Peer networking as a Personal Autonomic Computing Enabler

David Bantz Chatschik Bisdikian Steve Mastrianni Archan Misra

[1]IBM T. J. Watson Research Center, Yorktown Heights, NY 10598, USA

{dbantz, bisdik, stevemas, archan}@us.ibm.com

EXTENDED ABSTRACT

Computer communications based on peer-to-peer (P2P) technologies has seen measurable growth over the last few years. While a concrete definition of a P2P communications system does not yet exist it typically refers to situations where groups of users form virtual communities in an ad-hoc manner to share compute resources directly among themselves without the need of an elaborate infrastructure. The Napster model is an example of a highly centralized P2P system developed for the sharing of music files by distributing queries regarding the content of these files (the file metadata) through a central peer server. The Gnutella system is an example of a highly distributed P2P system where both content and metadata are distributed among all the nodes participating in the peer system. The SETI@Home project is another example of a peer network designed to support highly distributed computing by tapping into the computation power of idling computers.

Peer networking solutions have been developed which span the Internet and provide convenience to computer users in performing certain tasks. Autonomic computing, however, is being developed to benefit computer system administrators to perform their task more efficiently. Using auto-configuration and self-healing techniques, autonomic computing aims to simplify if not eliminate, human intervention during the operation of a computing system, such as a server cluster providing Web hosting services. Personal autonomic computing is an effort to extend the benefits of autonomic computing to personal computers and their users. Personal autonomic computing has become more important with the ever-increasing use of notebook computers that are used by a highly mobile enterprise workforce. While an application such as a slide presentation or word-processing application is being executed, it behaves in a fashion that is independent from the physical location of the computer. The same cannot be said about the computer itself. Personal computers may exhibit different behavior when they operate in different physical environments such as different office rooms, conference rooms, highly-distributed enterprise branch sites, homes, on the road, or in a hotel They also exhibit different behavior if they are operating in battery mode or unplugged into an AC socket.

Knowledge of location and context provides information of ever-increasing importance in the operation of a personal computer. For example, let us consider Jane during her visit to a remote branch office of her corporation. During her visit, she finds it difficult to acclimate herself to the new locale and so does her notebook computer. Jane can obtain little, if any, information about her new working location, such as the location or name of the local file or print server, and her computer may not be able to connect to the Internet. Information about networked services is usually acquired using service discovery technologies such as Jini or UPnP. However, typical service discovery procedures depend on an established support infrastructure (which itself has to be discovered). Although several standards for network discovery have been introduced, none of them has emerged as the preeminent standard accepted industry wide.

To address these issues, we have decided to exploit peer networking technology to develop an easily deployable system to support personal autonomic computing for mobile computers. Instead of having to ask her colleagues for help in configuring her notebook computer at the new work site, Jane's computer will configure itself autonomically by searching for and connecting to peer computers at the new site to learn how to configure itself. The computer does so without any prior knowledge of the location or existence of any other computers. Information about local printers, their physicallocation, SOCKS server settings, browser configuration parameters, and important Web bookmarks can then be obtained from a local peer system.

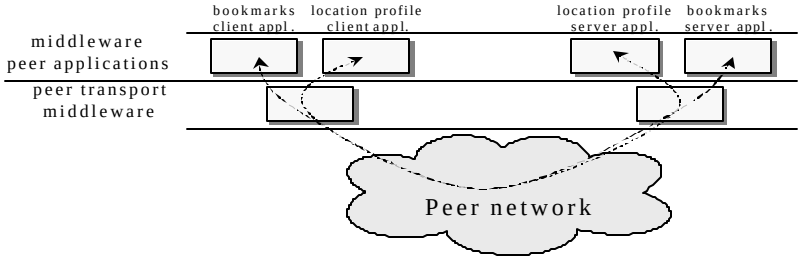

Figure 1: Peer networking-based personal autonomic computing.

Figure 1 summarizes the objectives of our endeavor in this space. It includes developing the necessary transport mechanisms over a peer network and applications in support of personal autonomic computing. This will establish a new paradigm for using peer networking by providing users with the ability to quickly adapt their system to a new location. Configuration information, such as the location of certain programs, files, printers, and other resources can be easily transferred to the user's computer. We believe this approach has great promise, since, quite often, valuable context information regarding computing operations is not centrally available, but rather distributed across individual personal computing devices.

Our initial work focuses on the creation of a flexible transport middleware layer and on the development of a few target applications. The transport middleware layer comprises a communications protocol running on top of a peer protocol stack; the peer protocol stack supports the ad-hoc creation of virtual communities of personal computers possessing a common binding relation. To instantiate our approach, we have selected the Jxta peer networking technology (http://www.jxta.org), which is an open source effort for developing a peer communications platform via a collection of protocols that provide peer discovery, peer membership, and pipe binding.

When Jane visits the remote branch office, her computer may need to create a peer network connection with her colleagues'computers. The first challenge is the creation of a peer network with other computers that: (a) are verified to belong to the same company (imagine creating a peer network with a notebook computer owned by a visitor); (b) support the autonomic applications Jane's computer is looking for; and (c) do not require an established infrastructure to support the above (cost is always of concern) or any prior knowledge of the existence of the other computers in the network.

When the desired peers have been located and the needed information has been collected, we need assess the validity of this information. For example, when site-related information is sought, security and privacy rules may first restrict what kind of information is shareable from a peer. Also, as the requested pieces of information arrive from other peers, some of that information may be more important than the information previously received from others. Therefore, analysis of any received data must be performed prior to making it available to the user. Since the objective is to support personal autonomic computing, the challenging task is to devise new rules for the analysis of the received information, and to classify that information and present it to a user in a meaningful form if so needed. These rules should be easily created and modified, and should optimize the behavior of the personal computer based on the context information obtained from other peers.

Peer networking technologies hold great promise in simplifying the way users deal with applications and computers. This project is focused on exploiting this new usage paradigm for peer networking and investigating the lengths to which it can be taken to make personal computing even more user-friendly.