TWC SBES: Small:

Allnet: Trustworthy Ubiquitous Connectivity

A sharing mechanism and secure protocol for delay-tolerant ad-hoc message exchange

1. Goals

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, however, could be configured to perform the communication directly. Even when my handset is not directly within range of my peer's handset, principles of mesh networks, ad-hoc networks, and delay-tolerant networks may be used to perform the communication by sending the data through other devices that may be located between the two endpoints.

We propose to enhance the current infrastructure-based communication among individuals with a secure infrastructureless peer-to-peer network that may be able to provide support whenever 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 as well as 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 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 making it available at all times.

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.

We propose to call this system AllNet, because it is intended to provide networking to all people at all times.

Allnet is designed to be particularly useful when other forms of networking fail or are not available. We also expect that some of the peer-to-peer communication features of AllNet may be sufficiently useful that they will be used even when Internet or other connections are available.

With AllNet, individual communication devices (cellphones, smartphones, 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.

While the goal is ambitious, the technical requirements are relatively modest. Ad-Hoc networks have been studied very thoroughly, and may not need much innovation at the lowest protocol levels. AllNet does leverage existing encryption technology to provide innovative privacy and security for the communications, such as to thwart not only the reading of messages, but also to some extent traffic analysis and denial of service.

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 more challenging, but still feasible, for web access.

Beyond these technical challenges, creating AllNet poses novel design requirements as well. 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 emailwhen there is no direct Internet connection, perhaps 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 plan 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.

After reviewing some of the related work that inspires the design of AllNet, we present our preliminary ideas for addressing these challenges.

2. Related Work

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.

2.1 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, the primary author of this proposal has 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 our 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.

2.2 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.

2.3 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.

2.4 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.

2.5 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.

3. 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 on a regular basis 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 any peer-to-peer system, 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 address.

AllNet addresses will have several different forms, typically including one or more of the following: an IPv4 or an IPv6 address, a domain name, a telephone number, some or all of the bits from the user'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 decrypting the messages, and then the challenges of routing to addresses that are not topologically based. The remainder of this section analyzes the security of AllNet and some of the applications we intend to support.

3.1 Decrypting Messages

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 revealing their identifier to all intermediate nodes, and thereby lessens the threat of traffic analysis.

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

In general, the identification of the user or session can be kept confidential, and all that is needed in the clear is enough information to reduce the number of decrypting attempts needed by each user. This is similar to a Bloom filter, where false positives are acceptable, that is, set membership may be declared incorrectly, as long as there are no false negatives, that is, if someone is a member of the set, the Bloom filter will always confirm membership. In this case, the cost to the recipient of a false positive is a decryption operation followed by discarding of the packet if the decryption does not produce a meaningful packet. To simplify this detection, it is easy to include a magic number at a specific position within a packet. The magic number can be a default value, or can be chosen at random for a specific session.

This mechanism is general enough to support either one-on-one communication or group communication. A session may be, as is currently conventional, a login to a web site, or could be set up for a fixed period of time, for example a week, among a group of friends or coworkers. As part of this project, we expect to develop a secure distributed chat application where the key changes on a regular basis, and members have to re-authenticate to obtain the new key. This supports secure dynamic group membership, where each individual only has access to the messages sent while the person’s group

membership is active.

3.2 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 would also be 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, we plan to at least characterize the usefulness of AllNet under different scenarios for resource availability on different nodes.

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.

- Users are allowed to tag some transmissions as being for emergencies, and these transmissions would normally be given higher priority and perhaps more resources. In case this feature is abused systematically, it can be turned off.
- 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.

The general form of routing in AllNet is unicast to an initial recipient, followed by limited broadcast from this initial recipient. 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, telephone number, 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 routing 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 routing, 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 “phone” 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, messages can be tagged with an expiration date, and nodes may store and keep forwarding messages until they expire. This supports the "data mule" concept where movement of the node itself helps carry the data to the intended destination. Again, to lessen the incentive for misuse, data with shorter expiration dates can be given higher priority than data with longer expiration dates.

Using this scheme, it would be common for a message to be delivered more than once. This means the message must contain a distinctive sequence number to eliminate duplicates. This sequence number can be in the encrypted part of the message, or may be in the clear, depending on whether the sender wishes for only the final recipient to discard duplicates, or intermediate nodes as well.

3.3 Security: Alice and Bob

To examine the security provided by AllNet, consider AllNet user Alice at a location where she does not have Internet access, but she can reach another AllNet user, Bob, that does have access and is willing to share resources. Bob may not be directly in range, but can perhaps be reached via a limited broadcast.

Assuming that Alice is trying to reach the web site charlie.com, she will use limited broadcasting to send packets to Bob, who will then act as a NAT and forward the packets as required. Any packets returned by charlie.com on the port selected by Bob for this communication, will be sent back to Alice. If the connection is reliable (no packets dropped), multihop ad-hoc unicast transmission can be used instead of the limited broadcast.

In this example, we note that Alice does not need to establish a TCP connection to charlie.com, or even to Bob. Instead, Bob can perform the TCP connection setup and management with charlie.com. We intend to design a low-overhead protocol for web requests and responses that is suitable for a multihop wireless ad-hoc network, in this case, between Alice and Bob. In particular, a normal HTTP request could be compressed and (usually) be sent as a single, relatively short packet. If the (compressed) reply is sufficiently short, it could also be sent as a single packet. The reply can also include any embedded content. If Bob saves the exchange for a little while, Alice can request retransmission and receive it directly from Bob. To allow users to reload the page, each request must include a counter which Alice will increment each time before reloading the same page. This allows Bob to distinguish a duplicate request, which can result in a retransmission of the original response, from a request to reload the page, which will result in additional communication with charlie.com.

For HTTP traffic, since Bob sends the request to charlie.com and receives the reply, Alice encodes the request with a key shared with Bob, and Bob can see everything that Alice and charlie.com communicate to each other. If charlie.com supports TLS and Alice connects using an https address, then Bob is no longer able to see the contents of the communication.

Now suppose that Alice wants to communicate with Diane, who is able to communicate with Elsie, who is connected to the Internet at IP address 1.2.3.4. Somehow, perhaps through a mechanism similar to Dynamic DNS [Vixie, 1997], Alice finds out Elsie's IP address. She can then broadcast or unicast the message to Bob, who sends it to Elsie's IP address, who broadcasts or unicasts the message to

Diane.

It is clear from these examples that AllNet messages need a way to specify how they are to be sent across the Internet. The mechanism itself need not be any more complicated than the four address fields in IEEE 802.11.

In general, we nodes on AllNet that are offering a specific service should be able advertise this willingness. Internet connectivity is a very important specific service at this time, but in the future other services may be just as essential, and so AllNet should support service discovery in general.

If multiple nodes are providing the same service, any given user will generally attempt to use more than one service provider, to lessen the load on any one provider, and also lessen the likelihood that any one provider may have access to too much data. On the other hand, if one provider gives better service than the other, there is nothing wrong with using the more favorable provider.

3.4 Protection against attacks: Alice, Bob, and Mallory

AllNet is intended to be secure by design. For example, the near-pervasive encryption and the user addressing scheme offer some protection against snooping and traffic analysis.

To use a simple example, AllNet should be secure enough to allow anonymous communication within a classroom, without one's classmates being able to determine who is communicating with whom, nor being able to disrupt communication among individuals, and without being able to determine the contents of private communication or to spoof messages from other individuals. Of course a teacher (or repressive government) may prohibit any use of electronic devices, but once these devices are permitted, it should be really hard to control or even find out who is communicating with whom.

AllNet is in part a mixing network though, unlike TOR, packets are not re-encrypted at each node. The traffic sent by a node may or may actually have originated at another node. In addition, except when Bob is transmitting web traffic in the clear for Alice, traffic is generally encrypted. Finally, if there are multiple equivalent paths to a destination, different packets from the same source will often take different routes.

All of these make it a lot harder for an attacker, Mallory or Eve, to either read the contents of Alice's transmissions, or change them.

There are some mechanisms Mallory might use to attempt to affect Alice's transmissions. One is to be near Alice, and advertise her own address as being very well connected to the Internet and to different geographic areas. Alice would then normally want to use Mallory to route her data. This would give Mallory the opportunity to perform traffic analysis, and potentially also provide a black hole by dropping some or all of Alice's packets.

The defense against traffic analysis attacks is that Alice might send her messages with few if any cleartext bits of Bob's and her own identification. Mallory would then have to guess which messages originate from Alice and are being sent to Bob, and which messages Alice is forwarding for other nodes.

The defense against black holes has at least two strategies. Over the long term, Alice and Bob can try different paths, and generally use more the ones that are more reliable. In a situation where node connectivity changes frequently, however, it is hard to collect good statistics, so the other defense simply consists in sending all data through multiple gateways. If this duplication is done judiciously, it will be a lot harder for an attacker to perform a denial of service attack.

3.5 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, usability might not be a big issue.

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. We propose to develop the compression algorithm and the protocol, and to do limited testing to evaluate its performance.

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 suppport small email message exchange.

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.

We 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), 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.

4. Resource Management

We assume that nodes will forward their own traffic before forwarding any other nodes' traffic. We also assume that nodes with some spare capacity are willing to devote some resources to forwarding AllNet traffic. These resources fall into three major categories:

- network bandwidth, and the energy needed to transfer data
- memory and storage for messages
- the CPU power, and corresponding energy, needed to decrypt incoming messages, and to a lesser extent, to encrypt outgoing messages

All of these (except energy) can be used for AllNet on an "as available" basis. For example, if there is no storage available for incoming packets, the packets should be discarded. If the CPU is not idle, it should be used for user- or system-centered tasks rather than for AllNet.

Given these requirements, system administrators, which may be the users themselves, or whoever selects the default configuration for the device, must decide how much of the available resources to devote to AllNet. Of particular concern for mobile devices is energy use. Further, bandwidth that appears to be "free" to the mobile node may not appear to be free to network operators. As a result,

even when resources appear to be available (e.g., while a mobile device is charging), AllNet should be as conservative as possible about resource usage.

As an initial design goal, we would like to explore a 1% rule, where AllNet is constrained to use at most 1% of the available CPU, memory, and bandwidth, and at most 1% of the average power consumption of a device. Under many circumstances, we hope to show that even with such a low resource budget, AllNet will prove to be useful as a backup to conventional connectivity, and also for local, chat-like communications among peers who are in a social relationship. As an alternative to this 1% rule, we plan to explore different rules and see how much they affect the user experience and service availability.

5. Broader Impacts

At its core, AllNet is a message-passing system designed to support social networks and communication among individuals and small groups. Being infrastructure-less and decentralized, control of the communication rests with the individuals involved and the intermediaries who forward the messages. This merging of peer-to-peer and social networking technologies can provide more opportunities for communication than are available today.

Fundamentally, the idea for AllNet came from the desire to support emergency communications, and the realization that, although hardware capable of peer-to-peer decentralized communication is already widely available, we must also provide the software and training needed should the infrastructure fail. By providing a networking technology that is useful in everyday life, users will wilingly train themselves to install the necessary software and use the required user interfaces.

The ability to exchange messages with peers even without infrastructure support can make a dramatic difference in a variety of scenarios, from road trips, to wilderness travel, to people who occasionally avoid communicating due to the cost of using commercial systems. 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.

Like most research grants, this one will also help train the next generation of researchers. One student, Caterina Desiato, in our Communications and Information Sciences interdisciplinary Ph.D. program, has already contributed to this proposal and will be hired if it is funded. She is ideally suited for this project, having a background in both Philosophy and Computer Science.

In addition, the budget includes funds to hire undergraduate students. 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 recently at a seminar, and found out he is now on the faculty at a local community college.

6. 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.

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 indeed. And if it ends up providing services to many people in their daily lives, its usefulness and worth will be even greater.