and Accounting, including functions such as the validation of identity, control of local resources, membership and policy negotiation, auditing of activity, and the provision of feedback for trust evaluation.

Such a local comunity also provides an environment that stimulates the provision of local neighbourhood services for coalition members. Such services may extend beyond connectivity sharing and include storage, web, data repositories, entertainment services, instant messaging or communication services. These services may be exported also between coalitions. (e.g. a local coalition directory exported to provide information to remote-coalition members). This opens up the possibilities for inter-CPD peering.

Policy mechanims also need to be examined carefully. Some more mature groups are already beginning to establish simple policy-based approaches [1].[4]

There is also the obvious problem of CMs ‘sniffing’ on each others’ traffic as it transits their CIF or CEF. However, this security problem is not specific to the use of a CPD and measures that are already in existence could be used if this is seen as a real threat by the users.

4 Related Work

The ‘MAR commuter mobile access router’ [5] provides an architecture that is somewhat close to the approach proposed in this paper, in terms of connectivity aggregation. However a key difference is that MAR focusses 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. The coalition-based approach proposed in this paper focuses instead on a distributed wide-area connectivity model that allows arbitrary numbers of existing wide-

---

[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/)

area links to be aggregated yet allowing coalition members to maintain autonomy and local contol.

The 7DS Peer-to-Peer Information Dissemination and Resource Sharing system [6] 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 HDNet system [7] focusses on a highly dynamic multi-hop wireless network model in which clustering is used to allow higher powered ‘mobile base stations’ to forward data on behalf of lower powered ‘mobile hosts’. The relative mobility aspects introduced by the HDNet system may be mapped to long-lived connectivity scenarios involving mobile nodes within a CPD.

The DIRAC software-based wireless router [8] 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 routing 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.

5 Summary And Conclusion

The formation of community networks is a growing trend that has particularly been aided by recent advances in local-area wireless network technologies making them much more affordable. We have presented an architectural outline that would enable groups or communities of users to better utilise their wide-area connectivity and resources through collaboration via their local-area wireless capabilities. This is achieved by adding structure, the **coalition 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.)

We have examined a number of tussles that may arise as a result of such a potentially disruptive practice and we have shown that in each case, there are new opportunities and incentives for its adoption by all actors concerned. The proposed approach requires investigations into a number of existing research areas including addressing, routing and peering; trust and security policy; and performance and resource utilisation within fixed and mobile wireless environments.

Although we have not dedicated a separate discussion on the subject of security, we have highlighted the security-related concerns. Without a basic level of trust between peers, local peering agreements and thus coalition peering domains cannot be formed. Through a combination of the human-level involvement in coalition establishment, and the distribution of administrative responsibility across the coalition peering domain, varying degrees of trust and security are ensured by the autonomy of individual coalition members who maintain complete 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 connectivity resources in a controlled manner but there are a number of technical issues that should be researched further.

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

[3] David D. Clark, John Wroclawski, Karen R. Sollins, and Robert Braden. Tussle in cyberspace: defining tomorrow’s internet. In *SIGCOMM ’02: Proceedings of the 2002 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 347–356. ACM Press, 2002.

[4] Julian Priest. The State of Wireless London, 31 March 2004. http://informal.org.uk/people/julian/publications/the_state_of_wireless_london/.

[5] 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.

[6] 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.

[7] Ricardo Sanchez, Joseph Evans, and Gary Minden. Networking on the Battlefield: Challenges in Highly Dynamic Multi-hop Wireless Networks. In *IEEE Military Communications Conference*, Atlantic City, NJ, Oct. 1999.

[8] 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.