

Geographic Addressing, Routing, and Resource Discovery with the Global Positioning System

Tomasz Imielinski and Julio C. Navas
{imielins,navas}@cs.rutgers.edu
Computer Science Department
Rutgers, The State University
Piscataway, NJ 08855

October 19, 1996

Introduction

In the near future, Global Positioning System (GPS) cards will be deployed in each car and possibly in every user terminal. A user's *location* will become information that is as common as the *date* is today, getting input from GPS, when outdoors, and other location providing devices, when indoors. Availability of location information will have a broad impact on application-level as well as on network-level software.

Possible new services and functionalities include geographic messaging, advertising, and resource discovery. Geographic messaging constitutes the ability to send a message selectively only to specific subareas defined by latitude and longitude. For example, sending an emergency message to everyone who is currently in a specific area, such as a building, train station, or a highway. Once given the ability to send a message to a distinct geographical area, it would then be possible to perform geographically-targeted advertising over the Internet. For instance, a business may wish to advertise a given service only to clients who are within a certain geographic range (say within 2 miles). Conversely, users could use geographic messaging to discover services or resources within a geographical region, such as in the direct proximity of the user. One can imagine a "Who is around" service which finds out who is currently present in a specific geographic area. If we assume that terminals are also equipped with a camera, a user may point his terminal in a specific direction and obtain annotation (links) to the objects displayed by a camera viewer. In this way the whole external world can become one large "web page". Thus, a building may have a link explaining its function. Links can also be attached to mobile objects appearing on the camera viewer.

To support such applications, location should become a first class citizen in networking protocols, such as IP or ATM, and in the application layer. Routing protocols for geographic messages should be developed that would allow routing to a specific area defined by a polygon of geographic coordinates. Location should also be a parameter in World Wide Web (WWW) access protocols to provide pages located on servers within a specific distance from the user. *Distance-based web bookmarks* can be used to restrict the relevance of web pages by using distance as additional relevance criterion when accessing material on the web.

The main objective of this paper is to show how new services and new network functionalities will emerge as a consequence of having location universally available to mobile terminals. We address the question of protocol support for location aware services such as geographic messaging, geographic service discovery, and geographical service advertising. Geographic routing is a key support which is necessary. Thus, the exact routing mechanisms to make it happen are discussed in this paper.

We will first discuss *geocasting* - broadcasting to geographical areas defined as arbitrary polygons. Next, we will follow with a brief discussion of *geo-multicasting*, multicasting to geographic areas.

Related Work

Linking an IP Address with a geographical location has been of interest for quite some time. The first attempt to design a system that actually routes packets according to their geographic destination and the work that is closest to ours is Cartesian Routing by Gregory G. Finn [5]. Xerox's PARC lab also pioneered early work on location dependent services [9].

The recent redesign of the Internet Protocol (IP) [2] and the advent of the Global Positioning System [10] [11] [12] has given a new stimulus for this work. In the proposed redesign of IP [2], IP address type space was specifically allocated for geographic addresses [3] [8]. IP addresses would be assigned to subnets and hosts based on topological criteria, such as geography. In [3] and [8] the sender of a "geographic message" would be unicasting messages only to such hosts which have geographic addresses. The methods in this paper attempt to provide the more general ability of sending a message to all recipients within a geographical area, regardless of whether or not the hosts have geographical addresses.

Addressing Model

Two-dimensional geographic positioning offers latitude and longitude information as a two dimensional vector:

$$\langle \text{latitude, longitude} \rangle$$

where longitude ranges from -180 (west) to 180 (east), and latitude ranges from -90 (south) to 90 (north). Thus $\langle 40.48640, -74.44513 \rangle$ is an example of the geographic coordinates for the town of New Brunswick, New Jersey, U.S.A.

Assuming the use of single precision floating-point numbers, four bytes of addressing space are necessary to store latitude and four bytes are also sufficient to store longitude. Thus a total of eight bytes are necessary to address the whole surface of the earth with precision down to 0.1 mile!

A destination geographic address would be represented by some closed polygon such as:

- point
- circle(center point, radius)
- polygon($point_1, point_2, \dots, point_{n-1}, point_n, point_1$)

where each vertex of the polygon is represented using geographic coordinates. This notation would be used to send a message to anyone within the specified geographical area defined by the closed polygon.

For example, if we were to send a message to city hall in Fresno, California, we could send it by specifying the geographic limits of the city hall as a series of connected lines that form a closed polygon surrounding it. Therefore the address of the city hall in Fresno could look like:

```
polygon([36.80,-119.80], [36,85,-119.76], ... )
```

In a hypothetical usage scenario, a user will interact with a zoom-able map with a graphical user interface. The address of the message will be specified as a polygon on the map. Then, the polygon will be translated into geographic coordinates and the message will be sent to all clients who are located within the bounds of that polygon. The figure 1 illustrates such a scenario where a polygon is drawn around the banks of a river:




Figure 1: Interacting with a Zoom-able Map Interface.

Routing Geographically

Let us now describe the three suggested solutions for delivering a message to any geographical destination. The solutions described are the Geographic Routing Method, the Geographic-Multicast Routing Method, and the Domain Name Service Method. The solutions were chosen so that the necessary geographic routing infrastructure in the Internet varies from very little (Domain Name Service Method) to medium (Geographic-Multicast Method), to significant (Geographic Routing Method). Currently, we are in the process of evaluating a prototype implementation of the Geographic Routing Method and setting up an experimental network capable of routing the geographic messages. This is a DARPA-sponsored Integrated Technology Demonstration (ITD) within the GloMo (Global Mobile Information Systems) program.

All of these methods assume that a user is able to determine his location in some way. While outdoors, the user can make use of the Global Positioning System to determine his location. When indoors, a different method needs to be used. One possible solution is to have each room contain

a radio beacon placed on the ceiling. Each radio beacon will have its own geographic address associated with it which it will broadcast. The geographic address of the mobile hosts will be set to be the same as the beacon's. In this manner, the mobile user can have a geographic address associated with him even though he is indoors and his GPS module is useless.

Geographic Routing Method

The Geographic Routing Method (GEO) uses the polygonal geographic destination information in the geographic message header directly for routing. Geographic routing is going to be implemented in the Internet Protocol (IP) Network layer and the Application layer in a manner similar to the way multicast routing was first implemented. That is, a virtual network which uses geographic addresses for routing will be overlayed onto the current IP internetwork. We would accomplish this by creating our own geographic address routers. These routers would use IP tunnels to transport data packets through areas which do not support geographic routing.

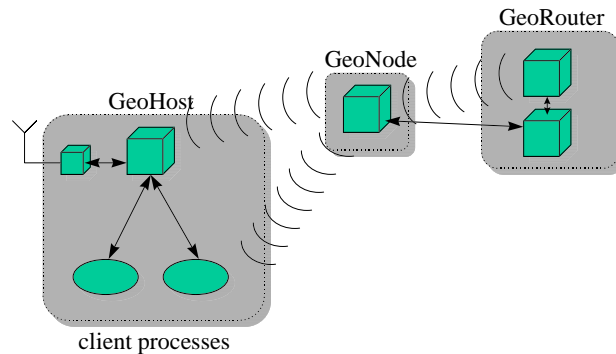


Figure 2: All of the Components of the Geographic Routing System

The system is composed of three main components: GeoHosts, GeoNodes, and GeoRouters as shown in figure 2.

The GeoHost is located on all computer hosts which are capable of receiving and sending geographic messages. Its role is to notify all client processes about the availability of geographic messages, the host computer's current geographic location, and the address of the local GeoNode.

A GeoNode is an entry/exit point for the routing system. The main function of the GeoNode is to store incoming geographic messages for the duration of their lifetimes and to periodically multicast them on all of the subnets or wireless cells to which it is attached. Each subnet and each wireless cell will have at most one GeoNode. The lifetime of a geographic message is specified by the sender of the message. Message lifetimes are necessary because the receivers of geographic messages

may be mobile and may possibly arrive at the message destination just after the geographic message first arrives.

Since, most likely, there will be several geographic messages residing in a GeoNode at one time, the multicasting of the various messages will be scheduled. The scheduling algorithm will take into account the size of the message, the priority of the message, and the speed of the subnet's transport medium. Clients wishing to receive geographic messages would then tune in to the appropriate multicast group to receive them.

Geographic routers (GeoRouter) are in charge of moving a geographic message from a sender to a receiver. GeoRouters are essentially routers which are geographically aware. Each router is charged with performing geographic routing functions for a small number of GeoNodes whose geographic areas of operation are contiguous. GeoRouters keep track of the geographic area that they service (called its service area) by calculating the union of the geographic areas covered by its GeoNodes. Its service area is represented as a single simple closed polygon whose vertices are denoted by geographic coordinates. GeoRouters build their routing tables by exchanging service area polygons. For scalability reasons, in order to reduce the size of the routing tables, GeoRouters are arranged in a hierarchical fashion with each layer corresponding to a distinct geographic area, such as a state or a city. For example in Figure 3, the campus routers for Busch and College Ave. both perform routing for two groups of three GeoNodes connected to wireless cells. Above them in the hierarchy is the county router whose job is to route messages between the local GeoRouters. In the current implementation, the hierarchy must be configured by human administrators. However, future research will examine methods of automatically configuring a hierarchical geographic routing network.

Sending a geographic message involves three steps: sending the message, shuttling the message between routers, and receiving the message. In order to send a geographic message, the programmer would use the Geographic Library routine `SendToGeo()`. The function will first contact the local GeoHost Daemon and query it for the IP address of the local GeoNode. It will then send the message directly to the GeoNode which, in turn, will simply forward the message to the local GeoRouter.

Once it receives a geographic message, a GeoRouter must first determine if it services any part of the area of the destination polygon. To do this, the router determines if the destination polygon and the router's service area polygon intersect¹ [7] each other. If not, then the router simply sends the message to its parent router. However, if the polygons intersect or contain each another, then the router does service the area described by the intersection polygon. First, if the polygons only partially intersect, then the router sends a copy of the message to its parent router. The router now tests each GeoNode's or child router's service area to see if it intersects or contains the destination polygon. Those that do will be sent the geographic message.

In Figure 3, a user on Busch Campus wishes to send a message to the destination polygon on the College Ave. Campus. The message is passed to the Busch Campus router. By using polygon intersection, the router determines that it does not service the target area, so it forwards the message to its parent, the county router. Using the same algorithm, the county router decides that its child router, the College Ave. router, services the destination area and forwards the message to it. The College Ave. router, in turn, forwards the message to the two GeoNodes which control the

¹Detecting polygon intersection takes $O(n)$ time.

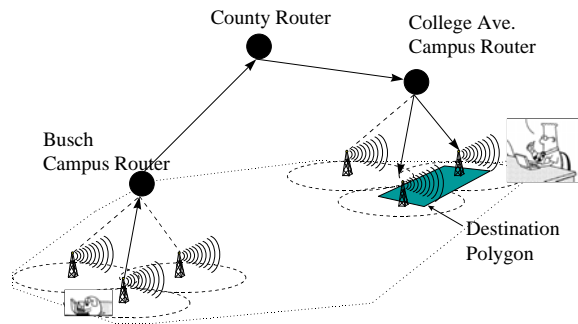


Figure 3: Geometric Routing Example

target area. These two GeoNodes then deliver the message to all of the users in the target area.

The router keeps a cache of the next-hop destinations of the most recent geographic message packets. When a router receives a geographic message packet, it will use the incoming packet's Message Id as a key into the cache. If this is not the first packet to arrive for this destination and if the timer on the cache entry has not yet expired, then the cache will return a list of all of the next hop addresses to which copies of the packet must be sent.

Once a geographic message has been sent to a GeoNode from a geographic router, the receive process can begin. The GeoNode will store the message locally and assign a multicast group to it. Periodically it will multicast the message to its multicast group and advertise its presence using ICMP messages on a well-known channel. The GeoHost daemons will receive the ICMP messages and determine if the host computer is located inside the message's destination polygon. When a client process executes a `RecvFromGeo()` call from the Geographic Library, the function will join the appropriate multicast address and receive the geographic message itself.

Geographic-Multicast Routing Method

Another geographic routing method, Geographic-Multicast Routing Method, leverages the power of multicast to transport the geographic messages to their destinations. This method uses the Domain Name Service (DNS) to determine which multicast address to use given a geographic region.

In this section, we will use the following two terms: *atoms* and *partitions*. Atoms are the smallest geographical areas which have geographic-multicast addresses. Partitions, on the other hand, are larger geographical areas which will also have a geographic address. A state, county, town, etc. may constitute a partition. Partitions and atoms are arranged in a hierarchical fashion. Each partition contains either a whole number of atoms or a whole number of smaller partitions.

The size and shape of the atoms and partitions would be determined by the density of subnets and wireless cells in a particular geographic area.

Each partition and atom will have a geographic-multicast address which will be used by routers. By a geographic-multicast address we mean that each partition and atom is mapped to a unique multicast address. Every GeoNode is required to join the multicast groups for the atoms and partitions which intersect its geographic range. Thus, a GeoNode needs to know not only its own range but also information about those partitions which intersect its range. This information can be obtained upon GeoNode installation or from the geographic database stored as a part of the DNS. The key idea here is to approximate the destination polygon with the smallest partition or atom which contains it and use the multicast address corresponding to that partition/atom as the address of that message. Since the partition/atom which is used is only an approximation to the destination polygon, some GeoNodes outside the destination polygon will erroneously receive the geographic messages. In order to counter this, the original destination polygon is inserted into the multicast packet body. The GeoNodes then use the destination polygon to determine if they should have actually received the message. If not, the message is simply ignored.

However, the multicast group information has to be propagated very carefully. Because of the large number of atoms and partitions and the resulting large number of multicast groups, we will modify the Protocol Independent Multicast Sparse Mode (PIM-SM) [4] which is slated by the Internet Engineering Task Force to become the future standard multicast protocol. PIM-SM is meant to be used in wide-area networks, networks which are bandwidth poor, or multicast groups which have few or widely-scattered members. It assumes that not everyone wants to receive the multicast packets and relies on explicit join messages from group members. As a result, PIM-SM has the advantage that it will only send multicast packets where they have been requested, and will not broadcast the initial packets as the current multicast protocol does. The PIM-SM protocol is similar to Core-Based Multicast Trees [1] in that it uses a Rendezvous Point (RP) to arrange for the senders and receivers of a multicast group to meet. This RP then also becomes the root of a sparse multicast tree with the multicast group members being the leaves of the tree. All senders ship their packets to the RP for distribution. The current PIM proposal calls for the RP to be selected by the first member of the group. Alternative RP's are also selected in case the primary RP fails. The PIM-SM specifications call for the multicast group address and its list of primary and alternative RP's to be broadcast to all PIM routers. However, we will change this specification by implementing the following intuition: *The smaller the size of the partition/atom, the more locally the information about that partition/atom is propagated.* Thus, only multicast group membership for very large partitions or atoms will be propagated across the whole country.

Domain Name Server Solution

In this subsection we sketch a solution which relies heavily on the Domain Name Service (DNS). Here, the geographic information is added to the DNS servers. These will provide the full directory information down to the level of the IP address of each GeoNode and its area of coverage.

A new first level domain - “.geo” is added to the set of first level domains. The second level domain names represent states, the third counties, and, finally, the fourth polygons of geographic coordinates. We can also allow polygons to occur as elements of second or third-level domains to enable the sending of messages to larger areas. Thus a typical geographic address can look like city-hall-Palo-Alto.San-Mateo-County.California.geo or Polygon.San-Mateo-County.California.geo

where Polygon is a sequence of coordinates.

This geographic address is resolved into a set of IP addresses of the GeoNodes which cover that geographic area. Depending on the size of the message, it may now be transported to the GeoNodes in one of two ways. If the message size is small, then it will be sent as a set of unicast messages to all of the GeoNodes corresponding to the addresses returned by the DNS. Alternatively, given a large message size, it is more efficient to first ask all of the GeoNodes to join a temporary multicast group for the geographic area specified in the message. The message content is then sent to that multicast group.

Geographic Email

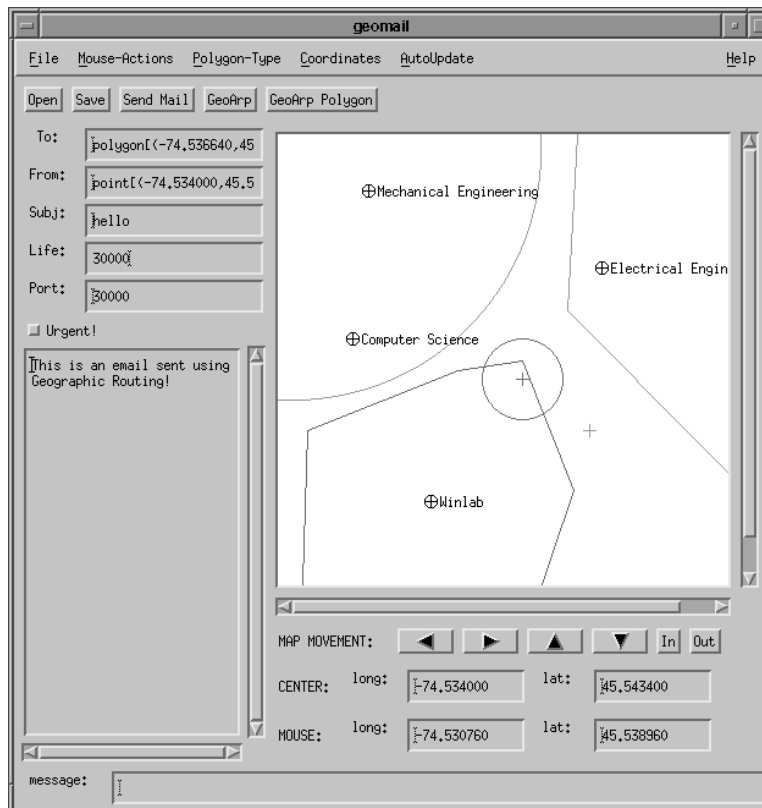


Figure 4: Geographic Email User Interface

The Geographic Email (GeoMail) application, shown in figure 4, demonstrates the use of geographic routing. It allows a user to send a text message to any geographic destination. The GeoMail program also displays all available geographic messages in relation to the user's current position. When the user's position intersects the destination polygon of a geographic message, then the GeoMail program will receive the message and display its contents.

Geo-multicasting

While geocasting is an important service, it is more likely that we will multicast rather than broadcast into the geographical areas. For example, we will be interested in reaching all *motorists* on a specific highway, or all *police cars*, rather than reaching *everybody*. This will be accomplished by geographically directed multicast. Both geocasting methods described before can be modified to accommodate geo-multicasting. The hierarchy of geo-routers can be used also to maintain information about multicast group memberships. Our multicasting solution can also be easily extended to handle arbitrary multicast groups in conjunction with partitions and atoms. Finally, other solutions are also possible, based on a concept of *area codes* analogous to the ones used in telephony today. Geo-multicasting will be described in more details in a forthcoming paper.

Geographic-Based Service Querying and Advertising

The future mobile users will need information that is pertinent to their location. Such information could include maps of the locale area, traffic and tourist information, as well as what restaurants or other establishments are available locally. With the rapid growth of the volume and diversity of data, the user of the future will find it increasingly difficult to discover or know in advance the correct servers to go to in order to obtain the location-dependent information he seeks. As businesses become directly connected to the Internet, the sought-after information sources will be collocated with the business that is providing it. In such a future, a geographic routing enabled Internet would allow for the existence of services based on geography or distance.

Geographic-based services would entail distributing or finding information within a geographical area. Note that since we assume that the information servers will be collocated with the individual or business that provides the information, we assume that the information contained in those geographically-close servers will contain information that is of greater relevance to the user's current location. Therefore, geographic-based services optimize the relevance of the information gathered rather than network resources used in order to obtain that information since the geographic proximity of the information does not imply that the information servers are nearby in terms of the network topology. Such geographic-based services would include advertising to a specific geographical region such as only within a certain distance from the server. A restaurant or store could advertise itself or its sales to travelers on nearby streets and roadways. Clients could also request services in a specific geographic region such as only within a certain distance from their current location. For example, a user could request local tourist information such as local maps or directions to monuments and buildings.

Acknowledgments

We would like to acknowledge the anonymous referee whose many insightful comments helped improve the flow and presentation of this paper.

Conclusions

We have addressed the impact which the universal deployment of GPS will have on various levels of Internet protocols, starting from the network layer (IP) and proceeding to the application layer (WWW). We have described several approaches to geographic routing, including the one which we

are currently implementing within the DARPA sponsored GloMo program. Finally, we discussed various approaches to service discovery based on location and distance.

Our assumption is that the *location* service will be as universal as the *date* service is today. Consequently, *distance* will become as important a search parameter on the World Wide Web as keywords are today. The geographic messaging project introduces location as a “first class citizen” both in message addressing as well as in service discovery.

The protocols described above are a part of the geographic addressing and messaging Request For Comments, published in November 1996 [6].

References

- [1] A. J. Ballardie and P. F. Francis and J. Crowcroft, *Core Based Trees*, Proceedings of the ACM SIGCOMM, San Francisco, 1993.
- [2] Deering, S., and R. Hinden, *Internet Protocol, Version 6, Specification*, RFC 1883, Xerox PARC, Ipsilon Networks, December 1995.
- [3] Deering, S., and Hinden, R., Editors, *IP Version 6 Addressing Architecture*, RFC 1884, Ipsilon Networks, Xerox PARC, December 1995.
- [4] Deborah Estrin et al., *Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification*, Internet Draft, March 1997. Work in progress.
- [5] G. Finn, *Routing and Addressing Problems in Large Metropolitan-scale Internetworks*, ISI Research Report ISI/RR-87-180, University of Southern California, March 1987.
- [6] T. Imielinski and J. Navas, *GPS-Based Addressing and Routing*, RFC 2009, Computer Science, Rutgers University, March 1996.
- [7] J. O'Rourke, C.B. Chien, T. Olson, and D. Naddor, *A new linear algorithm for intersecting convex polygons*, Computer Graphics and Image Processing 19, 384-391, 1982.
- [8] Rekhter, Y., and T. Li, *An Architecture for IPv6 Unicast Address Allocation*, RFC 1887, cisco Systems, December 1995.
- [9] — Mike Spreitzer and Marvin Theimer, “Providing location information in a ubiquitous computing environment”, 14th ACM Symposium on Operating System Principles, Vol 27, No 5, Dec 1993; also in “Mobile Computing”, Ed by T.Imielinski and H.Korth, Kluwer 1996. pp 397-423
- [10] —, *A Technical Report to the Secretary of Transportation on a National Approach to Augmented GPS Services*, [http:// www.navcen.uscg.mil/ gps/ reports/ reports.htm](http://www.navcen.uscg.mil/gps/reports/reports.htm)
- [11] —, *GPS SPS Signal Specification, 2nd Edition*, [http:// www.navcen.uscg.mil/ gps/ reports/ sigspec/ sigspec.htm](http://www.navcen.uscg.mil/gps/reports/sigspec/sigspec.htm), June 2, 1995.
- [12] —, *The USCG Differential GPS Navigation Service*, [http:// www.navcen.uscg.mil/ dgps/ dgeninfo/](http://www.navcen.uscg.mil/dgps/dgeninfo/)