CarNet: A Scalable Ad Hoc Wireless Network System

Robert Morris John Jannotti Frans Kaashoek Jinyang Li Douglas Decouto
{rtm,jj,kaashoek,jinyang,decouto}@lcs.mit.edu
MIT Laboratory for Computer Science
545 Technology Square, Cambridge MA 02139

Abstract

CarNet is an application for a large ad hoc mobile network system that scales well without requiring a fixed network infrastructure to route messages. CarNet places radio nodes in cars, which communicate using Grid, a novel scalable routing system. Grid uses geographic forwarding and a scalable distributed location service to route packets from car to car without flooding the network. CarNet will support IP connectivity as well as applications such as cooperative highway congestion monitoring, fleet tracking, and discovery of nearby points of interest.

1 Introduction

The Internet has evolved in a way that sacrifices dynamism in favor of scale: it groups nodes into an addressing and routing hierarchy that inhibits movement, and it depends on fixed physical infrastructure that inhibits rapid deployment. We are designing a scalable and dynamic network architecture called Grid, which will enable new kinds of applications and will be easier to deploy than existing technology.

We desire several kinds of dynamism. First, adding a new node to the network should require no human involvement beyond placing the new node within radio range of an existing node; all the configuration involved should be done automatically. Second, the network should not rely on any fixed infrastructure: such reliance would hinder deployment. Third, nodes should be able to move. Fourth, it should be easy for applications to interact with a changing set of nearby resources. Fifth, underlying protocols should provide APIs and use algorithms that take advantage of rich and changing network topologies. Our main research goal is to provide these forms of dynamism without sacrificing scalability.

This research is partially supported by DARPA contract N66001-99-2-8917.

These goals are similar to those of mobile ad hoc networking (MANET) systems. Grid uses a key MANET idea: network nodes use radios to talk to immediate neighbors, and reach distant destinations by forwarding each others' packets. Existing MANET routing systems typically find distant destinations by flooding location or topology information across the entire network. Global flooding allows nodes to attach to the network anywhere, but it scales badly beyond a few hundred nodes. We wish to build a system that can scale to hundreds of thousands of nodes. No existing network technology scales to that size without relying on static hierarchical techniques.

Though ad hoc networks may be attractive, they are more difficult to implement than fixed networks. Fixed networks take powerful advantage of their static nature in two ways. First, they proactively distribute network topology information among the nodes, and nodes precompute routes through that topology using relatively inexpensive algorithms. Second, fixed networks embed routing hints in node addresses because the complete topology of a large network is too unwieldy to process or distribute globally. Neither of these techniques works well for networks with mobile nodes because movement invalidates topology information and permanent node addresses cannot include temporary location information.

One topological assumption does work well for radio-based ad hoc networks: nodes that are physically close are likely to be close in the network topology, that is, connected by a small number of radio hops. Grid uses this assumption to provide scalable ad hoc routing.

To test Grid we are designing and implementing CarNet, a network of cars equipped with Grid nodes. We picked the CarNet application because it is incrementally deployable, involves mobility, suggests novel user applications, and doesn't force undue focus on issues such as battery life and equipment weight. Our vision is that all cars in a large geographic region form a network fabric that can support new applications such as cooperative highway congestion monitoring, fleet tracking, and discovery of nearby points of interest. Of course, as battery life and component weights decrease we hope that laptops and PDAs might be similarly equipped, allowing even more interesting applications.

Much of CarNet is speculative. We have started the design and implementation of the CarNet system, but the primary focus has been on scalable ad hoc routing. In this paper we outline the tentative overall design and Grid, CarNet's routing system.

2 Architecture

The foundation of Grid's scalability is geographic forwarding. A source node using geographic forwarding annotates each packet with the location of its destination. The packet moves, hop by hop, through the network, forwarded along by cooperating intermediate nodes. At each node, a purely local decision is made to forward the packet to the neighbor that is geographically closest to the destination. The fact that forwarding does not involve any global information helps geographic routing scale well and cope well with mobile nodes.

2.1 The Grid Location Service (GLS)

Geographic forwarding requires sending nodes to discover the locations of destinations. That is, the network must provide a database that maps each node's permanent ID to its current geographic location. The database ought not to depend on any special fixed infrastructure that might hinder easy deployment or scalability; instead, it should be distributed over all the nodes. The database ought to be robust: the failure of a few individual nodes should affect the reachability of at most a few other nodes. Finally, the location database should be efficient: the closer two nodes are, the more quickly they should be able to learn each others' locations.

GLS, Grid's distributed location service, fulfills these requirements, using the following general approach. All nodes in the system agree on a distributed algorithm $f(i)$ that maps each node identifier to a list of physical locations, expressed as latitudes and longitudes. f has hash-like qualities that ensure that its values for different arguments are spread out over physical space. The locations produced by $f(i)$ act as node i's location servers. Whenever node i moves, it uses geographic forwarding to send position updates to the locations specified by $f(i)$; the nodes close to those locations remember node i's position. When a node j wants to find i, it sends queries to the locations in $f(i)$; nodes close to those locations will know i's location and respond.

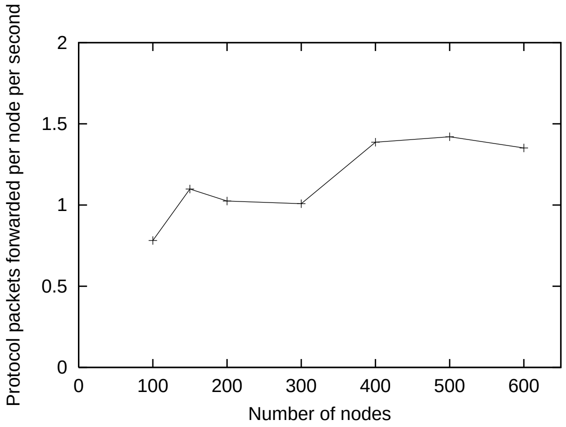

Figure 1: The number of Grid protocol packets forwarded per node per second as a function of the total number of nodes. (Taken from [5].)

A previous paper [5] describes the details of GLS's $f(i)$ and shows that it scales well. The per-node cost of GLS, in both storage and messages forwarded, is proportional to the log of the total number of nodes. Figure 1 shows GLS's message overhead over a range of network sizes, and Figure 2 indicates the fraction of data packets successfully delivered. These graphs were produced by simulations involving a fair amount of node mobility. Movement is the primary reason that Grid sends protocol messages and fails to deliver packets. The simulation scenarios are similar to those used in other studies of ad hoc networks [6], and show that Grid scales to large numbers of nodes more gracefully that other protocols.

2.2 Managing Density

Geographic forwarding may fail when the network is insufficiently dense, so that packets encounter "holes" in the topology. That is, a packet may arrive at a node which has no neighbors in radio range that are closer to the destination. On the other hand, a radio network may also be *too* dense. Radios that lie within radio range of each each other must share the limited spectrum. Therefore as node density increases, the bandwidth available to each node decreases.

Grid will route around holes in the distribution of nodes using techniques like those proposed by Karp [10] and Bose [4]. The basic idea is to have nodes agree on the conventions required to send pack-

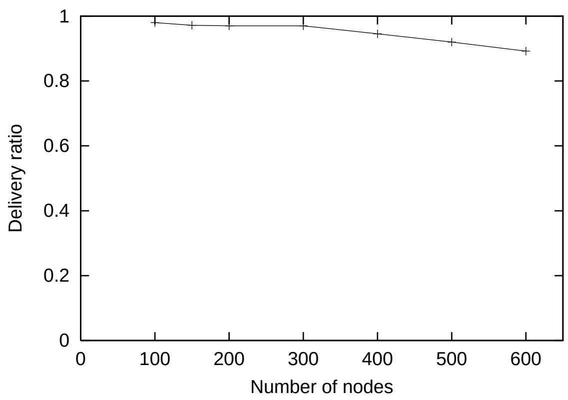

Figure 2: The fraction of data packets that are successfully delivered by Grid as a function of the total number of nodes. (Taken from [5].)

ets around the perimeters of hole until they either make progress or return to the starting point on the hole.

Grid will also use variable-power radios to help with variable node densities. The idea is to vary the transmit power in order to keep an approximately constant number of nodes in radio range. This helps keep the network connected at low densities, and allow good spatial reuse of spectrum at high densities. It also yields a power saving effect similar to minimum energy routing [16], since many short hops use less total power than a few long hops.

3 Applications

We are building the CarNet system around Grid to explore how Grid will interact with applications, link-level hardware, and the stresses of large-scale deployment. Each CarNet car will have a node consisting of an embedded Linux computer, an IEEE 802.11 radio, a GPS receiver, and displays for the driver and/or passengers. The nodes will be able to run standard Internet applications, such as web browsers, using CarNet to contact the wired Internet. We are also planning to produce applications that take advantage of CarNet's special features, as well as applications to address privacy concerns that a geographic routing system is likely to present.

3.1 Resource Location

Resource location is a common general-purpose application provided by mobile networks. In a static network it is common to configure nodes with information about local services, such as printers and web caches. In an ad hoc network, it must be possible to locate nearby services dynamically.

GLS can be used unaltered to locate nearby resources. The idea is simply to associate a standard name to a resource, such as "printer", "web cache" or "Internet access point." This name is hashed to obtain an ID. The resource then participates in the GLS protocol using that ID. GLS is robust in the face of multiple nodes using the same ID, and, in fact, the end result is exactly what is most desirable. When performing a location lookup, nodes will "see" the resource closest to them.

3.2 IP Connectivity

IP connectivity to the Internet at large is an important goal of CarNet. In the simplest case, this consists of a single wireless Grid network and a fixed Grid-to-Internet gateway. Each Grid node is given an IP address, a subnet mask, and the IP address of its default router — the Grid-to-Internet gateway.

Wireless nodes participate in the Grid protocol using a hash of their IP address as their Grid ID. This simplifies IP connectivity. If a Grid node wishes to send IP packets to another Grid node, instead of an ARP phase to determine the destination's MAC address, a location query obtains the destination's location. To send a packet to the Internet, a Grid node determines the location of the gateway and forwards the IP packet to it. This process uses the standard netmask and routing table mechanisms to distinguish between local and remote destinations. The only difference is that ARP is replaced by a location query. One may think of the destination's location as its "hardware address" when running IP over Grid.

Extra gateways can easily be added to this simple setup. The addition of redundant gateways might be used to provide better connectivity in a large metropolitan Grid network. Is is desirable for packets to spend as little time as possible in the wireless network, therefore it is desirable for nodes to use the nearest gateway as their default router. This conserves the relatively precious wireless bandwidth. This can be arranged for "outgoing" (Grid-to-Internet) packets quite easily as a special case of resource location. All gateways participate in the Grid protocol using the *same* IP address. Thus, when a wireless node transmits packets, the location query will find the nearest gateway.

Similar efficiency on the return path is harder to achieve: the Internet routing system will deliver packets to the Grid-to-Internet router closest to the source, not closest to the destination Grid node. This inefficiency is likely tolerable in a campus or metropolitan network, since the extra delay will be small. If it is not tolerable, the Grid-to-Internet gateways could use the wired network to forward packets among themselves in a manner similar to the inter-MSS forwarding in the Columbia Mobile*IP system [8]. Such forwarding could also help heal partitions in the Grid network itself.

3.3 CarNet Specific

In addition to straightforward services like resource location and IP connectivity, we hope that the unique properties of a wireless network in which nodes know their own locations may inspire other interesting applications. Some examples include:

- Location directed multicast – Messages may be sent to all nodes in a given area efficiently, to query for or advertise the existence of physical services.
- Traffic congestion monitoring – Cars can exchange speed and location information, which can be assembled and overlaid on a map to provide a picture of which roads are congested.
- Fleet tracking – A delivery service could optimize pick-up routes in real time, adjust drop-off time predictions, and collect maintenance information from vehicle diagnostic computers.
- Over-the-horizon radar detection – Cars encountering speed traps could alert approaching cars before they are in radar range.
- Chat – The success of CB radio suggests that drivers like to exchange views on traffic conditions and other issues.
- Marine – Fixed infrastructure is particularly troublesome to deploy at sea. BoatNet might be a cost effective alternative.

3.4 Privacy

CarNet's use of geographic routing presents a problem for users concerned about the privacy of their location and movements. Using GLS any node may locate any node whose ID is well-known. We believe that two techniques could prevent this fundamental fact from becoming an issue of concern for most users. Nodes could change IP addresses (and therefore Grid IDs) every few minutes. A thousand nodes might share a pool of a thousand IP addresses, thereby making the tracking of individuals difficult. Unfortunately, this technique would make opening connections to such a node impossible because its current ID would be unknown. The node can only open outgoing connections. This is the same constraint that many Internet users find themselves working under when placed behind a NAT (Network Address Translator), so we believe it to be satisfactory for many users.

Periodically changing one's IP address protects a user from tracking by a third party, but does not prevent nodes contacted by the user from learning its location. To address this concern, users may wish to employ a proxy for certain sensitive connections. Untrusted nodes would receive connections from the proxy instead of the node, thereby insulating the wary user from tracking. Even on the wired Internet a wide variety of such proxies are commonplace [15]. CarNet merely raises the stakes, since normal connections would reveal a user's current location rather than his IP address.

On the other hand, Grid's distributed design happens to afford CarNet users a degree of *protection* from privacy invasions. In the cellular telephone network a phone's location must be tracked for call routing and billing purposes. A phone's current location must be known at all times at a central location, therefore it is quite natural to maintain long term logs of a phone's movement history. In Grid, however, a node's location servers change constantly, and no one server is likely to have a large portion of an individual's movement history.

4 Related Work

A few examples of deployed ad hoc routing systems exist, particularly the Metricom Ricochet [1] and Nokia Rooftop [2] systems. Though the protocols these systems use are not public, neither appears to support mobile nodes (in the sense of mobile participants in the ad hoc protocol). Shepard [16] considers spectrum and medium access control issues involved in scaling such systems up to metropolitan sizes.

A number of ad hoc routing algorithms exist, such as DSR [9], AODV [13], DSDV [14], and TORA[12]. As shown in [5], simulations suggest that these systems do not scale as gracefully as Grid to large numbers of nodes. The fundamental reason is that Grid's geographic forwarding avoids global flooding of queries or topology information. Some previous routing systems do use geographic routing, but either don't discuss scalable location services [7], or use global communication to determine node locations [11, 3].

5 Summary

CarNet is a test bed to help explore protocols for large ad hoc networks. We expect that deployment of CarNet will bring to light new problems and suggest new solutions in a number of areas. CarNet will handle scaling to large numbers of nodes using geographic forwarding and the GLS distributed location service. The problem of varying node density will require adaptive algorithms to manage radio spectrum and power levels. Finally, CarNet will facilitate the creation of new applications suited to geographically aware networks.

References

[1] Metricom Ricochet wireless modem. `http://www.metricom.com`.

[2] Nokia Rooftop wireless routing system. `http://www.nwr.nokia.com`.

[3] Stefano Basagni, Imrich Chlamtac, Violet R. Syrotiuk, and Barry A. Woodward. A Distance Routing Effect Algorithm for Mobility (DREAM). In *Proc. ACM/IEEE MobiCom*, pages 76–84, October 1998.

[4] P. Bose, P. Morin, I. Stojmenovic, and J. Urrutia. Routing with guaranteed delivery in ad hoc wireless networks. In *Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications (DIALM99)*, 1999.

[5] Douglas S. J. De Couto, John Jannotti, David Karger, Jinyang Li, and Robert Morris. A scalable location service for geographic ad hoc routing. In *Proc. ACM/IEEE MobiCom*, August 2000.

[6] Samir Das, Charles Perkins, and Elizabeth Royer. Performance comparison of two on-demand routing protocols for ad hoc networks. In *Proc. IEEE Infocom*, pages 3–12, March 2000.

[7] Gregory G. Finn. Routing and addressing problems in large metropolitan-scale internetworks. ISI/RR-87-180, ISI, March 1987.

[8] John Ioannidis, Dan Duchamp, and Gerald Maguire. IP-based protocols for mobile internetworking. In *Proc. ACM SIGCOMM Conference (SIGCOMM '91)*, pages 235–245, September 1991.

[9] David B. Johnson. Routing in ad hoc networks of mobile hosts. In *Proc. of the IEEE Workshop on Mobile Computing Systems and Applications*, pages 158–163, December 1994.

[10] Brad Karp and H. T. Kung. GPSR: Greedy perimeter stateless routing for wireless networks. In *Proc. ACM/IEEE MobiCom*, August 2000.

[11] Young-Bae Ko and Vaidya Nitin H. Location-Aided Routing (LAR) in mobile ad hoc networks. In *Proc. ACM/IEEE MobiCom*, pages 66–75, October 1998.

[12] Vincent D. Park and M. Scott Corson. A highly adaptive distributed routing algorithm for mobile wireless networks. In *Proc. IEEE Infocom*, pages 1405–1413, April 1997.

[13] Charles Perkins, Elizabeth Royer, and Samir R. Das. Ad hoc On demand Distance Vector (AODV) routing. Internet draft (work in progress), Internet Engineering Task Force, October 1999. `http://www.ietf.org/internet-drafts/draft-ietf-manet-aodv-04.txt`.

[14] Charles E. Perkins and Pravin Bhagwat. Highly dynamic Destination-Sequenced Distance-Vector routing (DSDV) for mobile computers. In *Proc. ACM SIGCOMM Conference (SIGCOMM '94)*, pages 234–244, August 1993.

[15] Michael K. Reiter and Aviel D. Rubin. Crowds: Anonymity for web transactions. *ACM Transactions on Information and System Security*, 1(1):66–92, November 1998.

[16] Timothy Shepard. A channel access scheme for large dense packet radio networks. In *Proc. ACM SIGCOMM Conference (SIGCOMM '96)*, pages 219–230, August 1996.