Globally Distributed Datacenters: A Collaborative Peer-to-Peer Approach

H. M. N. Dilum Bandara and Anura P. Jayasumana
Department of Electrical & Computer Engineering
Colorado State University, Fort Collins, CO 80523

Large-scale and collaborative distributed processing and resource sharing, beyond file sharing, have been among the key promises of Peer-to-Peer (P2P) systems since early 2000's. While initiatives such as BOINC [1] and OurGrid [2] have demonstrated tremendous success in tapping into globally distributed, residual computing resources, they mainly execute independent batch jobs. With the advances in Web 2.0, ubiquitous high-speed networks, cloud computing, and widespread use of strong social networks, we envision that P2P systems will play greater role in distributed resource collaborations and collaborative applications. Emerging collaborative P2P systems [3] look for diverse peers that bring in unique resources and capabilities to a virtual community thereby empowering them to engage in greater tasks that cannot be accomplished by individual peers, yet are beneficial to all the peers/users. Collaborative P2P systems are applicable in a wide variety of contexts such as grid/cloud computing, distributed collaborative adaptive sensing [4], opportunistic computing [5], Internet of things, and mobile social networks. These systems share a variety of *resources* such as processor cycles, storage capacity, network bandwidth, sensors, special hardware, middleware, scientific algorithms, application software, services (e.g., web services and spawning nodes in a cloud), and data to not only consume a variety of contents but also to generate, modify, and manage those contents. For example, with increased levels of integration, even systems within and across datacenters exhibit attributes of distributed systems where groups of resources such as processing nodes, storage, bandwidth, and special hardware (e.g., GUPs and FPGAs) may be grouped to execute complex applications. P2P provides a scalable approach for these tasks. However, it is nontrivial to discover, aggregate, and utilize heterogeneous and dynamic resources that are distributed and characterized by multiple static and dynamic attributes.

We identify seven key phases of resource collaboration in P2P systems namely: advertise, discover, select, match, bind, use, and release [3]. Ability to support these key phases is essential for any resource discovery/aggregation solution and job scheduler, to achieve high performance, low latency, and quality of service. Significant progress has been made in multi-attribute resource discovery in grid computing and P2P [3, 6]. However, these solutions primarily focus on the first three phases. Moreover, in the absence of information, data, and understanding of real-life, multi-attribute resource and query characteristics, designs and evaluations of existing solutions have relied on many simplifying assumptions. Through the analysis of four real-world systems, we observed that their characteristics diverge drastically from conventional assumptions [6]. Consequently, existing solutions incur significant overhead and are prone to considerable load balancing issues when applied to real workloads [7]. These findings point to the need for novel solutions that can cohesively support all key phases of resource collaboration while being efficient, scalable, and load balanced. Such solutions can harness the collective power of P2P communities and their underutilized/unused resources to build a globally distributed datacenter with limitless number of applications.

[1] D. P. Anderson and K. Reed, "Celebrating diversity in volunteer computing," Hawaii Int. Conf. on System Sciences, Jan. 2009.
[2] N. Andrade et al., "Peer-to-peer grid computing with the OurGrid community," 23rd Brazilian Sympo. on Computer Networks, May 2005.
[3] H. M. N. D. Bandara and A. P. Jayasumana, "Collaborative applications over peer-to-peer systems – Challenges and solutions," Peer-to-Peer Networking and Applications, 2012, DOI: 10.1007/s12083-012-0157-3
[4] D. McLaughlin et al., "Short-wavelength technology and the potential for distributed networks of small radar systems.," Bull. Amer. Meteor. Soc., vol. 90, Dec. 2009.
[5] M. Conti et al., "From opportunistic networks to opportunistic computing," IEEE Communications Magazine, vol. 48, no. 9, 2010.
[6] H. M. N. D. Bandara and A. P. Jayasumana, "Characteristics of multi-attribute resources/queries and implications on P2P resource discovery," 9th ACS/IEEE Int. Conf. On Computer Systems And Applications (AICCSA '11), Dec. 2011.
[7] H. M. N. D. Bandara and A. P. Jayasumana, "Evaluation of P2P resource discovery architectures using real-life multi-attribute resource and query characteristics," IEEE Consumer Communications and Networking Conf. (CCNC '12), Jan. 2012.

This work is supported by the Engineering Research Center program of the National Science Foundation under NSF award number 0313747.