

Floating Content: Information Sharing in Urban Areas

Jörg Ott, Esa Hyytiä, Pasi Lassila
Aalto University, Comnet

Tobias Vaegs
RWTH Aachen University, ComSys

Jussi Kangasharju
University of Helsinki

Abstract—Content sharing using personal web pages, blogs, or online social networks is a common means for people to maintain contact with their friends, colleagues, and acquaintances. While such means are essential to overcome distances, using infrastructure services for location-based services may not be desirable. In this paper, we analyze a fully distributed variant of an ephemeral content sharing service, solely dependent on the mobile devices in the vicinity using principles of opportunistic networking.

The net result is a *best effort* service for *floating content* in which: 1) information dissemination is geographically limited; 2) the lifetime and spreading of information depends on interested nodes being available; 3) content can only be created and distributed locally; and 4) content can only be added, but not explicitly deleted. We present our system design, summarize its analytical modeling, and perform extensive evaluation for a map-based mobility model in downtown Helsinki to assess the operational range for Floating Content. We finally report on our Java-based implementation on Android mobile phones.

I. INTRODUCTION

Online social network applications for sharing content, opinions, and experiences (Facebook, Flickr, YouTube, Twitter, etc.) are widespread means for people to foster their relationships irrespective of physical distance. Increasing mobile Internet use has made sharing experiences from mobile devices popular, as the authors witness daily from their acquaintances, e.g., on Facebook. Context- and location-aware services, such as digital graffiti [1], and (to some extent) Google Maps and Google Earth, have been constructed around mobile users, however, relying on services in the network infrastructure: to store and maintain data as well as to determine geographic proximity. While network-based social applications are essential to overcome distances and connect people around the world, relying on infrastructure services for location-aware services may often not be desirable:

- *Location privacy* concerns arise since the location typically needs to be communicated at least to some level of accuracy to obtain the right context information.
- *Content privacy* issues occur since shared information will be available at some “central” location and thus easily subject to censorship (a mixed blessing).
- *Connectivity* to the infrastructure becomes a prerequisite that is often limiting, especially for traveling users who may face high roaming charges, unavailability of data services, or simply no network coverage.
- *Geographic validity*: locally relevant content may be of little concern to the rest of the world; storing it in some

well-accessible location may not be of much use.

- *Temporal validity*: Information may only be valid for a limited amount of time; yet, centrally stored content rarely is tagged using expiry information, leading to the content never being deleted—yet, quite frequently never being read either (WORN, write-once, read never).
- *User identification* of some kind is usually applied to limit the amount of data posted and creates some sense of responsibility towards the service provider.

In this paper, we propose and analyze a fully distributed variant of an ephemeral content sharing service, solely dependent on the mobile devices in the vicinity using principles of opportunistic networking, built upon our earlier reported work-in-progress [2]. Any user may create content and define its geographic *origin*, and *validity radius*—that together define the *anchor zone* in which content is meaningful—plus an *expiration time*.¹ The creator’s device starts disseminating the content to its neighbors within the validity radius, as do other nodes, reducing the replication and increasing the deletion probability as a function of the Euclidean distance from the origin. As long as there are enough supportive mobile devices around in the anchor zone to replicate and store a piece of content, it *floats*. Other mobile nodes interested in some type of content will be able to obtain a copy when they get “in range” of a particular piece of content and have either a copy disseminated to them or obtain one by means of a one-hop query. When the node density becomes too low (even temporarily), the content will disappear (“sink”).

The net result is a *best effort* service for *floating content* in which: 1) information dissemination is geographically limited 2) the lifetime and spreading of information depends on interested nodes being available in the target zone of a particular content item; 3) traffic can only be created and caused locally, thus limiting DoS effects; and 4) content can only be added, but not deleted, so that security issues are kept outside the system.² We expect that such as simple mechanism will allow creating more sophisticated services on top.

¹Moreover, content may be tagged with metadata (e.g., content channels [3]) to support filtering and selective forwarding and replication.

²We expect that the validity of local information be checked at low cost, e.g., in the real world by validating that there is really free beer in a particular bar and posting annotations in case a “story” is not true. In addition, within groups content postings may be signed leveraging the users’ identities in other domains (e.g., of SIM cards [4]).

II. RELATED WORK

Concepts similar to our floating content have been introduced in research already a few years ago, such as in [1], [3], [5]–[13]. Some, e.g., [3], [7], [11], focus on distributing content over an ad hoc or DTN-like network, using the wireless network only as a kind of a cache for Internet content. They do not consider the case of managing content *purely in the wireless domain*. Others, e.g., [8], [9], consider a purely wireless case, but focus on directed communications between users (unicast or multicast). In contrast, our work is not about directed communications, but about content floating in a particular location, i.e., the applications are different.

In [5], [10] the authors present models for hovering information, which both bear similarities with our floating content model. However, our work differs from theirs in an important way. Their focus is on defining the general model, whereas our focus is on *evaluating the general feasibility* of floating content systems. By feasibility, we understand aspects such as node density, required communication range, etc. Our work explores the parameter space of floating content systems to identify under which conditions such systems are feasible.

Work in a recent publication [12] presents a system concept very similar to our work in [2]. Their focus is mainly on implementation aspects and protocols and hence complementary to the work in this paper with a focus on evaluating the operating conditions in realistic settings for floating content systems.

Another very recent paper [13] presents a system called Locus. Conceptually, the work only re-iterates already published concepts on floating content [2], [5], [10], [12], [14], [15]. Locus allows remote nodes to query for information placed in a given anchor zone, which makes it different from related work in this area. However, we believe this property is actually a severe weakness, since it allows for remote nodes to perform denial-of-service attacks anywhere in the network. In pure floating content work [2], [5], [10], [12], [14], [15] such attacks require physical presence in the targeted anchor zone, making them far less vulnerable, if not even impervious, to these kinds of attacks. Our work pre-dates that of Locus and already our published paper [2] contains a more extensive feasibility study. This paper goes far beyond that, performing an extensive evaluation of floating content in a more realistic urban setting and taking more key parameters into account.

The floating content concept is an example of delay tolerant networks. The fundamental properties of routing in such networks have been recently analyzed, see [16] and the references therein. However, the floating content service requires modeling the notion of information availability in the anchor zone. We have analytically developed analytical models for this in [15]. Here we validate the result for the so-called non-spatial model in [15] in a realistic mobility scenario.

III. FLOATING CONTENT DESIGN

We assume that all users are mobile nodes and that there is no supporting infrastructure for the system.³ They use

³Content dissemination from fixed access points would lead us towards the PodNet model [3] paired with geographically limited distribution.




Fig. 1. An anchor zone of an item, mobile nodes and their communication ranges: the content item gets replicated across and deleted from nodes as a function of the distance from the anchor point. The probability of a node carrying an item (black nodes) tends to 1 inside the anchor zone r and decreases until, after an availability threshold a , no more copies are found.

mobile phones or similar devices to communicate, so that ample storage capacity is available of which a sufficient fraction is set aside for floating content. The devices use wireless interfaces (e.g., Bluetooth or WLAN) for ad-hoc communication. Real-world traces have shown that median contact durations between mobile users, e.g., range from 15 s or less as in RollerNet [17] to some 120 s in conference environments [18].⁴ Experiments using Bluetooth for inter-device communication yielded some 10–15 s channel setup delay and a net data rate of some 50 KB/s [18]. Thus, exchanging data up to several megabytes per contact appears feasible so that floating modestly sized pieces of contents could be supported (from text messages to photos); it also suggests that the communication capacity during a contact is likely to become the system bottleneck and mobile nodes may not need to reserve a lot of storage capacity for floating content.

The devices also need to be location-aware, e.g., by using GPS, triangulation-based methods using WLAN access points or cellular base stations, or any other method offering reasonable accuracy; suitable location APIs are commonly available in modern smart phones. Since the floating content system is probabilistic, there are no strict requirements on the accuracy of positioning techniques; nodes are only required to agree on basic measurement parameters and the overall operation to determine the extent of anchor zones. Finally, nodes need roughly synchronized clocks to time out content items; both GPS and cellular networks may provide local time.

Mobile users originate (“post”) content items with a defined anchor zone and TTL (and must reside in the anchor zone when doing so). Other users are interested in these items⁵ and will accept copies to store and (probabilistically) replicate further inside the anchor zone. We explicitly allow information items to disappear from the system and provide no guarantees about their availability. If, e.g., no (or too few) nodes are around to replicate a content item and the creator leaves the anchor zone, the corresponding items may disappear. Once an item’s lifetime expires it will be deleted by all nodes.

⁴See, e.g., the CRAWDAD archive for a collection of contact and mobility traces at <http://crawdad.cs.dartmouth.edu/>.

⁵The content can be organized into groups or channels similar to [3], but we restrict our considerations in this paper to a single channel.

A. System Operation

A node generates an information item I with a certain lifetime (TTL) and assigns an anchor zone defined by its center and two radii, r and a , as shown in figure 1: r defines the *replication range* within which nodes always try to replicate the item to other nodes they encounter; a defines the *availability range* within which the content item is still kept around with limited probability while outside a no copies of the item are to be found.

If two nodes A and B meet in the anchor zone of an item I , and A has I while B does not, then A replicates item I to B . Since replication is based purely on the location of nodes, in a simple case, every node in the anchor zone should have a copy of the item. Nodes leaving the anchor zone are free to delete their copy of the item.

In practice, the replication and deletion works as follows. Consider a node A having an item I , with an anchor zone defined by center point c and radius r . Let h denote the distance of node A from point c . When node A meets another node B , A will replicate item I to B with probability $p_r(h)$:

$$p_r(h) = \begin{cases} 1 & \text{if } h \leq r \\ R(h) & \text{if } r < h \leq a \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

where $R(h) \in [0, 1]$ is some (decreasing) function that gives the probability of replication outside the replication range but within the availability range. We define the deletion probability $p_d(h)$ in a similar way, $D(h) \in [0, 1]$:

$$p_d(h) = \begin{cases} 0 & \text{if } h \leq r \\ D(h) & \text{if } r < h \leq a \\ 1 & \text{otherwise} \end{cases} \quad (2)$$

The deletion function essentially serves the purpose of early prioritization to prevent buffers from filling up; it is evaluated upon each encounter with another node. If there is still a need to free buffer space the oldest messages are discarded.

These definitions yield the area outside the replication but within the availability range of the anchor zone as a buffer zone that offers smooth degradation of availability over distance. This definition of an anchor zone with two ranges is beneficial, because it provides protection against items disappearing when nodes move outside the anchor zone for a brief moment and then return. Beyond the availability distance a (see figure 1) copies are deleted (immediately or upon encountering the next node, depending on the implementation).

In our evaluation presented below, we use $R(h) = 0$ and $D(h) = 0$, i.e., neither replication nor active deletion take place for $r < h \leq a$, rendering the buffer zone entirely passive. Note that if $r = a$ $R(h)$ and $D(h)$ become immaterial.

B. Floating Content Protocol

A floating content message m is identified by a unique message id Id_m and carries its anchor zone spec, i.e., anchor point (P_m, r_m, a_m) ; and its lifetime T_m in its headers; the content I is in the message body. We use a 4-phase protocol for exchanging floating content messages:

1) Nodes continuously send neighbor discovery beacons to discover peers.

2) After discovering a peer, i.e. receiving a beacon, a node sends a summary message that includes a vector of the content items available for replication, i.e., all items for which $p_r(h) > 0$. Per item, the vector contains Id_m , $s(m)$, (P_m, r_m, a_m) , and T_m . The summary may be limited to what fits into one MTU size packet (or otherwise to a message sufficiently small to be exchanged efficiently so that most of the per-contact communication capacity is left to exchange content); if a node has more content to share, the list of content is spread across multiple summary messages in a round-robin fashion.

3) As soon as a device is aware of what a neighbor has to offer, it requests a subset of the content items: the receiver requests all those items for which evaluating $p_r(h)$ so suggests. It may prioritize the order in which elements are replicated according to a replication policy. The replication policy determines in which order messages are replicated when two nodes meet: FIFO preserves the order in which messages were created/received at a node, RND randomizes the order, and three further algorithms use ascending order by anchor zone ($A = a$)⁶, floating volume ($V = A \times s(I)$), and the product of volume and TTL ($V \times T$); we abbreviate these algorithms as *Smallest Area First (SAF)*, *Smallest Volume First (SVF)*, and *Smallest Total resource consumption First (STF)*. Replication only takes place within r .

4) Messages containing the requested items are then exchanged until the nodes lose contact or the batch is completed. In the former case, the receiver discards all incomplete messages; in the latter, the protocol returns to phase 2).

All phases are fully bidirectional so that message exchanges take place in both ways simultaneously. The protocol does not restrict message exchange to two nodes at a time (even though some link layer technologies may). Beacons are continued throughout the message exchange process so that new nodes may be discovered while a node is already exchanging messages with another. The protocol phases and the incremental summary messages allow nodes to append messages they just received to the list of those offered to other nodes while already exchanging messages.

Nodes delete messages following one of two deletion policies when $p_d(h)$ suggests deleting a message m . With $D(h) = 0$, no messages are deleted within a . Outside a , *immediate* deletion leads to m being discarded once a node leaves a , whereas *upon-encounter* deletion performs this action only when the next node is met, allowing a node to wander out of a and back in while keeping m . From an implementation perspective, the *upon-encounter* deletion policy is more sensible as it is triggered by an external event. Deletion takes place before messages are offered to other nodes.

C. Operational Considerations

We allow the user generating a content item to define the extent of the anchor zone. We only require the user to be in

⁶At this point, we consider just the anchor zone radius and do not use the anchor zone area $A = a^2 \times \pi$.

the anchor zone at the time of creation but do not impose any limits on the anchor zone size. This may naturally cause problems because there is no incentive for users to limit the anchor zone. It would thus be easy to spam the system by inserting items with “infinite” anchor zones, limited in practice only by numeric representation of the radius in the protocol messages. But flooding items throughout the network is problematic, since it is known that flooding quickly exhausts the system resources (both channel and buffer capacity).

Without infrastructure and with potentially sparse and frequently disconnected node populations, accounting or reputation mechanisms like [19] do not appear feasible, e.g., due to their vulnerability against the Sybil attack [20] because of the lack of central authentication of identities. We therefore define a simple mechanism for resource management to discourage unlimited content distribution: at any given point, we prioritize items inversely with respect to a) their expected resource consumption and b) the distance from their anchor. We realize a) by means of the three resource-aware replication policies SAF, SVF, and STF that take into account the anchor zone size A , $A \times s(I)$, and $A \times s(I) \times T_m$, respectively, that we evaluate below. We realize b) by means of $p_r(h)$ and $p_d(h)$.

When a node needs to replicate or store more items than it is able to, it gives preference to items with the smallest anchor zones and evicts the items with the largest anchor zones. It is possible to post items with very large anchor zones, but their availability is likely to be quite low.

It is possible for a spammer to move and create items with small anchor zones everywhere. We provide no mechanisms to guard against this. Instead, we consider the effort of having to move around to be a sufficient deterrent to most spammers.

Our approach has the advantage that it does not require security infrastructure or any degree of mutual trust. Of course, nodes could spam their immediate neighborhood or perform DoS attacks, but they could achieve the latter using physical layer interference anyway.

IV. ANALYTICAL MODEL

Analytically the floating content service has been studied in [15]. A fundamental objective of this work has been to derive the so-called *criticality condition* guaranteeing the availability of the information within the anchor zone with a high probability. This depends heavily on the mobility patterns of the users, especially inside the anchor zone. Next we present the so-called non-spatial model from [15].

Consider the anchor zone just as a locale where the nodes enter and then spend some time and finally exit. Inside the locale, the nodes are exchanging information between each other. We assume that the nodes move randomly within the locale and the time spent there is relatively long so that the exact points of entrance and exit do not have any bearing. The whole population is assumed “well mixed” so that the proportion of information-carrying nodes is everywhere constant.

During the sojourn time, a node encounters randomly other nodes. Consider a system consisting solely of two nodes staying permanently in the locale and denote by ν the frequency

at which they come in contact with each other, i.e., within each others’ transmission ranges. Now, if the total population of nodes in the locale is N , then there are $\frac{1}{2}N(N-1) \approx \frac{1}{2}N^2$ pairs and the total rate of encounters is $\frac{1}{2}N^2\nu$. The fraction $2p(1-p)$ of these encounters are such that a node without the information gets it, and the total rate of such events is $p(1-p)N^2\nu$. This is the rate at which the size of the population of information-carrying nodes tends to increase. Let $1/\mu$ denote the node’s mean sojourn time in the anchor zone. Then the total exit rate of nodes is $N\mu$, and in particular, the exit rate of information-carrying nodes is $Np\mu$, so that their net growth rate is

$$N \frac{d}{dt} p = N^2 p(1-p)\nu - Np\mu. \quad (3)$$

In equilibrium, the two terms on the right hand side are equal leading to the stationary value $p^* = 1 - \mu/(\nu N)$. Existence of a positive solution $p^* > 0$ requires that,

$$N\nu/\mu > 1. \quad (4)$$

This is the *criticality condition*. Note that the quantity on the left hand side is the *average number of encounters a randomly chosen node experiences during its sojourn time*. By considering the sign of the derivative (3) one easily sees that the solution is stable. Below the criticality, the derivative is everywhere negative and the solution is driven to $p = 0$, i.e., the information cannot be sustained even in the fluid model. In addition, the number of information-carrying nodes, Np , must be large to avoid information from disappearing due to stochastic fluctuations.

Different mobility models can be analyzed with respect to (4), see [15]. However, here we focus on the above non-spatial model, as its underlying assumptions are closest to the realistic mobility scenario considered in this paper.

V. SIMULATION ENVIRONMENT

While our initial evaluations [2], [15] focused on basic feasibility in a static and two idealized mobility settings (using random-waypoint and Manhattan mobility), for this paper, we choose the more realistic Helsinki City Scenario (HCS) [21] based upon a city map (4500 m \times 3400 m) of downtown Helsinki for evaluation (see figure 2). HCS models two types of nodes: most roam the city area following streets and walkways when moving to randomly chosen points on the map following a shortest path using pedestrian (0.5 – 1.5 m/s) or car (10 – 50 km/h) speeds; some follow a set of three predefined routes as trams with their own characteristic speed (25 – 35 km/h).

We consider this kind of tourist-style mobility of restless users as we are only interested in their local interaction within different anchor zones and do not try to achieve city-wide routing functionality. As a consequence, nodes can be treated as “new” once they left an anchor zone and, therefore, daily routines and social relations would not add to the model (beyond possible density fluctuations over the day) so that we deliberately focus on simpler modeling.

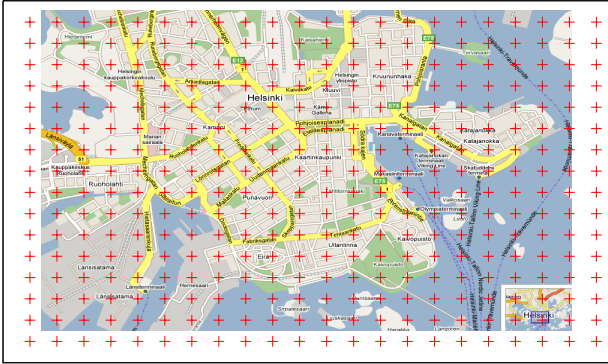


Fig. 2. Simulation area and reference anchor locations every 200 m.

We use the ONE simulator [21] for which we implemented: 1) a *FloatingApplication* class initiating content postings at different locations with configurable parameters (a , r , TTL, size); 2) a *FloatingRouter* that implements the five replication and two deletion policies described in section III-B above; and 3) Two classes *FloatingAppReporter* and *ContactConditionReport* provide information about the node density and encounter frequencies as input parameters to our analytical model. The class *FloatingMessageReport* delivers statistics about creation, replication, and deletion of floating messages.

We create six different mobility scenarios as shown in table I: We use 2 Mbit/s data rate for the wireless links for 10 m (Bluetooth) and 50 m (WLAN ad-hoc mode) radio ranges and vary the number of nodes from 126 (small) to 252 (medium) to 504 (large scenario). Unless stated otherwise, we use 5 MB buffer space per node, except for trams that use always 50 MB.

Reference	radio range	#nodes	ped's	cars	trams
$S10(r, a)$	10 m	126	80	40	6
$S50(r, a)$	50 m	126	80	40	6
$M10(r, a)$	10 m	252	160	80	12
$M50(r, a)$	50 m	252	160	80	12
$L10(r, a)$	10 m	504	320	160	24
$L50(r, a)$	50 m	504	320	160	24

TABLE I
OVERVIEW OF SIMULATION SCENARIOS

We conduct most of our simulations with two different anchor zone sizes: $a = r \in \{200 \text{ m}, 500 \text{ m}\}$. To assess the feasibility of floating content in downtown Helsinki, we choose fixed anchor locations every 200 m horizontally and vertically across the entire simulation area. The resulting 352 anchor points are shown as crosses in figure 2.

We run all simulations over a period of 24 hours, use a message TTL of 1h and correspondingly a cooldown period of 1 hour. We report the mean values for a given scenario a) across all anchor zones in the entire simulation area, using an average weighted by the number of messages generated per anchor zone and b) restricted to a *core area* of $1.8 \times 1.8 \text{ km}^2$ that includes diverse terrain (including water front) but where we do not expect any boundary effects due to simulations.

Figure 3 shows the mean number of node contacts per anchor zone per hour for the scenario $S50(500, 500)$ (that we

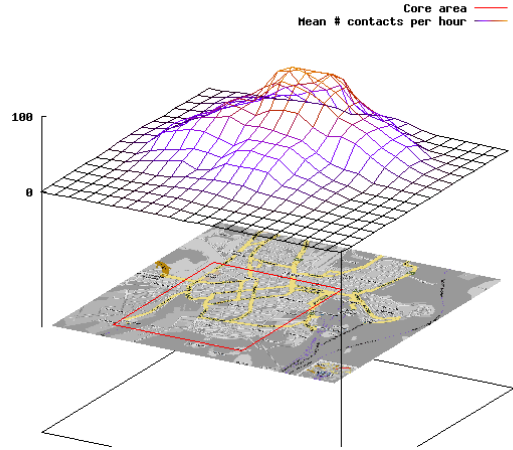


Fig. 3. $S50(500, 500)$: Mean number of contacts per hour per anchor zone

will also use in section VI-B). The core area is indicated as a square on the map. Expectedly, across all scenarios, the contact density increases with radio range, the number of nodes, and the anchor zone size from some 230 contacts/h for the densest spot in $S10(200, 200)$ to 17,500 contacts/h in $L50(500, 500)$.

VI. EVALUATION

We carry out two series of simulations: an initial set of simulations uses minimal message sizes and sufficiently large buffers to establish under which of the above mobility models floating content is feasible and how it matches to the predictions of our analytical model (VI-A). We then evaluate the floating content concepts and different algorithms for more diverse parameter settings (VI-B).

A. Feasibility Considerations

We run simulations individually for each of the 352 anchor points with two different radii for the six mobility scenarios introduced above. Individual runs are used to isolate anchor zones from the impact of neighboring ones. In these scenarios, we use tiny messages so that message transfers are never aborted because nodes move out of range. Each node generates a message when it comes within $0.5r$ of the anchor point, but message generation per node is limited to one message per 30 min (this yields from 1350 messages for $Sx(200, 200)$ to 9350 messages for $Lx(500, 500)$ per 24 hours simulation time in the busiest spots. Next we consider the spatial distribution of the probability that content floats for up to one hour for all $S10$ scenarios. We additionally validate criticality condition (4), which defines a condition under which a message will stay available with a high probability. To this end, the parameters for the mean number of customers N , the mean sojourn time $1/\mu$ and the mean contact rate ν for each anchor zone were also estimated from the simulation data.

Figure 4 depicts the value of the criticality factor $N\nu/\mu$ as well as the spatial probability that the content is available up to one hour in each anchor zone. Left column corresponds to the case with $r = a = 200 \text{ m}$, and the right column to $r = a = 500 \text{ m}$. In both columns, the left figure shows

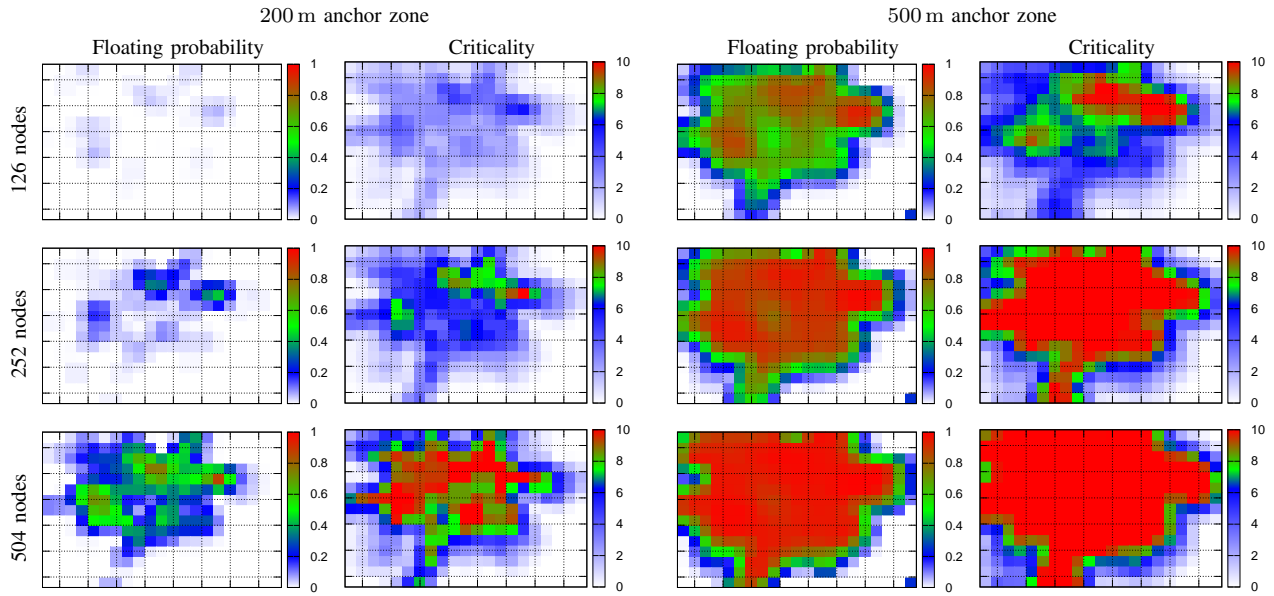


Fig. 4. The observed floating probability and the corresponding criticality condition in the whole area per anchor location for $r = a = 200$ m and $r = a = 500$ m when using 10m range Bluetooth radio. From top to down, the number of nodes is 126 (*S*10), 252 (*M*10) and 504 (*L*10). Similarity between the estimate (criticality condition) and the observed floating probability is striking in each case.




Fig. 5. Floating duration as fraction of TTL: from left to right scenarios *S*, *M*, and *L*.




Fig. 6. Floating duration for the $1.8 \times 1.8 \text{ km}^2$ core area excluding borders and the Sea as fraction of TTL: from left to right scenarios *S*, *M*, and *L*.

the observed floating probability (fraction of messages that survived the TTL of 1h), and the right figure the corresponding criticality factor. From top to down, the number of nodes increases: 126 nodes, 252 nodes and 504 nodes.

When the criticality factor is less than 1, or close to one, the expectation is that the messages will not stay alive (i.e., loosely speaking floating does not happen). As can be seen, the white/light blue areas in Figure 4 match nicely (floating probability vs. criticality). Also, by looking in more detail at the case with 126 nodes and $r = a = 200$ m we observe that the small light green areas, reflecting a somewhat better chance of floating, nicely correlate with the small islands in the corresponding floating probability figure. Increasing the number of nodes to 504 raises the value of the criticality

factor substantially above 1 (i.e., the criticality threshold) in some areas and it is reflected in a corresponding increase in the floating probability. Finally, when $r = a = 500$ m (right column) the criticality condition is well above the threshold almost everywhere except in the areas corresponding to the Baltic sea, which also closely resembles the pattern of the corresponding very high floating probability estimates.

Using a larger radio range yields further improvement as shown in Figures 5 and 6, especially for larger r and a . These latter two figures compare the mean floating time as a fraction of TTL for a piece of content across all scenarios considering the entire simulation area (except for places in the Sea where there are no messages generated) and only the core area, respectively. The mean values are weighted by the number

of messages (running from a few tens to some 17,000 per hour). We find many scenarios for which floating is basically feasible, i.e., content availability $p(t = T_{ttl}) \gg 0$. Especially the large anchor zones show a trend towards slowly decaying $p(t)$. This is an indication that the system operation point is well above criticality which ensures that messages even in the simulated stochastic system will float with a high probability for a long time, say tens of hours. We also see that in those cases when messages do not float (but *sink*), the majority sinks during the first 10–20% of their lifetime. This is an important finding as a system might be able to warn a user still close by that her message is likely to fail.

B. Characterizing Floating Content Performance

We choose scenario $S50(r, a)$ (with $r = a = 500$ m in most cases) for further evaluation as this provides on one hand sufficient floating success rates for messages to realize best-effort content sharing, yet leaves room for improvement.

One important element is understanding where limitations of floating performance come from. Given an encounter pattern, opportunistic communications can be limited by buffer size or by communication capacity when nodes encounter each other (determined by link rate and contact duration). We vary the buffer sizes available to the 120 mobile nodes and plot the floating success rate over the message size in figure 7a. We find that for messages smaller than 100 KB the system is limited by the encounter capacity, for large messages buffer space becomes an issue. We also find that 50 MB message buffer is sufficient for the system not to be buffer-limited.

In a separate set of simulations run (not shown), we chose message and buffer sizes and the link rate so that the entire buffers of two nodes could be exchanged even in the shortest possible contact (0.1s). We found that a buffer capacity of 100 messages was sufficient to maximize floating performance.

The floating content algorithm provides for two distance parameters, the *replication* radius r and the *availability* radius a . While all simulations so far assumed $r = a$, we now evaluate the impact of $a > r$. In two cases, we keep r constant at $r \in \{200 \text{ m}, 500 \text{ m}\}$ as above and vary a in $[r, r + 500 \text{ m}]$ in steps of 100 m. In the third case, we keep $a = 500$ m and vary r in $[100 \text{ m}, 500 \text{ m}]$. Figure 7b shows the resulting floating performance plotted over the difference $a - r$. We find that, for constant a (first curve), varying r has only a marginal effect; i.e. a small replication area can be sufficient as long as content is not deleted too quickly. This is confirmed by the third curve: even with a small r , increasing a quickly improves performance significantly which approaches 1. When r is already large (second curve), increasing a further provides additional gain, but this is not as pronounced.

We also experimented with two different deletion policies, immediate and upon-encounter using the same choices of r and a as shown in figure 7b. We find that the floating performance shrinks by up to 0.3–0.4 when applying immediate deletion compared to upon-encounter deletion. This shows the importance of keeping content even when outside the anchor zone. This is in line with practical implementation

considerations: a node would rather act upon an external trigger (such as encountering another node) than permanently scan its buffers for data items to expire.

Finally, we simulate a mixed scenario based on $S50$ with 50 MB buffers: the nodes move as before and generate content items when they are inside the *core area* at mean rates of 1 (single), 2 (double), and 4 (quad load) messages per node per hour. Each generated message is anchored at the generating node’s present location. Anchor zone sizes $r = a$ are chosen from $[500 \text{ m}, 2000 \text{ m}]$ in 100 m steps, message sizes from $[100 \text{ KB}, 1000 \text{ KB}]$ in 100 KB steps, and TTLs from $[30 \text{ min}, 180 \text{ min}]$ in 30 min steps, all uniformly distributed. We use upon-encounter deletion and simulate the five different replication policies introduced in section III-B: FIFO, RND, SAF, SVF, and STF.

Figure 7c shows the overall floating performance across all content items. We observe that the single load already appears to saturate the floating capacity of the network and that increased load leads to performance degradation. We also note that with lessor saturation, the impact of the replication algorithm is less pronounced. We dissect the impact of the replication algorithm further in figure 8, plotting the floating performance as a function of message size, TTL, and anchor zone. Across all plots, we can see that the two simple replication algorithms, FIFO and RND, perform worse than those considering at least one property of the content items. The difference becomes more pronounced with increasing load and their respective prioritization become visible: we can influence that less resource consuming messages (smaller size or TTL) shall be preferred. The impact of conscious replication is most visible for the different anchor zone sizes. FIFO and RND do not support small anchor zones well. But also the algorithms aware of r and a do not manage to obtain similar performance across all anchor zone sizes because their respective radius (paired with the mobility model) determines the number of contacts available for replication.

Overall, SAF (taking into account only a) performs best across all anchor zones while not adding bias towards smaller or shorter-lived messages. Adding conscious preference for smaller or short-lived content pieces comes at the expense of messages with small anchor zones, simply because we presently use a non-weighted scalar product of the message parameters for prioritization: a large (or high TTL) message with a small anchor zone will thus be treated similarly to a small (or low TTL) message with a large anchor zone. Experimenting with different weights would be subject to future investigations; but before attempting to determine “optimal” weights or exploring more complex algorithms, we believe that broader studies on different mobility models and loads are necessary to avoid optimizing for the wrong cases.

VII. IMPLEMENTATION

To explore the practicality of the concept we have developed a Java-based prototype for the Android™ platform for initial experimentation. The implementation uses the DTN bundle protocol [22] implemented in our DTN forwarder [23] as




Fig. 7. Floating performance in scenario $S50(500, 500)$ as a function of (a: left) the node buffer size for different message sizes, (b: middle) when varying a and r , and (c: right) for different offered loads and replication algorithms.

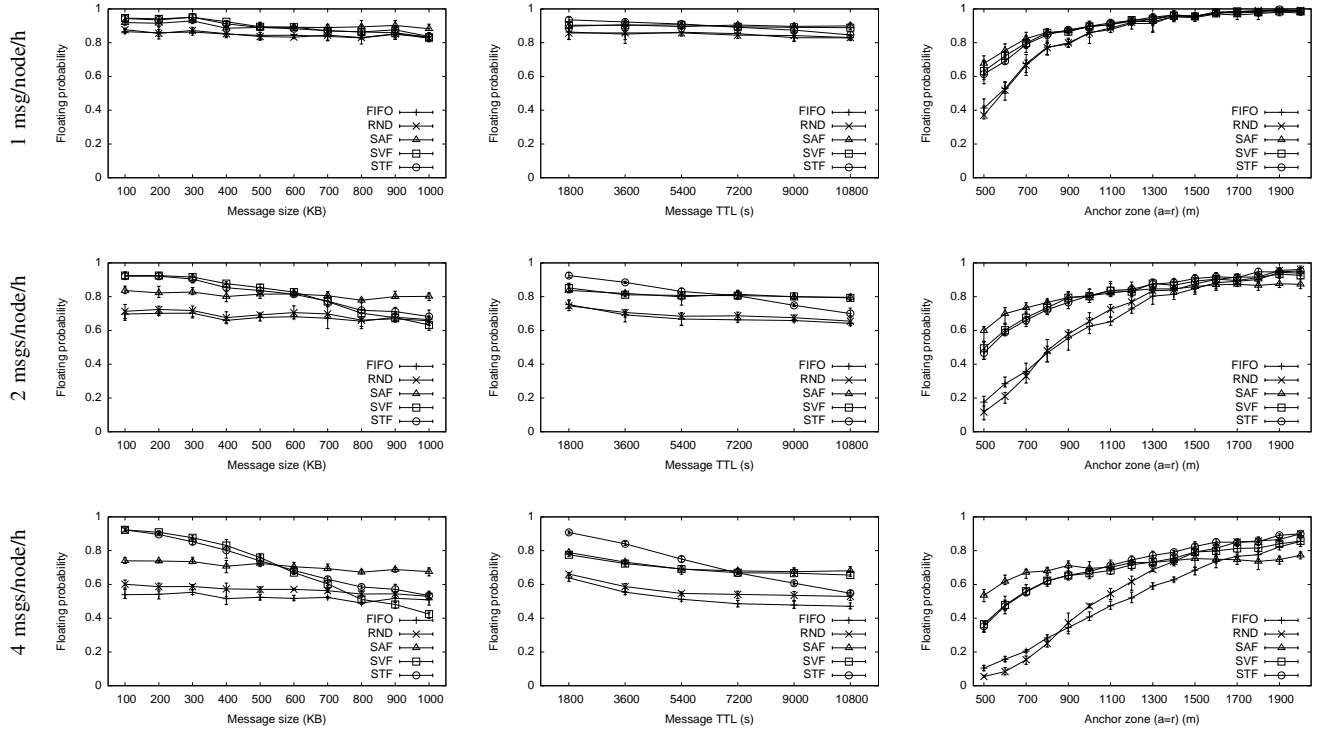


Fig. 8. Floating performance as a function message size, TTL, and anchor zone for five different replication algorithms

underlying communication layer. Each summary message, request message, and single content item is sent in its own bundle.

The implementation offers a simple interface for users to create content to float: a photo taken with the built-in camera can be posted automatically using the present coordinates and default values for r_m , a_m , and T_m . Floating content is stored in a dedicated directory of the flash memory. The user is alerted about received content and can browse the items.

We use WLAN for communication between devices, both in infrastructure and ad-hoc mode. Devices recognize each other's presence by means of an ongoing IP neighbor discovery process and then exchange their summary messages containing only the IDs and sizes of content associated with anchor zones in which the devices are currently located.

We assume the anchor zones of the exchanged information items to be large compared to the communication range of

the devices. Therefore, from being aware of a communication peer's existence a device can reason that the location of this peer is approximately the same as its own, and inaccuracies of positioning systems can be considered negligible. The offering node orders the information in the summary messages so that items are prioritized when they are small, short-lived and for a smaller anchor zone, according to the SAF, SVF, or STF replication policies. The receiver of summary messages from a neighbor will determine which items it does not yet have and is able to receive. It then requests those items it does not yet have and that fit the local storage capacity and policy, ordered following the same prioritization (we presently assume all devices to be configured with consistent policies).

We choose MIME as content format for the floating messages, as we can easily encode data of arbitrary types in them with reasonable overhead. Thus, users can send written text, taken pictures or videos, recorded sound or even attached

binary files by embedding it into the floating information item.

Each node maintains an infobase stored in flash memory including all available information. A buffer manager runs in the background to make sure that this infobase does not exceed the configured storage capacity. In need of space, expired items, items whose anchor zone the node has left ($h > a$), and those whose anchor point is furthest away from the device's current location are removed from the infobase in this order.

While the basic protocol is operational for pairs of homogeneous devices (HTC Desire), support for a wider platform diversity (e.g., Nokia N900) and larger scale tests and measurements of the floating content behavior in practice are subject of our present work. Evaluation in real-world settings will require overcoming some platform limitations (such as lack of proper multicast support on Android in WLAN ad-hoc networks). Future work will especially include exploring the power consumption implications of continued location tracking as well as neighborhood sensing to determine appropriate tradeoffs for floating content operation.

VIII. CONCLUSION

We have presented the concept of *Floating Content*, a fully distributed variant of an ephemeral content sharing service, solely dependent on the mobile devices in the vicinity that uses principles of opportunistic networking. Our evaluation using the Helsinki City Scenario suggests that floating content in urban environments can be feasible even with modest numbers of nodes supporting this application. Sufficiently large anchor zones are required for the node densities we investigated so that, while geographically limited sharing works, tightly constraining the availability may not be feasible unless the node density is very high. Indicatively, it seems that anchor zones should span multiple blocks in a city for sufficient replication, but a systematic exploration is for further study. We also find that—expectedly—floating content has limitations in border areas across all node densities when the movement of nodes into and out of an area is restricted, e.g., at the waterfront.

We could also confirm the predictions of our analytical model for the probability that content would float in a given location. This is particularly important as the input parameters to our model can be estimated locally (possibly in cooperation with other nodes) so that a predictor could be calculated that can indicate the floating expectation to the user or an application. This is one direction of further study.

While we are aware of the limitations of synthetic mobility models (and of HCS), we believe that macroscopic mobility and social context and interaction patterns may be of lesser significance for localized content sharing—unless selective support based upon (closed) groups or content channels come in—so that we expect our findings to hold for other scenarios.

Our current and future work includes on one hand a validation of our analytical models and algorithms for a broader set of mobility models and mobility traces with diverse node densities and system loads. On the other hand, we are targeting experimentation in lab scenarios and in various places of Helsinki to gain an initial understanding of floating content

performance in different environments. In the mid-term, it will be interesting to understand which more sophisticated applications could make use of such a best-effort service.

REFERENCES

- [1] S. Carter, E. Churchill, L. Denoue, and J. Helfman, "Digital graffiti: public annotation of multimedia content," in *Conference on Human Factors in Computing Systems*, 2004, pp. 1207–1210.
- [2] J. Kangasharju, J. Ott, and O. Karkilahti, "Floating Content: Information Availability in Urban Environments," in *Proc. of IEEE Percom 2010, Work in Progress session*, March 2010.
- [3] V. Lenders, M. May, G. Karlsson, and C. Wacha, "Wireless ad hoc podcasting," *ACM/SIGMOBILE Mobile Comp. and Comm. Rev.*, 2008.
- [4] N. Asokan, K. Kostianen, P. Ginzboorg, J. Ott, and C. Luo, "Applicability of identity-based cryptography for disruption-tolerant networking," in *Proc. 1st Intl. ACM MobiSys Workshop MobiOpp*, 2007.
- [5] A. Villalba Castro, G. Di Marzo Serugendo, and D. Konstantas, "Hovering information: Self-organizing information that finds its own storage," School of Computer Science and Information Systems, Birkbeck College, London, UK, Tech. Rep. BBKCS707, Nov. 2007.
- [6] D. Corbet and D. Cutting, "Ad loc: Location-based infrastructure-free annotation," *ICMU 2006*, 2006.
- [7] I. Leontiadis and C. Mascolo, "Opportunistic Spatio-Temporal Dissemination System for Vehicular Networks," in *Proc. 1st Int. ACM MobiSys Workshop MobiOpp*, 2007.
- [8] W. Gao, Q. Li, B. Zhao, and G. Cao, "Multicasting in delay tolerant networks: A social network perspective," in *Proceedings of the tenth ACM international symposium on Mobile ad hoc networking and computing*. ACM New York, NY, USA, 2009, pp. 299–308.
- [9] P. Hui, J. Crowcroft, and E. Yoneki, "Bubble rap: Social-based forwarding in delay tolerant networks," in *Proc. ACM MobiHoc*. ACM, 2008, pp. 241–250.
- [10] A. Wegener, H. Hellbruck, S. Fischer, C. Schmidt, and S. Fekete, "AutoCast: An adaptive data dissemination protocol for traffic information systems," in *2007 IEEE 66th Vehicular Technology Conference, 2007. VTC-2007 Fall*, 2007, pp. 1947–1951.
- [11] G. Karlsson, V. Lenders, and M. May, "Delay-tolerant broadcasting," *IEEE Transactions on Broadcasting*, 2007.
- [12] O. R. Helgason, E. A. Yavuz, S. T. Kouyoumdjieva, L. Pajevic, and G. Karlsson, "A mobile peer-to-peer system for opportunistic content-centric networking," in *ACM SIGCOMM MobiHeld Workshop*, 2010.
- [13] N. Thompson, R. Crepaldi, and R. Kravets, "Locus: A location-based data overlay for disruption-tolerant networks," in *Workshop on Challenged Networks*, Chicago, IL, Sep. 2010.
- [14] O. Karkilahti, "Distributed location-aware hovering information systems," Master's thesis, University of Helsinki, Oct. 2009.
- [15] E. Hyttiä, J. Virtamo, P. Lassila, J. Kangasharju, and J. Ott, "When does content float? characterizing availability of anchored information in opportunistic content sharing," Aalto University, Department of Communications and Networking, Finland, Tech. Rep., July 2010.
- [16] P. Jacquet, B. Mans, and G. Rodolakis, "Information propagation speed in mobile and delay tolerant networks," *CoRR*, vol. abs/0903.1157, 2009.
- [17] P. U. Tournoux, J. Leguay, F. Benbadis, V. Conan, M. D. de Amorim, and J. Whitbeck, "The Accordion Phenomenon: Analysis, Characterization, and Impact on DTN Routing," in *Proc. IEEE INFOCOM*, April 2009.
- [18] A. K. Pietiläinen, "Experimenting with Opportunistic Networking," in *Proc. of the ACM MobiArch Workshop*, June 2009.
- [19] S. Buchegger and J.-Y. Le Boudec, "A robust reputation system for p2p and mobile ad-hoc networks," in *Proceedings of Workshop on Economics of Peer-to-Peer Systems*, Cambridge, MA, Jun. 2004.
- [20] J. R. Douceur, "The Sybil attack," in *Proceedings of International Workshop on Peer-to-Peer Systems*, Cambridge, MA, Mar. 2002.
- [21] A. Keränen, J. Ott, and T. Kärkkäinen, "The ONE Simulator for DTN Protocol Evaluation," in *SIMUTools '09: Proceedings of the 2nd International Conference on Simulation Tools and Techniques*. New York, NY, USA: ICST, 2009.
- [22] K. Scott and S. Burleigh, "Bundle Protocol Specification," RFC 5050, November 2007.
- [23] P. Ginzboorg, T. Kärkkäinen, A. Ruotsalainen, M. Andersson, and J. Ott, "Dtn communication in a mine," in *Proceedings of the 2nd Extreme Workshop on Communications*, September 2010.