RaDP2P

Towards Generic and Extensible Support for Adaptation in P2P Systems

Daniel Hughes, Geoff Coulson, Ian Warren
Computing Department, Lancaster University, Lancaster, UK
d.r.hughes@lancaster.ac.uk | *geoff@comp.lancs.ac.uk* | *ian-w@cs.aukland.ac.nz*

Abstract

Peer-to-peer environments are highly heterogeneous and are likely to become more so due to the proliferation of mobile Internet access technologies and the development of novel peer-to-peer applications. Adaptation is essential in such heterogeneous systems in order to exploit the resources available on diverse nodes and tailor network services to meet the needs of diverse applications. We argue that existing support for resource awareness and adaptation in peer-to-peer systems is inadequate and we propose a framework for developing peer-to-peer systems with generic and extensible support for resource awareness and adaptation. A hybrid peer-to-peer model similar to Structella is used in this framework as a novel mechanism for supporting resource awareness and adaptation.

1. Introduction

P2P environments are typically highly heterogeneous. The nodes which compose them have very different capabilities and requirements; they may range from mobile nodes with slow CPUs, small amounts of memory and low-bandwidth mobile connections to modern high-specification workstations connected to the Internet via broadband or LAN connections.

Just as the requirements of the nodes which compose a P2P network are diverse, so are the requirements of the applications which run on these networks. Popular P2P applications such as file sharing [1], Internet telephony [2] and distributed computation [3] each have specific and different requirements of the underlying P2P network.

By adapting the role that each node plays in the network to better exploit its capabilities, we can maximize the contribution that each node makes to the system as a whole. Similarly, by adapting the behaviour of the network we can maximize its suitability for supporting different node and application classes. In this way, adaptation is key to the efficient operation of P2P systems in heterogeneous environments; particularly those seeking to support different application classes on the same network.

We argue that the need for adaptation will only increase, as will the heterogeneity of the nodes which compose P2P networks. While the processing power and bandwidth available to *typical* home PCs around the edge of the network is continually increasing, so too is the use of mobile devices such as PDAs and smart-phones, which have more restricted resources. We argue that, in order to provide viable services for such diverse nodes and efficiently exploit the resources they offer, P2P networks need to become more adaptive.

2. Existing Support for Adaptation

We classify the adaptation demonstrated in P2P systems into three discrete levels; network restructuring adaptation, routing behaviour adaptation and service selection adaptation.

1. **Network restructuring adaptation** adapts the relative position of nodes on the network through selective (re)connection. For example: in a peer-to-peer resource sharing network, a node may wish to modify its position in the network so that it is closer to the content it is seeking. This may be accomplished by

reconnecting to the network in a position which is closer to such resources. Examples of network restructuring include GIA [4], which uses network restructuring for content placement, Brocade [5] which uses network restructuring to improve the scalability of structured overlay networks and super-node schemes [6] which use network restructuring to improve the scalability of unstructured overlays. Network restructuring has also been suggested as an incentive scheme in the recent paper 'From Selfish Nodes to Cooperative Networks – Emergent Link-based incentives in P2P Networks' [7].

2. **Routing behaviour adaptation** adapts the routing behaviour of nodes on the network. For example, if the load of a neighbour peer is known, a node may choose to route more, or less messages to that peer based upon this information. Examples of routing behaviour adaptation include AGnuS [8], which uses routing behaviour adaptation to support content based routing and load balancing on the Gnutella [1] network.

3. **Service selection adaptation** adapts which service a peer selects following the resource discovery phase. For example, a node may discover several peers offering a service they desire. Meta information provided about these peers may be used to inform the decision about which service to select. Examples of service selection adaptation include Gnutella [1] wherein service selection adaptation is used to select the fastest available peer for file-transfer and "IHD - the interactive help-desk system" [9], where adaptation is used to select a user with the most appropriate knowledge of a given problem.

Despite these examples of adaptation in P2P systems, we argue that current P2P systems are not adaptive *enough*. These systems are all limited in the scope of adaptation they allow; typically engaging in only one class of adaptation behaviour (network restructuring, routing behaviour or service selection adaptation) and in each case the policy used to inform adaptation is fixed. Furthermore, the resource awareness that existing systems offer is not extensible; restricting the factors that can be used to inform adaptation. For example: AGnuS implements load balancing using routing behaviour adaptation informed by awareness of the load on neighbour peers, it does not support routing behaviour adaptation based upon any other.

We argue that a generic framework for building adaptive P2P systems is required and that it must provide:

- Support for each class of adaptation.
- Support for multiple adaptation policies.
- Extensible support for resource awareness.

3. Our Approach

Research has shown the potential of both structured and unstructured decentralized P2P networks for supporting specific application domains. Recent research [10] has also shown the potential of using hybridized schemes (discussed further in section 3.1).

We suggest that such hybrid schemes can be used as a powerful tool for supporting adaptation. Using such a scheme, we are currently implementing a framework for developing reflective and dynamic peer-to-peer networks (RaDP2P), capable of supporting the requirements outlined in section 2.

3.1 Network Architecture

Walkerdine et al. [11] isolate two categories of decentralized P2P networks; those which demonstrate structured communication and those which demonstrate unstructured communication.

Structured decentralized networks include examples such as Pastry [12] and CAN [13]. These two networks provide efficient, reliable routing of messages and are highly scalable; however, resource discovery in such networks is problematic [11] as they do not provide support for complex queries. This is a shortcoming typical of structured decentralized overlays.

Unstructured decentralized networks include Gnutella [1] and Fasttrack [14]. These networks are well suited to ad-hoc resource sharing as they provide inherent support for complex queries;

however, these networks are less scalable and reliable than structured overlays.

Both structured and unstructured decentralized overlays have been successfully deployed in P2P systems. Gnutella [1] and Fasttrack [14] are used to support several popular file-sharing applications, while systems such as Pastry [12] have been used in archival storage systems such as PAST [12] and group messaging protocols such as SCRIBE [15]; both architectures have advantages in specific application domains.

Recent research has studied the benefits of hybridized systems wherein unstructured decentralized networks are layered upon structured P2P substrates to provide support for resource discovery through complex queries. Structella [10] is one such hybridized network. It layers Gnutella [1] over the Pastry [12] routing substrate and demonstrates the advantages of such schemes.

Recent projects such as "Open Overlays" [16] and "PROST – A Programmable Structured P2P Overlay Network" [17] take hybridization further; abstracting overlays in such as way that many different technologies can be used together. For example: the Open Overlays abstraction model allows any unstructured network to be layered on top of any structured overlay; providing efficient resource discovery upon a scalable structured routing substrate.

We believe that such hybrid schemes can be used to support multiple levels of adaptation. As with Structella [10], RaDP2P uses a structured decentralized network to provide an efficient routing substrate upon which an unstructured decentralized layer is overlaid to support complex queries. In contrast to Structella, we intend to exploit the inherent structure of this routing substrate. By implementing our resource discovery network over a structured overlay, rather than directly over TCP/IP, it is possible to perform fine-grained adaptation on multiple levels. Case studies exploring this further are presented in section 4.

3.2 Supporting Adaptation

This section discusses how RaDP2P provides support for each level of adaptation described in section 2. Key allocation in RaDP2P differs from the mechanisms used in most structured overlays. The value assigned to RaDP2P keys is used to *reflect* information about each node. This information is used to support adaptation.

Network restructuring adaptation support: Reflective keys are used to support network restructuring adaptation. Meta-information harvested from each node is used to generate the most significant bits of each node's key. In RaDP2P, the structured routing substrate is ordered by key value. As the most significant bits of the key are derived from meta-information, the network restructuring policy defines each nodes position in the overlay, collocating nodes with similar meta-information. This approach is far more fine-grained in nature than traditional network restructuring approaches, such as used that used in GIA [4] as RaDP2P makes it is possible to precisely define each node's position in the overlay.

Routing behaviour adaptation support: Reflective keys also support routing behaviour adaptation. Meta-information harvested from each node is used to generate the least significant bits of each node's key. In this case the goal is not to alter the relative position of each node in the overlay network; rather it is used simply to mark nodes for differential treatment by their peers. Based upon the value of a node's routing tag, its neighbours will modify the manner in which they route messages to it. Encoding this information in the nodes key is advantageous as it avoids slow query-response interrogation such as that used in AGnuS [8].

Service-selection adaptation support: As with existing P2P systems; service selection adaptation occurs directly between peers following the resource discovery phase. Unlike existing systems, RaDP2P offers extensible support for resource awareness, allowing service-selection adaptation to be informed by any factor. This is in contrast to the resource awareness information offered in systems such as Gnutella [1] or The Interactive Help Desk [9], where the factors used to inform service selection are fixed.

Network restructuring and routing behaviour adaptation are controlled by global adaptation policies (the implementation of which is described on the following page). The use of global policies allows developers to precisely dictate the conditions that should cause nodes to adapt and the form that adaptation should take.

4. Design

The RaDP2P architecture is separated into three primary concerns; awareness and adaptation, network abstraction and applications.

The **awareness and adaptation** sub-system is responsible for the adaptation behaviour of each node, which is defined by a global adaptation policy and informed by extensible monitoring components as described in section 3.2.

A *Network restructuring policy* defines a monitoring component (which harvests meta-data used to form the most significant bits of the key), a period of adaptation (the interval at which the monitoring component should be inspected and the key regenerated), and how this meta-data should be used in key manufacture. In this way, a network restructuring policy defines the desired structure of the network and the level of dynamicity.

A *Routing adaptation policy* defines a monitoring component (which harvests meta-data used to form the least significant bits of the key), the action to be taken, (e.g. lower the volume of messages being routed on this connection) and an interval of adaptation. In this way, each node remains tagged for the most appropriate routing treatment.

Any factor which can be monitored may be used to inform adaptation at each level, providing flexible, extensible support for adaptation and resource awareness.

The **network abstraction** sub-system maintains the hybrid network architecture described in section 3.1.

Applications interact with the system through the network services abstraction, which provides a high-level interface to the underlying network, supplying common functionality such as connection, search, broadcast and point-to-point message delivery, thus abstracting over the specific *resource discovery* and *routing abstraction* used.

The **applications** themselves are considered outside of the RaDP2P model and independent from the underlying substrate and adaptation components. This makes applications portable across systems which may use different routing and resource discovery overlays and we hope that it will encourage the rapid development of both novel applications and adaptation strategies.

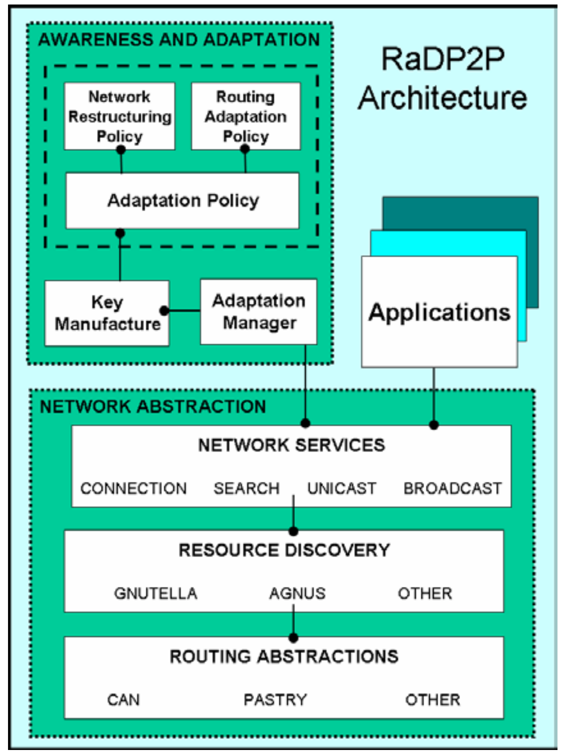

[Figure 1 – RaDP2P Architecture]

The Framework is being developed in Java. Adaptation policies and monitoring components are implemented based upon provided interfaces. Policy and monitoring components are loaded at run time using the Java reflection API.

5. Case Studies

In this section we describe two applications which are used as cases-studies to illustrate the potential of the RaDP2P framework. The case-studies used are ad-hoc mobile chat and geographically-aware service location.

- The ad-hoc mobile chat application illustrates how routing behaviour adaptation can be used to compensate for the highly variable capabilities of nodes, maximizing the performance of the network as a whole and also the benefit accrued by individual nodes.

- Geographically aware service location is an example of a of novel peer-to-peer application which makes use of network restructuring functionality.

5.1 Ad-Hoc P2P Chat

Consider an ad-hoc P2P chat application which operates over an unstructured decentralised network infrastructure (similar to Gnutella). While such networks are excellent for forming ad-hoc groups and simple resource discovery, the bandwidth consumed due to message passing can be prohibitively high for mobile nodes.

In unstructured decentralised networks, all message-passing is handled by the peer-nodes themselves. As peers must route all network messages, participating in this kind of community on expensive, narrowband, mobile connections, such as GPRS and GSM, would be extremely expensive and consume a significant fraction of each nodes available bandwidth. This makes participation for such nodes unfeasible.

In order to reduce the cost of participation in such networks for mobile nodes, a routing adaptation policy could be defined which tags nodes based on their connection type (mobile or fixed). Tagging nodes is accomplished by manufacturing the least-significant bits of their key from meta-information which reflects their connection type as described in section 3.

If any node has a directly connected peer possessing such a tag, it will not use it to route messages that are destined for other peers, instead making the next best selection from it's routing table. In this way, mobile nodes will receive all messages intended for them, but will not participate in routing messages destined for other nodes.

As messages are not routed through mobile nodes, the typical hop-count between origin and destination nodes will increase, potentially reducing performance, however, the performance decrease caused by an extra few hops when delivering messages may well be counterbalanced by the increased reliability of message delivery, as less reliable mobile nodes are no longer taking part in the message-routing process.

In situations where a very large number of nodes on the network are tagged as mobile, a situation could potentially arise where some nodes become unreachable or message delivery times become unacceptably long. In cases such as these, mobile nodes could be forced to route messages, ensuring they will still be delivered.

The effect of this kind of routing adaptation is to allow nodes on low-bandwidth mobile connections to use the ad-hoc chat service cost effectively and without having their limited bandwidth flooded by messages they are routing to other nodes. Furthermore, the network may well benefit from the exclusion of slow, unreliable peers from the routing process.

5.1 Geographic Service Location

We hope the ability to modify network structure based on a wide range of meta-information will lead to some novel applications and adaptation strategies.

Consider the example of a peer-to-peer communications network designed to support mobile emergency workers. This system uses mobile devices participating on an ad-hoc peer-to-peer network.

A network restructuring policy may be defined, wherein meta information harvested about the nodes geographical location (for example from GPS hardware), is used to restructure the network, such that nodes which are geographically closest are located close to each other on the routing overlay network.

By maintaining the network structure so that it reflects the geographical position of nodes, queries can be efficiently directed to those peers who are geographically closest to the sender (and thus able to assist most rapidly in the case of an emergency).

As queries are only broadcast to those peers who are geographically close and hence able to respond within a helpful time-frame, bandwidth consumption would be low. This is particularly important in low-bandwidth mobile environments.

6. Conclusions

P2P networks are typically highly heterogeneous and the heterogeneity of such networks is likely to increase due to the growing use of mobile devices and improving mobile Internet access technologies.

We argue that adaptation in peer-to-peer systems is vital; both to exploit the resources of heterogeneous nodes and to tailor network services to meet the needs of nodes with diverse requirements. We believe current peer-to-peer networks have inadequate support for adaptation and that a generic framework for supporting adaptation in peer-to-peer systems is required.

We propose a framework for developing Reflective and Dynamic peer-to-peer systems (RaDP2P) which uses a hybrid network structure similar to Structella [11] to provide generic and extensible support for adaptation in peer-to-peer systems. Through the use of this framework we hope to promote the development of novel applications and adaptation strategies.

7. References

[01] Gnutella, http://gnutella.wego.com

[02] Skype, http://www.skype.com

[03] Seti@Home, http://www.setiathome.com

[04] *"GIA: making Gnutella-like P2P Systems Scalable"* Chawathe Y., Ratnasamy S., Breslau L., Lanham N., Shenker S. published in the proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications Karlsruhe, Germany August 2003.

[05] "*Brocade: Landmark routing on overlay networks"* Zhao B., Duan Y., Huang L., Joseph, A., Kubiatowicz J. published in the proceedings of 1st International Workshop on Peer-to-Peer Systems (IPTPS), Linoping, Sweden 2001.

[06] *"Constructing a Balanced, log(N)-Diameter Super-peer Topology"* Pyun Y., Reeves D., published in the proceedings of the 4th IEEE International Conference on Peer-to-Peer computing (P2P'04). Zurich, Switzerland, August, 2004.

[07] *"From Selfish Nodes to Cooperative Networks Emergent Link-based incentives in Peer-to-Peer Networks"*, Hales D., published in the proceedings of the 4th IEEE International Conference on Peer-to-Peer computing (P2P'04). Zurich, Switzerland, August, 2004.

[08] *"AGnuS: The Altruistic Gnutella Server"* Hughes D., Warren I., Coulson G., published in the proceedings of the 3rd IEEE International Conference on Peer-to-Peer computing (P2P'03). Linköping, Sweden, September, 2003.

[09] *"The Interactive Helpdesk System",* http://polo.lancs.ac.uk/p2p

[10] *"Should be build Gnutella on a Structured Overlay",* Castro M, Costa M, Rowstron A. published in the proceedings of the 2nd Workshop on Hot Topics in Networks (HotNets-II). Cambridge, MA USA, November 2003.

[11] *"Dependability Properties of P2P Architectures"*, Walkerdine J., Melville L., Sommerville I. proceedings of the 2nd IEEE International Conference on Peer-to-Peer computing (P2P'02). Linkoping, Sweden, September, 2003.

[12] "*Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems."*, Rowstron A, Druschel P. available at http://research.microsoft.com/ antr/PAST/, 2001.

[13] *"A scalable content-addressable network"* Ratnasamy S., Francis P., Handley K., Karp R., and Shenker S. published in the Proceedings of the Internet Measurement Workshop (ACM SIGCOMM), San Francisco, USA, November 2001.

[14] The Fasttrack Protocol, www.p2pwatchdog.com/packet_fasttrack.html

[15] "SCRIBE: The design of a large-scale event notification infrastructure (2001)", published in networked Group Communication, 2001, pp. 30--43.

[16] The Open Overlays Project, http://www.comp.lancs.ac.uk/computing/research/mpg/projects/openoverlays/

[17] "*PROST: A Programmable Structured Peer-to-peer Overlay Network"* Portmann M, Ardon S, Senac P, Seneviratne A., proceedings of the 4th IEEE International Conference on Peer-to-Peer computing (P2P'03). Zurich, Switzerland, September, 2003.