

Dear Dr. Allman,

we thank you and the referees for reviewing our paper. Below we provide detailed answers to the reviewers' comments. Our responses are marked by (*). We highlight the changes in the manuscript using blue italic font.

Best regards,
Ionut Trestian
Supranamaya Ranjan
Aleksandar Kuzmanovic
Antonio Nucci

Reviewer(s)' Comments to Author:

Reviewer: 1

Comments to the Author

This paper describes a strategy for handling increases in data downloads on mobile networks called Drop Zones. A Drop Zone is an access point in a mobile network with significantly higher backend connectivity. The idea is that since users often delay uploading data to the Internet (for various reasons, which are explored in a user study in the first half of the paper), operators can take advantage of that fact by incentivizing users to wait until they are in range of a Drop Zone before uploading their content. The advantage to the operator is that they can selectively upgrade their network without upgrading all of it, and yet still deliver much of the content that users produce.

(*) We agree with the reviewer's description of the Drop Zone approach. One more aspect that deserves being pointed out is that operators already selectively upgrade their infrastructure and our architecture only points out the ones that would be upgraded first.

Overall, I thought this was a good and well written paper. The idea of moving content transfer to a smaller part of the network makes a lot of sense (caching and CDNs are certainly successful in other networking contexts, and Drop Zones are similar in a sense to these). My main comments are that 1) at times I didn't really understand what datasets you were using, and 2) the paper needs to address two trends that might interfere with Drop Zones: interactive applications and user growth. I'll address point 1 a bit later, but first point #2.

(*) We have addressed the reviewer comments regarding both issues. First, we have better described the datasets we used in each section and provided a citation to our previous work that explains the extraction process. The reviewer can also note that we incorporated the Appendix regarding the dataset in the text as he requested. With regards to interactive applications, we have acknowledged their use in the paper. With regards to user growth we have incorporated details about user activity along the

suggestions given by the reviewer.

It seems that Drop Zones are really best suited to bulk/batch processing (Flicker, uploading to YouTube, etc). However, it seems (at least to me) that a lot of bandwidth is sucked up by interactive applications. IPTV, Facetime, Skype, games, etc. So wouldn't networks have to upgrade their entire system anyways to support these applications? It seems like that could threaten the use of Drop Zones, but perhaps they are complementary. It would be great if you could address this.

(*) It is rather obvious that our trace is only a sample of user activity across applications (content uploads mainly) yet the best we could obtain at such a large scale. Indeed, as the reviewer notes, user activity is diverse, users use applications such as Skype, and Facetime to communicate in real time.

User-generated-content should not be neglected however. One can note how important such content is by looking at the Facebook usage statistics.

<http://www.facebook.com/press/info.php?statistics>

Three facts are important:

"More than 30 billion pieces of content (web links, news stories, blog posts, notes, photo albums, etc.) shared each month", "there are more than 250 million active users currently accessing Facebook through their mobile devices", and "people that use Facebook on their mobile devices are twice as active on Facebook than non-mobile users". These statistics, and our own statistics from Table 1 can give an understanding about the scale of the data we are dealing with.

Another thing to note is that by moving certain uploads to better connectivity regions, the user experience of the users that choose to use applications in poorer connectivity regions is improved.

We have included the following paragraph in Section II.B, now paragraph 7.

"It is rather obvious that our trace is only a sample of user activity across applications yet the best we could obtain at such a large scale. Indeed, user activity is diverse, users use applications such as Skype, and Facetime to communicate in real time. User-generated-content should not be neglected however. One can note how important such content is by looking at the Facebook usage statistics [8], mobile users being twice as active as desktop users."

Second, how are the benefits of selectively upgrading certain zones affected by changes

in user population? People get faster phones, people switch providers, people move around. It would be great to have a bit more of an explicit discussion of how sensitive your system is to changes in those variables. If 10% of the people use their phones twice as much, does that change things by 20%? Or 200%? Seems like there could be some unexpected affects there.

(*) As can be seen from Figure 6 and the information in the surrounding text only a small percent of users is responsible for the majority of the traffic (10% of the users are responsible for 54% of the traffic, the heavy uploaders). This is in accordance with what operators are witnessing to go on inside their networks; a minority of users doing most of the traffic. The same behavior exists for delaying content and movement.

To reflect the different behaviors in the implications section (Section IV.F.1) we have added the following 2nd paragraph.

"By considering only the heavy uploaders in the above scenario, we still require close to the same amount of infrastructure (just 8% less) to serve them in the first place without considering content increase. The reason is that the heavy uploaders are not particularly clustered in just a few locations. When considering content increase we still witness the same substantial increase in infrastructure needs at 5 orders of magnitude. So the heavy uploaders are the ones driving the infrastructure changes."

The other major thing I'd like to see addressed is being more explicit about which datasets you use. In section II.B.3, you describe a trace, than say "Our trace, different than the above..." Which trace are you using? Also something seems wrong with your data. You crawl Flickr for mobile photos, but only find < 50,000 photos over three years? You must be sampling I'm assuming--can you indicate the sampling factor here? Also the file sizes seem really small. 65KB for audio, and 191KB for video? I've got an iPhone and routinely use much more than that, though maybe that's an anomaly?

(*) A detailed description of how our trace was processed can be found in our previous work, reference [51]. In section II.B.3 we were describing the data crawled from Flickr. The sentence "Our trace, different than the above" was meant to contrast the Flickr picture uploads to the actual mobile photo uploads from our mobile trace.

Further, we are not sampling the Flickr dataset. As we have stated previously in the paper, we targeted two Flickr mobile photography groups (linked at references [9] and [15]). They contain around 50,000 photos. We did not crawl all of Flickr, just the mentioned mobile photo groups.

About the picture sizes, they are sizes of individual files. As

can be seen from the Flickr mobile photography groups, most of the pictures that users take using their mobile phone cameras are not particularly diverse in terms of colors or resolutions. They also do not have such big sizes.

We have replaced the sentence pointed out by the reviewer to: "Our mobile trace, different than the above Flickr trace".

Below is the paragraph we are referring to related to the Flickr crawling.

"In particular, we have crawled Flickr mobile photography groups where users upload pictures taken via their camera phones [9, 15]."

Line edits/nits/small comments follow:

- Intro paragraph 1: "photos, videos" should be "photos and videos"

(*) We have fixed this.

- Intro paragraph starting "Our Drop Zone approach"... how can you change where users upload their content without changing their behavior?

(*) We are not changing user behavior. As we have pointed out in the paper the Drop Zone idea is based on a mobile daemon that, when the user desires to upload content, he is queried if he can tolerate an amount of delay. If not, the content is uploaded straight away with the help of the existing infrastructure at that location. From the point of view of the user, he is done with uploading the content where and when he wants. As we have pointed out, he can be incentivized to do so by pricing schemes but this is beyond the scope of our paper. We have pointed out the area where this is addressed in the paper (Section V. paragraph 5).

"On the device we have a simple uploader application. The user can select a list of services that his files can be uploaded to (Facebook, Twitter, MySpace etc.). Whenever he wants to upload a set of files he goes to the application, selects the files, and sets a maximum delay he tolerates or if he wants them to be uploaded instantly."

- 963 base-stations. Out of how many?

(*) We agree with the reviewer that this would be an interesting

piece of information but we are not at liberty to disclose the total infrastructure that the provider has across the United States.

- Section II... the text here seemed pretty repetitive. In general, you can really tighten up the text in these first few pages. There were a lot of forward references that could probably be removed.

(*) We agree with the reviewer and we have shortened the text in this area in order to save space.

- Section II.A: "user-affine" sortof came out of nowhere--can you explain what this means in this context?

(*) We have added a short explanation for this.
"we call them user- affine since users are seen spending a majority of their time here"

- Section II.B: A footnote and appendix is not really necessary. I'd recommend incorporating Appendix A back into the text, as well as the footnote.

(*) We agree with the suggestion of the reviewer and we have proceeded to incorporate the footnote and Appendix A back in the text.

- Actually, the end of II.A could be tightened as well

(*) We agree with the reviewer and we have reduced the end of the section in an effort to save space.

- II.B.1: "Analyzing deeper these results" --> "Analyzing these results more closely, ..."

(*) We have fixed this.

- II.B.1: "home and work". What about school? Seems like most mobile users might be students?

(*) We agree with the reviewer and we have incorporated the change by referring to both work and school throughout the paper.

- II.B.2: Isn't the area of a cell already controlling for population? That would explain why there is uniformity across sizes of cells

(*) The reviewer is correct that the area controlled by a cell is somewhat related to the population. Providers indeed place cells with smaller radius in busy location. This is one of the possible explanations for such an effect and we have not claimed different in the paper.

- II.B.3: By choosing file names with dates in them, you restrict yourself to particular mobile phone vendors/models, and possibly providers as well. This creates an unaddressed and uncontrolled confounder: perhaps people who use that type of phone act differently than other users? This is especially true for Blackberry and iPhone users.

(*) We too have thought about this. One way we tried to see if this sample is indeed representative was by comparing the distribution of picture sizes for filenames with dates in them with the overall distribution of picture sizes. The two are very similar. We have added the following paragraph in the paper (Section II.B.3 end of first paragraph).

"In order to determine if this sample is representative for all the pictures, we compared the distribution of picture sizes for this sample with the overall distribution of picture sizes. The two are indeed similar."

- II.B.4: The end of the paragraph starting "Putting all the pieces together, we argue" isn't a summary, but rather just a restatement of the general idea of the paper. I think it could be removed.

(*) We agree with the reviewer and in order to save on space we have indeed removed the end of the paragraph.

- II.B.1: you say "details omitted for space considerations", yet this is a journal version. Why not include these details here? If they aren't particularly interesting, maybe just remove that discussion?

(*) The method we have used is described in one of our previous papers, citation [51]. We have placed a citation marker for the paper at the location instead of dismissing the explanation.

- III.A.1: "time in to discrete" --> "time into discrete"

(*) We have changed this as pointed out by the reviewer.

- III.A.1: $m_c \in \{0,1\}$ --> shouldn't the variable 'u' be in here somewhere? You mention that it corresponds to user 'u' but 'u' is not in this expression (unless I misunderstood here)

(*) The reviewer is correct that the variable m corresponds also to users being covered at different base stations. However, since there is a direct mapping between content chunks and the users that generated them we prefer to express the problem in terms of a cover over content and for the sake of simplicity we dropped the users from the formulation (again consider that there is a mapping from content to users). Alternatively, another formulation of the problem could consider covering users at specific time instances instead of pieces of content. The inputs to our algorithm have also been constructed in such a way.

- III.B: In para "We further explain". Algorithm 18 isn't in the paper, yet there is a ref here

(*) This is just a typo. It should have referred to algorithm 1 which is in the paper. We have changed this.

- III.B: Para starting "We further explain" was really good--nice example

(*) We agree.

- III.C: Max data seen was 3.5MB... seems low?

(*) This was actually in accordance to the data we have collected from Flickr where the maximum size of a picture was the order of MB. Note also that most of the uploads in our dataset are photos and in a 3G network, uploading a 3.5MB piece of content can take on the order of minutes.

- IV.D: "Here, we take a user view and analyze how users" -> "Here, we analyze how users"

(*) We have modified this.

- Fig 11(d): the different series (50%, 60%, ...90%, 100%) determines the number of drop zones, and so in a sense this graph really shows two effects. Not necessarily a huge deal.

(*) The reviewer is correct that the more content we intend to absorb the more Drop Zones are required. With the different series/curves we have tried to capture also by what amount the number of zones encountered by a user changes with a significantly bigger infrastructure.

- IV.D: "Given that the more delivered content necessarily" -> awkward

(*) We agree with the reviewer and we have changed this to the following:
 "A larger amount of content delivered implies a larger number of Drop Zones."

- VI: Maybe a UI to ask them if they're deviating from their normal routine?

(*) We agree with the reviewer that this would be an interesting enhancement and we have included the following note:
 "Several other enhancements could be included such as a user interface that asks the user if he intends to change his normal routine."

Endnotes:

- [23]: 3g should be "3G"
- [29]: needs a year and publisher
- [47]: needs a year, volume, issue, etc
- [49]: "connecting" -> "Connecting"; "3g" -> "3G"
- generally, you should use full references, not shortened "in IMC 09" type references

(*) We have fixed this throughout the references section.

Appendix B, Section B, equation (3):

- what is x_b ? What is N_c ? Maybe give some intuition or explain the terms a bit?

(*) Both were already explained in the text already and we have now highlighted the corresponding regions.

Appendix:

"The variables x_b from $\{0,1\}$ describe whether a Drop Zone is placed at base-station b (i.e., $x_b=1$) or not (i.e., $x_b=0$)."

Section III.A.1

"Furthermore, let n_c^i from $\{0,1\}$ indicate whether content chunk c from C was generated at time t_i (i.e., $n_c^i=1$) or not (i.e., $n_c^i=0$)."

Reviewer: 2
Comments to the Author

The paper is good in many ways, but altogether not of the quality and coherence that a ToN paper should have.

(*) We agree with the reviewer that the paper has good parts and we plan to further improve it by addressing his comments.

Primary issues I had with the paper are:

- The motivational portions of the paper do not really discuss whether the problem being examined is a feature of the current 2.5/3G infrastructure, and may be relieved by ongoing rollout of 4G becoming more widespread, used of WiFi as supplementary, or even future technologies. Unless content grows similarly to the network capacity, this should help the issue making drop zones a crutch for legacy portions of the infrastructure only. However, this was not adequately (or at all) discussed in the paper, making it difficult or impossible to gauge the potential impact of the contribution.

(*) The reviewer is correct in the sense that there is a finer point at work. Indeed, both data consumption and network capacity are increasing. Data consumption increases significantly at finer time scales while network capacity increases significantly once every several years. Even though currently providers see their networks overwhelmed by traffic and we are currently at a point where providers are discouraging heavy data usage by moving from unlimited data plans is this a general trend on the long run?

It is however generally accepted that growth in mobile data demand outpaces the growth in capacity provided through upgrades. One only needs to take a look at the Cisco Global Mobile Data Traffic Forecast and can see that the CAGR for Data Consumption is at about 92% depending on device while the increase in connection speeds has a CAGR of around 60% (from 2010 to 2015).

http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-520862.html

We have added the following paragraph at the end of Section VI.

"Also, it is generally accepted that growth in mobile data demand outpaces the growth in capacity provided through upgrades. One only needs to take a look at the Cisco Global Mobile Data Traffic Forecast [6] and can see that the CAGR for data consumption is at about 92% depending on device while the increase in connection speeds has a CAGR of around 60% (from 2010 to 2015). So, on the long run, providers will most likely have to upgrade their networks again and again making our Drop Zone approach a viable first step to conduct upgrades."

- The paper (and approach) do not discuss the user's expectation of having content available immediately after posting it, even though it does consider in good detail the fact that users frequently delay posting of content. This is a major adjustment to user expectations that would need to be made in order for this to be successful.

(*) As can be seen in Figure 4, the reviewer is correct. There are users that indeed want to have the content available straight away. However what we pointed out is that there is a large number of users that already delay content uploads. While we cannot predict how user expectations change—it might be that users would find it easy to spend 3 seconds in a 4G cell and upload a photo right away, compared to a few minutes in a 2.5G cell. Or they would still prefer to delay the uploads since 4G usage is a lot heavier on the battery. This research is not targeted at this (as we have acknowledged) and pricing based incentives (not our focus) for users delaying content complement our paper.

In Section II.B.3 paragraph 3 we have added the following sentences.

"We do not expect all users to postpone content uploads. Indeed, some users have the expectation of having content available immediately after posting it. It is however hard to predict which way user behavior will change as there are factors working both ways (battery life, capped data plans, smaller upload delays, incentives that providers might offer)."

- The Section 4.F "What-If Scenarios" were good, but incomplete, missing several of the most relevant cases, and the claim of 14 years lifetime is dubious given these other what-if conditions, some of which are near certainties. These include: large increases in the number and percentages of users with smart devices, an increase in the number of mobile smart devices per-user, increases in the desire for real-time high-fidelity audio/video communications, etc. This claim was repeated in the conclusions without qualification.

(*) We agree with the reviewer that the mobile world contains high variability in terms of devices, applications, content sizes, etc. As such we can only state that our claim was made in the context of the assumptions given in the paper and cannot be taken out of context.

We have further included the following phrases at the beginning of Section IV.F.

"One needs to understand that there are multiple (often unpredictable) ways in which the mobile world can evolve. In particular, some trends might involve large increases in the number and percentages of users with smart devices, an increase in the number of mobile smart devices per-user, increases in the desire for real-time high-fidelity audio/video communications, etc."

"We agree that studied in isolation some of the above scenarios can be seen as limited, however we consider them a useful exercise for predicting future growth to a limited extent."

- Section 5 is absolutely critical to understanding the proposed technology, and yet happens to be the worst part of this paper, dragging down several of the other sections which were much better. This section needs to explain a technical overview of how the system would be integrated with existing infrastructure and applications and actually operate, yet it seems to largely be hand-waving. For instance, it mentions Mobile IP in an offhand way, without motivating or explaining why or what MIP would be used for. The configuration with FAs implies MIPv4, which is of course, not compatible with IPv6 and inconsistent with the portions of 3GPP2 specifications that are built on IPv6. Further, the statement that a 2nd interface on the drop zone equipment is used to sniff all traffic going in and out of the base station seems to be completely unmotivated in the paper, and also completely unnecessary, since the description later says that clients would detect whether base stations are drop zone enabled (but forgets to mention how this is technically done!). This section needs major work in order to show that the solution is technically capable of being implemented and integrated in the way proposed. The impact on applications most of all needs to be very well explained (whether only new or legacy applications would work, etc.). This is the primary reason why I recommend for rejection rather than revision/resubmission, because it seems that this requires fundamental technical investigation and rework by the authors beyond that involved in minor revisions.

(*) The reviewer makes a few points in that we further summarize:

- use of Mobile IP.
- reason for the 2nd sniffing interface and how the Drop Zone is technically detected.
- impact on applications

First of all what we failed to point out is that this was a case

study for an existing network and describes how a Drop Zone architecture would be **currently** integrated. The network currently uses MIP and is configured with a FA and HA. We have introduced the following paragraph to address this (Section V. 2nd paragraph).

"Note that this design was obtained from an existing (IPv4 based) cellular network in which we have integrated our Drop Zone architecture. The usage of a Foreign Agent in IPv4 implies MIPv4 that is currently not compatible with IPv6 and inconsistent with the portions of 3GPP2 specifications that are built on Ipv6. Switching to IPv6 would imply changing the some Drop Zone functionality also."

The reason for the 2nd sniffing interface is that we do not wish to change the way the base station functions. It still receives and handles traffic for the mobile phone yet the traffic is stopped from going over the backhaul link (which in newer wireless technologies is a potential bottleneck, see referenced Cisco report).

Detecting if a Drop Zone is present can be done as simple as sending a packet and obtaining a reply from a special IP address such as 0.0.0.0 (the Drop Zone equipment would provide the reply).

We have added the following sentences in the paper in Section V, paragraphs 5 and 7 respectively:

"The reason for the sniffing interface is that this way, we do not need to change the way the base station functions. It still receives and handles traffic for the mobile phone yet the traffic is stopped from going over the backhaul link."

"detecting if a Drop Zone is present can be done as simple as sending a packet and obtaining a reply from a special IP address such as 0.0.0.0 with the Drop Zone equipment providing the reply"

With regards to applications, we have already explained what the applications are. Applications such as Skype or Facetime would remain unaffected. Only applications that upload or download bulky data (videos, photos, noncritical patches) that can tolerate delays will be changed. Even now the photo camera of some phones is integrated with Facebook. Users are directly being asked upon snap if they want to share the photo via Facebook or Flickr. They will also be asked if they can tolerate upload delays in which case the picture upload is handled by a daemon that uploads it only when a Drop Zone is available. This is especially important since phone and camera resolutions are expected only to increase bringing about an increase in traffic. Newer services can also be integrated with the phone camera and given the same options for upload.

We have added the following paragraph in Section V, paragraph 6.

“Even now the photo camera of particular phones is integrated with Facebook. Users are directly being asked upon snap if they want to share the photo via Facebook or Flickr and the photo is uploaded when network connectivity is available. In our scenario they will also be asked if they can tolerate upload delays in which case the picture upload is handled by a daemon that uploads it only when a Drop Zone is available. Newer services can also be integrated with the phone camera and given the same options for upload.”

Some smaller nits are below:

- Figure 5's caption may be misleading, as this doesn't show statistical correlation between movement and uploading, but really just shows the user distributions of upload volume and movement independently on different axes

(*) We agree with the reviewer that the caption is somewhat misleading. We have enhanced it with the following explanation:

“Each of the curves corresponds to a different y axes therefore the figure does not show statistical correlation between movement and uploading.”

- In Section 4, "ILP" isn't defined, though I think it should be "Integer Linear Programming"?

(*) We have adapted Integer Linear Program throughout the paper.

Reviewer: 3 Comments to the Author

The authors present an approach, drop zones, to allow wireless providers to focus their infrastructure updates to keep costs low while improving the user experience. These zones are chosen using a greedy algorithm, which shows that if the user is willing to accept some small degrees of delay, the network could easily accommodate projected increases in demand without the full cost of system-wide upgrades.

(*) Although we did not directly point this out in the paper, as it is a secondary result, the reviewer is correct that the user experience could be improved. By certain users not uploading their content in poorer connectivity regions, the user experience of the users that choose to upload their content there is

improved.

I am pretty upbeat about the idea and paper. It's perhaps a little obvious as a solution, but it makes sense and yields good results. The analysis to back it up seems pretty good.

(*) We agree with the reviewer, and thank him for the thorough review. The paper does not only introduce the idea but it is also the first paper that examines the tradeoff between content delay and infrastructure reduction (that can be obtained by analyzing natural user movement patterns and targeting the infrastructure to specific areas).

The authors failed to even acknowledge the existence of femtocells, which surprised me in an otherwise well-considered paper. The idea of higher capacity channels has been explored out-of-band. This is particularly useful for home users or businesses with poor reception. Since these are considered the "comfort zones" for users, what affect would this have on drop zone coverage? I imagine the effect would be limited, since these cells would only service a few users, but real analysis would be worthwhile. I would condition acceptance on at least acknowledging the existence of femtocells, but I believe the authors could really strengthen the work with greater consideration of/comparison with the approach.

(*) We were aware of the existence of femtocells and the omission was not on purpose. We do consider them an important of the mobile landscape. Our paper however considers system-wide network upgrades and as such, femtocells were not part of the formulation. However, we agree with the reviewer that the existence of femtocells needs to be acknowledged and we have proceeded to incorporate the following paragraph in the paper in Section II.A paragraph 3:

"Note that businesses and sometimes even users decide to take matters into their own hands and employ femtocells (small cellular base stations) in order to get bigger capacity and increased reliability for themselves. They are however complementary to our Drop Zone upgrades as they are targeted at covering a small number of users; unlike the Drop Zones that we envision that will cover significantly larger areas and numbers of users. We also do not have femtocell statistics for the provider that we have analyzed and the results we provide should be considered with this in mind."

What is going on in section 3A1? The authors' notation is challenging to understand based on how it is presented. While it may be the most succinct and precise way of describing the system, the authors should consider whether there are better ways to

ensure reader comprehension.

(*) We have made an effort to improve the presentation with further explanations.

There are other little things the authors can do to improve the work:

1) Take a look at the usage of \cite. Some instances have a space between the word before and the citation and others do not. The former approach looks better.

(*) We have fixed this.

2) The abstract's 1st paragraph ends with "address this problem" and talks about "solutions" but the authors never defined the problem. A sentence like "this increase in demand will overwhelm capacity and limits the providers' ability to provide the quality of service demanded by their users" would indicate what the problem is.

(*) We agree with the reviewer and we have inserted the exact sentence suggested at the specific point in the abstract.

"This increase in demand will overwhelm capacity and limits the providers' ability to provide the quality of service demanded by their users."

3) In the 2nd to last paragraph in Section 1, the authors say "nationwide infrastructure." What nation? The authors are clearly assuming the United States throughout, but this should be stated somewhere.

(*) This was stated in the abstract (now pointed out). We have also introduced it in the place requested by the reviewer.

4) In Section 2A, I kept thinking about all the smartphones that also use wifi and prefer it when such a signal is present. Could we at least acknowledge wifi usage? Would this show up in your data? Is that a weakness?

(*) Since our trace is collected from the content billing part of the cellular network, WiFi usage does not show up in our trace. We did not mean to hide this fact since it is rather obvious that capturing both cellular and WiFi usage would imply a user agent installed on phones that would be hard to deploy at such a scale. We have added a paragraph that acknowledges WiFi usage in section II.B (now paragraph 4) where we describe the trace.

"Note that in our trace there exist smartphones that also use WiFi and prefer it when such a signal is present. Since our trace was collected from the content billing part of the cellular network, WiFi transfers are not recorded. Even though it would be desirable to analyze WiFi effects (transfers meant for Drop Zones could be scheduled when WiFi is available), WiFi usage does not directly impact our study."
