Pseudo-IP: Providing a Thin Network Layer Protocol for Semi-Intelligent Wireless Devices

Kevin C. Almeroth
Dept of Computer Science
University of California
Santa Barbara, CA 93106-5110
almeroth@cs.ucsb.edu

Katia Obraczka
Information Sciences Institute
Univ of Southern California
Marina del Rey, CA 90292
katia@isi.edu

Dante De Lucia
Computer Science Department
Univ of Southern California
Los Angeles, CA 90089-0781
dante@usc.edu

ABSTRACT

In the near future users will be able to move freely and still have seamless network and Internet connectivity. We envision that the Internet of the future will interconnect mobile or stationary clouds into the existing IP infrastructure. While many of the Internet protocols are immensely successful in traditional networks, we believe they will be inappropriate for communication among limited-capability devices in amorphous clouds. The question we are trying to address is whether the network layer services provided by IPv4 and IPv6 [1] are necessary and sufficient for supporting heterogeneous devices in these highly dynamic, arbitrarily dense environments. The problem with an IP infrastructure is that, for certain applications, it adds unnecessary complexity. The proposed research is based on the specification of a new network layer protocol for intra-network communication in clouds containing devices with limited processing and communication capabilities. Our *Pseudo-IP* protocol is designed to operate among devices in the farthest branches/leaves of an internet while providing inter-network connectivity with other clouds and the existing IP infrastructure. Our goal is to extend the scope of IP to environments containing devices that cannot handle the extra complexity introduced by routing, error detection/recovery, optional headers, and even addressing. More intelligent devices interacting in the cloud will be responsible for interoperating with the existing IP infrastructure including IP-based devices that happen to be roaming locally.

1. Introduction

In the near future users will be able to move freely and still have seamless network and Internet connectivity. Portable computers and hand-held devices will be for data communication what cellular phones are now for voice communication: they will keep users connected at all times. In addition to being continually connected over time, the concept of “universal connectivity” also means that a variety of “unconventional” devices will be connected to the Internet. A variety of devices, including sensors, home appliances, light switches, etc., will be interconnected forming *clouds*. We envision that the Internet of the future will interconnect these (mobile or stationary) clouds and the existing IP infrastructure. As users roam among clouds, they will encounter, and be required to communicate with, a range of devices varying in processing, mobility, and communication capabilities.

While many of the Internet protocols are immensely successful in traditional networks, we believe they will be inappropriate for communication among limited-capability devices in amorphous clouds. The question we are trying to address is whether the network layer services provided by IPv4 and IPv6 [1] are necessary and sufficient for supporting heterogeneous devices in these highly dynamic, arbitrarily dense environments. The problem with an IP infrastructure is that for certain applications it adds unnecessary complexity and overhead. One example scenario is a cloud of thousands of sensors transmitting small pieces of data. The data portion, assuming traditional IPv4 or IPv6 headers and lower layer protocol headers, will be a very small percentage of the overall packet size. Besides the inefficiencies of payload bytes versus header bytes, there is also the issue of the devices' limited processing and communication capacity. The numerous sensor devices may simply transmit their data and the overall system would rely on a roaming user carrying a more intelligent device. Alternatively, the sensor cloud could contains one or more data collection devices the roaming user's computing device is able to locate and query. Whatever the scenario, there will likely be devices with limited functionality that still have important data to communicate.

The premise of this paper is to introduce a new network layer protocol for intra-network communication in clouds containing devices with limited processing and communication capabilities. Our *Pseudo-IP* protocol is designed to operate among devices in the farthest branches/leaves of an internet while providing inter-network connectivity with other clouds and the existing IP-based Internet infrastructure. Our goal is to extend the scope of IP to environments containing devices that cannot handle the extra complexity introduced by routing, error detection/recovery, optional headers, and even addressing. More intelligent devices interacting in the cloud will be responsible for interoperating with the existing IP infrastructure including IP-based devices that happen to be roaming locally.

While there has been a great deal of interest in wireless protocols, much of the work has focused on providing higher bandwidth and more sophisticated communications services. Our approach is directed towards simplification. We believe that our Pseudo-IP protocol will be an important enabling technology for the Internet of the future since it will allow limited-capability devices to be connected to the existing IP infrastructure.

One research initiative related to the proposed ideas is the Daedalus/BARWAN project [2] at UC Berkeley. More specifically, they have proposed an architecture that supports adaptive client device's functionality to new services that are discovered/located as the client moves [3]. Their architecture is based on the existing IP infrastructure. Other research initiatives in this area are beginning and should be jump started by DARPA interest and funding.

The remainder of this paper is organized as follows. Section describes a number of potential applications and their requirements. Section gives an overview of the concept behind Pseudo-IP. Section summarizes the research challenges and how they relate to our Pseudo-IP protocol.

2. Applications and Requirements

The motivation for Pseudo-IP is grounded in a number of cloud-based applications and scenarios. We describe some of these scenarios and their specific protocol requirements below.

- **Home spaces**: Homes of the future will be full of devices ranging in intelligence from dumb to semi-intelligent to fully-intelligent devices. Some examples include dumb devices like light switches; more intelligent appliances with embedded circuitry like TVs and microwaves; and programmable, network-aware devices such as PCs. The least intelligent devices may only be capable of broadcasting state information. Some may have additional functionality which allows them to receive and process commands and then change state. The most intelligent devices would be responsible for state collection, device control, and management. Consider the commonly referenced example of a fully connected home. All electronic devices are connected together in a controllable network. Voice commands are detected, translated, and executed through some voice recognition system. Many household devices will have wireline connectivity but some may be wireless. In either case, these devices must be capable of receiving, executing, and acknowledging commands. Another important service required in this scenario is security. It will likely be an important issue, especially from from the need to authenticate commands and prevent unauthorized users outside the home from controlling devices in the home.

- **Highway spaces**: Highway environments offer a slightly different set of requirements from the home environment. Objects in the highway environment are likely to be more intelligent but may be more transient. Current projections suggest vehicle-based communications systems providing services other than cellular telephony will be able to communicate using IP. So in several of the highway scenarios, many of the key objects will be capable of speaking IP. But because of the limited range capabilities of these devices, an intra-cell network protocol will likely not require IP-style services. One common application consists of drivers receiving information about road conditions, restaurants, shopping, service plaza amenities, etc. via broadcast transmissions. This type of data can be broadcast repeatedly so reliability (other than through non-ACK based techniques like forward error correction[4]) is not required. However, some applications may provide transaction-based services like ordering food before arriving at a highway exit. In these scenarios, the network protocol would have to facilitate reliable, secure transactions.

  A second set of applications in the highway environment is more similar to those discussed above for the home. Dumb, sensing devices unique to a highway environment might provide data about traffic conditions like congestion, flow patterns, weather, etc. These devices would likely only need to continually transmit a best-effort sampled data stream. These sensor devices may need to execute basic commands or may simply transmit a continuous flow of data.

- **Inhospitable environments**: There are a number of inhospitable environments that would benefit from an array of very simple devices that use a lightweight network protocol to communicate. Generally, we are talking about a large number of wireless sensors spread over an area to provide continuous feedback. The key difference between this class of applications and the others is that the devices would have to self-organize into a network and work together to communicate necessary information to points on the periphery of the network. A prototypical application is the blanketing of a battlefield with thousands of sensors. Sensors would collect and communicate reconnaissance information to the edges of the cloud where some intelligent agent would gather, possibly process, and likely relay the information through some traditional network. To be truly successful sensoring devices would have to be simple, cheap, and plentiful.

Other applications included in the inhospitable environments class of applications are conditions monitoring, disaster relief assistance, and search and rescue efforts. For example, seismic sensors could be scattered over a collapsed building and used to detect the motion of trapped survivors. Devices might even be installed into the building infrastructure during construction and used for search and rescue efforts in a collapsed building if it is destroyed by disasters like earthquakes or fires. These devices might also server to collect nominal environmental data like building air quality, temperature, etc.

3. Overview of the Pseudo-IP Protocol

The goals of Pseudo-IP are (1) to reduce the overhead and complexity of a full network layer protocol, (2) be flexible enough to interoperate with different medium access layer protocols, and (3) be flexible enough to provide network service in a variety of environments. Obviously the most common network layer protocol is IPv4[1]. In considering what functionality Pseudo-IP should provide, we should examine what functions IPv4 provides. They include the following[5]:

- Packet length – 1 bytes
- Identification/Sequence number – 2 bytes
- Fragmentation/reassembly – approximately 2 bytes
- Time to live – 1 byte
- Upper layer protocol identifier – 1 byte
- Header checksum – 2 byte
- Source and destination addressing – 8 bytes
- Miscellaneous other bits – 2 bytes
- Options and variable length headers – variable

In addition to the overhead associated with the IP header, there is processing overhead required to implement protocol functionality. For example, to properly support IP, the Internet Control Message Protocol (ICMP) should be implemented. Furthermore, protocols to provide translation between medium access control layer addresses and network layer addresses requires two resolution protocols: the Address Resolution Protocol (ARP) and the Reverse Address Resolution Protocol (RARP). And finally, all of this overhead is in addition to whatever overhead is required by the medium access control protocol. If devices implement a wireless protocol like the IEEE 802.11 wireless LAN standard, fewer Pseudo-IP functions will be required because 802.11 provides its own addressing and checksum mechanisms; uses collision avoidance; and has provisions for acknowledgments[6].

Other, simpler medium access control protocols might have to be used like Aloha[7] and slotted-Aloha[8]. Specifically then, what Pseudo-IP should provide is (1) a lightweight interface for communication among dumb or semi-intelligent devices, and (2) protocol translation between Pseudo-IP and traditional IP. Much of the routing functionality will be based on radio transmissions, i.e., the fact that these unconventional devices are usually broadcast-capable.

4. Research Challenges

The goal of Pseudo-IP is to provide basic network layer functionality while still allowing higher layer protocols to provide services like reliability, congestion control, authentication, etc. Dumb devices should only have to implement the minimum number of functions to achieve connectivity. Furthermore, simple devices should not incur overhead penalties for functions they cannot or do not wish to perform. Our research plan is based on creating a lightweight network layer protocol. Conceptually, our Pseudo-IP protocol can be compared to the relationship between UDP and TCP. UDP, when compared to TCP, is a lightweight protocol providing almost no transport layer services.

Pseudo-IP will eliminate most of the fields of both IPv4 and IPv6. We will investigate the issues raised by having no addressing, no routing information, no fragmentation/reassembly function, no error detection, and no sequence numbers. The basic paradigm for this simplest case will be random broadcast of data. Packets will still have a Medium Access Control (MAC) protocol which will provide framing, and probably some form of identification and basic error detection.

The research challenges we plan to explore are associated with the issues raised by providing communication using Pseudo-IP. Part of this challenge includes investigating how to build additional network services, like reliability for control functions, on top of Pseudo-IP. A second challenge is how to interconnect Pseudo-IP clouds with the existing IP infrastructure. A brief description of some specific research issues

[1] We also will consider how IPv6 differs from IPv4 but the philosophy of IPv6 is similar enough that we can concentrate on IPv4 at this point.

include the following:

- **Data Flow**. Straightforward data flow should likely only require the simplest version of Pseudo-IP. For example, in the case of simple sensors, data will flow one way, not even requiring return information or feedback to the sensor. Sensors will periodically broadcast their sensor information and not care if it is ever received. These devices should be very cheap compared to an IP-capable device. Given the potential environments, there are two types of network topologies that these types of devices might have to communicate in. The first topology would not require any routing because all devices can communicate directly with the desired receiver. Devices either have a wired connection to the receiver or operate in a wireless environment where the receiver is known to be in range. A medium access control protocol would be responsible for implementing collision avoidance functionality.

  The second topology assumes that all data should be delivered to a single remote receiver and not all transmitters are within range of the receiver. This type of topology requires basic routing functionality and represents a significant jump in complexity. The additional complexity includes the following components:

  - How to do routing? Given that devices may be expected to perform in inhospitable environments, the network topology may actually change frequently. Furthermore, running a complex route discovery protocol is unlikely to be feasible given the nature of the devices. Our preliminary assertion is that some sort of optimized random routing or intelligent flooding algorithm should be used. A second consideration associated with broadcast-based routing is the need to remove old packets from the network. IP uses a monotonically decreasing time-to-live (TTL) field that causes a packet to be discarded when the TTL value reaches 0.
  - Addressing can range in importance from critical to not necessary. For some applications, like a blanket of sensors dropped in an inhospitable environment, the actual location of a device may need to be known. Sensors may need a GPS-based locating system.
  - In some cases, strict timing information will be required. Time stamps might have to be taken and then used as sequence numbers. The medium access control protocol might provide some part of this function, for example through a hardware address or through a fully pre-configured arrangement like Time or Frequency Division Multiplexing (T/FDM).

- **Semi-Reliable Feedback**. One potential problem with not having a way to transmit feedback to a large set of dumb sensors is the total lack of control a management station would have. The problem occurs as more and more sensors are brought on-line and/or when using more capable devices (e.g., mobile sensoring devices). The period between each sensor's transmissions may be too short and a large number of collisions may result reducing the effective data rate. A control station might want to communicate with the array of devices using some very simple feedback channel. In this example, the control station might select a specific interval to broadcast to all sensors. These semi-reliable control functions can be achieved by again using a broadcast paradigm. By repeating the broadcast multiple times, all or most of the sensors will eventually receive the control information.

- **Reliable Transactions**. There will likely be intelligent enough devices that will want to exchange information reliably. For instance, a host might want to reliably control a light switch, or other simple device in a room. In order to do so, some information needs to be passed between the host and the device. In IP-based communication, this is accomplished using a series of messages. In Pseudo-IP, this could be accomplished by exploiting MAC layer mechanisms such as TDMA slots and MAC-level ACKs/NACKs. While this breaks the traditional model of layered network protocol design, it greatly enhances the ability to accommodate devices that are only capable of sending small messages. The idea is to perform authentication, sequencing and reliability based on lower-layer mechanisms, rather than relying on extra network layer bits.

- **System Control**. Control functions are built on top of reliable transactions but the additional challenge is determining what parameters are available for control and identifying a way of communicating the control function. Ideally, we would like not to have to define a standard for information exchange (e.g., identifying that information is coming from a light switch and not the dish washer). Standards are susceptible to politics, ineffi ciencies due to aging, and additional problems of interacting with devices that do not support the standard or may support different versions of it.

- **Security**. Reliable transactions require secure channels. End-to-end encryption can be used since devices engaging in transactions are likely to be more powerful in terms of processing and communication capabilities. When dealing with the more simple devices, physical security may be the only option. In the home environment, for instance, devices could rely on line-of-sight communication, requiring physical proximity to the devices. Remote control could be enabled through the use of an intelligent IP gateway which could provide remote authentication services.

- **Inter-Cloud Routing**. In order to allow communication between simple devices in different clouds, edge devices would act as intelligent gateways. Although devices might not be able to address other devices directly, gateways could collect Pseudo-IP packets and encapsulate them in IP packets for remote distribution. Edge devices could have static IP addresses, while intra-cloud communication is via local dynamic addressing.

- **Directory Services**. The problem of discovering local services in an area becomes problematic as the number of devices grows. Since bandwidth is limited, there must be some way to gather information without consuming all the bandwidth with service announcements. This can be accomplished by proving directory servers in a region. Although this should not be a requirement of the system (devices should still disseminate their status periodically) they can provide an optimization when improved network performance is required. Directory server would collect data on the local region and provide this information upon request. For example, and vehicle entering a cloud could request information on available services from a directory of services in the area, rather than having to wait for all services to announce their availability. If such a directory server were not available, the vehicle would have to wait longer to obtain such a list, but it could still be obtained.

One problem in designing a protocol like Pseudo-IP is the uncertainty in knowing what specifications the devices meant to use Pseudo-IP will actually have. There are questions about processing capability, bandwidth, transmission range, storage capacity, duration of operation, etc. In addition to device specifications, there are questions about the environments these devices will have to operate. There are also questions about the type and size of data collected, the real-time requirements of data delivery, number of devices in a region, environmental hazards, etc. Our proposed research agenda will focus on designing a number of detailed scenarios and then specifying a protocol to most efficiently and effectively address the network needs.

References

1. S. Deering and R. Hinden, "Internet protocol, version 6 (IPv6) specification." Internet Request for Comments RFC 1883.
2. Daedalus Project Team, "The daedalus project home page." http://daedalus.cs.berkeley.edu/index.html, January 1996.
3. T. D. Hodes and R. H. Katz, "Composable ad-hoc location-based services for heterogeneous mobile clients." Submitted for review, Wireless Networks Journal special issue; also available from http://daedalus.cs.berkeley.edu/publications/services-WINET.ps.gz, November 1997.
4. J. Nonnenmacher, E. Biersack, and D. Towsley, "Parity-based loss recovery for reliable multicast transmission," in *ACM Sigcomm 97*, (Canne, FRANCE), August 1997.
5. A. Tanenbaum, *Computer Networks, 3rd Edition*. Upper Saddle River, New Jersey: Prentice Hall, Inc., 1996.
6. V. Hayes, "IEEE standard for wireless LAN medium access control (MAC) and physical layer (PHY) specifications," Tech. Rep. IEEE 802.11-1997, Draft 6.1, IEEE Computer/Local & Metropolitan Area Networks Group, June 1997.
7. N. Abramson, "Development of the ALOHANET," *IEEE Transactions on Information Theory*, vol. IT-31, pp. 119–123, March 1985.
8. L. Roberts, "Extensions of packet communication technology to a hand held personal terminal," *Proceedings of Spring Joint Computer Conference, AFIPS*, pp. 295–298, 1972.