Mobility and Service Discovery in Opportunistic Networks

Mikko Pitkänen, Teemu Kärkkäinen, Jörg Ott

Aalto University, Department of Communications and Networking (Comnet)

Abstract—Past research on opportunistic communications has demonstrated the power of mobile phones as information mediators and content providers to co-located devices. As an extension of opportunistic communication paradigm, opportunistic computing and services research leverages the powerful capabilities and pervasive nature of modern smartphones to provision services to nearby nodes. When an opportunistic network is viewed as a provider of services rather than as a simple message routing system, node and service discovery come to play a critical role. This paper discusses several factors that affect the discoverability of pervasive services with current existing devices in light of their practical limitations. Moreover, the impact of these limitations to pervasive service discovery is studied in evaluation results.

I. INTRODUCTION

Opportunistic networking and direct peer cooperation show great promise for enabling entirely new types of applications and new ways of using old applications. These applications can rely purely on direct communication between devices [9] while also exploiting static infrastructure components, such as network access points, when available [10] [4]. The central idea in many of the proposed applications is to substitute a fixed routing infrastructure by an opportunistic message routing medium built on, for example, pervasive mobile devices. Such a routing medium focuses on moving data from sources to destinations (e.g., from sensor nodes to database servers) or on providing access to certain data (in content centric networking). We can, however, widen our view and consider the opportunistic network, and the devices within it, as a potential provider of complex services. These services can be composed from the resources available within the variety of components that make up the opportunistic network [2].

Producing services as a part of an opportunistic network places requirements on the nodes that differ from those found in opportunistic message forwarding and information dissemination networks. In the case of pervasive services, the nodes may invoke the services on devices which they are connected to for a brief time and never encounter again. Therefore, it becomes important to perform service discovery efficiently so that the chances of getting service are maximized. This paper discusses service discovery among co-located nodes and discusses the mechanisms that are often seen in practical implementations and evaluates their impact on service discoverability in realistic opportunistic scenarios.

In this paper we focus on the two main communication technologies that are ubiquitous in the devices that make up opportunistic networks: Bluetooth and WiFi. While both technologies provide the same basic capability – transmitting and receiving data between devices – they both have their own sets of practical limitations. These limitations stem largely from the radio and protocol layer power-saving optimizations that extend the battery life, while other limitations are a result of the original design assumptions and scenarios.

In general we can identify two device-to-device communication paradigms: *direct*, and *infrastructure-assisted*. Direct communication technologies allow nodes to discover and communicate with each other simply by being within radio range. These include Bluetooth, WiFi ad-hoc mode, WiFi-Direct and WiFi-Opp. On the other hand, many mobile devices are used primarily for connecting to an infrastructure network through, for example, WiFi access points. This infrastructure may enable nodes to communicate directly with others within the same local network segment, providing opportunities for discovery and communication between co-located devices.

This paper discusses practical limitations seen in real world mobile devices and access points. An evaluation in an emulated environment sheds light on how these limitations interplay with user mobility and service discovery in pervasive scenarios. Observations cover how limiting the activation of the device radio interfaces and power saving mechanisms both in mobile nodes and access points increase the node discovery time, and therefore decrease the potential service time in access point-assisted communications. Observations demonstrate the impact of radio communication range on the discovery and service times with different node mobility characteristics – both in synthetic and real mobility traces. Moreover, the needs of service discovery process add steps to resolving network and security information during connection establishment. Statistical properties of the contact durations provide information on how this affects the service discovery process.

II. PERVASIVE SERVICE DISCOVERY

The following discusses pervasive service discovery where mobile users move in an urban area and create a network between the co-located devices. Quite often in the forward looking research papers an assumption is that the devices can communicate with co-located devices by means of close range radio such as WiFi in ad hoc mode. However,

as practice has shown such radio service is not often easily available in commercial mobile devices, which come without the access rights for the user to modify the radio firmware or the related software configurations. Therefore, we look into how the service discovery can be executed when we take some important practical limitations into account. Specifically, we account for the limits in smartphone Bluetooth communication stacks as well as in WiFi access points and evaluate their effect on the discovery process.

Our discussion centers around the following concepts. By *device discovery* we denote the process of two devices finding out that they are in the radio communication range with each other. *Service discovery* refers to the process that begins after the device discovery and completes when the devices have successfully exchanged service information. In this context, the service information includes at least the protocol description to connect to the service, the service descriptions, and the necessary operations for security.

There are two basic scenarios for service discovery: 1) co-located nodes within each other's radio range form an opportunistic contact and no third parties are required, this can be achieved by using *direct* device-to-device communication technologies. 2) *infrastructure-assisted* nodes co-located within the range of the same access point and they use the access point to mediate the discovery messages. The following summarizes the main differences between the two considering our scenarios of interest.

The most widely deployed *direct* communication technology, Bluetooth, was designed for interconnecting personal devices, such as a headset and a mobile phone. Importantly, Bluetooth was not designed for building opportunistic networks. This is reflected by characteristics such as shorter radio range, secure pairing, and lack of broadcast support in smartphone Bluetooth stacks. The lower radio coverage as well as lack of support in smartphone Bluetooth stacks for broadcasting and simultaneous connections limits their use in opportunistic networking. Due to these limitations, the minimal requirements for pervasive service discovery with Bluetooth include opening socket communications after the device discovery to negotiate the service parameters.

Beyond Bluetooth, there are many WiFi based solutions for direct connectivity, including WiFi ad-hoc mode, WiFi-Direct, and WiFi-Opp [12]. WiFi ad-hoc mode is an alternative mode to the normal infrastructure mode that enables devices to communicate directly using the WiFi radio. However, it is not widely used due to poor support by operating systems, inability to use it in parallel with infrastructure mode, and lack of support for address allocation. The new arrival, WiFi-Direct, overcomes several of the above limitations by reusing the widely spread WiFi infrastructure mode hardware in combination with new peer-to-peer features that allow nodes to dynamically act as access points. However, WiFi-Direct shares its design decisions and use cases with Bluetooth, and has not been designed with opportunistic networking in mind. WiFi-Opp is a similar solution to WiFi-Direct, but it has been designed particularly for opportunistic networking rather than for interconnecting personal devices such as phones and printers. All WiFi based solutions in mobile devices are limited by power saving features, that call for the radio to be inactive as much as possible.

Access points provide an interesting *infrastructure-assisted* alternative as enabler for communication between mobile nodes. They tend to have more capabilities, such as support for address allocation, and unlimited power supply, which allows parts of the discovery responsibilities to be shifted away from the mobile nodes. However, the access points have implications to the aggregate node mobility in two major ways. First, access points have fixed locations, which makes their placement critical in how successfully they will be able to serve the mobile nodes. Second, using access points as mediators implies that in order to communicate, two nodes must be within the same hotspot coverage area – it is not sufficient for the nodes to simply be co-located at any given point in time. On the other hand, large local area network broadcast domains spanning multiple hotspots might mean that two nodes may communicate even when not co-located.

The access point-assisted service discovery provides a way for nodes to use the powerful WiFi radio. Wide availability of WiFi hotspots, coupled with devices configured to actively search and connect to them results in frequent opportunities to discover and interact with other nodes. From the point of view of discovery mechanisms, WiFi has desirable broadcasting properties including ability to multicast datagrams to multiple nearby receivers. For example, simple service advertisements can reach all nodes connected to an access point without the requirement of pair-wise connection establishment of Bluetooth. However, our practical experience with hotspots have shown that the energy saving mode where packets are buffered in the access point and sent out in bursts (to allow client radios to sleep for predetermined periods) often prevents immediate propagation of packets, introducing significant delays (tens of seconds) in service discovery. Moreover, the delays in dynamic host configuration protocol (DHCP) may add significantly to service discovery time when first entering a hotspot, however, optimizations are possible and can significantly speed up the process.[1] Therefore, we will discuss in Section V how these limitations affect the access point-assisted service discovery.

III. RELATED WORK

For energy aware neighbor discovery there are numerous schemes available [7], [5], [1] to schedule the discovery beaconing (or node sleeping) independently in nodes so that simultaneous wake schedules for the nodes can be

---

[1]http://cafbit.com/entry/rapid_dhcp_or_how_do

calculated. The calculated schedule then guarantees that the node discovery happens within a given maximum time if the nodes are in the communication range. Wang et al. [14] have proposed an adaptive probing mechanisms that adjusts to the arrival process of the contacts during the normal node operation and saves significant amount of energy compared to constant probing. Practical experiments with opportunistic networking using Bluetooth [8] report discovery issues leading to missed connections and suggest that with parallel connections interference reduces transfer capacity harmfully.

The Jyotish framework [13] is a proposal where combination of WiFi and Bluetooth scan metrics are used to predict the movement of the mobile nodes and the social mobility patterns. Such prediction about the movement and stay schedules can provide input to adaptive scanning mechanisms. SociableSense [11] framework presents social aware adaptive scanning framework for collecting general sensory information beyond wireless probing. Proactive scanning [15] for WiFi access points can be implemented as a client-only software making it a possible mechanisms for the access point-assisted opportunistic communication.

IV. FACTORS IN SERVICE DISCOVERY

This section discusses key factors that affect effectiveness of the service discovery in opportunistic scenarios. The process is observed from the service user's standpoint and mechanisms for both Bluetooth and access point-assisted service discovery are discussed. The discussion draws from our practical experience on working with opportunistic communications and computing software.

A. From Node Discovery to Service Discovery

The service discovery usually starts after the nodes have come into physical communication range and negotiated a radio contact through which communication can be established. This can happen either directly between the nodes or as assisted by an AP. The purpose of the service discovery is to find knowledge about existing service and their contact information (protocols and ports) as well as to secure the nodes and the access to the services. This process includes steps from assigning network addresses by means of dynamic host configuration or establishing Bluetooth channels to setting up convergence layers to establish overlay networks, for example TCP convergence layer is used in DTNs [3].

The service discovery is significantly different in broadcast and multicast media where the common practice is to query and announce the services to all nodes in proximity. This approach is used for example in Multicast DNS (or Bonjour) service discovery mechanisms and it allows to use broadcast datagrams to distribute service information to all nodes connected to an access point. In this case, all nodes that have a registered interest for a service, but may not yet be known to the provider, are able to discover the service by listening to multicast messages without requiring the two-way communication establishment.

The common smartphone stacks for Bluetooth do not, however, support broadcast or multicast messaging. Therefore, in order to exchange information about opportunistic services (beyond Bluetooth SDP that advertises simple connectivity services such as the type of device), the Bluetooth devices need to first establish a mutual communication context in order to exchange information about the services. This context is formed by establishing a socket to another device. If the device is not part of the opportunistic network, the resources used for scanning and socket formation procedures are wasted. This might suggests blacklisting the devices to prevent future wasting of resources. However, from the client's standpoint it may be difficult to differentiate whether the device was not an opportunistic service provider, whether it was simply busy serving another client, whether it just had its service stopped for a moment, or whether the connection failed for another reason. In case of a successfully established socket communication, the process continues by first exchanging the service and security information with a custom protocol. After, a convergence layer can be established to allow the communication in an overlay network. In such a case, the peer device might be "whitelisted" as a peer known to work.

When the basic services provided by a node are likely to stay unchanged for the node over a period of time, a straightforward optimization follows from caching the discovered service information for nodes in the community.

B. Application Driven vs. Network Driven

Application developers often use software development kits (SDKs) to program for the smartphones. The SDKs abstract many of the lower layer operation system events from the developers. When programming for such devices, the application developers may decide to periodically start to scan for the radio devices based on changes in, for example, time or location. However, many platforms provide asynchronous notification mechanisms, which provide notifications immediately after the network state changes, for example, when radio interface notifies the operating system about an event. Relying on these notifications can effectively speed up the discovery and make it an adaptive process from the application point of view. This effect is seen in the simulations with different scanning intervals in Section V.

Although designing the applications to leverage the system level notifications can be beneficial, their use is limited in access point-assisted communications. The node that is connected to an access point does not automatically get notifications about the other joining nodes from the operating system (listening to ARP broadcast traffic could provide this information, but this is not commonly available in smartphones today). Therefore, the applications have to

be designed so that they actively probe for the devices periodically, and when joining to a new network.

C. Node Discovery and Disassociation

Experience from the practical implementation work has shown that the access point-assisted service discovery mechanisms suffer from the energy saving mode, where the access point may buffer datagram packets for an extended period of time. This prevents immediate packet transmissions between the mobile devices, and can significantly delay service discovery when using discovery protocols that rely on multicast such as Bonjour. In an example case, all UDP packets that were send over a period of few beacon intervals got buffered and were flushed to multicast destinations after a longer period had elapsed. This period was often much longer than the node mobility in opportunistic scenarios would have suggested. This has again effects similar to increasing scanning interval.

The delay due to packet buffering in access points has an even more severe effect on detecting the contact cutoff. Any discovery protocol that works on top of the unreliable UDP messages needs to account for number of packet losses and should decide that node has disappeared only after number of beacons was not heard. This can be particularly harmful, for example, if the node would like to terminate the current connection to start seeking service from another node that it has discovered to be in within the communication distance. The effect of this can be estimated from the scanning interval parametrization and the contact duration distributions that are shown in the evaluation results.

D. Node Movement

Taking into consideration the typical variable walking speed of humans, the range of possible values for velocity are limited. The evaluation traces of this paper use velocities drawn from real mobility traces recorded in KAIST campus or draw velocities from random distributions for synthetic mobility models.

In addition to walking nodes, we place several stationary nodes into our scenarios, which model wireless access points and have the obvious zero velocity. The use of the hotspots is interesting from mobility standpoint as they allow to measure the contact patterns when two mobile nodes have to meet simultaneously under a constrained location (within radio range to same access point).

V. Evaluation

The following simulation results are obtained with real-world and synthetic mobility models in ONE simulator [6]. We run the simulations without traffic in the network and focus on the effect of scanning frequency, radio range and node mobility to discovery process. The nodes initiate a scan when a fixed scan interval has elapsed since the last scan. The scan completes immediately and reports all nodes that are within the radio range of the device. The location of nodes is updated once every second. The assumption about instant discovery is unrealistic, but the recorded contact duration distribution allows to reason about the effect of the discovery process to contact durations, that is, potential service times. Radio range is fixed to all directions and all nodes have uniform range within each simulation.

The KAIST campus trace follows real world movement of 92 pedestrians at 30 second intervals. The ONE simulator interpolates linear locations for the pedestrians between the recorded locations so that node speed is according distance moved during the interval. Toevaluate the access point-assisted connectivity, we place WLAN APs at central locations on the campus area with 100 meters distance in a grid topology.

For the synthetic mobility scenario we choose the Helsinki City Scenario (HCS) where 166 pedestrians move around city streets of the downtown area. In HCS the access points are placed around the city [10] based on locations of existing hotspots provided by a community WiFi service. To evaluate the impact of pedestrian speed, we run all the simulations with slow walking speed drawing speed for each leg is from uniform distribution 0.5-1.5m/s, and fast walking speed 1.5-2.5m/s respectively. KAIST mobility follows recorded locations and provides no meaningful to vary the speed.

We plot two graph types for each simulation: The first graph is simple bar graph that plots the number of contacts within the simulation. The access point-assisted contacts are reported for duration two nodes spent under the same access point. This provides absolute number of contacts, which is scaled to highest contact count in each scenario.

The second graph plots the cumulative distribution function of contact durations. This allows to observe the impact of the node and service discovery process to contact duration or potential service time. Estimated service discovery time in x-axis provides the share of the contacts that last less than the discovery time and thus reduce to unusable for services.

A. Reference: Mobile Ad-hoc Discovery

We first turn our attention to direct contacts between two mobile devices using ad-hoc networking as usually assumed for mobile opportunistic communication. This primarily serves as a reference for investigating subsequently the impact of practical limitations on service discovery performance, but we can make some initial observation concerning service discovery.

Figure 1 shows the total number of contacts and the contact durations for the above three scenarios. On the left, we can see the that—expectedly—the number of contacts grows with the radio range and reduces with increasing scanning intervals. On the right, the figure shows three graphs detailing the fraction of contacts that last for a certain duration, as this is the determining factor for successful service discovery and interaction: if service discovery *after*

encountering a contact would take another 15 s one may easily lose 20-50% of the contacts in the KAIST trace and even more for short radio range in HCS. The loss becomes naturally more pronounced with growing scanning intervals because a fraction of the potential contact duration will already have elapsed when two nodes encounter each other.

B. Limited Bluetooth Discovery

When we apply the limitation of only a single contact at a time to the above simulations. We find that the contact time *distribution* does not change noticeably compared to the one shown in figure 1. However, the *magnitude* of the contacts decreases significantly. Figure 2 shows the numbers of contacts (next to the Y axis) and their relative distribution for different radio ranges and scanning intervals. Comparing the absolute numbers we see that limiting simultaneous Bluetooth interactions to at most one peer cuts the number of contacts roughly in half (both HCS) or reduces them to about one fifth (KAIST).

Moreover, the Bluetooth-based discovery simulations show first an increase (from 10 to 30 m radio range) and then a decrease in the number of contacts when moving to 60 m, the latter of which appears counterintuitive to our findings above. However, the slowly increasing CDF in figure 2 shows that a large fraction of contacts last longer with 60 m radio ranges. But since we allow only one simultaneous contact for Bluetooth devices, the longer-lasting connections occupy the two involved nodes and thus prevent them from discovering further nodes. This is seen as decrease in number of connections with large 60 meter radius. This observation suggests that it may be advisable for nodes to cut the connections short rather than sticking to them until they break, so that they can establish new ones: for example, if heterogeneity of service providers and diversity of services is more important than long service connections.

C. Access Point-Assisted Discovery

Figure 3 shows the performance of access point-assisted discovery. On the left, we see the strong impact of radio range on the number of contacts. While access points are capable of serving many nodes and nodes may discover multiple peers at a time, the fact that the WLAN access points are stationary and two nodes have to be at the same time in one of the right places influences the contact pattern even stronger. Using only (unrealistically low) 10 m radio ranges yield virtually no contacts at all, but the number of contacts grows fast with increasing radio range. Nevertheless, the total number of contacts remains well below the reference scenario and even 10–20% below the limited Bluetooth setup.

The distribution of contact durations yields marginally steeper slopes for KAIST, especially leading to short contacts, and significantly steeper ones for the artificial HCS mobility models. The increase of short duration contacts

poses challenges to service discovery because the fraction of usable contacts quickly decreases with increasing discovery time. For example, with 15 s discovery time, 30–50% of the contacts are lost. As noted above, increasing the scanning interval reduces the number of contacts but does not change the distribution of contacts.

We finally note two further practical aspects concerning WLAN APs: 1) The most active areas of node movement and encounters may not happen under the coverage of the hotspot as the hotspots locations are not chosen to suit walking users but rather to accommodate stationary ones (often indoors) and their (outdoor) coverage that could assist peerwise interactions, to date, would be more coincidence than planned. 2) Access points may not support multicast between mobile nodes (without prior authentication or at all), limiting the opportunities to run distributed service discovery protocols in the first place. In both cases, when WLAN hotspot operators become aware that their APs' could possibly assist mobile users in opportunistic communication, this might lead to supportive extensions (better range, more APs) or more restrictive AP configurations.

VI. DISCUSSION AND CONCLUSION

In this paper, we have reviewed some practical limitations of peer discovery. Determining the presence of another node will be insufficient since we cannot assume homogeneous node populations. Higher layer service discovery is needed to determine which services (if any) two peers have in common, which takes time to complete and shortens contacts. Since not all peers may run a discovery service, declaring failure efficiently is important—and thus essential for getting the tradeoff right between accidentally missing capable nodes and waiting too long for incapable ones to answer (missing other opportunities in the meantime).

The two dominant networking technologies found in today's mobile devices, WLAN and Bluetooth, impose further limitations. Our observations complement the experimental findings for three conference scenarios [8] and provide insights how typical simulation-based setups tend to overestimate the number of *usable* contacts, irrespective of the scanning interval for which numerous approaches for adaptation were put forward (e.g., [14], [7]).

With a limited Bluetooth stack, we "lose" contacts because, once a node has discovered a set of peers, it needs to pick one peer to connect to and will then remain unavailable for other while talking to the chosen one. This opens up another dimension of research that has received little attention in opportunistic networking: cutting contacts deliberately short. While protocol studies generally assume that two nodes will exchange all the information they possibly can while the contact lasts, a different notion of prioritization and preemption appears necessary: routing protocols should not just determine the order in which to replicate messages but also when to stop and turn to a potential next peer, without
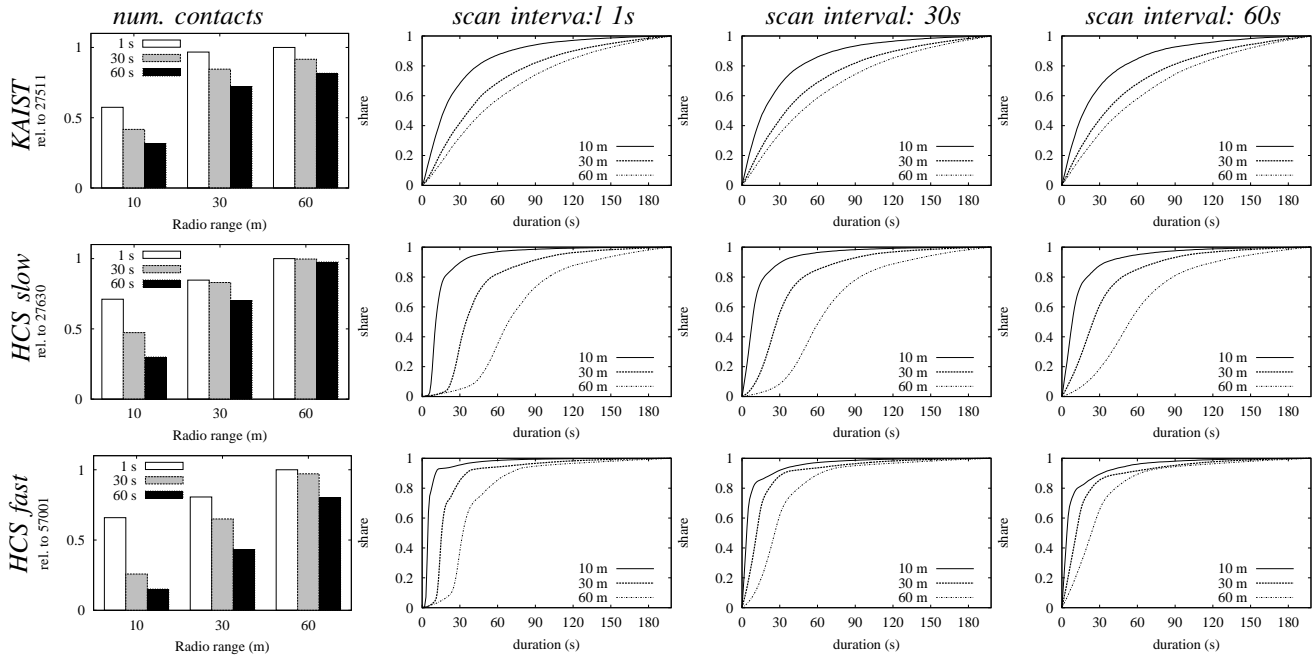
Figure 1. Radio range limited discovery with mobile nodes only: number of contacts (the leftmost column) with different radio ranges and scan intervals. Distribution (three columns on the right) of contact durations with different radio ranges and scan intervals for each scenario.
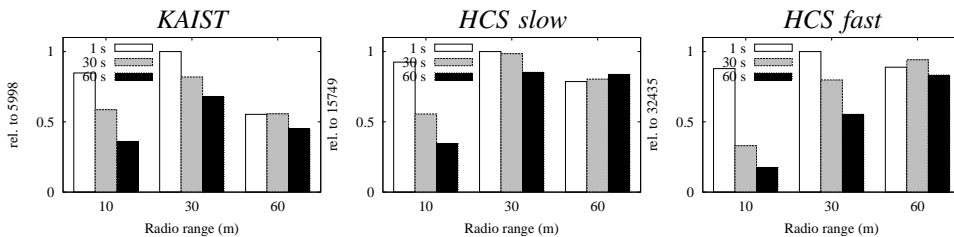

Figure 2. Bluetooth discovery with mobile nodes only: number of contacts with different radio ranges and scan intervals.

knowing if such a peer is (still) around, thus adding a similar dimension as for adaptive scanning.

With WLAN, in spite of the diverse approaches currently explored for mimicking hotspots in mobile devices, physical hotspots can provide a supplement for ad-hoc interactions between mobile nodes—as long as service discovery mechanisms function within their realm. While they clearly limit the places where mobile devices can meet and communicate (which would restrict opportunistic communication to rather urban environments), they solve an important issue of efficient address allocation and (can) preserve communication among multiple peers at once, both essential for effective service discovery. Investigating how well present commercial hotspots fulfill such functions is subject of our ongoing research. Given the widespread availability of WLANs in urban areas, a really nice feature for access points would be offering a complementary unprotected "virtual" WLAN without Internet access for the sole use of opportunistic

networking. With no Internet access required, one could even go further to deploy micro-APs that just have a DHCP server and a singe WLAN interface to facilitate opportunistic communication wherever power is available.

REFERENCES

[1] B. J. Choi and X. Shen. Adaptive asynchronous sleep scheduling protocols for delay tolerant networks. *Mobile Computing, IEEE Transactions on*, 10(9):1283 –1296, sept. 2011.

[2] M. Conti, S. Giordano, M. May, and A. Passarella. From opportunistic networks to opportunistic computing. *Communications Magazine, IEEE*, 48(9):126 –139, sept. 2010.
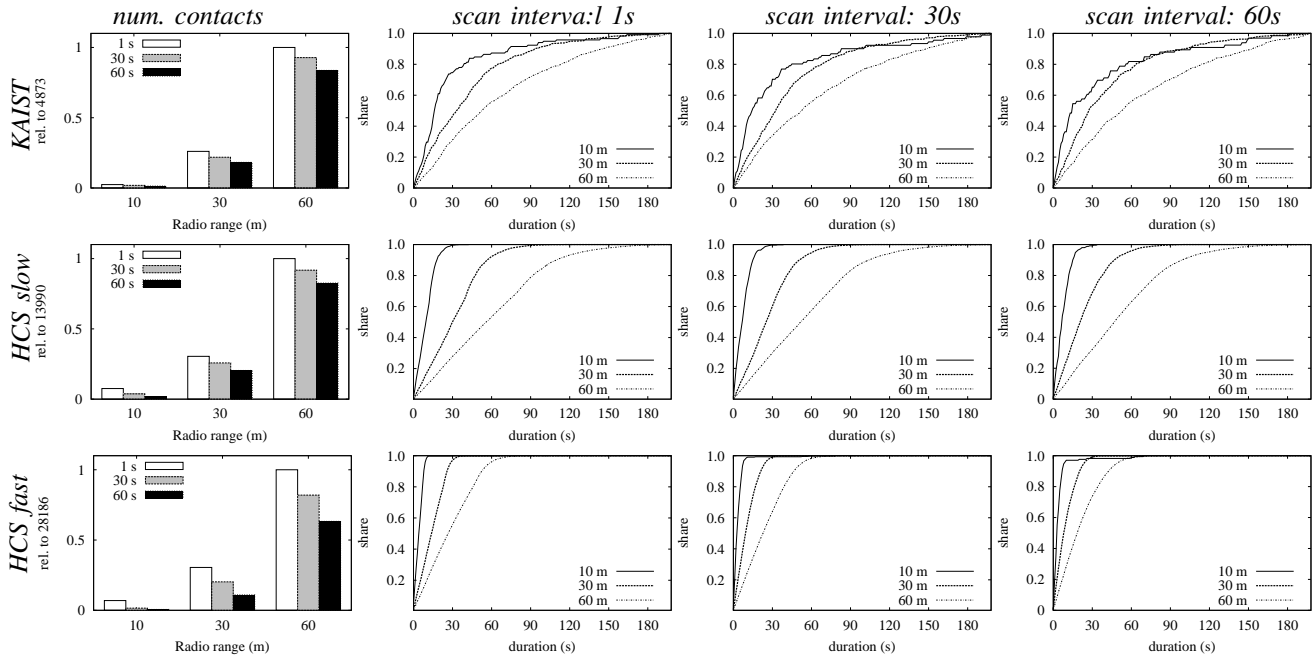
Figure 3. Access point-assisted discovery, number of contacts (the leftmost column) with different radio ranges and scan intervals. Distribution (three columns on the right) of contact durations with different radio ranges and scan intervals for each scenario.

[3] M. Demmer and J. Ott. Delay Tolerant Networking TCP Convergence Layer Protocol. Internet Draft draft-irtf-dtnrg-tcp-clayer-01.txt, Work in Progress, February 2007.

[4] S. Dimatteo, P. Hui, B. Han, and V. O. Li. Cellular Traffic Offloading through WiFi Networks. In *Proc. IEEE International Conference on Mobile Ad-Hoc and Sensor Systems (MASS)*, 2011.

[5] A. Kandhalu, K. Lakshmanan, and R. R. Rajkumar. U-connect: a low-latency energy-efficient asynchronous neighbor discovery protocol. In *Proceedings of the 9th ACM/IEEE International Conference on Information Processing in Sensor Networks*, IPSN '10, pages 350–361, 2010.

[6] A. Keränen, J. Ott, and T. Kärkkäinen. The ONE Simulator for DTN Protocol Evaluation. In *Proc. of the 2nd International Conference on Simulation Tools and Techniques (SIMUtools)*, 2009.

[7] J. A. B. Link, C. Wollgarten, S. Schupp, and K. Wehrl. Perfect difference sets for neighbor discovery: Energy efficient and fair. In *Proceedings of the 3rd Extreme Conference on Communication (ExtremeCom)*, 2011.

[8] A.-K. Pietiläinen and C. Diot. Experimenting with opportunistic networking. In *Proc. ACM MobiArch*, 2009.

[9] M. Pitkänen, T. Kärkkäinen, J. Greifenberg, and J. Ott. Searching for Content in Mobile DTNs. In *Proc. IEEE PerCom*, 2009.

[10] M. Pitkanen, T. Karkkainen, and J. Ott. Opportunistic Web access via WLAN Hotspots. In *Proc. IEEE PerCom*, 2010.

[11] K. K. Rachuri, C. Mascolo, M. Musolesi, and P. J. Rentfrow. SociableSense: Exploring the trade-offs of adaptive sampling and computation offloading for social sensing. In *Proc. ACM MobiCom*, 2011.

[12] S. Trifunovic, B. Distl, D. Schatzmann, and F. Legendre. WiFi-Opp: ad-hoc-less opportunistic networking. In *Proc. of the ACM CHANTS workshop*, 2011.

[13] L. Vu, Q. Do, and K. Nahrstedt. Jyotish: A novel framework for constructing predictive model of people movement from joint wifi/bluetooth trace. In *Proc. IEEE PerCom*, 2011.

[14] W. Wang, V. Srinivasan, and M. Motani. Adaptive Contact Probing Mechanisms for Delay Tolerant Applications. In *Proc. ACM MobiCom*, 2007.

[15] H. Wu, K. Tan, Y. Zhang, and Q. Zhang. Proactive Scan: Fast Handoff with Smart Triggers for 802.11 Wireless LAN. In *Proc. of IEEE INFOCOM*, 2007.