

Pocket Switched Networks and Human Mobility in Conference Environments

Pan Hui^{*#}, Augustin Chaintreau[#], James Scott[#],
Richard Gass[#], Jon Crowcroft^{*} and Christophe Diot[#]

^{*} University of Cambridge
firstname.lastname@cl.cam.ac.uk

[#] Intel Research Cambridge
{augustin.chaintreau, james.w.scott, richard.gass, christophe.diot}@intel.com

ABSTRACT

Pocket Switched Networks (PSN) make use of both human mobility and local/global connectivity in order to transfer data between mobile users' devices. This falls under the Delay Tolerant Networking (DTN) space, focusing on the use of opportunistic networking. One key problem in PSN is in designing forwarding algorithms which cope with human mobility patterns. We present an experiment measuring forty-one humans' mobility at the Infocom 2005 conference. The results of this experiment are similar to our previous experiments in corporate and academic working environments, in exhibiting a power-law distribution for the time between node contacts. We then discuss the implications of these results on the design of forwarding algorithms for PSN.

Categories and Subject Descriptors

C.2.1 [Computer Systems Organization]: Computer-Communication Networks—*Networking Architecture and Design*

General Terms

Measurement, Experimentation, Human Factors

Keywords

Mobile networking, Delay-tolerant networking, Network measurement, Wireless networking

1. INTRODUCTION

Internet users (and in particular mobile users) can report two very different experiences with networking services. In regions where Internet connectivity is available to their device, the experience is generally good, and the user can access all Internet applications. On the other hand, in between

such "Internet connectivity islands", users find themselves unable to send or receive any information, even if the party that the user wishes to communicate with is within range of his/her wireless device and has wireless communication capabilities.

This is due to how current network services have been deployed: they rely on managed infrastructure (e.g. DNS, DHCP, centralised servers, etc.), and use network protocols requiring continuous connectivity between the communicating parties throughout a transmission (e.g. TCP). These constraints can not be met easily by mobile devices. Therefore, in the absence of end-to-end connectivity, the state of the art in network communication is the "memory stick hand-shake" protocol, well known to all networking nomads.

We propose *Pocket Switched Networking (PSN)* aimed at enabling network services for mobile users even when they are not in reach of Internet connectivity islands. PSN falls under the more general space of Delay Tolerant Networking (DTN), which looks at enabling communication in the absence of end-to-end connectivity.

This paper has three main contributions. First, we describe the ideas behind Pocket Switched Networking and some of the research issues that it encompasses. Second, we describes the latest research in that project, which was a deployment of small Bluetooth devices to 54 participants of IEEE Infocom 2005 in order to measure the implications of human mobility patterns for PSN. The latter follows on from our previous work [1], which presents similar deployments in corporate research lab and university settings. Third, we describe some implications of these results for the design of PSN forwarding algorithms.

2. POCKET SWITCHED NETWORKING

Delay Tolerant Networking (DTN) explores networking in the presence of networks with challenged networking conditions, such as links which are often disconnected, or are subject to long delays. To date, activities in the DTN area have addressed various communication environments where standard Internet protocols would be difficult to use or would provide very poor performance, e.g. networking using buses following predictable routes [2], interplanetary networking [3], interfacing with sensor networks [4], and using mobile nodes to bridge data between remote village networks and the Internet [5, 6].

Another scenario in which DTN can be useful is in net-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGCOMM'05 Workshops, August 22–26, 2005, Philadelphia, PA, USA.
Copyright 2005 ACM 1-59593-026-4/05/0008 ...\$5.00.

working for devices carried by users of mobile and portable devices. For example, mobile workers move between connectivity islands (e.g., WiFi at home and work). Outside these islands, end-to-end connectivity becomes expensive, slow, or simply unavailable. Moreover, many communication services rely on access to centralised resources such as the DNS. That prevents, for example, two users sitting beside each other from easily exchanging data.

There is a huge amount of untapped resource in portable networked devices such as laptops, PDAs and mobile phones, including in local wireless bandwidth (e.g. 802.11 and Bluetooth), storage capacity, and CPU power. The only scarce resource is power, but advances in power engineering and battery technologies have meant that mobile phones now last for a week on a single charge, while remaining in constant network contact. We expect that this innovation will continue, allowing devices to participate in wireless networks while minimising power consumption.

We envision a world where these resources can be used to provide networking functionality alongside access networks, and where users' applications make use of both types of bandwidth transparently. This is the goal of Pocket Switched Networking.

We now discuss some of the key features of the PSN problem space.

2.1 Human mobility

Mobility is a double-edged sword. User mobility has the potential to increase network bandwidth, as large amounts of data can be carried around the network using device storage. However, mobility also makes it challenging to communicate with users, as forwarding paths may be unstable and device reachability may be highly variable.

While previous work has looked at wireless network proximity in mobile animals [7, 8], little work has been done to determine this for human-carried devices. We present results in this area in Section 3.

2.2 Opportunistic networking

Internet routing assumes that a contemporaneous path between two nodes exists. The mobile ad-hoc networking community has focused on techniques for determining end-to-end contemporaneous paths in a network of mobile nodes. However, using contemporaneous paths is restrictive, and excludes the possibility of data transport using local connectivity and node mobility, which has in the past been referred to as data muling [9], or store-and-haul forwarding [10].

In PSN, we do not try to find or build end-to-end paths. Instead, data is forwarded hop-by-hop, taking advantage of any opportunities in the course of device mobility. One type of opportunity is found in local network connectivity (using wireless or otherwise). Whenever two PSN nodes come into contact, they must detect each other and determine what to transfer in each direction.

In addition to forwarding data using local connection opportunities, a PSN node may, at some times and in some places, have global connectivity, e.g. via WiFi access infrastructure or GPRS. Note that it is not always the case that global connectivity is more useful than local connectivity. Local networking can provide better service than Internet access if the corresponding party is nearby, either because the correspondent does not themselves have global connectivity, because the traffic requires a high bandwidth or low

latency, or because global connectivity is expensive (e.g. in an airport lounge).

On the other hand, for traffic with no local destination, forwarding it towards the Internet (either using global connectivity yourself or by forwarding locally towards nodes with global connectivity) may be the fastest method of reaching the recipient. The tricky part may be in the return path: how does a node on the Internet forward messages to a mobile node which does not itself attach to the Internet, but which passes near other nodes which are. This is an open problem.

The use of both local and global opportunities allows PSN to provide highly-robust networking for users, as it can transparently switch over to local connectivity when global connectivity is unexpectedly lost. This all-too-often situation can be brought on by hardware failure, software bugs, or misconfiguration. It may also be useful in situations of natural or human-made disasters, in which users' mobile phones and laptops become useless for communications purposes, but opportunistic networking would allow the physical movement of people to transfer important messages.

2.3 Personal devices

Pocket switched networking targets devices which are always-on and always-carried by users. Such devices include mobile phones, PDAs, and laptops. This focus on personal devices has implications for the networking requirements. User-facing applications such as web browsers and email/messaging clients must be able to communicate the status of the network (e.g. expected performance for a large file transfer) and of individual transactions (e.g. whether an instant message was received).

With personal devices, the primary aim of networking conducted by the device must be the support of its owner's tasks. Providing support for others' networking (e.g. by storing and forwarding data for them) is possible when there is spare resource (CPU, storage, networking bandwidth, and battery). Establishing and maintaining trust relationships and the provision of incentives may be important considerations, since selfish behavior would naturally cause each node to decline service to other nodes. Security and privacy must also be examined, since data may traverse many uncontrolled and potentially malicious nodes. These are more open issues in the design of PSN.

3. MEASURING HUMAN MOBILITY

As we have described, human mobility plays a key role in PSN, as it is mobility which gives rise to local connection opportunities when access infrastructure is not available. In order to explore this further, we chose to conduct real-world deployments of devices to members of various communities, allowing us to determine the effects of users' mobility patterns on the prevalence of networking opportunities. In this section we present an experiment conducted within a group of conference attendees. This experiment follows on from similar experiments presented in [1], which involved both students at universities and members of a research lab.

3.1 Experimental setup

The devices used to collect connection opportunity data and mobility statistics in this experiment is the Intel iMote. This is a small platform designed for embedded operation,




Figure 1: Picture of an iMote with battery.



Figure 2: iMotes packaged for the experiment.

comprising an ARM processor, Bluetooth radio, and flash RAM, and are shown with a CR2 battery in Figure 1. We packaged these devices in a dental floss box, as shown in Figure 2, due to their ideal size, low weight, and hard plastic shell.

Fifty-four of these boxes were distributed to attendees at the IEEE Infocom conference in Miami in March 2005 (which had eight hundred attendees in total). The volunteers were asked to keep the iMote with them for as much of their day as possible, with most carrying the iMote in their pockets, and they were given a small gift in compensation. Volunteers were chosen to belong to a wide range of organizations — more than thirty were represented. To assure the participants of their anonymity, we did not record the MAC address of the iMotes that they were given, instead only recording an uncorrelated number printed on the outside of the box, so that we could perform the logistics of distribution and collection. Of the fifty-four iMotes distributed, forty-one yielded useful data, eleven did not contain useful data because of various failures with the battery and packaging, and two were not returned.

The iMotes were configured to perform a Bluetooth baseband layer “inquiry” discovering the MAC addresses of other Bluetooth nodes in range, with the inquiry mode enabled for five seconds. Despite the Bluetooth specification recommending that inquiry last for ten seconds, preliminary experiments showed that five seconds is sufficient to consistently discover all nearby devices, while halving the power-expensive inquiry phase. Between inquiry periods, the iMotes were placed in a sleep mode in which they respond to inquiries but are not otherwise active, for a duration of 120 seconds plus or minus twelve seconds in a uniform random distribution. The randomness was added to the sleep interval in order to avoid a situation where iMotes’ timers were in sync, since two iMotes performing inquiry simultaneously cannot see each other. However, we still expect iMotes to fail to see each other during inquiry around four percent of the time (the duty cycle).

The results of inquiry were written to flash RAM. Since flash capacity is limited (64K for data), we could not store the full result of each inquiry without running the risk of exhausting the memory. Instead, we decided to record “contact periods”. This is achieved by maintaining an “in-contact” list comprising the Bluetooth MAC addresses of the nodes that are currently visible. When device on this list stops responding to inquiries, we store a record of the form {MAC, start time, end time}. Preliminary tests revealed the following problem: Bluetooth devices on a specific brand of mobile phone did not show up consistently during inquiries (and increasing the inquiry period to ten seconds did not help). Therefore, a small number of nodes were causing the memory to fill too quickly. To avoid this problem, we keep a device in the “in-contact list” even if it is not seen for one inquiry interval. If it comes back in-contact on the next interval, nothing is stored. If it does not, a record is stored as normal. This solves the problem, at the expense of not being able to detect actual cases where a node moved out of range during one two-minute period, and back into range for the next two-minute period.


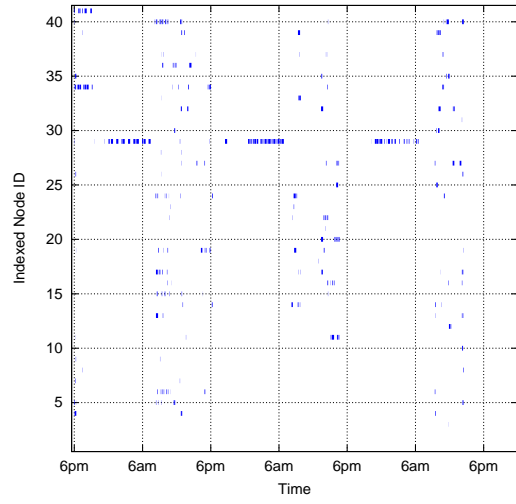


Figure 3: Histograms of contacts seen by an iMote: other iMotes (left) and all other device types (right).





Figure 4: Distributions of contact times (left) and inter-contact times (right) for pairs of nodes.





Figure 5: Distribution of any-contact times (left) and inter-any-contact times (right).

4. ANALYSIS OF CONFERENCE MOBILITY PATTERNS

In Figure 3 we show a graphical representation of the log from a typical iMote. These plots show a unique ID number on the Y axis, with time on the X axis. We order the IDs such that the first 41 are iMotes, and all others are “external” devices, i.e. Bluetooth devices which are not iMotes. The left plot shows just “internal” iMote-iMote sightings for iMote number 4, (whose MAC address ends with “6057”), while the right plot shows all sightings of external addresses from the same iMote.

In addition to the 41 iMotes from our experiments, we recorded 182 external devices. While external devices are in greater number, they are seen significantly less. We recorded 22459 contacts between iMotes, and 5791 contacts between an iMote and an external device. We conjecture that many of the external devices were not actively carried in the same way as iMotes, either because they were only powered on (and discoverable) for some portion of the experiment time, or because they were left by their owners in static locations. Another factor is that many of the external nodes may have been seen by an iMote at some distance from the conference, and therefore only seen once. The maximum number of other iMotes an iMote had in range at one time was twelve, while the maximum number of external nodes in range at one time was seven.

4.1 Contact and inter-contact times

For a given pair of nodes A and B, the timeline can be divided into two regions, “contact times” and “inter-contact times”. The contact times are when A and B are in range of one another, and could therefore have sent data if they had wished to. Inter-contact times are simply the times between the contact times, when data is not transferrable directly between A and B.

The distribution of contact durations is given in Figure 4 (left). As with the previous experiment sets, the contact durations follow an approximate power law. We do not pursue the contact durations further in this paper, though they

would be relevant to a discussion of the bandwidth available in various implementations of a PSN, which is left for future work.

The behaviour of inter-contact times is important when considering the delay experienced by packets in a PSN. This is the time a node has to wait to get in contact with a specific node (as seen immediately after losing contact with that node). In Figure 4 (right), we display the distribution of inter-contact times aggregated across all iMotes, for internal sightings and external sightings. As with our previous experiments, the inter-contact distribution exhibits a strong heavy tail property, which can be observed on the plot as an approximate power law for the time scale [2min:1day], with coefficient 0.4. The step effect to the left of the graph are caused by the quantisation imposed by our measurement period, while the fall-off towards the right of the graph is due to the duration of the experiment limiting the possibility of seeing long inter-contact times. The diagram also shows that external sightings behave very similar to internal sightings. This is encouraging since external nodes were not introduced by ourselves but were already present, indicating that our measurements are representative of real-world connection opportunities.

The results presented here confirm and generalise our observations in [1] to another trace of human mobility in a different environment. In this dataset, the power law is even more apparent than in the previous experiments.

5. CONSEQUENCES FOR FORWARDING

As we have shown above, the human mobility characteristics measured in the conference environment are very similar to those from our previous data sets in corporate and academic environments using both Bluetooth and WiFi. In [1], we describe how many existing forwarding algorithms perform badly in the presence of the power law mobility profile for inter-contact time, particularly for coefficients less than one, which is the case for the traces above. One problem is that much of the evaluation of such algorithms has only been in reference to a mobility model without heavy-tailed inter-

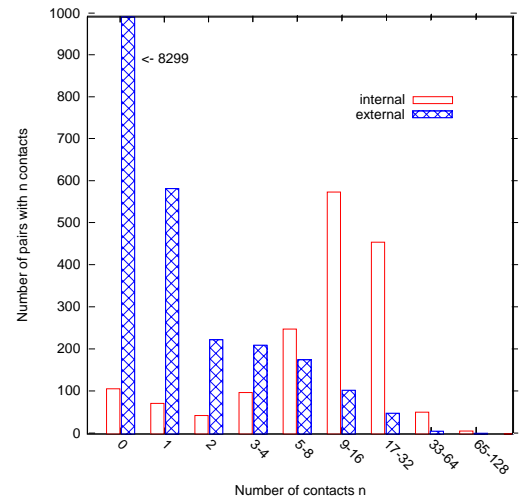



Figure 6: Distributions of the number of sightings of a device (left) and the number of contacts between pairs of devices (right).




Figure 7: Distribution of the number of iMotes met by each device in the experiment.




Figure 8: Distribution of inter-contact times during various parts of the day.

contact characteristics, such as random waypoint or random walk, and with identically-behaving nodes. In this section, we examine the data gathered further in order to determine some of the factors that forwarding algorithms must take into account to cope with human mobility conditions.

We chose not to conduct trace-based evaluation of forwarding algorithms for this paper. This is mainly because, although we have a trace of network conditions, we do not have data on a traffic model for such a network, which is particularly important when, as we show below, the network performance is highly sensitive to the time of day and to the traffic matrix. Such evaluation is the subject of future work.

5.1 Contact with a group of nodes

In addition to looking at the times when two particular nodes are in contact, which was looked at above, it is also useful to examine the periods that nodes (internal and external) spend in contact with at least one of the iMotes. We refer to such periods as “any-contact” times, and the intervals between them as “inter-any-contact” times. Results are shown in Figure 5. We observe that, as expected, any-contact times are longer than contact times, but with the same distribution shape. For inter-any-contact times, we observe that the best fit power law coefficient increases from 0.4 to 0.6 for external nodes, and 0.9 for internal nodes.

One implication of this difference is that, if a node wanted to communicate with a member of a group of other nodes¹, the forwarding possibilities become much better. Since our previous work has shown that different power-law coefficients enable different forwarding algorithms to be useful, this result indicates that applications needing to communicate with one of a group of nodes may be able to use different algorithms to applications which require unicast networking.

The idea that groups of nodes are usable for application-layer traffic is appealing - take for example the groups corresponding to “a member of a particular community of people”, “anyone who has Internet connectivity right now”, or “nodes with information on train timetables”.

5.2 Distribution of contacts among nodes

We now study three measures of the distribution of contacts among the nodes (internal and external), to see whether this distribution is uniform (as many simulation models assume), or whether there is large variability.

Figure 6 provides two views on this data, with the left plot showing the distribution of the number of times a node was seen by an iMote in the experiment (a “sighting”), and the right plot showing the number of times that particular pairs of (iMote, iMote) and (iMote, external) nodes had contact with each other.

The left plot shows that each iMote was sighted between 257 and 946 (the maximum) times during the experiment. This large range is contrary to mobility models in which all nodes have similar mobility profiles. This observation leads to the possibility that Pocket Switched Networking forwarding algorithms could detect and preferentially use highly-mobile nodes. Only one external device, on the other hand, was sighted over 512 times, with the modal number of sightings being 2, and the vast majority of external devices

¹In the IP world, this is known as “anycast”, with one difference being that, in the case we discuss, the group of nodes we are interested in may not be well-known.

having less than 10 sightings. This contrast can be explained by the fact that the iMotes are distributed to humans who form a close-knit community during the period measured, while external addresses could be seen by an iMote who happened to be distant from the conference, and thus not likely to be seen by any other iMote. This leads us to consider algorithms that keep track of user communities (either by automatic detection or by user configuration), and exploit this knowledge by preferring other community nodes for forwarding.

The right plot in Figure 6 shows a large variability in the number of times particular pairs of iMotes saw each other. This is again contrary to many mobility models which treat nodes as equally likely to meet one another, and indicates that forwarding algorithms could usefully keep track of the frequency with which a node sees another, and route towards nodes with a higher frequency for the destination. The plot also highlights the different patterns of visibility with external devices, with the vast majority of potential pairings (8299) never coming into contact.

Figure 7 shows the number of iMotes who saw a particular device (no matter how many times). This plot differs from the others in that it does not count multiple sightings by the same nodes as significant, thus, it shows the inter-connectivity over the whole three days, but not how often this connectivity occurred. This diagram also clearly shows a difference between the community of iMote carriers and most others, with iMote carriers seeing all but a few other iMotes, while external devices were typically only seen by a couple of iMotes. A few external devices, however, are seen by 35 or more of the iMotes; a good automatic community formation system would identify these devices as members of the community.

Unlike with the previous graphs we have presented, the results for external addresses are very different to the results for iMotes only in this illustration of the data. Most of the external addresses are seen fewer than 30 times, and a given external address in our study is extremely unlikely to meet a given iMote in our study (as is the case for 8339 of such potential iMote-external pairs). This contrast can be explained by the fact that the iMotes are distributed to humans who form a close-knit community during the period measured, while external addresses could be seen by an iMote who happened to be distant from the conference, and thus not likely to be seen by any other iMote. This leads us to consider algorithms that keep track of user communities (either by automatic detection or by user configuration), and exploit this knowledge by preferring other community nodes for forwarding.

5.3 Influence of the time of day

The previous section studied how contact patterns differed for different nodes. This section complements the observations made above by looking at the way contacts are distributed over the time of day.

Figure 8 shows the distribution of the inter-contact time during various three-hour periods of the day. This illustrates that there is also a time dependence in the contact distribution, with the most obvious being the diurnal cycle — daytime periods have a greater power law coefficient than night periods. This suggests that forwarding algorithms would benefit from reacting to temporal patterns found in the user’s behaviour. A highly related area is that of the

variation in the user's networking demands over the course of a day, as one might expect that busy contact opportunity periods might coincide well with busy networking periods. This is an ongoing topic of research.

6. CONCLUSIONS

Pocket Switched Networks (PSN) provide mobile networking support for users with mobile devices. By using a combination of local wireless networking and the use of human mobility, they extend network functionality to include scenarios where connection to access infrastructure is not available. We have described PSN in the context of the larger field of Delay-Tolerant Networking (DTN), and presented real-world measurement results for the mobility of attendees in a conference environment. The (anonymised) results are freely available to researchers — visit <http://www.cambridge.intel-research.net/haggle/>.

These results are in agreement with a previous study using workplace and university scenarios in exhibiting a power law for inter-contact time. We then explored characteristics the collected data which can be used to design improved forwarding algorithms, with four main results. Firstly, when forwarding to any one of a group of nodes, the power law coefficient increases significantly. Secondly, we saw that the nodes are not equal — some individuals are much more active, and some pairs see each other more often than others. Third, we showed that the nodes had very different frequencies of connection opportunities with each other than with external nodes. This leads to the notion that identifying shared communities (in this case, the community of conference attendees) can help greatly when forwarding data between two members of such a community. Finally, we observed that forwarding algorithms could the presence of different contact patterns at different times of the day.

In future work, we intend to continue collecting and publishing traces of human mobility measurements in various networking environments. We will continue with the mathematical analysis of such environments, and wish to create mobility models which are representative of human movement patterns. We also intend to design and evaluate forwarding algorithms for PSN, based on the insights from our measurements of human mobility.

7. REFERENCES

- [1] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, and J. Scott. Pocket switched networks: Real-world mobility and its consequences for opportunistic forwarding. Technical Report UCAM-CL-TR-617, University of Cambridge, Computer Laboratory, February 2005.
- [2] Kevin Fall. A delay-tolerant network architecture for challenged internets. In *Proceedings of ACM SIGCOMM*, 2003.
- [3] Scott Burleigh, Adrian Hooke, Leigh Torgerson, Kevin Fall, Vint Cerf, Bob Durst, Keith Scott, and Howard Weiss. Delay-Tolerant Networking: An approach to interplanetary internet. *IEEE Communications Magazine*, July 2003.
- [4] Alan Mainwaring, Joseph Polastre, Robert Szewczyk, David Culler, and John Anderson. Wireless sensor networks for habitat monitoring. In *ACM International Workshop on Wireless Sensor Networks and Applications (WSNA '02)*, Atlanta, GA, September 2002.
- [5] Alex (Sandy) Pentland, Richard Fletcher, and Amir Hasson. DakNet: Rethinking connectivity in developing nations. *IEEE Computer*, 37, January 2004.
- [6] S. Jain, K. Fall, and R. Patra. Routing in a Delay Tolerant Network. In *Proceedings of ACM SIGCOMM*, 2004.
- [7] P. Juang, H. Oki, Y. Wang, M. Martonosi, L-S Peh, and D. Rubenstein. Energy-efficient computing for wildlife tracking: Design tradeoffs and early experiences with zebranet. In *Proceedings of the 10th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, pages 96–107, 2002.
- [8] T. Small and Z. J. Haas. The shared wireless infostation model - a new ad hoc networking paradigm (or where there is a whale, there is a way). In *Proceedings of the 4th ACM International Symposium on Mobile Ad Hoc Networking and Computing*, pages 233–244, 2003.
- [9] R. C. Shah, S. Roy, S. Jain, and W. Brunette. Datamules: Modelling a three tiered architecture for sparse sensor networks. In *Proceedings of the First IEEE International Workshop on Sensor Network Protocols and Applications (SNPA)*, 2003.
- [10] James Sterbenz, Rajesh Krishnan, Regina Rosales Hain, Alden Jackson, David Levin, Ram Ramanathan, and John Zao. Survivable mobile wireless networks: Issues, challenges, and research directions. In *Proceedings of the ACM Wireless Security Workshop (WiSE)*, 2002.