

Trust and Mobility Aware Service Provision for Pervasive Computing

Liam McNamara, Cecilia Mascolo and Licia Capra

Department of Computer Science, University College London
Gower Street, London, WC1E 6BT, United Kingdom
{l.mcnamara | c.mascolo | l.capra}@cs.ucl.ac.uk

Abstract. Successful provision of services in pervasive networks represents a major challenge. The requester of a service or resource should be able to decide whom to rely on, not only depending on the apparent QoS that various providers offer, but also on the *trustworthiness* of the providers, as well as their *mobility* patterns. Existing approaches offer a rather simplistic solution to the problem: either resources and services are discovered and selected purely based on their name/category, or, at most, by looking at the QoS promised by the various providers. In this paper we present a service provision framework that reasons about mobility and trust to enable effective service selection in mobile ad-hoc networks. We illustrate our ideas through a scenario of service sharing on public transport where passengers are able to select other passengers' devices for information transfer, based on a novel combination of quality attributes (e.g., network connectivity), trust and mobility.

1 Introduction

Mobile devices and embedded computers are gradually becoming pervasive in modern life. An extremely large percentage of people already possesses devices such as portable music players and mobile phones. These devices are quickly growing both in terms of supported functionalities and computational capabilities. Crucially, an increasing number of these devices are being networked, enabling services such as file sharing, interactive games and information exchange, among these mobile peers. To truly realise Mark Wieser's ubiquitous computing vision, the management and operation of these devices cannot be explicitly human controlled. Rather, spontaneous communication and co-operation must be performed with the burden of configuring devices being abstracted away from the user, thus fostering a simpler and more satisfactory user experience.

The concept of *service orientated* computing is very useful in wireless environments, where there is little fore-knowledge of which services and peers will be available at any given time and location. The ability to seamlessly search for the desired service, and to select the one which has specific *qualities* that will maximise the user's utility is crucial. Two major issues underlie this challenge: first, the ability to assess whether a provider will indeed deliver a service at the promised QoS (service providers may inflate their service advertisements). We

thus need to provide the device with the ability to reason about the *trustworthiness* of the various providers, so as to make accurate predictions about actual QoS (as opposed to advertised/promised QoS), and consequently maximise the chances of a satisfying service delivery. Second, the service selection mechanism must account for an orthogonal parameter, that is, *mobility*. For example, we must maximise the chances that the service requester and provider will be co-located for a sufficiently long period of time, in order to complete the service delivery.

Some solutions in this space have been presented, especially on intelligent QoS-aware routing in multi-hop networks [7], though less attention has been given to higher level tasks such as *service selection*. In order to gain security and robustness, maintaining reputations about the performance of hosts in a system will discourage misbehaviour. A centralised repository is used in CONFIDANT [2] to track global reputations, leading to a single point of failure and load, limiting scalability. Distributed global trust methodologies have also been proposed [6], but connectivity to a sufficient number of non-malicious peers is required, which may not be guaranteed in mobile situations. Some local reputation systems have been suggested, as used by Dewan et al [5], where each peer maintains its own opinion of other peers, which we believe offer greater flexibility and robustness. One approach which seems quite close to ours is [4]. In that paper, a model for trust based service interaction in decentralised systems is presented. However, we further consider mobility as an essential qualitative measure on which the service provider must be chosen; we have not seen any approaches making use of this information so far. The trust model used in [4] is also quite limited in the sense that the long-term storage of reputations is centralised. In this paper, we propose a framework where each peer can independently reason about the QoS and reliability of other peers to select the ‘best’ service provider available without relying on information from other hosts. This paper also builds on previous work [3], where we outlined a QoS based service discovery and selection framework, now with the addition of mobility and trust reasoning linked to quality of service.

2 Scenario

In big cities, such as London, a very large number of people commute between suburbs and the centre on public transport (e.g. buses and trains). Commuters on these vehicles are usually in quite close proximity. Most carry handheld devices with one or more network interfaces (e.g. 802.11, Bluetooth, GSM), their patterns of mobility are quite “seasonal” (in the sense that they travel usually at the same time, repeating the same path day after day), and they tend to stay on the vehicle for quite a prolonged period of time¹. In addition, devices are often diversely equipped and, as a result, a wide variety of services could be offered or shared among people, through their devices:

¹ This is also true for people in coffee shops as indicated by [4].

- Location or time information sharing: a device with a GPS receiver could be serving this information to others.
- News headlines, stock market levels: someone able to access the Internet through a GPRS phone could forward updates to others.
- Gaming: devices could participate in a shared game during their trip.
- Software components: new applications/functionalities could be shared and downloaded from a peer.
- Information about traffic and delays: commuters traveling in different directions could inform each other.




Fig. 1. Train scenario diagram

To illustrate a possible scenario (Figure 1), let us assume that *Alice* gets on a busy train with her PDA and starts her ‘Travel Planner’ program, causing her PDA to broadcast a request for location information. Many other passengers have mobile computing devices, including *Bob*, *Carol* and *Daniel*. Any device that receives this request, and that is providing a matching service, can potentially respond to her. Let us assume that *Bob*, *Carol* and *Daniel* have a GPS receiver (i.e., they are able to provide the service), and that they send back a response, containing all pertinent attributes of their service (i.e., accuracy and age). Moreover, each device maintains information about the user’s mobility pattern, based on his/her past journeys; this information is also sent back to *Alice*, so that her device may estimate for how long she will be co-located with each service provider, thus being able to make a selection that minimises resource expensive switches of service providers. After receiving responses from co-located devices, *Alice* can then decide whom to rely on. This will involve pruning all unacceptable responses that do not satisfy her requirements (e.g., accuracy lower than a minimum threshold). The remaining services are all feasible to use, but *Alice* should select the ‘best’, according to her own utility.

To this end, let us assume that *Alice* is mainly concerned with accuracy of the location data. A first ranking of the short-listed service providers will then

be done based on the advertised accuracy. Let us also imagine that, besides accuracy, Alice is interested in minimising switches among service providers and having a recent location. The ranking will thus be influenced by the probability of each service provider remaining co-located with Alice for as long as she is expected to remain on the train and the location age. As a result, it may be that Carol promises better accuracy than everyone else; however, if Alice’s trust in her is low as a result of past interactions, Carol will actually go down in the ranking. Similarly, it may be that Daniel promises higher accuracy than Bob, but he is expected to get off at the next station, while Bob is predicted to leave after Alice. In such a situation, our framework would put Bob at the top of the ranking.

The interaction with Bob can then be attempted, and the communication initiated using the relevant encoding and protocols. Depending on the outcome, Alice’s trust in Bob will face either an increase or a reduction. The magnitude of any reduction will depend on the difference between promised and perceived service quality; in doing so, a problem that may not be the intentional fault of Bob should not be punished as much as an obvious QoS inflation on his part. The application would be able to disregard a location value that is wildly different from the previous reading (if taken recently), and signal to the middleware a failure has occurred. Failure would then lead to the second highest ranked provider being used. Conversely, if the experienced QoS is very close to the promised level, Alice’s trust in Bob will increase.

2.1 Requirements of the Scenario

From the scenario(s) described above, we can elicit the following requirements for a pervasive service selection framework:

- Ability to encode both functional attributes and QoS requirements in service requests, with the latter including the ability to express conditions on the execution context of the device (e.g., power level, predicted co-location).
- Ability to judge the outcome of a service request, with respect to the relevant attributes. Monitoring of service attributes can be automated by a middleware, but assessing responsibility is a much harder task, and application logic must be employed to discern most problems.
- Ability to share the mobility patterns of a peer with other hosts, to give an indication of how long both peers will be co-located (it is not the goal of this paper to discuss how these can be shared anonymously to protect user’s privacy).
- Ability to reason about a device’s trustworthiness in different and sometimes new environments.

3 Our Approach

We intend to create a framework that satisfies these requirements and that is also lightweight enough to not congest the network or overly load resource-constrained devices. Internet-based resource discovery approaches (such as HTTP

URIs [1]) rely on explicit naming of required resources; however, in mobile networks, this is less than ideal. Rather, the type (e.g., *printer*) and attributes of a required resource should be specified, to allow the searching and binding to a sufficiently suitable provider within communication range. This kind of search requires a shared ontology between requesters and providers. Ontologies have been investigated in OWL [8], which is predominantly an Internet technology. In this paper, we do not elaborate on the ontology to use, but, rather, on how the service description (QoS and mobility based) can be used for selection, how it can be monitored with respect to the actual service provision, and how it can be integrated with trust and reputation information. We will now give a description of the QoS and mobility aspects of the proposed framework, and then focus on trust maintenance and sharing.

3.1 Quality of Service and Mobility Based Selection

In highly dynamic, sometimes faulty, networks, it is necessary to ascertain which peers can provide acceptable QoS levels in a given context. QoS constraints vary from application to application, and from context to context. In Q-CAD [3], we proposed using *context-aware* QoS specifications of services. For instance, a host may forgo requesting encryption of communication if their battery level is low, to conserve power. We now enrich QoS specifications with mobility considerations, so that a service description can also be evaluated based on the expected co-location time of client and server. This information is vital for services that will take a significant amount of time to complete, like file downloads.

In an open network with heterogeneous nodes, it is sensible to use an easily parseable format, such as XML, for the representation of QoS requirements. The metrics included would have to be measurable locally by the peer, but could include network qualities, such as congestion and co-location. Local metrics would most likely be battery power and load. Each application thus has an associated *Application Profile* that defines the attributes of importance to that specific program. When a resource/service is searched for, it is the middleware's responsibility to evaluate the profile of the requesting application (i.e., application-dependence), to monitor the operating context (i.e., context-dependence), and to then broadcast the appropriate request. In the scenario we are considering (see Section 2), mobility is specified in the profile as an estimate of the time a host will be present in the current environment. This is quite a naïve way to measure mobility; for example, it does not take into account many of the finer points of possible movement. It is, however, a useful metric to easily estimate and compare hosts' co-location, which is the mobility aspect we are currently interested in. Note that a minimum trust level may be required, to discard untrustworthy providers for which a reliable QoS prediction cannot be made. The application profile is communicated to the middleware upon startup, but it remains accessible to the application for run-time modification. The same application may register different services with the middleware, to allow fine-grained customisation. An example *Application Profile* from our target scenario is shown in Figure 2. Each *Attribute* has an acceptable value specified, and the *Operations*

[Attributes]	[Operations]
service_type: location	
accuracy: 50m	<=
mobility: 2100s	>
trust: 0.1	>
age: 360s	<
[Rank]	
accuracy: 3	
mobility: 2	
trust: 1	
age: 1	

Fig. 2. Application profile

section defines whether it should be maximised or minimised to be improved. The *Rank* section defines the weighted importance that each attribute has to the application.

```

<REQUEST>
  <SOURCE name='aliceID' />
  <TYPE name='location' />
  <ATTRIBUTE name='accuracy' op='lessThanEqual' value='50m' />
  <ATTRIBUTE name='age' op='lessThan' value='360s' />
  <ATTRIBUTE name='mobility' op='greaterThan' value='2100s' />
</REQUEST>

```

Fig. 3. Service request example

A service request is shown in Figure 3 using a simple XML format. The requester (*aliceID*) is a unique identity generated by Alice’s device, for example, by hashing her public key (the corresponding private key would sign each message). The service attributes are then listed, each with their own values and operators, not including the locally stored measures (such as trust). This request will only match location providers with an accuracy of 50 meters or less, delivering information which is under 6 minutes old, and that should not move (e.g., get off the train) for at least 35 minutes.

The response (see Figure 4) should confirm the service type and Bob’s advertised service attributes. As argued before, Bob’s claims may not reflect the quality that he can/will actually offer. Rather than simply believing the claims of another peer, we thus include trust reasoning about the service providers.

```
<RESPONSE>
  <SOURCE name='bobID' />
  <DESTINATION name='aliceID' />
  <TYPE name='location' />
  <ATTRIBUTE name='accuracy' value='30m' />
  <ATTRIBUTE name='age' value='100s' />
  <ATTRIBUTE name='mobility' value='15000s' />
</RESPONSE>
```

Fig. 4. Service response example

3.2 Trust based Selection

Hosts build up an opinion about every other device/service they interact with; these opinions will gradually become more authoritative with each interaction among the same pair of hosts. Initially, peers that have not been encountered before will usually have a neutral reputation, neither positive nor negative. This value would be increased after successful interactions, while appropriately decreased following unsatisfactory service deliveries. We use a trust range of values between $[-1 = \textit{HighlyDistrusted}, +1 = \textit{HighlyTrusted}]$. Rather than just change in a linear fashion, trust should increase with exponentially decaying increments, and decrease multiplicatively. A reasonably low default trust value would combat susceptibility to Sybil attacks, where malicious hosts can simply generate more identities to avoid being punished for past misbehaviours. An application will use the trust in a service provider to estimate the accuracy of the promised QoS; moreover, it may specify that only providers possessing a trust value higher than a certain (application-specific) threshold should be considered.

Trust values are not only useful to clients to rate services, they can also be used by the providers to prevent their misuse. Attempting to provide service to all peers who ask, can leave the providers open to DOS attacks. Ignoring peers with a poor reputation will provide an incentive for all peers to behave correctly. Old or unimportant opinions should be discarded according to certain host specific constraints.

Devices should be able to share their opinions about other hosts. This is particularly important to retrieve information about providers we have never interacted with before. Shared opinions (i.e., recommendations) should only be trusted as much as the host announcing them. A peer that provides a good opinion of another, only for that peer to act maliciously, should suffer a reduction in trustworthiness. This should discourage groups of malicious peers from colluding and giving false-positive opinions. Additionally, opinions from hosts that are distrusted should be ignored, to avoid malicious peers spreading lies. The importance assigned to recommendations should be much lower than a peer's own experience; particularly suspicious hosts may even decide to ignore others' opinions (thus assigning a zero weight to recommendations).

More powerful devices could also maintain a transaction history (promised and delivered service level) to share, rather than the subjective numeric value.

This would allow hosts to make judgments based on their own criteria. Transaction repositories could also be run by certain trusted parties such as the train operator, to increase the passenger's satisfaction, or by a completely separate third-party for profit.

4 Conclusion

This paper suggests the combination of trust, mobility and QoS estimations to provide a more reliable and rewarding pervasive service experience in mobile ad-hoc networks. The decentralised trust management model allows the dynamic calibration of the service selection, based on a history of service provisions; this should in turn promote co-operative behaviours among the various peers. The combination of mobility patterns with QoS information gives insightful details on the foreseen service availability. Predicting a host's mobility is an extremely hard task, even with knowledge of position, time and historical movement patterns. An effective lightweight metric needs to be devised to allow communication of expected future movements, a subject of further work.

We now intend to formalise the syntax and semantics of application profiles and service requests. We will then realise our framework in a component-based architecture and implement it in a lightweight middleware. It is our plan to then evaluate our ideas in terms of increased reliability in ad-hoc service provisioning, by means of simulation of realistic scenarios, which include real mobility traces.

References

1. T. Berners-Lee. Uniform Resource Identifier (URI): Generic Syntax, Jan 2005.
2. S. Buchegger and J.-Y. L. Boudec. Performance analysis of the CONFIDANT protocol. In *Proceedings of MobiHOC*, Lausanne, June 2002. IEEE.
3. L. Capra, S. Zachariadis, and C. Mascolo. Q-CAD: QoS and Context Aware Discovery Framework for Mobile Systems. *ICPS*, July 2005.
4. R. Chakravorty, S. Agarwal, S. Banerjee, and I. Pratt. MoB: A Mobile Bazaar for Wide-Area Wireless Services. In *MobiCom 2005*, pages 228–242. ACM Press, 2005.
5. P. Dewan, P. Dasgupta, and A. Bhattacharya. On using reputations in ad hoc networks to counter malicious nodes. In *ICPADS*, pages 665–, 2004.
6. S. Kamvar, M. T. Schlosser, and H. Garcia-Molina. The Eigentrust algorithm for reputation management in p2p networks. 2002.
7. S. R. Medidi and K. H. Vik. QuaSAR: QoS-Aware Source Initiated Ad-hoc Routing.
8. W3C: Web Ontology Working Group. OWL Web Ontology Lang Ref, 2004.