Incentives For Community-Oriented Distributed Connectivity

Manish Lad *M.Lad@cs.ucl.ac.uk*

Saleem Bhatti *S.Bhatti@cs.ucl.ac.uk*

Peter Kirstein *P.Kirstein@cs.ucl.ac.uk*

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

Keywords: community networks, mobile networking, ad-hoc networking, routing

Current availability and choices of connectivity make possible pervasive networking: "everywhere, for everyone, at any time and for a wide range of applications". However, to allow users to exploit fully the available connectivity, users should collaborate to share resources and so make use of the statistical multiplexing gain in shared scenarios. We take as an example a **community network**, a neighbourhood of users who are close to each other geographically, have wireless connectivity (e.g. using IEEE 802.11 standards) and have permanent connectivity to the Internet (e.g. using DSL or cable-modem). Specifically, we examine incentives for cooperation for the users and the impact on service providers. There is a win-win situation: users sharing connectivity creates new opportunities for service provisioning.

Sharing: control of service usage

The shift in the model of connectivity access takes control away from service providers and places it in the hands of consumers. Relatively cheap wireless access routers coupled with extension antennae, have allowed some consumers to connect several personal machines and create networks on an ad-hoc basis to form local neighbourhood communities. 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 – a **local peering domain (LPD)**. This model of peering can be referred to in general terms as a peer-to-peer architecture. However, as the number of such peering arrangements increases, and some of the LPDs begin to overlap, we can talk of a **coalition** within the community and the formation of a **coalition peering domain (CPD)**. In Figure 1, the LPD is between individual nodes; between two base-stations, a base-station and a mobile unit or two mobile units. Thus, multiple sets of peering agreements join to form a single overall coalition or community. Such coalitions may be formed upon specific premises (for example basic peering, traffic forwarding, resource provision, resource sharing/pooling etc.).

Developing this further, consider a local neighbourhood community pooling its resources together, including use of the wired Internet connectivity owned by a subset of community members (shown on Figure 1). This pool can be shared such that some or all community members may benefit from higher data-rate Internet connectivity than their respective individual connections. All community nodes that have wired access are said to reside at the edge of the CPD and act as **edge coalition forwarders (ECF).** ECFs distribute outgoing packets, sending some out through their broadband interfaces and 'spraying' (distributing) some to other ECFs within their LPD. Thus outgoing traffic is distributed among all coalition edge nodes, potentially enabling higher upload speed by pooling all the community 'uplinks'. This is especially useful where available capacity between LPDs is greater than the capacity of individual uplinks at the ECF.

Taking this further, the coalition-internal nodes (such as the mobile units in Figure 1) may act as **internal coalition forwarders (ICFs)**, forwarding coalition-outbound traffic toward their nearest ECF for coalition egress. This traffic could be sprayed across peer ECFs by the receiving ECF as described above; the ICFs may also use a mechanism for load balancing and take responsibility for spraying directly to multiple ECFs. In Figure 1, mobile units are ICFs; hosts collocated with base-stations are **edge host (EH)** systems and are owned by the same person.

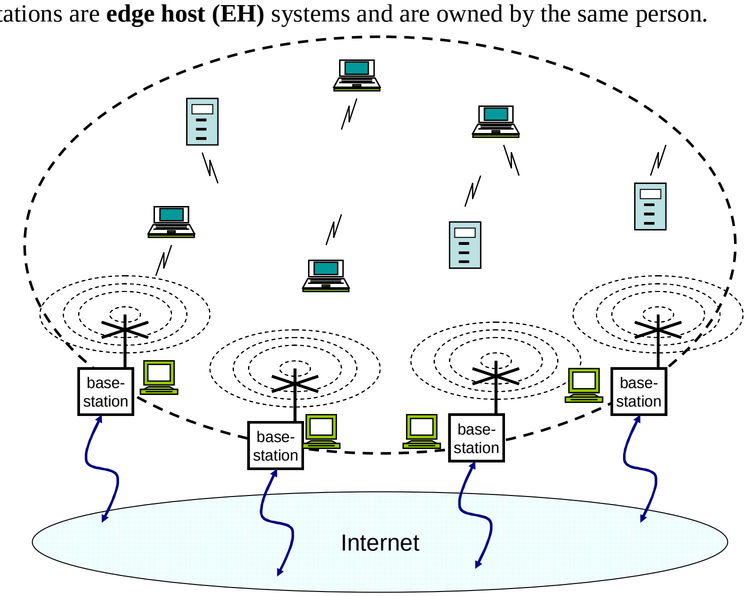

Figure 1

Addressing, routing and architecture

This does however lead to a number of issues for addressing, routing and for the operation of transport layer protocols that rely upon the network layer IP address as part of the transport protocol state (such as TCP).

If NATs are used by ECFs (for example), the packet spraying mechanism could potentially lead to the violation of end-to-end connectivity, as the receiver will receive packets belonging to a single stream but with inconsistent paths back to the source. One solution for this would be to use the packet spraying method (whether by ECFs or load balancing via ICFs) solely to increase community-egress bandwidth. This would be possible through loose source routing mechanisms used by the ICFs, but we intend to investigate further. Attempting to increase community-ingress bandwidth would require the receiver to have sufficient knowledge to spray return traffic across the receiving community's edge nodes.

Provisioning used solely to increase coalition-egress bandwidth is more useful for community members who may have asymmetric outgoing connectivity with lower uplink speeds (e.g. ADSL and cable-modem connectivity). However the receiver's responses would travel along a single path back to the source via a single ECF, which may congest a single coalition-ingress connection (e.g. a single ADSL line). This would be offset through load balancing and the natural asymmetry of the connectivity.

Incentives for service providers

Such scenarios would not benefit ISPs who are providing specific members of the community with Internet access. Instead of achieving maximum sales by selling their broadband access services to all the members of a community, the members are instead forming coalitions and pooling together to make the most of a single or a very few connections. However this could in fact open up opportunities for new models of service provision. ISPs could target certain connectivity services specifically with communities in mind, encouraging coalition members to purchase particular quantities of products, providing in return incentives guaranteeing specific levels of service for coalition-egress traffic.

In addition, ISPs may be able to provide some localised (coalition area) access point capable of handling coalition-ingress packet spraying. This may have to rely upon all coalition members using the same ISP, which may be an undesirable dependency for users. However, it could be used as a selling point for ISPs in order to market their services to local area communities.

Additionally, with coalition based connectivity, there is an incentive for all coalition members to upgrade their service provision from an ISP in parallel, as it increases the capacity available for the coalition overall.

Many community application architectures are delay-tolerant, and the users often form ad-hoc communities with other such users. It is these types of architectures and applications that would benefit most from the above scenarios. Examples include file transfer, file sharing and downloading (especially through peer-to-peer protocols), software and systems updates; information services (news, stock quotes, weather, etc.)

Wider applicability of coalitions

Although we have discussed the packet spraying mechanism within a fixed-node community environment with Internet and wireless communications, it is also applicable more widely to any scenario where the local connectivity speed between a number of peers is greater than or equal to their individual connectivity speeds to a common remote entity (e.g. the Internet), *and* their individual connectivity is not fully loaded.

For example two Bluetooth-enabled mobile phones will have a local connectivity speed between them of 768Kb/s, and individual connectivity via GPRS of 34Kb/s. On the assumption that their individual GPRS connections are not running at 100% utilisation, by pooling together they will each be able to achieve a connection speed to a common entity reachable via GPRS that is greater than their individual GPRS connection speeds of 34Kb/s. Further it may be possible to introduce a third node, a Bluetooth-enabled PDA perhaps, with a local connectivity speed of 768Kb/s to one or both of the above two mobile phones, but without GPRS. The PDA may connect to one or both mobile phones via Bluetooth and benefit from onward connectivity via GPRS that is greater than a single GPRS connection. Adding another degree of complexity to this example, each of the devices may not have the same remote connectivity speed (e.g. one mobile phone may be connected via GPRS and a laptop may be connected via a 56Kb/s modem).

Further work

The mechanism for spraying of packets to increase coalition-egress connection speed will also be attractive to mobile and wireless environments. The downside is that this will raise a number of issues about billing, specifically for non-flat-rate tariffs. There may be dissimilar incentives for people with dissimilar costs to share their connectivity. An important question for the assessment of the system is, “At what point does the addition of ICF nodes bring a collapse of the overall gain in coalition capacity?”

Acknowledgement

The initial concept of ideas on coalition-based networking have arisen from the work of Defence Research & Development Canada (DRDC) on coalition-based VPN infrastructures, focussing upon absolute local control on each coalition member site. Various discussions and deployment activities for this system were undertaken at UCL during 2003-2004.

Vitae

Manish Lad is a Research Fellow in the Department of Computer Science, University College London (UCL). He is also studying for a PhD within the department. He has five years experience in Data Communication Networking and Distributed Systems, including IPv6 protocol stack development for a core Internet router at Lucent Technologies Inc., and involvement in two pan-European collaborative projects: 6NET IPv6 pilot network deployment and SEINIT Security Expert Initiative. His work has involved deployment of dynamic IPv6-enabled overlay and VPN systems, and contributions to the design of a trusted security framework capable of working across heterogeneous devices and networks.

Saleem N. Bhatti is currently a Senior Lecturer in the Department of Computer Science, University College London (UCL). Since 1991 he has been part of the Research Staff at UCL, involved with several pan-European collaborative projects. The project partners have included telcos, equipment manufacturers, CATV companies, and academic institutions from all over Europe. His work has involved aspects of multi-service networking, tele-working, multicast, network and systems management, network security, IPv6, the consideration of QoS (Quality of Service) adaptability support for Internet applications and adaptive systems, and Grid networking & Grid computing.

Peter T. Kirstein is a Professor in the Department of Computer Science, University College London, and was appointed as head of the department from 1980 until 1994. As a director of Research and since 1970 he has lead multiple research groups on projects concerned with aspects of computer networks. He is a Fellow in a number of professional bodies such as the Royal Academy of Engineering, British Computer Society, Institute of Physics, Institution of Electrical Engineers, Senior Member Institution of Electrical and Electronic Engineers, Governor International Computer and Communications Committee. Prof. Kirstein was awarded the ACM SIGCOMM award and the IEE senior award - both in 1999. He was made an Honorary Foreign Member of the American Academy of Arts and Sciences in 2002 and made a Distinguished Fellow of the British Computer Society in 2003. He was awarded a CBE in the Birthday Honours list and received the Internet Society's Postel Award, both in 2003.