

Context Discovery Using Attenuated Bloom Filters in Ad-Hoc Networks

Fei Liu and Geert Heijenk

University of Twente, P.O. Box 217,
7500 AE Enschede, The Netherlands
{fei.liu, geert.heijenk}@utwente.nl

Abstract. A novel approach to performing context discovery in ad-hoc networks based on the use of attenuated Bloom filters is proposed in this paper. In order to investigate the performance of this approach, a model has been developed. This document describes the model and its validation. The model has been implemented in Matlab, and some results are also shown in this document. Attenuated Bloom filters appear to be a very promising approach for context discovery in ad hoc networks.

1 Introduction

Ad-hoc networks are non-infrastructure wireless networks in which most of the terminals are both mobile and power-consumption constrained. When one needs to obtain a service or context information from other devices, querying and fixing the location of the service or context information source might generate a lot of traffic. In a network with a high query rate, such traffic can be rather heavy. As a result, the terminals consume quite an amount of power and bandwidth for querying. An efficient context discovery mechanism needs to be developed for such situations.

This paper describes the development of a discovery mechanism for networks that are context aware. These networks utilize context information to improve their operation, or to enrich the services provided to users. We propose a novel approach to discover context information sources in an ad-hoc network based on the use of attenuated Bloom filters, which represents a decentralized space-efficient discovery method. Instead of broadcasting full information about the type and location of context information, nodes send attenuated Bloom filters which contain context type information for all the reachable nodes up to a certain number of hops away. Moreover, Bloom filters have a special feature of false positive probability, which leads to probabilistic querying. Queries are only forwarded in the directions which possibly contain the required information. Our analysis reveals that this type of probabilistic discovery method can substantially reduce the network load compared to discovery using traditional approaches.

This paper is structured as follows. Section 2 will introduce our novel approach for performing context discovery using attenuated Bloom filters. In Section 3 we will describe an approximate model for the transmission costs of our method. Further we compare it with context discovery without attenuated Bloom filters. In Section 4 we will provide numerical results. Finally in Section 5, we present our conclusions and proposed future work.

2 Context Discovery Using Attenuated Bloom Filters

2.1 Related Work on Context Discovery

Context discovery has a lot of resemblance to service discovery. Service Discovery Protocols (SDPs) can be classified into centralized and decentralized architectures. In an Ad-hoc network environment, nodes are both mobile and mostly battery-powered. Those characteristics fit well with some features of decentralized architectures. The choice for a proactive or reactive SDP in decentralized architectures depends substantially on the network and service context and on the interaction with the underlying routing protocol [1].

Service descriptions are different for various SDPs. The most popular format is the attribute-value structure [2]. For instance, the Service Location Protocol (SLP) [3] uses service templates which predefine the attributes in a template document readable by humans and machines. Service agents (SAs) advertise the location of one or more services; directory agents (DAs) store service location information centrally. Whenever necessary, user agents will look for the required services at SAs and DAs. A client/server structure is used in the Bluetooth Service Discovery Protocol. Bluetooth SDP [4] defines a service record consisting of the entire list of attributes, which is then stored in the SDP server. Clients will send requests to the SDP server to obtain the required services.

Further, hierarchical attribute-value pairs, which mostly rely on eXtensible Markup language (XML), are also used in some protocols, such as Global Service Discovery Architecture (GloSev) [5] and Group-based Service Discovery protocol (GSD) [6]. GloSev is proposed for worldwide and local area network usage. Services are described and categorized hierarchically by using the Resource Description Framework (RDF) which is based on Uniform Resource Identifiers (URI) and XML. This hierarchical service architecture is similar to the DNS domain name architecture. GSD is a distributed service discovery protocol for Mobile Ad hoc NETWORKS (MANETs). Services are described based on DARPA Agent Markup Language (DAML+OIL). Advertisements are sent periodically to nodes within a maximum number of hops. Each node has peer-to-peer caching to keep a list of local and remote services that a node has received from advertisements. Services are also grouped to ease service discovery by selectively forwarding queries.

In some protocols, such as Jini [7], attributes are described as Java objects. Service objects are registered in service registries, which are also used to look up services. A client needs to download the service object and invoke it to access the service.

Among the protocols mentioned above, GloSev was developed for wide area networks. SLP and Jini were designed for local area networks. GSD and Bluetooth SDP were specifically for MANETs. Further, it is clear that whichever method is used to describe services, sending the complete service attributes causes heavy traffic. It is inefficient in a high-density mobile ad hoc network with many services to be advertised and/or queried. Due to the limited battery power of terminals, a simple, efficient context description and discovery mechanism is required. Clearly, a mobile ad hoc network is less suitable for a centralized structure due to the mobility of the nodes. Nodes should not depend on other specific nodes to reach the required context information. Context discovery using attenuated Bloom filters can solve this problem in ad hoc networks.

2.2 Brief Introduction of Attenuated Bloom Filters

A Bloom filter [8] is a data structure for representing a set in order to support membership queries. It can denote a set simply and efficiently, with a small probability of false positives. Bloom filters can be used in various network applications, such as distributed caching, P2P/overlay networks, resource routing, packet routing, and measurement infrastructure [9]. Bloom filters were also proposed to be used as an efficient approach for lossy aggregation and query routing for a Secure Service Discovery Service in [10]. Recently, researchers have explored the applications of Bloom filters to Ad hoc networks, such as speeding-up cache lookups [11], group management [12], and hotspot-based trace back [13].

A Bloom code can represent a set of context information types. Each context type will be coded by using b independent hash functions over the range $\{1 \dots w\}$, where w is the width of the filter. The default value for each bit in the Bloom code is 0. The bits of positions associated with the hashes will be set to 1. Our approach uses attenuated Bloom filters, each of which consists of a few layers of basic Bloom filters. The first layer of the filter contains the context type information for the current node, while the second layer contains the information about the nodes one hop away, and so on. In other words, a node can find the context type information i hops away in the i^{th} layer. When querying for a certain type of context information, the same hash functions are performed. If all positions in a Bloom filter indicated by one of the hashes contain a 1, the presence of the queried context type is likely (but not certain). Otherwise the context type is not present. The use of these attenuated Bloom filters introduces the possibility of having false positives, which will be resolved during a later stage of the context discovery process. By using attenuated Bloom filter consisting of multiple layers, context sources at more than one hop distance can be discovered, while avoiding saturation of the Bloom filter by attenuating (shifting out) bits set by sources further away.

For example, we assume a 6-bit Bloom filter with b equal to 2. If location information is hashed into $\{1, 3\}$ and temperature information is mapped into $\{2, 5\}$, we obtain the filter shown in Fig. 1.

The filter will give a positive answer to queries for location or temperature information. It definitely does not contain presence information which is hashed into $\{0, 3\}$. Nodes may also think humidity information $\{1, 5\}$ is contained in this filter, but actually it is not. This situation is termed false positive [9].

0	1	2	3	4	5
0	1	1	1	0	1

Fig. 1. A simple 6-bit Bloom filter

Context aggregation can be simply implemented by attenuated Bloom filters. When a node A receives incoming Bloom filters $filter_B$ and $filter_C$ from neighbors B and C respectively, it shifts all the contents of $filter_B$ and $filter_C$ one layer down and

discards the last layer. The first layers will be filled with 0s. An OR operation will be done to those new filters, $filter_B'$ and $filter_C'$, and the first layer will be filled by (first layer of) $filter_A$. Consequently, the Bloom filter of node A is updated such that the first layer represents the local information from node A; the second layer contains the information from neighbor B and C; the third layer covers the information two hops away which can be reached via B or C. Fig. 2 shows the process of context aggregation in a node.

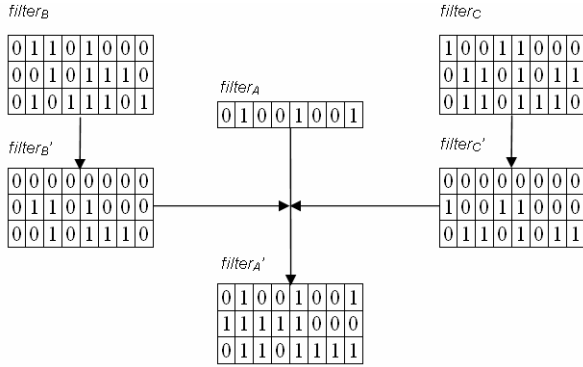


Fig. 2. Context aggregation

2.3 Context Discovery with Attenuated Bloom Filters

Context discovery by using attenuated Bloom filters is a method combining proactive and reactive discovery mechanisms. Nodes can obtain an overview of available context types and their distribution by exchanging Bloom filters. When a node enters a new environment, it first sends a Bloom filter with its own context information types to its neighbors. When neighbors receive this filter, they will merge it to their existing attenuated Bloom filters by shifting the incoming filter and doing an OR operation, as explained in Section 2.2, and rebroadcast the updated filter. The new node stores the incoming Bloom filters separately for each neighbor, and also generates a new filter aggregating all the reply filters with the local one. This new filter will be sent to the neighbors in the next advertisement. The Bloom filters will be exchanged periodically, or when they change.

When there is a query, the node will first check locally whether the required context information exists. If not, the query will be encoded into a Bloom code using the same hash functions and compared with the locally stored filters. If there is a match at any layer in any of the local filters, a query message containing the Bloom code will be sent to the neighbor from which the stored filter was received. The same action will be taken by a node receiving a query. When a node receives the same query again, it will drop the query. Further the path of the query message will be recorded by each node. If an exact match is finally found, a Context-Available (CAL) reply will be sent back along the same path. If no match is found, e.g.,

because of a false positive match in an earlier node, the query will be discarded. In this way, query messages will be filtered out as early as possible, depending on the stored attenuated Bloom filters. The originating node will set up a connection to the destination node based on the best CAL reply. Note that a hop counter is used to restrict the query range. Queries will only be sent a limited number of hops away, based on the depth of the Bloom filters. If no CAL is received by the originating node after a time out, it understands the required information is unavailable in the current range. It will make a choice to enlarge the discovery area or send the query again after a certain time.

Note that recording the path of the query message requires maintaining (soft) state routing information in the nodes. This can put a burden on these nodes. Alternatively, the return path for CAL replies can be stored in the query messages. This will result in increased transmission costs. The detailed network architecture is described in [14].

3 Performance Modeling

Queries due to false positives can potentially contribute significantly to the costs of context discovery. In order to reduce the number of unnecessary queries, we have to reduce the rate of false positives. However, the minimum false positive rate, which is by definition equal to 0, will result in large Bloom filters. We believe there is a balance to be struck between a reasonable false positive rate and the size of Bloom filters to achieve the optimal network cost. The size of Bloom filters is decided by their width, depth, the number of hash functions, and the cardinality of the represented set. The goal of the model is to find the optimum balance between those parameters to minimize the network cost. Further comparisons between context discovery with and without Bloom filters will be made.

3.1 Assumptions and Related Vital Parameters

Ideally, the two-dimensional radio coverage of a mobile terminal is a circle. Therefore, a circular structured network is assumed in our modeling. The communication range of nodes will be generalized as a set of concentric circles. Nodes located in the inner circle are the ones reachable within one hop from the center node. Nodes in one ring outside the inner circle are the ones reachable in two hops, and so on.

Each node in the network is supposed to have the same number of context (or service) information types, s . Further all context types are supposed to be unique, and taken out of an infinitely large set of possible context types. The same width, w , and depth, d , of attenuated Bloom filters, and the same b hash functions are used for the entire network. Queries are forwarded at most d hops, based on the depth of Bloom filters.

We assume that the communication range, r , of a node is 30 meters. The density of the network is n nodes/m² in average. Assuming 4 neighbors per node, we obtain $n \approx 0.0007$ [14].

General notation is listed in the Table 1.

Table 1. Notation

General		Bloom Filter	
Notation	Description	Notation	Description
s	number of services (context information types) per node	w	the width of the filter
μ	advertisement (update) rate	d	the depth of the filter
λ	query rate	b	number of hash functions
n	network density (nodes/m ²)		
r	communication range		

3.2 Model

There are two kinds of traffic in the network: advertisements and queries. Three types of advertisement messages can be identified: normal advertisement, updates, and maintenance. We assume that advertisements are broadcast periodically at a constant rate. Among queries, there are also two types based on the different answers: positive query and false positive query. Note that there is no false negative in Bloom filters. Therefore, the cost for a node is defined as the sum of the advertisement cost, positive cost and false positive cost. We determine the cost considering the transmission cost of a single node in the network. The cost is expressed in bits per second per node:

$$cost = adcost + pcost + fpcost \quad (1)$$

Since advertisements are broadcasted periodically, $adcost$ can be derived by:

$$adcost = \mu \cdot adpack \quad (2)$$

where μ is the advertisement (update) rate; $adpack$ is the advertisement packet size.

The positive cost is denoted as $pcost$. To simplify the problem, we assume that queried services are not available in the network, which indicates all queries will result in a negative answer, or a false positive, i.e. $pcost = 0$.

Then $fpcost$ represents the false positive cost of a node. In order to be able to determine these costs, let us first determine the probability of generating a false positive as a result of a query. We define $P_{fp,j}$ as the probability of a false positive for layer j ($1 \leq j \leq d$) of an attenuated Bloom filter, where x_j represents the number of services available in layer j . In this paper, we assume that the hash functions we choose are perfectly random. The probability that a specific bit is 0 is equal to:

$$\left(1 - \frac{1}{w}\right)^{bx_j} \approx e^{-bx_j/w} \quad (3)$$

So we have:

$$P_{fp,j} = \left(1 - \left(1 - \frac{1}{w}\right)^{bx_j}\right)^b \approx \left(1 - e^{-bx_j/w}\right)^b \quad (4)$$

Formula (4) shows that the false positive probability depends on the width (w) and number hash functions (b) of the Bloom filters, and the number of services contained inside the filter. We define the number of services in j^{th} hop as:

$$x_j = s \cdot n\pi(j \cdot r)^2 \quad (5)$$

The assumption is that all services that are reachable in j or fewer hops are represented at the j^{th} layer. A service that is represented at a certain layer is also assumed to be represented at all layers below, because of likely alternative paths with more hops to the node containing that service, which will impact the broadcast attenuated Bloom filters.

We have to consider all transmission costs for false positive queries incurred by a query initiated in the node under consideration. Transmission of such query messages can take place on all links up to d hops away from the node under consideration. Thus, we can denote the false positive cost as:

$$fpcost = \lambda \cdot \sum_{i=1}^d cost_{fp,i} \quad (6)$$

We assume queries are performed with a certain rate λ , i.e., λ queries will be initiated per second per node. $cost_{fp,i}$ denotes the total cost of all false positive queries transmitted to nodes i hops away from the node under consideration.

In order to obtain this false positive query cost to the i^{th} hop, we have to count the possible number of query transmissions sent by nodes $i-1$ hops away to their neighbors, $numofTransmission_{fp,i}$. Such a transmission is indeed done, with a packet size $qpack$, if the attenuated Bloom filter received from the intended receiver of the query gives a false positive in any of the layers 1 to $d-i+1$. A false positive in a layer beyond $d-i+1$ does not result in a query message being transmitted, because that would lead to a query being transmitted more than d hops from the originating node. Note that false positives at multiple layers of the Bloom filter will result in multiple transmissions being counted for the relevant link. We are still investigating if this can be avoided in our system, but also assume that the false positive probability is so small, that we can neglect this effect. The resulting false positive query cost to the i^{th} hop can be given as:

$$cost_{fp,i} = \sum_{j=1}^{d-i+1} P_{fp,j} \cdot numofTransmission_{fp,i} \cdot qpack \quad (7)$$

For determining $numofTransmission_{fp,i}$, let us suppose that each node can reach all the nodes within communication range, which are $(n\pi r^2 - 1)$ nodes excluding the querying node itself. When a node forwards a query to the next hop, there are potentially $(n\pi r^2 - 1)$ neighbors to transmit the query to. These transmissions to the i^{th} hop will potentially be done by all nodes that can be reached within $i-1$ hops from the node under consideration, excluding those that can also be reached within $i-2$ hops. Therefore, there will be $(n\pi(i-1)^2 r^2 - n\pi(i-2)^2 r^2) = (2i-3)n\pi r^2$ ($i > 1$) nodes involved in transmission. These results in the following potential number of transmis-

sion sent in the i^{th} hop. Note that this also includes transmissions to nodes that have already received the query before. These duplicates will be discarded upon reception.

$$\text{numofTransmission}_{fp,i} = (2i-3)n\pi^2 \cdot (n\pi^2 - 1) \quad 2 \leq i \leq d \quad (8)$$

For the first hop, the original querying node will send the query to all the nodes in range, so that the number of transmissions sent in 1st hop equals:

$$\text{numofTransmission}_{fp,1} = n\pi^2 - 1 \quad (9)$$

So the false positive query cost at hop i is:

$$\text{cost}_{fp,i} = \begin{cases} \sum_{j=1}^d (1 - e^{-bx_j/w})^b \cdot (n\pi^2 - 1) \cdot \text{qpack} & i = 1 \\ \sum_{j=1}^{d-i+1} (1 - e^{-bx_j/w})^b \cdot (2i-3)n\pi^2 \cdot (n\pi^2 - 1) \cdot \text{qpack} & 2 \leq i \leq d \end{cases} \quad (10)$$

For the sizes of the advertisement and query packets, we assume that the context discovery protocol using Bloom filters is running on top of UDP. For both advertisements and queries, besides the header of Bloom Filters protocol, the headers of UDP, IP, and MAC layer will be attached. The advertisements and queries packet size are defined as follows:

$$\text{adpack} = \text{header}_{MAC} + \text{header}_{IP} + \text{header}_{UDP} + \text{header}_{AD} + w \times d \quad (11)$$

$$\text{qpack} = \text{header}_{MAC} + \text{header}_{IP} + \text{header}_{UDP} + \text{header}_Q + w \quad (12)$$

3.3 Two Extreme Cases

To evaluate the performance of context discovery using attenuated Bloom filters, we have compared it with two alternative discovery solutions: complete advertisement and non advertisement.

Complete advertisement floods all network nodes within d hops with complete descriptions of all context information types. Nodes had the complete map of the network, which indicates how nodes can send queries directly to the destination. It is a proactive protocol. The advertisement cost is the main concern in this situation. We assume that each context information type can be presented in c bits, so we have:

$$\text{cost} = \mu \times \text{numofTrans} \times (\text{header}_{MAC} + \text{header}_{IP} + \text{header}_{UDP} + \text{header}_{AD} + s \times c) \quad (13)$$

where numofTrans denotes the number of transmissions for the entire advertisement within d hops. We suppose each node up to $d-1$ hops away to broadcast the advertisement, so that we have:

$$\text{numofTrans} = \left(1 + n\pi((d-1) \cdot r)^2\right) \quad (14)$$

In the non-advertisement case, nodes do not advertise context information types. When a query comes, nodes forward it to all the neighbors. It is a reactive protocol.

Nodes do not have any idea about the network. The queries are spreading around the whole network, up to d hops from the originator. There is no way to stop forwarding queries, even though the query node has already received an answer. The cost for querying is counted as the cost for sending queries to the network, which is:

$$cost = \lambda(header_{MAC} + header_{IP} + header_{UDP} + header_Q + c)(1 + n\pi(d \cdot r)^2) \quad (15)$$

4 Experimental Results

The model described above has been implemented in Matlab 6.5. Using the model, four sets of experiments are done. The following tables and figures show the results of the experiments. In all the experiments below, we assume the update frequency $\mu = 0.1$ and each context information type can be represented in 32 bits, i.e., $c = 32$ bits. The sizes of headers are assumed as follows [15]: $header_{MAC} = 160$ bits; $header_{IPv6} = 320$ bits; $header_{UDP} = 64$ bits; $header_{AD} = 32$ bits [14]; $header_Q = 192$ bits [14].

There are 4 sets of experiments done in this section. Experiment 1 is used to achieve the optimal cost by choosing the proper width, w , of the Bloom filter and the number of hash functions, b , with given depth, d , of the filter, query rate, λ , and number of services, s , per node. Experiment 2 shows the influence of the query rate, λ , on the network cost for given d and s . The influence of the query range, d , is evaluated in Experiment 3. In the final experiment, we show the impact of density of services, s , in the network.

4.1 Experiment 1

In this experiment, we assume query rate $\lambda = 0.1$, $s = 1$. For each given value of depth of filter, the experiment result shows there exists a certain value of w and b which leads to the minimum network cost. The result is shown in Table 2. It is also compared with the complete and non advertisement under similar situations.

As we see from Table 2, for each depth of the filter, the proper width and number of hash functions leads to a minimum network cost which is much lower than for the cases of a complete advertisement and a non-advertisement. The difference becomes larger as query range d increases. The final column shows the maximum number of services that are covered by one Bloom filter based on the related size of Bloom filter.

Table 2. Optimal BF cost for certain depth d compared with complete and non advertisement

d	w (bit)	b	BF cost (bit/s)	Complete Advertisement (bit/s)	Non Advertisement (bit/s)	Maximum number of services in BF
3	128	5	114	547	1459	18
5	512	8	328	2006	3917	49
7	1024	8	809	4438	7603	97
10	2048	8	2231	9910	15437	198

4.2 Experiment 2

Using Bloom filters, we can reduce the packet size by using simple and efficient coding. However, false positives also create redundant traffic. We expect that there exists a point at which the traffic generated due to false positives is much more than the benefit of using Bloom filters. In contrast, if there are only few queries in the network, it does not pay to broadcast the context information to the entire network. A non-advertisement protocol can perform better in this case. This experiment is going to discuss the suitable range of using Bloom filters for context discovery to achieve the minimum network cost.

We set μ as a reference, and change the value of λ . Here we talk about λ/μ . The experiments show that the suitable range of λ/μ decreases when the depth of the filter, d , increases. When each node has only one service ($s = 1$), the Bloom filter context discovery algorithm performs better than the non-advertisement algorithm when λ is at least 0.1 times μ . When $d = 3$, the Bloom filter algorithm performs better than the complete advertisement algorithm even if λ is 10^8 times μ . When $d = 10$, this decreases to 10^4 times. Fig. 3a shows the situation when $d = 5$.

When each node has 4 services ($s = 4$), the network requires larger Bloom filters to contain more information. The results show that for $d = 3$, the proper range of λ/μ is (0.1, 100); for $d = 10$, the proper range of λ/μ is (0.8, 1.6). Fig. 3b shows the result when $d = 5$.

We found that in practical situations the Bloom filter algorithm has a better performance. Therefore, it is a promising algorithm for mobile ad hoc networks. Note that the axes in Fig. 3 are represented in log scale.

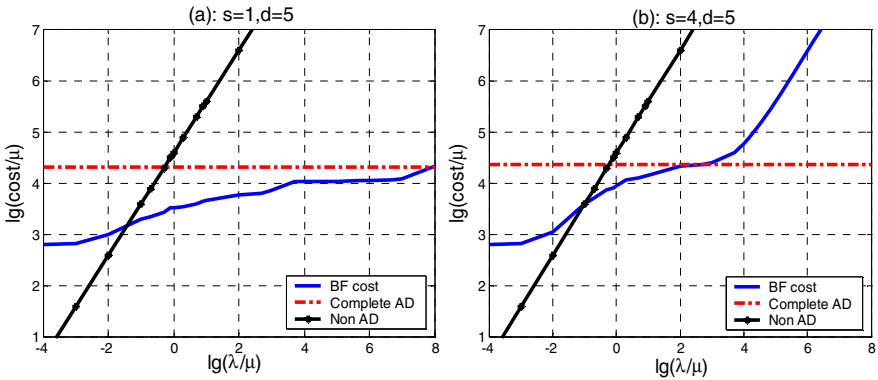


Fig. 3. Performance results for λ/μ when $s = 1$ (a) and $s = 4$ (b)

4.3 Experiment 3

With a larger search range (larger d), there are more context information types available within the range. On the other hand, a larger d also leads to larger Bloom filter. In this set of experiments, we would like to see the impact of d .

We set the depth of the Bloom filter, d , from 3 to 10, and compare the performance with different values of s and λ (fixed $\mu = 0.1$). The results show that, in general, the

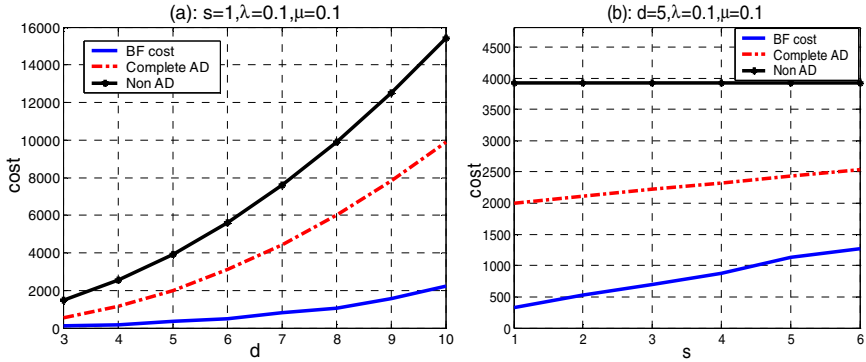


Fig. 4. when $\lambda = 0.1$: (a) impact from change of d ($s = 1$); (b) impact from change of s ($d = 5$)

Bloom filter algorithm has better performance than complete and non-advertisement algorithms. There is a limit to the number of services and the query rate for which the algorithm has the best performance. When exceeding that limit, the performance of Bloom filter becomes worse. When the number of services within the range (this depends on both d and s) and the query rate is quite high, the cost of using the Bloom filter algorithm increases significantly. For instance, this happens when $s = 4$, $\lambda = 1$, and $d > 7$. Fig. 4a shows the curve when $s = 1$, $\lambda = 0.1$.

4.4 Experiment 4

From the experiment above, we find that the number of services per node also has some influence on the network cost. In this set of experiments, we would like to investigate it in detail. We do this for fixed d and λ . The results show that s has some influence from s , but not much. When s increases from 1 to 6, i.e., the number of context sources increases from 0.0007 to 0.0042 per m^2 , a Bloom filter still gives the best result among three alternative algorithms. The network cost of using a Bloom filter increases only a little bit faster for than the complete advertisement algorithm. We can expect the Bloom filter algorithm to perform worse when s is really large, which will seldom happen in reality (for given d and λ). Fig. 4b shows the curve for three alternative algorithms when $d = 5$, $\lambda = 0.1$.

5 Conclusions and Further Work

The use of attenuated Bloom filters for advertising available context types in ad-hoc networks is very promising. Results obtained from the model presented in this paper reveal the combined cost of advertising and doing unsuccessful queries due to false positives. There exists a proper size of Bloom filters to achieve optimal network cost. The performance of Bloom filters also highly depends on the ratio of query and advertisement rates, and query range of nodes. Density of network context information sources also has some influences. For a fully distributed ad hoc network in practical

situations, this approach requires significantly less traffic load than advertising a full map of all available context types, or broadcasting queries when no advertisements are used. As such, it is a very promising compromise between these two extremes.

Future research should include further refinement of the model, e.g., finding a more thorough investigation of parameters and scenarios. A next step is to develop the idea further, by specifying a protocol, and testing this in a detailed, discrete event simulator and/or prototype. Security issues are also subject to future research. Finally, an interesting idea to explore is to use the broadcasting of attenuated Bloom filters to execute directed route requests (instead of undirected broadcasts) for ad hoc routing protocols such as AODV. Such an approach would allow ad-hoc nodes to establish routes only to other nodes with relevant context information, rather than establishing multiple routes first, and then finding out where the relevant context information is.

Acknowledgements

This work is part of the Freeband AWARENESS project (<http://awareness.freeband.nl>). Freeband is sponsored by the Dutch government under contract BSIK 03025. We would also like to thank Patrick Goering for his helpful comments.

References

1. Jeroen Hoebeke, Ingrid Moerman, Bart Dhoedt, Piet Demeester, "Analysis of Decentralized Resource and Service Discovery Mechanisms in Wireless Multi-hop Networks", *Third International Conference, Wired/Wireless Internet Communications 2005 Xanthi*, Greece, May 2005, Proceedings.
2. R. Marin-Perianu, P. H. Hartel, J. Scholten, "A Classification of Service Discovery Protocols", *Centre for Telematics and Information Technology, Univ. of Twente, The Netherlands*, Technical report nr. TR-CTIT-05-25, June 2005.
3. E. Guttman, C. Perkins, J. Veizades, M. Day, "Service Location Protocol version 2", *IETF, RFC 2608*, June 1999.
4. Bluetooth Consortium, "Specification of Bluetooth System Core Version 1.0b: Part e, Service Discovery Protocol (SDP)", November 1999.
5. Knarig Arabshian, Henning Schulzrinne, "GloServ: Global Service Discovery Architecture", *MobiQuitous*, pages 319-325, IEEE Computer Society, June 2004.
6. Diapanjan Chakraborty, Anupam Joshi, Tim Finin, Yelena Yesha, "GSD: a Novel Group-based Service Discovery Protocol for MANETs", *4th IEEE Conference on Mobile and Wireless Communication Networks (MWCN)*, September 2002.
7. Sun Microsystems, "Jini Architecture Specification Version 2.1", Nov 2005.
8. Burton H. Bloom, "Space/Time Trade-offs in Hash Coding with Allowable Errors", *Communications of the ACM 13(7)*: 422-426.
9. Andrei Broder, Michael Mitzenmacher, "Network Applications of Bloom Filters: A Survey", *Internet Math 1(2003)*, no.4, 485-509.
10. S. E. Czerwinski, B. Y. Zhao, T. D. Hodes, A. D. Joseph, R. H. Katz, "An Architecture for a Secure Service Discovery Service", In *Proc. Of MobiCom-99*, pages 24-35, N. Y., August 1999.

11. E. Papapetrou, E. Pitoura, and K. Lillis, "Speeding –up Cache Lookups in Wireless Ad-Hoc Routing Using Bloom Filters", *the 16th Annual International Symposium on Personal Indoor and Mobile Radio Communications (PIMRC 2005)*, Sept 11-13, 2005, Berlin, Germany, 2005.
12. J. Liu, F. SAILHAN, D. Sacchetti, V. Issarny. "Group Management for Mobile Ad hoc Networks: Design, Implementation and Experiment", in *Proceedings of the 6th IEEE International Conference on Mobile Data Management (MDM'2005)*, May 2005.
13. Yi-an Huang, Wenke Lee, "Hotspot-Based Traceback for Mobile Ad Hoc Networks"., in *Proceedings of the ACM Workshop on Wireless Security (WiSe'05)*, September 2005.
14. Fei Liu, Geert Heijenk, "Context Discovery Using Attenuated Bloom Codes: Model Description and Validation", Technical Report University of Twente, TR-CTIT-06-09, ISSN 1381-3625, March 2006.
15. James F. Kurose, Keith W. Ross, "Computer Networking: A Top-Down Approach Featuring the Internet", *Addison Wesley Longman, Inc.* ISBN 0 2-1 47711 4, 2001.