Mobility and Address Freedom in AllNet

Edoardo Biagioni
University of Hawai'i at Mānoa
esb@hawaii.edu

ABSTRACT

Mobile devices can be addressed through a variety of means. We propose that each device select its own addresses, we motivate this choice, and we describe mechanisms for delivering data using these addresses.

Hierarchical point-of-attachment addresses are not effective with mobile devices. To support mobility, the network has to maintain a global mapping between addresses and locations. Given that such a mapping must be maintained anyway, there is no reason for structure in mobile addresses.

Mobile addresses may therefore be unstructured and even self-selected. The advantages of self-selected addresses include the ability of devices to join and form a network without any need for prior agreement or central coordination, and the ability to choose a personal, memorable address.

Self-selection does require a mechanism to resolve ambiguity when the same address is selected by multiple devices. Of the many possible ways to provide such disambiguation, we investigate disambiguation that relies on the authentication and encryption used for communications security.

The principles presented here are general, and complemented by examples of unstructured and self-selected addresses used within the AllNet network protocol.

Keywords

Network Architecture, Infrastructureless Communication, Interpersonal Communication, Routing, Ad-Hoc Network, Delay-Tolerant Network, DTN, Networking Protocol

1. INTRODUCTION

Ubiquitous communication cannot be provided only by expanding the fixed infrastructure. For the foreseeable future, there will always be areas that are not covered and people who cannot afford to pay for service.

Instead, mobile devices can themselves forward messages when there is no infrastructure to do so. In particular, end devices may act as traditional routers to forward messages destined for other devices, so messages can be delivered well beyond the range of the radio of the original sender, a technique known as Ad-Hoc networking. Similarly, devices may store messages destined for other devices, and deliver them later when the destination device is in range, a technique most commonly referred to as Delay-Tolerant Networking or DTN.

1.1 Motivation and Background

Even forwarding just a few messages per second would be very useful for text messages, and can make a substantial difference in both emergencies and daily life situations. There are many systems that have started to explore and implement this functionality [18, 10, 11]. In AllNet, messages may be forwarded both within ad-hoc networks and DTNs, and across the Internet by any devices with an Internet connection [4, 2].

1.1.1 Mobility, Forwarding, and Address Structure

Since both ad-hoc networks and DTNs are designed to support dynamic changes in network connectivity, addresses within such networks do not need to reflect network structure. Instead, the networking protocol must include a discovery phase that allows any address to be located[15, 6]. This discovery phase necessarily relies on a broadcast, usually limited to devices in a relatively small network, which is similar to both ARP (MAC addresses also have no useful structure) and to routing within IP networks.

Broadcasting does not scale well to larger networks[1]. Instead, such networks usually implement a database mapping addresses to locations.

For example, IP requires that the network part of the address reflect the point of attachment to the network. In Mobile IP [12] [13] [14] a *"mobile node is also associated with a care-of address, which provides information about its current point of attachment to the Internet."* [13]. IP datagrams are routed to the destination network, where the home agent forwards the datagrams to the destination through an IP tunnel.

As another example, devices using mobile telephony and mobile data protocols (i.e., mobile public networks) have to connect to a local broadcaster and announce

[1]Except that at least one communication network, inspired by the BitCoin distributed ledger, broadcasts messages to all nodes in the BitMessage network [20], using proof-of-work to limit the number of messages sent.

their presence at their current location. Similar to Mobile IP, this announcement updates the Home Location Register at the cellphone's home location. This Home Location Register is then used to route incoming calls [9]. As an optimization, the mobile device's location is usually cached in the Mobile Switching Center's Visitor Location Register of its current location.

The collection of home agents across the Internet in mobile IP, and the collection of home location registers and visitor location registers in mobile public networks, each amount to a distributed database mapping a structured address to a location that is unrelated to the address.

Such distributed databases are centrally managed in the sense that the database is implemented by a small number of trusted participants (mobile public networks) or by participants trusted according to a hierarchical trust delegation model (mobile IP). In these networks, the IP address or mobile telephone number identifies which servers contain the desired translation. However, this is a matter of management rather than technology. If both the network and addresses lack a hierarchical structure, alternative technologies, such as Distributed Hash Tables [16, 19, 17, 21], can provide distributed mappings

To summarize, the structure in addresses that is designed for fixed, hierarchical wired networks is not helpful in forwarding messages to mobile devices.

1.1.2 Forwarding, Security, and Addresses

Intermediate devices forwarding user messages in an ad-hoc network or DTN typically do not all belong to the same individual or organization, so messages should be encrypted and authenticated. Also, most device owners would prefer to limit resource consumption. If the sender is cryptographically known to the intermediate device, messages can be forwarded eagerly. If not, intermediate devices may limit the rate at which messages are forwarded.

For both communication security and to preferentially forward messages from known devices, each device must securely keep the public keys of the contacts of the owner of the device, and its own corresponding secret keys. Given a message to forward, a device may then verify whether it was signed by one of the public keys for a known contact, and forward it preferentially only if it is from one of those devices. Signature verification is relatively quick (less than $200\mu s$ on a 3GHz CPU, see also [3]), so a mobile device can efficiently discard messages not sent by a known contact.

Signature verification may be combined with encryption to determine whether a message is for this device. If a message is from a known contact, and can be decrypted using the corresponding secret key, it is for this device. This kind of cryptographic verification can provide disambiguation when different devices have the same address. Even assuming processors much slower than 3GHz, the cost should be less than $1ms$ per message. In other words, secure ad-hoc networks and DTNs work well even if all devices have the same address, or equivalently, even without using addresses.

1.2 Summary of Contributions

The remainder of this paper presents the benefits of using self-selected addresses or IDs, in cryptographically secure and authenticated networks where addresses are not required for correct delivery, and where mobility means that there is no benefit to having addresses reflect network structure.

Specific contributions of this paper include:

1. A discussion of the principles of the design of AllNet self-selected addresses and their applicability beyond AllNet (Section 2). AllNet is a secure network that, unlike BitMessage[20], supports devices that are not continuously connected to the Internet and does not require broadcasting to all, or even a large subset, of all participating Internet-connected devices.

    - AllNet devices that are connected to the Internet, spontaneously organize into a Distributed Hash Table (DHT) that, analogous to a collection of redundant email servers, allows devices with intermittent Internet connectivity to store and retrieve messages while connecting to relatively few DHT nodes. The DHT provides a distributed store for messages sent to and from arbitrary self-selected addresses.
    - Such addresses need not be unique. Messages sent to an address shared by different devices are generally delivered to all these devices. Devices then use the cryptographic principles described in Section 1.1.2 to select the ones truly destined to themselves. This works even if, as in IPv6 [7], each device selects multiple addresses.

2. Self-selected addresses correct many of the issues and disadvantages of structured addresses. In particular, when stability is desired, such self-selected addresses may be permanently used by a single individual or organization, even when all the devices are mobile. When anonymity is preferred over stability, self-selected addresses may be changed at will. This pattern is also supported by IPv6. Section 2.3 discusses these and other benefits of self-selected addresses.

3. The AllNet addresses just described are arbitrary bitstrings that carry no meaning for normal human users, and in this way are analogous to IP

addresses. AllNet also supports a different kind of address, a human-readable string designed to be memorable, self-selected, yet provide a degree of security. These Allnet Human-Readable Addresses (AHRAs) are more analogous to Domain Names or email addresses, and are described in detail in Section 3.

2. ALLNET ADDRESS DESIGN

2.1 AllNet Overview

AllNet is a protocol designed primarily to support secure interpersonal communication. The novelty of AllNet includes the fundamental support for security combined with the equally fundamental support for mobility, all without relying on centralized servers.

AllNet relies on two main underlying means of communication: the Internet, and direct ad-hoc and delay-tolerant networking between mobile devices. Each of these is used as available.

Key Exchange: Authenticated key exchange is provided either through communication among devices, or by the exchange of Allnet Human Readable Addresses (AHRAs). The first mechanism requires users to enter a short string, to prevent spoofing by other devices. The second mechanism leverages any authenticated side channel to exchange AHRAs. This side channel need not be encrypted: examples include a telephone conversation or a business card. In both cases, the recipient of a public key has interpersonal assurance that the public key received belongs to the other party in the communication.

Priority: Because a device may receive more messages than it is willing or able to forward and cache, each device has an automatic prioritization mechanism based on information available in each message. The prioritization may be changed by each device as appropriate, but by default gives priority to messages recognizably to or from existing contacts, and to a lesser extent, to messages that will consume the least device and network resources. At any give time, the highest priority messages are forwarded and cached.

Resource Management: One of the goals of AllNet is to forward messages for unknown people using at most a small fraction of available resources, typically 1%. Subject to the priority and resource limitations, communication using the Internet and using localized broadcasts should deliver messages whenever allowed by the resource constraints.

Distributed Hash Table: When forwarding data on the Internet, messages are sent to DHT nodes corresponding to the destination address. Again, only the highest priority messages are forwarded, but this may include all packets if traffic is sufficiently low.

Acknowledgements: AllNet messages may be acknowledged. To request an acknowledgement, the sender includes in each message a 128-bit random string called the message ACK. As the message is encrypted, the message ACK is encrypted along with the message itself. The sender also includes in the message the first 128 bits of the unencrypted hash of the message ACK. This hash is called the message ID. Only the intended recipient will be able to decrypt the message and recover the message ACK. When the recipient returns this ACK to the sender, every other node can hash the ACK and, if it corresponds to the message ID of a previously seen message, delete the corresponding message from its cache.

Social Network: AllNet is used primarily for interpersonal communication, so AllNet maintains a list of contacts and the public key for each contact. That is, each user's social network is kept only on the user's device, with enough information to allow secure and authenticated communication with each person in the social network.

Additional details about AllNet are available from prior papers [4, 2, 1].

2.2 AllNet Addresses

Section 1.2 explained some of the ways in which AllNet addresses resemble and differ from addresses used in other networks. Like some other networks, AllNet has two kinds of addresses: addresses used for data routing and delivery, and AllNet Human Readable Addresses, or AHRAs. The latter, discussed in Section 3, provide some of the benefits of email addresses or domain names, without some of the drawbacks of each. This section describes the AllNet addresses used in data delivery.

We begin by describing a simplified version of AllNet that uses no addresses whatsoever. In this simplified network, every participating node receives every message. Since most AllNet messages are signed and encrypted[2], every receiver of a message checks to see if the message is signed by a sender known from the social network. In the absence of addresses, this can be done only by trying to verify the signature with the public key of every contact in the social network. If one of these verifications succeeds, then the receiver attempts to decrypt the message. If the decryption succeeds, then with overwhelming likelihood, the message was encrypted with this receiver's public key.

Since message addresses can be used by an attacker for Traffic Analysis, sending and receiving without addresses is supported by AllNet. Each packet carries both a 64-bit sending and receiving address, and also the number of bits of each address that are meaningful. If the number of bits (of either source or destination ad-

[2]AllNet also supports broadcast messages, which are signed but not encrypted.

dress, or both) is zero or a small number, that address cannot be used for Traffic Analysis.

Messages with no addresses can be useful on relatively small networks or where security is sufficiently important that traffic analysis should be thwarted and the overhead of verifying a packet against all the signatures in the social network is acceptable. However, this is not practical on most mobile devices, so most AllNet messages carry a number of address bits > 0, currently typically 16 bits.

AllNet devices store the public key for each contact together with a local and remote address. The number of bits in the address is decided independently by each peer. When peers exchange messages, they may specify any number of bits. A source or destination address matches the remote or local address for a peer when the number of matching bits is at least the number of bits specified in the packet, or at least the number specified in the original exchange. For example, if Alice gives to Bob an 8-bit address, and Bob gives to Alice a 16-bit address, then Bob may attempt to verify and decrypt any packet with s source bits and d destination bits as long as the first $\min(s, 8)$ bits of the source address match Alice's address, and the first $\min(d, 16)$ bits of the destination address match the address that Bob gave to Alice.

This flexibility means each communication can be configured to trade off resource consumption with security. When more address bits are specified, each device needs to verify and decrypt fewer packets, which can be important for energy-constrained devices. With fewer bits, more verifications and decryptions must be carried out to find out which the packets are meaningful to a particular user, and more DHT nodes may need to be contacted to find the packet in the first place.

In practice, the current version of AllNet (3.2.1) always uses 16-bit addresses. Even with a relatively large network, 16-bit source and destination addresses mean that typically 2^{-16} of the public keys in a user's social network will be tried for any received packet.

Because the bit pattern in the address is not meaningful to AllNet, each device can select any suitable address, without coordinating with any central authority or even a group consensus. Instead, each device can get online without prior authorization, and may even (and in the current implementation, usually does) choose different addresses for communication with different peers.

The ability to self-select an address fits well with the AllNet emphasis on distributed communication where each device is an equal peer, and with enabling communication whenever each device can possibly communicate, without waiting for registration or approval.

Whenever addresses are self-selected, there is the possibility that two systems will select the same address. However, as illustrated by the simplified example of AllNet using no addresses, duplicate addresses only require additional verification or decryption attempts, and do not lead to incorrect behavior.

2.3 Advantages of AllNet Addressing

AllNet addresses have no relation to the device's point of attachment. Addresses need not change when the mobile device actually moves or for any other external reason, and may be retained forever. Conversely, addresses may be changed at will. The only party that needs to be informed when an address changes is the personal contact in the communication that uses the address. In this way, AllNet addresses function as self-selected ID only meaningful within one's social network.

Other advantages of AllNet addresses were mentioned in Section 2.2. These include the option of masking most or all of the bits of the address to foil traffic analysis, and the competing benefit of using addresses help filter out packets for which verification and decryption need not be attempted.

Self-selected addresses can also be used while still defeating traffic analysis. Any two devices that can agree on a sequence of addresses may use a different address for each message. As long as the set of expected addresses is limited, this retains many of the benefits of AllNet addresses, while making it hard for an attacker to use traffic analysis.

When AllNet is used for wireless ad-hoc and delay-tolerant message transmission, addresses are also useful to prioritize specific outgoing messages. Since wireless spectrum is sometimes a valuable commodity, and in any case using onboard radios consumes energy, it is beneficial to prioritize the transmission of packets that are of use to the recipient. In the handshake at the beginning of an AllNet ad-hoc transmission, receivers may indicate addresses they are particularly interested in receiving, and senders may prioritize such messages.

3. HUMAN-READABLE ADDRESSES

Section 1.2 introduced the notion of AllNet Human-Readable Addresses, or AHRAs. Unlike regular AllNet addresses, AHRAs are designed specifically to exchange among humans, in a way analogous to current email addresses. This section describes AHRAs in detail.

The purpose of an AHRA is to identify a public key in a manner analogous to a certificate, but in a way that humans can exchange painlessly and without relying on certificate authorities.

The basic format of an AHRA resembles an email address:

`personal_name@word_pair.word_pair`

The personal name (PN) is a public address for this individual. To minimize the effect of mis-spellings and confusing fonts, Allnet maps each letter in the personal name to 4 bits, and multiple, sometimes indistinguish-

able letters are mapped to the same pattern. For example, the letters "1", "l", "L", "i", and "I" are all mapped to the same bit pattern `0001`.

An address including only the PN (`name@`) is valid, but may not be unique.

The word pairs (WPs) are a way of encoding in a memorable way a hash of the public key. An AHRA may have zero or more WPs.

Each word pair encodes 14 bits of the hash, with successive pairs encoding successive bits. The 14 bits are encoded as two seven-bit parts, each taken from a dictionary of 128 common words. For example, in English, the first few words in the dictionary used for the first part of a word pair includes the words "the", "be", "of", "to", "a", and so on. The dictionary for the second word in the pairs includes the words "time", "people", "year", "well", "work", and so on. Word pairs might then be "of-time", "to-work", and so on.[3]

When a user wishes to create an AllNet address, the user puts his or her device to work creating keys. Keys must have certain properties, so in general a suitable key will only be found after generating a large number of public/secret key pairs. The generation of AHRAs is the only mechanism in AllNet resembling proof-of-work.

All the keys that satisfy the properties needed of an AHRA are saved, and the user then chooses one for actual use. For example, AllNet has a time server that sends a time message once an hour. The AHRA for this time server is `allnet-hourly-time-server@if-wish.think-past.get-future`. This was one of many AHRAs generated for the PN "allnet hourly time server", and was selected by the author who hopes these word pairs are memorable in the context of this PN. A similar server that sends messages once a day has the AHRA `daily_time@go_car.for_computer.do_future`.

An AHRA with n word pairs can always be used with only the first $m < n$ word pairs. This lessens security, but makes the AHRA easier to communicate and remember. `daily_time@go_car` identifies, less securely, the same key as
`daily_time@go_car.for_computer.do_future`.

What makes a public key valid for use in an AHRA is that the cyphertext from encrypting the PN with the public key must contain a minimum of n 16-bit strings taken sequentially from the hash of the PN.

For example, consider a PN that hashes to a value ending[4] with (hex) `5518 22B5 5D7C` . Then, if $n = 3$, the only acceptable public keys for this PN are those where the PN, when encrypted with the key, contains all of the 16-bit strings `5518`, `22B5`, and `5D7C`. If $n = 2$, only the last two are required.

[3]Just like "1" and "l", hyphen, underscore and space are used interchangeably.

[4]The beginning of the hash is used as the local address, so the bit strings are taken from the end of the hash.

Assuming that the PN, when encrypted using the public key, is no longer than $2^{14} = 16,384$ bits, then only 14 bits are needed to encode the position of each of the bit strings found. Each word pair encodes one of these positions. With this scheme,

- verifying that a public key matches a given AHRA only requires an encryption and a few lookups,
- generating a key given a PN is slow,
- but still much faster than finding a key to match a complete AHRA.

To see that the last is true, consider that a PN encrypted using a valid public key with n word pairs must have n matches of the last b-bit (in AllNet, $b = 16$) strings of the hash of the same PN. Given that the encrypted PN has l bits (and assuming that the hash length $> n \times b$), approximately $(2^b/l)^n = 2^{bn}/l^n$ keys must be examined to find one with n word pairs. In contrast, an attacker must examine approximately 2^{bn} keys to generate a new AHRA for a given PN and n word pairs. The ratio of the work needed to steal a key to the work needed to generate a key is then l^n.

For example, with a $w = 3$ word-pair AHRA using an $l = 4,096$-bit encrypted PN, an attacker must do at least l^w or $68,719,476,736 > 10^{11}$ more work than the legitimate owner.

While this may be feasible for a determined and powerful attacker, it is a deterrent to most attackers. Note also that breaking an AHRA does **not** compromise keys already exchanged, only future key exchanges, and that duplicate public keys matching given AHRAs could be detected "in the wild" on the AllNet network.

This analysis has assumed that (a) public key encryption is independent of hashing, so that the chances of finding the bit strings of the hash in the ciphertext is random, (b) exhaustive search is the fastest way for an attacker to find a public key matching a given AHRA, and (c) the effort to compromise one AHRA cannot be re-used in compromising unrelated AHRAs.

4. SUMMARY AND CONCLUSION

We have shown that in mobile systems the address space is essentially flat. Effective addressing in such systems must either rely on broadcasting, or on a global database mapping addresses to locations.

The design of AllNet takes advantage of the flat address space to dispense with the need for addresses to be assigned hierarchically. Since AllNet already provides encryption and authentication, these are leveraged to support non-unique addresses. Addresses can then be self-selected, and used both as an optimization to avoid having to verify and decrypt every message, and as a key for locating messages in a distributed hash table.

The AllNet Human-Readable Addresses combine personal names with sets of word pairs identifying public keys, and can be used as the equivalent of email addresses or URLs.

Everything described here has been implemented in AllNet version 3.2.1 [1], with the exception of the the message selectivity for the Distributed Hash Table. The implementation runs on Linux, Windows, OSX, and (with limited functionality) iOS. An Android implementation is planned. The wireless ad-hoc portion of AllNet is currently functional between Linux systems, using 802.11 ad-hoc mode.

4.1 Conclusion

Hierarchically assigned point-of-attachment globally unique addresses have been and continue to be useful, and indeed the current Internet would be unimaginable without them. However, these addresses simply don't work well for mobile nodes. Either the mobile node must change address every time it moves, or the address cannot identify the point of attachment.

This limitation, requiring both updating of the mobile database and message redirection whenever the mobile node moves, is also liberating, providing opportunities to self-select addresses rather than requiring addresses to record the person's or device's position in a hierarchy.

The indirection required to support mobility can be an opportunity for new designs. We have presented some of the design choices for AllNet, but many other choices are possible, and we hope this paper inspires others to take advantage of this new freedom.

5. REFERENCES

[1] `http://alnt.org/`

[2] E. Biagioni, "Ubiquitous Interpersonal Communication over Ad-Hoc Networks and the Internet", 47th Annual Hawaii Convention, January 2014, Waikoloa, Hawaii, USA.

[3] E. Biagioni, Y. Dong, W. Peterson, K. Sugihara, "A Protocol for Secure Electronic Remote Voting". IFIP International Conference on Network and Service Security (N2S), Paris, France, June 2009.

[4] Edoardo Biagioni, "A Ubiquitous, Infrastructure-Free Network for Interpersonal Communication", 4th Int'l Conf. on Ubiquitous and Future Networks (ICUFN), July 2012.

[5] Ian Clarke, "A Distributed Decentralized Information Storage and Retrieval System", Div. of Informatics, Univ. of Edinburgh 1999, `https://freenetproject.org/papers/ddisrs.pdf`

[6] Clausen and Jacquet, Editors, "Optimized Link State Routing Protocol (OLSR)", Experimental RFC 3626, Oct. 2003.

[7] S. Deering and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, Dec 1998.

[8] R. Dingledine, N. Mathewson, P. Syverson, "Tor: The Second-Generation Onion Router", Usenix Security 2004.

[9] "Digital cellular telecommunications system (Phase 2+); Universal Mobile Telecommunications System (UMTS); LTE; Location management procedures" 3GPP TS 23.012 version 11.2.0 Release 11, Jan. 2013.

[10] "goTenna Uses Smart Protocols + Radio Waves To Send Messages Off-Grid", online at `https://gotenna.com`. Unlike AllNet, goTenna requires a separate hardware device, with the benefit of increased range.

[11] "Open Whisper Systems", `https://whispersystems.org`.

[12] C. Perkins, ed., "IP Mobility Support", RFC 2002, Oct. 1996.

[13] C. Perkins, "IP Mobility Support for IPv4, Revised", RFC 5944, Nov. 2010.

[14] C. Perkins, ed., D. Johnson, J. Arkko, "Mobility Support in IPv6", RFC 6275, Jul. 2011.

[15] C. Perkins and E. Belding-Royer and S. Das, "Ad hoc On-Demand Distance Vector (AODV) Routing", RFC 3561, July 2003.

[16] Ratnasamy, Francis, Handley, Karp and Shenker, "A scalable content-addressable network", ACM SigCOMM 2001.

[17] Rowstron and Druschel, "Pastry: Scalable, decentralized object location and routing for large-scale peer-to-peer systems", 18th Int'l Conf. on Distributed Systems Platforms, 2001.

[18] Tom Simonite, "The Latest Chat App for iPhone Needs No Internet Connection" Technology Review, March 28, 2014.

[19] Stoica, Morris, Karger, Kaashoek, and Balakrishnan, "Chord: A scalable peer-to-peer lookup service for Internet applications", ACM SigCOMM 2001.

[20] Jonathan Warren, "Bitmessage: A Peer-to-Peer Message Authentication and Delivery System" Nov. 2012, `https://bitmessage.org/bitmessage.pdf`

[21] Zhao, Huang, Stribling, Rhea, Joseph, Kubiatowicz, "Tapestry: A Resilient Global-Scale Overlay for Service Deployment", IEEE JSAC, vol. 22, 2004.

[22] Philip Zimmermann, "Why I Wrote PGP", in the Original 1991 PGP User's Guide (updated in 1999).