Evaluation of Peer-to-Peer Overlays for First Response

Dirk Bradler
TU Darmstadt
Darmstadt, Germany

Jussi Kangasharju
University of Helsinki
Helsinki, Finland

Max Mühlhäuser
TU Darmstadt
Darmstadt, Germany

Abstract

Efficient communication architectures are vital for handling larger scale first response scenarios, and existing mechanisms have several shortcomings due to the heterogeneity of first response groups. In this paper, we consider using peer-to-peer-based communication architectures for first response scenarios. We evaluated different network overlays in environments found in first response -like scenarios. A first response P2P system needs to address both, reliability in high churn situations in network infrastructure mode as well as reasonable usability in wireless ad-hoc networks. Our evaluation shows that a superpeer architecture based on peer capabilities has superior performance when compared with purely unstructured or DHT-based overlays.

1 Introduction

Peer-to-peer (P2P) technologies are already established in several application domains, e.g. IP telephony and file sharing. In file sharing, P2P technology removes the necessity for central storage space and saves bandwidth, since file transfers are performed directly between peers. The self organizing capabilities of P2P enable fast deployment without the need for manual configuration of individual peers. Especially in an unreliable and fast changing environment with limited infrastructure, P2P technology can be used as a basis for communication infrastructures.

In this paper we consider a first response communication system, which is intended to be used after the existing communication infrastructure has been heavily damaged, e.g., after an earthquake or a hurricane. The need for reliable first response communications was one of the findings of the 9/11 attack on the Pentagon [9]. Using a P2P-based communication approach, we are able to relieve the overloading of communication channels and improve interoperability. In order to get realistic requirements and address the needs of first response helpers, we have analyzed findings of disaster reports and interviewed first response helpers [2]. The main contribution of this paper is an evaluation of different P2P overlay networks in terms of their suitability for first response communications. We consider aspects such as resistance to churn and reliability of communications in our evaluation, and find that a superpeer architecture based on geographical proximity strikes the best tradeoff between these two properties.

This paper is organized as follows. In section 2 we present the needs of first response communications. Section 3 presents the evaluated overlays. Section 4 presents an evaluation of our system. Section 5 discusses related work and Section 6 concludes the paper.

2 Scope and Scenario

We present an overview of the basic requirements for a P2P first response communication system. We have analyzed the different use cases and requirements in our earlier work [2]. The requirements are based on reports of first response cases as well as extensive telephone interviews with first response professionals. The main use of P2P communications in first response is in organizing on-site communications in situations where most of all of the existing infrastructure has failed. The on-site teams need a reliable group communications which support the organizational hierarchy of the response teams. Communication takes place over wireless links and between small handheld devices carried by the personnel. The networks are typically augmented with access points or more powerful communication nodes (e.g., communication vans). Our goal is to augment the classical first response communication approach, which is often supported by foot messengers and replace these with a P2P-based communication system.

3 Overlay Networks

In contrast to Internet-based P2P networks, a communication system for on-site communications for first

0-7695-3113-X/08 $25.00 © 2008 IEEE
DOI 10.1109/PERCOM.2008.14

response has a relatively small amount of users, on the order of hundreds or a few thousand at most. First response communication needs to deal with heterogeneous devices, small PDAs with limited bandwidth and limited CPU power, as well as highly available servers offered by operation control. Current first response communications and planned future systems are based on wireless LAN or WIMAX [5,10].

However, we cannot assume a wireless access point infrastructure right after a catastrophic event. Therefore, the network must also support a wireless ad-hoc mode at least until a more stable communication infrastructure is available. In a typical first response case, the personnel carries small, handheld devices. Even sensor technology might be used. In addition, more powerful equipment is also present, e.g., at the command post. These nodes generally have much more power and bandwidth available. In addition these fixed spots are able to provide reliable wireless LAN infrastructure which greatly improves transfer rates.

In this paper, we consider four different overlay architectures: a fully unstructured network, a network based on a distributed hash table, and two kinds of superpeer networks, one with random superpeer selection and another with superpeer selection based on actual peer capabilities. Our goal is to identify the strengths and weaknesses of each of these overlays for the particular needs of a first response scenario. The expected results are to find a reasonable range in which the different P2P overlays perform as intended and to find the borders at which an overlay becomes unusable.

The unstructured network is a randomly connected network where each peer has 3 neighbors. We evaluated two superpeer network configurations. In the first, all peers had the same resources and some were selected as superpeers. In the second, we gave some peers higher uptimes and selected them as superpeers. For the DHT, we selected Tapestry [12] (the results apply for other DHTs as well). We initially assume that the actual network is able to provide fully reliable ad hoc routing from any node to any other node. This lets us evaluate the resistance to churn. We revisit the assumption about reliable ad hoc routing in Section 4.4.

4 Evaluation

We now present the results from our evaluation, which are averaged over 20 simulation cycles. All four evaluated connection topologies were simulated using Peersim (see http://peersim.sourceforge.net) and perform perfectly in the case of no churn. In each case, we simulated the disconnection of some fraction of peers, and varied this fraction from 0 to 1. The failed nodes were randomly picked. Analyzing the resulting network graph, we determined the number of clusters and the size of the largest cluster. These are vital metrics for ensuring reliable delivery of critical messages in a first response scenario. If there is no path between two nodes, there is no possibility to send a message. If a path exists, it should be possible to forward all messages to their right recipients under the assumed reliable ad hoc routing (see Section 4.4). The goal of our first evaluation is to compare the different overlay networks and determine which of them performs best at different node failure ratios.

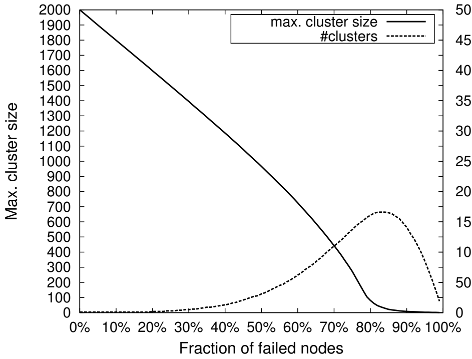

Figure 1. Clustering in unstructured network

4.1 Unstructured Network

Figure 1 shows the results for the unstructured network case. The network consisted of 2000 peers. X-axis shows the percentage of failed nodes, left y-axis the size of the largest connected component (cluster), and right y-axis shows the number of clusters. The random topology is created with an average node degree of 3 and non-directed connections. Advantage of a random topology is the low clustering factor. Even with a 30% rate of failure, only 2.5% of all remaining peers are not reachable. Further increasing churn rate, shows complete collapse of the topology occurs at 70%–80% failure rate.

4.2 Superpeer Network

We consider two cases for the superpeer network. In the first case (random select), all nodes are homogeneous and some of them are designated as superpeers. The second case, sensitive select, simulates a typical first response scenario after the first larger wave of responders have arrived with more communications equipment. In this case, we select a small subset of peers, make them more robust against failure, and designate them as superpeers. In both cases we assume

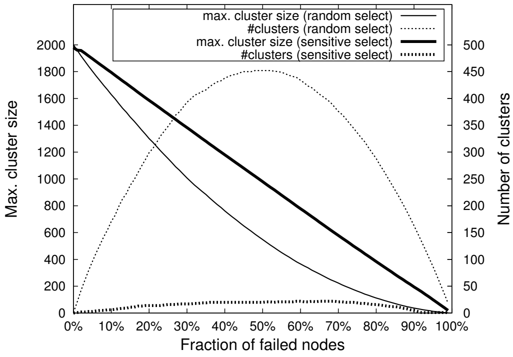

Figure 2. Clustering in random and context sensitive superpeer selection

that the superpeers form a full mesh between them, but other network topologies are also possible [3]. Selecting good peers as superpeers can be performed online by the peers themselves. The details are outside the scope of this paper.

From the 2000 peers, we designated 200 peers as superpeers and the rest of the peers were evenly distributed between the superpeers, with each peer connected to only 1 superpeer. (Connecting to several superpeers in range might be desirable in a real-world deployment.) As mentioned above, the superpeers form a full mesh between themselves. In the first case, all nodes had equal failure probabilities. In the second case, more powerful peers as superpeers, the superpeers had a failure probability which was 1/30 of the failure probability of the normal peers. Note that it is still possible for superpeers to fail; it is just much less likely than for normal peers. Since we evalute a worst case scenario, we did not provide backup links for client peers, i.e. if a superpeer fails all connected clients are considered as disconnected.

Figure 2 shows the size of the largest cluster and number of clusters for the two superpeer cases. The thin lines show the random select case and the thick lines show the sensitive select case where superpeers had longer uptimes.

The sensitive select strategy, i.e., selecting superpeers with better uptimes, improves performance significantly. The size of the largest cluster is almost the same as the number of remaining peers, meaning that in most cases, the complete network remains connected. The number of clusters is also minimal. In contrast to the randomly selected superpeers, the number of clusters is cut by a factor of 10.

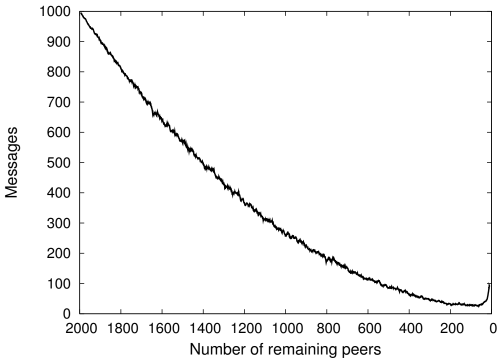

Figure 3. Message reachability in Tapestry

4.3 DHT

We also investigated the Tapestry DHT in the same scenario. We observed that in all of the cases, the Tapestry network remains connected, thus in principle every message should always reach its intended recipient. We evaluated the successful messages by choosing 1000 times a random pair of peers and letting one of them perform a lookup of the other. The results are shown in Figure 3. X-axis plots again the fraction of failed nodes and y-axis shows the fraction of messages which reached their intended recipient. Note that the sender and receiver were chosen from the set of nodes that were still alive after the failure. In other words, 100% reachability was theoretically always possible.

Our results show a message success rate of less than 30% if about 50% of peers fail. In other words, even though a graph of Tapestry is not clustered, delivery rate of messages dramatically decreases even at medium churn rates. The main reason for the low success rate is that Tapestry relies on certain routing table entries for forwarding messages. Note that we did not run any topology maintenance in any of the overlays after the failure. Each of the overlays is able to recover full connectivity after the usual periodic maintenance, but the costs are different. Unstructured and superpeer networks have relatively low costs, but Tapestry needs to rebuild its routing tables, which presents a considerable overhead.

Although our evaluation was made with Tapestry, similar results could be expected from any other DHT, with small differences in terms of connectivity and reachability. However, all DHTs are based on the assumption that any peer is able to send a message to any other peer. Under the reliable ad hoc routing assumption from above this is feasible, but under more realistic conditions, this assumption is no longer valid. Next we will focus on evaluating the impact of unreli-

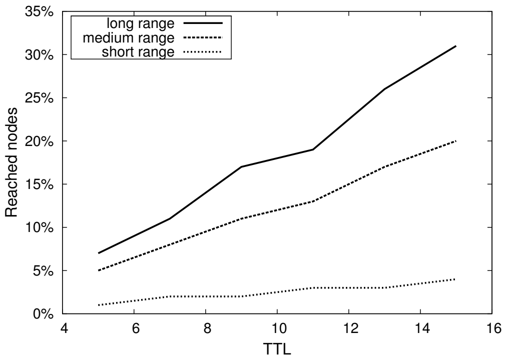

Figure 4. Reached nodes in mobile scenario

able ad hoc routing on the overlay performance.

4.4 Routing Reliability

Our goal is to evaluate the chance one peer might have on a square shaped field with a fixed edge length and randomly placed peers to talk to arbitrarily selected nodes. Note that this evaluation is completely independent from the chosen overlay.

We model the probability of a successful one-hop communication as in [11], and the probability of a message being sent successfully over a distance x is $p(x) = \frac{x^2}{D^2} - \frac{2x}{D} + 1$, where D is the theoretical maximum transmission range. This model corresponds with already observed behavior of TCP performance in mobile ad-hoc networks [4]. Our network distance model is applied for each hop in the underlay separately.

Figure 4 shows the fraction of nodes *reached* as a function of the number of hops the message is forwarded over. There were a total of 1000 nodes scattered uniformly over a square. The three ranges refer to the values of D which were set at 100% of the square's diagonal for "Long", 50% for "Medium", and 25% for "Short". These different ranges reflect different kinds of devices and different sizes of areas.

In general, only a very small fraction of nodes can be reached, because the probability drops quadratically as a function of the distance. This implies that any overlay which relies on communicating with arbitrary nodes is likely to fail. A standard random graph, such as the one evaluated in Section 4.1 requires a node to be able to pick its neighbors freely, i.e., communicating with arbitrary nodes. However, such an unstructured network could also be constructed using only connections to immediate neighbors, although the properties of the overlay topology would likely no longer be as good as those of a truly random graph. However, the network would be able to function.

A DHT-based network also requires communications with arbitrary nodes and would also likely fail. In contrast to an unstructured network, a DHT has strict rules on how nodes choose their neighbors, and thus it is likely that a DHT could not be made to work on top of an unreliable ad hoc network. Note that this is not specific to Tapestry; any DHT has the same problem.

A superpeer network with sufficiently powerful superpeers is able to handle the communications and the only potential problem is that the communications between normal peers and superpeers still follow the unreliable probability model. Assuming there are enough superpeers, the overlay network should be relatively immune to the underlay. If there are not enough superpeers, then a hybrid solution where normal peers outside the range of a superpeer use other normal peers as intermediaries to reach a superpeer seems appropriate. Such a network would combine the properties (good and bad) of superpeer and unstructured networks.

4.5 Summary of Results

Both the unstructured network and Tapestry are very good at keeping the network connected, i.e., minimizing the number of clusters. The superpeer network with correct superpeer selection is able to get similar performance. However, both the unstructured network and Tapestry require a fully reliable ad hoc network to deliver the service, and as the evaluation in Section 4.4 shows, this will severely impact their performance.

In Table 1 we summarize the behavior of the overlay networks as a function of the fraction of failed nodes. We consider two cases: *operational* where 70–100% of the messages are routed to the correct recipients, and *failed* where less than 50% of the messages arrive correctly. For each overlay network, we show how many nodes in the overlay are allowed to fail for the network still to meet the performance targets. The reason for failure is that either the network has clustered or is unable to find a correct routing path (DHT). The unstructured topology performs better than randomly elected superpeers (SP/Random); the failure occurs ar 75% node failure, while in the superpeer case a 50% is sufficient to push the network to a failed state. This is because when a superpeer fails, all its normal peers become individual clusters. Tapestry, as discussed in Section 4.3 has very poor performance. All values are derived from the measured clustering which is shown in figures 1, 2 and 3. For example in the unstructured case 72% of the nodes (1440 peers) need to fail in order to keep 70% of the remaining nodes (392 peers) operational, i.e. 30% (168 peers) of the network are not part of the giant component. If superpeers are selected

Topology	Operational	Failed
Superpeer	Always	Never
SP/Random	27% failed	50% failed
Unstructured	72% failed	76% failed
DHT	15% failed	30% failed

Table 1. Effect of node failures

based on their uptime the network stays always in operational state. This is because of two reasons. First, as shown in Figure 2, there is only minimal amount of clustering. Second, the superpeer network does not suffer from the problems of ad-hoc routing (Section 4.4) and is thus able to keep good performance.

In summary, good selection of superpeers yields significant performance gains over normal unstructured networks. Our results also show that simply selecting superpeers is not sufficient to obtain good performance.

5 Related Work

Another P2P requirement analysis is done by University of Virginia [1]. They identified three main issues in current first response approaches developed a prototype for a P2P based first response solution. Nevertheless their approach focuses mainly on usability for the first response team, while our contribution evaluates the underlying technical approach. Further implementation of their P2P solution [1] is done using hypercast, GPS capabilities, multicast streaming video and access control mechanism. Nevertheless the goal was to develop a prototype implementation, large scale simulations were not conducted. DHTs have in general found to be unsuitable for dynamic environments. In particular, at higher churn rates DHTs become completely unusable [8]. Churn rate behavior is both theoretically examined [7] and simulated. As our results on Tapestry show, DHT performance drops rapidly with even small amounts of churn and thus it is questionable whether a DHT-based communication system is able to meet the requirements of first response scenarios. Our evaluation shows that superpeer networks are more robust under churn. Solar [3] is a middleware layer for transporting context information in emergency response scenarios. Solar is based on a superpeer architecture where the superpeers are connected over the Pastry [6] DHT. Although the superpeers have more transmission power, their interconnection network is still subject to the effects from Section 4.4, and it is not clear that a DHT is the correct choice for the superpeer interconnection.

6 Conclusion

Working communications are extremely important in disaster management and first response scenarios. The heterogeneity of devices means traditional communication systems have to overcome several hurdles to be effective. Peer-to-peer-based communication architectures can alleviate many of these problems, by exploiting the inherent heterogeneity of the devices. We have evaluated different overlay networks in terms of their suitability for first response, and have found out that a superpeer overlay with good selection of superpeers is the most promising candidate. An unstructured network is able to handle medium amounts of churn, but breaks down under heavy churn. A DHT-based network, although it remains connected, is not able to transmit messages to correct recipients, making it unsuitable for first response scenarios.

References

[1] A.S. Bahora et al. Integrated peer-to-peer applications for advanced emergency response systems. part i+ii. concept of operations. In *SIEDS, IEEE*, 2003.

[2] D. Bradler, J. Kangasharju, and M. Muehlhaeuser. Systematic First Response Use Case Evaluation. *MODIES*, 2008.

[3] G. Chen and D. Kotz. Context-aware resource discovery. In *PerCom*, 2003.

[4] G. Holland and N. Vaidya. Analysis of TCP Performance over Mobile Ad Hoc Networks. *Wireless Networks*, 8(2):275–288, 2002.

[5] Intel. Intel pledges 1500 PCs, wireless access points, technical support for hurricane Katrina disaster relief efforts. www.intel.com/pressroom/archive/releases 20050902corp.htm, 2005.

[6] A. Rowstron and P. Druschel. Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems. In Middleware 2001.

[7] S. Krishnamurthy et al. A statistical theory of chord under churn. In *IPTPS*, pages 93–103, 2005.

[8] S. Rhea et al. Handling churn in a DHT. In *USENIX*, June 2004.

[9] Titan Systems Corporation. Arlington county after-action report on the response to the september 11 terrorist attack on the pentagon. www.911investigations.net/document793.html, 2002.

[10] Nortel Government Solutions. Positioning paper wimax for government-grade secure mobility. www.actgov.org/actiac/documents/pdfs/wimax.pdf, 2006.

[11] Y. Yamao et al. Multi-hop radio access cellular concept for fourth-generation mobile communications system. *PIMRC*, 2002.

[12] B. Y. Zhao and et al. Tapestry: A resilient global-scale overlay for service deployment. *Journal on Selected Areas in Communications*, 22(1):41–53, Jan. 2004.