===
==
SIGCOMM 2010 Review #53A
Updated Sunday 14 Mar 2010 9:53:44pm PDT

Paper #53: Taming User-Generated Content in Mobile Networks via Drop Zones
----

Overall merit: 4. Weak reject
Reviewer expertise: 4. Knowledgeable (I am knowledgeable about the problems and solutions in the area)
Evaluation confidence: 5. Excellent

===== Paper summary =====

Using cellular network traffic traces, the authors determine that most users upload their content from one of three dominant locations and that the majority of uploaded content is uploaded in a deferred manner. Thus, they propose that network capacity be upgraded in a few select locations identified as "drop zones" from which all users upload their content as they go about their daily business.

===== Paper strengths =====

Measurement study of cellular network traces identifying when and where users upload data is interesting.

===== Paper weaknesses =====

They do not make a compelling case that the current architecture does not work effectively.
The case for their proposed drop zone-based architecture is weak.

===== Comments for author =====

I enjoyed reading about the measurement study of user upload behavior. The fact that majority of uploads are deferred and that uploads happen at a few locations such as a given user's home/workplace matches well with ones intuition.

However, the leap from the measurement study to a drop-zone based architecture is not at all convincing. I agree with your observation that since uploads happen from most (all) locations, the *aggregate* bandwidth requirements is impacted at most (all) cell sites but is this really a problem -- high aggregate bandwidth requirement spread over a long time interval (e.g. all night) can easily be satisfied!   Since cell sites are provisioned for a given peak capacity and there is lot of idle capacity available at nights from homes, how big of a problem is this really? How about a simpler alternative where uploads are staggered/scheduled throughout the night from these phones -- will this still cause a bottleneck at today's cell site provisioning level?

Also, instead of deploying/adding capacity at drop-zone base stations, operators could consider alternative mechanisms (e.g. pricing based incentives) to offload these uploads from their congested cells to WiFi-based hotspots.

In the alternate drop zone-based architecture that you propose, the scheduling of all uploads has to be accomplished in the relatively short interval that the users pass these drop zones (compared to all night that is available at their home sites) -- wouldn't that be more expensive? Also, what would such an architecture do to the battery life of the device since the users may prefer to do their uploads from home when the phone is charging rather than during the day.

In your evaluation, you assume 75Mbps throughput at cell sites -- 4G LTE trials today see only up to 50Mbps max and on average, depending on user location, 12Mbps downlink and 5Mbps uplink. So your numbers appear to be off by an order of magnitude for uplink.

Finally, the whole proposal of deploying drop-zones is for optimizing uploads of user generated content -- how about downloads? If cell capacity needs to be upgraded at most locations because of bottlenecks on the downlink, couldn't these upgraded base stations be used in a time-multiplexed manner to cater to these deferred uploads?

===

SIGCOMM 2010 Review #53B
Updated Tuesday 16 Mar 2010 7:45:34am PDT

Paper #53: Taming User-Generated Content in Mobile Networks via Drop Zones

----

Overall merit: 6. Borderline accept
Reviewer expertise: 4. Knowledgeable (I am knowledgeable about the problems and solutions in the area)
Evaluation confidence: 4. High

===== Paper summary =====

This paper uses a measurement study to justify a basic but interesting idea: given that users uploaded modestly old (approximately one day old) content, they can be prompted to upload the content at locations with greater capacity.

===== Paper strengths =====

The paper offers a clever, albeit straightforward idea. Through a user study and a bit of intuition, users can be incentivized to better utilize the existing infrastructure.

===== Paper weaknesses =====

The idea while clever, is fairly simplistic. Further, the decision on where to place any network capacity and what kinds of cellular-based services to offer have are only modestly driven by technical issues. For example, AT&T's decision to spend 3B to upgrade its North American network is as much driven by publicity as anything else. This paper really only describes an engineering solution that prioritizes what locations to upgrade first. This is a paper that could have made its point in much less than 14 pages.

===== Comments for author =====

The fact that this paper is driven by a data set is a very nice component.

This is not quite a "new cellular network architecture."

Otherwise, the paper seems reasonable, it just does not have the kind of contribution to be a Sigcomm paper.

==

SIGCOMM 2010 Review #53C
Updated Wednesday 17 Mar 2010 12:13:04am PDT

Paper #53: Taming User-Generated Content in Mobile Networks via Drop Zones

----

Overall merit: 6. Borderline accept
Reviewer expertise: 3. Familiar (I am familiar with work in the area)
Evaluation confidence: 3. Average

===== Paper summary =====

The paper analyzes the data transfer workloads of smartphone users. They dataset used is impressive with 2 million users, and content information regarding several important applications. Based on these measurements, the paper proposes the use of drop zones, locations with better connectivity, and having users defer transfers until they are in a drop zone. The paper shows that this can offload significant load and not have transfers delayed for too long.

===== Paper strengths =====

The measurement study is significant and provides some valuable insights.

Very well written.

===== Paper weaknesses =====

I just couldn't appreciate the drop zone idea. Seems like an overly painful solution for a simple problem.

===== Comments for author =====

I really liked parts of this paper and really disliked others parts. The observations and even algorithms were nice. However, the system design and motivation made no sense to me.

You measure mms traffic from users, however I assume that users now use 3g connectivity for most of their significant data transfers. Does this bias the type of data being observed? Also, I would think that user behavior would be tightly linked to other user properties such as plan purchased, type of device, age of user. It would be good if the measurement part of the paper highlighted some of these issues.

Since the primary use location is home and work, why not just use 802.11 and wired connectivity. For example, how many smartphone users don't have this access in these environments. If the goal is simply to reduce average utilization of the cellular infrastructure, this seems like a simple solution. One concern might be that 802.11 is an energy hog. However, given that it is easy to identify these zones, you could enable 802.11 just when needed. This seems like a much better choice than delaying my transfer by a day or more.

Fig 5 was very difficult to follow. You do need something that provides a scenario-like view of your system operation -- but this was too painful.

It would be good to evaluate your system under a more "evolutionary" setting. Are the locations chosen for particular delivery interval a superset of those chosen for a longer interval. I.e. is there a way to incrementally add coverage that fits with your scheme? It would seem that this is possible given your description of the greedy heuristic - but I couldn't tell for sure.

Fig 9b's explanation was hard to follow.

===

SIGCOMM 2010 Review #53D
Updated Sunday 4 Apr 2010 8:56:52am PDT

Paper #53: Taming User-Generated Content in Mobile Networks via Drop Zones

----

Overall merit: 7. Weak accept
Reviewer expertise: 4. Knowledgeable (I am knowledgeable about the problems and solutions in the area)
Evaluation confidence: 4. High

===== Paper summary =====

This paper studies upload of user generated content (through MMS) in a mobile network. Justifying through measurement that the time between content creation and upload is in fact non-negligible, the proposes an architecture where a given set of base stations are selected as drop zone, and a user uploads its content only when it is inside one of this drop zone. The paper presents a drop-zone selection problem, providing a centralized and a posteriori greedy algorithm to select drop zone, which is shown to outperform a simple popularity heuristic by 10-15%, and approaches the optimal closely.

It is a very difficult paper to judge. On the one hand, it certainly generates interest, at least for the novelty of the problem addressed and observations made from an interesting data set. On the other hand, a large part of the paper is analyzing a drop zone placement problem, which seems natural to pose but is not addressed well: the solution presented is, in practice, infeasible, it does not present very impressive improvement or novelty over obvious heuristic, the claims and simulations are somewhat misleading w.r.t. base stations capacity. In conclusion, if we were to define the non-debatable contribution of this paper, it is probably not much more than presenting an interesting new problem setting. How much of that is a contribution justifying to be presented at SIGCOMM is difficult to judge. I prefer to err on the side of favoring novelty in the program, therefore I weakly recommend this paper for publication in SIGCOMM.

===== Paper strengths =====

- interesting new topic: to the best of my knowledge, this paper is the first to study upload of user generated content in a mobile setting with a delay-tolerant architecture. As pointed by the authors the notion of drop zone could be more general.
- Some of the empirical observations are interesting for reference: (1) content is uploaded by all users in a reasonably large amount of locations, whereas each user typically uploads its content in a small number of location, (2) content sent has been, for a significant fraction, created some time ago in the past.

===== Paper weaknesses =====

- The drop zone placement problem and algorithm proposed do not stand out as a significant contribution: (1) it has almost no practical interest as it essentially assumes that all future movement and uploads are known, it may serve as a guidelines or approximation, but that it not made in the paper, and this may greatly reduce the actual improvement seen. (2) it has no speciic novelty: greedy heuristic for set cover and infrastructure deployment are well known.
- The limited capacity of a base station is a fact which is poorly addressed by the paper: (1) it leads to wrong claim regarding algorithm theoretical performance, (2) it is not really tested in the simulations (as traffic volume is not well specified, capacity is essentially infinite in all setting), (3) it is not motivated on the given scenario w.r.t. 3G capacity.
- The paper is easy to read, unfortunately most of the claims are justified by a single figure, with little effort to analyze them in details or connect them together (in particular section 2 and 4). In contrast, a large amount of space (about 2 pages and a half) are spent on rather obvious observations dressed as apparent paradox.

===== Comments for author =====

General comments:

I) About the drop zone placement algorithm:

a) This algorithm may be called "a posteriori". Given that we have recorded the positions and upload of every user in a mobile network, the operators can decide which drop zone would have been the best to activate in this period.

Is that useful? yes, perhaps as a building step. But unfortunately this paper is presenting this algorithm as a practical method (comparing the direct performance of this method to simpler one). This is unfortunately not very convincing: performing 10% or even 20% better, in terms of traffic served, than an aggregate heuristic which use only basic information (popularity of base stations for uploa etc), is rather modest for such a big difference in the input.

I would recommend to at least run a few simulations were the decisions of drop zone is made in a separate training time window, and then observe what is the magnitude of the improvement, which is the only practical way I can see of using the method presented in this paper.

b) The presence of a maximum capacity per drop zone during a time slot is interesting, but it is currently misleading.

- First of all, some of the technical claims are incorrect:
p.6 "It has been
shown [29,32,41], that the worst case approximation ra-
tio achieved by our Greedy algorithm is H(s)"
With an additional constraint in terms of capacity per drop zone per time slot, it is not clear at all that a greedy heuristic presents an appromixation.

- Second, this adds a new dimension to the problem: the fact that a coverage may simply be impossible. It's easy to come up with a case where even by selecting all base stations as drop zone, the demand cannot be served due to capacity. Assuming that this is not the case (which seems like a safe assumptions), it is still possible that the way drop zone are selected influence the capacity of the system, which
(1) proves that the H(s)-approximation claimed above does not make sense (what is an approximation given that the coverage can be imperfect?),
(2) indicates that sometimes greedy performs badly.
As an example, if a base stations is the most busy and a bottleneck in terms of capacity, then it seems clear that the upload which can be served elsewhere should not be served here. Alg.1 is doing precisely the opposite: it first select this base stations and then schedule to upload on it all the chunks which are the least seen on this particular base stations (which are the one more likely to be served by another base stations). For this reason one can create an example where Alg.1 is in fact (i) serving only a fraction of what could have been served, and even (ii) uses all base stations as drop zone.

I understand that these are worst cases, unlikely to be of any practical importance, but if the authors wants to claim to have any approximation or algorithmic contribution then they have to clarify this point.

- Third, and not surprisingly, this capacity is simply completely ignored in the simulations, since as shown in Fig. 10, content size may increase by more than 100 before any difference is noticed on the output of this current parameters. This indicates that
(1) although never stated, the simulations results are essentially the same as if capacity \zeta^max were infinite for all base stations. In this case, Alg.1 simplifies to a well known incremental greedy heuristic of set cover.
Note that this contradicts some statement:
p.10 "Note that the Greedy algorithm takes a global ap-
proach when determining which content should be up-
loaded from which Drop Zone. For the actual deliv-
ery, we assume that users opportunistically deliver their
postponed content upon encountering the first Drop
Zone with available capacity to deliver the content."
That does not make sense (if you have more user with possible upload, how can you claim that users

simply drops the first time they since they cannot? If this case does not occur, then why even consider a capacity as this is exactly the same as infinite capacity)
p.12 "The service provider selects the Drop
Zone with the most unused capacity to upload the con-
tent."
The paper claim that this is an "extension" of the method of this paper but this is in fact a part of what is made in Alg.1 (so it's not an extension it's already written). However, the authors are right that the current simulations are assuming infinite capacity, a case for which this extension is of no difference.
(2) In practical scenario, if the capacity becomes a limit, then it is clear that no significant reduction can be hoped for through drop zone. Not surprisingly we observe on Fig.10 that once the capacity is limiting, the number of drop zones grows linearly with the number of users.

c) The algorithm itself is not clear: it introduces unused notation (why do we need to keep track of the number of bits generated at different time slot Delta, since it is always 1 in the chunk creation). Why do we keep track of Rij_c while it would be zero whenever i is not the creation time (or otherwise it's just a difference which does not depend on c). Similarly, delta^ij_c in appendix B is not using the index i. It looks like the notions of chunks is very unclear. We hear that uploads created series of chunks, but then the size is chosen to be the maximum upload size.

II) Observations to make clearer

- The observation that users are individually highly concentrated (they spend most of their time in a small number of cells), but collectively spread out, is interesting. It is not really surprising, which is why quantifying this effect, as opposed to simply quote it, would be important.

Right now the paper does that poorly.
1) Fig.1 and Fig.2 are not attempting at using a similar metric, whereas it can be done. This would be much more convincing to present the contrast between these two scales.
2) Fig.2 ends after 400 base stations, whereas most of the schemes used later typically needs about 1,000 drop zones more or less. It would be important to know the profile of popularity. Also it seems much more appropriate here to use CDF, in order to know what fraction of the base stations are needed.
3) None of these observations, which are key to justify drop zones, are put in perspective with later results. I am surprised that, for instance, we do not see a discussion of different contents, that the number from Fig.2 are not connected to improvement made later.
p.3 "we validate the above hypothesis, that fundamen-
tally impacts the feasibility of the Drop Zone approach."
yes, and sadly this impact is simply not analyzed precisely

Some of the claims are empty:
p.3 "Independently from where a user may generate
the content, we find that with a high probability, the
user uploads the content from a certain set of locations."
This claim is not justified, you did not compare to the actual place where the content were produced. It may be that most of the content is created in these locations, or it may not be, but the paper does not prove anything here.

p.3 "In particular,
independently from the number of locations that users
visit in their daily commute, they tend to upload their

content from the top three locations."
Again, no comparison is made to make this claim precise.

p.4 "Third, the relative ratio among dif-
ferent content types stays nearly constant for most base-
stations, that further implies similar upload trends at
most locations."
No it does not imply anything. If that it true, then say it, or show a Figure, but there is no implication which you can use it to justify.

p.5 "At the same
time, we have demonstrated that large portions of user-
generated content is uploaded in a postponed manner,
but still from top locations."
again not justified, it has not been shown that the postponed content is coming from top locations.

- Similarly the fact that users upload their content later should be explained better. These are MMS and hence it may not be appropriate to sent the content immediately. It's possible to make a case for automatic upload to the cloud (to decide later to publish or not), but that point is not made. Another issue is that we do not know what is the actual fraction of messages that could be studies this way (the paper only mention a "subset" but what fractions of users/messages are included there is unclear).

III) Scenario to clarify

Considering alternative architecture for 3G is interesting. This paper presents one, and show that it's possible due to user movement. On the other hand, it does not show why it is needed.

- As a general remark, the traffic used for this study has been served on 3G, and hence it is not clear why an operator would consider an alternative. I think extrapolation like that of Fig.10 are important here to show the potential of drop zone. Currently, the paper's objective is to replace an architecture that works to provide the exact same traffic. Not uninteresting but perhaps not very convincing except for a theoretical exercice.

In particular, at no point is the capacity of LTE compared to the current capacity of 3G to make the case better.
p.3 "For ex-
ample, this could be WiMAX [15] or LTE [7], for which
base-station ranges can be roughly matched among 3G,
2.5G and WiMAX and LTE. We discuss these issues in
more detail later in the text."
You do not discuss these issues anytime later (you simply quote a number for LTE to give as input in Alg.1) that does not explain why it's appropriate.

- Another issue is that there is a very simple alternative:
p.2 "Indeed, we find that an individual
user is likely to upload ¿heavy¿ content only from a small
subset of locations, typically corresponding to his home
or work locations."
If indeed, this is the case, then having a WiFi option on your cell phone and use it only at home or work

in order to upload your content is going to be sufficient. This is likely to be possible, at no cost for the provider.
What makes drop zone more powerful? faster? extends upload service significantly? It's not 100% clear.

- p.2 "the analyzed provider can become capable of absorb-
ing 50% of user-generated content delivered in a post-
poned manner as part of the user daily movement rou-
tine."
What do you do with the remaining 50%. This is still an important question as it means you need almost the same 3G infrastructure in place.

Minor comments:

Fig.4 would be better presented with data points or steps rather than linear interpolation.

p.6 "Furthermore, in the solution, we also
allow for content chunks to be uploaded in parallel, even
if that was not really the case as per the trace."
not clear what you mean here why would a chunk be uploaded in different places at the same time

p.6 "since a user
could visit the same location at a slightly different time
next day."
or it may not visit it at all.

p.9 "and hence, almost infinite gains
can be obtained from user mobility."
not true from the Figure at least.

p.10 "As the
number of Drop Zones increases, users who are already
close to existing Drop Zones are further covered, while
the larger number of Drop Zones singles out the users
who are further away. Hence, the larger distance."
very unclear. There is no parafox. It just comes from the fact that you take an average (whereas you should take a distribution of median time of meeting the drop zone, which allows to include value infinity in the distribution.

p.11 "2 kilometers away from the given Drop Zone
placements,"
how many Base stations does this distance typically cover?

p.11 "Hence, perturbing their move-
ment slightly brings significant gains."
We see 10% improvement for 2km. Is that going to motivate people to do it?

p.12 "This is due to the increased time in-
terval (72 hours vs. 6 hours) that actually brings more

content closer to a smaller number of Drop Zones."
This seems obvious, if the delay is larger, you can hope to deliver the same content, in fact for radius=0km they are equal for this reasons.

p.12 "This is
consistent with insights from Figure 4, that shows that
around 50% of users upload their content within the
first hour"
not clear what you mean precisely, why is the amount of exactly an hour relevant here?

p.12 "This happens because the shorter postponed deliv-
ery interval incurs more Drop Zones"
This is again a fake paradox, it is obviously coming from the amount of drop zones which is the only relevant parameter here. It's not clear what this result adds to the paper at all. I think this space should be spent on strenghtening the actual contribution.

p.12 "That is not our aim and
we can only note that cellular providers are already se-
lective when upgrading infrastructure."
I really do not see the message here?

p.12 "Indeed, a
lot of services have become successful just by reducing
the effort required from the user."
Not clear the drop zone can be easily explained this way.

===

SIGCOMM 2010 Review #53E
Updated Saturday 10 Apr 2010 5:09:21pm PDT

Paper #53: Taming User-Generated Content in Mobile Networks via Drop Zones

----

Overall merit: 5. Borderline reject
Reviewer expertise: 3. Familiar (I am familiar with work in the area)
Evaluation confidence: 3. Average

===== Paper summary =====

The paper proposes a cellular network architecture where the capacity of certain locations (drop zones) is upgraded. The idea is that mobile users delay their uploads until they are within the drop zones and that the drop zones are within the natural movement patterns of a large number of users. The paper reports on measurement study which indicates that users naturally delay uploading content and significant uploads are made at least one day after generation. The paper gives algorithms for infrastructure placement which find, among candidate locations, the minimum number of drop zones under constraints of maximum delivery postponement and maximum drop zone delivery capacity. The paper reports on the performance of its algorithms.

===== Paper strengths =====

The paper is generally well written and the motivation for the work is good. The paper bases its conclusions on user behavior based on the analysis of mobility and upload patterns of two million users in a 2.5G and 3G mobile networks. The paper develops interesting insights from the analysis of its algorithms.

===== Paper weaknesses =====

The concept of drop zones based on upload delay relies on expected user behavior/desires that is not well tested and requires user behavior modification. This may not resonate with the idea of future communications at anytime and from anyplace. The paper neither proposes a better (efficient) solution to an existing problem with basically same constraints (to users and network), nor does it give a new network design with new requirements (costs) but better services (less constraints to users). The explanation of the results and the conclusions are not clear in some places.

===== Comments for author =====

I think you should address the impact of concentrating traffic in drop zones on the upstream network. What are the ramifications (positive/negative) beyond the air interface?

I did not see upfront a simple comparison of the performance of a network without any upgrades and a system with upgrades and delay varying from zero upwards. I have doubts about the advantage of concentrating the traffic vs serving it distributedly.

i think it is agood ida to discuss an incentive scenario for users to delay their uploads, for example based on a charging scheme.

The comments on figures 9(b) related to average distance add confusion rather than explanation.