Peer-to-Peer Streaming

1. Research Team

Project Leader:	Prof. Cyrus Shahabi, *Computer Science*
Other Faculty:	Prof. Leana Golubchik, *Computer Science* Prof. Roger Zimmermann, *IMSC and Computer Science*
Graduate Students:	Farnoush Banaei-Kashani
Undergraduate Students:	Gerran Ueyama

2. Statement of Project Goals

The goal of the **Kerena** project is to study, design, and develop distributed resource-sharing frameworks with two distinct applications:

1. To serve objects to a large set of geographically dispersed clients; and
2. To allow a large set of peers share resources distributed among them.

These frameworks enable applications such as video-on-demand, employee training, distance learning, and peer-to-peer stream-sharing on a large scale. At IMSC, we plan to incorporate and evaluate **Kerena** within both the *2020Classroom* and *ImmersiNET* projects.

3. Project Role in Support of IMSC Strategic Plan

Kerena provides a general peer-to-peer resource-sharing environment. This is an essential requirement of both the *2020Classroom* project and the *ImmersiNET* project, where multiple remote users/peers would join a collaborative environment dynamically to share resources such as educational material in various modalities and high fidelity audio, video, and image content.

There are many challenges in realizing such a large-scale and dynamic resource-sharing environment. Among these challenges, *search* (or resource discovery) is one of the most fundamental issues. With a fully decentralized data management framework, the time required to find a shared resource becomes proportional to the network size; hence, not scalable. In the context of the *2020Classroom* project (potentially with thousands of users), locating an object (e.g., a video clip) requested by one of the users requires querying all other users. The goal is to search for the requested object *efficiently* in a distributed manner, without incorporating a centralized resource management unit such as a directory server. Centralized management is not tolerable in a dynamic and distributed learning environment because each student can potentially create his/her own content autonomously. We adopt lessons learned from other disciplines, such as social-network studies (e.g., the small world effect), and receive inspirations from natural complex systems, such as biological networks, to design effective and scalable search techniques for unstructured communities of autonomous users.

4. Discussion of Methodology Used

Distributed resource-sharing environments can be deployed in two fundamentally different ways. In the context of streaming applications, for instance, on one hand there are *content distribution systems* that are organized sets of content servers used to provide streaming service to a large set of geographically dispersed users. Such systems are usually designed, developed, and maintained mainly by a single authority, e.g. a company, and are mostly static, homogeneous, and robust, but possibly fragile. Since content distribution networks are organized systems, well-designed algorithms can perform resource management and object discovery tasks in these systems efficiently. Particularly, with known object placement policy, search for objects in such systems is mapped to source-to-destination routing. On the other hand, autonomous users with personal content servers can generate distributed streaming networks to share their objects. As compared to content distribution networks, these networks are usually more dynamic, heterogeneous, and resilient. Management tasks in these self-organized systems are ad hoc.

With **Kerena**, we are designing distributed resource-sharing frameworks with *peer-to-peer* management architectures that suit the types of environments described above. **Ace** (network-Aware Community-based strEaming network) is a content distribution network with *structured* peer-to-peer management architecture inspired by the community model of the social networks. **Ace** supports server heterogeneity and object replication, and considers network locality to minimize the redundant overhead of the management with the overlay peer-to-peer network. Also, we are designing **GlobeStream**, an *unstructured* peer-to-peer network of autonomous users to share resources. We consider **GlobeStream** as a "complex system", a system with macroscopic behaviors that emerge from simple microscopic interactions between the clients. With **GlobeStream**, we investigate various techniques to optimize the search procedure, which is the most dominant management task in such systems. We are also developing analytical models that capture the dynamic behavior of such unstructured peer-to-peer networks as instances of complex systems.

5. Short Description of Achievements in Previous Years

During previous years, we studied various distributed resource management techniques (see [1]) by designing and simulating a preliminary generation of **Ace**, termed GMeN. Typically a distributed content server is designed as a pure hierarchy (tree) of centralized servers. We proposed a redundant hierarchical topology for GMeN, termed RedHi [2]. RedHi relaxes the hard degree-1 parent connectivity restriction of a pure hierarchy to be degree-2 or more. Consequently, there is a higher potential for load-balancing among nodes and links of the network; hence, higher resource utilization, cost-efficiency, and reliability result as compared with a pure hierarchy. Moreover, we developed a distributed resource management scheme that exploits the characteristics of the RedHi topology to realize its main potential advantage, i.e., higher utilization of system resources [1]. Our proposed resource management system is based on a fully decentralized approach to achieve optimal scalability and robustness.

In addition, we also developed an efficient search mechanism for our unstructured peer-to-peer network, **GlobeStream**. Without structure, location queries cannot be efficiently routed towards the node that holds the corresponding location metadata because there is no global consensus on the location of the metadata. Therefore, with a naive approach the location query is flooded to the network so that by exhaustive coverage eventually the node holding the corresponding location metadata receives the query and responds. Obviously, this search approach is not scalable. We consider an unstructured peer-to-peer resource-sharing network as a "Complex System" [3] and use percolation theory [4] to formalize the main problem with normal flooding search, propose a remedial approach with our *probabilistic flooding* technique, and find the optimal operating point for probabilistic flooding rigorously, such that it improves scalability of the normal flooding by 99% [5.6]. We also applied this novel search approach to provide a distributed discovery service for Web Services, as an instance of peer-to-peer networks [7].

5a. Detail of Accomplishments During the Past Year

During the past year, we developed the second generation of **Ace,** SWAM [8.9]. SWAM is a family of access methods for efficient similarity search in peer-to-peer networks. With SWAM, we further relax the hierarchical topology of the first-generation **Ace** (i.e., RedHi) to an arbitrary graph to improve its performance characteristics. To develop SWAM, we were inspired by the principles of the small-world models, which are proposed to explain the phenomenon of efficient communication in social networks [10.11]. First, in order to formalize the problem of similarity search in peer-to-peer networks, we 1) modeled the problem, 2) defined a set of metrics to evaluate the efficiency of the access methods for peer-to-peer networks, and 3) introduced a basic access method to set a lower efficiency bound for similarity search. Second, to define the basic properties that characterize the SWAM family, we generalized the efficiency principles of the small-world models. These properties ensure that the index structure effectively partitions the data space for efficient filtering, closely co-locates the nodes with similar content for batch content-retrieval, and properly interconnects the nodes for fast traversal of the data space. Finally, we introduced a Voronoi-based instance of SWAM, termed SWAM-V, which for a peer-to-peer network with N nodes has query time, communication cost, and computation cost of $O(\log N)$ for exact-match queries, and $O(\log N + sN)$ and $O(\log N + k)$ for range queries (with selectivity s) and kNN queries, respectively. SWAM-V proposes a non-hierarchical distributed index structure that indexes the peer-to-peer networks with multi-attribute objects. SWAM-V also respects the autonomy of the network nodes and self-configures the topology of the peer-to-peer network (i.e., the index structure) based on the nodes own content. Consequently, it avoids unnecessary content migration and replacement, supports object replication, and adapts to the object distribution as new nodes join the peer-to-peer network with the new content. We have performed a comparative study via simulation to verify the efficiency of SWAM-V versus our basic access method, as well as a version of CAN [12] that we extended over the original DHT to support range and kNN queries. Our experiments show that unlike the basic access method, SWAM-V achieves logarithmic query time with limited resource usage, and consistently outperforms CAN in query time and communication cost.

6. Other Relevant Work Being Conducted and How this Project is Different

To address the scalability problem with search in unstructured peer-to-peer networks, the research community has taken initial steps to gain a better understanding of the characteristics of these systems by performing measurement studies [13]. Moreover, assuming partially cooperative nodes, some has suggested replication of the metadata [14] or caching query responses [15] for more efficient flooding. Besides, researchers have tried different variations of classical search techniques, other than normal flooding: directed BFS and iterative deepening [16], expanding ring [16], and random walk [17.18.19]. Although some of these techniques improve the scalability of search, due to enormous complexity of the unstructured networks as *large-scale*, *self-configure*, and *dynamic* systems, these techniques often fail to characterize the unstructured networks with models that are both accurate and applicable, at the same time. They either assume oversimplified models that fail to capture important characteristics of the unstructured networks, or apply highly complex models that render analytical evaluations impossible.

We propose recognizing unstructured peer-to-peer networks as “complex systems” (e.g., social networks, biological networks, etc.) and employing the accurate statistical models extensively used to characterize these systems for formal analysis and efficient design of peer-to-peer networks. Since complex systems are large-scale and dynamic, taking account of individual entities of these systems results in highly complex and impractical models. Instead, statistical models (such as the percolation model) consider the complex systems as probabilistic systems and capture their statistical properties. Due to pure statistical effects originating in the law of large numbers, these models are able to characterize properties of large-scale networks accurately, while being applicable. With these theoretical models, we can formally characterize the unstructured networks, understand the problems with search in these systems, provide design guides for new search algorithms, and analyze the performance of the proposed search algorithms, and find the optimal operating points for their parameters.

7. Plan for the Next Year

We plan on continuing our research in two ways:

1. To extend the SWAM family by introducing a new member, SWAM-P, which as compared to SWAM-V enforces less constraining assumptions and supports peer-to-peer applications with specific restrictions/requirements. With some peer-to-peer applications (e.g., peer-to-peer streaming), strict enforcement of the neighbor selection rules to construct the index structure is either impossible or inefficient. SWAM-P is an access method with probabilistic index topology and flexible neighbor selection policies that allow network nodes to exercise their autonomy in selecting their neighbors as they join the peer-to-peer network.
2. To develop **GlobeStream** based on the probabilistic flooding search approach. We intend to evaluate the developed framework both analytically and via simulation. Finally, we study **GlobeStream** empirically in the context of the *2020Classroom* project and the *ImmersiNET* project.

8. Expected Milestones and Deliverables

- **GlobeStream-S**: a Discrete-Time Event Simulation engine to study large-scale peer-to-peer networks (NB: The first version of the simulator, which was developed last year, is thread-based. To be scalable, the new version is object-based.)

9. Member Company Benefits

Currently, **Kerena** is not sufficiently mature to attract potential partners from industry yet. Nevertheless, we anticipate tremendous attention from numerous industries such as content distribution networking companies (e.g., Akamai [20]) and peer-to-peer computing companies (e.g., SETI@HOME [21] and Entropia [22]) in near future.

10. References

[1] C. Shahabi, F. Banaei-Kashani. *Decentralized Resource Management for a Distributed Continuous Media Server*, IEEE Transactions on Parallel and Distributed Systems, Vol. 13, No. 6, June 2002.

[2] C. Shahabi, M.H. Alshayeji, S. Wang. *A Redundant Hierarchical Structure for a Distributed Continuous Media Server*. In Proceedings of the IDMS'97, September 1997.

[3] J.M.Carlson and J.Doyle, *Complexity and Robustness*, Proceedings of National Academy of Science of the USA, vol. 99, pp. 2538-2545, February 2002.

[4] D. Stauer and A. Aharony, *Introduction to Percolation Theory*, Taylor and Francis, second edition, 1992.

[5] F. Banaei-Kashani and Cyrus Shahabi, *Criticality-based Analysis and Design of Unstructured Peer-to-peer Networks as Complex Systems,* Third International Workshop on Global and Peer-to-Peer Computing (GP2PC) in conjunction with CCGrid'03 , Tokyo, Japan, May, 2003.

[6] Farnoush Banaei-Kashani, Cyrus Shahabi, *Brief Announcement: Efficient Flooding in Power-Law Networks*, Twenty-Second ACM Symposium on Principles of Distributed Computing (PODC'03) , Boston, Massachusetts, July, 2003.

[7] F. Banaei-Kashani, Ching-Chien Chen, and Cyrus Shahabi, *WSPDS: Web Services Peer-to-peer Discovery Service*, Submitted for review.

[8] Farnoush Banaei-Kashani and Cyrus Shahabi, *Searchable Querical Data Networks*, International Workshop on Databases, Information Systems and Peer-to-Peer Computing in conjunction with VLDB'03 , Berlin, Germany, September, 2003.

[9] Farnoush Banaei-Kashani and Cyrus Shahabi, *SWAM: A Family of Access Methods for Similarity Search in Querical Data Networks*, Submitted for review.

[10] D.J. Watts and S.H. Strogatz, *Collective Dynamics Of Small World Networks*, Nature, 393, 440-442, 1998.

[11] J. Kleinberg, *The Small-World Phenomenon: An Algorithmic Perspective*, Proceedings of the 32nd ACM Symposium on Theory of Computing, May, 2000.

[12] S. Ratnasamy and P. Francis and M. Handley and R. Karp and S. Shenker, *A Scalable Content-Addressable Network*, Proceedings of ACM SIGCOMM '01, August, 2001.

[13] S. Saroiu, P.K. Gummadi, and S.D. Gribble, *A Measurement Study of Peer-to-peer File-sharing Systems*, Proceedings of Multimedia Computing and Networking (MMCN'02), January 2002.

[14] E. Cohen and S. Shenker, *Replication Strategies in Unstructured Peer-to-peer Networks*, Proceedings of ACM SIGCOMM 2002, August 2002.

[15] A. Crespo and H. Garcia-Molina, *Routing Indices for Peer-to-peer Systems*, Proceedings of the 22nd International Conference on Distributed Computing Systems, July 2002.

[16] B. Yang and H. Garcia-Molina, *Efficient Search in Peer-to-peer Networks*, Proceedings of the 22nd International Conference on Distributed Computing Systems, July 2002.
[17] Q. Lv, P. Cao, E. Cohen, K. Li, and S. Shenker, *Search and Replication in Unstructured Peer-to-peer Networks*, Proceedings of the 16th International Conference on supercomputing (ICS '02), June 2002.
[18] L. Adamic, R.M. Lukose, A.R. Puniyani, and B.A. Huberman, *Search in Power-law Networks*, Physics Review Letters, , no. 64 046135, 2001.
[19] Q. Lv, S. Ratnasamy, and S. Shenker, *Can Heterogeneity Make Gnutella Scalable?*, Proceedings of the 1st International Workshop on Peer-to-Peer Systems (IPTPS '02), 2002.
[20] http://www.akamai.com/
[21] http://setiathome.ssl.berkeley.edu/
[22] http://www.entropia.com/