

Articles

PEER-TO-PEER INTERACTIONS IN WEB BROKERING SYSTEMS

Ubiquity, Volume 2002 Issue May, May 1- May 31, 2002 | BY [GEOFFREY FOX](#)

 [SHRIDEEP PALLICKARA](#) 

|

Full citation in the ACM Digital Library 

Global structure and local dynamic messaging support a wide range of applications

The peer-to-peer (P2P) style interaction model facilitates sophisticated resource sharing environments between "consenting" peers over the "edges" of the Internet; the "disruptive" impact of which has resulted in a slew of powerful applications built around this model. Resources shared could be anything from CPU cycles, exemplified by SETI@home (extraterrestrial life) and Folding@home (protein folding), to files (Napster and Gnutella). Peer "interactions" involves advertising resources, search and subsequent discovery of resources, requests for access to these resources, responses to these requests and exchange of messages between peers.

Since peer interactions, in most P2P systems, are XML-based, peers could be written in any language and can be compiled for any platform. Driven entirely on demand a resource may be replicated several times; a process that is decentralized and one over which the original peer that advertised the resource has sometimes little control over. The greater the demand for a resource, the greater the system initiated replications for the resource, resulting in fewer bottlenecks and faster accesses to that resource.

Services proliferating the Internet can be abstracted as a distributed system of "clients" comprising "users," "resources" or proxies thereto. These Web services are usually hosted on a "broker" or a network of brokers, where the term server is avoided to distinguish it clearly from application servers. Clients access Web services by connecting to a broker, where the client's online presence is neither time constrained nor is it tied to a specific device, geographic location or communication channel. Clients then maintain active connections to the hosting broker throughout the duration that they use the service.

Integrating support for P2P systems, such as Sun Microsystems' JXTA (from juxtaposition), within the brokering framework have several benefits that can be accrued by both systems. This integration would allow Web-brokering systems to leverage P2P interactions and applications while providing peers access to legacy systems. Intelligent routing of P2P interactions, which would be accorded to peers in the integrated system, would result in improved latencies at peers and efficient overall network usage. Furthermore, the integrated system would enable the development of "hybrid-clients" and applications that could utilize features derived from these different frameworks. In such integrated systems not every hybrid-client interaction needs to be funneled through the broker network since such hybrid-clients can sometimes service each other's requests while peer accesses to shared resources can be reliable and error-free. Support for other messaging frameworks such as the Java Message Service (JMS) within the brokering system and bridges to, and from, other P2P systems would facilitate the

efficient development of rich complex applications.

In traditional brokering systems the request/response model is fixed (one response for a request) with the client initiating the request having no ambiguity on how the request would be interpreted while also being aware of the template that the response would conform to. In contrast, in P2P systems, not every request goes through, and even if it does, there could be zero or more valid responses. Peers anticipate neither the template that the responses would conform to nor the order in which these responses would be received. Furthermore, responses are not identical with each responding peer processing any given request based on the resources at its disposal and its interpretation of the request. Communication between a requesting peer and responding peers is facilitated by peers en route to these destinations. These intermediate peers are thus made aware of capabilities that exist at other peers. This discovery of services offered by other peers constitutes dynamic real time knowledge propagation. Peer requests are sometimes satisfied through cached responses, and peers generally have a choice on whether to accept these cached responses or not. It's up to the peer to discard responses that it deems are not correct.

P2P interactions are self-attenuating with interactions dying out after a certain number of hops. These attenuations in tandem with traces of the peers, which the interactions have passed through, eliminate the continuous echoing problem that results from loops in peer connectivity. However, attenuation of interactions sometimes prevents clients from discovering certain services being offered. Peers hosting these services could not respond because the request attenuated before it could have reached the hosting peer. This results in P2P interactions being very localized. Of course, if the peer at the edge of the attenuation had cached the response or the resource itself, for a similar request that it made at an earlier time, that particular peer can respond to the requesting peer. Unlike clients in Web brokering systems that interact via the broker network, peers in P2P systems interact directly with each other and sometimes use other peers as intermediaries in interactions. Specialized peers are sometimes deployed to enhance routing characteristics. Nevertheless, sophisticated routing schemes are seldom in place and interactions are primarily through simple forwarding of requests with the propagation range being determined by the attenuation indicated in the message.

The aforementioned issues of discovery of services and intelligent routing can be addressed very well in distributed brokering systems. Single broker solutions would entail P2P interactions being funneled through the lone broker leading to bottlenecks. The associated drawbacks in scaling, increased latencies due to queuing delays and the single point of failure that such a scheme constitutes are among the reasons favoring a distributed model being in place. Strategies for integrating P2P interactions need to ensure that minimal changes need to be made to the brokering core. Furthermore, this integration should entail neither any changes to the peers nor a straitjacketing of the interactions that these peers could have had prior to the integration. A solution to this problem is via P2P-proxies, which would provide an interconnection bridge between the two systems. The P2P-proxies are part routing peers and part clients of the Web brokering system. Peers would interact with the P2P-proxies as they would with any other peer, while the P2P-proxy inherits guarantees accorded to clients of the brokering system. Since every P2P-proxy is treated as a peer by the other peers it is part of the communication and dynamic discovery schemes associated with original peers. In its dual role as a client of the brokering system, the P2P-proxy forwards these interactions to the broker that it is attached to in order to aid in their efficient disseminations. This strategy involves no changes to peers neither does it entail major changes within the existing distributed broker network. The broker network would be used primarily as a delivery engine, and a pretty efficient one at that, while locating peers and propagating interactions to relevant peers.

The most important aspect in P2P systems is the satisfaction of peer

requests and discovery of peers and associated resources that could handle these requests. The broker network needs to forward these requests only to those peers that it believes can handle the requests. Peer interactions in most P2P systems are achieved through XML-based data interchange. XML's data description and encapsulation properties allow for ease of accessing specific elements of data. Individual brokers routing interactions could access relevant elements, cache this information and use it subsequently to achieve the best possible routing characteristics. Routing for P2P interactions can be optimized since individual brokers, being aware of interconnectivities within the broker network, the broker topology and individual client/peer interests; would arrive at routing decisions to ensure that the events would be routed to the P2P-proxies best suited to handle the P2P interaction. These P2P-proxies would then locate peers that would service the request or the peer that the response was intended for or peers that are en route to the destination peer(s). The brokering system, since it is aware of advertisements, can also act as a hub for search and discovery operations. These advertisements when organized into "queryspaces" allow the integrated system to respond to search operations more efficiently.

Resources in Web brokering systems are generally within the purview of the broker network. P2P systems replicate resources in an ad hoc fashion, the availability of which is dependent on the peer's active digital presence. Some resources, however, are best managed by the brokering system rather than being left to the discretion of peers who may or may not be present at any given time. An understanding of the network topology and an ability to pinpoint the existence of peers interested in that resource are paramount to efficient replications of a resource. The distributed broker network, possessing this knowledge, best handles this management of resources while ensuring that these replicated resources are "closer" and "available" at locations with a high interest in that resource. Furthermore, the broker network is also better suited, than a collection of peers, to eliminate race conditions and deadlocks that could exist due to a resource being accessed simultaneously by multiple peers. Broker networks can be responsive to changes in peer concentrations, volumes of peer requests, and resource availability. Brokers and associated interconnections can be dynamically instantiated or purged to compensate for affected routing characteristics due to changes in peer interactions.

The integrated system could also be used to connect islands of peers together. Peers that are not directly connected through the peer network could be indirectly connected through the broker network. Peer interactions and resources in the P2P model are traditionally unreliable, with interactions being lost or discarded due to peer failures or absences, overloading of peers and queuing thresholds being reached. Guaranteed delivery properties existing in traditional brokering systems can augment peer behavior to provide a notion of reliable peers, interactions and resources.

Such an integrated brokering solution would also allow for hybrid interaction schemes to exist alongside each other. Applications could be built around hybrid-clients that would exhibit part peer behavior and part traditional client behavior. P2P communications could be used for traffic where loss of information can be sustained. Similarly, hybrid-clients needing to communicate with each other in a "reliable" fashion could utilize the brokering system's capabilities to achieve that. Sometimes, hybrid-clients can satisfy each other's requests, in which case they would, obviating need for funneling interactions through the broker network. The broker merely serves as an efficient conduit for supporting interaction between different applications (clients, peers or hybrid).

A typical application of this integrated solution is distance education. This often consists of multiple linked classrooms where the participants in each classroom are individually linked to collaborative environment. Here a peer-to-peer model can be used in a classroom to give fast dynamic response to shared input control while the guaranteed delivery capability is used to multicast between classrooms. More generally this combination of global

structured and local dynamic messaging scales to support a wide range of applications.

Geoffrey Fox is a professor in the Departments of Computer Science and Physics at Indiana University. Shrideep Pallickara is a post doctoral fellow in the Computer Science Department at Indiana University. The research prototype of the strategy outlined in the article is available at <http://grids.ucs.indiana.edu/ptliupages/projects/narada/>.

COMMENTS

POST A COMMENT

Your Name (Required)

Your E-Mail address (Required)

Comment (Required - HTML syntax is not allowed and will be removed)

Post

