A Lightweight Protocol for Interconnecting Heterogeneous Devices in Dynamic Environments

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 people will be able to move freely and still have seamless network and Internet connectivity. One implication of this new style of interaction is an increased diversity in the types of devices which will be used to maintain connectivity. While many of the Internet protocols have proven successful and long-lived in traditional networks, we believe they will be inappropriate for use in communicating among limited-capability devices. In this paper, we introduce a new network layer protocol for intra-network communication in clouds containing devices with varying processing and communication capabilities. Our proposed protocol, called Pseudo-IP, is designed to operate among devices in the farthest branches/leaves of an intranet while providing inter-network connectivity with other clouds and with the existing IP-based Internet infrastructure. Pseudo-IP will accommodate a wide range of devices with varying power, processing, and communication capabilities, while supporting a variety of applications.

1. Introduction

In the near future people will be able to move freely and still have seamless network and Internet connectivity. Portable computers and hand-held devices will do for *data* communication what cellular phones are now doing for *voice* communication. Users will be able to roam and still be connected. One implication for this new style of interaction is an increased diversity in the types of devices which will be used to maintain connectivity. Furthermore, not only will people and their devices become more closely connected, but additional, unconventional devices will begin to be connected as well. A whole variety of devices, including sensors, home appliances, light switches, etc., will be interconnected forming diverse new network clouds. We envision that the Internet of the future will interconnect these clouds into the existing IP infrastructure. As people move in and among these clouds, they will encounter and communicate with, a range of devices varying in processing power, mobility, and communications capabilities.

While many of the Internet protocols have proven successful and long-lived in traditional networks, we believe they will be inappropriate for use in communicating among limited-capability devices. The question we are investigating is whether the network layer services provided by IPv4 and IPv6[1] are necessary and sufficient for supporting the heterogeneous devices that will exist in these highly dynamic network environments. We believe there are numerous problems with using IP in these networks. Primarily, for several classes of applications, IP adds an unnecessary and sometimes prohibitive amount of complexity and overhead. Besides the inefficiencies of payload bytes versus header bytes, there is also the issue of the devices' limited processing and communication capacity.

The premise of this paper is to introduce a new network layer protocol for intra-network communication in clouds containing devices with varying processing and communication capabilities. Our proposed *Pseudo-IP (PIP)* protocol is designed to operate among devices in the farthest branches/leaves of an intranet while providing inter-network connectivity with other clouds and with the existing IP-based Internet infrastructure. Our goal of connecting heterogeneous devices to the network can be broken into two sub-goals: protocol flexibility and efficiency. These goals translate into a design requirement that states that PIP's overhead, both in terms of per-packet overhead and protocol complexity, should be proportional to the functionality needed by the application. In other words, protocols should only require complexity proportional to the difficulty of implementing a particular protocol function.

This paper is organized as follows. Section 2 describes different scenarios and applications in which PIP might be used. In Section 3 we provide a protocol description of PIP and in Section 4 we discuss research issues. Section 5 relates PIP to other work. The paper is concluded in Section 6.

0-7695-0253-9/99 $10.00

2. Scenarios, Applications, and Functions

The motivation for using a lightweight network protocol is based on understanding scenarios in which PIP might be used. In this section we refine the concept of a large-scale network environment into a set of scenarios, applications, and network functions. We also offer several detailed scenarios to help motivate the need for PIP. Finally, we identify the types of data delivery that PIP will have to support.

2.1. A Model for Network Environments

The scenarios described in this section are typical of next-generation scenarios which could greatly benefit from simplified connectivity in devices with limited capabilities. The goal is to refine these scenarios into applications and services which can then be refined into basic protocol functions. This hierarchy is shown in Figure 1.

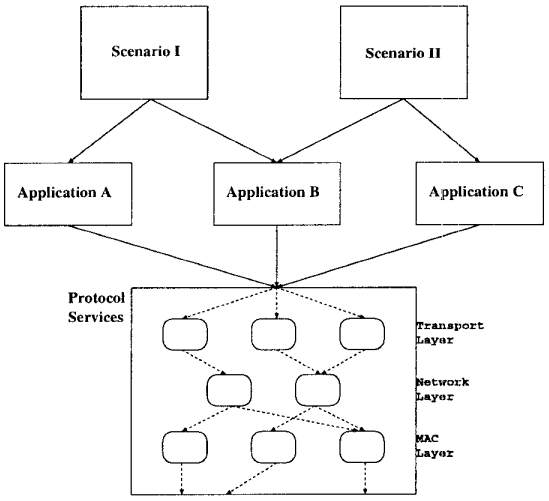

Figure 1. A model to decompose scenarios into applications and protocol functions.

Top-level scenarios can be refined into potentially overlapping sets of applications or functions. Applications and services are then refined into network services which map directly onto particular protocols. For a given application, typical transport layer services might include user authentication, reliable transactions, data aggregation, etc. These transport functions would then require lower layer protocol functions like reliability, error detection, flow control, etc.

2.2. Scenario Classes

In this section, we describe three scenarios in which the availability of a lightweight network protocol could offer significant complexity and performance savings.

Home environment: Homes of the future will be full of devices ranging in capability from very limited functionality to fully functional computers. Further categorizing the list of potential devices results in the following:

- **Sensors**: Uni-directional, broadcast-only devices.
- **Binary-state devices**: These two-state devices will be controlled from a remote station. Typical devices include lights, locks, switches, etc. Devices will broadcast state information as well as receive and process simple commands.
- **Poly-state devices**: These devices have more than two states and potentially some limited embedded processing capability. Examples include kitchen appliances, home entertainment components, etc.
- **Fully capable devices**: These devices have significant communication and processing capabilities. Examples include PCs and PDAs. One function will be to manage other devices in an environment.

The home environment has been the subject of futuristic scenarios describing how everything in the home may be controlled remotely. Tasks such as making coffee, turning on the dish washer, adjusting room temperature could all be controlled via a management station. The major drawbacks are the cost of connecting many devices together, and the complexity of providing communications functions like reliable data transfer and security.

One of the challenges that must be overcome before we are likely to see a fully connected home environment is the cost/complexity tradeoff. Consider the impact of adding communications capability to devices that are very inexpensive and currently have little compute capability. Attempting to implement the full version of a protocol like IP and then build in transport protocol services would add several orders of magnitude in complexity and cost. The cost to produce a simple device like a light switch would go from a few cents to a few dollars.

One of the primary applications required in the home environment is transaction-based state control for simple devices. In fact, this application turns out to be one of the more difficult to provide. It requires both a high level of complexity, and yet, it is expected to work on some of the simplest devices. For example, consider the application of turning on a light switch. This application seems straightforward but is not. The light switch will likely have to support dynamic addressing, reliable data exchange, and some form of security. Reliability is necessary to guarantee proper operation, and security is required if the communications path uses a wireless medium or the wireline path is accessible from the outside world.

Highway spaces: Highway environments offer a different set of requirements than the home environment. Our first highway scenario is based on a geographical area with fixed, mobile, and transient devices. A common example might be a street intersection or a roadway. As drivers approach an area, they may want to learn about gas stations, car repair shops, restaurants, or other businesses in the

area. Businesses will want to announce their existence either through broadcast advertisements or via some directory service. Furthermore, customers and businesses may want to interact including applications like reservations for service, pre-payment for products, etc. The highway scenario is much different than many of the traditional transient, cell-based communication problems in that it is not only about the movement of objects among cells but the provision of communications services in a diverse environment.

Three types of applications that will likely occur in the highway environment are transaction-based applications, broadcast-based information delivery, and data collection from simple sensing devices. These applications are similar to those in the home environment but differ based on the types of devices and distances over which communication will take place. Transactions are again important, and so is simple data collection. A new application is the straightforward, unreliable broadcast of information to all receivers in the network cloud. This is similar to data collection but instead of many-to-one it is one-to-many. Some types of data that might fall under this paradigm include weather information, road conditions, traffic conditions, etc. This type of information, because it is dynamic and repeated frequently, need not use an acknowledgment-based protocol.

Inhospitable environments: One of the most interesting environments in which a protocol like PIP would be well suited is in inhospitable environments. One scenario would use sensors to collect data for disaster relief or search and rescue efforts. Seismic or sound sensors might be scattered throughout a collapsed building to detect motion as a means of locating trapped survivors. These sensors might even be built into the building infrastructure and activated in the event of an emergency. Sensors would have to communicate data to the "edges" of the network because searchers could not get very far into the collapsed building. Furthermore, the collapse of the building would likely destroy some of the sensors and the thick rubble might completely isolate some sensors or strictly limit the communication range. Network communication would have to be flexible enough to handle numerous malfunctioning devices.

A related scenario involves trying to find survivors during other natural catastrophes like floods, fires, volcanos, etc. These events are characterized by an inability of searchers to reach certain areas or to sufficiently search for survivors over an affected area. Sensing devices could be scattered across a large geographical area. Depending on the sensitivity of the devices and their range, several thousand or even tens of thousands of devices could be scattered over a very large geographical area. In the distant, or even not-so-distant future, these devices may even have video capability triggered based on detected motion. A personal computer located somewhere at the periphery of the sensor network would listen for any sensor feedback that might give searchers a better idea where to search for survivors. Because so many of these devices might be used in a single search effort, the need to make them robust yet inexpensive is a critical requirement.

2.3. Data Transmission

Given the types of scenarios in which PIP might be used, it is now useful to describe the types of data delivery services that will have to be supported. We envision the need for PIP to handle three different kinds of data delivery:

Unreliable, sensor transmissions. PIP will have to transfer data between a large number of low-power, low-capability devices and one or more fully-capable devices targeted with data collection and aggregation.

Semi-reliable data delivery. Some scenarios and applications will not necessarily need full reliability all of the time. For these applications a "better-than-best-effort" model may be sufficient. In these cases, Forward Error Correction (FEC)[2] or cyclic transmission[3] may be used.

Reliable data delivery. Reliable data transfer will still be a critical service. Reliability can be implemented at different layers in the stack and includes both hop-to-hop reliability and end-to-end reliability.

Many-to-many transmissions. The need to disseminate *and* aggregate information suggests many different communication paradigms will be needed. In addition to one-to-many, many-to-one, and many-to-many communication, higher functionality like anycasting[4] may be needed to facilitate data aggregation or directory services.

3. Overview of the Pseudo-IP Protocol

The motivation for PIP is the design of a lightweight protocol to support a variety of applications among devices with widely varying capabilities. There are three design goals:

1. Attempt to minimize protocol overhead and complexity at the network layer. In particular, we use IPv4 (and IPv6) as the standard by which to compare PIP.
2. Attempt to provide flexibility in terms of protocol support for both simple and complex network functions. Furthermore, overhead should be proportional to the complexity of the function, i.e. simple functions should require less overhead and more complex functions will necessarily require more overhead.
3. Attempt to provide flexibility and efficiency when interacting with various data link layer and transport layer protocols.

Figure 2 shows the flexibility of PIP. At one extreme, the PIP design allows it to replace the functionality of the data link, network, and transport layers. Alternatives to this

allow PIP to be used underneath almost any existing transport or even network layer protocol. It is even conceivable, though unlikely, to run IP on top of PIP. Finally, PIP may run on top of a data link layer protocol like IEEE 802.11[5].

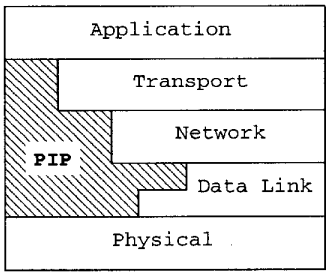

Figure 2. Position of PIP in the protocol stack.

When comparing the performance of IP to PIP, there are two types of overhead to consider. First, there is per-packet overhead in terms of bits required for the header. Second, there is computational overhead in terms of the complexity for implementing required protocol functions. For applications transmitting at high packet rates, the per-packet overhead is typically the more significant disadvantage. For applications running on devices intended to be simple and low cost, the cost of implementing a complex protocol stack is the more significant disadvantage. Because these two tradeoffs depend very heavily on the types of applications and the types of devices, it is difficult to have a single, efficient, fixed-header protocol. For this reason, the PIP header is completely variable.

Our proposal for a PIP packet format is not a true specification, but more of a design. A true specification is beyond the scope of this paper and left to the relevant standards groups. Our proposed header has the following two parts:

Meta-Header. The meta-header is a map of the header fields included in the header. A representation of the meta-header is shown in Figure 3. The meta-header is an array of bits divided into 8-bit words. The first bit of each word is called the *continuation bit*. A "1" in the continuation bit position means there is an additional 8-bit word after the remaining 7 bits. The meta-header can be made arbitrarily long by stringing together as many 8-bit words as needed. The 7 remaining bits contain information about what header fields actually follow the meta-header. If a bit value for a particular header field is "0", the particular field is not included in the packet. This scheme provides significant flexibility. All fields are considered optional and may not be included. The decision can be made on a per-application, per-connection, or per-packet basis. Some header fields which might be included are: source address, destination address, packet type, packet length, protocol type, time-to-live, source digital signature, encryption information, etc.

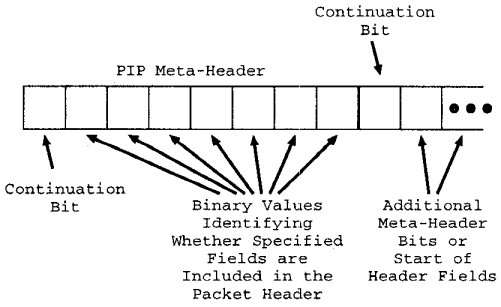

Figure 3. The PIP header format.

Header Fields. The remainder of the header contains the fields described by the meta-header. The ordering of fields is defined by the protocol specification and the length of each field is variable using the same continuation-bit scheme that was used for the meta-header.

We believe the continuation bit scheme offers an excellent trade-off between overhead and flexibility. For example, consider a fixed-length address field. Experience tell us there will eventually be an insufficient number of addresses. But, trying to solve the problem by reserving a very large number of bits for the address field adds too much overhead for networks with only a few devices. Given that variable size fields are a must, the continuation bit scheme offers better flexibility than using other variable length techniques. For example, a protocol might use a meta-field to specify length. But the meta-field is likely to be fixed length and will eventually limit the header's flexibility. The continuation bit scheme has no limits, and furthermore, the actual overhead is proportional to the size of the field. Each field has exactly 12.5% overhead no matter how long the field.

The minimum header size for a PIP packet is 8 bits, i.e. a packet with a continuation bit of "0" in the first meta-header word and no header fields. We envision applications which communicate between devices in which no header information is ever needed, and each packet contains a very small amounts of data. For example, if a fixed-assignment data link layer protocol is used, source and destination addresses may be implicit based on the slot used.

In addition to the simplified header, PIP also offers a major advantage in terms of requiring fewer support protocols. Most Internet devices are required to support additional protocols to facilitate address translation and other control functions. For example, the Internet Control Message Protocol (ICMP) is an important support protocol for IP. Another example is the need to convert MAC layer addresses and network layer addresses through the Address Resolution Protocol (ARP) and the Reverse Address Resolution Protocol (RARP). Given that PIP offers significantly reduced network layer functionality, support protocols will likely be unnecessary.

In the Internet there is often significant overlap in functionality and redundancy of information. The flexibility of the PIP header and PIP functions allow existing protocols, both at lower and higher layers, to be efficiently integrated. For example, PIP will likely function well with a variety of MAC layer protocols. This is an important characteristic given the wide variety of devices and environments proposed in Section 2. For example, if devices implement a wireless protocol like the IEEE 802.11 wireless LAN standard, fewer PIP functions will be required because 802.11 provides its own addressing and checksum mechanisms; uses collision avoidance; and has provisions for acknowledgments[5].

4. Research Issues and Approaches

The goal of PIP is to provide simple yet flexible network layer functionality to be used by higher layer protocols providing services like reliability, congestion control, authentication, etc. Dumb devices should only have to implement the minimum number of functions to achieve their communications requirements. Furthermore, simple devices should not incur overhead penalties for functions they cannot or do not wish to perform. Conceptually, the functionality of PIP can be compared to IP in the same way UDP can be compared to TCP. UDP is a lightweight protocol providing almost no transport layer services while TCP provides many more services and is more complex.

The research challenges we plan to explore result from issues raised by building a general networking infrastructure based on PIP. Part of this challenge includes investigating how to build additional network services, like reliability, on top of PIP. A second challenge is how to interconnect PIP clouds with the existing IP infrastructure. In the following paragraphs, we describe some specific research issues related to PIP. The issues are organized bottom-up starting at the physical layer and moving up the OSI protocol stack.

Power: For battery powered devices, power control is essential. For extremely power limited devices, data transmission will be the dominant power drain. This is in contrast to IP networks where throughput and delay are the primary considerations at routers. Power considerations will affect the choice of MAC and routing style.

In very low power environments, it may be advantageous to use multi-hop routing as opposed to increasing the transmit power to extend link range since power requirements increase at a rate of $O(r^4)$ as the distance, r, increases.

MAC Layer Integration: How PIP handles routing given consideration for power will be a critical issue. Traditionally, MAC layer functions have been separated from the network layer. The resulting inefficiencies are not significant in a high powered, high bandwidth, wired environment. However, in low powered, low bandwidth environments, these inefficiencies become non-negligible. Overall efficiency may be improved if network layer protocols understand services provided by the MAC layer.

Another important consideration is that many of these wireless devices will have to coexist with other wireless devices using unlicensed spectrum in the ISM bands. Instead of using a contention-based protocol like the IEEE 802.11 standard, the MAC layer may incorporate both TDMA and CSMA features. TDMA may provide quality of service functions, while CSMA may provide access for devices with less well defined traffic characteristics. By exposing this structure to the networking layer, better utilization can be made of available transmission resources saving power and improving network throughput and delay.

Mobility: Mobility impacts both the routing and MAC layer considerations. In a static environment, routing is relatively straightforward. In a mobile environment, considerable effort must be expended to maintain network connectivity. Routing protocols for mobile environments need to be able to deal with unstable links and a changing set of neighbors. Device mobility and heterogeneity will affect the provision of routing services.

Addressing: The type of addressing employed can have a significant impact on operation. The issue is ease-of-use and the two options are static or dynamic addresses. Statically assigned addresses can either be globally unique as in the case of Ethernet, or assigned by the user prior to deployment as is done with IP addresses. Globally unique addresses will tend to be long and would consume valuable bandwidth in isolated networks. User assigned addresses could be optimized to the local network, but would increase the administrative complexity.

Dynamic addresses are advantageous in that they can be optimized to support the requirements of a local network. A network with a small number of devices can use very short addresses, thus saving valuable header space and hence bandwidth. The drawback is the added complexity required. In a single-hop, broadcast network, assignment is relatively straightforward, but will be more complex in a multi-hop network. In the case of interconnected PIP networks, gateways will be responsible for address translation.

Data Flow: Data can flow can be broken into the following categories: (1) *One-to-One* in which data is sent from a single source to single destination; (2) *Unidirectional One-to-Any* which allows simple devices to transmit readings but without ability to specify a target and without ability to receive feedback; (3) *Bidirectional Many-to-Many* which is basically the IP multicast model where data is sent unreliability by any group member. The key issue is how information flow can be achieved given the challenges of mobility, link quality, and device heterogeneity.

Reliability. Reliability can be addressed at either the link layer or the transport layer. In a single hop network, reliable delivery can be ensured by using link-level acknowledgments. In a multi-hop network, link-level acknowledgments do not suffice. While they can provide a greater level of assurance, they do not guarantee delivery (e.g. in the case of transit node failure). In the most general case, reliability is a transport level function. In specific cases, MAC layer functions can be leveraged to provide reliability thus saving valuable bandwidth and power.

Security. Many of PIP's applications will require security, but adherence to flexibility goals demands that it not be explicitly required. The issue is whether existing MAC and transport security functions can be leveraged. In addition, environmental characteristics may be useful. For example, physical security may provide line-of-sight limitations. Intelligent gateways like firewalls may also be used.

Gateways. PIP networks are not global; they are clouds containing localized addressing and communication. PIP clouds may be connected via gateways to the global Internet. A PIP device in one cloud is not expected to communicate with a PIP device in another network without the aid of an intermediate gateway handling address and protocol conversion. Gateways can provide complex services for local devices as well as provide interoperability with IP networks.

5. Related Work

In a July 1998 DARPA-sponsored workshop, several research efforts in the area of smart, user-aware working environments, or *smart spaces*, were presented. Our position paper on PIP[6], as well as a white paper on the implications of using large numbers of computers per person[7] were among the invited papers. In addition, since the early 1990s, several areas related to smart spaces have been the subject of research. One example is the Daedalus/BARWAN project (http://daedalus.cs.berkeley.edu/) at UC Berkeley. They have proposed an architecture that supports adaptation to new functionality and services. These new services are discovered as the client moves through the network[8].

Another effort, Piconet[9] is a radio network developed at the Olivetti and Oracle Research Laboratory (ORL). Piconet provides connectivity to a range of portable and embedded devices including sensors, appliances, laptops, and PDAs. Researchers at ORL have also conducted work in the Active Badge location system which uses tags called active badges to locate personnel in environments like offices, hospitals, and campuses[10]. To support their Active Badge location system, ORL built an infra-red (IR) network[11] connecting active badges and IR base sensors (which are also interconnected through a wired backbone). More recently, IR interfaces have been added to other mobile devices like equipment tags and different types of personal computers. Using the active badge location technology, Harter et al.[12] built a distributed service that provides access to location information. This location service has been used as a building block to implement several location-aware applications for the active office or home environments.

6. Conclusions

In this paper, we introduced the PIP protocol whose goal is to interconnect devices with varying power, communication, and processing capabilities. We motivated the need for PIP by describing several application scenarios and their corresponding functional requirements. PIP's approach is to accommodating heterogeneous devices participating in applications of varying complexity by providing as much simplicity and flexibility as possible. PIP allows for a various amount of overhead depending on the application requirements and device capabilities. This paper also lists several research challenges and some potential solutions. Finally, research work related to PIP shows that there is truly a need for this type of protocol.

References

[1] S. Deering and R. Hinden, "Internet protocol, version 6 (IPv6) specification," Tech. Rep. RFC 1883, Internet Engineering Task Force (IETF), December 1995.

[2] J. Nonnenmacher, E. Biersack, and D. Towsley, "Parity-based loss recovery for reliable multicast transmission," in *ACM Sigcomm 97*, (Canne, FRANCE), August 1997.

[3] K. Almeroth and M. Ammar, "Scalable delivery of web pages using cyclic best-effort (UDP) multicast," in *IEEE INFOCOM '98*, (San Francisco, California, USA), July 1997.

[4] C. Partridge, T. Mendez, and W. Milliken, "Host anycasting service," Tech. Rep. RFC 1546, Internet Engineering Task Force (IETF), November 1993.

[5] 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.

[6] K. Almeroth, K. Obraczka, and D. DeLucia, "Pseudo-IP: Providing a thin network layer protocol for semi-intelligent wireless devices." USC Computer Science Technical Report 98-680, August 1998.

[7] J. Heidemann, R. Govindan, and D. Estrin, "Configuration challenges for smart spaces." USC Computer Science Technical Report 98-677, July 1998.

[8] T. D. Hodes and R. H. Katz, "Composable ad-hoc location-based services for heterogeneous mobile clients," *Wireless Networks Journal (to appear)*, November 1997.

[9] F. Benner, D. Clarke, J. Evans, A. Hopper, A. Jones, and D. Leask, "Piconet: Embedded mobile networking," *IEEE Personal Communications*, vol. 4, no. 5, pp. 8–15, October 1997.

[10] R. Want, A. Hopper, V. Falcao, and J. Gibbons, "The active badge location system," *ACM Transactions on Information Systems*, vol. 10, no. 1, pp. 91–102, January 1992.

[11] A. Harter and F. Bennett, "Low bandwidth infra-red networks and protocols for mobile communicating devices." Olivetti Research Laboratory Technical Report.

[12] A. Harter and A. Hopper, "A distributed location system for the active office," *IEEE Network*, vol. 8, no. 1, January 1994.