COALITION-BASED PEERING FOR FLEXIBLE CONNECTIVITY

Manish Lad [1], Saleem Bhatti [2], Steve Hailes [1] and Peter Kirstein [1]
[1] Department of Computer Science, University College London, Gower Street, London WC1E 6BT, UK
[2] School of Computer Science, University of St Andrews, St Andrews, Fife, KY16 9SX, UK
{m.lad, s.hailes, p.kirstein}@cs.ucl.ac.uk, saleem@dcs.st-andrews.ac.uk

ABSTRACT

Mobile devices available today provide users the ability to communicate using a number of different wireless network interfaces. However, these devices do not yet fully exploit the potential for multi-homed and multi-path communication allowing them to better utilise all the connectivity that is available to them. We present here the *Coalition Peering Domain (CPD)*, an architecture that supports collaborative networking relationships between mobile devices. This improves the speed and robustness of communication through more flexible use of all available connectivity.

I. INTRODUCTION

The nature of inter-network and inter-device connectivity is evolving. Not only are data rates increasing significantly, but so is the ease with which users may inter-connect multiple machines or devices. By doing this, they are able to utilise more efficiently and flexibly both their local resources and their access to the wide-area. Such flexibility has increased user expectations and fuelled the emergence of new and diverse types of high-bandwidth multimedia applications and services, including video and audio streaming, Internet telephony and multimedia conferencing.

A. *The Changing Face of Device Connectivity*

However, there still remains some disparity between the relatively high data rates that may be achieved within the local-area and the relatively lower data rates available within the wide-area. This is particularly accentuated within the mobile arena, putting a limit on the types of applications and services that users are able to access.

Many mobile devices now provide users with the choice of multiple types of wireless network connectivity including 3G, IEEE 802.11b/g and Bluetooth. However, these devices do not yet exploit fully all the connectivity that may be available to them. For example, a group of devices that support both 802.11b and 3G could use any untapped 802.11b connectivity between them as a 'backplane' for forwarding wide-area traffic through all the 3G links. By doing this, they would benefit from the statistical multiplexing gain of, effectively, aggregating their wide-area connectivity.

Indeed, generalising from this, devices could exploit any form of higher data rate, *local-area*, inter-device connectivity, to share their lower data rate, *wide-area* connectivity.

B. *Driving Forces: The "Killer Apps"*

With this in mind, our vision is to enable devices that have multiple network interfaces, to discover and peer with each other so as to better utilise the connectivity that is available collectively to them.

This has wide-reaching benefits for a diverse range of scenarios. For example, even though commuters on a train may be using their own individual 3G connections, it is unlikely that the nature of their usage and running applications make continuous and full use of these. This means that there will be periods when some portion of their connections are idle. By aggregating their wide-area connectivity, these commuters could tap-in to this unutilised capacity.

By far a more dynamic but resource-limited environment is one in which relief workers find themselves following a natural disaster. In such environments, relief workers have access to highly heterogeneous, frequently resource-limited, devices. Likewise, there are varying forms of connectivity that are partly determined by the devices' network interfaces and partly by which parts of the infrastructure remain operational. Added to this, such situations always have a time-critical element to them: the need to rescue survivors and to provide appropriate help to them. To achieve this, a range of information must be disseminated as quickly as possible. This includes, amongst other things, information about the potential hazards, the potential locations of survivors, the aid-items in short supply and the location of medical centres. The quicker such information is disseminated, the quicker aid can be deployed and the more effectively resources can be targeted. The more information that is disseminated using whatever networking capability happens to be available, the greater the opportunities for effective management of the situation. Conventional ad hoc networking mechanisms are inadequate because not all the parties to rescue efforts come from the same organisation or use a consistent set of networking technologies.

Our solution enables the dynamic formation of a ***Coalition Peering Domain (CPD)*** between devices for a common goal: to improve the speed and robustness of connectivity for all members of the CPD. By co-operating to form a CPD, the various aid agencies arriving at the scene of a disaster can pool their resources and make the most of all *potential* wide-area connectivity that is available collectively to them. They gain, additionally, an improved level of robustness in their communication because, even if some of their wide-area connections were to fail, CPD formation would ensure that all CPD members continue to make the most of the collective connectivity.

The viewpoint presented here proposes the Coalition Peering Domain as a network architecture solution to support collaborative networking relationships between mobile devices.

In the next section, we present the case for coalition-based peering. We outline the challenges in realising the CPD architecture and related work in this area. In Section III. we present part of our ongoing work into coalition-based peering. We detail a protocol to enable devices to negotiate and manage the local peering agreements that form a CPD. We then present in Section IV. our work towards validating the mechanisms for forming a CPD. Finally, we conclude in Section V. and outline future work.

II. The Case for Coalition-Based Peering

We present here our architecture to enable devices to peer with each other while allowing administrative responsibility to remain distributed. This ensures that users retain local control of their own resources.

A. An Outline of the Coalition Peering Domain

Fig. 1 illustrates a number of collaborative relationships or 'local peering agreements' between pairs of devices within proximity of each other. These peerings may be either as simple as links interconnecting different pairs of devices, or more complicated associations controlled through policy defined locally by the device owners. As the numbers of such local peering agreements begin to increase and form connected topologies between different devices, we refer to the creation of a *coalition* within the community of devices and the formation of a *Coalition Peering Domain (CPD)*.

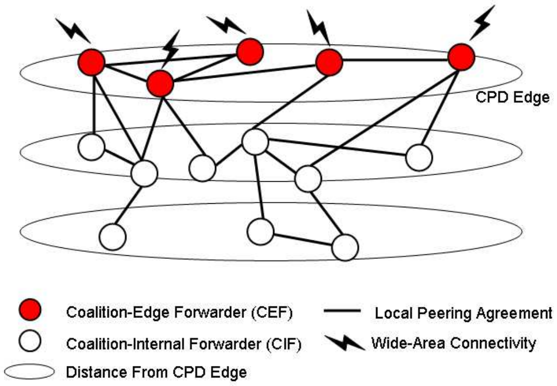

Figure 1: CPD Architecture

Each *Coalition Member (CM)* may represent an individual device or a local-area/personal-area network. Local peering agreements are negotiated and maintained directly between CMs. They enable each CM to declare the resources that they are willing to offer, to accept or reject the resources that are being offered by potential peer-CMs, and to amend any of the peering parameters after the local peering agreement has been formed.

Coalition members who have wide-area connectivity (or more generically, connectivity outside the CPD) collectively form the edge of the CPD and act as *Coalition-Edge Forwarder (CEF)* nodes; they are the CPD ingress–egress points, allocating some proportion of their external connectivity for this purpose. Together, their external connectivity provides the CPD with a higher *potential* data rate to the wide-area than any individual is able to achieve. To efficiently utilise this for CPD-egress traffic, CEFs forward some proportion of outgoing packets on their own CPD-egress links, but forward the remaining egress traffic by 'spraying' (distributing) it across the 'CPD edge', via their local peering agreement links to neighbouring CEFs. The neighbouring CEFs do the same. Such distribution of outgoing traffic across multiple CEFs enables the aggregation of CEF egress links, providing higher CPD-egress data rates and robust connectivity through multiple connections.

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)*. CIFs contribute to the CPD by forwarding traffic to other CMs, including to/from CEFs. They need simply to forward CPD-outbound traffic by directing it towards their 'nearest' CEF for CPD egress. 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.

B. The Challenges facing Coalition-Based Peering

The formation of a CPD poses a number of technical challenges related largely to addressing and routing but also impacting the operation of upper layer protocols. We list here briefly the main challenges that we are investigating currently. There are also a number of other challenges (including security, authentication and incentive models for co-operation [7, 10], as well as policy definition, service models from providers, regulation, and effects on higher layer protocols), but we leave these for another discussion.

The CPD is a *collection* of collaborating systems. Administrative responsibility is *distributed* across all CMs. Thus, a CPD does 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. Existing mechanisms for addressing and routing are not optimal in this context because they have been designed for environments in which administrative responsibility is hierarchic in nature. Addressing and routing among the collaborative group of peering systems needs to be managed in a distributed manner.

The CPD is a *multi-homed virtual* edge entity. It resides at the edge, but remains connected to the Internet. As all the available wide-area connections are utilised, traditional intra-domain ad hoc routing mechanisms [2] are not really appropriate in this context. They model the domain or ad hoc network as a single AD and focus 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).

The CPD requires *new* control plane protocols. Existing BGP mechanisms [8] are too heavy-weight to be employed within the mobile and resource-poor scenarios described earlier. Also, allocating Autonomous System (AS) numbers to each CM would lead very quickly to an explosion in the lengths of autonomous system paths and the numbers of routing table entries. The impact of such back-pressure onto the core network would be made much worse by even a handful of misconfigurations [5], let alone the potentially transient nature of CMs. This would lead to greater instability throughout the infrastructure.

C. Related work

Previous work aimed at better utilising connectivity and resources focuses on traditional models of resource sharing and discovery. It assumes openness among peers, treats them as a single edge network, and provides only single-path source-to-destination routing. The possibility of utilising multiple connections simultaneously is not taken into account. A good example is the 7DS Peer-to-Peer Information Dissemination and Resource Sharing system [6]. 7DS does provide also some load balancing mechanisms, but these are based on the selection of single (least loaded) gateways rather than distribution across multiple CEFs as undertaken by the coalition-based approach.

Although the 'MAR commuter mobile access router' [9] does aggregate multiple wide-area wireless interfaces, it focuses on a *hotspot* model (albeit multi-homed) of access. All local users gain wide-area access via a single provider (i.e., a MAR device placed in a moving vehicle), representing a single point of failure.

The CUWiN project [1] moves substantially towards easing the effort required to establish oneself as a member of a community-area network. However, its software requires higher-powered machines. This limits the range of possible deployment scenarios, excluding any involving the types of lower-powered mobile devices described earlier. Although CUWiN simplifies and automates all technical configuration, this automation has a price: a lack of local control.

III. Operation of a Coalition Peering Domain

We present here a summary of our work so far, focusing on the management of devices joining and leaving a CPD.

A. Solving the Challenges of Addressing and Routing

CMs need a mechanism to communicate with each other both before and after the formation of local peering agreements. While the latter implies that some form of addressing scheme should be employed across a CPD, the former holds the stronger requirement that an addressing scheme is employed already across all potential CPD members. 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 Network Address Translation (NAT) devices may provide a solution for community-area networks, they pose a number of problems that may limit the overall usefulness of a CPD. There is a chance that the arbitrary use of private addressing may lead to clashes between peering CMs. Moreover, they limit the operation of some types of applications, they limit the ability to apply security at the IP layer, and ultimately they introduce unnecessary complexity in configuration and maintenance.

We present here a solution that uses a combination of IPv6 link-local addressing, IPv6 multicasting and globally reachable IPv6 address block allocations.

1) Negotiation using Link-Local Addressing

We assume that all CMs listen on a "well known" IPv6 multicast address and port on their local-area network interfaces. A coalition member CM_A that wishes to form a local peering agreement with a neighbour, transmits a *CPD Peering Request (CPD_PREQ)* to the well known address. This CPD_PREQ is sourced by the IPv6 link-local address of CM_A, and contains the resource parameters that CM_A is willing to offer for peering. This initiates a standard three-way handshake between CM_A and any recipient CM_B wishing to agree to the formation of a local peering agreement. Any CM_B not willing to agree to a local peering agreement formation may either transmit a *CPD Peering Response (CPD_PRESP)* rejecting the request, or simply ignore the CPD_PREQ and allow CM_A to timeout awaiting response.

A coalition member CM_B that wishes to agree to the formation of a local peering agreement then transmits back a CPD_PRESP accepting the request. This response message contains a proposed CPD_ID, an identifier for the CPD within which the local peering agreement is to operate. This CPD_ID may be either a newly generated identifier (thus forming a new CPD), or the identifier for a CPD within which CM_B already is a member (thus adding the new local peering agreement to the existing CPD). The CPD_PRESP contains also the resource parameters that CM_B is willing to offer in return. On receiving a CPD_PRESP from CM_B, CM_A transmits a *CPD Advertisement (CPD_ADV)* message if it wishes to continue a the local peering agreement formation with CM_B. Alternatively, it may either respond with a CPD_PRESP that rejects the local peering agreement formation, or simply ignore the CPD_PRESP from CM_B and allow CM_B to timeout awaiting an advertisement.

2) Address Allocation ***within*** the CPD

We assume further that all CEFs have in addition an IPv6 globally routed address block allocation. Each CEF then splits this address block into a number of sub-allocations. The granularity of this split is determined by the CEF's owner or local policy. A CEF CEF_I then assigns one sub-allocation to each of the CIFs CIF_j with which it has a local peering agreement. All CIFs CIF_j appear externally as sub-networks of CEF_I, so the routing of traffic to CEFs takes place through standard IPv6 routing. This behaves well also with reverse path or CPD-ingress traffic, which is forwarded to the CEF that owns the address block and can then be propagated down to the recipient CIF.

Depending on the granularity of the address block split, the process of splitting is repeated by CIFs. Thus a CIF CIF_J splits its sub-allocation into further sub-allocactions

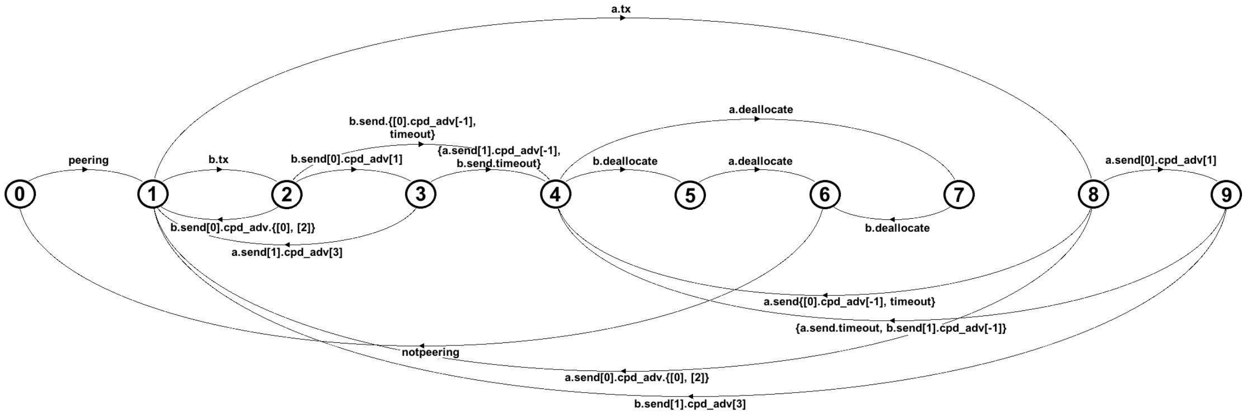

Figure 2: Local Peering Agreement State Machine

and assigns them to each of the CIFs CIF_k with which it has a local peering agreement. All CIFs CIF_k appear to CEF_I as sub-networks of CIF_J, and appear esternally as sub-networks of CEF_I.

3) Signalling to Sustain Local Peering Agreements

Once a local peering agreement has been established between two CMs CM_A and CM_B, it must be sustained through regular advertisements. Each peering CM transmits to the other, at regular intervals, a CPD_ADV. If at any time CM_A wishes to terminate the local peering agreement, it may either transmit a CPD_ADV that contains a 'Terminate' flag to end the local peering agreement, or simply stop transmitting and allow CM_B to timeout awaiting further advertisements.

To simply sustain its local peering agreement in its present form, CM_A transmits to CM_B a CPD_ADV containing a 'NoChange' flag and, after transmission, the CM_A returns to its default peering state.

If CM_A wishes to change its resource parameters, it transmits a CPD_ADV containing the new parameters and a 'ChangeParam' flag, and again returns to its default peering state. If CM_B does not wish to sustain a local peering agreement with the new parameters being offered by CM_A, then it may either transmit a CPD_ADV that contains a 'Terminate' flag to end the local peering agreement, or simply stop transmitting and allow CM_A to timeout awaiting further advertisements.

If CM_A wishes to change state from being a CEF to a CIF, this affects any originally negotiated local peering agreement conditions between CM_A and CM_B because CM_A no longer provides direct CPD-egress capacity and CPD-edge spraying for CM_B. If CM_A wishes to change state from being a CIF to a CEF, this also may affect any originally negotiated local peering agreement conditions between CM_A and CM_B because CM_A now forms part of the CPD-edge and has direct spraying capability. Therefore, in the case that CM_A changes state, CM_A transmits a CPD_ADV containing a 'ChangeState' flag and any new resource parameters that it is willing to offer. It then enters an intermediate state where it must await confirmation from CM_B. If CM_B wishes to sustain the local peering agreement under the new conditions, it transmits a CPD_ADV containing a 'ChangeConf' confirmation flag to CM_A. If CM_B does not wish to sustain the local peering agreement under the new conditions then it may either transmit a CPD_ADV that contains a 'Terminate' flag to end the local peering agreement, or simply ignore the new state and parameters from CM_A and allow CM_A to timeout awaiting a confirmation.

4) Routing traffic through the CPD:

To route traffic to the edge, CIFs need only use their address block allocation and route up towards a more general prefix.

CEFs operate a link state routing protocol between themselves (across the CPD edge), using IPv6 link-local addresses as next-hop values. CEFs that receive CPD-inbound traffic destined for a CM that is not within their own addressing hierarchy, route that traffic appropriately via their local peering agreement links based on the link-state information that they have received from neighbouring CEFs.

B. Further Considerations for 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 subject to a proportionally greater volume of return path CPD-ingress traffic. However, the asymmetry of most wide-area connectivity technologies may be sufficient to offset this inequality in the short term. Alternatively, the function of reverse spraying may be placed on either the remote party, or a provider-controlled device located beyond the CPD edge. Multi-path routing does, however, have implications for delay-intolerant applications and for higher layer protocols that assume certain behaviour form the underlying routing.

De Couto *et al* [3] 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, may not be an efficient mechanism for the routing of CPD-outbound traffic when the CPD is 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.

IV. Validating the CPD State Machine

We present here an outline of our work so far to validate the correctness of the states through which CMs transition.

A. Modelling a Local Peering Agreement

To validate the CPD local peering agreement processes that we have presented here, we have modelled the negotiation and signalling mechanisms with Finite State Process (FSP) algebra using the Label Transition System Analyser (LTSA) [4]. Fig. 2 illustrates the states through which CMs CM_A and CM_B transition while they sustain a local peering agreement between each other. For the time being, we have assumed that there is no packet loss between CM_A and CM_B, and for simplicity, that each CM operates a single threading process where it either transmits (in which case the other listens) or vice versa.

Our model for negotiation and signalling has been checked for the absence of deadlocks and liveness progress violations. This demonstrates that for the simplified case where no packet loss is assumed, the CPD negotiation and signalling described here provide a valid mechanism.

B. Future Validation

The most important next step for future validation is to remove the initial assumption that there is no packet loss. As one would expect, this shall increase the number of alternative transitions that CMs may take at each state. Although the complexity shall inevitably be greater, we expect that the increase in complexity shall be confined to the number of transitions, and not in an explosion of states themselves.

Moving forward from modelling individual local peering agreements, the next step will be to model multiple local peering agreements from both an overall CPD perspective, and also from the perspective of an individual CM that is a member of multiple CPDs simultaneously. In parallel, we are pursuing a prototype implementation of the CPD.

V. Conclusions and Future Work

We have presented the Coalition Peering Domain, an architecture that supports collaborative networking relationships between mobile devices. This architecture allows users to make the most of all connectivity that is available collectively to them by pooling the heterogeneous types of connections that are provided by their devices.

We have outlined a number of challenges towards realising the CPD architecture and we have highlighted why existing mechanisms for addressing and routing are not adequate within this context.

Our solution defines a protocol to enable the dynamic formation of a CPD. It allows individual members to negotiate and maintain the collaborative relationships that form the CPD. Coalition Members negotiate an initial set of local peering agreement resource parameters. They then retain local control by using local peering agreement signalling to vary these parameters as required.

Our approach provides an improved flexibility, speed and robustness of connectivity that is valuable for a number of scenarios, particularly those in which wide-area connectivity may be limited.

There remain a number of open issues to be investigated further. In addition to the challenges of addressing and routing presented here, our work on STRUDEL [7] tackles the "tragedy of the commons" dilemma within a CPD to isolate malicious coalition members and, thus, minimise connectivity disruption.

VI. Acknowledgements

The authors would like to acknowledge support from the IST-RUNES project, and to thank Daniele Quercia and Adam Greenhalgh for their valuable feedback on this work.

References

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

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

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

[4] Jeff Magee and Jeff Kramer. *Concurrency. State Models and Java Programs.* Wiley, 1999.

[5] 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.

[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] D. Quercia, M. Lad, S. Hailes, L. Capra, and S. Bhatti. STRUDEL: Supporting Trust in the Dynamic Establishment of peering coaLitions. In *Proceedings of The 21st Annual ACM Symposium on Applied Computing (SAC 2006)*, 23-27 April 2006.

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

[9] 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.

[10] D. Zhu and M. W. Mutka. Promoting Cooperation Among Strangers to Access Internet Services from an Ad Hoc Network. In *Proceedings of the Second IEEE International Conference on Pervasive Computing and Communications (PerCom 2004)*. IEEE Computer Society, 14-17 March 2004.