CS 856 Paper Reviews
W Anthony Young - 20161423
October 17, 2004

Paper Title: Tapestry
Section I. Overview

A. Reader Interest

1. Which category describes this paper?

( ) Practice / Application / Case Study / Experience Report
(X) Research / Technology
( ) Survey / Tutorial / How-To

2. How relevant is this paper to the readers in the area? Please explain your rating under III. Detailed Comments.

( ) Very Relevant
(X) Relevant
( ) Interesting - but not very relevant
( ) Irrelevant

B. Content

1. Please explain how this paper advances this field of research and / or contributes something new to the literature. Please explain your answer under III. Detailed Comments.

2. Is the paper technically sound? Please explain your answer under III. Detailed Comments.

(X) Yes
( ) Appears to be - but didn't check completely
( ) Partially
( ) No

C. Presentation

1. Are the title, abstract, and keywords appropriate? Please explain your answer under III. Detailed Comments.

(X) Yes
( ) No

2. Does the paper contain sufficient and appropriate references? Please explain your answer under III. Detailed Comments.

(X) References are sufficient and appropriate
( ) Important references are missing; more references are needed
( ) Number of references are excessive

3. Does the introduction state the objectives of the paper in terms that encourage the reader to read on? Please explain your answer under III. Detailed Comments.

(Y) Yes
( ) Could be improved
( ) No

4. How would you rate the organization of the paper? Is it focused? Is the length appropriate for the topic? Please explain your answer under III. Detailed Comments.

(X) Satisfactory
( ) Could be improved
( ) Poor

5. Please rate and comment on the readability of this paper.
Please explain your answer under III. Detailed Comments.

( ) Easy to read
(X) Readable - but requires some effort to understand
( ) Difficult to read and understand
( ) Unreadable

Section II. Summary and evaluation

A. Summary (provide here a summary of the paper)

This paper presents an overview of the Tapestry system: a distributed location and routing protocol. Tapestry is built with the intention of routing messages and data in the presence of an unstable network. The authors present related work in the field of peer networks, discussing systems such as Gnutella and CAN.

Tapestry uses datagram-style communication between nodes to route messages. Objects are stored at local nodes and referenced using Global Unique ID's (GUID's). Nodes maintain a list of neighbour links and IP addresses, and neighbour node ID's. Messages are forwarded to the neighbour with the node ID closest to the node we are trying to communicate with. Nodes query the network by routing messages requesting object O. When a server S has the object, it returns a message back to the requester. Each intermediate server caches the (O, S) mapping for future queries. The authors present algorithms for node insertion, node deletion and node failure. Further, a discussion of the component architecture and functions of components is presented in detail.

The authors present some evaluations of their work. An implementation of Tapestry was tested on the PlanetLab system, which is under constant load from other users, etc. Thus, they believe that their tests accurately reflect how Tapestry would behave under "real world" conditions. The results presented are fairly good, with most delays being attributed to node virtualization and JVM problems. Tapestry appears fairly resilient to massive node failures, allowing the network to quickly recover from node losses as well as route messages in the face of failures.

As the paper winds down, the authors briefly discuss some applications that are implemented on the Tapestry architecture, as well as how well Tapestry works with the internet and the challenges posed by it for security, etc. The paper concludes with a short summary.

B. Evaluation

Please rate the paper. Please explain your answer under III Detailed Comments.

( ) Award Quality
(X) Excellent
( ) Good
( ) Fair
( ) Poor

Section III. Detailed Comments

This paper is quite relevant to researchers in the area of overlay networks and lookup routing protocols. This paper presents a new system and protocol for routing lookups between nodes in an overlay. The paper provides many implementation details, including a solid technical discussion of routing protocols and system components. Some questions remain, such as:

-Can anyone un-publish a listing? Doesn't this make the data particularly vulnerable to attack?
-What is "republish traffic"? This is not defined. What characterizes it and how does it work?

As well, a more detailed description of the tradeoffs considered with the amount of space required for the routing table is desirable. Doesn't the neighbour table take up a lot of space? It seems like there will be a significant number of entries in each table. Also, don't we use a lot of space caching mappings between GUID's and servers? I would think this should be discussed.

The paper's introduction presents motivation for development of Tapestry, as well as a brief overview of the Tapestry system. This is enough to get the reader excited about a new routing protocol. As well, the abstract and keywords provide a concise synopsis of the contents of the paper, touching on many pieces of information contained in the sections of the paper.

Sections of the paper are well focused on one specific topic. Authors provide some transition between topics, but a discussion of how some topics interrelate would have been useful. Also, some discussions were slightly hard to understand and required re-reading several times to grasp. Some care should be taken to present the paper in simpler terms.

Overall, I believe this paper is very well written, and contributes some novel routing mechanisms to the field.

Paper Title: Pastry

Section I. Overview

A. Reader Interest

1. Which category describes this paper?

( ) Practice / Application / Case Study / Experience Report
(X) Research / Technology
( ) Survey / Tutorial / How-To

2. How relevant is this paper to the readers in the area? Please explain your rating under III. Detailed Comments.

( ) Very Relevant
(X) Relevant
( ) Interesting - but not very relevant
( ) Irrelevant

B. Content

1. Please explain how this paper advances this field of research and / or contributes something new to the literature. Please explain your answer under III. Detailed Comments.

2. Is the paper technically sound? Please explain your answer under III. Detailed Comments.

(X) Yes
( ) Appears to be - but didn't check completely
( ) Partially
( ) No

C. Presentation

1. Are the title, abstract, and keywords appropriate? Please explain your answer under III. Detailed Comments.

(X) Yes
( ) No

2. Does the paper contain sufficient and appropriate references? Please explain your answer under III. Detailed Comments.

(X) References are sufficient and appropriate
( ) Important references are missing; more references are needed
( ) Number of references are excessive

3. Does the introduction state the objectives of the paper in terms that encourage the reader to read on? Please explain your answer under III. Detailed Comments.

(X) Yes
( ) Could be improved
( ) No

4. How would you rate the organization of the paper? Is it focused? Is the length appropriate for the topic? Please explain your answer under III. Detailed Comments.

(X) Satisfactory
( ) Could be improved
( ) Poor

5. Please rate and comment on the readability of this paper.
Please explain your answer under III. Detailed Comments.

( ) Easy to read
(X) Readable - but requires some effort to understand
( ) Difficult to read and understand
( ) Unreadable

Section II. Summary and evaluation

A. Summary (provide here a summary of the paper)

This paper presents the Pastry peer routing protocol. Pastry nodes organize themselves into a locality-aware overlay network to route messages. Each node has a random nodeID.

Messages with keys are routed to the numerically closest node to the key by forwarding the message through intermediate nodes that have greater and greater commonalities with the key we are searching for. Each node has a routing table, leaf set and neighbourhood set of nodes. The routing table is used to route messages through a network. On average, $\log_2^b(N)$ entries are required for a network with N nodes (b is a tuning parameter). The neighbourhood set contains a list of those nodes that are spatially closest to the peer and is used to maintain locality. The leaf set contains the peers that have numerically closest nodeID's to this node. Routing algorithms, the Pastry API, node join and leave algorithms, and failure recovery algorithms are all presented.

The authors present a very large analysis of their protocol. Pastry is capable of routing over 3000 messages per second (with their experimental setup). As well, only $\log_2^b(N)$ nodes need to be contacted to route a message to its destination in the worst case. Pastry is also fairly robust to failures, but does work well in situations where nodes are acting maliciously. More importantly, Pastry scales well both in state space and routing efficiency.

The paper concludes with a discussion of related work and a brief outline of the contents of the paper.

B. Evaluation

Please rate the paper. Please explain your answer under III Detailed Comments.

( ) Award Quality
( ) Excellent
(X) Good
( ) Fair
( ) Poor

Section III. Detailed Comments

This paper presents an overview of the Pastry protocol. This protocol is relevant to any group studying overlay networks or routing in peer systems. The novel idea presented is of a self-organizing network overlay that takes into account the physical location of nodes in the network. To the author's knowledge, no such system presently exists. This paper advances the field by providing a protocol that makes more efficient use of physical network topologies than a traditional peer system.

The authors provide an abstract that briefly summarizes the contents of the article. However, no keywords were presented. These should be added to aid in content-based searches and to provide more information to the reader. The introduction hooked this reader and made him want to read about this novel new method of overlay organization. The authors' presentation appears to be technically sound. However, some information was missing:

-How did the authors choose the values of L and b? There is no discussion of how these values were determined.
-What happens if delivery fails (i.e. a message is routed to a node that has failed, etc.)? Does the sender receive a notification that their query failed?

The authors present their paper well. Sections are well organized, and there is a sense of flow between subsections. However, some missing words and grammar mistakes make certain sentences hard to follow. As well, some join and repair algorithms are hard to follow, as was the initial discussion of the routing protocol (section 2.2 was much clearer). The authors should have presented these concepts more clearly. A complete set of references was provided, and it appears that the authors took the time to perform background research before writing their paper.