Discovering Service Dependencies in Mobile Ad Hoc Networks

Petr Novotny,[†] Bong Jun Ko,[‡] and Alexander L. Wolf [†]

[†]Imperial College London
London, UK

[‡]IBM T.J. Watson Research Center
Hawthorne, New York, USA

ABSTRACT

Abstract

The combination of service-oriented applications, with their run-time service binding, and mobile ad hoc networks, with their transient communication topologies, brings a new level of complex dynamism to the structure and behavior of software systems. This complexity challenges our ability to understand the dependence relationships among system components when performing analyses such as fault localization and impact analysis. Current methods of dynamic dependence discovery, developed for use in fixed networks, assume that dependencies change slowly. Moreover, they require relatively long monitoring periods as well as substantial memory and communication resources, which are impractical in the mobile ad hoc network environment. We describe a new method, designed specifically for this environment, that allows the engineer to trade accuracy against cost, yielding dynamic snapshots of dependence relationships. Through extensive simulations, we evaluate the performance of our method in terms of the accuracy of the discovered dependencies, and draw insights on the selection of critical parameters under various operational conditions.

1 Introduction

Understanding the dependencies among the components of a distributed system is critical to making good operational and maintenance decisions. For example, fault localization and change impact analysis are tasks enabled by accurate and timely data on component dependencies. The importance of dependence information increases with the complexity of the system, both in terms of the number of interacting components required to carry out a given computation and the nature of the environment in which the system operates.

In service-based systems, such as those based on the Service-Oriented Architecture (SOA) or Web Services frameworks, computations are structured as a set of services that respond to requests, where a request typically originates at a user-facing client. The computation required to fulfil each request results in a cascade of further requests across some subset of the services; services make requests on other services in order to fulfil the requests made upon them. Obtaining dependence information in such a system is made difficult by the inherent loose coupling of services, as many dependencies are unknown at design time, and only established at run time through a dynamic service binding mechanism (so-called "service discovery"). The consequence is that the dependencies among run-time instances of services are not something that can be reliably specified before execution, but instead must be discovered during or after execution.

Existing dependence discovery methods focus on statically structured systems operated in fixed networks [1, 2, 3, 4, 5, 7, 8, 10, 11, 12]. A critical assumption made by these methods is that the dependence data, although changing, is relatively stable over time. The significance of the stability assumption is that the methods can make use of statistical techniques based on data collected over long execution periods. Furthermore, by operating in the context of a fixed-network environment, the methods can assume no practical limits on the storage, computational, and communication resources needed to support those statistical techniques.

The context for our work is instead service-based systems deployed on mobile ad hoc networks (MANETs). Mobility and ad hoc networking bring increased dynamicity to service dependencies, beyond those caused by the basic service-binding regime of SOA or Web Services. Moreover, the MANET environment is typically characterized by severe limitations on the resources available for dependence discovery. Existing methods based on the stability assumption cannot adequately cope with such high levels of dynamicity nor stringent resource constraints.

We have formulated a relatively simple method that is nonetheless more suitable for services operating in the challenging MANET environment. Our intuition is that dependence discovery must be focused on capturing snapshots of dependence data relevant to each service request of concern, rather than on the tradition of determining statistical averages for long-term, system-wide dependencies as a whole. Furthermore, the method must be lightweight in its resource usage, which to our thinking means that dependence data should be collected locally, aggregated locally, and drawn to some central location only when and if needed.

Our approach is based on the use of *monitors* deployed onto the mobile hosts. The monitors collect dependence data by observing the message traffic between services and extracting relevant information. The data collected by the monitors provide only a local view of the dependence information. When a more global picture of the dependence relationships among the services is required (e.g., to carry out some particular analysis), the monitors are contacted by a central discovery element charged with integrating the data. Importantly, only the monitors relevant to a particular analysis question typically need to be contacted, and therefore communication can be reduced. Moreover, the monitors can aggregate the data they collect, and can impose limits on the amount of data they store.

In order to employ our method, three basic prerequisites must be met. First, to obtain complete dependence information, the monitors should be deployed on the mobile hosts that are either the source or the target of service messages; intermediate hosts in the network used only to store and forward network-level messages are not involved in data collection. Second, the monitors need access to synchronized clocks to allow consistent time-stamping of the collected dependence data. Clock synchronization in MANETs is a well-researched topic, with techniques available to achieve precision of tens or even single microseconds [13]. The shortest time period we use to time stamp aggregated data is 6 milliseconds, well within this precision. Third, the monitors must be able to observe service messages and obtain certain information from those messages, including such things as client and service identifiers. On the other hand, there is no need for the

monitors to have access to the payload of messages. This kind of general information is typically available and visible, since it is used by the underlying service infrastructure to manage service interactions.

Clearly, our method is subject to inaccuracies due to such effects as the delay between data collection and data analysis, storage constraints that might require monitors to "forget" some data, and failures in hosts and links. The essence of the present work is to understand how these effects impact the accuracy of dependence discovery, subject to various tuning and environmental parameters. In the longer-term, and beyond the scope of this paper, we are also studying how inaccuracies in dependence information can be accommodated in higher-level analyses such as fault localization.

In this paper we introduce our dependence discovery method and evaluate its sensitivity to a distinguishing aspect of the MANET environment, namely *time-dependent behavior*. The evaluation is carried out through a series of simulation-based experiments under various scenarios that represent a range of mobile-network dynamics, service connectivity and dynamics, and critical parameter settings. Because no benchmarks yet exist for MANET-hosted service-based systems, we develop synthetic data to explore the space of independent variables. The dependent variables in these experiments are the *true positives* and *false positives* in the discovered dependence relationships.

We next review prior and related work. We then introduce our discovery method and its underlying dependence model. This is followed by a presentation of the experiments used to reveal and evaluate the time-dependent behavioral properties of the method. We conclude with a summary of our observations and a look at our on-going and future work.

2   Background and Related Work

Existing dependence discovery methods can be generally classified as to whether they operate at the network level or at the (application) service level. Network-level discovery [1, 2, 3, 4, 8, 10] focuses on coarse-grained dependencies between network hosts. Network-level dependencies are usually described in terms of IP addresses and port numbers. They can be augmented with additional information, such as port mappings [7] or a classification of client applications [11]. On the other hand, service-level discovery [5, 6] focuses on the detection of fine-grained relationships between services. Services are hosted in application containers, such as J2EE and .NET, and typically associated with application identifiers, such as URLs.

Our method, although informed by the network level, operates at the service level in the sense that we wish to discover service dependencies that can be used for fault localization and impact analysis in service-based software systems. It is important to point out that we do not assume the availability of prior information, such as a port mapping, application categories, or even a specification of the services, nor do we require changes in existing software components to support discovery as found in other methods [4, 6].

Many of the service-level discovery methods apply statistical techniques to traffic traces. The statistical techniques correlate network packets or service-level messages and identify co-occurrences of messages across different services. While the particular statistical techniques may differ (e.g., correlation based on a time window [1, 2, 3], delay distribution [7], time difference of messages [5], or timing and frequency of packet flows [8]), all of the methods share the same limitation: they require a long period of time to collect statistically stable data and, therefore, are inappropriate in highly dynamic environments such as MANETs.

Alternatives to statistical approaches do exist. For example, in Macroscope [11] a subset of packets and network connections is sampled and analyzed to identify relationships between network flow data and applications. Lu et al. [10] collect system log data and correlate the events in the logs using data-mining techniques. Magpie [4] instead correlates the events based on input from (human) network operators. These methods, however, require the transfer of large amounts of trace data from the collection points to a central analysis element, which is prohibitive in resource-constrained MANETs. In contrast, our method transfers dependence information selectively and on demand, contacting only the monitors that are potentially relevant to the events of interest (e.g., the possible receivers of a failed service request). In addition, the monitors actively summarize and aggregate dependence information before that information is transferred (e.g., they may transfer only the number of messages exchanged between two services, rather than sending the content of those messages).

Pinpoint [6] uses a technique for correlating messages that is similar to the one found in our method. The technique makes use of identifiers to uncover flows of end-to-end processing through servers in a fixed network. Pinpoint requires application servers to insert the identifiers into the headers of messages sent over extensible protocols such HTTP or SOAP. Monitors deployed in the network record requests into traces from which the paths are reconstructed, post hoc, and further aggregated into dependencies. Their technique is a general one, not restricted to service-based systems.

Similarly, we make use of what are called *conversation identifiers* in service-based systems. In service-based systems, a conversation is the set of all messages exchanged during the processing of a single client request. Typically, the conversation is identified across the messages using tags inserted into message headers; the tags contain a unique conversation identifier. Of course, conversation identifiers must be implemented by service programmers. Fortunately, this is a relatively routine task, as there are several existing standards for doing so, including WS-Addressing,[1] WS-SecureConversation,[2] and WS-Coordination,[3] that are increasingly used in current-day applications.

3   Dependence Discovery and Representation

In this section we present our dependence discovery method. We begin by describing the two types of dependencies in which we are interested. We then describe how we discover dependencies, represent dependencies, and construct the representation.

3.1   Service Dependencies

In service-based systems, a *dependence* is a relation between services defined by the message flow induced by a client request. (As an edge case, a dependence is also the relation between a client and a service. Without loss of generality, we mainly focus here on relations among services.) When a dependence relation exists between two services S1 and S2, one service is considered the *source* and the other the *target*. In general, sources issue requests on targets, thus defining a directionality to the dependence.

We are concerned with two types of dependencies over a given set of services: *inter-dependencies* and *intra-dependencies*. An inter-dependence is the basic dependence relation that exists between the requester of a service and the receiver of that request. Figure 1 illustrates a simple set of inter-dependencies, where the arrows indicate the directionality of the dependencies, from sources to targets. For instance, service S1 is directly dependent upon services S3 and S4, and indirectly dependent upon services S6, S7, and S8. Each inter-dependence in which a particular service is engaged can be classified as either *incoming* or *outgoing*. Service S3 has an incoming inter-dependence with S1 and outgoing inter-dependencies with S6 and S7.

An intra-dependence is a more complex relation between services that relates an incoming inter-dependence to an outgoing inter-dependence. In this sense, intra-dependencies reflect more detailed insight into the nature of the dependencies between services than do the basic inter-dependencies. This is illustrated in Figure 2, which shows the dependencies among S1, S2, S4, S7, and S8 resulting from the four given inter-dependencies (solid arrows) and the three given intra-dependencies (dashed arrows). It is instructive to compare the information gained from Figure 2 to that available in Figure 1. We can see that S1 is (indirectly) dependent upon S7 and S8, while S2 is only (indirectly) dependent upon S8. This is not evident from Figure 1.

3.2   Discovering Dependencies

As mentioned above, dependencies arise from the flow of messages among services. To discover dependencies, we must therefore track these flows. Because our aim is to be minimally intrusive, we restrict ourselves to observing the message traffic (i.e., messages that contain service requests and responses) as it occurs. Our method makes use of *monitors* deployed within the network to observe messages and record information

---

[1] http://www.w3.org/2002/ws/addr/
[2] http://www.ibm.com/developerworks/library/specification/ws-secon/
[3] http://www.ibm.com/developerworks/library/specification/ws-tx/

Figure 1: Inter-dependencies of Client.

Figure 2: Intra-dependencies of service S4.

about the flows. A convenient place to deploy a monitor is within a service's container. The monitor is then easily aware of the associated service's identity, as well as being provided a context in which to execute.

The main advantage of an approach based on monitors is that it allows us to discover dependencies instantaneously and precisely, with minimal delays between dependence occurrence, detection, and the availability of the dependence information. Moreover, we can do so without having to modify the services themselves. Monitors can also minimize data storage and communication requirements, since they can actively aggregate and summarize the information. Thus, our approach can be thought of as a process for collecting evidence of dependencies, which is in sharp contrast to methods that require storage and transfer of large amounts of data for later statistical analysis.

Inter-Dependence Discovery.   Pairs of source and target services that induce inter-dependencies can be identified from the flow of messages exchanged between the services. In a service-based system, services are uniquely identified by application-specific identifiers, such as the URIs of the Web Services framework. Although the specific type of information provided within messages differs with the service platform and standard used, request messages always contain identifiers for the requested services.

Since a monitor is aware of the identity of the service with which it shares a container, it can record outgoing inter-dependencies simply by extracting the identifiers of target services from any outgoing request messages originating at the service. The target service identifier is an essential field present in all request messages, such as plain HTTP or SOAP requests. Discovering incoming inter-dependencies is a bit more involved, as it requires request messages to contain the unique identifier of the requesting service. This is satisfied by most service standards (e.g., the WS-Addressing standard provides the fields `wsa:To` and `wsa:From`). With source and target identifiers present in messages, monitors can detect all incoming and outgoing inter-dependencies.

Intra-Dependence Discovery.   To expose the correspondence between the incoming and outgoing inter-dependencies of a service, we rely on the presence of *conversation identifiers* within messages, as discussed at the end of Section 2. Discovering an intra-dependence then reduces to having a monitor relate incoming and outgoing messages using the conversation identifiers appearing in both.

The aggregation of observed service interactions into dependencies can cause the monitors to "forget" some of the details necessary to reconstruct accurate dependence information. This is the case when internal behaviors of a service that are not visible to the monitor can generate seemingly identical interactions in aggregate. As an example, consider the incoming inter-dependence of S1 on S4, and the outgoing inter-dependencies on S7 and S8, shown in Figure 2. The incoming inter-dependence could in fact be the result of aggregating two separate conversations between S1 and S4, where the first resulted in a message to S7 and the second resulted in a message to S8. The aggregated dependence information held by the monitor will

4

Figure 3: Architecture of dependence discovery.

not, in our current approach, record the true dependencies regarding these individual conversations.

3.3   Dependence Graph

A dependence graph (DG) is a directed acyclic graph constructed from a set of nodes representing services and a set of edges representing direct inter-dependencies. The direction of an edge represents the direction of the inter-dependence, from source to target. Each node can be annotated with intra-dependence information, conceptually adding directed edges between the incoming and outgoing inter-dependencies of the service.

The DG maintains information concerning a specific *time window*, reflecting only the dependence information collected by (or perhaps available from) monitors during that period. The time window is a property of the interaction between the application behavior, the network behavior, and the information accessible to monitors. The size of the time window has many effects on the results of analysis. For example, a small time window serves to reduce the size of the DG, but some critical service interactions might be missed. A large time window provides a more complete record of dependencies, but might include stale or irrelevant interactions (e.g., those belong to conversations other than the target conversation for the analysis).

Conceptually, a DG could be used to represent the full set of dependencies of an entire application system. In practice, many analysis techniques only require a subgraph of the full dependence graph related to a specific node or subset of nodes. For example, a failure impact analysis might examine only the nodes that can reach (i.e., are dependent upon) a given node, and a fault localization analysis might examine only the nodes that are reachable from a given node.

3.4   Dependence Graph Construction

Our method uses a set of distributed monitors that provide information to a central dependence discovery element, as illustrated in Figure 3. The intent of this architecture is to minimize resource utilization, while still providing timely data. The monitors perform continuous dependence discovery and maintain aggregated dependence data.

Dependence data are stored with a sliding expiration window such that only a limited history is maintained, using a data structure representing sliding time slots. Entries for each time slot maintain Boolean data about whether or not a given dependence occurred within that time slot. Of course, the size of the time slot affects the precision of the data maintained. For example, a slot size of 0.1 seconds will provide up to 10 times more data compared to a slot size of 1 second, but will require 10 times more space to store those data. Beyond the slot size, the size of the whole history can be controlled through the pruning of expired time slots.

The central discovery element will construct a DG on demand, querying the relevant monitors to harvest their local dependence data for the time window of interest. The harvesting algorithm is designed to incrementally construct the DG—typically a subgraph of the full application dependence graph—by visiting only the monitors considered relevant based on the data seen to that point. In this way, the amount of data transmitted over the network can be significantly reduced compared to existing methods. Of course,

the most common case of DG construction is for a particular client, revealing the services upon which that client depends directly or indirectly. Conceptually, the data are harvested by a walk rooted at the monitor associated with the client. For certain analyses, we may additionally want to limit the dependence discovery to a specific conversation, something particularly useful in fault localization. The concrete design of the harvesting algorithm is relatively straightforward, so we do not detail it here.

4 A Sample Application

Dependence information forms the basis for a variety of important system-level analyses. As an example, we have experience in using the dependence information obtained from the discovery method presented here to develop a new *fault localization* aid for the operators of service-based systems deployed on MANETs. Describing this experience in detail is beyond the scope of the present paper. However, for purposes of illustration we give a sketch of its basic design.

Fault localization in general refers to a technique for identifying the likely root causes of failures observed in systems formed from components. Fault localization in systems deployed on MANETs is a particularly challenging task because these systems are subject to a wider variety and higher incidence of faults than those deployed in fixed networks, the resources available to track faults are severely limited, and many of the sources of faults in MANETs are by their nature transient (e.g., a mobile host that temporarily moves out of radio range). Further complicating the situation is that faults at the network level may not even manifest themselves as failures at the service level, since some of those faults may occur outside of the time during which the relevant services are communicating, or involve communication outside of the relevant conversation. Moreover, service-based systems are typically designed to tolerate certain kinds of faults through mechanisms for dynamic service (re)binding. Therefore, fault localization must be especially adept at sifting through noisy data.

Our technique uses service-level dependence relationships as the basis for focusing the analysis to a first approximation. We envision being given a failure report in terms of a failed client conversation. Then, as described in the previous section, the data needed to construct the dependence graph are collected by using the client as the root of a walk. The technique also makes use of *fault symptoms* recorded in logs, as well as a *fault propagation pattern* built from a classification of the failures that can be experienced by a client. The pattern indicates how the symptoms can be propagated through the system from root causes to clients. The specific pattern associated with a given failed conversation is then mapped onto the discovered dependence graph corresponding to the conversation, resulting in a *fault propagation model* (FPM) that represents how specific faults propagate through the individual services involved in the conversation. The FPM is combined with the occurrence times of the relevant symptoms. This assignment allows the technique to reconstruct the possible fault propagation paths, both in time and in space, and thus to form a set of plausible hypotheses about the causes of the failure. The hypotheses are represented as a *causality graph*.

After the causality graph is constructed, a ranking algorithm is applied to the faults in the hypothesis set based on their likelihood of being the root cause of the failure. We have devised three alternative approaches to ranking. The first is a timing-based approach that ranks hypotheses based on the time difference between possible root causes and the failure. The second is a probabilistic method that uses a Bayesian network to infer independent probabilities of individual root-cause hypotheses. We obtain the Bayesian network by combining the FPM with fault propagation rules that represent conditional probability distribution parameters. The algorithm infers the posterior probability of a hypothesis being a root cause given the evidence of the failure. The third approach is a hybrid of the first and the second. Here we successively apply the Bayesian inference method to the causality graphs obtained for multiple, incremental time windows preceding the time of the failure; the inferred probabilities of the hypothesis for each time window size are adjusted so that higher probabilities are assigned to those calculated from shorter time windows. Our results show that the third approach consistently provides the best results.

Figure 4: Experimental framework.

5 Experimental Setup

The evaluation of our dependence discovery method is based on a simulation framework for service-based systems operated in MANETs. As shown in Figure 4, the framework first generates message and fault traces according to various simulation parameters and scenario configurations. These traces are stored in a database to hold the results of the simulation runs. Dependence discovery is prototyped as a query over this database. In addition, we derive the "ground truth" used in Section 6 to evaluate our method.

Without loss of generality, we make the assumption that faults occur only at the network level, and that the services themselves are error free. From an observational point of view, any non-Byzantine service-level fault can be simulated by a corresponding network-level fault (e.g., whether the absence of a response message is due to a service fault or a network fault is indistinguishable to a service waiting for that response). Therefore, this assumption does not materially impact our results, but does allows us to concentrate on simulating the effect of the network on service interactions.

The simulation engine is built on top of the discrete-event network simulator NS-3,[4] which provides a comprehensive network simulation using detailed implementations of low-level network protocols. However, the highest abstraction provided by NS-3 is that of sockets and network packets. We therefore extended NS-3 with higher abstractions for simulating service entities and their interactions. We did this by first encapsulating the socket layer into a *messaging layer* that provides the abstraction of service-level messages. The messaging layer itself is then encapsulated into a *service layer* that provides abstractions for services, clients, and their service-oriented interactions. Finally, we provide a means to configure and run the extended simulation.

Service-level scenarios are created using several basic models and associated parameters.

Entity models: Entity models provide the building blocks of service-based systems: clients and services. Clients behave autonomously, making requests on specific services at configurable times (see the workload model, below). Services provide methods that can be requested by clients or other services. Each method contains a specification of its behavior, consisting of delays to simulate computations, and a set of steps for making further requests on other services.

Deployment model: The deployment model specifies where in the network the clients and services are hosted.

Interconnection model: The interconnection model defines the (dynamic) binding of clients to the methods of services, and of each method to the methods of other services.

Message model: Three types of messages are exchanged between entities: requests, responses, and exceptions. Requests are used to invoke methods in other services, while responses are sent by services back to the requester upon completion of the requested processing. Exceptions are used to propagate fault symptoms caused by failures. As described in Section 3, the flow of messages related to an original client request is called a *conversation*. In the simulation, all messages contain a conversation identifier, possibly available to the dependence discovery element (according to what we assume about the discovery process), but always available to establish the ground truth for dependence relationships.

Workload model: The workload is determined by the rates at which clients send requests to services. Currently, our simulator implements workloads in which clients, at uniformly random times, request a method chosen uniformly at random from its set of currently bound methods.

---

[4]http://www.nsnam.org/

Fault propagation model: Two mechanisms are used to propagate a fault from its root cause back to the client that initiated the failed conversation: exceptions and timeouts. Exceptions are explicit messages sent by services to their requesters to notify them of a failure to fulfil a request. Typically, they engender a cascade of such messages. Timeouts are implicit indications of faults detected by the requesters upon the lack of a response from services. Timeouts engender an implicit propagation cascade.

The simulation scenarios are constructed by a configuration generator. NS-3 defines the available network parameters, while the models above provide the service parameters.

6   Evaluation

In our method, dependence graphs are constructed on demand by a discovery element. The graphs are rooted at a given client, beginning at a given time instant, and for some time window. The data provided to the discovery element include both relevant and irrelevant information, since any given monitor will provide data about *all* interactions traversing its associated service during the time window. These interactions involve not just those of the given conversation of interest, but also those of others.

Under such circumstances it would be difficult for a dependence discovery method to provide a perfect result. Moreover, the method by design loses information (e.g., monitors retain only aggregate data, not individual messages) and is sensitive to the dynamics of the operational environment. Thus, the evaluation questions of interest center mainly on the accuracy of the resulting dependence graph. In particular, we examine four factors that should influence accuracy in this setting: degree of service connectivity, dynamics of service interconnection, network mobility, and time window size. Notice that the first two are application properties, the third an operating environment property, and the fourth a tuning parameter for the method. We also compare our method to other methods, as well as examine the data storage and data transfer requirements of our method.

6.1   Evaluation Metrics and Parameters

Our experiments focus on a particular hypothetical conversation C. A good result for our method would be that it can discover as many dependencies of C as possible, while not including the dependencies of other conversations. We use two metrics to characterize the quality of our results, namely the ratio of true positives (TP) and the ratio of false positives (FP), defined as follows:

$$TP\ ratio = \frac{|D(C) \bigcap GT(C)|}{|GT(C)|}$$

$$FP\ ratio = \frac{|D(C) - GT(C)|}{|D(C)|}$$

where $D(C)$ is the set of discovered dependencies, $GT(C)$ is the set of ground-truth dependencies, true positives are in the intersection of these two sets, and false positives are in the set difference. The TP ratio represents the fraction of discovered dependencies as compared to the actual dependencies in the conversation, and is therefore equivalent to the *recall* metric of information retrieval (taking "ground truth" as "relevance"). The FP ratio represents the fraction of discovered dependencies not belonging to C, and is therefore the complement of *precision*. We assume $D(C)$ and $GT(C)$ are non-empty.

Table 1 summarizes the basic network-layer parameters used in our simulations. These are standard settings used widely in the networking community and embodied in the NS-3 simulator. Specifically, we use the log distance model with path-loss exponent 3 for wireless signal propagation, reproducing a network operated in an urban area [9]. We set the spatial mobility bounds to a 75 meter square, as larger regions induce long-term network partitioning. Another important parameter is the mobility speed of the mobile hosts. For most of the experiments we set the mobility speed of all nodes to 10 m/s. This is a challenging

8

Number of nodes	50
Spatial bounds	75m x 75m
Mobility speed	10 m/s
Mobility model	RandomDirection2dMobility
Propagation delay model	ConstantSpeedPropagationDelay
Propagation loss model	LogDistancePropagationLoss/α3
WiFi standard	80211b
WiFi rate	11Mbps
Routing protocol	OLSR and Ipv4StaticRouting
Protocol stack	UDP/IPv4

Table 1: Network-layer parameters.

Number of clients	50 (one per node)
Number of services	30
Invokable methods per service	2
Workload (client request rate)	10s
Service dependence switch rate	5 minutes – 15 minutes
Message size	500 bytes – 1500 bytes
Response timeout	60s

Table 2: Service-layer parameters.

	Low	Medium	High
Inter-service, method-to-method connectivity probability	0.0125	0.025	0.05

Table 3: Service-connectivity parameter.

	Low	Medium	High
Client-to-service dependencies	250	250	250
Service-to-service dependencies	30	59	130
Dependencies in graph (avg.)	2.03	3.21	7.9
Dependencies in graph (std. dev.)	1.2	2.28	3.63

Table 4: Dependencies induced by connectivity.

scenario that allows us to simulate environments in which message exchanges between the nodes are materially affected by the temporary disruptions in communication at the network layer. In Section 6.4 we report results in which we vary this parameter from 0 m/s (i.e., a "stationary" wireless network) to 15 m/s.

Table 2 summarizes the basic service-level parameters used in our simulations. The message sizes and timeouts are derived from standard values found in SOA and Web Services implementations. Note that the number of methods in each service is not significant from a simulation point of view, as long as we have at least two methods available so that we can examine the impact of intra-dependence information on dependence discovery.

For the interconnection model, we use a 2-tiered topology. The first tier consists of the connections between the clients and a set of "front end" services, while the second tier consists of the connections between the services themselves. In our simulations, we use 50 clients, five front-end services, and 25 "back end" services. When starting a conversation, each client invokes a method selected uniformly at random from all methods provided by the five front-end services. We experimented with other topologies, including a single-tier topology in which there are no designated front-end services, and with different numbers of clients and services, but found that the results were consistent. We therefore only report results based on the 2-tier topology.

Of particular importance is the degree of connectivity among the services. Table 3 shows how we configure the simulations to capture three different connectivity scenarios, denoted as "Low", "Medium", and "High". The connectivity degree is induced by the probability that a method in one service invokes a method in another service; the higher the probability, the denser the interconnection topology.

The effect of different connectivities on the resulting ground-truth dependencies for each scenario is shown in Table 4. Because the 50 clients will always invoke the five front-end services over the course of an experimental run, there will be 250 client-to-service dependencies in all scenarios. However, the total number of unique service-to-service dependencies increases with the connectivity probability, from 30 dependencies for low connectivity to 130 dependencies for high connectivity. Also increasing with the connectivity probability is the resulting number of dependencies per conversation (i.e., the number of arcs appearing in the ground-truth dependence graph rooted at a given client). The average and standard deviation for the conversations occurring in the three scenarios are given in the table. Notice that a high standard deviation in this case indicates a good mix of different kinds of conversations.

An artifact of having more services involved in a conversation due to a higher degree of connectivity is that it increases the time it takes to complete the conversation. This is borne out by Figure 5, which gives the cumulative distribution of the conversation lengths in our simulations for the three connectivity scenarios. For example, 80% of the conversations in the low connectivity scenario are completed in under 0.2 seconds, whereas the top 20% of the conversations in the high connectivity scenario take longer than 4 seconds.

Service dependencies may change not only due to normal computational progress, but also to optimize
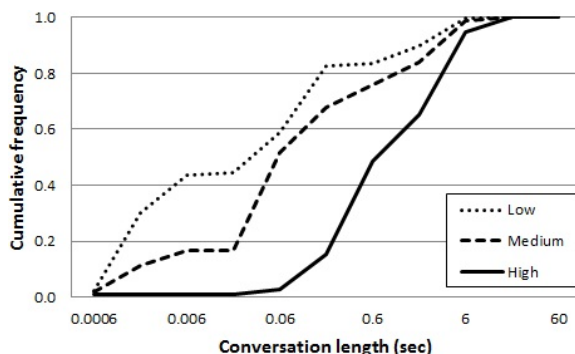

Figure 5: Cumulative distribution of conversation lengths (based on request messages) for three connectivity scenarios.

10

Figure 6: Accuracy of inter- and intra-dependence discovery methods for different time window sizes given as (a) TP ratios, (b) FP ratios, and (c) a trade off between TP and FP ratios.

the use of otherwise equivalent services. To see the impact of this form of dynamic service rebinding, we introduce into our experiments a *service-switching* behavior. Every service is given two designated alternative services, and every dependent service then switches among these three services after a randomly selected time period. In Section 6.3 we report results in which we vary the switching-time parameter in increments between 10 seconds on average (representing an extremely dynamic behavior) and infinite (representing a non-switching behavior).

We collect our results from 30 minutes of simulated execution time after excluding 30 seconds of warm up. Each combination of parameters in our experiments results in thousands of conversations occurring during the simulated 30-minute execution. For instance, the low, medium, and high connectivites combined with 10-minute switching-time periods and 10 m/s mobility speeds result in 8440, 8139, and 7165 conversations, respectively. The results given below are averages over the data collected from these conversations, where each conversation is then a statistical sample subject to the random variables.

6.2 Impact of Time Window Size and Service Connectivity

We first look at the impact of the time window size and connectivity degree on the accuracy of the results. We hypothesize that as the time window size grows, so too should the TP ratio, since more dependencies will be captured. However, increasing the time window size should also increase the FP ratio, since there is a greater chance that messages belonging to other conversations are included in the dependence graph. For a given time window size, we expect the TP ratio to be negatively correlated with the connectivity degree, since a higher connectivity increases the conversation length, which in turn increases the chances that some dependencies are missed. Similarly, we would expect the FP ratio to be higher in densely connected service configurations, since dependencies in other conversations are more likely to overlap those of the subject conversation.

We calculate the TP and FP ratios for both inter- and intra-dependence discovery separately. Figure 6 depicts the results, where each data point is the ratio averaged over all conversations. Here we consider only fixed service bindings (i.e., non-switching behavior). The variances of the TP and FP ratios are small, and therefore omitted from the figures. For example, the largest 95-percentile confidence intervals for TP and FP ratios in the medium connectivity scenario are 0.006 and 0.0053, respectively.

As shown in Figure 6a, increasing the time window size increases the TP ratio, both for inter- and intra-dependencies. However, increasing the time window size also increases the FP ratio, as shown in Figure 6b. The same figures also confirm our hypotheses about the impact of the connectivity degree: the TP ratio decreases and the FP ratio increases as the service topology becomes denser. Notice, too, that intra-dependence discovery has a significantly lower FP ratio than inter-dependence discovery. This is due to the fact that it can precisely correlate incoming and outgoing inter-dependencies, something to which inter-dependence discovery is blind (recall Figure 2). To directly display the trade off between the TP and FP ratios under various time window sizes, we plot them against each other in Figure 6c, where each point represents the given time window size.

6.3 Impact of Dynamic Service Rebinding

We now investigate the sensitivity of the method to dynamic service rebinding. For a given time window size, we would expect a high switching rate to cause many dependencies belonging to irrelevant conversations to be included in the discovered dependence graph, thereby increasing the FP ratio. In contrast, we would expect the TP ratio to be insensitive to the switching rate because once a message is seen, any additional messages, relevant or irrelevant, should not increase that ratio.

We ran experiments using seven different average switching-time periods, from 10 seconds up to a case in which no switching occurs; a lower value results in a faster switching rate. Figure 7 shows the intra-dependence FP ratio in the medium connectivity scenario. (As hypothesized, the TP ratio is essentially unaffected by the switching rate, so we do not show that result here.) The effect on FP ratio is most noticeable when the switching-time period is on the same order as the time window size, which indicates possibly many service switches happen within the time window, rendering the resulting dependencies irrelevant from the discovery element's point of view.
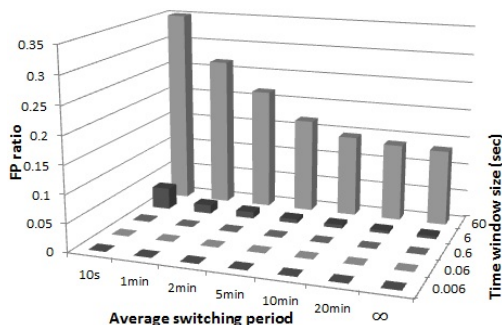

Figure 7: FP ratio for intra-dependence discovery under medium connectivity for various switching-time periods and time window sizes.

6.4 Impact of the Mobility Speed

The next set of experiments investigate the impact of the mobility speed on the inter-dependence discovery TP and FP ratios. We work with four different mobility speeds: 0 m/s (amounting to a fixed network), 5 m/s, 10 m/s, and 15 m/s.

We would expect that as the speed of the mobile hosts increases, the quality of the links between them should deteriorate and, consequently, the failure rate of message exchanges should increase. Higher message failure rates should in turn increase conversation lengths, as more messages must be resent in an attempt to complete the conversations. Therefore, increasing the mobility speed should have the effect of decreasing the TP ratio, since messages require longer time periods to be exchanged and detected. This, however, should not impact the FP ratio, since the proportions of relevant and irrelevant messages remain unaffected.

The results are reported in Figure 8. As expected, the TP ratio generally has a positive correlation with mobility speed, with the fixed network (0 m/s) exhibiting the highest ratio. Furthermore, we observe that the lengths of the conversations increase with mobility speed (and network failure rate), which is visible in the TP ratios between the 0.006s to 0.6s time windows. However, for longer conversations, the time windows of 6s and 60s are long enough to capture all cases. The FP ratios are virtually unaffected. This coincides with our understanding that the ratio of relevant and irrelevant messages remains the same.

6.5 Comparison with Existing Methods

We now compare the accuracy of our dependence discovery method to that of existing methods. We make this comparison by implementing two alternative methods to represent the two major classes of existing
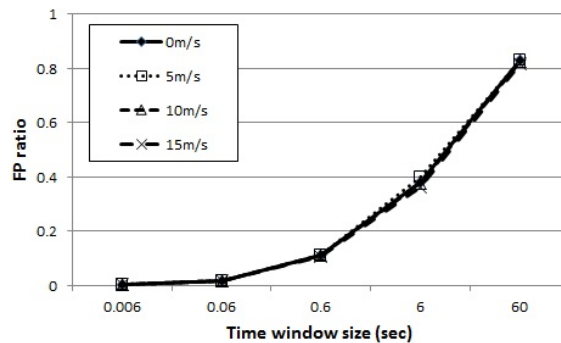
Figure 8: TP ratios (a) and FP ratios (b) for inter-dependence discovery under four mobility speeds.

approaches: those that perform discovery at the service level [5, 6] and those that perform discovery at the network level [1, 2, 3, 4, 8, 10]. These implementations are, like our own method, prototyped as queries over the trace database (see Section 5). The service-level alternative discovers a global system dependence graph by observing all the service messages exchanged over the whole execution period and, from this, builds dependence graphs for the individual client conversations. The network-level alternative works similarly, but only observes the flow of messages by inspecting the information contained in the headers of packets exchanged over the relevant IP ports. It then builds dependence graphs using external information provided to it about the deployment of clients and services on hosts.

We compare our method against the two alternatives using the medium connectivity scenario, a 10 minute service-switching time period, and a 10 m/s mobility speed. The average number of ground-truth dependencies is 3.21 (see Table 4). We configure our method to use a 60s time window size, which is large enough to capture all such dependencies (see Figure 5). The comparison then reduces to one based on the false dependencies appearing in the discovered dependence graphs.

As we discuss in Section 2, the existing service- and network-level methods are designed for use in fixed networks and for relatively stable service configurations. Therefore, since both these alternative methods build dependence graphs from long-term observations, we do not expect them to adequately filter out stale dependencies caused by the dynamics of the scenarios, resulting in higher false positives than with our discovery method. Moreover, the network-level method should include even more false positives than the service-level method because it builds dependence graphs from coarser-grained information.

The results are reported in Figure 9, where a vertical line is used to separate the results for our method on the left from the results for the alternative methods on the right. We give the FP ratio, as well as a count of the false dependencies appearing in the dependence graph. Both are computed as the average over the total number of conversations (8139) occurring in the 30-minute execution period. As expected, our method provides dependence graphs having a significantly lower FP ratio than the alternative methods.

13

Furthermore, although the FP ratio for our inter-dependence discovery is similar to that of the alternative methods, the actual number of false dependencies is significantly lower.
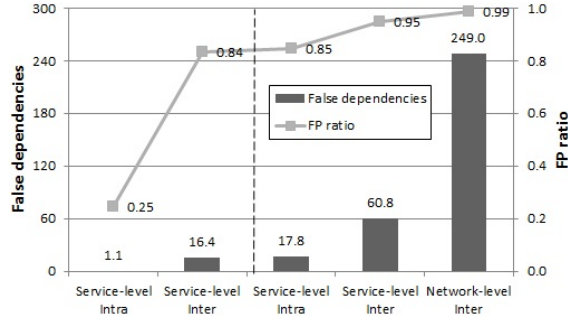

Figure 9: Comparison of methods in medium connectivity scenario. Results for existing methods are to right of dashed line.

6.6 Estimates of Data Storage and Data Transfer Needs

We conclude our evaluation by investigating the data storage requirements placed on the monitors, and the data transfer requirements placed on the network. We do this by positing a representation for the data, analyzing the space requirements of that representation, and computing the total for a particular experimental scenario.

A monitor will maintain the history of the dependencies it has seen for some bounded period of time, divided into time slots. Each dependence is associated with a set of those time slots, such that when the monitor detects the occurrence of a dependence, it signifies this by setting a 1-bit flag in the corresponding time slot. It also records identifying information about the source and target of the dependence.

The length of each time slot can be configured according to a desired level of resolution. When combined with the time period and the size of each dependence entry, the data storage needs for each monitor is determined. We estimate the required bytes of storage space S per monitor as follows:

$$S = \frac{T_h}{T_s} \times \frac{(N_{inter} + N_{intra})}{8} + (B_{inter} N_{inter} + B_{intra} N_{intra})$$

where T_h is the length in seconds of the time period, T_s is the length in seconds of a time slot, N_{inter} and N_{intra} are the average number of inter-dependencies (both incoming and outgoing) and intra-dependencies recorded by an monitor during the time period, and B_{inter} and B_{intra} are the maximum sizes of the identifiers of inter- and intra-dependencies.

Consider the medium connectivity scenario, where the average number of inter- and intra-dependencies per service is 12 and 53, respectively. Assume a maximum size for identifiers of 256 single-byte characters, which corresponds to the Web Services framework's use of the maximum URL length. With a time slot length of 0.01 seconds (an especially precise configuration) maintained for a time period of 10 minutes, the required data storage will amount to approximately 506 kilobytes per monitor.

The estimates above are in some sense worst case. In practice the identifiers will be smaller (certainly smaller than the maximum allowed URL length). Moreover, the storage used for intra-dependence identifiers can be significantly reduced, simply by maintaining them as references to the corresponding incoming and outgoing inter-dependencies.

While the monitors locally store data describing all dependencies detected during the time period, the data transferred over the network to the discovery element contains only the data for the desired time window. This time window is typically much shorter than the stored time period.

Consider again the medium connectivity scenario. The typical data bundle transferred by an individual monitor contains data for approximately one outgoing inter-dependence and one intra-dependence. This

14

amounts to a maximum of 512 bytes of data per service. The average number of services involved in a client request is 3.21, so the discovery element will issue on average 4.21 requests (including the initial request to the client's monitor) for dependence data in order to construct the dependence graph. This results in a total of approximately 1.6 kilobytes of data transferred over the network, which is substantially less than the total amount of dependence data stored by the monitors.

6.7 Discussion of Results

The experiments presented above establish how the degree of service connectivity, dynamics of service interconnection, network mobility, and time window size affect the accuracy of our dependence discovery method. Connectivity, interconnection dynamics, and mobility speed are properties of the application and network, while time window size is an operator's tuning parameter. This represents a broad exploration of the behavioral characteristics of our method.

In addition to the results reported here, we have experimented with several other parameters, including the client workload and rates of service faults. However, these parameters do not seem to have a significant impact on the accuracy of the method.

In general, the results for intra- and inter-dependence discovery indicate consistently equivalent accuracy as measured in terms of TP ratio. On the other hand, the FP ratio is significantly lower for intra-dependence discovery compared to that for inter-dependence discovery. Both the TP and FP ratios are significantly affected by the selection of the time window size: as the time window size increases, so do the TP and FP ratios. The TP ratio, however, tends to reach a maximum value at a certain time window size, whereas the FP ratio degrades (i.e., increases) unabated beyond that value. This implies that the time window can be selected such that it will yield a discovery result striking a good balance between true positives and false positives.

We can also see that the FP ratio is significantly affected by the service-switching time period. This factor depends on the service discovery and selection technique used in the system. In general, however, as TP ratio is relatively insensitive to the service interconnection dynamics, one can select the time window size after taking into account the service switching rate, which in a MANET is typically related to the dynamics of the network.

In contrast to time window size, the accuracy of the dependence results appears to be less sensitive to network mobility, assuming the service interconnection dynamics are at a given level. This implies that our method is relatively robust to variations in operational conditions.

The trade off between TP and FP ratios leads us to consider approaches based on *adaptive time windows*, which would necessitate a small amount of additional information (e.g., the completion time of a conversation) to be retained and communicated by the monitors. Presentation of these approaches is beyond the scope of this paper, suffice it to say that our initial results have been very promising.

The comparison with existing methods validates our hypothesis that MANET-hosted service-based systems require a fundamentally different approach to dependence discovery. Particularly striking is the difference in FP ratio between our method's intra-dependence and the other methods. While our method provides approximately one false dependence on average per dependence graph, the existing approaches yield dependence graphs significantly larger and constructed mostly from false dependencies. This is especially evident in the network-level approach, which yields dependence graphs containing almost whole the entire set of services, rendering dependence discovery virtually useless.

Finally, our estimates for the method's storage requirements fall well within the capabilities of today's mobile hosts. Moreover, the method's aggregation of data results in an extremely low requirement for data transport.

Threats to Experimental Validity

The results reported are a selection from the experiments we have conducted. This selection focuses on particular aspects of the method, as detailed above. The results withheld support the results reported.

The threats to the validity of the results derive from the three prerequisites listed in Section 1 and the experimental parameter values detailed in Tables 1–3. We have chosen these values based on the joint experiences of the NS-3 community of users, and on the recommendations of commercial SOA and Web Services providers. Where specific community experience was lacking in the choice of values (e.g., the absence of a benchmark) we attempted to broadly sample in the space of values. We made use of uniformly random distributions in several instances (e.g., the request behavior of clients), which are a simplification that is a methodologically accepted practice at this stage of investigation. This yields statistically sound

results, as each data point in our results is an average over all arising conversations, each of which is a statistical sample subject to the random variables.

We make no claim that the results generalize beyond the reported experiments, but overall they give us confidence that the method is a viable approach to service dependence discovery in the challenging MANET environment.

7   Conclusion

We have presented a run-time method to discover the dependencies among services operated in the highly dynamic and resource-constrained environment of MANETs. Unlike existing approaches, the method does not require stable dependence relationships, nor does it require that large amounts of evidence data be collected over long periods. Through an extensive set of simulation-based experiments, we have evaluated the accuracy of the method in terms of operational factors characteristic of both service-based systems and MANETs. The method exhibits good behavior when subjected to the stress of a changing underlying network topology. Furthermore, its data storage and data transfer requirements scale well with the number and connectivity of the services involved.

Dependence information is not particularly useful in and of itself, but instead serves as an indispensable building block for important analysis capabilities. We are currently developing such analyses, including those for probabilistic fault localization, as briefly sketched in Section 4, and cross-layer performance anomaly diagnosis.

Acknowledgments

References

[1] P. Bahl, P. Barham, R. Black, R. Ch, M. Goldszmidt, R. Isaacs, S. K, L. Li, J. Maccormick, D. A. Maltz, R. Mortier, M. Wawrzoniak, and M. Zhang. Discovering dependencies for network management. In *Fifth Workshop on Hot Topics in Networks*, Nov. 2006.

[2] P. Bahl, R. Chandra, A. Greenberg, S. Kandula, D. A. Maltz, and M. Zhang. Towards highly reliable enterprise network services via inference of multi-level dependencies. In *Proceedings of the Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, pages 13–24. ACM, August 2007.

[3] P. Barham, R. Black, M. Goldszmidt, R. Isaacs, J. MacCormick, R. Mortier, and A. Simma. Constellation: Atomated discovery of service and host dependencies in networked systems. Technical Report MSR-TR-2008-67, Microsoft Research, 2008.

[4] P. Barham, A. Donnelly, R. Isaacs, and R. Mortier. Using Magpie for request extraction and workload modelling. In *Proceedings of the 6th USENIX Symposium on Operating Systems Design and Implementation*. USENIX Association, 2004.

[5] S. Basu, F. Casati, and F. Daniel. Toward web service dependency discovery for SOA management. In *Proceedings of the IEEE International Conference on Services Computing*, pages 422–429. IEEE Computer Society, 2008.

[6] M. Y. Chen, A. Accardi, E. Kiciman, J. Lloyd, D. Patterson, A. Fox, and E. Brewer. Path-based faliure and evolution management. In *Proceedings of the Symposium on Networked Systems Design and Implementation*. USENIX Association, 2004.

[7] X. Chen, M. Zhang, Z. M. Mao, and P. Bahl. Automating network application dependency discovery: Experiences, limitations, and new solutions. In *Proceedings of the 8th USENIX Symposium on Operating Systems Design and Implementation*, pages 117–130. USENIX Association, 2008.

[8] D. Dechouniotis, X. Dimitropoulos, A. Kind, and S. Denazis. Dependency detection using a fuzzy engine. In *Proceedings of the 18th IFIP/IEEE International Workshop on Distributed Systems: Operations and Management*, number 4785 in Lecture Notes in Computer Science, pages 110–121. Springer-Verlag, 2007.

[9] I. K. Eltahir. The impact of different radio propagation models for mobile ad hoc networks (MANET) in urban area environment. In *2nd International Conference on Wireless Broadband and Ultra Wideband Communications*. IEEE, Aug. 2007.

[10] J.-G. Lou, Q. Fu, Y. Wang, and J. Li. Mining dependency in distributed systems through unstructured logs analysis. *SIGOPS Operating Systems Review*, 44:91–96, March 2010.

[11] L. Popa, B.-G. Chun, I. Stoica, J. Chandrashekar, and N. Taft. Macroscope: End-point approach to networked application dependency discovery. In *Proceedings of the 5th International Conference on Emerging Networking Experiments and Technologies*, pages 229–240. ACM, 2009.

[12] S. Wang and M. A. M. Capretz. A dependency impact analysis model for web services evolution. In *Proceedings of the IEEE International Conference on Web Services*, pages 359–365. IEEE Computer Society, 2009.

[13] D. Zhou and T.-H. Lai. An accurate and scalable clock synchronization protocol for IEEE 802.11-based multihop ad hoc networks. *IEEE Transactions on Parallel and Distributed Systems*, 18(12):1797–1808, Dec. 2007.