NeTS: Small:

Allnet

Networking Technology for Ubiquitous Communication

1. Objectives for the Period of the Proposed Work

At present, if I want to send an SMS to a person in the same room, the message has to travel to my service provider, then to my peer's service provider, and from there to my peer's handset. While this admirable service works remarkably well, it is not always available.

The hardware in most mobile devices can be configured to perform the communication directly. Current technology provides WiFi communications in ad-hoc mode and through WiFi Direct. In the future, opportunistic communications may provide direct connections with higher data rates, longer range, or both. Even when my handset is not directly within range of my peer's handset, the message might be relayed from one mobile device to another, utilizing principles of mesh networks, ad-hoc networks, delay-tolerant networks, and peer-to-peer networks.

We propose to enhance the current infrastructure-based communication among individuals using mobile devices with a secure peer-to-peer network that may be able to provide support when the infrastructure is not available.

When most end users use an Internet connection, most of the time, the connection is not carrying packets. We propose to take advantage of this unused bandwidth for low bit-rate, delay insensitive, low priority traffic that could prove useful for both emergency and routine communications.

Many researchers have observed that wireless devices are now commonplace and could be used to create ad-hoc networks in emergency situations. However, in general, such research focuses on special-purpose devices or software that must be deployed specifically for the emergency. It seems more natural to employ the wireless devices that people carry routinely. We propose to train people to use this facility by giving them a convenient texting facility that they may use at all times. Together with a low-bandwidth Internet access facility, users should be motivated to learn to use AllNet and continue to use it even when the infrastructure is not available.

From a technical perspective, we propose to add an infrastructure-less mode to the wireless portion of the Internet. This new mode can be used for two main purposes: to support local and delay-tolerant communication among nodes within an ad-hoc network, and to provide low-bandwidth Internet access by all nodes when at least some of the nodes in the ad-hoc network do have Internet access. Because of the very limited resources available to most ad-hoc wireless networks, and especially to such a network relying on the goodwill of others to carry messages, AllNet prioritizes messages and usually only delivers the highest priority messages. This priority may vary from device to device, and is highest for packets originated by the device itself. For packets originating elsewhere, the priority is highest for those packets likely to consume the least resources for both the device, and the network as a whole.

Allnet is designed to be particularly useful when other forms of networking fail or are not available, and also sufficiently useful when the Internet is available that it people will employ AllNet routinely.

As in most forms of ad-hoc networking, with AllNet, individual communication devices (cellphones, smartphones, tablets, laptop computers, etc) can automatically form a network whenever they are in range of each other. If people are accustomed to using AllNet for daily communications, they will not hesitate to use it under emergency conditions. The incentive to use the system on a daily basis is provided by AllNet's provision of interpersonal communication and low-bandwidth free Internet access at any time when the alternative Internet infrastructure fails provide these. The peer-to-peer nature of the network and the need to prioritize communications make it natural to teach the devices the social network of the devices' owners. This is enhanced by the security and authentication mechanisms of AllNet.

While the goal is ambitious, the technical requirements are feasible. Ad-Hoc networks have been studied thoroughly, and encryption technology is used daily in securely accessing web pages. AllNet goes beyond these in providing convenient and reliable mechanisms for secure key exchange, and using these keys for confidential and authenticated communication among people who know each other.

Further, AllNet plans to remap existing Internet applications to AllNet, wherever possible supporting infrastructure-less designs rather than server-centric designs. This is simple for chat applications when routing or limited broadcast is effective, and still feasible for web access and, when the bandwidth is available, for multimedia exchanges.

Beyond these technical challenges, creating AllNet poses novel design requirements. We intend to motivate ordinary people to use AllNet for a variety of purposes, and to motivate manufacturers to include AllNet in their systems. We must also motivate ordinary people to agree to share their connectivity when it costs them in battery life, and could cost them small amounts of money if their Internet service charges proportionately to data transferred. We intend to motivate users by providing a variety of incentives, from selfish to altruistic, and by carefully designing both the protocol and the implementation to minimize security threats and annoyances to anyone using AllNet.

To restate, there are two classes of challenges in designing AllNet.

1. Purely technical challenges. Given that participants are willing to provide some of their bandwidth for AllNet, we need to design the protocols and algorithms, and to verify their performance both in simulation and in practice. Where possible we must design protocols for chat and for accessing the web and exchanging email when there is no direct Internet connection, when necessary in a Delay Tolerant mode. The software and protocols must manage resources to minimize impact on the network and on the platforms, provide fully distributed security, and provide power cycling of the radio equipment to minimize power consumption while providing acceptable performance. We need to provide at least initial application software to use AllNet so others will be able to experiment with and adopt the concept.
2. Human challenges. Currently, TCP congestion control is a truly remarkable example of distributed cooperation to achieve a shared goal. AllNet needs to support generous contributions while at the same time supporting users whose motives may be more selfish. We therefore have started to design a variety of incentives for people to cooperate in carrying others' traffic. We also plan to make it easy for untrained people to make good decisions about managing the resources on the devices they own, as well as about resources, such as spectrum, potentially shared for many uses.

2. Expected Significance of the proposed work.

When I tell people that their mobile devices could talk directly to each other, without going through cellphone towers, they invariably get excited at the thought. This is easy to explain. While the infrastructure we have is extremely impressive and generally reliable where it is installed, it is almost always associated with a gatekeeper that seeks to recoup its investment by charging users. As a result, for many users, our impressive technology comes with mixed feelings of having to pay, or even more negative feelings in case of inability or unwillingness to pay.

Furthermore, the infrastructure is not available everywhere, and it might be impractical and very expensive to attempt to extend it to many of the areas where it is not available.

AllNet is not a replacement for the infrastructure. What AllNet provides is a means for very low-bandwidth communication whenever technically possible, specifically whenever other mobile devices are in range and could deliver the message. That even such limited communication can be useful is evidenced by the success of services such as Twitter and SMS, which although limited to 140 characters per message, are used to alert people and convey information including facts and emotions.

I have begun the implementation of AllNet, and intend to continue, and to try and involve students, whether or not the project is funded by NSF. However, NSF funding would help produce a better technology by adding the efforts of graduate and undergraduate students, would lead to quicker implementations on popular platforms, and overall, would make it more likely for AllNet to succeed.

I see the involvement of students as essential both for the project itself, and to train students in thinking about and creating novel technologies that can change the world. By contributing to creating a system that will change people's lives, I believe students become highly motivated to learn more and become researchers themselves. Indeed, I have seen the enthusiasm in the students I have collaborated with so far, whose involvement has been limited due to my lack of funds. One undergraduate student did exploratory work on AllNet as part of a for-credit project and another has expressed strong interest in collaborating once I return from my current travels. A graduate student participated in writing last year's version of this proposal, highly motivated by the idea of a distributed system where everyone contributes to creating a better world.

On an earlier project funded by DARPA, undergraduate students were hired who later went on to earn M.S. degrees. While we did not keep careful track of all of these students, the PI met one of them at a seminar, and found out he is now on the faculty at a local community college. As well as their contributions to the project, these students are contributing to our society both by the knowledge they have acquired, and by their willingness to contribute this knowledge to the next generation and to apply it in technical and other contexts.

As well as training students, I expect the AllNet system to train users and the population in general. AllNet is designed so that collaborative behavior is rewarded, and non-collaborative behavior is not. Although this design may need more refinement, it is a goal of this project to make sure that users understand that their participation is a requirement for the success of AllNet. In other words, AllNet is designed to help build a community, and the community only exists when the members collaborate and actively participate.

3. Relation to Longer-Term Goals of the PI's Project

I have been involved with wireless ad-hoc networks since the year 2000, when DARPA funded the Pods project, in which, with two other PI's (one of them a botanist) we designed and built wireless sensor networks to monitor endangered plants in Hawaii. We were among the first to realize that the

microenvironment surrounding plants is significant and had not been carefully studied. Since then, many studies have focused sensor networks on plant environments.

Since the Pods project, I have continued to work in the field of wireless networks, but also undertaken research in the field of secure communications. Colleagues and I have a protocol to support anonymous but voter-verifiable voting across the Internet.

Besides its potential to change the world, this project is very inspiring to me in part because it combines my interest and substantial experience in wireless networks with my more recent expertise in secure communications. AllNet is designed to provide seamless and easy-to-use secure communications to the average user of a mobile device, a capability that I feel is otherwise being developed very slowly, and often hidden from the awareness of the users themselves.

4. Relation to the Present State of Knowledge in The Field

There are several classes of related work that influenced and inspired our design so far. The first three are related to the technical challenges, and the latter to the more human challenges of this design.

4.1 Ad-Hoc and Delay-Tolerant Networking

There has been much research on Ad-Hoc networking. As a co-editor for the series on Ad-Hoc and Sensor networks of IEEE Communications magazine, and as past recipient of a grant to research wireless sensor networks, I have been exposed to much of this research.

Delay-Tolerant Networking (DTN) is ad-hoc networking with the added challenge that individual nodes or parts of the network may be disconnected from other parts of the network for periods of time. Common algorithms include Epidemic [Vahdat and Becker, 2000], and Spray-and-Wait [Spyropoulos, Psounis, and Raghavendra, 2008]

While ad-hoc networking is fairly well established and effective, and DTN algorithms exist and can be effective in many cases, their use so far has been more limited than the technology would afford. It is my understanding that this is because most equipment with WiFi technology defaults to only using infrastructure mode, and particularly because the human factors required for a wider deployment are missing. Experts can configure ad-hoc mode, but experts are few and far between.

While existing algorithms themselves are distributed, most current research and deployments assume that a single entity or a group of cooperating entities will control the nodes of the network. We would like AllNet to support untrained users, each with independent goals, but willing to cooperate to a greater or lesser degree.

4.2 Peer-to-peer Networks

In contrast to ad-hoc networking, promoters of peer-to-peer (P2P) networks have done a wonderful job of getting people interested in the technology, often in the face of significant and powerful interests attempting to throttle back, control, or suppress the deployment of P2P. There are many reasons for this, including a natural desire for freedom from external controls, the economic desire to save money by copying content, and the ease of deployment of P2P technology.

Because of their distributed nature, P2P designs have also struggled with issues of fairness. In particular, nodes that allow others access to useful material should be rewarded in some way by giving them better access to other material. Anonymity has often been a feature of these networks, and providing incentives while maintaining anonymity can be challenging.

4.3 Anonymizing Networks

In the popular distributed P2P networks, participants are identified by the IP address they use. While an IP address may or may not identify an individual, in many cases, it is preferable not to reveal even such limited information to any potential eavesdroppers. Systems such as Freenet [Clarke, 1999] and TOR [Dingledine, Mathewson and Syverson, 2004] have been designed to hide as much information as possible. TOR does this by encrypting all communications, and only revealing the next hop to the intermediate node that needs to know the next hop. In Freenet, content is transmitted without any indication of the source of the content, is stored encrypted, is replicated in proportion to the demand for the content, and is deleted independently by each node when that node needs to reuse the space.

We plan for AllNet to anonymize and encrypt communications for several reasons. The first is that there is no reason to believe that the intermediate nodes would be trustworthy, and without encryption, these nodes might tamper with the data transfer. Likewise, pseudonymous anonymity offers some protection against both traffic analysis and accidental or malicious detection of a particular sender's packets. Unlike TOR, if AllNet is used for Internet access, the anonymity would be limited to the ad-hoc part of the network, since encryption over the Internet is only feasible for https access, and in any case, the node(s) connected to the Internet are privy to the entire communication. For the ad-hoc part of the network, like TOR, in AllNet intermediate nodes need not have access to any of the data that is being transferred. It is for this reason that trust needs to be developed among anonymous partners.

4.4 Game Theory and Intrinsic Motivation

Why would anybody offer to freely share their paid Internet connection with total strangers? The motivation to share resources in a network can vary depending upon personality and socio-technical context.

Distributed Computing, P2P networks and Open Source Software present numerous successful projects that rely on voluntary resource sharing. Typically, studies in these areas have relied on Game Theory to identify and design reward-based incentives. By focusing on external rewards, these studies rest mainly on extrinsic motivation (action driven by external and usually measurable benefits).

However, several psychological studies [Glucksberg, 1962; McGraw and McCullers, 1979] show that external rewards may be ineffective or even hinder ("crowding effect") intrinsic motivation, the action driven by personal satisfaction in performing the action itself. Moreover, intrinsic motivation is shown to be very powerful and often more effective, stimulating more sustainable creativity than extrinsic motivation.

AllNet will draw from both Game Theory and Psychology to design a diversity of incentives (as in the Maniac Challenge [DaSilva and MacKenzie, 2007 and 2009]) that support both extrinsic and intrinsic motivation. Although the two types of motivation may seem in conflict, it has been shown that the "crowding effect" can be avoided by taking factors of intrinsic motivations into account in the design of the incentive system [Kunz and Pfaff, 2002].

Although from a different angle, both Game Theory and Psychology underline the importance of social relations in motivating co-operation. Therefore, AllNet is likely to benefit from the development of a social awareness interface [Erickson and Kellogg, 2000], which allows users to be aware of who is sharing bandwidth with whom. A social awareness interface would introduce psychological and social incentives of networking, reputation building and feeling of contributing (altruism), which are shown to be among the strongest motivators of online resource sharing [Holohan and Garg, 2005]. In this way, AllNet can integrate and make visible a diversity of individual, mutual and network benefits of sharing bandwidth with others.

4.5 Digital cash and Bitcoin

There are many different forms of digital cash, including bitcoin [Nakamoto, 2008]. Most of the technical ideas in these systems are not directly applicable to AllNet, because Internet connectivity is required for the system to work. However, the incentive structure of some of these systems, and also P2P systems such as Bittorrent, are applicable in a setting where not everyone is connected. In particular, BitCoin's proof of work idea, and Bittorrent's favoring users who contribute more, is similar to the idea that if an individual has helped others communicate, that individual should get better service from others.

5. Relation to Work in Progress

5.1 Work in Progress by the PI

Taking advantage of my sabbatical, I have invested substantial personal time and effort in the design and initial implementation of AllNet. As well as a few publications, I have implemented two versions of the system, available at http://alnt.org/ , and am working on improving the software and porting it to different platforms. By the time of the projected start of the grant, I am hoping that a chat system built on top of AllNet would allow users on different platforms to communicate with each other.

The current system is effective at the scale at which I have tested it. For example, AllNet supports different forms of addressing, including limited broadcasts and IPv4 and IPv6 addresses for destinations. AllNet already provides end-to-end encryption and secure key exchange.

There is much more that needs to be done, from improving the fundamental peer-to-peer protocol itself, to exploring different forms of addressing, to making sure all the technical decisions made can scale to a larger user base.

5.2 Work in Progress Elsewhere

As mentioned in Section 4, many projects have studied user motivations, devices for emergency communications, peer-to-peer networking, ad-hoc networking, DTN, and similar technologies. Yet nothing similar to AllNet has emerged to date.

The closest system to AllNet that I have observed is the “Personal hotspot” offered by many smartphones. This system generally uses WiFi, bluetooth, or a USB wired connection to allow a computer to access the Internet through the data connection of a smartphone.

This personal hotspot can be extremely useful, but compared to AllNet, has several limitations. It does not focus on prioritizing communications over a low-bandwidth channel, leaving it to the data connection to arbitrarily drop packets that exceed the supported data rate. The owner of the mobile device is responsible for turning off the hotspot when the battery is depleted. Finally, the system uses WiFi security (or no security), which is easy to use once set up, but a little cumbersome the first time it is used.

As the name suggests, personal hotspots are not meant to be used by others, and do very little to build a community of users.

6. General Plan of Work

The activities to be undertaken broadly include:

- Protocol design and evaluation. Further design work on the protocol is needed in response to evaluation of the performance of the current prototype, and considerably better evaluation is needed as well.
- Software development. Further development is needed to turn the current prototype into a more mature product, adding good user interfaces, and to port the current implementation to different platforms.
- Mechanism design to incentivate users. A paper leading to a conference presentation, “AllNet: using Social Connections to Inform Traffic Prioritization and Resource Allocation ”, explores the social network that can be built as a result of exchanging public keys, and how users may be incentivated to contribute resources to AllNet by the knowledge of the social distance between themselves and other users of the network. Further work may lead to very fruitful results in this area.
- Security evaluation. Although the basic mechanisms in AllNet are simple and well-tested, building a secure system requires careful evaluation of the entire system. During my travels this sabbatical, I have met colleagues who may well be interested in performing such an evaluation.

7. Broader Impacts

Section 2 described the expected significance of the proposed work, including how the project will integrate research and education.

One of the broader impacts of AllNet will be to provide wireless communications and help build communities among people who currently cannot afford to do so. Among these are many individuals from low-income groups, who in the U.S. are frequently from underrepresented ethnic groups. Further, people living in rural communities are less likely to be served by the current infrastructure, and thus as a geographic group more likely to benefit from Allnet and the ability to communicate at least locally.

The results of the project will be disseminated widely both through our own work and advocacy, and through the permissive BSD license, which allows manufacturers to integrate the results of this work into their own products.

Ultimately, society at large will benefit from AllNet whenever anyone communicates who would not have been able to communicate otherwise, from the dramatic example of people finding each other after a disaster, to the more mundane issue of communication among hikers.

Conceivably, many applications could be built on top of AllNet, perhaps partly supplanting existing wireless technologies, such as for taxi dispatching, and freeing up the corresponding spectrum.

Being voluntary, free, and independent of providers, but occasionally reliant on the goodwill of strangers, AllNet provides a cooperative, more individualized model of communication than is commonly available at this time. The closest existing parallels to AllNet are CB and ham radios, both of which are used under normal circumstances, but occasionally prove invaluable in emergencies. AllNet would expand these concepts and make them available to the public at large, to be used freely.

In a slightly different context, AllNet somewhat resembles wikipedia in relying on the goodwill of others to provide a very valuable service and in building a community of people with a common goal. As for wikipedia, the goal of AllNet is better communication among people. Unlike wikipedia, the goal is achieved by providing a more reliable, trustworthy, and secure communications medium, and applications to take advantage of the medium’s unique properties.

The remainder of this document describes technical details of AllNet that may help reviewers trust in the feasibility of the project and better understand the contributions to knowledge that funding the project will bring.

8. Algorithms and Protocols

AllNet is designed for low-rate communication when other forms of connectivity fail. Encrypting small amounts of data should require little enough energy and computational power that it is acceptable to require that all private communication be encrypted using public-key encryption.

There are many ways of selecting, communicating, and certifying public keys. In a social network, it seems reasonable that individuals communicating with each other in person might directly exchange a public key, or exchange a public key through friends, obviating the need for certifying authorities and the consequent security vulnerabilities seen in the DigiNotar episode [Electronista, 2011]. The secret and public key pair are chosen at random, and, as in peer-to-peer systems, any individual may hold any number of key pairs.

If every user has at least one secret/public key pair, this can be used to lessen the need for passwords when communicating with web sites, as the secret key can be used to identify the user. The secret key can be kept encrypted on local storage. If individuals choose to have more secure and less secure key pairs, they can do so, and the level of protection can be different for different keys.

Assuming that randomly selected public keys rarely or never collide, they can be used as pseudonymous user identifiers and also as part of the destination address of AllNet packets.

AllNet addresses currently have several different forms, typically including one or more of the following: an IPv4 or an IPv6 address, a domain name, some or all of the bits from the recipient's public key, a randomly chosen identifier, and physical coordinates. Addresses can also specify every recipient within a number of hops or a geographical distance around one of the above addresses.

With such an addressing scheme, there are two main challenges: routing and decrypting. We first discuss the challenge of key exchange and decrypting the messages, and then the challenges of routing to addresses that are not topologically based. The remainder of this section then looks at some of the applications AllNet is designed to support.

8.1 Decrypting Messages and Initial Key exchange

Decrypting is a challenge because each node in the network will receive many packets, only some of which are addressed to itself. To detect which packets are addressed to itself, either the message must be decrypted, which is computationally expensive, or the message must carry, in the clear, an identifier (possibly the public key) identifying the intended recipient. We favor a flexible solution, in which a number of bits of the identifier are carried in the clear with each packet. We call this a "hint" of the identifier, and it can be between 0 bits (no hint) and the number of bits in the identifier (unique identification, assuming no collisions). We expect that the hint will often be between 8 and 16 bits. This allows recipients to only attempt to decode a small fraction of the packets, but still protects recipients against directly revealing their identifier to all intermediate nodes, and thereby lessens the threat of traffic analysis.

As is common (for example in SSL/TLS), for large transfers the public key is used to encrypt a randomly selected AES key, which can be used for an entire session.

Exchanging the public key is simple if two individuals are near enough to each other that they can exchange a short confidential string. As currently defined and implemented in AllNet, one of the two users, Alice, must enter the string (currently 12 alphabetic characters) displayed on the other user's device. Alice's device sends to the other user's (Bob's) device its public key, a nonce, and a hash of the (public key, nonce, and confidential string). Bob's device replies with a similar message containing his public key. While these messages are in the clear, an attacker can gain little knowledge from this exchange other than the two parties' public keys. The shared string prevents acceptance of unknown keys, and the nonce adds randomness to the exchange in case it is needed.

It has been suggested to me that this key exchange resembles an HMAC, and a conventional HMAC may provide better security. I am currently studying this, and if appropriate, will modify the design correspondingly.

8.2 Forwarding and Routing

Assuming that a node is willing to carry AllNet traffic, it must make decisions about the resources to be allocated to each packet. In particular, it must decide whether to forward the packet at all, and if so, with what priority.

Since AllNet is designed for low bit-rate and delay insensitive communication, it is usually better for a node to forward all its own traffic before forwarding any AllNet traffic. This means that only nodes that are not congested will forward AllNet traffic. This is appropriate for a low priority service. We note that protocols such as TCP often cause temporary congestion, so if one or more AllNet packets are buffered in a low priority queue, they can be transmitted whenever the channel becomes available, e.g. due to TCP backoff.

It is reasonable to put a maximum cap on the total amount of resources devoted to AllNet. While the amount of bandwidth, storage, and processing power devoted to AllNet is an individual decision, the design goal is for AllNet to work well while consuming at most 1% of any one resource on a mobile device.

Since there are at least two levels of priority, normal traffic and AllNet, we believe it would be useful to have multiple priority levels within AllNet itself. While ultimately each node must determine its own forwarding policies, we have a few suggestions that might maximize the usefulness of forwarding each packet.

- Packets with a better defined destination are normally favored over packets with a less well defined destination. For example, packets addressed to a specific IP address should get priority over packets addressed to a large geographic area. Packets with a destination defined by many bits of the identifier, and therefore which need to be decoded by only a small number of recipients, should be forwarded with higher priority than packets whose decryption would be attempted by a large fraction of the recipients.
- Smaller packets should be favored over larger packets, and sources that send less data may be favored over sources that send more data. Since identifiers can be manufactured at will, the priority must be applied on a hop-by-hop basis among neighbors. Also, this rule should be applied probabilistically rather than strictly, to avoid offering any incentive to break larger amounts of data into small packets just to get around this rule. The point is, transmissions should get better service if they require fewer resources, but even large transmissions may be supported on a very limited basis.
- Packets from friends should get better service than packets from strangers, with the level of service monotonically decreasing with the social distance between the sender and the owner of

the forwarding device.

The general form of routing in AllNet is a loose source route that may include limited broadcast and unicast addresses. The broadcast is limited by hop count, geographic area, or both (nodes that are not location aware do not participate in geocasting). The initial recipient may be identified by IP address, domain name, or anycast to any node within a given geographic area. While the hop count limitation could in theory be used on wired networks such as the Internet, at least initially we expect it would only be used in wireless ad-hoc networks, such as WiFi used in ad-hoc mode, or 802.15.4/Zigbee based networks. Such limited broadcast and many of its variants have been studied thoroughly, for example by El Fawal et al. [El Fawal, Le Boudec, and Salamatian, 2007].

The unicast portion of the source route can be used to send a message across the Internet, to a node that is believed to be able to reach the intended destination using a limited broadcast.

In order to provide the necessary forwarding, each node may tell its neighbors about its IP address(es), its geographic coordinates, its domain name, and whether it is willing to forward packets over the Internet or a wired or wireless data link. As long as nodes choose to forward this information from other nodes, the wireless ad-hoc network can set up one or more gradients, indicating neighbor nodes that may be used to reach the Internet.

As in Delay Tolerant Networks, the limited broadcast need not happen immediately. In fact, nodes may store and keep forwarding messages for a while, and forward them when the opportunity arises. This supports the "data mule" concept where movement of the node itself helps carry the data to the intended destination.

Using this scheme, it would be common for a message to be delivered more than once. This means nodes must be designed to discard duplicates, received no matter where from. Messages based on Bloom Filters can be very helpful for fast detection that a message has not been seen before.

8.3 Application-level Protocols

There are three major applications that we wish to provide special support for. They are web access, email access to an existing mailbox, and chat. Web access, in turn, includes both HTTP and SSL/TLS. Email might only include mail access protocols such as POP and IMAP, and particularly their secure versions, since anonymous SMTP makes it too easy to generate unsolicited commercial email. Finally, there is a wide variety of chat systems, which we consider to span the range from SMS to Twitter as well as conventional chat. These all have enough characteristics in common that a generic protocol to support chat might provide many of the benefits of all of these systems.

Web Access

A simple web request header is usually on the order of hundreds bytes or more. In a resource-constrained network, this is a rather substantial number of bytes for what is often a very simple request. When the requested URL is long, this must of course be transmitted, but the remaining fields of the header could usually be compressed to a fraction of their size.

In addition, a web request might sometimes be a HEAD request or an "If-modified-since" request, or a simple request for a small amount of data. If TCP connection setup is expensive, as it could easily be the case in an ad-hoc network, then there is little reason to set up a TCP connection for what is, after all, an exchange of two packets. If no response is obtained, the request can be retransmitted at a fixed rate a limited number of times (to avoid causing congestion collapse), or can be resent after establishing a TCP connection.

We therefore propose a new web access protocol, that would provide the following:

- header compression for both request and response headers
- UDP delivery of simple request-response messages when no established TCP connection would be effective.

These features somewhat resemble the strategy used in the DNS protocol, to use UDP for simple exchanges, and TCP when reliability is essential and for larger transfers.

One advantage of using UDP is that limited web service could be provided even to nodes that are only connected through a DTN network, where the response might take hours or days rather than fractions of a second. Users would have to learn how to use such an asynchronous web service, and perhaps even be given the ability to cancel outstanding requests that are no longer needed, but given the variety of services that users have learned to use over the past decade, this seems feasible.

Ultimately, if we are successful, many browsers and servers would incorporate such a protocol, and might use sophisticated strategies to determine when to use UDP and when to use TCP. I propose to develop the compression algorithm and the protocol, and implement them in a proxy that could be used by the browser on the local system.

Email

Email can be viewed as a mechanism to asynchronously transfer arbitrary amounts of data to an arbitrary recipient.

Interestingly, partial delivery of an email message is rarely useful, except to technically proficient recipients. Though delivery of the attachments of an email can often be avoided, even an attachment usually needs to be delivered as a whole. So, we would suggest that each email be either completely delivered, or discarded. All-or-none delivery of asynchronous objects ("messages") over a resource-limited network suggests favoring smaller messages over larger messages and messages from less frequent senders over messages from senders that send more frequently. In addition, as mentioned above, targeted messages to a single recipient will be favored over messages sent to a larger number of potential recipients.

If the network is connected, and especially for large messages, TCP and the usual email protocols should be used, since they are efficient for large content. On delay tolerant or intermittently connected networks, we also wish to support the exchange of small email messages and mailbox summaries.

In such circumstances, again we can use UDP to send messages. The message may be lost, and even if it is not, the sender might not get confirmation of delivery. While reliable delivery is preferred, unreliable delivery is better than no delivery at all.

Splitting up a large message into smaller messages would only be useful if the individual pieces have a meaning on their own. Otherwise, if resources are limited, there is a high likelihood that at least one of the smaller messages would be discarded. This would require retransmission. While all of this may ultimately be successful, the incentives are to transmit small messages rather than large messages, and this is exactly the kind of incentive a resource-limited network should be offering.

I plan to implement at least one protocol to support asynchronous email access to an existing mailbox.

Chat

The major difference between email as described above and chat systems is that most chat systems already limit the size of individual messages sent. Also, users of chat systems generally have an

expectation that messages will be delivered promptly, and may expect a confirmation of delivery. Experience with existing chat systems shows that good feedback systems, such as used in the Blackberry Messenger (BBM) or WhatsApp, can help communicate to users exactly how much the system knows about message delivery, and this can help adjust users' expectations to the reality of the situation.

In this project, we expect to implement at least one chat protocol that works well in either a DTN, or a connected low-resource network. A very simple prototype of this application is available for Linux in Version 1.0 of the AllNet source code distribution, available at http://alnt.org/

Cellphone Walkie-Talkie

Where devices are within range of each other, it is perfectly sensible to expend as much energy as their owners desire to provide high-quality wireless communications. A walkie-talkie application with optional video or image transfer can be useful in a variety of scenarios, but particularly when hiking or when multimedia transfer is needed in more mundane situations within a building.

Interesting, most of the mechanics of implementing such an application are provided by the mobile devices themselves, and the major challenges are detecting that the devices are in range and providing the user interface needed for the exchange. The first challenge will be solved routinely by AllNet, with the current versions providing an early version of mutual discovery when within wireless range.

9. Summary: Connecting Intermittently Connected Devices

While the Internet has amazing and wonderful properties, good support for intermittently connected devices is not one of them.

This proposal seeks to improve this situation, both at the network layer, and at the application layer for selected applications: chat, web access, and email access. Assuming that devices with limited or intermittent connectivity are also limited in their resources, we focus on supporting transmission of small amounts of data where they can be most useful. To encourage users to trust the network, we encrypt as much as possible of the communication, and provide mechanisms to lessen the opportunities to perform traffic analysis or denial of service.

The ultimate goal of this project is to provide a network that is useful for low bit rate communications among arbitrary parties when the infrastructure is unavailable, and especially during emergencies. In order to train people to use this network, and to motivate them to use it, we hope to make this network useful in a variety of common situations. As well, we plan to support high-bandwidth communications where the underlying technology makes it practical.

If universally or at least widely implemented, for example on most laptops and WiFi-capable smartphones, AllNet can provide emergency communication, private interpersonal and group communication, and low bit-rate communication with the Internet. This will most useful when circumstances do not permit access through regular infrastructure, for example, when people travel outside the area served by their usual provider, but we expect it will prove useful on a daily basis as well.

It has been said that the Internet was designed to be a network that would survive nuclear war. It has, of course, turned out to be much more than that. If this proposal only provides connectivity in cases of emergency, it will be very useful. And if it ends up providing services to many people in their daily lives, its usefulness and worth will be even greater.