Active Collaborations for Trustworthy Data Management in Ad Hoc Networks

Anand Patwardhan, Filip Perich, Anupam Joshi, Tim Finin and Yelena Yesha
Department of Computer Science and Electrical Engineering, UMBC, Baltimore, MD 21250
{anand2, fperic1, joshi, finin, yeyesha}@cs.umbc.edu

***Abstract*— We propose a trust-based data management framework for enabling individual devices to harness the potential power of distributed computation, storage, and sensory resources available in pervasive computing environments. Available resources include those currently present in the fixed surrounding infrastructure as well as those resources made available by other mobile devices in the vicinity. We take a holistic approach that considers trust, security, and privacy issues of data management in these environments. We focus on collaborative mechanisms to provide a platform for trustworthy data management for devices in ad hoc networks. A fundamental aspect of our framework is a pack formation mechanism for enabling collaborative peer interaction in the pervasive computing environments based on context information and landmarks. A pack provides a routing substrate for enabling devices to find reliable sources of information. A pack also provides a platform for coordinated pro-active and reactive mechanisms that can detect and respond to malicious activity. Consequently, a pack can be used for providing a foundation to distributed trust management and data intensive interactions. We describe our proposed data management framework with an emphasis on forming packs in mobile ad-hoc networks and present preliminary results from our simulation of collaborative data management using packs.**

I. Introduction

Advances in technology and growing demand for wireless connectivity are fueling a proliferation of wireless capabilities in everyday appliances. Popularity of personal mobile devices, like Apple's iPod [6], have surpassed cult standards and have evolved into cultures in their own right, where unforeseen applications like pod-casting [4] are emerging. Another example is the recently released Sony Playstation Portable(PSP) [7] intended as a portable gaming console that comes equipped with 802.11b [2], [5] InfraRed capabilities. This trend of continuous improvements in capabilities of small form factor devices and the consequent evolving applications are harbingers of a highly connected networked society that integrates into data-intensive and resource-rich environments.

Network connectivity in pervasive environments is typically local. Devices communicate via ad hoc networks. Infrastructure based continuous connectivity to the Internet is either unavailable or expensive. As such, data and information is provided and consumed locally using peer interactions with other devices and resources in the vicinity. Existing data management architectures [12], [21] usually assume that all the encountered devices are trustworthy, and so is the information obtained from them. This assumption, however, is not entirely realistic in a pervasive environment [17]. As such, assurances of quality and accuracy of the retrieved information need to be provided. Due to limited Internet connectivity it is unfeasible to resort to conventional security paradigms for determining trustworthiness of other devices available for interaction. For example, communicating with a trusted authority in order to verify the credentials of other entities encountered in the vicinity may be impossible. Moreover, the nature of these environments requires these devices to be able to function independently and be capable of making autonomous decisions about trustworthiness of peers and accuracy of information received. Metrics for evaluating the reliability of data and trustworthiness of peers are the foundations for achieving secure and trustworthy data management networks. With the lack of centralized authorities in ad hoc networks, trust evaluation and reputation management are useful mechanisms that will allow devices to function autonomously with minimal user intervention. To achieve these goals it is necessary to have a holistic approach for addressing issues related to device security, secure routing, peer discovery, data management, and

This research was supported by NSF awards IIS 0209001 and CNS 0203958

reputation management.

In our previous work [20], [21] we have shown how profiles encoding Beliefs, Desires and Intentions (BDI) of the user, can be used by mobile devices to forage for information in pervasive environments. The user and device profiles are used for pro-active semantic data caching [22]. Also, in our MoGATU framework [19] we have shown how reputations can be employed to ascertain accuracy of information exchanged between mobile devices in pervasive computing environments. However, significant effort is required for evolving trust and maintaining reputations in terms of computation, storage, and messages interchanged.

We define a pack as a dynamic grouping of individual devices which agree to collaborate for executing individual and collective goals. In this paper, we propose a collaborative approach that addresses issues of trust and security in data management and present initial results from our simulations on pack formations. Our approach is based on providing primitives for trust evolution and pack formation. A pack formation will serve two purposes, namely, (1) a pack can provide a routing substrate for enabling devices to find reliable sources of information, and (2) a pack can provide a platform for coordinated pro-active and reactive mechanisms that can detect and respond to malicious activity. A pack, therefore, forms the basis for mobile devices to evolve, manage, and evaluate trust of their peers and information and services they offer.

In the following sections we first present a scenario that illustrates typical pervasive environments, which we use to identify the important constituents of pervasive environments and their characteristics. We discuss issues regarding to reputation management, pervasive trust, and data management. We then present our proposed trustworthy data management framework that employs pack formation to provide a data routing substrate. Next, we present a simulation environment and empirical results from our simulations that demonstrate the utility of collaborative searches using packs. We then conclude with ideas for future work.

II. Background and Related Work

Pervasive environments are in a constant state of flux. Finding useful services and data is serendipitous since continuous connectivity to reliable information sources cannot be guaranteed. Ability to form ad hoc networks depends on the participants' willingness to collaborate in relaying information on behalf of each other. Data reliability is also an important issue since prior trust relationships may not always exist. Moreover, finding reliable data on-demand in a serendipitous environment is another challenge. Past research has provided insights into how a profile-driven agent can use notions of trust and reputations for query processing. However, several challenges remain, e.g., how to find reliable information, use active collaborations, leverage abundant storage – to collate and coordinate data search, and how to improve latency of simplistic discovery mechanisms.

Continuing advances in compact storage technologies like semiconductor memory in CF, SD, MMC cards etc., and miniature harddisks and microdrives have led to mobile devices with substantial storage capacities. Abundant on-board storage will enable mobile devices to store significant amounts of information, to relieve the burden of requesting data from remote servers (thereby relying on connectivity). To be able to guide themselves, the devices will need to sense their contexts (e.g., location, intent of user). Reading local information from reliable sources, the devices can compose locally available services and use their knowledge bases to service their needs and be largely self-reliant. Moreover, all such devices will be capable of providing useful services to other (mobile) devices in their vicinity. The collective data storage capacities and unique sensory and effector capabilities of mobile devices will further enrich the pervasive environment. Long range wireless services are often not suitable for high data rates and at times are not cost-effective. We propose to harness the immense storage capacity of the mobile devices, optimize use of available connectivity to merely keep the knowledge base updated, and enable devices to function autonomously.

Using information from the surrounding pervasive environment introduces several trust and security issues. Due to the inherent nature of pervasive environments, conventional mechanisms of providing security are not suitable. Devices must be made self-reliant to make trust evaluations and use reputations to guide their behavior.

Since mobile devices are potentially innumerable, it is, however, not possible to be able to cache all possible identities, nor is it likely that all encountered devices will be cooperative. We propose that it will be sufficient to remember only devices that are of future potential value in forming social networks and those that will be most likely to cooperate.

Franklin [16] introduces the concept of "Data recharging," likening it to recharging a battery. Data recharging is the process of caching information relevant to the user's needs expressed in the user's profile. The profile is thus a form of longterm query that continuously processes relevant data whenever it is available. Mobile devices can then continue to function even with lack of connectivity to data source.

Past research by Cherniak *et al.* [11], [12] has shown how profiles can be used by mobile devices for client-

server based data recharging.

In pervasive environments, the data sources are varied and are usually fixed with respect to their location in the environment – providing streams of data, e.g., interfaces to sensor networks, traffic conditions, weather conditions, or gas prices. Such data is being created continuously and being disseminated in the neighborhood. Franklin [16] proposes *Adaptive Dataflow Query Processing* to address collecting data from such data streams. He proposes that query processing in such environments requires non-blocking queries and presenting immediate results even if only partial data is available. He notes that conventional database query optimization that static query plans computed from statistics on data and cost metrics, are not suitable for queries on streaming data, given the highly unpredictable nature of availability of sources and data streams.

Connectivity will be provided by available forms of network support (ad hoc networks, or other infrastructure based networks). Application requirements like data rates and acceptable costs will determine which form of connectivity amongst the available possibilities is most suitable.

When finding a resource of interest, the key is in finding the path of minimum cost and maximum benefit. Practical solutions can merely indicate the best path, though not guarantee it. Once data sources or other resources are located in the vicinity of the device (lack of prior arrangement or knowledge) the devices must then later have the option to either continuously monitor/use the resource (identified by its distinguishing source identity) or get updates of information of interest. This requires a close and continuous (minimal) interaction with surrounding nodes who provide assurance of trust about resources and volunteer or collaborate in availing the service. Since these devices are not continuously connected, such communication must be assumed to be intermittent at best. Devices collaborating to provide or extend the range of the services (by providing connectivity to the resource or intermediate storage) in use are only required to prove that they are providing data fidelity in their retransmission but not necessarily provide the trustworthiness of the data source itself (e.g. can relay whatever is transmitted without change, yet not be able to assess the quality of the data).

Such collaborations are similar to those in Peer-to-Peer (P2P) networks, in the sense that a data routing substrate enables resource pooling and scalability. In existing P2P deployments, connectivity is predominantly wireline and thus the overall network topologies are fairly static, which make it possible to efficiently maintain and update distributed hashtables [1], [3], [25] that provide fast lookups and enable resource pooling. However in pervasive environments, instead of a few devices leaving or joining the network – each device is mobile and connectivity is provided by ad hoc networking – complicating any hashtable update or maintenance. Also, note that unlike P2P systems where all nodes are directly addressable(reachable) by each other, in ad hoc networks mobile devices are restricted in their connectivity to only other devices in their neighborhood. Mobile devices have to rely mostly on locally available information.

However, motion of mobile devices is relative. From the perspective of a mobile device, the device can consider itself stationary and all available resources and device continuously moving. The following scenario presents a situation where some devices are relatively stationary with respect to the device. We use this scenario to exemplify and outline the characteristics of pervasive environments and the expected behavior of mobile devices.

Scenario *In Figure 1 mobile devices are housed in each of the vehicles. The vehicles are moving in the same direction. Though the vehicles may change relative positions their relative displacements while traveling in the same direction at comparable speeds. Changes in their relative positions are expected to be slow compared to their surface velocities.*

In the above scenario we expect the devices to utilize data and resource information from their local peers. Moreover, we assume that multiple devices will be able to cache information coming from the same source in the vicinity. Therefore, replication of the information is likely to increase data availability. Similarly, collaborative verification will help eliminate corrupt data. In the remaining sections of this paper, we explore the various possibilities and benefits of pack formation and the costs involved.

III. Trust-based Framework

A. Assumptions

Such scenarios are amenable to pack formation since mobility patterns, common intent, and activities can be associated with such a set of devices. Note that it is usually unlikely for any such aggregation of devices to have a majority of devices with a common affiliation or prior trust relationships. This scenario exemplifies situations that have scope for mutual collaboration. Such devices are in fact bound by a commonality in the physical world. Context information of such kind is of

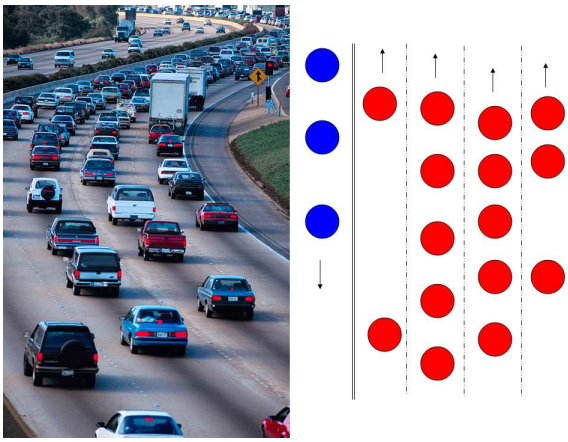

Fig. 1. Freeway example

crucial importance for pack formation since the situation provides a natural incentive to collaborate.

We assume that the the mobile user is not continuously interacting with the device and making decisions. Instead the device is largely autonomous in its functioning and decision making – guided by the user's profile. The user is alerted if the device cannot make a crucial decision or if some event of interest occurs. The device is thus provided a specification of the kind of data it should be looking for. At certain intervals the user may refer to the device and avail of services that the device can offer or locate within the environs. The device's job is to provide trustworthy (ensure fidelity of data) services. Also it may cooperate with other devices to perform negotiated collective goals or tasks again based on the user profile or explicit notification from the user. The objective of the data management framework is to help locate specific resources or services and rate the quality and trustworthiness of those resources. It is usually not realistic to determine the correctness of data (sensors may malfunction or be disabled). The data collection mechanisms must provide non-repudiation and indicate trustworthiness of the source. Data verification usually cannot be done unilaterally; federated or distributed trust mechanisms are required.

We assume that devices are able to sense their spatial and temporal contexts. We propose to use a set of landmarks or beaconing devices that advertise themselves and can be used by mobile devices to identify context. For example, certain constituents of the pervasive environments are stationary objects which can be identified as distinct entities associated with a particular geographical area.

Though movement of devices is unlikely to be restricted to particular geographical areas, they can be expected to be frequently seen in particular areas and their mobility patterns can be mapped to a large extent based on time, e.g., day of the week and time of the day [10], [13].

We also assume that devices' identities will be persistent, so that reputations can be associated with them. Such identifiers are required to be unique and possess capabilities for non-repudiation, e.g., cryptographically generated addresses [9] like Statistically Unique and Cryptographically Verifiable (SUCV) identifiers [18]. SUCVs provide an identity – a secure binding between an address (unique identifier) and a Public Key. They are particularly well-suited for such situations where centralized trust authorities like Certificate Authorities (CAs) or Key Distribution Centers (KDC) are often unreachable.

In the context of this document we refer to the overall individual assessment by an entity about another entity's integrity and worthiness to participate in collaborations – to access resources or services, either provided or shared – as *Trust*. For simplicity we refer to this assessment as 'trust', abstracting its possible manifestation as a complex multi-dimensional entity. Depending on the application and context, trust can be appropriately represented as a scalar or vector quantity.

We qualify trust to be a temporary assessment, having an instantaneous value, computed based on the reputation of the identity in question, the context, and the current intent. By reputation we mean a stateful quantity associated with an identity that stores information about past behavior, recommendations, accusations, associated neighborhoods etc. Reputation of a device is used in computing its trustworthiness.

B. Evolving, managing, and evaluating trust

Social networks will play an important role in both the data management and trust evaluation aspects of

pervasive computing. The context of a device is often used to guide its behavior [13], [14], [24]. We propose to additionally use associations of devices with the neighborhoods that they are seen in. For example, a student who frequently visits his University campus can notice other students on campus, and associate their identities with the context of the University campus. Given the finite number of potential neighborhoods that can be frequently visited, such a device can associate itself with particular neighborhoods of interest and time contexts – and also associate other devices with those neighborhoods. As shown in Figure 2(a), the device has its own view of the social network based on its own experiences and interactions. Figure 2(b) illustrates a global view of the social network showing the individual mobile devices, packs, and neighborhoods in the social network.

The trustworthiness of a peer or a resource in a neighborhood, can be computed from the history of prior encounters or can be provided by other trusted devices in the same neighborhood. Perich *et al.* have shown how relatively simple techniques that essentially involve incremental evolution of trust, based on positive and negative encounters, can provide a good first order approximation of trust [23]. When trust opinion is sought from peers, it can be weighted by one's degree of trust in them. Assuming that devices can be associated with particular neighborhoods it is logical to assume that at least some such "resident" devices are present who can provide trust assessments about other devices. Consequently sufficient information about trustworthiness of devices in the neighborhood already exists within the neighborhood, distributed amongst its frequent visitors. Devices can be categorized into: *trusted*, *offender* and *unknown*. Any device which is not a known offender is a candidate for admission into a pack. Membership can be later revoked if some other trusted sources later prove that device to be an offender.

A device which frequents a particular neighborhood will slowly build its reputation as a trustworthy resource. The degree of trust it enjoys in a particular neighborhood governs the level at which other devices and community resources are accessible or willing to cooperate. This knowledge of a device and its neighborhood associations thus serves as an enforcement for long term accountability. If a device exhibits malicious behavior, devices in the vicinity can effect an immediate response by denying further resources to the misbehaving entity. This response can be at all levels, from application to lower level networking layers. Furthermore, accusations of the recorded misbehavior are eventually propagated back to all its known home neighborhoods. Sufficient accusations from independent devices will lead to its reputation being tarnished in its home neighborhood. Such long term accountability will help deter malicious activity. Thus, we propose to provide reactive (reacting to activity deemed malicious) and pro-active (device with a history of misbehavior) measures for protecting packs and the mobile devices.

C. Pack dynamics and dominion

Segregation of mobile device identities by frequent encounters within areas of interest or specific neighborhoods will allow portable devices to leverage their large storage capacities to cache relevant information about potential trusted sources, yet at the same time minimize the storage requirements using attributes like location and frequency of encounters. Such users then have an incentive to collaborate towards common goals while they are in particular neighborhoods since the relationships in such environments are peer-to-peer. Additionally, since persistent identities will be associated with reputations, these entities are accountable for their actions within neighborhoods they frequently visit. Other mechanisms like recommendations from known sources can then be used to further increase the number of trusted sources. Maintaining reputation information will help maintain a set of devices likely to be encountered in particular neighborhoods that will have incentives to collaborate. Since a reputation is associated with a persistent identity, and reputations build up over time, malicious entities will not be allowed to participate in collaborations, or will be merely given less preference, effectively denying resources to such entities and limiting the damage that can be caused by malicious behavior. An entity changing its identity after committing a malicious act, will no longer enjoy the same level of reputation, thus undermining its ability to inflict further harm, by merely assuming a new identity. Since we assume reasonable parity in device capabilities, risk of Sybil attacks [15] can be mitigated by requiring potential members to solve computational puzzles before admission into packs.

D. Data routing substrate and storage management

Once social communities exist, it is possible for devices converging around a particular landmark to identify each other from their cached knowledge and trust relations. Pack formation is now possible from such a subset of entities willing to collaborate. Due to the inherent nature of the recommendation system, there are several incentives for such devices to collaborate, viz., (1) increased scope of search, (2) faster data retrieval, (3) updates related to trust information, (4) faster updates to existing data sets, (5) increased awareness of other trusted devices in the vicinity, (6) minimal effort in evaluating trustworthiness of providing

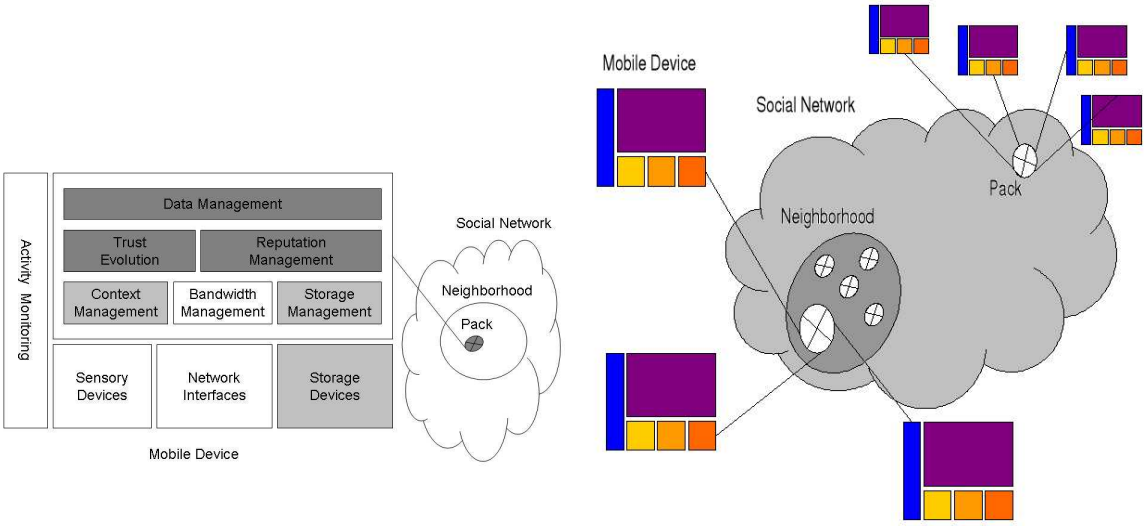

Fig. 2. (a) Device's view of a social network (b) Global view

source and accuracy of data, and (7) re-assertion of their own membership in the pack. Those devices that do not collaborate in these basic information exchanges risk losing their privileges by failing to reaffirm existing trust relationships. Also since devices will tend to cache the information related to the most recently encountered trustworthy and cooperative resources – uncooperative or passive devices stand to lose their established reputations by replacement in cache by other trusted sources.

With the existence of sufficient data for efficient trust evaluation of devices in the vicinity of any landmark, it is possible to form packs or groups that collaborate to achieve particular goals. Assurances of trust and accountability will allow devices to collaboratively perform specialized tasks depending on their current context and capabilities, and to perform far more complex operations than individually possible. Moreover devices can use their user and device profiles to selectively cache trust related information of only specific neighborhoods, thus minimizing individual storage requirements for trust related data, yet be largely self-reliant in assessing trustworthiness. Heterogeneous devices can form on-demand collaborations to share their unique capabilities and perform complex tasks with optimal planning. Such packs will also be more resilient to attacks from malicious entities, since they will have more knowledge of trustworthy devices in the vicinity, and will be able to isolate and ignore malicious entities.

Moreover, this formation will help create a data routing substrate for distributed storage management amongst the pack members collaborating to achieve (individual or collective) goals. We propose to implement such a data routing substrate that provide pack members with an inexpensive and efficient mechanism for locating reliable information sources and collaborative queries.

IV. Simulation

Our hypothesis is based on the assumption that collaboration amongst a set of trusted peers will improve the efficiency of search, increase scope of search, and decrease query response times. To test our hypothesis we conducted simulations using MoGATU's simulation environment implemented using GlomoSim [26]. The simulation parameters are provided in table I.

In order to allow devices to manage and reason over their profiles and profiles of their users, the model must enable devices to represent themselves and human users. To address this issue, we use MoGATU's model that represents devices, users and computational entities as intelligent entities – *agents*. By using the same representation for devices and users, the model allows mobile devices to express user's preferences and needs but also policies that restrict devices' actions.

The model is represented using OWL [8] and consists of a set of ontologies, i.e. vocabularies. Each ontology provides formal specification of the profile concepts and relationships among these concepts within a domain of discourse. The model enables mobile devices to express the needs and preferences of their users. It also facilitates mobile devices with exchanging and reasoning over the stored knowledge. Shared goals can then be inferred from these profiles, which consequently allow pack

formations – a group of devices collaborating with each other to solve individual and collective goals – based on trust relationships.

A. Pack formation and collaborative queries

Spatial Dimensions	150x150 m^2
Simulation Period	50 min
Simulation KB	Knowledge in vicinity, 40% – 100%
Mobile Devices	50
Mobility Pattern	Random Way-Point, 5 s waiting period, speed 1-5 ms^{-1}
Routing Protocol	AODV
Flooding Range	2 hops
Tx Range	$25m$
Tx Throughput	$2Mbps$
Device's Cache Size	250 kB, 50% of simulation KB
Device's Initial KB	100 questions, 100 answers not matching initial questions
Collaborators	4 to 9

TABLE I

SIMULATION ENVIRONMENT FOR COLLABORATIVE QUERIES

We simulated an environment with 50 nodes spread in random locations in a two dimensional square area. We present some of the interesting performance results from two separate sets of simulations. In the first case, each device assigned a task set of distinct questions, individually searches for answers. In the second case, the same set of devices with the same task set of questions search for the answers collaboratively. Since we also wanted to simulate the serendipitous nature of the environment, we varied the knowledge from 40% to 100% (40% knowledge means that answers to only 40% of each device's questions are present in the neighborhood, dispersed amongst the other devices).

In our simulation we assumed that some initial trust already exists to be able to form collaborative groups – packs. We experimented with pack sizes of 5 and 10, where each device was given a task set of 100 questions and the answers were randomly distributed amongst the total population of 50. We also varied the percentage of the knowledge base present in the neighborhood from 40% to 100% in increments of 20 percentage points. We ran the simulation using five different starting positions for the devices, for five runs of the simulation.

These simulations provided interesting results, some of which we present below. For simplicity, we assumed that all the sources of information were reliable and would only provide accurate answers. In the collaborative version, pack members help each other find answers to their questions. When an answer for a collaborator's question is found, the device tries to send it back to that collaborator.

The Figure 3(a) depicts 5 collaborators, each having a task set of 100 questions (not common with other collaborators). The devices themselves do not have answers to their own questions. The figure shows that the collaborative version is able to find twice as many answers in under a minute since the start of the querying process. In the non-collaborative version where devices independently try to query other devices in their radio-range, they manage to find approximately 50% of the answers and took upto 10 mins. In these simulations, 100% knowledge was available in the neighborhood.

Figure 3(b) shows the results, when there is only 40% knowledge in the vicinity, i.e., answers to only 40% of the questions are available. Here it is seen that the non-collaborative version was able to find no more than 5% of the answers, whereas the collaborating devices managed to find as many as 30% of the answers in less than 4 minutes.

Figures 4(a) and (b) show results of identical settings with 10 collaborators instead of 5. In case of 10 collaborators and 100% knowledge, more answers were found; yet the reponse times were slightly slower due to congestion introduced by pack members trying to relay back a large number of redundant answers over a short period of time.

All the simulations showed promising results in terms of faster responses and search effectiveness, in case of the collaborative models. We observed that as the pack size was increased from 5 to 10, the control overhead for communication between the pack members increased and introduced minor increase in latency to query responses, yet the number of successfully answered queries were consistently more than the non-collaborative version.

V. CONCLUSIONS AND FUTURE WORK

We described a holistic approach for addressing the highly interdependent issues of security, privacy, and trust related to data management in pervasive environments. Our proposed approach enables individual devices to harness the potential power of distributed computation, storage, sensory, and effector resources available in pervasive computing environments by evaluating the trustworthiness and accuracy of devices and their offered data and services. The key point of our approach is the pack formation mechanism, which allows devices to form social networks, which in turn allow

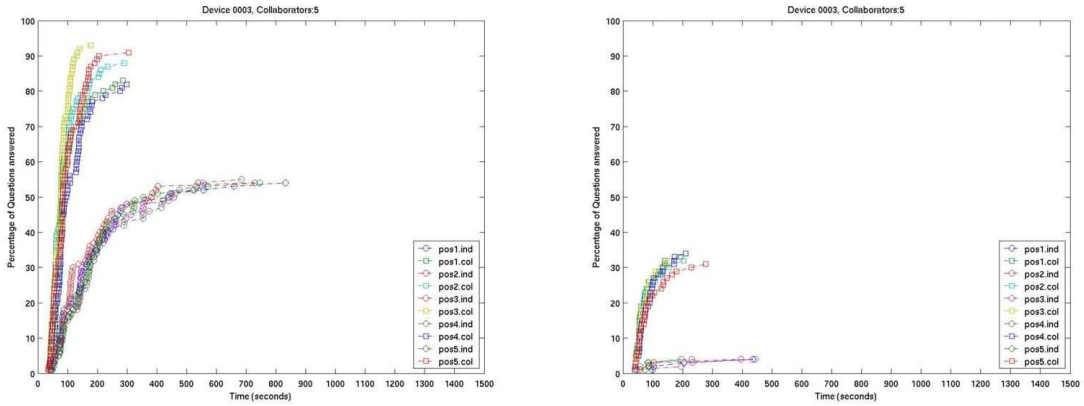

Fig. 3. (a) 5 collaborators, 100% knowledge (b) 5 collaborators, 40% knowledge

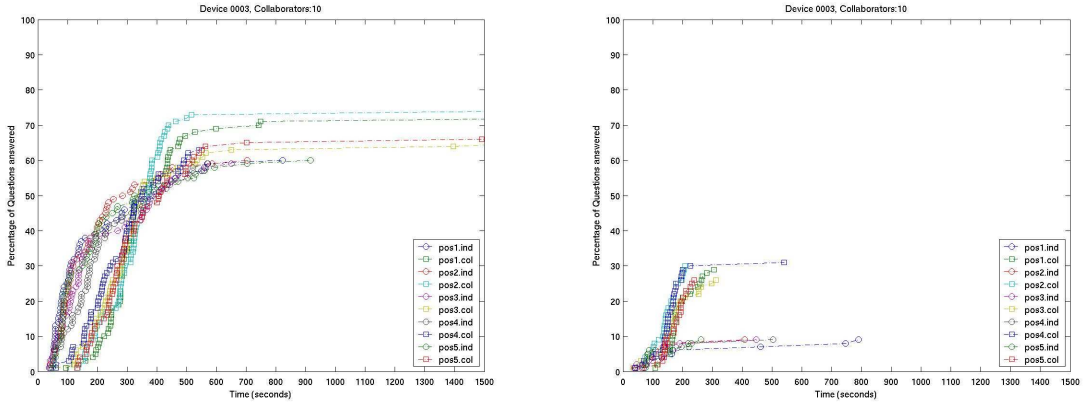

Fig. 4. (a) 10 collaborators, 100% knowledge (b) 10 collaborators, 40% knowledge

devices to detect and respond to malicious activities. We identified the important constituents of pervasive environments and their characteristics. We discussed issues regarding to reputation management, pervasive trust, and data management. We then presented our simulation environment and empirical results from our simulations that validate our framework as well as address the underlying issues. The benefits of pack formation are apparent from our preliminary results; for future work we plan to investigate trust evolution and reputation management using activity monitoring, and integrate those into the pack formation mechanisms.

REFERENCES

[1] g2dn: Gnutella2 developer's network. http://www.gnutella2.com/.

[2] IEEE Standards for Information Technology – Telecommunications and Information Exchange between Systems – Local and Metropolitan Area Network – Specific Requirements – Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications. http://standards.ieee.org/getieee802/download/802.11-1999.pdf.

[3] Open Source Napster Sever. http://sourceforge.net/projects/opennap/.

[4] Podcasting. http://en.wikipedia.org/wiki/Podcasting.

[5] Supplement to 802.11-1999,Wireless LAN MAC and PHY specifications: Higher speed Physical Layer (PHY) extension in the 2.4 GHz band. http://standards.ieee.org/getieee802/download/802.11b-1999.pdf.

[6] Apple iPod. http://www.apple.com/ipod/, February 2005.

[7] Sony Playstation Portable(PSP). http://www.us.playstation.com/psp.aspx, March 2005.

[8] Web Ontology Language (OWL). http://www.w3.org/2004/OWL/, February 2005.

[9] T. Aura. Internet Draft: Cryptographically Generated Addresses (CGA). http://www.ietf.org/proceedings/04mar/I-D/draft-ietf-send-cga-05.txt, February 2004.

[10] J. B. Begole, J. C. Tang, R. B. Smith, and N. Yankelovich.

Work rhythms: analyzing visualizations of awareness histories of distributed groups. In *CSCW '02: Proceedings of the 2002 ACM conference on Computer supported cooperative work*, pages 334–343. ACM Press, 2002.

[11] M. Cherniak, M. Franklin, and S. Zdonik. Expressing User Profiles for Data Recharging. *IEEE Personal Communications*, July 2001.

[12] M. Cherniak, E. Galvez, D. Brooks, M. Franklin, and S. Zdonik. Profile Driven Data Management. In *28th International Conference on Very Large Databases*, August 2002.

[13] M. J. Covington, W. Long, S. Srinivasan, A. K. Dey, M. Ahamad, and G. D. Abowd. Securing context-aware applications using environment roles. In *SACMAT '01: Proceedings of the sixth ACM symposium on Access control models and technologies*, pages 10–20. ACM Press, 2001.

[14] A. Dix, T. Rodden, N. Davies, J. Trevor, A. Friday, and K. Palfreyman. Exploiting space and location as a design framework for interactive mobile systems. *ACM Trans. Comput.-Hum. Interact.*, 7(3):285–321, 2000.

[15] J. R. Douceur. The sybil attack. In *IPTPS*, pages 251–260, 2002.

[16] M. J. Franklin. Challenges in ubiquitous data management. In *Informatics - 10 Years Back. 10 Years Ahead.*, pages 24–33. Springer-Verlag, 2001.

[17] L. Kagal, T. Finin, and A. Joshi. A Policy Language for A Pervasive Computing Environment. In *IEEE 4th International Workshop on Policies for Distributed Systems and Networks*. June 2003.

[18] G. Montenegro and C. Castelluccia. Statistically Unique and Cryptographically Verifiable(SUCV) identifiers and addresses. citeseer.ist.psu.edu/montenegro02statistically.html, 2002.

[19] F. Perich. MoGATU: Data Management in Pervasive Computing Enviroments. http://mogatu.umbc.edu/, 2001-2005.

[20] F. Perich. MoGATU BDI Ontology. http://mogatu.umbc.edu/bdi/, 2004.

[21] F. Perich, S. Avancha, D. Chakraborty, A. Joshi, and Y. Yesha. Profile Driven Data Management for Pervasive Environments. In *13th International Conference on Database and Expert Systems Applications (DEXA 2002)*, Aix en Provence, France, September 2002.

[22] F. Perich, A. Joshi, T. Finin, and Y. Yesha. On Data Management in Pervasive Computing Environments. *IEEE Transactions on Knowledge and Data Engineering*, May 2004.

[23] F. Perich, J. L. Undercoffer, L. Kagal, A. Joshi, T. Finin, and Y. Yesha. In Reputation We Believe: Query Processing in Mobile Ad-Hoc Networks. In *International Conference on Mobile and Ubiquitous Systems: Networking and Services*, Boston, MA, August 2004.

[24] Roy, A. and Das Bhaumik, S.K. and Bhattacharya, A. and Basu, K. Cook, D.J. and Das, S.K. Location aware resource management in smart homes. In *Proceedings of the First IEEE International Conference on Pervasive Computing and Communications*, pages 481–488, 2003.

[25] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan. Chord: A scalable peer-to-peer lookup service for internet applications. In *Proceedings of the 2001 conference on applications, technologies, architectures, and protocols for computer communications*, pages 149–160. ACM Press, 2001.

[26] X. Zeng, R. Bagrodia, and M. Gerla. GloMoSim: A Library for Parallel Simulation of Large-Scale Wireless Networks. In *Workshop on Parallel and Distributed Simulation*, 1998.