

Anycast Routing for Mobile Networking

Vincent D. Park and Joseph P. Macker
Naval Research Laboratory
Washington, DC 20375

Abstract—This paper considers the problem of locating and forwarding network traffic to any one of a set of distributed servers or service points—primarily in the context of mobile ad hoc networks. The advantages of providing such a capability in mobile networks through the use of anycast routing techniques at the network layer are discussed. The results of a simulation study are highlighted to demonstrate how anycast routing techniques can provide a one-to-any communication capability with greater efficiency than traditional unicast based techniques. The simulation results also indicate anycast routing simplifies required configuration and management and reduces connection setup latency and overall message packet delay. Potential applications of anycast routing technology in military networks are presented and related issues are discussed.

INTRODUCTION

A communication paradigm, known as *anycasting*, has recently been introduced within the networking community—primarily for the purposes of locating distributed services [1]. Anycasting essentially provides a means to locate and communicate with *any one* of a set of distributed servers or service access points within a network. This is analogous to directing an individual that needs to make a phone call to a public payphone. While there are potentially many points of service, the end user only needs to find one. In a networking context, anycasting facilitates more robust distributed system design and eases network configuration and management under a variety of scenarios.

Most related research to date has focused on the development of anycast techniques at the application layer [2, 3]. However, we believe greater communication efficiency and robustness may be realized through the use of anycast routing techniques at the network layer. While gains are realizable in quasi-static hardwired networks, they are of more critical importance in mobile wireless networks—which have more dynamics (e.g., rapid and unpredictably changing interconnectivity between routers) and are more bandwidth constrained than traditional hardwired networks. A far-forward military network is perhaps the best example of such an environment. Although a fixed infrastructure with hardwired links may form part of the networking infrastructure, significant portions of the network/internetwork will comprise mobile nodes and platforms relying on wireless communications.

While renewed research interest and progress is being made in the area of unicast (one-to-one) and multicast (one-to-many or many-to-many) routing for mobile ad hoc networks [4, 5], locating and managing mobile services for end users in such networks remains a largely unexplored topic. Anycast (one-to-

any) routing helps support and manage this required functionality. Within static networks, critical networking services are often centralized or distributed with preconfigured lists creating a fundamental adaptation, robustness, and location problem. Without robust mobile support for such networking services, end systems are severely handicapped in functionality and performance regardless of the available network connectivity or bandwidth. The prevalence of performance degradation, global scale, and potential denial of service in today's static network infrastructures has spawned a flurry of recent developments in more distributed network databases and services. Even when such distributed services are available problems of service location, transaction, and data collection are exacerbated by the full or partial inclusion of mobile network architectures. In these scenarios, the concept and use of anycast routing technology provides an important service enhancement by efficiently providing robust distributed location and collection services for the end users and easing mobile network configuration burdens.

ANYCAST ROUTING


Functionally, anycast routing in a datagram network can be described as follows. An anycast address corresponds to multiple receivers (or service points) within a given routing domain. The anycast routing protocol must establish and maintain the information necessary for forwarding datagrams with the anycast address to one of the corresponding receivers. Conceptually, the forwarding information maintained by the anycast routing protocol defines a “service area” for each receiver and dynamically adjusts these service areas to accommodate changes in the receiver set, network topology and potentially other aspects of the networking environment, Fig. 1.

In an earlier work [6], we described how the link-state, distance-vector and link-reversal classes of unicast routing protocols can be extended to provide efficient construction and maintenance of anycast routes. The techniques developed are readily adaptable to many existing routing protocols—e.g., Open Shortest Path First (OSPF) [7], Routing Information Protocol (RIP) [8] and the Temporally-Ordered Routing Algorithm (TORA) [9, 10]. While there are still open issues regarding the use of anycast routing in Internet Protocol (IP) based internetworks [1], our extended routing functionality approach has minimal impact on present open networking standards. Thus, the techniques provide a viable solution for anycast routing that is complementary to existing unicast and multicast routing.

Advantages of Anycasting in Mobile Networks

Anycast technology and related dynamic routing functionality provide significant improvements to mobile network architectures. A major difficulty that exists in today's military and commercial mobile networks is managing mobile nodes and

This work was supported by the Office of Naval Research under contract number N0001499WR20017.



services under dynamic changing conditions. Also, with progress of internetworking technology, more distributed services are being deployed and relied upon by end users and applications. Anycasting helps provide a robust means of dynamically managing the end user requirement of finding “one service point out of a set”.


Modern networks often include multiple levels and types of distributed services and applications that end users need to periodically contact, potentially exchange data with, and/or continuously provide data reports to. These services and applications may provide such functions as security key management, application directory services, name/address resolution or data collection and fusion. Wireless and mobile users more frequently experience partial network outages and dynamic topology changes making the job of locating and contacting alternate service points more difficult. Anycast routing provides a robust mechanism for communicating with distributed services and applications under such dynamic networking conditions. Reattachment of mobile users and end systems to new portions of a network also requires the dynamic use of alternate service points to maintain service and/or to optimize the network usage (e.g., minimize delay). Anycast routing provides a major benefit in the aforementioned situations by allowing the network routing to do most of the work in locating and tracking necessary distributed network services for end systems—thus, easing network configuration and management burdens. As mentioned, much of this can be accomplished with only minor changes to existing unicast and multicast routing technology already in place or planned for these environments.

Anycast Routing Simulation Study

While some of the primary advantages of anycast routing in mobile ad hoc networks are the potential to ease network configuration and facilitate more robust distributed system design, there are possible performance gains as well. We conducted a limited simulation study using the Optimized Network Engineering Tools (OPNET) to demonstrate how anycast routing techniques can provide a one-to-any communication capability with greater efficiency and robustness than traditional unicast based techniques. Herein we provide a brief overview of the results to illustrate the performance advantages of anycast routing.

We modeled the network using a *fixed* baseline topology with the ability to control the failure/recovery of individual links. This relatively simple model provides sufficient control of the networking environmental characteristics (e.g., rate of topological change, average network connectivity) and also permits simulation of relatively large networks in a reasonable time. For a given baseline network topology, each link in a given network continuously cycled between two states (ACTIVE and INACTIVE) independently of all other links. Once ACTIVE, the time a link remained ACTIVE was randomly determined based on an exponential distribution. The mean of the distribution (“mean-time-to-failure,” $1/\mu$) was an input parameter of the simulation. Essentially, a lower link mean-time-to-failure corresponded to a higher rate of topological change. The long-term average fraction of time each link would be operational, f , was also a simulation input parameter. Variation of this parameter affected the average overall network connectivity (i.e., when $f = 0.2$, on average 20% of the links in the baseline topology are operational at any given time). The parameter f was also used to determine the initial state of each link at the beginning of each simulation execution. Once INACTIVE, the time a link remained INACTIVE was also determined randomly by an exponential distribution. However, the mean of the distribution (“mean-time-to-repair,” $1/\lambda$), was computed from $1/\mu$ and f . The state transition diagram for this continuous-time Markov process, and the equation by which $1/\lambda$ is computed, are presented in Fig. 2. Operational links in the fixed topology simulation essentially represent radio connectivity between node pairs in a wireless network.

We implemented anycast extensions to link-state routing and compared its performance to a mechanism based solely on unicast routing. Nodes that were designated as providing the anycast service would receive packets sent to either their unique unicast address or the anycast address. When a node generated a packet destined for the anycast service, the packet was forwarded using the anycast address.



We compared this anycast routing technique to the use of unicast routing with the destination selected based on a prioritized list of the nodes providing the anycast service. Thus, each node maintained a prioritized list with the unicast addresses of the nodes designated as providing the anycast service. When a node generated a packet destined for the anycast service, it queried the list to determine the unicast address of the highest priority server to which a valid route was available and forwarded the packet using that unicast address.

The following results are based on a sequence of simulations conducted to investigate the performance tradeoffs as a function on network connectivity. The baseline topology was defined by a “complete” graph of 20 nodes (i.e., each node was connected to every other node by a direct link). Since this allows for the possibility that any two nodes may be able to communicate directly at some point in time (i.e., when the link between them is ACTIVE), this is perhaps the best representation of a mobile network given the limitation of our fixed baseline topology model. The fraction of time operational, f , was varied from 0.02 to 0.25 for successive simulation runs in the sequence—while the link mean-time-to-failure, $1/\mu$, was kept constant at 60 seconds. Three nodes were selected as anycast destinations (i.e., nodes providing an anycast service) and all other nodes in the network randomly generated message packets for the anycast service.

The entire simulation sequence was executed first using unicast routing (as previously described) to forward message packets to the nodes providing the anycast service and then repeated using anycast routing. For each simulation run the two approaches were subjected to an identical sequence of random events (e.g., topological changes and packet arrivals).

The amount of additional routing control traffic due to the anycast extensions was measured during each simulation run. In each case, the anycast extensions increased the number of routing control packets by approximately one to two percent. As expected, this increase corresponds approximately to the percent increase in the number of links represented in the link-state database (i.e., three virtual links added to the 190 physical links in the baseline topology).

While the anycast extensions increased the bandwidth utilization for routing control traffic, there was also a reduction in the bandwidth utilization for message traffic. The reduction in bandwidth utilization for message traffic was realized because message packets forwarded based on the anycast routing technique were delivered to the destination using shorter paths on average. The mean message packet hop count for both the anycast and unicast routing techniques is plotted as a function of average network connectivity in Fig. 3. The plot clearly illustrates the mean number of hops (i.e., transmissions) required for message delivery using anycast routing is less. While the anycast routing technique forwards to the nearest node providing the anycast service, the unicast technique forwards based on the server priority list and path availability to the servers. Thus, the unicast technique will forward to the primary server (if a path is available) despite the fact that the secondary or tertiary server may be available via a shorter path. Depending on the traffic load and networking environment, the reduction in bandwidth

utilization due to the use of shorter paths may significantly outweigh the increase due to the additional routing control traffic.

We also collected statistics regarding the availability of paths from traffic sources to the nodes providing the anycast service. For both the anycast and unicast techniques, upon generation of a message packet, if an available route could not be determined by the source (i.e., no next-hop forwarding information in the routing table), the packet was discarded. For the unicast routing technique, the source would check route availability in the order specified by the priority list and forward using the first valid route determined (i.e., the highest priority server for which valid next-hop forwarding information was available). In all cases when a valid route for a given destination was not available a statistic was collected.

Fig. 4 illustrates the route availability statistics for the unicast routing technique. This plot can be interpreted as follows. The lowest curve approximates the probability that a route is available to the primary server. The middle curve approximates the probability that a route is available to the secondary server, given that a route is not available to the primary server. Finally, the highest curve approximates the probability that a route is available to the tertiary server, given that a route is not available to either the primary or secondary server. Although not included on the plot, the route availability for anycast routing was essentially equivalent to the highest curve depicted. This illustrates the improvement in robustness achieved by increasing the number of nodes providing the anycast service. It also illustrates the difference in robustness that would be seen if the unicast routing technique were to only maintain a partial list of the nodes providing the anycast service. If the number of nodes providing the anycast service is large or the set of nodes is dynamic, maintaining a complete list of anycast servers at all nodes will be complex and potentially impractical. The anycast routing technique provides a mechanism to maximize the robustness with minimal configuration and management.




Fig. 3. Mean message packet hop count as a function of average network connectivity.

Finally, we combine these results in consideration of message packet delay and the effect on higher-layer protocols. The mean message packet delay for both the anycast and unicast routing techniques is plotted as a function of average network connectivity in Fig. 5. The difference in delay corresponds to the difference in hop count for the two approaches. That is, message packets forwarded based on the anycast routing technique experience less delay because they were delivered using shorter paths on average.




Fig. 4. Route availability as a function of average network connectivity.




Fig. 5. Mean message packet delay as a function of average network connectivity.

The third curve in Fig. 5 illustrates the potential effect that route availability has on higher-layer protocols. This curve was generated by adjusting the mean message packet delay for the unicast routing technique based on the route availability statistics depicted in Fig. 4 and an approximation of a retransmission timer. The retransmission timer was approximated as $2(\eta+2\sigma)$, where η is the mean message packet delay and σ is the standard deviation of the message packet delay. This illustrates the additional delay that may be experienced for connection setup or reliable packet delivery when route availability is not known or is not signaled to higher-layer protocols. The retransmission timer approximation is quite conservative; thus, in many applications the retransmission timers may be much larger—resulting in much larger delays than those depicted in Fig. 5.

APPLICABILITY TO MILITARY NETWORKS

It is clear that anycast technology has a role to play in future military networks, especially within heterogeneous and mobile architectures. This section further discusses issues regarding some specific applications of anycast routing.

Robust Distributed Service Location

Within future military networks, there is a requirement to provide survivable services within potentially highly dynamic environments. Distributed services help provide decentralized control resulting in more robust and scalable service provisioning across a network. However, a problem of managing and locating distributed services without some *a priori* configuration is non-trivial. Military networks require the ability smoothly operate through dynamic changes and failures with a minimal amount of user intervention and management required. Anycast routing technology helps achieve this goal by providing a robust communication mechanism that supports distributed service and application designs.

Information Location, Retrieval and Collection

Another set of related military network requirements is the distributed location, retrieval, and collection of information. Anycast technology helps these goals in several respects. First, anycast routing provides a simple means for locating information that may be distributed at a higher network layer across a set of services and/or applications. Second, anycast routing facilitates information retrieval by providing server location so that subsequent data retrieval requests and/or long lived connections can be established. A simple example of location and retrieval is the desire for a mobile network node to use a distributed directory or security-related service within a mobile architecture. Third, anycast routing provides direct functional support for dynamic, distributed network data collection. End users and/or systems can continuously provide data input to a higher-layer distributed data collection and fusion application by simply forwarding their data input towards the appropriate anycast address. Dynamic anycast routing will forward this data to “one of a set” of designated anycast endpoints. Other complementary technology (e.g., multicast) can play a role in maintaining connection between distributed anycast endpoints. Fig. 6 shows an example of how anycast technology can help support several important military network functions.




Fig. 6. Military Architecture Example

Routing to Gateways between Mobile Networks and Hardwired Infrastructure

Finally, another potential application of anycast technology is in supporting routing within mobile networks to gateways that serve as fixed network access points. Within mobile networks, end users and nodes may migrate across an architecture with the desire to sustain communications with the larger external network. Mobile nodes may form multi-hop ad hoc routing relationships with their neighbors but there likely remains a subset of nodes within the architecture that serve as fixed or pseudo-fixed access point to the more extensive non-local network. The ability to locate and dynamically migrate across different gateways and access points could be supported through the use of specific anycast addresses assigned for this purpose. In a sense, anycasting can provide a “default” or “external” route within a mobile network to set of gateways providing access to a hardwired infrastructure.

CONCLUSIONS

The anycast communication paradigm functionally provides the capability to locate and forward network traffic to any one of a set of distributed servers or service points that provide *equivalent* service. Such a mechanism facilitates more robust distributed system design, which will likely be critical in military networks. While there are many possible approaches to providing an anycasting capability, the use of anycast routing algorithms is perhaps the best-suited approach for the mobile wireless networking environment. It is more communication efficient and requires less configuration and management of end systems than most application-layer approaches.

Several different classes of unicast routing protocols can be extended to provide efficient construction and maintenance of

anycast routes. The techniques are readily adaptable to many existing networking technologies and provide an elegant solution for anycast routing that is complementary to existing approaches for both unicast and multicast routing.

The performance aspects of anycast routing have been compared to traditional unicast routing based techniques. We have shown that, depending on the traffic load and networking environment, the use of anycast routing can reduce the overall bandwidth utilization by forwarding message traffic over shorter paths. The simulation results also indicate that anycast routing can ease the configuration and management required to achieve a given level of robustness and can reduce connection setup latency and message packet delay.

While anycast routing has benefits even in quasi-static hardwired networks, the realizable gains are of critical importance for more dynamic networking environments such as a mobile ad hoc network. Anycast routing provides needed functionality in military networks, where there is a requirement to provide survivable services within a potentially highly dynamic environment. Although, open issues remain regarding the further adaptation and use of anycast routing in Internet Protocol (IP) based internetworks; the technology is readily applicable and should be further studied and developed.

REFERENCES

- [1] C. Partridge, T. Mendez and W. Milliken, “Host Anycasting Service,” *Internet RFC 1546* (November 1993).
- [2] J. Guyton and M. Schwartz, “Locating Nearby Copies of Replicated Internet Servers,” *Proc. SIGCOMM '95* (August 1995).
- [3] S. Bhattacharjee, M. Ammar, E. Zegura, V. Shah and Z. Fei, “Application-Layer Anycasting,” *Proc. INFOCOM '97* (April 1997).
- [4] S. Corson and J. Macker, “Mobile Ad hoc Networking (MANET): Routing Protocol Performance Issues and Evaluation Considerations,” *Internet RFC 2501* (January 1999).
- [5] Working documents of the Internet Engineering Task Force (IETF) Mobile Ad Hoc Networks working group, <http://www.ietf.org/html.charters/manet-charter.html>.
- [6] V. Park and J Macker, “Anycast Routing for Mobile Services,” *Proc. Conference on Information Sciences and Systems (CISS) '99* (January 1999).
- [7] J. Moy, “OSPF Version 2,” *Internet RFC 2328* (April 1998).
- [8] C. Hedrick, “Routing Information Protocol,” *Internet RFC 1058* (June 1988).
- [9] V. Park and M.S. Corson, “A Highly Adaptive Distributed Routing Algorithm for Mobile Wireless Networks,” *Proc. IEEE INFOCOM '97* (April 1997).
- [10] V. Park, J. Macker and M.S. Corson, “Applicability of the Temporally-Ordered Routing Algorithm for use in Mobile Tactical Networks,” *Proc. IEEE MILCOM '98* (October 1998).