

Annual Report for Period:09/2007 - 08/2008

Submitted on: 05/29/2008

Principal Investigator: Tran, Duc A.

Award ID: 0753066

Organization: U of Massachusetts Boston

Title:

CSR-PDOS: A Scalable and Self-Organizing System Architecture for Fast and Efficient Information Retrieval in Large-Scale Decentralized Networks

Project Participants

Senior Personnel

Name: Tran, Duc

Worked for more than 160 Hours: Yes

Contribution to Project:

Post-doc

Graduate Student

Name: Nguyen, Khanh

Worked for more than 160 Hours: Yes

Contribution to Project:

Khanh Nguyen is a PhD student under my supervision at UMass Boston. His contribution in the project includes the JAVA-based implementation of two systems: the Zigzag system (a decentralized communication overlay prototype based on Zigzag) and the EZSearch system (a decentralized system built on top of Zigzag). This contribution led to a paper presented at the NSF Workshop on Next Generation Software co-located with IEEE IPDPS 2008 in Miami.

Name: Pamidimukkala, Salini

Worked for more than 160 Hours: Yes

Contribution to Project:

Salini Pamidimukkala was a graduate research student under my supervision when I was with the University of Dayton. She involved in the research and development of a relevance-feedback image retrieval system based on multiple instance learning. That was also the topic for her Master thesis successfully defended in April 2007. The result of her research (algorithms and system implementation) led to an IEEE conference publication.

Name: Ponnappalli, Surya

Worked for more than 160 Hours: Yes

Contribution to Project:

Surya Ponnappalli was a graduate research student working with me when I was at the University of Dayton. He investigated incentive mechanisms in decentralized networks. He gave a presentation on his findings in the PI's Special Topics course on Distributed Information Retrieval in Winter 2007 semester.

Undergraduate Student

Technician, Programmer

Other Participant

Research Experience for Undergraduates

Name: Igbinovia, Stephen

Worked for more than 160 Hours: Yes

Contribution to Project:

Stephen is a undergraduate student at UMass Boston. He is a McNair scholar student and as required by the McNair program, he needs to pursue research with a faculty member. He chose to work with me on this NSF project. His assignment is to explore how the findings in the project can be applied in a wireless decentralized network such as sensor networks. He helps implement a sensor network testbed for search purposes using SunSPOT sensor nodes. His work in the project is funded with a REU supplemental grant in summer 2008.

Years of schooling completed: Junior

Home Institution: Same as Research Site

Home Institution if Other:

Home Institution Highest Degree Granted(in fields supported by NSF): Doctoral Degree

Fiscal year(s) REU Participant supported: 2008

REU Funding: REU supplement

Organizational Partners

Other Collaborators or Contacts

Under the umbrella of the project, I had the following collaborations:

- I collaborated with Professor Tinh Nguyen (Oregon State) to investigate the problem of hierarchical multidimensional search in P2P networks. Our collaboration results in a research article published in Elsevier's Journal of Computer Communications in 2008.
- I co-authored with Professor Nguyen and Professor Sen-ching Cheung (University of Kentucky) on the problem of source-constraint P2P networks. This work results in a research article published in IEEE Transactions on Multimedia in 2008.
- I collaborated with Dr. Khanh Vu (University of Central Florida) on the problem of dimensionality reduction in hyperbolic data spaces. We co-authored a research paper that was accepted and presented at IEEE Conference on Computer and Information Sciences in 5/2008.
- I also collaborated with Dr. Junichi Suzuki (UMass Boston) in editing a Special Issue on Information Retrieval in Sensor Networks with the Int'l Journal on Parallel, Emerging, and Distributed Systems, and proposing an IEEE Workshop on the same topic with IEEE CCNC conference in 2009.

Activities and Findings

Research and Education Activities:

Education activities:

- Before moving to UMass Boston, I advised two graduate research students at the University of Dayton on this project, both funded by the grant. One of them graduated with a Master degree in Computer Science in May 2007. Her thesis was based on her findings on the project, which also resulted in an IEEE conference paper. I also delivered in Winter 2007 a Special Topics course on Distributed Information Retrieval, which summarized the state-of-the-art techniques on Distributed IR as well as our own findings on the topic. The feedback from the students was very positive.
- Now at UMass Boston, I have been advising three PhD students on this project, one of them funded by the grant. I have also been advising a undergraduate student on the project. This student is an under-represented minority student. I have obtained an REU supplemental grant to support him this summer.
- The techniques as a result of the project were included in the teaching materials of CS646 Computer Communications Networks that I currently teach at UMass Boston. In fact, I assigned to the Fall 2007 students a class project based on the research of the NSF project.

Research activities

- The project aims to develop a unified system infrastructure, that is scalable, self-organizing, and facilitates a comprehensive set of fast,

accurate, and efficient information retrieval services for the users of decentralized networks.

- In Year 1, we completed the designs for the following main components: (1) communication, (2) indexing, (3) search. Most of this research was focused on Internet-based decentralized systems.

- In Year 2 (this report), we focused on three main research items. First, we completed the implementation for the above 3 components. Second, we started to investigate the problem of publish/subscribe search in P2P networks. Third, we investigated this problem in the context of sensor networks (which is an important form of decentralized systems). Unlike the traditional distributed search, which has received a lot of attention, publish/subscribe is another form of information retrieval and enabling it in a decentralized network remains an open problem for research.

- Our research publications during 2007-2008 include one book chapter, one article with IEEE Trans. on Multimedia, one article with Elsevier Journal of Computer Communications, and five conference papers, including one accepted at IEEE ICCCN 2008 (26% acceptance rate) and selected as a Top 5 paper (out of 399 submissions) and a Best Paper Award candidate. We acknowledged the NSF grant in all these publications.

Findings:

The following summarizes our findings in Reporting Year 2007-2008: (Previous findings were already presented in the report for Year 2006-2007.)

1) Zigzag and EZSearch's System Implementation

As reported in Year 1, Zigzag and EZSearch are two key components of our decentralized information retrieval architecture. Zigzag is an efficient communication overlay architecture for decentralized systems. Zigzag organizes peers into a robust hierarchy of logarithmic height with constant-bounded node degree. Zigzag has been used as a reference technique for comparison in many publications. Built on top of Zigzag, EZSearch is our indexing framework. It is one of very few hierarchical designs aimed at decentralized search in high dimensionality. Other distributed search frameworks tightly couple the indexing architecture with the communication architecture. EZSearch is a unique approach because it decouples these architectures while still gracefully supporting their cooperation. Because of this decoupling, the dimensionality of the data does not affect the communication efficiency, which cannot be achieved by most other frameworks.

We have implemented both Zigzag and EZSearch. The implementation is programmed in Java, whose instructions and source codes are made available on the project's website. The source codes can be modified to customize the application of EZSearch to different decentralized search applications.

2) Publish/Subscribe in P2P Networks

Search applications can be categorized into two models: request/response and publish/subscribe (in short, pub/sub). In the former, a search query is submitted on demand expecting the results to return immediately; if they do not exist, a response indicating so is returned. Contrarily, a query in the pub/sub model is submitted and stored in advance. The results may already exist, which will be returned immediately; otherwise, the query subscriber will be notified when the matching results later become available. Thus, a main problem for request/response search systems is to manage existing data in advance for fast data retrieval at a later time, while the corresponding problem for pub/sub systems is to manage the queries in advance for fast query matching in the future.

We are interested in the pub/sub service deployed in a P2P network. We particularly work on the case the network overlay is structured according to the CAN technique, one of the well-known Distributed Hash Tables (DHTs) designs for self-organizing and scalable P2P networks. In CAN, each node is given a unique identifier that represents a non-overlapping rectangular zone partitioned from a virtual multi-dimension space, called the CAN space. For routing purposes, each node has a list of neighbor nodes whose zones are adjacent to its zone. Consequently, the number of neighbors per node is proportional to the dimension of the CAN space, which is independent from the network size (i.e., the number of nodes in the network). Although the other DHT techniques have a node overhead increasing with the network size, CAN is favored over them only when the CAN dimension is low.

The motivation behind CAN is distributed search applications that follow the request/response model. Deploying a pub/sub service on top of CAN is not as straightforward. From the database perspective, because we typically model a data object as a point and a query as a range of points, we need to address the range indexing problem in pub/sub systems, which is more challenging than the point indexing problem as in request/response systems.

From the networking perspective, due to its range, a subscription query may be replicated at multiple nodes to wait for notification of all

possible matching data objects. Hence, the number of subscriptions stored in the network can be large, resulting in not only the communication cost to replicate the subscriptions, but also high storage cost for each node and long time to match an object against a subscription query. We need to minimize unnecessary replications, yet at the same time store the queries in the network intelligently so that data notification remains efficient. Such issues are not present in request/response systems.

Because of the low dimensionality of the CAN space, another challenge to a CAN-based pub/sub system is due to the mismatch between the CAN dimension and the data dimension. Data can be, and usually, of high dimension, such as in applications searching documents, multimedia, and sensor data, which normally are associated with many attributes. It is difficult to hash similar high-dimension data objects into zones in a low-dimension space which are close to each other, making the search for a continuous range of data highly inefficient. For the request/response model, in which data need to be preprocessed and queries are subsequent and on-demand, dimension mismatch can be resolved effectively by locality-sensitive dimensionality reduction techniques.

For the pub/sub model, because subscription queries are submitted in advance, we need to store similar subscriptions in nodes of close proximities so as to make the data notification process efficient. While the aforementioned reduction techniques apply only to singular points whose similarity can be defined by a distance metric, this metric is not applicable to queries which are ranges of points.

We explore the application of Random Projection in pub/sub networks that are specifically built on CAN. In specific, we have proposed an efficient technique that uses Random Projection to map the pub/sub data and queries into a CAN space of lower dimension. We have also devised a strategy aimed to avoid replicating subscriptions, guaranteeing that any query is matched with all possible data objects using an efficient notification process.

The result of this work is encouraging. We have published two full conference papers, including one full paper selected by the 17th IEEE Int'l Conference on Computer Communications and Networks (ICCCN 2009) as a Top 5 Paper and a Best Paper Award candidate.

3) Publish/Subscribe Search in Sensor Networks

Publish/subscribe services are natural applications of sensor networks as sensors are designed mainly to detect and notify upon events of interest. For example, in the application of disaster monitoring, a subscriber can submit to the network in advance a query specifying the early warnings of a wildfire, so that as soon as an event matching these warnings is detected by a sensor, the event is propagated to the network to notify the subscriber. In another example, we can help our soldiers navigate safely on a battlefield by deploying a sensor network that is notified upon events implying the existence of enemy forces.

Research on publish/subscribe mechanisms for sensor networks remains very challenging. It is understandable, though. Firstly, unlike publish/subscribe systems on the Internet, that designed for sensor networks faces unseen challenges due to limitations in sensor storage, processing, and communication capacities. Secondly, the publish/subscribe model is much different from the traditional search model which has widely been addressed for sensor networks.

For any publish/subscribe service in a sensor network, it is desirable that the following costs are low: the communication cost to disseminate a subscription query or event, the storage cost to replicate queries in the network, and the computation cost to check query/event matching conditions. In addition, it should take short notification delay; i.e., short time between when an event is published until the publication reaches its subscribing nodes.

However, the time requirement and the cost requirement are conflicting with each other. Indeed, to increase the chance for an event to meet its matching queries quickly, the query should be replicated at an sufficiently large number of nodes, which, however, results in large communication, storage, and computation costs. In the project, we focus on reducing the costs. The main approach to lowering these costs is to use a GHT method to map a query and an event to the same geographic coordinate of the sensor network if they match. Although GHT is well-known, its application for publish/subscribe services in sensor network has been proposed only for the case where queries are point queries not range queries. This is because when GHT is used to hash a range query, the corresponding geographic area can be large; hence, expensive replication and computation costs.

We have devised a GHT-like solution that works efficiently for multidimensional range queries. This solution is a modified version of our publish/subscribe technique for P2P networks. Our hashing method is based on random projections to the sensor location space. To reduce the replication cost, many publish/subscribe techniques take advantage of subscription coverings to remove unnecessary replications. However, not many coverings can be found in high dimension. We instead use the covering relationship among their projections, rather than the original-subscription covering relationships. To reduce the computation cost, the queries are stored in the network such that, given an event, the computation checking the matching conditions takes places only at nodes highly likely of storing the matching queries.

Another publish/subscribe problem we address for sensor networks is how to enable publish/subscribe services without knowing any information about sensor node locations. This problem is unique because techniques in the literature all assume that a sensor node knows its location. We have found a solution that consists of two components: (1) a network structure that assigns to each sensor node a virtual label (a binary string) such that the node labels form a prefix tree, and (2) a hashing method to map a query or event to a node on the tree. Using this solution, a query is replicated at only one node and an event finds all of its matching queries by following the path toward the root of the tree. The communication cost to disseminate a query or event is $O(\log\{n\})$ and the notification delay is also $O(\log\{n\})$.

The results of our work on publish/subscribe in sensor network are included in two papers, one under review and one under preparation.

Training and Development:

Student training and development is an important goal of our educational program. We have graduated 1 MS thesis and supported with the NSF grant 1 PhD student, 2 MS students, and 1 undergraduate student. Thus far, two female students and one under-represented minority student have contributed to the project.

This year alone, although the grant (with REU supplemental) provides support for only one graduate student and one undergraduate student, using my other sources of funding, I have been able to have three other graduate students participating in the project. Their research skills have been improved significantly as the project provides them with access to a state-of-the-art laboratory and significant research materials. We have a group meeting every week in which students present their in-progress work in the form of a research seminar. Two graduate students published papers this year. Three graduate students gave research presentations in my CS648 Wireless Networks class in Spring 2008. We expect one MS thesis next year based on the research on the project. The undergraduate student has expressed his plan to pursue graduate study after graduation. He is working in Summer 2008 under the support of an REU supplemental grant. These students' activities in the project have been helpful to them in terms of personal and professional development.

Outreach Activities:

- This NSF project was featured on Campus Report and UD Quarterly, two official periodical publications of the University of Dayton. These publications are delivered to all current and past members of the University of Dayton (with more than 90,000 living alumni).

- I presented a report on our NSF project at IPDPS 2008 - NSF Workshop on Next Generation Software in Miami (4/2008).

- I gave a talk highlighting the project at the UMass Boston College of Sciences and Mathematics faculty summit in Fall 2007

- I gave two seminar talks on the findings of the NSF project at the UMass Boston Computer Science colloquium, in Spring 2007 and Fall 2007, respectively.

- I promoted the research on information retrieval in decentralized networks, especially sensor networks, by collaborating with Dr. Junichi Suzuki (UMass Boston) to edit a Special Issue on Information Retrieval in Sensor Networks with the Int'l Journal on Parallel, Emerging, and Distributed Systems (publication in 2009), and propose an IEEE Workshop on the same topic with IEEE CCNC 2009.

Journal Publications

Tran, D. A. and Nguyen, T., "Hierarchical Multidimensional Search in P2P Networks", Journal of Computer Communications, p. 346-357, vol. 31(2), (2008). Published,

Nguyen, T., Kolazhi, K., Kamath, R., Cheung, S. , & Tran, D. A., "Efficient Multimedia Distribution in Source Constraint Networks", IEEE Transactions on Multimedia, p. 523-537, vol. 10(3), (2008). Published,

Books or Other One-time Publications

S. Pamidimukkala, "Development of a Relevance-Feedback Image Retrieval System based on Multiple Instance Learning", (2007). Thesis, Published

Bibliography: Master Thesis, Department of Computer Science, University of Dayton

Tran, D. A., & Nguyen, T., "Publish/Subscribe Techniques in Sensor Networks", (2009). Book Chapter, Accepted

Editor(s): S. Pierre

Bibliography: In "Encyclopedia of Next-Generation Mobile Networks and Ubiquitous Computing"

Tran, D. A., & Nguyen, T., "Publish/Subscribe Service in CAN-based P2P Networks: Dimension Mismatch and the Random Projection Approach", (2008). Conference Proceedings, Accepted

Bibliography: In Proceedings of IEEE Int'l Conference on Computer Communications and Networks

Tran, D. A., & Nguyen, K., "Multidimensional Information Retrieval in P2P Networks", (2008). Conference Proceedings, Published

Bibliography: In Proceedings of NSF Workshop on Next Generation Software - IEEE IPDPS 2008

Tran, D. A., Pamidimukkala, S., & Nguyen, P., "Relevance-Feedback Image Retrieval based on Multiple-Instance Learning", (2008). Conference Proceedings, Published

Bibliography: In Proceedings of 2008 ACIS/IEEE Conference on Computer and Information Science

Tran, D. A., & Vu, K., "Dimensionality Reduction in Hyperbolic Data Spaces: Bounding the Reconstructed Information Loss", (2008). Conference Proceedings, Published

Bibliography: In Proceedings of 2008 ACIS/IEEE Conference on Computer and Information Science

Tran, D. A., & Nguyen, T., "A Random Projection Approach to Subscription Covering Detection in Publish/Subscribe Systems", (2007). Conference Proceedings, Published

Bibliography: In Proceedings of 2007 IEEE Int'l Conference on Collaborative Computing

Web/Internet Site

URL(s):

<http://www.cs.umb.edu/~duc/research/nsf0615055/>

Description:

Other Specific Products

Product Type:

Software (or netware)

Product Description:

Zigzag Implementation in Java: This is a Java implementation of the Zigzag design. Once we run all peer instances, they will be connected according to the Zigzag hierarchy. We also include a video broadcasting component with this implementation to illustrate the use of Zigzag. The video is produced first by the server and follows the Zigzag hierarchy to reach all the peers in the network. Thus, there are two main components in the software development: the Zigzag component and the Streaming component. All the programs are written in Java. Streaming is enabled by calls to VideoLan (aka., VLC).

Sharing Information:

The source codes for the software are free to the research community. The project website provides a link to this software.

Product Type:

Software (or netware)

Product Description:

EZSearch:

This is the first implementation of a hierarchical multidimensional search system based on Zigzag -- a P2P overlay architecture known for its scalability and robustness under network growth and dynamics. The indexing architecture of EZSearch is built on top of the Zigzag hierarchy, that allows both k-nearest-neighbor and range queries to be answered with low search overhead and worst-case search time logarithmic with the network size.

Sharing Information:

The source codes for EZSearch are free to the research community. The project website has a link to this software.

Contributions

Contributions within Discipline:

- We publicize the implementation source codes of Zigzag and EZSearch on the project website so that other researchers can use or modify for their own applications. Zigzag/EZSearch is the first architecture that decouples the design of the communication overlay and that of the indexing. Thus, it can be used to run multiple decentralized search services on the same (possibly already existing) communication overlay.
- Our technique for CAN-based P2P publish/subscribe services is the first: (1) that can be used to deploy a publish/subscribe application on top of an existing CAN network, and (2) that is designed for high-dimensional publish/subscribe services in a low-dimension CAN network.
- Our publish/subscribe technique for sensor networks without location information is the first efficient technique of this kind. All other techniques assume that sensor node locations are known.

Contributions to Other Disciplines:

- The Zigzag prototype software can be used for broadcast/multicasting services such as video conferencing and IPTV broadcasting.
- Because of the applicability of sensor networks in many fields such as environmental monitoring, homeland security, and industrial safety, where notification services are highly important, our results on publish/subscribe services in sensor networks can make a direct contribution. Using our solution, we can save costs and at the same time improve the service efficiency.

Contributions to Human Resource Development:

This year, the project provides research and training opportunities to 3 PhD students, one female MS student and an African-American undergraduate student. Among them, one is fully funded by the NSF grant and one by a REU supplemental grant.

Contributions to Resources for Research and Education:

The project has a web site at <http://www.cs.umb.edu/~duc/research/nsf0615055>, which allows electronic access to our publications, talks, teaching materials, and software downloads relevant to this project. The software is free and open for the research community to use.

Contributions Beyond Science and Engineering:

Special Requirements

Special reporting requirements: None

Change in Objectives or Scope: None

Animal, Human Subjects, Biohazards: None

Categories for which nothing is reported:

Organizational Partners

Contributions: To Any Beyond Science and Engineering