Hypermedia by coincidence

Mark K. Thompson and David C. De Roure
Department of Electronics and Computer Science
University of Southampton, UK
Tel: +44 23 8059 4490
Email: {mkt|dder}@ecs.soton.ac.uk

ABSTRACT

We introduce an approach to linking hypermedia documents dynamically in a decentralised, peer-to-peer manner using resources that are available by coincidence, without explicit configuration. The particular approach presented utilises an open platform in combination with Southampton Distributed Link Service technology enabling dynamic hypertext generation.

KEYWORDS: Open Hypermedia; Transient Web; Dynamic Link Services; Peer to Peer systems; Pervasive Computing

INTRODUCTION

The past twelve months have seen an explosion of Internet applications enabling users to participate in compelling applications that break 'traditional' barriers for resource sharing without the need for central authority or service. Applications including Napster, FreeNET, and SETI@Home to name but a few, demonstrate a shift away from the client/server modality of the Web with a trend towards a peer-to-peer approach where clients of the system are as much providers of services as they are users.

Focusing in from Internet-wide applications, there are many scenarios where peer-to-peer relationships are desirable in localised arenas, utilising *local* resources for the dissemination of *local* information, for example, in business meetings, conference seminars or public gatherings in museums and galleries. The culmination of all the resources carried by each individual participant peer on their laptop, PDA or embedded in their clothing(!) adds to both the complexity and diversity of the information available to be discovered and potentially navigated.

The decentralised peer-to-peer model is disruptive because it offers *probabilistic* services over traditional *deterministic* services: the set of applications and resources that are available at any particular time is unknown and unpredictable: the same peer action (for example, a query for a document or a distributed search for a user) could yield different results when repeated moments later. It truly is the case that pure coincidence determines what resources are available at any given time.

P2P AND HYPERMEDIA

Aside from Instant Messaging and the need to share MP3 music files, the hypermedia aspect concerns how multimedia information can be meaningfully conveyed to and navigated by participant peers. Through utilising the open hypermedia model whereby hypertext links are treated as first class citizens — arcs referencing both source and destination resources — we are presented with the necessary means to enable navigation between the available resources.

Activities such as WAP[1] and PIE[1] cater for scenarios where devices other than just laptops are mobile with access to the Internet for published data. These approaches are somewhat limited in that they all assume connectivity to a wide-area network and do not address the hypertext issues concerning navigating (linking) between resources that co-incide in and around the user device's locality.

In a peer-to-peer context, the resources that may be present on users' devices may include copies of resources[2] that were originally published on the World Wide Web, on intranets or on personal filestores. The de facto method of choice for addressing these resources in their 'home' context is typically through URLs, or naming services that resolve well-known URLs such as PURL[5]. These naming approaches all refer to the original context of the resources, not the 'copy' that happens to be local and therefore more accessible, perhaps in a more appropriate format.

The issue of resource discovery for multimedia documents in this context is further complicated when pre-authored or generated linkbases (collections of links) are introduced. Typically, the source and destination anchors of links are authored to refer to the documents in their 'global' context. However, in the peer-to-peer scenarios alluded to here, those URLs are

[1] `http://www.wap.net/`

[2] *copies* here can be interpreted as verbatim, reformatted, summarised, annotated or augmented versions of resources

meaningless. Providing a centralised resource naming service, like the one suggested by [4], is not suitable given the ad hoc nature of the network of peers. Rather, a decentralised service that maps global URLs to identifiers to resources that are presently available is required.

P2P AND SOUTHAMPTON

The Southampton focus is on information applications that take advantage of a hyperstructure in these ad hoc and dynamic settings. The goal of our research is to provide an architecture of components that empower users to discover information that is available by virtue of the other peers around them in their locality, be they user agents or devices provided by the organisation controlling the venue. These interactions should be enabled in such a way as is familiar and easy to interact with, whilst involving as little, if any, manual configuration or policy definition as can be avoided.

The candidate system developed to explore this space is based on our Link Service Infrastructure[2] and Pervasive Information Fabric[3] research, comprising of a number of servent processes — code that acts as both a provider of a service and as a consumer. For the underlying peer-to-peer architecture, we have adopted code from Project Jxta[3], a recently announced open-source project. Jxta offers a stack of core services including discovery, communication and security provisioning. Jxta is also aiming to cater for different architectures other than PC desktops, for example PDAs and embedded devices, rendering it especially interesting given our on-going Pervasive Information Fabric activities.

In our system, we have a servent component through which applications such as Web browsers can navigate the documents present in the various peers' stores. This is effectively a window onto resources that can be local to the peer or to documents on other peers, dependent on who else is local. Instances of this component interact with a Link Service component that offers a web-proxy interface (similar to that in [3]) for link resolution and linkbase communication, adding hyperstructure to the available resources where appropriate. Currently, both servents share the responsibility for overcoming the resource naming issue through distributed hashes that map global URLs to peer resource identifiers. Both components are Jxta-enabled for discovery and communication such that, controlled through a GUI, the peer can be made aware of and interact with components that embody new sources of documents and new linkbases that are available by coincidence.

The upshot is, given the appropriate set of present peers, a user of the system can be presented more data than would otherwise be possible with standard document publishing techniques on local networks. To realise this in the 'traditional' web service modality would require centralised services, manually configured and reconfigured when new peers join the system.

This is only a proof-of-concept application at this stage. There are no quality controls on the amount and type of data shared amongst peers, nor is there yet any processing performed on the available data. Applications that adapt the view of the content provided to a peer depending on their capabilities, or level of skill are trivial extensions of this code. One can readily envisage Transcoder and Summariser servents that could perform these services in an ad hoc manner.

SUMMARY

We see the peer-to-peer trend of decentralised and on-demand information sharing to be a target community for enrichment with open hypermedia services. We have discovered that the critical hypermedia application issue when introducing hypermedia services in such environments is how to resolve URLs to local copies or representations of the resource in a meaningful way.

This paper alludes to a system of distributed components that make available peer resrouces (e.g. documents and linkbases) and provide distributed link services and name resolution technology over the Jxta peer-to-peer platform.

ACKNOWLEDGEMENTS

Work developed in part under the IBM University Fellowship Programme under the guidance of Adrian Colyer, to whom we are particularly grateful for discussion.

REFERENCES

1. B. Carmeli, B. Cohen, and A. Wecker. Personal Information Everywhere (PIE). In *HT'00 - Proceedings of the Eleventh ACM Conference on Hypertext and Hypermedia Systems*. ACM, May 2000.

2. D. De Roure, N. Walker, and L. Carr. Investigating Link Service Architectures. In *HT'00 - Proceedings of the Eleventh ACM Conference on Hypertext and Hypermedia Systems*. ACM, May 2000.

3. M.K. Thompson, D.C. De Roure, and D.T. Michaelides. Weaving the Pervasive Information Fabric. In S. Reich and K.M. Anderson, editors, *Proceedings of the Sixth Workshop on Open Hypermedia Systems and the Second Workshop on Structural Computing*, number 1903 in Lecture Notes in Computer Science. Springer Verlag, August 2000.

4. M. Tzagarakis, N. Karousos, D. Christodoulakis, and S. Reich. Naming as a Fundamental Concept of Open Hypermedia Systems. In *HT'00 - Proceedings of the Eleventh ACM Conference on Hypertext and Hypermedia Systems*. ACM, May 2000.

5. S. Weibel, E. Jul, and K. Schafer. PURLs: Persistent Uniform Resource Locators. *OCLC Newsletter*, November 1995. *http://purl.oclc.org/OCLC/PURL/SUMMARY*.

[3] `http://www.jxta.org/`