Draft version of manuscript to appear in the International Journal of Wireless Information Networks, April 2006.

Querying in Packs: 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 describe a trust-based data management framework enabling mobile devices to access the distributed computation, storage, and sensory resources available in pervasive computing environments. Available resources include those in the fixed surrounding infrastructure as well as services offered by other nearby mobile devices. We take a holistic approach that considers data trust, security, and privacy and focus on the collaborative mechanisms providing a trustworthy data management platform in an ad hoc network. The framework is based on a pack formation mechanism that enables collaborative peer interactions using context information and landmarks. A pack provides a routing substrate allowing devices to find reliable information sources and coordinated pro-active and reactive mechanisms to detect and respond to malicious activity. Consequently, a pack forms a foundation for distributed trust management and data intensive interactions. We describe our data management framework with an emphasis on pack formation in mobile ad-hoc networks and present preliminary results from simulation experiments.

I. INTRODUCTION

A. Motivation

Advances in technology and the growing demand for wireless connectivity are fueling a proliferation of wireless capabilities in everyday appliances. There is a large heterogeneity in the types of portable devices and their communication, computational, and storage capabilities. Moreover, computation capable embedded electronic devices equipped with sensors and actuators are transforming home, office and urban landscapes into resource-rich and data-intensive environments.

Peer interactions that use local ad hoc connectivity are most suited in locating and consuming services available in vicinity, since infrastructure based continuous connectivity to the Internet is often unavailable or expensive. Our focus is on building a trustworthy data management framework for personal mobile devices in pervasive environments which utilize wireless connectivity to access resources provided by self-organizing networks. We characterize such mobile ad hoc networks (MANETs) as having two kinds of nodes (or actors): (i) fixed assets businesses, including restaurants and stores, public service kiosks providing traffic, weather, road conditions; and (ii) mobile devices, including mobile phones or car computers. Fixed assets can thus be associated with geographical locations, whereas mobile devices are assumed to have certain mobility profiles.

Existing data management architectures [8], [21] assume that encountered devices and the information obtained from them are trustworthy. This assumption, however, is unrealistic in a pervasive environment [14]. Therefore, assurances of the quality and accuracy of the retrieved information must be provided. Limited Internet connectivity makes it infeasible to use conventional security paradigms for determining the trustworthiness of other devices. For example, communicating with a trusted authority in order to verify the credentials of nearby entities may be impossible. Moreover, the nature of these environments requires these devices to function independently and be capable of making autonomous decisions about the trustworthiness of peers and accuracy of their information. Metrics for evaluating the reliability of data and the trustworthiness of peers are critical to achieving secure and trustworthy data management networks. Given the lack of centralized authorities in ad hoc networks, trust evaluation and reputation management are essential mechanisms that allow devices to function autonomously with minimal user intervention.

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

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]. The MoGATU framework [19] demonstrated the utility of reputations in ascertaining the accuracy of information exchanged between mobile devices. However, significant effort is required for evolving trust and maintaining reputations in terms of computation, storage, and messages interchanged.

B. Approach Overview

Our approach to modeling trust is focused on risk management and coping with the inherent uncertainty arising from the serendipity in locating and communicating with other devices or resources. Also, it is necessary for devices to be able to dynamically respond to changes in the environment. For our purposes, the definition of trust as proposed by Luhmann [16] as a means for managing risk, and reducing the complexity of the environment, is most fitting.

In the context of MANETs, key factors in ascertaining *situational trust* [17] include knowledge of the environment regarding mobility, network topology, and - proximity, reachability, and availability of resources. Other factors include familiarity or knowledge of other node identities in the vicinity. Thus, we stress the need for devices to be aware of their context when making trusting decisions.

A very basic form of trust is assumed to exist amongst all the participants in an ad hoc network - the expectation that other devices in the neighborhood will participate in the routing process and help relay data traffic. Without the existence of a basic communication mechanism to relay messages, no further interactions are possible. Higher degrees of trust can exist when devices seek to cooperate beyond provisioning the basic connectivity either for incentives or out of necessity.

Reputations are used to estimate the risk of future cooperation based on past behavior. Recommendations are sought from other trusted peers when prior knowledge or reputation of the entity in question is not available. Our approach hinges on the use of normative behavior and activity monitoring in order to evolve trust relationships, maintain reputations, and detect malicious activity.

We define a pack as a dynamic grouping of individual devices which agree to collaborate in pursuit of individual, mutual and collective goals. In this paper, we describe a collaborative approach that addresses issues of trust and security in data management and present initial results from simulation experiments on pack formation.

Our approach is based on providing primitives for trust evolution and pack formation. A pack formation serves two purposes: (i) a pack provides a routing substrate enabling devices to find reliable sources of information, and (ii) a pack supports 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 begin with a scenario that illustrates typical pervasive environments and helps identify the important constituents of pervasive environments and their characteristics. We discuss issues regarding reputation management, pervasive trust and data management and present our proposed trustworthy data management framework that employs pack formation to provide a data routing substrate. We then discuss the performance results from our simulations of such scenarios and make the case for collaborative querying using packs. Finally, we conclude with a discussion on the cost and benefits of pack formations and ideas for future work.

II. BACKGROUND AND RELATED WORK

A. Trust and Reputation Management in MANETs

Pervasive environments are in a constant state of flux. Routing mechanisms in self-organizing networks, which provide the communication mechanism, are necessarily of a collaborative nature. The inherent uncertainty in locating and accessing resources (services and data) further necessitate collaboration. 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. Collaboration on these various fronts requires various degrees of trust in potential collaborators.

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 in order to make appropriate trust evaluations and use reputations to guide their behavior.

With a lack of security enforcing mechanisms or arbitrators, devices need to cope with the freedom of others and the possibility of malicious activity. There is a need to be able to adapt to the emergent properties in the environment and make autonomous decisions using observed behavior, available reputations, and recommendations.

Several trust and reputation management systems have been proposed for online transactions amongst buyers and sellers in e-Commerce [13], in Peer-to-Peer systems [15], for fostering trust and cooperation. Buchegger and Le Boudec [6] propose using a Bayesian approach for reputation management system for MANETs in which nodes maintain reputations for only specific nodes. Reputation management schemes depend on observed behavior and recommendations - positive and/or negative, based on expected "ethical" behavior. Since it is often impossible to enforce protocol fidelity, sanctions need to be imposed collectively to make deviant behavior expensive and to discourage reneging.

Due to the potentially innumerable number of devices, it is unlikely to be able to individually cache all possible identities and their reputations. Moreover, it will be infeasible to maintain a consistent view on trustworthiness amongst all devices. We propose that it is sufficient to remember only devices that are of future potential value in forming social networks and those that will be most likely to cooperate, and maintain reputations regarding only those.

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.

B. Data Management in Pervasive Environments

Franklin [12] 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.* [7], [8] 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 [12] has proposed *Adaptive Dataflow Query Processing* to address collecting data from such data streams. He suggests 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 is 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 have the option to either continuously monitor/use the resource (identified by its distinguishing source identity) or get updates of information of interest at a later time. 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 in use - through providing connectivity to the resource or intermediate storage - 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. In other words, they 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. Thus the overall network topologies are fairly static, which make it possible to efficiently maintain and update distributed hashtables [1], [2], [25] that provide fast lookups and enable resource pooling. 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, which limits any hashtable update or maintenance. Additionally, 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.

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.

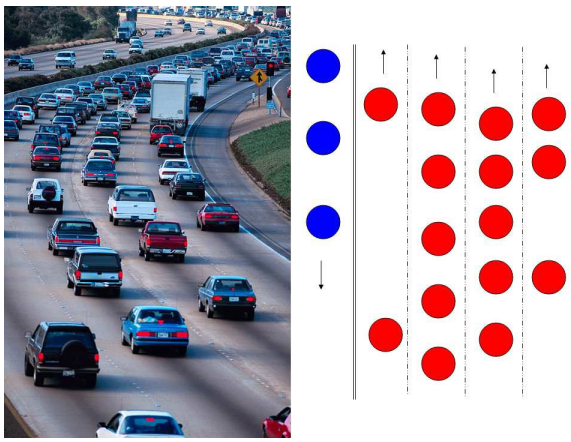

Scenario: *Mobile devices are housed in each of the vehicles. The vehicles are moving in the same direction. Although, the vehicles may change relative positions and relative displacements while traveling in the same direction at comparable speeds. Changes in their relative positions are expected to be slow compared to their surface velocities.*

Fig. 1 Freeway example

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.

Such scenarios are amenable to pack formation since mobility patterns, common intent, and activities can be associated with such a set of devices. It is unlikely for any such aggregation of devices to have a majority of devices with a common affiliation or prior trust relationships. However, being bound by a commonality in the physical world
these devices have a natural incentive to collaborate. Common context information from the physical surroundings is of crucial importance for pack formation since the situation provides nucleating points and thus the incentive to collaborate.

III. TRUST-BASED FRAMEWORK

A. Assumptions

We assume that a mobile user is not continuously interacting with the device and making decisions. Instead the device is largely autonomous in its functioning and decision making while 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 available services the device can offer or locate within the environment. The device should then provide trustworthy services by ensuring fidelity of the services and the data they provide. Also it may cooperate with other devices in order 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, as 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; instead, federated or distributed trust mechanisms are required.

We assume that devices are able to sense their spatial and temporal contexts. We 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, devices can be expected to be seen frequently 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 [5], [9].

We also assume that devices' identities are persistent. Therefore, reputations can be associated with them. Such identifiers are required to be unique and possess capabilities for non-repudiation, e.g., cryptographically generated addresses [4] 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.

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, for example, about past behavior, recommendations, accusations, or associated neighborhoods. 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 [9], [10], [24]. Additionally, we use associations of devices with the neighborhoods that they are seen in. For example, a student who frequently visits her 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.

Figure 2(a) shows a hierarchy of the components involved in the data management framework on each individual device. Further each device is shown to be a pack member in some neighborhood, each part of a larger social network. Each 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.

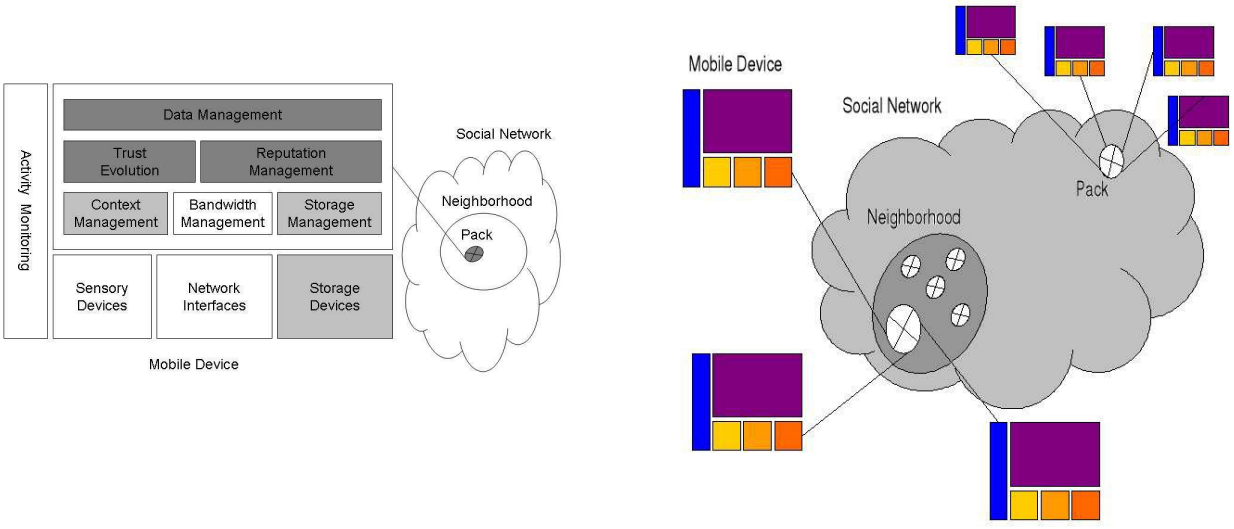

Fig. 2. (a) Constituents and component-hierarchy of the data management framework (b) Global view of the pervasive environment

Several definitions and formalisms of trust exist in literature [17]. We choose to simplify certain aspects of trust in order to use a specific formalism that allows us to build a reputation system. We choose to represent trust using a continuous value in the interval [0,1].

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. In our previous work [23], we have shown how relatively simple techniques that involve incremental evolution of trust, based on positive and negative encounters, can provide a good first order approximation of trust. 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. The necessity of cooperation in peer-interactions is related to the current degree of resource and data availability. The need is higher when availability is low, i.e., when trustworthy resources who can provide the required data are unavailable or limited. Our approach relies on monitoring the availability of resources and levels of cooperation of other peers for adapting to current conditions and to reciprocate peer-interactions. We observe positive and negative experiences in terms of events, such as validated or invalidated data or a misbehavior at lower networking layers. We consider these observations in the reputation management while evaluating trustworthiness of the peer and in order to determine a level of cooperation upto which to reciprocate in collaborative processes. This mechanism serves as an encouragement for cooperation - a non-cooperating peer will get diminishing attention from its peers while cooperativeness is automatically rewarded. We consider two aspects of trust viz. the accuracy of data provided (validation/verification) and the degree of cooperation.

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 helps to deter malicious activity. Thus, we 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 are associated with reputations, these entities are accountable for their actions within neighborhoods they frequently visit. Other mechanisms such as recommendations from known sources can then be used to further increase the number of trusted sources. Maintaining reputation information helps in maintaining a set of devices likely to be encountered in particular neighborhoods that 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 infiict further harm, by merely assuming a new identity. Since we assume reasonable parity in device capabilities, risk of Sybil attacks [11] 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 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. Since devices 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 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 are also more resilient to attacks from malicious entities, since they have more knowledge of trustworthy devices in the vicinity and can isolate and ignore malicious entities.

Moreover, this formation helps create a data routing substrate for distributed storage management amongst the pack members collaborating to achieve (individual or collective) goals. We are implementing such a data routing substrate that can 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 [3] 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

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. We present some of the results below. 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.

Spatial Dimensions	150 X 150 m^2
Simulation Period	50 *min*.
Simulation KB	Knowledge in vicinity, 40% - 100 %
Mobile Devices	50
Mobility Pattern	Random Waypoint, 5 *s* waiting period Speed 1-5 *m/s*
Routing Protocol	AODV
Flooding Range	2 hops
Tx Range	25 *m*
Tx Throughput	2 *Mbps*
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

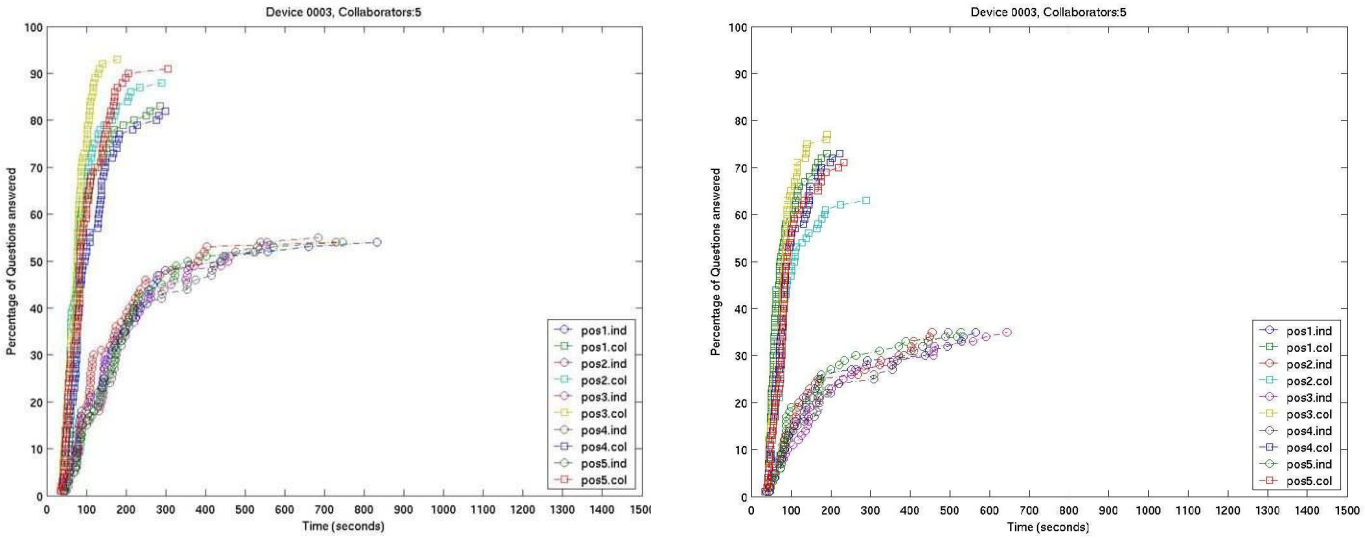

(a) 5 collaborators, 100% knowledge (b) 5 collaborators, 80% knowledge

Fig. 3.

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 radiorange, they

manage to find approximately 50% of the answers and took upto 10 minutes. In these simulations, 100% knowledge was available in the neighborhood.

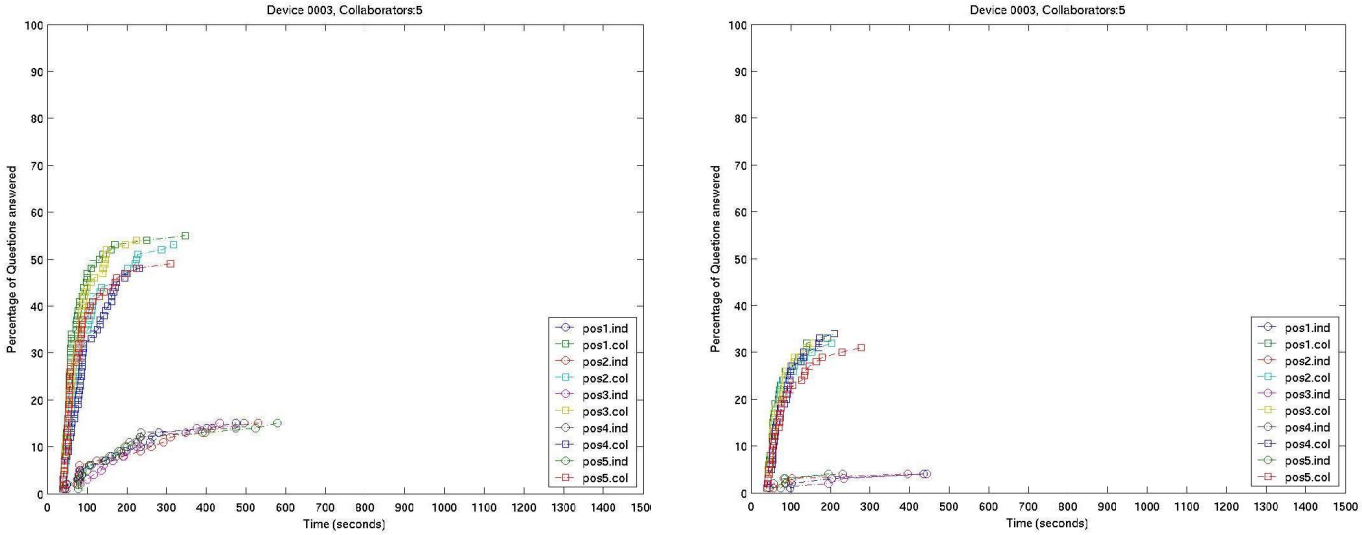

(a) 5 collaborators, 60% knowledge (b) 5 collaborators, 40% knowledge

Fig. 4.

Figure 4(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.

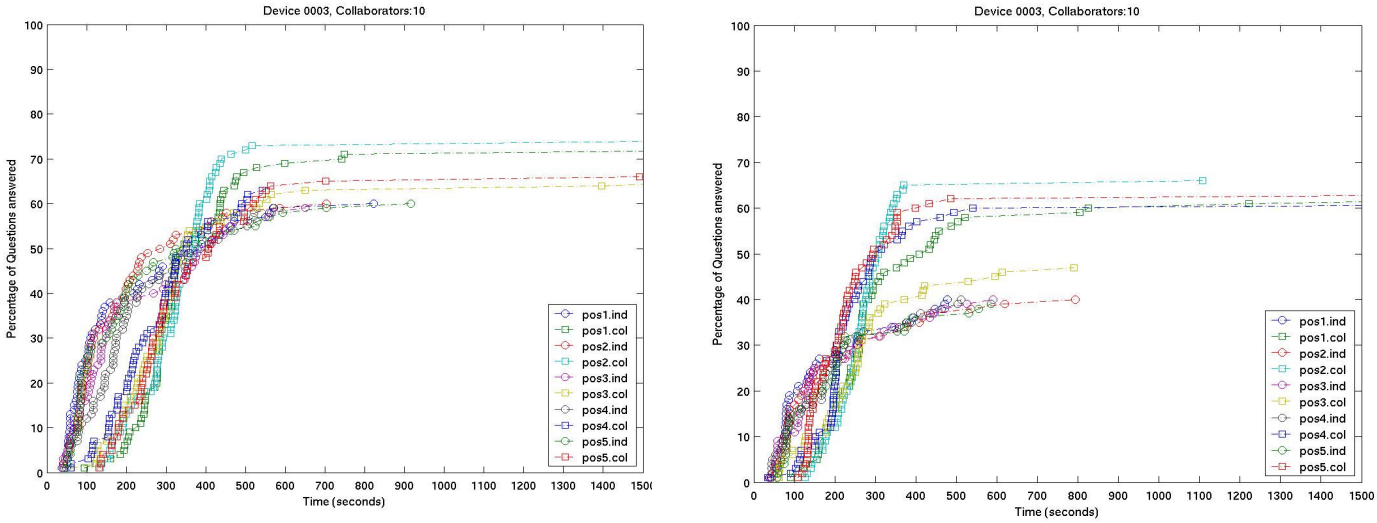

(a) 10 collaborators, 100% knowledge (b) 10 collaborators, 80% knowledge

Fig. 5.

Figures 5 and 6 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 response 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.

As illustrated in the figures, all simulations showed positive 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.

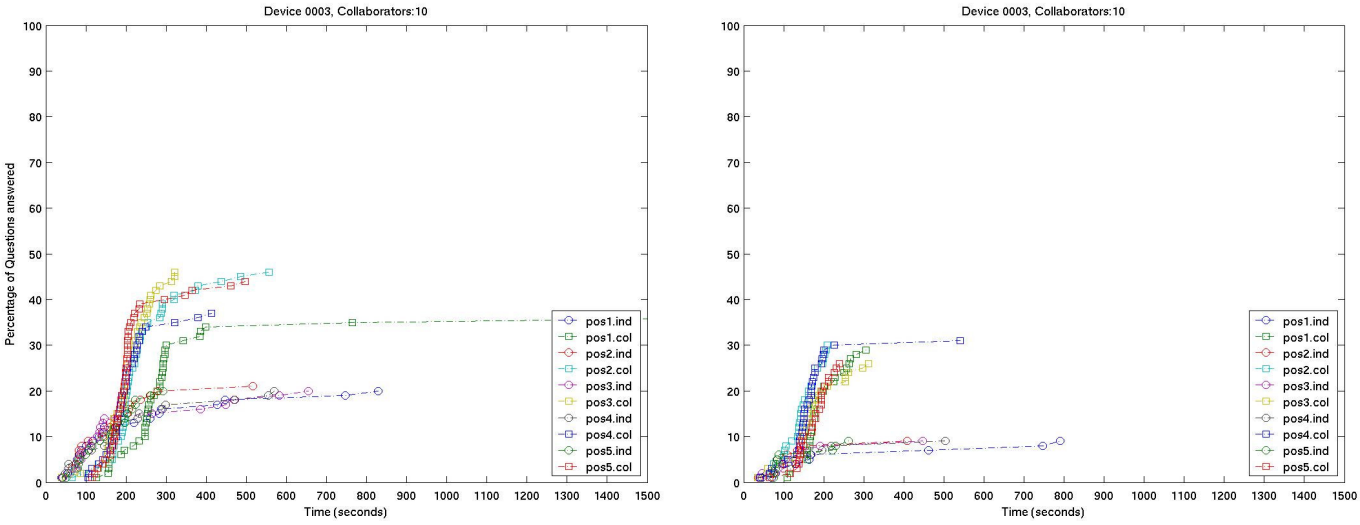

(a) 10 collaborators, 60% knowledge

(b) 10 collaborators, 40% knowledge

Fig. 6.

V. CONCLUSIONS AND FUTURE WORK

Our approach focuses on enabling efficient resource location and trustworthy data management using pack formations. In order to enable pack formations we suggest the use of fixed landmarks in the pervasive environment to serve as nucleating points for social networks of frequent visitors. We presented scenarios to elucidate the context and identify the various crucial aspects, in evaluating trust in pervasive environments. Collaboration in routing and data management is key to the functioning of MANETs. Moreover, time is of essence in making trust decisions while locating resources and utilizing services as access is usually short-lived. While reputation management is essential in fostering cooperation and making trusting decisions, pack formations further simplify the decision making processes and provide added stability to resource location mechanisms in the highly dynamic MANET conditions, enabling an essential platform for trustworthy data management.

Forming local packs as decentralized and redundant trust management authorities is helped by the use of the large on-board storage capacities, nucleating points for storing reputation, and other geographic context information.

We described a holistic approach for addressing the highly interdependent issues of security, privacy, and trust related to data management in pervasive environments. The 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.

We simulated this scenario with varying pack sizes and the results demonstrate the superior performance of collaborative querying. Collaborative querying yielded highly improved success rates and response times. The results validate our hypothesis that increasing the scope of search over multiple devices is vastly superior to repeated querying. Cost of querying increases as the number of collaborators making the same query increase, and is also seen to affect the response time. 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] Open Source Napster Sever. http://sourceforge.net/projects/opennap/.

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

[4] T. Aura. Internet Draft: Cryptographically Generated Addresses (CGA). http://www.ietf.org/proceedings/04mar/I-D/draft-ietf-send-cga05.txt, February 2004.

[5] 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.

[6] S. Buchegger and J.-Y L. Boudec. A robust reputation system for p2p and mobile ad-hoc networks, 2004.

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

[8] 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.

[9] 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.

[10] 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.

[11] J. R. Douceur. The Sybil attack. In *IPTPS*, pages 251-260, 2002.

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

[13] R. Guha, R. Kumar, P. Raghavan, and A. Tomkins. Propagation of trust and distrust. In *WWW '04: Proceedings of the 13th international conference on World Wide Web*, pages 403-412, New York, NY USA, 2004. ACM Press.

[14] 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.

[15] S. D. Kamvar, M. T. Schlosser, and H. Garcia-Molina. The eigentrust algorithm for reputation management in p2p networks. In *WWW '03: Proceedings of the 12th international conference on World Wide Web*, pages 640-651, New York, NY USA, 2003. ACM Press.

[16] N. Luhmann. *"Trust and Power"*. John Wiley & Sons, Inc., 1st edition, 1979.

[17] S. Marsh. Formalising trust as a computational concept, 1994.

[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.