

Towards a scalable ad hoc network infrastructure

Ansley Post
abpost@rice.edu

Rice University, Houston, TX, USA

Abstract

As wirelessly networked devices become more pervasive, large scale mobile ad hoc networks will become increasingly important for providing network services. In addition to basic communications, higher level services must be provided in order to make these networks useful for a large community. When this goal is met, the necessity of deploying dedicated network infrastructure may no longer be required. Ad hoc networks are highly dynamic and decentralized in much the same way as p2p networks, this similarity may make p2p ideas well suited for application in this area. This paper discusses the issues of applying p2p ideas to provide scalable services on a large ad hoc wireless network.

1 Introduction

Ad hoc wireless communication is a powerful technology allowing self organizing connectivity and network services with no preexisting infrastructure. Due to the fact that communication is not tied to any dedicated infrastructure, ad hoc networks are potentially more resilient and pervasive. This flexibility allows networks where there is no place to put wiring or the cost of installing infrastructure is prohibitive. In the case of a disaster where the existing infrastructure is damaged, ad hoc networking would allow communication that would otherwise not be possible. Additionally, ad hoc communication allows the network to grow with the number of people using it, not requiring new infrastructure to be built.

Currently, mobile ad hoc networks (MANETs) allow the federation of a few hundred nodes to provide mutual connectivity. The only service provided by the network is basic connectivity; services such as NTP, DNS, and other high level applications are not available or could not be deployed in an efficient manner. The underlying network does not lend itself well to centralized applications since the broadcast medium would be quickly saturated around the server. While this lack of services may be acceptable for the small networks possible today, larger systems will need these applications to be useful.

In order for larger ad hoc networks to be practical, scalability problems must be solved and self organizing distributed

applications must be developed to replace those traditionally provided by infrastructure servers. There have been several solutions proposed to the problem of scalability, most notably those based on landmarks or coordinate systems. The coordinate system imposed by these networks is strikingly similar to those used in structured overlays. MANETs and structured p2p overlays share the properties of dynamic network membership and a decentralized nature with no administrative control.

These similarities may help solve the problem of developing scalable ad hoc applications. Many decentralized applications have been built upon scalable p2p overlays and these same ideas may be useful in developing higher level services in an ad hoc setting. This idea seems promising, but there are several key challenges that makes the environment more difficult than the traditional Internet. The networks are highly dynamic with mobile nodes, and small cross sectional bandwidth introduces challenges not seen in a wired network.

The goal of providing higher level service on top of large MANETS require a number of concerns be addressed. In this paper we explore these concerns from the standpoint of implementing p2p style applications on top such a network and discuss adaptations that are necessary.

The structure of the paper is as follows. Section 2 provides background information on peer to peer overlays and mobile ad hoc networks. Section 3 describes the issues that must be addressed in a wireless environment. Section 4 concludes.

2 Background

In this section, we briefly provide background and describe ad hoc networks and peer-to-peer networking.

2.1 Mobile Ad Hoc Networks and Landmark Routing

A number of ad hoc network routing protocols use a routing table to maintain a path from any source to any destination, called *proactive routing* [17, 14]. However, as the nodes are mobile, a lot of control overhead is generated to maintain these frequently broken routes. *Reactive routing* protocols

[12, 18] for ad hoc networks are designed so that routing information is acquired only when needed. The reactive protocol discovers a new route to the destination on-demand, and tends to maintain it as long as the path is still being used for data traffic. This reduces control overhead to zero if the nodes are stationary. The above approach works well for small sized networks but for large networks routing without having any infrastructure is a big challenge. Routes might break as soon as they are discovered, wasting a lot of bandwidth without getting any data across.

Scalable routing algorithms based on landmark hierarchies were described in a seminal paper by Tsuchia [26, 24, 25]. Landmark nodes self-organize themselves into a hierarchy, such that landmarks at a given level in the hierarchy are an approximately equal number of network hops apart. Each node maintains routes to the nearest landmark node at each level. The address of a node consists of the sequence of ids of the nearest landmarks, from highest to lowest level. During routing, a node extracts from the destination address the highest level landmark id that differs from its own node address, and forwards the packet towards the landmark with that id. The mapping from node identifiers to their current address is maintained in a distributed fashion. Landmark routing achieves scalability by dramatically reducing the size of per-node routing tables at the expense of somewhat longer routes. Since node addresses are maintained dynamically, landmark routing also lends itself to mobile hosts.

LANMAR [16] is a variation landmark routing that targets mobile ad hoc networks with groups of nodes that are related in function and mobility, combining some of ideas from Landmark routing and from Fisheye State Routing [15] to make MANETs' more scalable. L+ [5] is an adaptation of landmark routing for mobile ad hoc networks.

2.2 Peer-to-peer systems

In this section, we briefly provide some background on peer-to-peer systems. Peer-to-peer systems are self-organizing, decentralized overlay networks, in which participating nodes contribute resources and cooperate to provide a service.

Structured p2p overlay networks conform to a specific graph structure that allows them to locate objects reliably and in a probabilistically bounded number of routing hops [19, 23, 21, 27]. In these overlays, every node and every object is assigned a unique identifier randomly chosen from a large identifier space, referred to as a *nodeId* and *key*, respectively. Given a message and a key, the overlay routes the message to the live node whose *nodeId* is numerically closest to the key in $\log N$ hops, where N is the number of nodes in the network. Structured overlays scale to millions of nodes and they are highly resilient to node failures.

One of the abstractions that can be implemented upon a structured overlay is a distributed hash table (DHT). Like a conventional hash table, a DHT stores (key, value) tuples,

but does so in a highly scalable, decentralized, and fault-tolerant manner. DHTs can be used, for instance, to provide decentralized storage and distributed lookup services [6, 22]. Structured overlays can also provide scalable, end-system-based group communication [28, 20, 3]. Scribe [3], for instance, supports group multicast and anycast, and scales to large numbers of groups, highly dynamic groups, and groups sizes ranging from one node to the entire Internet. It has been used to support scalable event notification [4], high-bandwidth content distribution [2], and decentralized resource location.

Existing structured overlay networks were designed for the Internet, with its mostly stationary hosts and its high-capacity backbone. Non-local communication is the norm in these overlays, overlay routing is completely separate from network routing, and overlay maintenance overhead can be substantial. As a result, these systems are not directly applicable to ad hoc wireless networks, where cross-section bandwidth is scarce and the network topology may change rapidly [10, 13, 9]. We are currently developing a system that combines a Landmark like structure and p2p ideas to provide scalable network services in a wireless environment.

3 Issues

In this section the issues of adapting p2p overlays for use on top of a MANET to provide the basis for scalable network services are discussed. Once these problems are addressed, implementing p2p applications to provide critical networks services is much more practical.

3.1 Traffic Characteristics

In the Internet, there are high bandwidth long distance channels that connect physically far away parts of the network. In a MANET, all of the interfaces share the same channel throughout the network. There is a very limited cross sectional bandwidth because there are no high capacity links. Traffic that is sent across the network will have to traverse the links between many nodes, thus using much of the network's capacity. Cooperative applications designed for the Internet, such as DHTs, must be reexamined before being transferred to a MANET as bandwidth optimization is now important.

It would be ideal if all traffic crossed a very few links, but the predominant mode of communication in a large MANET is unknown and must be studied. In addition to studying the general traffic patterns in the network, it would also be useful to observe the traffic patterns of particular applications. Some applications may be more intrinsically suitable for deployment in a wireless network. Studying the traffic patterns of applications may also help determine ways applications could be better adapted to the wireless environment.

Some application's data has natural physical locality and is

only useful to a group of users who are geographically close to each other. Limiting the scope in which data can be retrieved to a smaller number of nodes than the whole system will naturally limit long distance traffic. Some applications have natural locality such as a community bulletin board system or class newsgroups. There should be a mechanism for inserting data so that limits the scope in which it is viewable, perhaps similar to [8]. This could help conserve bandwidth in the system, but there should be an incentive for an application to do this.

While applications can be designed to minimize the amount of traffic that travels very far in the network, it can not always be avoided. To prevent such traffic from degrading performance of the entire network there should be mechanisms in place to either make such traffic less damaging, or to throttle it. If a metric is developed for how much a packet would cost to send in terms of bandwidth, perhaps a packet can be buffered in order to conserve bandwidth until the networks is idle and then sent. It also might be possible to combine packets headed in the same direction from several sources into one and then split them when their paths diverge so that all the space in one maximum transfer unit is utilized.

Once basic network connectivity is established, it would be possible to run a traditional centralized service such as DNS on an ad hoc network. This would not be a scalable solution for several of the reason mentioned above, the broadcast channel around the nodes running these services could easily be saturated, and the distances to the root servers may be significant. To build a viable service on top of an ad hoc network it should balance the load equally among the nodes and be aware of the wireless environment it is running in, something that traditional services do not do.

3.2 Fault Model

Structured overlays are often built in a manner that makes failures unlikely to affect the overall system. This is possible because in the wired internet failures are less likely then in a wireless environment and the network is fairly well connected. In the wireless environment the failures will be more drastic, thus aggravating many of the problems of correlated failure [11]. In an ad hoc environment there may be one node holding two halves of a network together, and if that node fails the two halves of the network will each see half the nodes as having failed.

Additionally there are also many other factors that can cause wireless nodes to fail at a much higher rate. Many devices using wireless interfaces are powered by batteries, and will turn off to conserve batteries. Wireless devices are subject to disconnection due to the peculiarities of signal propagation, a device moving between two rooms may get varying degrees of connectivity. The loss rate of the wireless channel is also much higher than that of a wired ethernet channel resulting in the need to retransmit more packets.

A structured overlay designed on top of an unreliable system such as this must be implemented to handle these types of failures gracefully and not produce much additional overhead. The system must be resilient to failure as extreme as large network partitions that may make many nodes unreachable. For successful scalable network services to be built, these types of faults must not greatly degrade the services being provided.

3.3 Utilizing Existing Infrastructure

Having the technology to build large scale MANETs may eliminate the need for dedicated network infrastructure, but it is unknown whether the problems of low bandwidth and reliability can be overcome. One possible way of overcoming this obstacle is to utilize wired infrastructure when possible. In the case of a disaster some small islands of wired connectivity may exist and by attaching these to the ad hoc network, some critical services may be restored.

Utilizing existing wired infrastructure would allow islands of wireless connectivity to be connected forming a much larger network than would otherwise be possible. Inside a wireless island there maybe a wired connections that connects two distant points; it maybe more efficient for a packet to be routed through this connection towards its destination instead of the wireless channel.

The presence of wired infrastructure may also help eliminate several problems of reliability, as the wired channel is more reliable than the wireless. A network that has some underlying wired connectivity is less likely to suffer partitions as it has the redundant connection.

The best way to integrate these two very different environments needs to be studied. Simply using conventional bridging to attach the two types of networks leaves several questions unanswered about how to most efficiently use wired connection. For example it may be better to travel more hops if some of these hops are wired connection and then to take less wireless hops. Also the wired nodes are likely to not be mobile, there may be additional reliability assumptions that can be made about them, and there must be a mechanism to advertise their presence and capabilities. These issues must be addressed before existing infrastructure can be leveraged to provide a better level of service. While utilizing existing infrastructure may make some problems easier, simply wiring everything is not a viable option as it is expensive and should not be relied on as a solution.

3.4 Data Replication

In order to guarantee availability of data, it is often replicated over a configurable number nodes nodes in the system. The availability requirements of the data must be weighed against the bandwidth cost to store the data in nodes far away from

the inserter. Storing data in a diverse set of nodes will decrease the likelihood that all will fail but it will increase the cost of inserting and retrieving the data. If the nodes close to the inserter are chosen the cost of inserting the data will decrease but in a correlated failure, such as a partition, it will be likely that data is unavailable to the network.

This tradeoff must be studied carefully in order to provide highly available distributed data. A combination of caching, storing a small number of replicas far away and some close to the inserting node may yield satisfactory availability. It is possible that in highly dynamic network such as this that scalable storage and high availability may not be possible [1]. This issue must be studied in more depth before systems requiring availability guarantees can be implemented in this environment.

3.5 Id Assignment

In a traditional Internet environment ids are assigned either by generating a random id from a very large id space or assigned by a trusted authority. The possibility of collision is made very small by having the id space be very large like 160 bits. This may not work as well in the wireless environment. A node has a coordinate in the system, as well as a node id. A node's id may be its coordinate for each landmark level and this entire coordinate string uniquely identifies the node. Since the wireless channel is low bandwidth having many bit identifiers, which must be included in each packet, will add overhead to every packet sent. The length of each coordinate must be balanced for efficiency versus the possibility of collision. Alternatively the node may have a separate node id and coordinate, if this is the case there has to be a distributed lookup mechanism to map node ids to coordinates. This lookup may result in a higher overhead for communication.

If ids are generated by location and not by a secure mechanism, the system may be vulnerable to id based attacks. In particular a Sybil attack [7] could be mounted by a hostile party simply moving a large number of mobile nodes into close proximity to a target node. This could allow the attacker to become an intermediate hop in all traffic being sent to the target. An attacker could also not forward control packets that contain information about the topology, causing the target to have an incorrect view of the network.

3.6 Id Change

A complicating factor in a mobile environment is that when a node moves its coordinate relative to other nodes will change. This can cause several problems, making this environment much more challenging. First, there is the problem of finding out where the node one wants to communicate is at any given time. This problem can be solved by a distributed look up service similar to the one in L+ [5]. This allows a

node's locations to be determined by knowing an well known identifier for that node such as an IP address.

Secondly, since nodes can move arbitrarily coordinates that are looked up may become stale. In the event that this happens it may not be detected until no response is received from the destination node. This results in wasted traffic that consumes part of the limited bandwidth. To counteract this, coordinates can be updated in the DHT more often but this itself requires bandwidth. This tradeoff needs to be studied to find the sweet spot where most of the coordinates are right but they are not being updated too frequently. This has important consequences when caching coordinates, which will be very important in reducing traffic from lookups.

3.7 Preliminary Ideas

In this paper the issues that must be confronted to use p2p ideas to scale ad hoc networks have been enumerated. While each of the pose a specific problem and must each be individually addressed, there is still reason to be optimistic. A system that takes advantage of the broadcast nature of the wireless environment could use extensive caching to possibly overcome some of the bandwidth concerns. Information about the network could be exposed to the application level so that decisions about how to most effectively use the limited bandwidth resource can be made. In addition to leveraging existing wired infrastructure, these techniques will bring us closer to practical, usable, large scale ad hoc networks.

4 Conclusions

The demand for wireless connectivity is increasing, it is now possible to check your email at the local coffee shop and to surf the web while in meetings. Current technology works by providing base stations that connect to the wired Internet. Hobbyists are experimenting with different types of antennas and higher powered base stations in order to increase the range of wireless networks to cover larger and larger areas. This is ideal for people who are inside these coverage areas but when they leave, they lose coverage. Using the wireless interface inside of mobile devices it is possible to extend the range of wireless base stations or to even eliminate them.

In a citywide MANET, it would be possible for the population to cooperatively provide wireless coverage to an entire city, if current scalability and the lack of network services are overcome. In this environment there would be no need for infrastructure companies to charge for services as they would be available to everyone who had a wireless interface. In disaster situations surviving infrastructure in addition to wireless technology could be used to provide communication and networks services and help coordinate relief efforts. In order for these applications of wireless technology to be successful the issues in this paper must be addressed.

References

- [1] C. Blake and R. Rodrigues. High availability, scalable storage, dynamic peer networks: Pick two. In *Proc. HotOS IX*, Kauai, Hawaii, May 2003.
- [2] M. Castro, P. Druschel, A.-M. Kermarrec, A. Nandi, A. Rowstron, and A. Singh. Splitstream: High-bandwidth content distribution in cooperative environments. In *Proc. of the 19th ACM Symposium on Operating System Principles (SOSP 2003)*, Bolton Landing, NY, October 2003.
- [3] M. Castro, P. Druschel, A.-M. Kermarrec, and A. Rowstron. SCRIBE: A large-scale and decentralized application-level multicast infrastructure. *IEEE JSAC*, 20(8), Oct. 2002.
- [4] M. Castro, M. Jones, A.-M. Kermarrec, A. Rowstron, M. Theimer, H. Wang, and A. Wolman. An evaluation of scalable application-level multicast built using peer-to-peer overlay networks. Technical report, Submitted for publication, 2002.
- [5] B. Chen and R. Morris. L+:scalable landmark routing and address lookup for multi-hop wireless network. Technical Report MIT-LCS-TR-837, Laboratory for Computer Science Massachusetts Institute for Technology, 2002.
- [6] F. Dabek, M. F. Kaashoek, D. Karger, R. Morris, and I. Stoica. Wide-area cooperative storage with CFS. In *Proc. ACM SOSP'01*, Banff, Canada, Oct. 2001.
- [7] J. R. Douceur. The Sybil attack. In *Proceedings for the 1st International Workshop on Peer-to-Peer Systems (IPTPS '02)*, Cambridge, Massachusetts, Mar. 2002.
- [8] M. J. Freedman and D. Mazieres. Sloppy hashing and self-organizing clusters. In *Proceedings for the 2nd International Workshop on Peer-to-Peer Systems (IPTPS '03)*, Berkeley, CA, Feb. 2003.
- [9] N. Gupta and S. R. Das. A capacity and utilization study of mobile ad hoc networks. In *Proceedings of the Wireless Local Networks (WLN) Workshop*, November 2001.
- [10] P. Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Transactions on Information Theory*, 40(2):388–404, March 2000.
- [11] N. Harvey, M. Jones, S. Saroiu, M. Theimer, and A. Wolman. Skipnet: A scalable overlay network with practical locality properties. In *Proc. USITS 2003*, Seattle, WA, 2003.
- [12] D. B. Johnson and D. A. Maltz. Dynamic Source Routing in Ad Hoc Wireless Networks. In T. Imielinski and H. Korth, editors, *Mobile Computing*, chapter 5, pages 153–181. Kluwer Academic Publishers, 1996.
- [13] J. Li, C. Blake, D. S. J. D. Couto, H. I. Lee, and R. Morris. Capacity of ad hoc wireless networks. In *Proceedings of the 7th Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom 01)*, July 2001.
- [14] V. D. Park and M. S. Corson. A highly adaptive distributed routing algorithm for mobile wireless networks. In *Proceedings of INFOCOM'97*, pages 1405–1413. IEEE Computer Society and Communications Society, Apr. 1997.
- [15] G. Pei, M. Gerla, and T.-W. Chen. Fisheye state routing in mobile ad hoc networks. In *ICDCS Workshop on Wireless Networks and Mobile Computing*, pages D71–D78, 2000.
- [16] G. Pei, M. Gerlan, and X. Hong. Lanmar: Landmark routing for large scale wireless ad hoc networks with group mobility. In *SIGCOMM'88*, Stanford, CA, 1988.
- [17] C. E. Perkins and P. Bhagwat. Highly dynamic Destination-Sequenced Distance-Vector routing (DSDV) for mobile computers. In *Proceedings of the SIGCOMM '94 Conference on Communications Architectures, Protocols and Applications*, pages 234–244. ACM, Aug. 1994.
- [18] C. E. Perkins and E. M. Royer. Ad-hoc on-demand distance vector routing. In *Second IEEE Workshop on Mobile Computing Systems and Applications*, pages 90–100, Feb. 1999.
- [19] S. Ratnasamy, P. Francis, M. Handley, R. Karp, and S. Shenker. A scalable content-addressable network. In *Proc. ACM SIGCOMM'01*, San Diego, CA, Aug. 2001.
- [20] S. Ratnasamy, M. Handley, R. Karp, and S. Shenker. Application-level multicast using content-addressable networks. In *NGC*, Nov. 2001.
- [21] A. Rowstron and P. Druschel. Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems. In *IFIP/ACM Middleware 2001*, Heidelberg, Germany, Nov. 2001.
- [22] A. Rowstron and P. Druschel. Storage management and caching in PAST, a large-scale, persistent peer-to-peer storage utility. In *Proc. ACM SOSP'01*, Banff, Canada, Oct. 2001.
- [23] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan. Chord: A scalable peer-to-peer lookup service for Internet applications. In *ACM SIGCOMM'01*, San Diego, CA, Aug. 2001.
- [24] P. F. Tsuchiya. The landmark hierarchy: description and analysis. Technical Report Technical Report MTR-87W00152, MITRE corporation, June 1987.
- [25] P. F. Tsuchiya. Landmark routing: architecture, algorithms, and issues. Technical Report Technical Report MTR-87W00174, MITRE corporation, May 1988.
- [26] P. F. Tsuchiya. The landmark hierarchy: a new hierarchy for routing in very large networks. In *Proc. of MobiHOC*, Stanford, CA, Aug. 2000.
- [27] B. Zhao, J. Kubiatowicz, and A. Joseph. Tapestry: An infrastructure for fault-resilient wide-area location and routing. Technical Report UCB//CSD-01-1141, U. C. Berkeley, April 2001.
- [28] S. Zhuang, B. Zhao, A. Joseph, R. Katz, and J. Kubiatowicz. Bayeux: An architecture for scalable and fault-tolerant wide-area data dissemination. In *NOSSDAV*, June 2001.