Modeling Epidemic Information Dissemination on Mobile Devices with Finite Buffers*

Christoph Lindemann and Oliver P. Waldhorst
University of Dortmund
Department of Computer Science
August-Schmidt-Str. 12
44227 Dortmund, Germany

http://mobicom.cs.uni-dortmund.de/

ABSTRACT

Epidemic algorithms have recently been proposed as an effective solution for disseminating information in large-scale peer-to-peer (P2P) systems and in mobile ad hoc networks (MANET). In this paper, we present a modeling approach for steady-state analysis of epidemic dissemination of information in MANET. As major contribution, the introduced approach explicitly represents the spread of multiple data items, finite buffer capacity at mobile devices and a least recently used buffer replacement scheme. Using the introduced modeling approach, we analyze seven degrees of separation (7DS) as one well-known approach for implementing P2P data sharing in a MANET using epidemic dissemination of information. A validation of results derived from the analytical model against simulation shows excellent agreement. Quantitative performance curves derived from the analytical model yield several insights for optimizing the system design of 7DS.

Categories and Subject Descriptors

C.4 [Performance of Systems]: *modeling techniques, design studies.* I.6.5 [Simulation and Modeling] Model Development - *modeling methodologies*. C.2.1 [Computer Communication Networks]: Network Architecture and Design - *distributed networks, wireless communication*.

General Terms

Algorithms, Design, Experimentation, Performance.

Keywords

Performance-oriented design and evaluation studies of distributed systems, mobile ad hoc networks, peer-to-peer data sharing, analytical performance modeling.

* This work was partially founded by the German Research Council (DFG) under Grant Li 645/13-1.

1. INTRODUCTION

Epidemic algorithms have recently been proposed as an effective solution for disseminating information in large-scale peer-to-peer (P2P) systems and in mobile ad hoc networks (MANET). Epidemic algorithms mimic the spread of a contagious disease. Like infected individuals pass on a virus to those with whom they get into contact, each mobile device in a MANET may relay data items to the other devices in its transmission range. Basically, every node of the MANET buffers messages it receives up to a certain buffer capacity and may forward such messages a limited number of times. A mobile device forwards a message to some deterministically determined other devices or to a randomly selected set of other devices of limited size. Important parameter of an epidemic algorithm constitute the number of devices to which the message is forwarded denoted as the fan-out of the dissemination, the number of times a message is forwarded, and the buffer size and management scheme at mobile nodes.

Infostations have been introduced as a communication model in which mobile users can connect to the network in the vicinity of their access points [7]. Infostations trade connectivity for capacity. Furthermore, Grossglauser and Tse demonstrated a capacity/mobility trade-off by showing that the average throughput per source-destination pair of nodes can be kept constant for increasing the density of nodes in a multihop wireless network with mobile devices [8]. Moreover, they presented a formal proof and simulation results showing that mobility increases the per-session throughput in a MANET for delay-insensitive applications. This observation exploits epidemic dissemination of information in MANET and has triggered a number of novel algorithms for MANET; e.g., for routing [9] and P2P data sharing [13], [14].

Mathematical models for the spread of epidemic diseases have been widely studied (see e.g., [1]). Typically, such models consider the spread of one disease (i.e., one data item). Such models analyze the transient behavior for determining the fraction of infected individuals (i.e., mobile nodes) among a finite population (i.e., the number of nodes of a MANET) at some time and the probability with which the entire population is infected after a given time (i.e., all nodes have received the data item) [6]. However, comprehensive modeling of epidemic information dissemination among nodes of a MANET (e.g., for implementing P2P data sharing) requires considering the long run behavior for system with large number of data items as well as finite buffer size with a buffer management scheme like least recently used (LRU).

In this paper, we present an analytical modeling approach for epidemic information dissemination (EID) systems in MANET. As major contribution, the introduced approach explicitly represents the spread of multiple data items, finite buffer capacity at mobile devices and an LRU buffer replacement scheme. The introduced modeling approach extends the approach of Dan and Towsley [5] for modeling LRU buffer management. A major difference between [5] and our modeling approach lies in that due to intermittent connectivity of a MANET, we have to distinguish between the local buffer and the buffers of other remote devices.

Papadopouli and Schulzrinne introduced seven degrees of separation (7DS), a system for P2P Web document sharing between mobile users based on epidemic information dissemination [13]. Using the introduced modeling approach, we analyze the performance of four variants of 7DS. In particular, we perform a comparative evaluation of systems comprising of mobile devices running 7DS with and without power conservation as well as with and without support of fixed infostations. A validation of results derived from the analytical models against simulation results shows excellent agreement. In fact, the analytical results almost always lie in the 99% confidence interval of the corresponding simulation results.

The performance study for 7DS shows that neither the transmission range nor the selected variant of 7DS has a significant impact on the fraction of dataholders in the long run. However, for high transmission ranges the selected variant of 7DS has a significant impact on the hit rate. Depending on the 7DS variant and the buffer size, hit rates between 0.48 and 0.92 can be achieved. Furthermore, a reduced transmission range of 115 m yields higher hit rates than using aggressive power conservation at 230 m transmission range.

The remainder of this paper is organized as follows. Section 2 summarizes related work on modeling and performance evaluation of epidemic information dissemination in MANET. Section 3 recalls the main features of the 7DS system and the modeling approach for LRU buffers. The novel analytical modeling approach for epidemic dissemination of information and an analytical performance model for 7DS are introduced in Section 4. In Section 5, we present a comprehensive performance study of 7DS using the analytical model. Finally, concluding remarks are given.

2. RELATED WORK

Several recent papers introduced design proposals for P2P data sharing in MANET using epidemic dissemination of information and presented performance models and simulation studies of their design. Papadopouli and Schulzrinne [13] introduced 7DS and presented performance results. They proposed several design variants for 7DS. In particular, 7DS systems with and without power conservation as well as with and without infostations (denoted as P, NP, FIS-P, and FIS-NP) have been introduced. Using simulation they evaluated these design variants by considering the percentage of dataholders and the delay for receiving a data item in a finite time horizon. They represent the popularity of this data item by varying the number of mobile devices, which are querying this item. They also presented a simple analytical model based on a diffusion-controlled process for analyzing the transient behavior of a hybrid scenario consisting of mobile devices and a fixed infostation. Both their simulation study and the analytical model assumed unlimited buffers, an unlimited number of query repetitions until a data item is received and just considered the spread of a single data item.

We first extend the design of 7DS by introducing limited LRU buffers at mobile devices and by limiting the number of re-queries for a data item, since we feel that these features are essential for P2P data sharing in MANET. In contrast to [13], our performance study considers the long-run performance of 7DS with these extensions rather than the transient phase and is entirely based on an analytical model. Our analytical model is applied for a comparative evaluation of the long-run performance of four 7DS variants.

Small and Haas [14] proposed an epidemic algorithm for collecting information in a hybrid network consisting of mobile nodes (i.e., whales) and fixed infostations (i.e., buoys) [14]. Their architecture, denoted as shared wireless infostation model (SWIM), actively transfers information among wireless nodes on each contact, until information is unloaded to one of the infostations. According to [14], SWIM is a marriage of the infostation concept [7] with the ad hoc networking model. They assume unlimited buffers and just consider the spread of a single data item. For the analysis of SWIM, they used Markov models for which two input parameters must be determined off-line by simulation depending on the mobility model. By transient analysis of the Markov chain, they determine the time that a data item is present in the system as well as the probability that a data item can be uploaded to an infostation. As for [13], the main difference of their work to ours lies in that we take into account multiple data items and limited LRU buffers, and do not rely on off-line simulation for determining model parameters. Furthermore, our approach allows the derivation of an approximate analytical system model for SWIM with LRU buffers for analyzing its long-run performance behavior by appropriately tailoring the equations as presented for 7DS.

Dan and Towsley presented an approximate analytical model of the LRU and FIFO replacement schemes of a stand-alone cache under the independent reference model [5]. Building upon this work, Dan, Dias, and Yu [4] analyzed the effects of skewed data access on the buffer hit ratio in a distributed data sharing environment connected by a wired network. In particular, they observed two properties recalled in Section 3: the conservation of flow and the relative push down rate. Our analysis of P2P data sharing in a MANET based on epidemic data dissemination builds upon their work. We show that the conservation of flow property still holds such mobile data sharing environments and derive the corresponding equation. Furthermore, we observe that the relative push down rate is given by the same equation as in [4].

3. BACKGROUND

3.1 Description of 7DS with Finite Buffers and Limited Query Repetition

The 7DS system [13] has been designed to enable Web browsing on mobile devices with a non-persistent Internet connection. When an Internet connection is not available for a mobile device, Web pages can be retrieved from other mobile devices or an infostation within the spatial proximity. 7DS is designed for devices with IEEE 802.11 wireless network interfaces and a memory capacity typical for today’s laptops or personal digital assistants (PDAs).

To keep the paper self-contained, we recall the basic features of 7DS. Each mobile device running 7DS maintains a local buffer, i.e., a directory on the local disk, for storing Web pages. To query the buffer of remote devices, 7DS uses query and report message. A mobile device that searches for a Web page sends a query message. The message contains the URL of the Web page and the address of the inquiring device. A query message is sent to all devices in the transmission range, using single hop multicast in the IEEE 802.11 ad hoc mode. When a mobile device or an infostation receives a query message, it searches its buffer for a

Web page with the given URL. On a hit, it returns a report message to the inquiring mobile device. The inquiring device selects the most appropriate download source from all report messages and, subsequently, transfers the Web page using a direct HTTP connection. Subsequently, this Web page is available to other mobile devices within the former device's radio coverage implementing epidemic information dissemination.

Besides this P2P data sharing among mobile clients, the 7DS system allows mobile devices to switch between ON and OFF periods for power conservation. Furthermore, among other additional features [13] outlines how 7DS can support server-client (S-C) data delivery using fixed or mobile infostations as well as hybrid forms comprising of P2P data sharing among mobile clients and additional S-C data delivery by fixed or mobile infostations.

The design of 7DS presented in [13] assumes an unlimited buffer and leaves it to the user to remove Web pages that are out-of-date or no longer of interest from the buffer in order to free disk space. However, for mobile devices like PDAs or even Smart Phones, buffer capacity is scarce. Thus, we model and analyze EID systems with a finite buffer of size B and the least recently used (LRU) replacement scheme that manages the buffer as a stack. Upon access to an item in the buffer that is located at stack position j, the item is placed at the top of the stack, and all items on positions $1,\ldots,j$-1 are pushed down by one position. If an item that is not found in the buffer is inserted into the stack, it is inserted at the stack top, and all items are pushed down by one position, where the item on position B (i.e., the least recently used item) is removed from the buffer. A straightforward way to incorporate limited buffers in 7DS lies in updating the LRU stack on each access to the local buffer, i.e., on either a local or remote query. The least recently used Web page is replaced, if on a local request a Web page is retrieved from a remote node and buffer capacity is exceeded.

Furthermore, the design of 7DS according to [13] considers no limit on the number of query repetitions for a desired data item; i.e., queries are sent periodically until the data item is retrieved from another mobile device or an infostation. For the practical deployment of a 7DS system, the number of repetitions should be limited. Furthermore, assuming movements of pedestrians, increasing the time between repetitions will improve the likelihood of a success. Thus, we consider 7DS extended by a limited number of repetitions and a binary exponential back-off.

3.2 Modeling LRU Buffer Management

Obviously, the hit rate of the buffers is the key quantity for the long-term performance behavior of an EID system like 7DS. Thus, performance modeling of EID systems requires the derivation of LRU buffer hit rate in steady-state. To keep the paper self-contained, we recall a computationally efficient approximate approach for determining the steady state buffer hit rate for LRU buffers introduced by Dan and Towsley [5]. The approach cannot only be employed for stand-alone caches, but has also been applied for distributed data-sharing environments in wired networks [4].

Consider an application, which retrieves data items that are identified by keys. The application defines a set of distinct data items $\mathbb{D}$ with cardinality $|\mathbb{D}|=D$, and a set $\mathbb{K}$ of distinct keys with cardinality $|\mathbb{K}|=K$. For ease of exposition, we identify the keys with the numbers, i.e., $\mathbb{K}=\{1,2,\ldots,K\}$. Each key $k\in\mathbb{K}$ matches a fraction $\beta(k)$ of the data items, where each item is matched by exactly one key. Thus, the set $\mathbb{D}$ can be divided into K partitions $D(k)$, each with size $\beta(k)D$ where $1 \le k \le K$ and $\sum_{k=1}^{K}\beta(k)=1$. Each key k has access probability $\alpha(k)$ where $1 \le k \le K$ and $\sum_{k=1}^{K}\alpha(k)=1$. Queries are issued according to a Poisson process with rate λ and follow the independent reference model. That is, for a query q_i holds $P\{q_i = k\}= \alpha(k)$, regardless of the history of previous queries. The independent reference model is commonly used for studying buffer behavior.

As key concept, the approach of Dan and Towsley approximately determines the expected number of items of partition $D(k)$ in the top j positions of the LRU stack, denoted by $b(k,j)$. When a query is issued for key k, the hit probability in the top j stack positions is given by $b(k,j)/\beta(k)D$. Let $r(k,j)$ denote the rate for pushing down items of partition $D(k)$ from stack position j to stack position j+1 where $1 \le k \le K$ and $1 \le j \le B$. [5] argues that under steady-state-conditions the long-term rate for pushing down an item from position j to stack position j+1 is equal to the rate for inserting the item into the top j stack positions. Thus, with the assumption that on a miss an item can always be brought into the buffer, the rate $r(k,j)$ is equal to the miss rate for key k in the top j stack positions. This miss rate is given by the product of the query issue rate, the access probability for key k, and the miss probability for key k. Thus, we have:

$$r(k,j) = \lambda\alpha(k)\left(1-\frac{b(k,j)}{\beta(k)D}\right) \tag{1}$$

This observation is denoted as *conservation of flow* in [4]. Let $p(k,j)$ denote the probability that an item of partition $D(k)$ is located at position j in the LRU stack where $1 \le k \le K$ and $1 \le j \le B$. As shown in [5], the probability $p(k,j)$ can be derived by exploiting the fact that if an item is pushed down from stack position j in the LRU scheme, then it moves to position j+1. Subsequently, the probability $p(k,j)$ is given by the probability for pushing down an item from position j-1 to position j under the condition that the current query triggers such a buffer movement from position j-1 to j in the LRU stack. The probability $p(k,j)$ can be approximated closely by the ratio between the rate for a buffer movement of an item of $D(k)$ from position j-1 to j and the rate for such a buffer movement of any item of $\mathbb{D}$. Thus, using (1) we have:

$$\begin{aligned} p(k,j) &\approx \frac{r(k,j-1)}{\sum_{n=1}^{K} r(n,j-1)} \\ &= \alpha(k)\left(1-\frac{b(k,j-1)}{\beta(k)D}\right) \Big/ \sum_{n=1}^{K}\alpha(n)\left(1-\frac{b(n,j-1)}{\beta(n)D}\right) \end{aligned} \tag{2}$$

This observation is denoted as *relative push down rate* in [4]. Note that the probability $p(k,j)$ is independent of the query issue rate. Subsequently, for each partition $D(k)$, the expected number of items of D(k) in the top j positions, $b(k,j)$, can be determined by the summation of the probabilities for finding an item of D(k) at a position less or equal to j in the buffer. Thus, we have:

$$b(k,j) = \sum_{n=1}^{j} p(k,n)\ , \ k=1,\ldots,K,\ j=1,\ldots,B \tag{3}$$

Combining (2) and (3) together with $b(k,1) = \alpha(k)$ for $1 \le k \le K$ leads to an iterative scheme for the approximate computation of $b(k,j)$ for $k=1,\ldots,K$, and $j=2,\ldots,B$. This iterative scheme has the computational complexity $O(KB)$. As noted in [5], due to the approximative computation, values of some $b(k,j)$ may exceed $\beta(k)D$ for some k and j. In this case, all $r(k,n)$ are set to 0 for $j < n \le B$. That means, it is assumed that all items of partition k are

Table 1. Definitions of key measures for modeling the 7DS system.

Expression	Meaning
$p_{local}(k,j)$	Probability for a hit for key k in the top j positions of the local buffer.
$p_{local}(k,B-j\mid j)$	Conditional probability for a hit for key k in the bottom B-j positions of the local buffer given that no hit for key k occurs in the top j positions.
$p_{origin}(k)$	Probability for retrieving an item matching key k from the origin device.
$p_{remote}(k)$	Probability for retrieving an item matching key k from a remote device other than the origin device.
$r_{local}(k,j)$	Rate for pushing down items of partition $D(k)$ from stack position j to $j+1$ due to query responses from the local buffer.
$r_{remote}(k,j)$	Rate for pushing down items of partition $D(k)$ from stack position j to $j+1$ due to query responses from remote buffer.

located in the top j locations of the buffer and are never pushed down beyond position j. Having computed the quantities $b(k,B)$ using (3), the hit rate of an LRU buffer is given by:

$$HR_{LRU} \approx \sum_{k=1}^{K} \frac{\alpha(k)b(k,B)}{\beta(k)D} \qquad (4)$$

In the next section, we show how to generalize (1) to (4) to EID systems on a MANET with finite LRU buffers.

4. THE ANALYTICAL MODELING APPROACH

4.1 General Approach for Modeling Epidemic Information Dissemination

We assume a MANET with N mobile devices. Each device has a transmission range R. We assume that a transmission between two devices is successful, if and only if for the distance d between the devices holds $d \le R$. We assume that the MANET is deployed in a square area A. These N devices are initially distributed within A according to a uniform distribution. To incorporate device mobility, which is essential for EID, we make two: (i) the mobile devices have a steady-state spatial distribution that is equal to a uniform distribution within the area and (ii) the set of mobile devices in the vicinity of a device changes substantially between two successive queries, i.e., the nodes that receive a query are independently chosen from a uniform distribution for each query. On the first view, these assumptions seem to be very restrictive towards the mobility model, since most mobility models to not exhibit these properties. However, we will show in Section 5 that the analytical model closely matches simulation results, even if the mobility model employed in the simulation model violates these assumptions.

To incorporate a generic model for the EID system consistent with the approach presented in Section 3.2, we assume that the EID systems distributes D distinct data items that are matched by K distinct keys $k \in \mathbb{K}$. Each item d is initially stored by exactly one of the N mobile devices denoted as the origin device of item d. The selection of the devices storing d is performed according to a uniform distribution. A data item is available from its origin device at all time, even if it is not stored in the buffer of any other device.

Each device issues queries according to a Poisson process with rate λ. We refer to a device that issues a query q_i as the inquiring device for query q_i. The key k in the query is selected with probability $\alpha(k)$ according to the independent reference model, for $1 \le k \le K$. Besides querying the local buffer, the inquiring device sends the query to a set of other devices. Throughout this paper, we assume that the wireless transmission range R solely determines this set of other devices. A device that receives the query will generate response messages, if it is either the origin device of a data item matching q_i, or stores matching data items in its buffer. After retrieving an item from a remote device, the inquiring device inserts the item into its buffer.

We consider a finite buffer of size B with LRU replacement for storing data items at each mobile device. We assume that data items do not change or expire and that their popularity distribution remains constant over time. It is easy to see that the system can be modeled by a Markov chain with a number of states that grows exponentially with N, B, and K. The states are given by the content of all buffers, where the state of each buffer is a permutation of a subset of $\mathbb{K}$. The state transition probabilities are determined by the content of the buffers in the current state, the access probabilities $\alpha(k)$, $1 \le k \le K$, and the probability that a mobile device with an item from partition $D(k)$ in the buffer is in vicinity. We omit a formal proof of convergence, but state that it is easy to show that the Markov chain is finite, irreducible and aperiodic. Thus, it has a limiting state probability, which is approximated by our modeling approach.

Recall that the conservation of flow property used in (1) assumes that on a miss an item is always inserted into the buffer. This assumption does not hold for data sharing in a MANET. In a MANET, on a miss an item can only be inserted into the buffer with some probability depending on:

(i) the probability that the item can be retrieved from the (remote) origin device and

(ii) the probability that the item can be retrieved from the buffer of any other remote device.

Subsequently, we show how to extend (3) for P2P data sharing on mobile devices with LRU buffers using EID. We assume that all buffers are stochastically identical, i.e., the expected number of data items of partition D_k found in any buffer is given by $b(k,B)$. We note that the number of data items of partition $D(k)$ in the top j positions depends on both the number of data items of all partitions in the top j-1 positions of the local buffer, $b(k,j\text{-}1)$, and on the number of items in remote buffers at any position, $b(k,B)$. Analogous to (3), we can write:

$$\begin{aligned} b(k,j) = f_k\big(&b(1,j-1), b(2,j-1), \ldots, \\ &b(K,j-1), b(1,B), b(2,B), \ldots, b(K,B)\big), \end{aligned} \qquad (5)$$

$$\text{for } 1 \le k \le K,\ 1 \le j \le B,$$

In (5) in general the f_k are non-linear functions in $b(k,j)$ which opposed to (3) cannot be solved by simple back-substitution. In the following, we propose an effective iterative scheme for the approximate computation of $b(k,B)$, $1 \le k \le K$. For ease of exposition, we define $b(k,0) = 0$. We start with the initial assignment of the remote buffers:

$$b^{(0)}(k,B) = \frac{B}{K},\ 1 \le k \le K. \tag{6}$$

Then, we iteratively compute

$$\begin{aligned} b^{(i)}(k,j) = f_k\Big(&b^{(i-1)}(1,j-1),\dots,b^{(i-1)}(K,j-1),\\ &b^{(i-1)}(1,B),\dots,b^{(i-1)}(K,B)\Big) \end{aligned}, \tag{7}$$

for i = 1,2,… and j = 1,…,B

until the termination condition

$$\left|b^{(i)}(k,B) - b^{(i-1)}(k,B)\right| < \varepsilon \tag{8}$$

is satisfied for $1 \le k \le K$.

Note that customizing the iterative scheme to a particular EID system requires derivation of f_k. To prove the convergence of the iterative scheme, note that (7) describes a simple fix point iteration. We can easily use the f_k to construct a function $F{:}\mathbb{R}^d \to \mathbb{R}^d$ with $d = (B+1)K$. The function operates on a vector $\boldsymbol{b} \in \mathbb{R}^d$ with $\boldsymbol{b}(jK+k) = b(k,j)$. Following the contraction mapping principle, the iteration $\boldsymbol{b}^{(i+1)}=\mathbf{F}(\boldsymbol{b}^{(i)})$ for i= 1,2,… converges, if holds:

$$\left\|F(\boldsymbol{b}) - F(\boldsymbol{b}^*)\right\|_\infty \le c\left\|\boldsymbol{b} - \boldsymbol{b}^*\right\|_\infty \tag{9}$$

for some constant c with $0 < c < 1$. The convergence of (9) must be shown for each F given by the functions f_k for a particular EID system.

4.2 Conservation of Flow for 7DS

For customizing the approach presented in Section 4.1 to the 7DS system, we have to derive an expression for the push down rate of a data item of partition $D(k)$ from position j to position j+1 in the LRU stacks of 7DS. Using the conversation of flow observation, this is equal to finding an expression for the rate of inserting an item of partition $D(k)$ into the top j positions. In 7DS, an item can be inserted into the top j positions of the buffer either on a local or a remote request. That is:

(i) on a local request, the item is not already located in the top j positions and it is either located in the bottom B-j positions of the local buffer or it can be retrieved from the origin device or from any other remote device.

(ii) on a remote request, the item is not already located in the top j positions of the local buffer, but is located in the bottom B-j positions.

To derive an expression for the relative push down rate for 7DS, we introduce additional measures in Table 1 for $1 \le k \le K$ and $1 \le j \le B$. Subsequently, we split the push down rate $r(k,j)$ into two rates $r_{local}(k,j)$ and $r_{remote}(k,j)$ and derive these rates as follows:

(i) *Local request*: Using the notation introduced in Table 1 to derive the rate for pushing down an item on a local request, we have:

$$\begin{aligned} r_{local}(k,j) = \alpha(k)\lambda&\left(1 - p_{local}(k,j)\right)\\ &\cdot\left(1 - \left(1 - p_{local}(k,B-j\,|\,j)\right)\right.\\ &\left.\cdot\left(1 - p_{origin}(k)\right)\left(1 - p_{remote}(k)\right)\right) \end{aligned} \tag{10}$$

Since all devices issue queries according to a Poisson process with identical rate λ, the probability that a request is a local request is $1/N$.

(ii) *Remote request*: For the rate for pushing down an item k on a remote request, we have:

$$r_{remote}(k,j) = \alpha(k)\lambda\left(1 - p_{local}(k,j)\right)p_{local}(k,B-j\,|\,j). \tag{11}$$

Under the assumption that a device sends a query message even if it stores a matching item in the local buffer, analogous to case (i) the probability for a remote request is $(N\text{-}1)/N$.

By weighting $r_{local}(k,j)$ and $r_{remote}(k,j)$ with their probability of occurrence, the overall push down rate for an item of partition $D(k)$ is given for $1 \le k \le K$ and $1 \le j \le B$ by:

$$\begin{aligned} r(k,j) &= \frac{1}{N}r_{local}(k,j) + \frac{N-1}{N}r_{remote}(k,j)\\ &= \alpha(k)\left(1 - p_{local}(k,j)\right)\cdot\\ &\quad\cdot\left(\frac{1}{N}\left(1 - \left(1 - p_{local}(k,B-j\,|\,j)\right)\left(1 - p_{origin}(k)\right)\right.\right.\\ &\quad\left.\left.\cdot\left(1 - p_{remote}(k)\right)\right) + \frac{N-1}{N}p_{local}(k,B-j\,|\,j)\right) \end{aligned} \tag{12}$$

Note that (12) constitutes the conservation of flow for the 7DS system with LRU buffers operating in a MANET analogous to (1) for standalone LRU caches. Subsequently, we will show how to derive the individual terms of (12). Similar to [5], the probability for a hit in the top j positions of the local LRU buffer is given by:

$$p_{local}(k,j) = \frac{b(k,j)}{\beta(k)D} \tag{13}$$

For computing the probability for a hit in the bottom B-j positions, recall that in (10) and (11) this probability is used under the condition that no hit has occurred in the top j stack positions. Thus, we derive the conditional probability as:

$$p_{local}(k,B-j\,|\,j) = \frac{p_{local}(k,B) - p_{local}(k,j)}{1 - p_{local}(k,j)} \tag{14}$$

Neglecting border effects in the simulation area, the probability for a successful transmission to a particular device can be determined by the ratio between the area covered by the wireless transmission range of a device (i.e., πR^2) and the total considered area A. Since the origin device of key k is with probability $1/N$ equal to the local device and with probability $(N\text{-}1)/N$ an arbitrary remote device, the probability for retrieving an item from the origin device is given by:

$$p_{origin}(k) = \frac{1}{N} + \frac{N-1}{N}\frac{\pi R^2}{A} \tag{15}$$

The computation of the probability for receiving an item from a remote buffer can be broken down to the probability that exactly n remote devices receive the query from the inquiring device and at least one of them stores a matching piece of information in its local buffer:

$$p_{remote}(k) = \sum_{n=1}^{N-1} \binom{N-1}{n} \left(\frac{\pi R^2}{A}\right)^n \left(1 - \frac{\pi R^2}{A}\right)^{N-1-n} \cdot \left(1 - \left(1 - p_{local}(k,B)\right)^n\right) \quad (16)$$

Putting (13) to (16) into (12) completely specifies the overall push down rate for an item of partition $D(k)$. Subsequently, following (2) we derive the probabilities:

$$p(k,j) \approx \frac{\frac{1}{N} r_{local}(k, j-1) + \frac{N-1}{N} r_{remote}(k, j-1)}{\sum_{n=1}^{K} \left(\frac{1}{N} r_{local}(n, j-1) + \frac{N-1}{N} r_{remote}(n, j-1)\right)} \quad (17)$$

Equation (17) constitutes the relative push down rate for the 7DS system with LRU buffers operating in a MANET analogous to (2) for standalone LRU caches. Note that as in (2) the rate λ of the query arrival process cancels out in (17). Subsequently, the $b(k,j)$, can be determined by summation of the probabilities $p(k,j)$ according to in (3).

Finally, the overall hit rate HR_{7DS} of the 7DS system can be determined using the probabilities that upon a query for key k the items matching k can be retrieved either from the local buffer, the origin device, or the buffer of a remote device. That is:

$$HR_{7DS} = \sum_{k=1}^{K} \alpha(k) \left(1 - \left(1 - p_{local}(k,B)\right) \cdot \left(1 - p_{origin}(k)\right)\left(1 - p_{remote}(k)\right)\right) \quad (18)$$

Each iteration of the iterative scheme (7) requires the computation of $b(k,j)$ for $k = 1,\ldots,K$ and $j=1,\ldots,B$. Furthermore, each iteration of (7) requires one evaluation of (16) with N-1 summations for $k = 1,\ldots,K$,. Thus, the complexity of each iteration of the scheme is $O(KN + KB)$. Assuming that the number of iterations of (7) until reaching the termination condition (8) is bounded by some constant, the iterative scheme for calculating the hit rate of the local buffer has an overall complexity of $O(KN + KB)$. The subsequent computation of the overall hit rate of the EID system using (18) has the additional complexity $O(KN)$. Putting it altogether, the overall computational complexity of the iterative scheme is $O(KN + KB)$. We omit a formal prove of the convergence following (9) and refer to Section 5 for an experimental validation. For all 7DS system models analyzed in Section 5, the corresponding scheme converges in less than 40 iterations.

4.3 System Models for the 7DS Variants

In this section, we will show how to extend the basic model presented in Section 4.2 to represent the different variants of 7DS introduced in [13]. As first extension of the model, we include server-client (S-C) scenarios using fixed and mobile infostations. For a 7DS variant with fixed infostations (FIS), we assume that the data is not provided by the mobile device, but by N_{FIS} fixed infostations that are placed within the area A with an optimal spatial distribution. That is, the area covered by the infostations does not overlap. Note that retrieving a data item from an infostation requires bi-directional communication to exchange query and response messages. Subsequently, the infostations have an effective transmission range that is equal to the transmission range of the mobile devices, R, even if it possible to operate infostations with much higher transmission ranges in practical applications due to a much better power supply. Assuming an optimal spatial distribution, each infostation linearly increases the probability for retrieving an item from the origin device. Thus, (15) changes to:

$$p_{origin}(k) = N_{FIS} \cdot \frac{\pi R^2}{A}. \quad (19)$$

The assumption of optimal spatial distribution does not hold for mobile infostations (MIS). Therefore, with N_{MIS} mobile infostations a data item can be retrieved if any of the mobile infostations is reachable, where the radio coverage of the mobile infostations may overlap. In this case, (15) changes to:

$$p_{origin}(k) = N_{MIS} \cdot \left(1 - \left(1 - \frac{\pi R^2}{A}\right)^{N_{MIS}}\right) \quad (20)$$

Consistent with [13], we will use N_{FIS} = N_{MIS} = 1 in our performance studies. Note that in this case (19) is equal to (20).

When using 7DS in a P2P variant instead of a S-C variant, the mobile devices must supply data items. For example, in a mobile Web browsing application, a mobile device may download Web pages using a fixed network connection before joining the 7DS system. Subsequently, the mobile device will supply the Web pages to other devices. In such scenario, it is not realistic to assume that each data item is stored by exactly one mobile device. We will rather assume that the number of origin devices for a data items from partition $D(k)$ reflects the popularity of the items $\alpha(k)$, $1 \le k \le K$. Thus, (15) changes to include the increased quantities, while ensuring that each data item is available from at least one mobile device by taking the ceiling for values smaller than 1:

$$p_{origin}(k) = \frac{\lceil \alpha(k)N \rceil}{N} + \left(1 - \frac{\lceil \alpha(k)N \rceil}{N}\right)\left(1 - (1 - \frac{\pi R^2}{A})^{\lceil \alpha(k)N \rceil}\right) \quad (21)$$

Beside the S-C and P2P variants without power conservation denoted as FIS-NP and NP, respectively, [13] considers 7DS variants that use power conservation. These variants are denoted as FIS-P and P, respectively. With power conservation, the mobile devices divide the operation cycle in ON periods of duration P_{ON} and OFF periods of duration P_{OFF}. During an OFF period, the mobile device switches off the wireless communication interface and does not respond to any queries. Since the ON and OFF periods are not synchronized among the mobile devices, the probability for contacting a mobile device in an ON period is P_{ON} / (P_{ON} + P_{OFF}). To model power conservation in P2P mode, (16) and (20) can easily be extended:

$$p_{remote}(k) = \sum_{n=1}^{N} \binom{N-1}{n} \left(\frac{\pi R^2}{A}\right)^n \left(1 - \frac{\pi R^2}{A}\right)^{N-1-n} \cdot \left(1 - \left(1 - \frac{P_{on}}{P_{on} + P_{off}} p_{local}(B)\right)^n\right) \quad (22)$$

$$p_{origin}(k) = \frac{\lceil \alpha(k)N \rceil}{N} + 1 - \left(\frac{\lceil \alpha(k)N \rceil}{N}\right) \cdot \left(1 - (1 - \frac{P_{on}}{P_{on} + P_{off}} \frac{\pi R^2}{A})^{\lceil \alpha(k)N \rceil}\right) \quad (23)$$

As last feature, [13] considers periodical query repetition in the case that a data item is not retrieved. Note that this feature is essential for the performance analysis of the transient behavior of a 7DS system with a single data item, since without query repetition no node will be interested in the item after some time and dissemination will stop. In realistic applications, a limited number of N_{RT} repetitions with an exponential back-off time is obviously preferable than unlimited query repetitions. For ease of exposition, we count the initial query as first repetition. It is easy to see that the number of repetitions changes the probability for an access to a key k, $\alpha(k)$. This is because a query for key k of partition $D(k)$ is likely to be repeated, if the probability for a miss in partition $D(k)$ is high. The probability for a hit $p_1(k)$ in the first repetition of a query for an item from partition $D(k)$ is given by the probability that an item can be either retrieved from the local buffer, the origin or a remote buffer, i.e.,

$$p_1(k) = 1 - \left(1 - p_{local}(k,B)\right)\left(1 - p_{origin}(k)\right)\left(1 - p_{remote}(k)\right) \quad (24)$$

In contrast, the item will not be in the local buffer on successive repetitions. Thus, the probability for a hit $p_2(k)$ is given by:

$$p_2(k) = 1 - \left(1 - p_{origin}(k)\right)\left(1 - p_{remote}(k)\right) \quad (25)$$

Thus, the expected number of repetitions of a query for an item of partition $D(k)$ is given by:

$$\begin{aligned} R(k) &= p_1(k) + \left(1 - p_1(k)\right)\left(1 + \sum_{n=1}^{N_{RT}-2}\left(n\left(\left(1 - p_2(k)\right)^{n-1} p_2(k)\right)\right)\right. \\ &\qquad \left. + \left(N_{RT} - 1\right)\left(1 - p_2(k)\right)^{N_{RT}-2}\right) \quad (26) \\ &= p_1(k) + \left(1 - p_1(k)\right)\left(1 + \frac{1 - \left(1 - p_2(k)\right)^{N_{RT}-1}}{p_2(k)}\right) \end{aligned}$$

The effective access probability $\hat{\alpha}(k)$ to partition $D(k)$ with N_{RT} queries can be determined by:

$$\hat{\alpha}(k) = \frac{R(k)\alpha(k)}{\sum_{n=1}^{K} R(n)\alpha(n)} \quad (27)$$

To model the 7DS variants considered in [13], we use (21) to model a P2P scenario without power conservation (NP) as well as (23) and (22) to model a P2P scenario with power conservation (P). For the S-C scenarios, we use (19) to model a scenario without power conservation (FIS-NP) as well as (19) and (22) to model a scenario with power conservation (FIS-P). Note that [13] considers passive querying in S-C scenarios, i.e., a mobile device only sends queries when it receives an advertisement from an infostation. Thus, hit ratio is always equal to 1. We will rather consider the no-trivial case of using infostations together with active queries. To avoid confusion, we will denote the resulting scenarios as FIS-NP* and FIS-P*, respectively. The impact of query repetition is only analyzed for the NP variant by using (21) and (27).

5. PERFORMANCE ANALYSIS OF 7DS WITH LIMITED BUFFERS

5.1 Modeling Assumptions

[13] presents a study of the transient behavior of 7DS considering the spread of a single data item among mobile devices with unlimited buffers. In the performance study presented in this section, we use assumptions identical to [13] as far as possible. However, to evaluate the long-term performance of 7DS with limited buffers and LRU replacement, we have to extend the system model used in [13]. Consistent with [13], we assume that N mobile devices move in a square area of 1000 m $\times$ 1000 m. All mobile devices have an identical transmission range R. Similar to [13], we consider three different transmission power levels in many experiments, i.e., 281.8 mW, 17.6 mW, and 1.1 mW. Using the radio model assumed in [13], these transmission powers result in a high transmission range of R = 230m, a medium transmission range of R = 115m, and a low transmission range of R = 57.5m, respectively.

In contrast to [13], we assume that K = 1000 popular data items are distributed in the buffers of the mobile devices. Since keys in 7DS constitute URLs that identify data items given by Web documents, we have a one-to-one matching between keys and data items. That is, $\beta(k)$=1 and $|D(k)|$= 1 for $1 \le k \le K$. [13] uses numbers of mobile devices $N \in \{5,10,\ldots,25\}$ to represent different levels of popularity for the single data item considered in each experiment. We assume a fixed number of mobile devices, N = 64 and represent the popularity of the data items by setting the access probabilities $\alpha(k) \cong k^{-\gamma}$ for $1 \le k \le K$. This is, the access probabilities follow a Zipf-like distribution. [3] analyzes the object popularity distribution in a mobile Web browsing application and reports a Zipf-like distribution with values of γ between 0.85 and 1. Thus, we will use γ = 0.9 in most of our experiments. [3] also reports that mobile devices repeat a significant fraction of requests. Thus, the independent reference model constitutes a realistic assumption for a mobile Web browsing application.

In many experiments, we will consider two different buffer sizes B = 64 and B = 256. In a typical 7DS application, data items constitute Web objects. Assuming an average object size of 20 KB, 256 data items constitute about 5 MB of data, which is the size of a typical Web browser cache that can be easily handled by a notebook. 64 data items constitute about 1.25 MB, a data volume that can be handled by a modern PDA. Table 2 summarizes the default values for all model parameters. These values are used in all performance experiments if not stated otherwise.

To validate the analytical model, we implement a detailed simulation model of the 7DS system. Most assumptions of the simulation model are identical to those used for the analytical model. Since we consider a 7DS system in operation, the mobile devices will start with full buffers in the initial state of the simulation. The initial distribution of items in the buffers depends heavily on the application scenario. Thus, beside initially empty buffers, we consider a best case and a worst case initial distribution in the simulation experiments:

(i) each buffer is filled with B data items chosen from $\mathbb{K}$ according to a Zipf-like distribution with parameter γ similar to the parameter of the request probabilities $\alpha(k)$. This constitutes a best-case scenario, since popularity of the data items is reflected in the initial content of each buffer.

Table 2. Default values for model parameters.

Parameter	**Value**
Number of mobile devices N	64
Transmission Range R	115 m
simulation area A	1000 m $\times$ 1000 m
Query rate λ	1/60
Zipf-parameter of query locality γ	0.9
Number of keys K	1000
Buffer size B	256

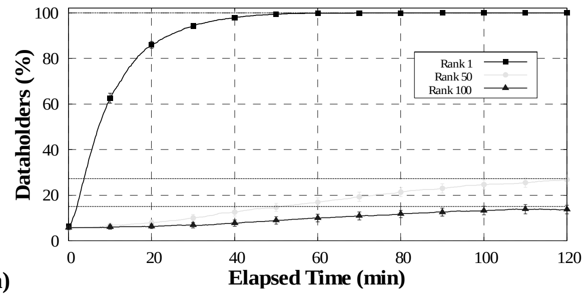

(a)

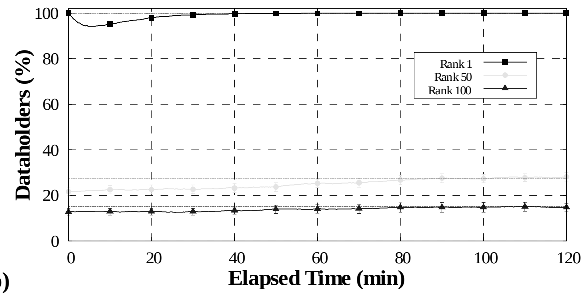

(b)

Figure 1. Transient phases of 7DS with initial distribution for buffer content drawn from (a) a uniform distribution and (b) a Zipf-like distribution.

(ii) each buffer is filled with B data items chosen from $\mathbb{K}$ according to a uniform distribution. That is, each item is equally likely to be contained in a buffer, which constitutes a worst-case scenario since popularity of the data items is not considered.

To avoid inaccuracies due to border effects when mobile devices are located close to edges of the simulation area, we use the toroidal distance model as proposed in [2]. That is, the flat square-like area becomes a torus, so that mobile devices located at one border of the simulation area can contact mobile devices at the opposite border. The mobile devices move according to one of the following three mobility models:

(i) *Random placement model* (RP): New positions of the mobile devices are chosen from the simulation area by a uniform distribution between two successive queries.

(ii) *Random waypoint mobility model with pauses* (RWP, [12]): A mobile device starts at a position chosen from the simulation area by a uniform distribution. It moves to a destination position that is chosen by a uniform distribution, too. The device speed is chosen uniformly from $(0,v_{max}]$. When the mobile device reaches the destination position, it holds for an amount of time chosen uniformly from $(0,T_{hold}]$, before choosing a new destination position and continuing the process.

(iii) *Reference point group mobility* (RPG, [10]): The mobile devices move in G groups that cover each a circular area with radius r_g. Groups move according to the random waypoint model with $T_{hold} = 0$. Each mobile device is associated with a reference point uniformly chosen from the area covered by the group. The mobile devices are placed at positions that are randomly chosen from a circular area with radius r_n around their reference point.

Note that the RP mobility model is not realistic, but very close to the assumptions of the analytical modeling approach. The RWP mobility model is quite realistic for mimicking the movement of individual pedestrians. However, it has been shown that the RWP model leads to a non-uniform distribution of mobile devices while they are moving [12]. Thus, it violates a key assumption of the modeling approach. Additionally, the RPG mobility model violates the assumption of a substantially change of the vicinity of a mobile device between two successive queries.

To determine the performance of 7DS in the simulation study, we count the number of responses received for each query and calculate the average hit rate. We conduct 30 batches each comprising of 10,000 queries and compute the 99% confidence interval for the hit rate using the batch means.

5.2 Validation of the Analytical Model

To provide evidence that a steady state analysis is important for evaluating the performance of EID systems with limited buffers, we derive the duration of the transient phase for a 7DS system starting with full buffers. The initial buffer content is chosen according to both the worst case and the best case model described above. Figure 1 plots the transient phase of the average fraction of dataholders, i.e., the fraction of nodes that store a particular data item, for the data items with rank 1, 50, and 100, respectively. The average fraction of dataholders for each item is calculated in intervals of 60 seconds. For all data items, the steady state fraction of dataholders is shown as a dashed horizontal line. For the experiments, we use the RWP mobility model with parameters v_{max} = 2 m/s and T_{hold} = 30 s. Figure 1 shows that even in the worst case scenario with a uniform distribution of buffer content, the fraction of dataholders reaches the steady state for a popular data item in less than one hour. For less popular data items, the fraction of dataholders is close to the steady state values in less than 2 hours. If the initial buffer content is chosen according to a Zipf-like distribution, the fraction of dataholders is close to the steady state fraction almost immediately. Note that due to the Zipf-like distribution of access probabilities, the hit rate of 7DS is determined by the number of hits to popular data items. Thus, hit rate is close to the steady state in below one hour. We

Table 3. Validation of the analytical model for different mobility models

Buffer Size	Analytical Model	Random Placement		Random Waypoint v_{max} = 2 m/s T_{hold} = 30 s		Random Waypoint v_{max} = 32 m/s T_{hold} = 30 s		Random Waypoint v_{max} = 2 m/s T_{hold} = 1200 s	
	Hit Rate	*Mean*	*Conf. Interval*	*Mean*	*Conf. Interval*	*Mean*	*Conf. Interval*	*Mean*	*Conf. Interval*
32	0.5216	0.519	[0.516 , 0.521]	0.526	[0.523 , 0.529]	0.525	[0.522 , 0.528]	0.519	[0.517 , 0.522]
64	0.6176	0.615	[0.612 , 0.617]	0.622	[0.618 , 0.625]	0.621	[0.618 , 0.624]	0.617	[0.614 , 0.621]
96	0.6770	0.668	[0.666 , 0.671]	0.674	[0.670 , 0.679]	0.682	[0.679 , 0.685]	0.675	[0.672 , 0.678]
128	0.7205	0.712	[0.709 , 0.715]	0.726	[0.723 , 0.729]	0.721	[0.717 , 0.725]	0.719	[0.715 , 0.723]
160	0.7551	0.749	[0.746 , 0.751]	0.760	[0.755 , 0.764]	0.761	[0.758 , 0.764]	0.754	[0.751 , 0.756]
192	0.7838	0.776	[0.771 , 0.780]	0.784	[0.781 , 0.788]	0.785	[0.782 , 0.789]	0.779	[0.776 , 0.781]
224	0.8084	0.802	[0.797 , 0.807]	0.810	[0.806 , 0.813]	0.806	[0.804 , 0.809]	0.805	[0.801 , 0.808]
256	0.8299	0.826	[0.823 , 0.829]	0.836	[0.833 , 0.839]	0.837	[0.834 , 0.841]	0.825	[0.820 , 0.830]

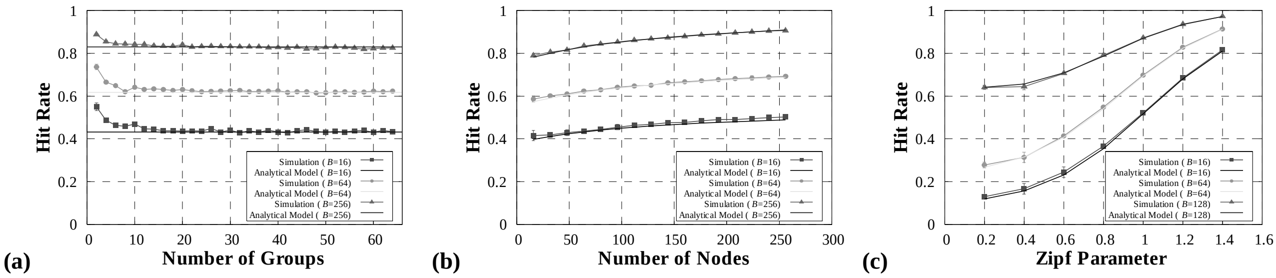

Figure 2. Validation of the analytical model for (a) different degrees of group mobility, (b) different system sizes, and (c) different degrees of locality in the queries.

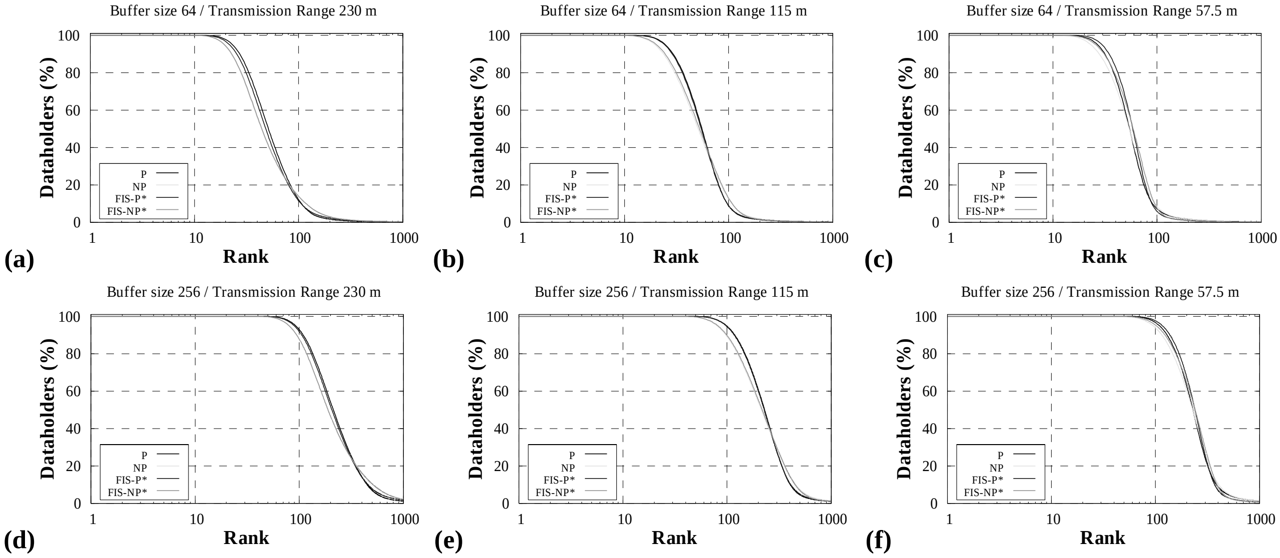

Figure 3. Long run behavior of dissemination of data items for different 7DS variants.

conclude from Figure 1 that a steady state analysis constitutes an important tool for the performance evaluation of EID systems with limited buffers, since a 7DS system with limited buffer capacity reaches steady state in a short amount of time. Furthermore, Figure 1 provides some evidence that the steady state behavior of a 7DS system is independent of the initial distribution of buffer content. In fact, we analyzed the steady state performance for various distributions of the initial buffer content in a series of simulation experiments not shown due to space limitations. These results clearly indicate that the steady state performance measures are independent from initial buffer size.

In the next experiment, we validate the analytical model against the results of steady state simulations. We use the RP mobility model as well as the RWP mobility model with three different parameter settings as shown in Table 3. Table 3 compares the analytical results for the basic variant of 7DS as described in Section 4.2 to the results of the simulation model for various buffer sizes. The table includes the 99% confidence intervals for the simulation results. Confirming the claim from Section 4, Table 3 shows that the simulation reaches a steady state for all mobility models. Furthermore, the steady state hit rates for all mobility models are almost identical for a given buffer size. The results of the analytical model almost always lie with the 99% confidence intervals of the simulation results except for small buffer sizes. However, the maximum difference to its mean value is 3%. Note that [5] reports a maximum error of about 3% for small buffers, too, which results from the approximation used in Equation (2). We conclude from Table 3 that the steady state hit rate of the 7DS system is almost independent of the mobility model and is accurately predicted by the analytical model. To gain further insight into the impact of the mobility model, we validate the analytical model against a simulation study that uses the RPG mobility model. Figure 2 (a) plots the hit rate as a function of the number of groups G for v_{max} = 2 m/s, r_g = 200 m, and r_n = 50 m. The mobile devices are distributed to the groups in a round robin fashion. In scenarios with few groups, the probability that the inquiring node and the origin node are in the same group is high, significantly increasing the hit rate above the prediction of the analytical model. With an increasing number of groups, the probability for reaching the remote origin converges against $\pi R^2/A$ and the analytical model closely matches the simulation results. We conclude from Figure 2 (a) that the analytical model can be applied for estimating hit rate even in scenarios with many small groups of mobile devices. In all further simulation studies, we use the RWP mobility model with the parameters v_{max} = 2 m/s and T_{hold} = 30 s.

We further validate the analytical model for different numbers of mobile devices N and different buffer sizes B in Figure 2 (b). The analytical model closely matches the simulation results, with an average relative difference of 0.73% for B = 64 and 0.25% for B = 256. Again, for small buffer sizes, i.e., B = 64, the analytical model underestimates the hit ratio with an average relative difference of 2.07% and a maximum of 4.3%. In a validation of the analytical model for different transmission ranges R ranging from 50 m to 250 m we observe average relative differences between the simulation results and the analytical model that range

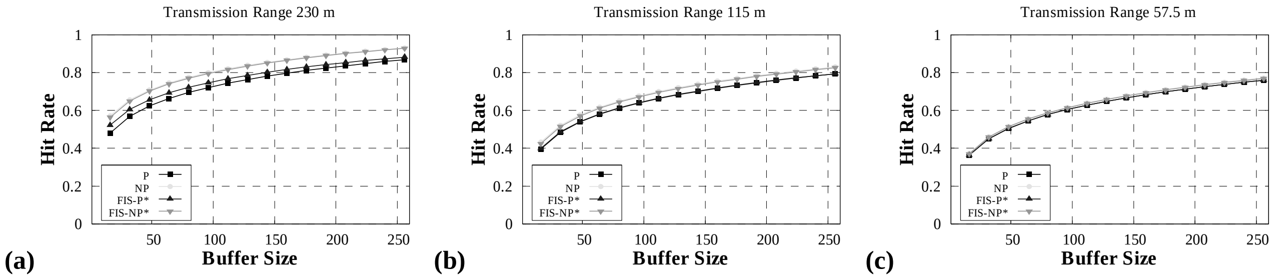

Figure 4. Comparison of hit rates achieved by different 7DS variants.

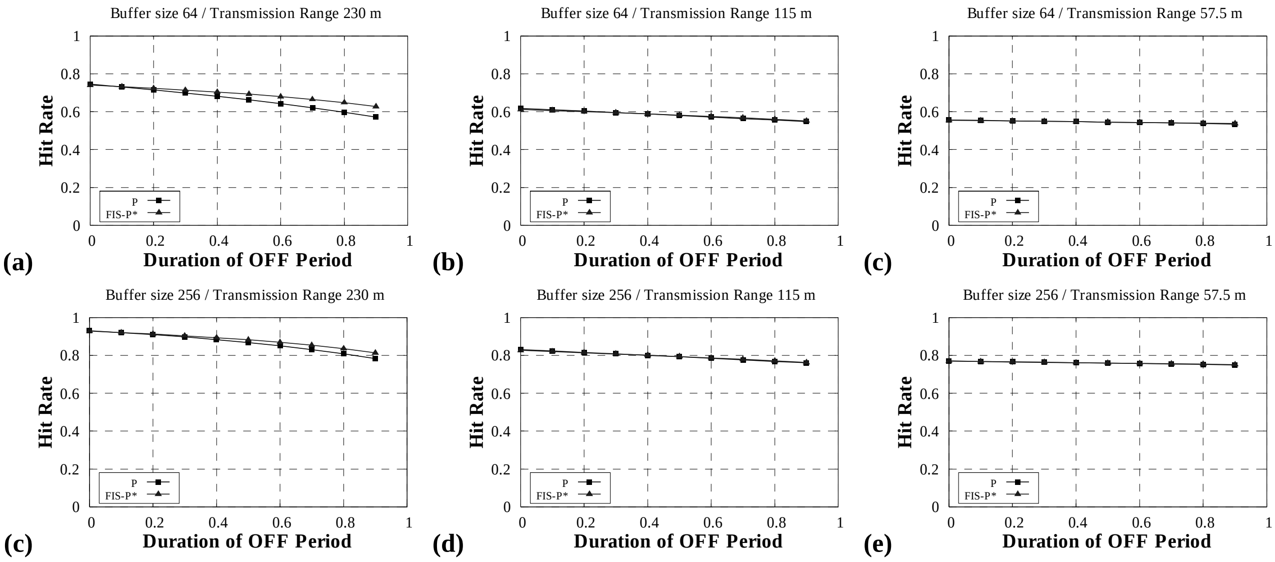

Figure 5. Impact of OFF period on the performance of the 7DS variants P and FIS-P*.

from 0.59% for B = 256 to 1.09% for B = 16. We omit a plot of the results due to space limitations. To validate the analytical model for different distributions of the access probabilities for the keys, we plot the hit ratio as a function of the locality in the query stream given by the Zipf parameter γ in Figure 2 (c). Again, we observe a slight underestimation of hit rate for small locality and small buffers, resulting in an average relative difference of 3.71% with a maximum of 8.54%. For larger buffers, the average relative difference decreases to 0.48%. Figures 2 (b) and (c) show that the analytical model provides reasonable estimates of the hit rate for various parameter settings.

We performed various experiments to validate the analytical model for other system configuration as well as for the 7DS variants NP, P, FIS-NP*, and FIS-P*. We do not include the results due to space limitations. However, we state that all validations have shown that the analytical model closely matches the simulation results.

5.3 Comparative Evaluation of 7DS Variants

To evaluate the impact of the different 7DS variants on data dissemination in the long run, we analyze the average fraction of dataholders for each data item. Figure 3 plots the average fraction for each data item for the selected variant of 7DS, i.e., NP, P, FIS-NP*, and FIS-P*, respectively. The data items are associated with their popularity rank. For both variants using power conservation, i.e., P and FIS-P*, we use a relative OFF time of 50%. [13] reports significant differences in the fraction of dataholders for data items with a different popularity during the transient phase. Supplementing this result, we find that all 7DS variants perform almost equal in the long run. This shows that both power conservation and infostations only affect the speed of the dissemination process in the transient phase, but not the long run behavior. [13] considers only popular data items, which are comparable to the items with rank 1 to 5 in our study. The experiments presented in [13] show that these items are stored by a significant fraction of the mobile devices after a transient phase of 25 minutes. Extending the observations of [13], we find that even in the long run an unpopular item is hardly stored by any mobile device. E.g., with small buffers more than 80% of the data items are stored by less than 5% of the mobile devices. Comparing the individual 7DS variants, we find that NP and FIS-NP* store slightly more unpopular data items than P and FIS-P*. However, since this observation is not very significant, we conclude from Figure 3 that neither the transmission power nor the selected variant of 7DS has significant impact on the long run behavior.

Recall that the performance of 7DS not only depends on the hit rate of the local buffer, but also on the hit rate of remote buffers. Since the long-term remote buffer hit rate cannot be analyzed using transient simulation, this component is omitted in [13]. We plot the hit rate as a function of the buffer size for all 7DS variants and three transmission ranges. Figure 4 shows that the hit rate grows with the buffer size in a log-like fashion with a minimum of 0.48 and a maximum of 0.92 depending on the 7DS variant. Decreasing transmission range from 230m to 57.5m results in hit rates between 0.37 and 0.77. That is a reduction of about 0.2 for NP and FIS-NP*, and by about 0.1 for P and FIS-P*. This confirms the observation of [13] that the transmission power has a significant impact on 7DS performance. Nevertheless, comparing

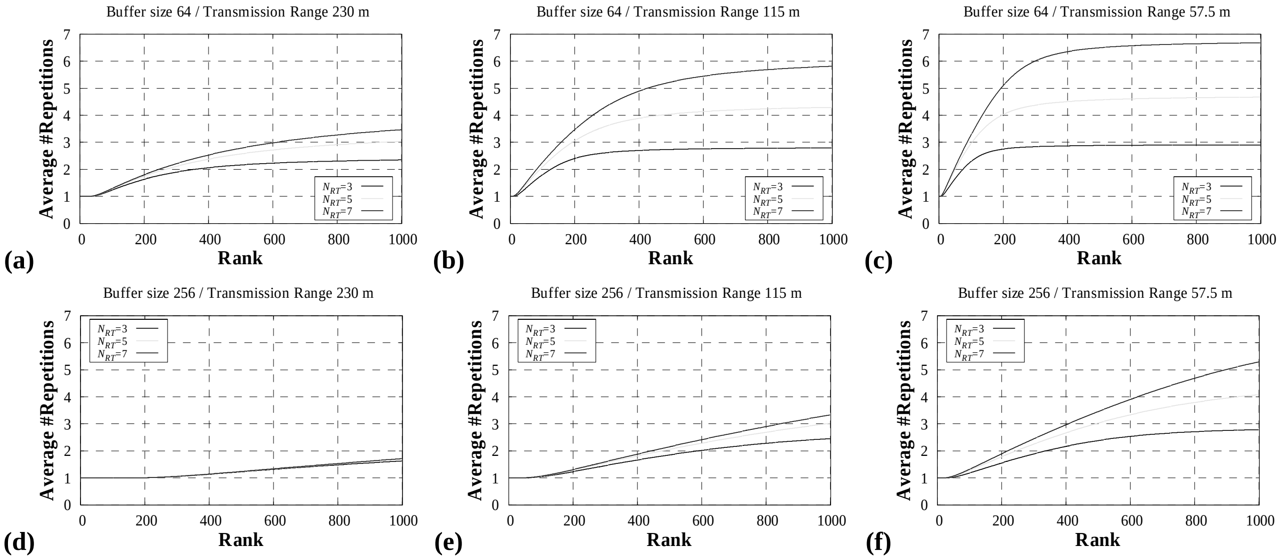

Figure 6. Average number of query repetitions for each data item.

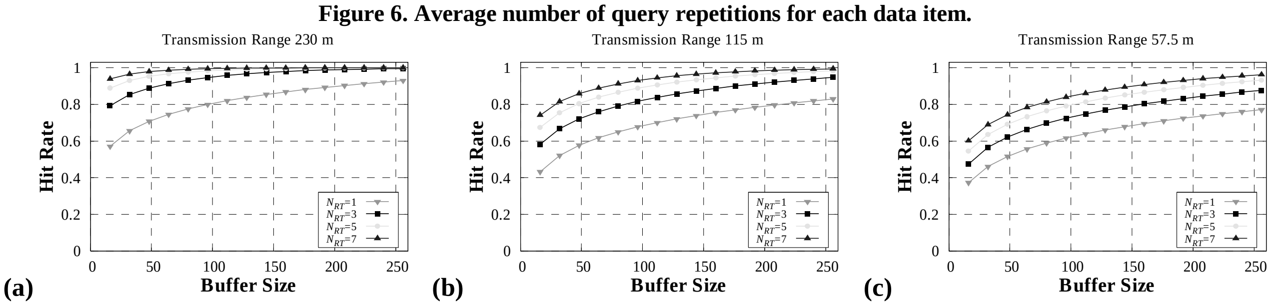

Figure 7. Hit rate of 7DS systems with query repetition.

Figures 3 and 4 shows a different rational behind this conclusion: the reduction of the hit rate is due to a reduced number of hits at remote buffers for low transmission ranges, and not due to a significant different distribution of data items in the buffers of the mobile devices. As a consequence, using the variant P and FIS-P* instead of NP and FIS-NP*, i.e., enabling power conservation, reduces the hit rate by up to 0.1 or 18% for high transmission ranges as shown in Figure 4 (a). Recall that power conservation reduces the probability for retrieving a data item from a remote buffer, since remote mobile devices will respond to queries only 50% of the time. Infostation-based 7DS variants are less sensitive to power conservation, since the infostations do not employ power conservation. The performance gain due to infostations is largest for very small buffers, where the hit rate is increased by 0.05. An additional evidence for the impact of remote hits on the system performance is given by Figures 4 (b) and (c). These figures indicate that limiting the transmission power reduces the performance loss caused by power conservation. In particular for a low transmission power, the selected variant of 7DS has no impact on the system performance.

5.4 Sensitivity to Power Conservation

In the next experiment, we calculate the hit rate as a function of the relative duration of the OFF period. The results for different buffer sizes and transmission ranges are shown in Figure 5. Confirming the conclusions from Figure 4, we find that power conservation affects the long-run performance of 7DS only if the transmission range is high. In these scenarios, aggressive power conservation degrades the hit rate by 22% for a pure P2P system and by 16% if infostations are used. This extends the results from [13], which reports a significant impact of power conservation for all transmission ranges in the transient phase. Consistently for both buffer sizes, Figure 5 shows that FIS-P* outperforms P for high transmission ranges, in particular if power conservation is used aggressively. Comparing Figure 5 (a) to (b) and (d) to (e), respectively, we find that an OFF period of 90% yields approximately the same hit rate as reducing the transmission range from 230m to 115m. That is, using OFF periods of 50%, we have a hit rate of 0.66 for high transmission ranges and small buffers with the 7DS variant P. Extending the OFF interval to 90% even reduces the hit rate to 0.57. In contrast, using a medium transmission range and no power conservation yields a hit rate of 0.62. Recall that this reduction of the transmission rage equals a reduction of transmission power from 281.8 mW to 17.6 mW, i.e., a reduction of 93.7%. Power conservation as proposed in [13] with an OFF period of 0.9 will reduce both the energy required to listen to remote queries and the number of generated response messages by 90%. Nevertheless, the node has to send an identical number of query messages. Since a reduced transmission power will save 93.7% of the energy consumption for each message, reducing the transmission range saves significantly more power than using long OFF periods. Thus, we conclude from Figure 5 that reducing the transmission power yields a much more efficient approach for reducing the power consumption than switching off the wireless interface.

5.5 Impact of Query Repetition

In a last set of experiments, we extend the results of [13] by calculating the average number of repetitions of a query for each data item in the long run. We consider a maximum number of repetitions of three, five, and seven. In this experiment, we use the 7DS variant NP, which achieves the best performance as shown in Figure 4. The results for small and large buffers and high, medium and low transmission ranges are shown in Figure 6. Figures 6 (a) and (d) show that for high transmission ranges the average number of query repetitions is significantly below the maximum number regardless of the buffer size and the popularity of a data item. For low transmission ranges and small buffers, the number of maximum repetitions is almost reached for 60% to 80% of the data items as shown in Figure 6 (c). Recall from Section 5.3 that in these scenarios, 80% of the items are stored in less than 5% of the buffers. Larger buffers significantly reduce the average number of repetitions at low transmission ranges as shown in Figure 6 (f). Here, the average number of repetitions is below six even for unpopular data items if the maximum number of repetitions is set to seven.

Figure 6 implies the conclusion that query repetition is particular useful for small buffers. To gain further insight into the tradeoff between buffer size and the maximum number of query repetitions, we calculate the hit rate as a function of the buffer size for high, medium and low transmission range. Here, a query is counted as a hit if a response is received before the maximum number of repetitions is reached. We find that for small buffer sizes, aggressive query repetition increases the hit rate by up to 61% depending on transmission range. However, hit rate is a concave function of the number of repetitions, so increasing the number of repetitions further than seven will not yield significant performance gain. Figure 7 (a) confirms the conjecture from Figure 6 that aggressive repetitions will not increase the performance of 7DS for high transmission ranges and large buffer sizes. That is, if the buffer size exceeds 120, a maximum number of three repetitions perform almost equal to five or more repetitions. To achieve a reasonable hit rate beyond 0.9, we recommend a maximum number of five repetitions for high transmission ranges and buffers smaller than 120 entries. For larger buffers, three repetitions suffice. If the transmission range is medium to small, a hit rate of 0.9 is not reached regardless of the maximum number of repetitions. Here, the buffer size must be larger than 100 items for a medium transmission ranges and larger than 155 for a small transmission ranges. In both cases, a maximum of five to seven repetitions is required to achieve a reasonable performance.

CONCLUSION

We presented an approximate analytical modeling approach for analyzing the performance of epidemic dissemination of information in MANET in the long run. The modeling approach explicitly represents the spread of multiple data items and finite LRU buffers at mobile devices. Previous work [13], [14] just considered the transient behavior of a single data item, assumed unlimited buffers and required off-line simulation for determining model parameters. We showed how to derive performance models for four variants of 7DS [13], a well-known system for P2P data sharing in MANET.

We presented a comparative evaluation of four 7DS variants P, NP, FIS-P* and FIS-NP* as well as investigated the impact of power conservation and query repetition. We found that neither the transmission range nor the selected variant of 7DS has a significant impact on the fraction of dataholders in the long run. However, for high transmission ranges the selected variant of 7DS has a significant impact on the hit rate. Depending on the 7DS variant and the buffer size, hit rates between 0.48 and 0.92 can be achieved. 7DS variants NP and FIS-NP* outperform variants P and FIS-P* by up to 18%. For low transmission ranges, hit rate lies between 0.37 and 0.77 regardless of the 7DS variant. Enabling aggressive power conservation in scenarios with high transmission ranges degrades hit rate by 22% for a pure P2P system, and by 16%, if infostations are additionally provided. Furthermore, a reduced transmission range of 115 m yields higher hit rates than using long OFF periods at 230 m transmission range. Repeating queries up to seven times increases hit rate by up to 61% for small buffers and low transmission ranges. Only for high transmission ranges and large buffers, query repetition yields marginal performance gains. Thus, the derivation of an optimized power management strategy for 7DS requires the comprehensive study of the trade-off between duration of OFF periods at mobile devices, transmission range, and number of repetitions.

REFERENCES

[1] N.T. Bailey, *The Mathematical Theory of Infectious Diseases and its Applications*, 2nd Edition, Hafner 1975.

[2] C. Bettstetter, On the Minimum Node Degree and Connectivity of a Wireless Multihop Network, *Proc. 3rd ACM MobiHoc 2002*, *Lausanne, Switzerland*, 80-91, 2002.

[3] F. Chinchilla, M. Lindsey, and M. Papadopouli, Analysis of Wireless Information Locality and Association Patterns in a Campus, *Proc. IEEE INFOCOM 2004*, Hong Kong, 2004.

[4] A. Dan, D. Dias, and S. Yu, The Effect of Data Access on Buffer Hits and Data Contention in a Data Sharing Environment, *Proc. VLDB 1990, Brisbane, Australia*, 419-431, 1990.

[5] A. Dan and D. Towsley, An Approximate Analysis of the LRU and FIFO Buffer Replacement Schemes, *Proc. ACM SIGMETRICS*, Boulder, CO, 143-152, 1990.

[6] P. Eugster, R. Guerraoui, A-M. Kermarrec, and L. Massoulie, Epidemic Information Dissemination in Distributed Systems, *IEEE Computer*, **37**, 60-67, 2004.

[7] D.J. Goodman, J. Borras, N.B. Mandayam, and R.D. Yates, Infostations: A New system for Data and Messaging Services, *Proc. IEEE Vehicular Technology Conference* (*VTC Fall)*, 969-973, 1997.

[8] M. Grossglauser and D. Tse, Mobility Increases the Capacity of Ad-hoc Wireless Networks, *IEEE/ACM Trans. on Networking* **10**, 477-486, 2002.

[9] M. Grossglauser and M. Vetterli, Locating Nodes with EASE: Mobility Diffusion of Last Encounters in Ad Hoc Networks. *Proc. IEEE INFOCOM 2003,* San Francisco, 2003.

[10] X. Hong, M. Gerla, G. Pei, and C. Chiang, A Group Mobility Model for Ad Hoc Wireless, *Proc. ACM Int. Workshop on Modeling, Analysis and Simulation of Wireless and Mobile Systems* (*MSWiM 99)*, Seattle, WA, 1999.

[12] J.-Y. Le Boudec, *Understanding the Simulation of Mobility Models with Palm Calculus*, Technical Report IC/2004/53, EPF Lausanne, 2004

[13] M. Papadopouli and H. Schulzrinne, Effects of Power Conservation, Wireless Coverage and Cooperation on Data Dissemination among Mobile Devices, *Proc. 2nd ACM MobiHoc 2001*, Long Beach, NY, 117-127, 2001.

[14] T. Small and Z. Haas, The Shared Wireless Infostation Model – A New Ad Hoc Networking Paradigm (or Where there is a Whale, there is a Way), *Proc. 4th ACM MobiHoc 2003*, Annapolis, MD, 233-244, 2003.