

Paper ID (the number in the name of the file): 13

Paper title: Measurement and Diagnosis of Address Misconfigured P2P Traffic

Authors (or the first author): Zhichun Li, Anup Goyal, Yan Chen and Aleksandar Kuzmanovic

Please thoroughly analyze the paper and provide **detailed** comments on the following:

*1: Is the paper relevant to the topic of the special issue?

No

Yes

More info if any:

The paper deals with Address Misconfigured P2P traffic. The topic is very interesting since this type of traffic is responsible for a considerable bandwidth consumption and consequent waste of resources. As the authors state, this traffic “consumes 7.9 Gb/s” and it is “mostly intercontinental”, which raises even more the costs associated to this kind of traffic. Therefore, it is a very relevant topic.

*2: How innovative is the paper?

5 (Very innovative)

4 (Innovative)

3 (Marginally)

2 (Not very much)

1 (Not)

0 (Not at all)

More info:

The authors present an interesting, and novel, analysis of the causes behind P2P misconfigured traffic. The authors also present a tool named *P2PScope* which is used to track and find the mentioned causes.

*3: How would you rate the technical quality of the paper?

5 (Very high)

4 (High)

3 (Good)

2 (Needs improvement)

1 (Low)

0 (Very low)

More info:

The technical quality of the paper is good. The authors present a deep knowledge of P2P systems and protocols and the data used for analysis is also adequate since it is composed of traffic collected by HoneyNet/Honeyfarm sensors.

*4: How is the presentation?

5 (Excellent)

4 (Good)

3 (Above average)

2 (Below average)

1 (Fair)

0 (Poor)

More info if any:

*5: How is the quality of the reported results?

- 5 (Excellent)
- 4 (Good)
- 3 (Above average)
- 2 (Below average)
- 1 (Fair)
- 0 (Poor)

More info:

The results presented are quite interesting. However, they are quite mixed with the text which makes it a bit confusing and difficult for the reader to find and understand the presented results. I suggest more tables and graphs. I also suggest that the authors make a better separation between the text and the achieved results.

*6: How is the appropriateness of references and of related work?

- 5 (Excellent)
- 4 (Good)
- 3 (Above average)
- 2 (Below average)
- 1 (Fair)
- 0 (Poor)

More info if any:

The section "Related Work" could present more similar studies and it is very short section

*7: Is the paper of interest to readers of IEEE Network?

- 3 (Yes)
- 2 (May be)
- 1 (No)
- 0 (Not applicable)

More info if any:

*8: Detailed comments to the authors (how the authors can improve the paper?)

The paper is very interesting as it deals with Address Misconfigured P2P traffic. The topic is very interesting since this type of traffic is responsible for a considerable bandwidth consumption and consequent waste of resources. However there are some technical aspects I would like to discuss that can, in my opinion, improve the paper:

- In section II, it is mentioned "We have data from 2004 through 2007": don't you need more recent data? P2P traffic has evolved so much in the last years with so many protocols disappearing and other appearing. It'd be interesting to use more recent data. It is also said "26 days of traffic from GQ honeyfarm": are 26 days of capture enough to infer a conclusion about misconfigured P2P traffic? Don't you need more traffic and more data?

- In section II-F, it is mentioned “We then use the average percentage from all the samples as a more representative result for the whole Internet”: can you explain this better? This sentence seems a bit bold.
- In section III-C, the authors say “Most of the events are caused by BitTorrent and eMule”: how do you reach this conclusion in the section where the design of the system is presented? Does this affect the design of your system? This sentence should be in the section Results.
- In section IV B 4) it is said “This is challenging because there are certain cases in which we cannot determine if they have...”. Why? It’d be interesting to determine the actual percentage of peers with this problem.
- In section IV C 1), the authors say “declare they have any file the remote peer wants...”: how did you find this? Can you explain it better?
- Section V Related Work: shouldn’t this be section II? It seems more interesting and structured to have the related work right after the introduction.

There are some typos:

- Abstract: “and tracking millions of peers, We find” -> “tracking millions of peers we found”
- Introduction: “In addition, from the ISPs” -> there shouldn’t be a change of line here
- Section IV C 1) “There are no events in between. anti-P2P events surprise us” -> “There are no events in between. Anti-P2P events surprise us”
- Conclusions: “Since such traffic is harmful for both users and ISPs” -> the sentence seems incomplete

* What is your confidence in your review of this paper?

- 2 (High)
 1 (Medium)
 0 (Low)

* Overall recommendation

- 3 (Accept: high quality, without changes)
 2 (Conditional Accept: minor revision)
 1 (Weak Reject: re-submit)
 0 (Reject: hopeless)

Justification of your recommendation (a few words, e.g., strength and weakness):

The paper is very interesting. Some aspects, mentioned above, should be addressed and some text re-written. But in the overall the paper is very good and the topic is interesting. It fits in this magazine issue of IEEE. Strong points are the analysis of the misconfigured P2P traffic, the results they obtained and the discovery of the main causes behind this phenomenon. Weak points are the data used which could be more recent (traces are from 2007) and the fact that more traffic data should be used when attempting to determine causes for Internet background traffic radiation.

Paper ID (the number in the name of the file): 13

Paper title: Measurement and Diagnosis of Address Misconfigured P2P Traffic

Authors (or the first author): Zhichun Li, Anup Goyal, Yan Chen, and Aleksandar Kuzmanovic

Please thoroughly analyze the paper and provide **detailed** comments on the following:

*1: Is the paper relevant to the topic of the special issue?

- No
- Yes

More info if any:

The paper is focused on the identification and measurement of address misconfigured peer-to-peer (P2P) traffic. The authors state that this kind of traffic represents a considerable share of the traffic in computer networks, which impacts negatively the available bandwidth.

Therefore, it is opinion of this reviewer that this paper is relevant for a special issue on Network Traffic Monitoring and Analysis.

*2: How innovative is the paper?

- 5 (Very innovative)
- 4 (Innovative)
- 3 (Marginally)
- 2 (Not very much)
- 1 (Not)
- 0 (Not at all)

More info:

The address misconfigured traffic is a well known problem addressed by several previous studies. Though, this paper focus on the part of this traffic that has origin in P2P systems and provide a few root causes for this kind of traffic.

*3: How would you rate the technical quality of the paper?

- 5 (Very high)
- 4 (High)
- 3 (Good)
- 2 (Needs improvement)
- 1 (Low)
- 0 (Very low)

More info:

Generally, the technical quality of the paper is good. The authors contextualized the problem addressed in the paper and used references to support some of the assertions made and the approaches followed.

However, there are several inconsistencies and a few conclusions may require stronger evidences.

For instance, in the second paragraph of section I, a few conclusions of the study that the paper describes are used as introduction to the topic, almost as a motivation. Although this is mostly an organization problem, it is not also very sound, from the technical perspective, to use the conclusions of the paper to justify the importance of the problem it addresses. Maybe the authors could use a good reference that has results concerning the measurement of this kind of *non-requested* traffic. If such reference does not exist (or even if it exists), then the results of this study should be presented as a contribution and conclusion, instead of describing them in the beginning of the introduction, almost as a way to show how serious the problem is and to justify the importance of solving it or measuring it.

Besides that, there seems to be too many assumptions or hypotheses.

The most important of them is the assumption that there are only two reasons that can be responsible for the P2P misconfiguration traffic (section III.A): software bugs and misconfiguration injection (which the authors automatically relate with anti-P2P companies).

Is there any evidence of this?

Is there any reference that can support this assumption?

Maybe there is an explanation that can be included in the paper.

Please consider an hypothetical case where a host is running a torrent based application that shares peers information through a peer exchange protocol, e.g., PEX.

If the application is stopped, will the host continue to receive requests?

If so, could this be another reason for the misconfiguration traffic?

Was this hypothesis considered?

I once inspected the traffic of my personal computer and I noticed that, after closing the torrent client application (*Transmission*), I kept receiving requests, for several days, in the same port that was defined in *Transmission* for the incoming connection. The requests were always rejected by the OS as there was no process listen on that port. Several days later, I opened *Transmission* again and chose another port for the incoming connections. After a few hours, I closed *Transmission* and noticed that I kept receiving requests again, but this time on the new port defined in *Transmission* for the incoming connections.

Do the authors think that this may be another reason for the misconfiguration traffic?

If so, isn't it possible that there are also other reasons that should be consider besides software bugs and traffic injection in order to have a stronger and sounder study?

What are malicious payloads?

Payloads whose data matches payload signatures of well known threats/worms/virus?

Is so, are those signatures included in *l7-filter*?

*4: How is the presentation?

5 (Excellent)

4 (Good)

3 (Above average)

2 (Below average)

1 (Fair)

0 (Poor)

More info if any:

The paper requires further improvement in terms of presentation and organization. It is not clear and easy to read. It lacks a consistent thread, or at least, a clear line of reasoning that could make it easy to read and to follow the ideas.

The organization is, indeed, one the weakest points of the paper.

Moreover, the organization of the paper does not privileges the need to provide tutorial content to make it understandable for a general audience.

There are also a few misconstrued sentences or expressions.

More details on this topic are provided in question 8.

*5: How is the quality of the reported results?

5 (Excellent)

- 4 (Good)
- 3 (Above average)
- 2 (Below average)
- 1 (Fair)
- 0 (Poor)

More info:

The results presented in the paper are interesting. However, there are a few important inconsistencies and technical inaccuracies that weaken the study. Some of them were described in question 3.

One of the concerns of this reviewer, regarding the results, is related with the generalization of the conclusions obtained. The authors used three different botnets in their study and the results and conclusions are presented as a generalization from the three botnets to the Internet.

However, there is not any consideration in the paper regarding how representative of the Internet those three botnets are.

How many protocols (P2P and non P2P) are simulated by each of them?

How many distinct attacks they can simulate?

Do they attract connections of each protocol in the same percentage as they are present in Internet?

Probably it is not easy to answer these questions. And, because of that, one should be carefully when extrapolating these results to the whole Internet.

A good example of this problem is table II.

Why are the results, in terms of percentage, so different from one dataset to the other?

If there are such a big difference between the results of the two datasets, how can one be sure that these botnets are representative of the Internet?

Without being sure of that, it is not possible to make a sound extrapolation of the results obtained.

In such scenario, one may have legitimacy to think that maybe these results depend on the botnets used in the study.

Also, from what it was possible to understand, although the authors evaluated the increase of the address misconfiguration traffic from 2004 to 2007, only the data from one of botnets spans the four years. In spite of this, instead of presenting the evolution of the P2P traffic in terms of connections, it may be more interesting to present its percentage in the whole address misconfiguration traffic. As it is, the reader cannot know if the P2P share has grown or if the all traffic increased. The same applies to the number of events with more than 100 sources that is shown in table II.

Another curious fact concerns the GQ Honeyfarm dataset. The dataset is referred and described in the text. However, no results regarding this botnet are provided, and the only reference to this dataset in the text, after being described (and besides table I), appears in the end of section II.E just to say that the address misconfiguration is prevalent across the Ipv4 space.

Why was this dataset used?

Did it contribute to the results obtained?

Were the results for this dataset consistent with the ones obtained for the remaining datasets?

Why didn't table II include any results for the GQ Honeyfarm dataset?

Maybe there are good explanations for these questions, but, in that case, those explanations should be included in the paper.

*6: How is the appropriateness of references and of related work?

- 5 (Excellent)
- 4 (Good)

- 3 (Above average)
- 2 (Below average)
- 1 (Fair)
- 0 (Poor)

More info if any:

The description of the related work and the references included are appropriated. Though, given the tutorial nature mentioned in the call for papers of this special issue, it may be more interesting to include references that can guide the readers to the concepts and approaches followed by this paper, instead of some of the references that support technical decisions or scientific conclusions.

A small reduction of the number of references may also improve the consistency of the paper.

*7: Is the paper of interest to readers of IEEE Network?

- 3 (Yes)
- 2 (May be)
- 1 (No)
- 0 (Not applicable)

More info if any:

The topic addressed by the paper, the work described, and the conclusions reached are clearly suitable for this special issue and it is opinion of this reviewer that the paper may be of interest to readers of IEEE Network.

However, as mentioned before, there is little tutorial content which can make the paper unappealing to readers outside the field of expertise of the article.

*8: Detailed comments to the authors (how the authors can improve the paper?)

As said before, the paper addresses an interesting problem and it presents good results. Because of such reasons, it is opinion of this reviewer that the paper has value to be published. However, it requires major work and improvement before being published.

One of the problems that makes the paper difficult to understand is the organization and the absence of a clear line of reasoning.

The scientific articles are not always read as a literally text, from the beginning to the end. It is important that one can have a summarily idea of the the paper just by reading the abstract, or the introduction, or even the conclusion. Likewise, the structure of the paper should facilitate the understanding of the paper and make it easy to find any specific information that the reader may want.

When reading this paper it is sometimes difficult to understand if the goal was to present a new tool/application/framework, to present the results of the study about the misconfiguration traffic, or both. These goals should be clearly exposed and explained, as well as the different steps of the approached followed and their conclusions.

The Introduction section would be more effective if it was better organized. Ideally, it should start by contextualizing the study and then it should describe clearly the problem that is being addressed (if necessary, using strong references). After that, it should explain how the authors propose to solve the problem, the procedures they developed, the experimentation/simulation they run, and finally a brief summary of the results obtained. The same structure could be used for the abstract but, obviously, in a more summarily and compressed version.

The Conclusions section is extremely short. It is almost impossible to figure it out what is the subject of the paper just by reading the Conclusions. It would be better if it could provide a brief summary of problem addressed and procedures made, an explanation of the results obtained, and a description of the conclusions of the study.

The tutorial content that could make the paper understandable for a general audience, as requested in the call for papers, is also very little.

For instance, the concepts of *botnet* or *honeypot* could be better explained in the paper, emphasizing the details necessary to understand the approach proposed.

The introduction of the acronyms is not consistent. In most cases, the long form of the acronym is never introduced. For instance: ISP, DoS, DDoS, AS, BGP, etc.

The long form of the acronyms should always be introduced in their first appearance. Ideally (this is to the consideration of the authors), they should also be reintroduced after the abstract.

The paper also requires a careful revision as there are few typos or misconstrued sentences and expressions.

In a few places, a lower case is used after a period, e.g.:

p. 3, "... a challenging problem. we have..."

p. 4, "if a peer is not running..."

p. 4, "if a peer is unroutable..."

p. 6, "traffic pattern. its traffic"

A period is missing in:

p. 1, "sources of Internet traffic Given the"

p. 2, "we use total number" ==> "we use the total number"

Several commas are missing through the text.

In section III.C, in the enumeration of root causes, after the colon it is used upper case and period at the end of each enumeration instead of semicolon or a comma.

It should be better:

"... root causes: (I) track the information flow within the suspicious P2P software; (ii) trace the Honeynet IP addresses propagated in the P2P systems; (iii) check the routability of the peers returned; (iv) ..."

Sometimes it is used the expression *UTorrent*, and other times it is used *uTorrent*.

The following sentence seems uncompleted:

"Since such traffic is harmful for both end users and ISPs."

Your confidential comments to the Guest Editors:

* What is your confidence in your review of this paper?

- 2 (High)
 1 (Medium)
 0 (Low)

* Overall recommendation

- 3 (Accept: high quality, without changes)
 2 (Conditional Accept: minor revision)
 1 (Weak Reject: re-submit)
 0 (Reject: hopeless)

Justification of your recommendation (a few words, e.g., strength and weakness):

It is my opinion that the paper follows an interesting approach and presents good results. The problem addressed is of relative importance and fits perfectly on this special issue.

The authors used real data in their analysis and justified their conclusions.

I consider that the paper has value to be published. However, it requires major modifications and a few of them are related with critical issues. For instance, in question 3, which concerns the technical quality of the paper, because of the inconsistencies described there, I was undecided about answering with 3 *Good* or 4 *Needs improvement*. The truth is that, although the idea and the study described in the paper are interesting and valid, the whole description made herein seems to lack some rigor and technical accuracy.

There are several issues that threaten the soundness of the study and the consistency of the results. The main problems are related with the organization of the paper, some technical inaccuracies, and a few weak assumptions in which some results are based.

These issues are better explained in questions 3 and 5.

Besides that, it contains little tutorial content and in some parts it almost looks like a report of the study.

Paper ID (the number in the name of the file): 13

Paper title: Measurement and Diagnosis of Address Misconfigured P2P Traffic

Authors (or the first author): Zhichun Li

Please thoroughly analyze the paper and provide **detailed** comments on the following:

*1: Is the paper relevant to the topic of the special issue?

- No
 Yes

More info if any:

*2: How innovative is the paper?

- 5 (Very innovative)
 4 (Innovative)
 3 (Marginally)
 2 (Not very much)
 1 (Not)
 0 (Not at all)

More info:

*3: How would you rate the technical quality of the paper?

- 5 (Very high)
- 4 (High)
- 3 (Good)
- 2 (Needs improvement)
- 1 (Low)
- 0 (Very low)

More info:

*4: How is the presentation?

- 5 (Excellent)
- 4 (Good)
- 3 (Above average)
- 2 (Below average)
- 1 (Fair)
- 0 (Poor)

More info if any:

*5: How is the quality of the reported results?

- 5 (Excellent)
- 4 (Good)
- 3 (Above average)
- 2 (Below average)
- 1 (Fair)
- 0 (Poor)

More info:

*6: How is the appropriateness of references and of related work?

- 5 (Excellent)
- 4 (Good)
- 3 (Above average)
- 2 (Below average)
- 1 (Fair)
- 0 (Poor)

More info if any:

*7: Is the paper of interest to readers of IEEE Network?

- 3 (Yes)
- 2 (May be)
- 1 (No)
- 0 (Not applicable)

More info if any:

*8: Detailed comments to the authors (how the authors can improve the paper?)

This paper tackles an interesting problem, address misconfigured P2P traffic, and presents how to detect and diagnose such unwanted traffic. However, regarding address misconfigured P2P traffic, the paper didn't provide any possible ways to fix the problem.

The reviewer could not understand the originality and significance of the authors' work summarized in the current paper. According to result of this paper, address misconfigured P2P traffic is mainly caused by coding bug, as well as the contributions of anti-P2P companies. Such measurement results seem to have limited significance for other researchers. What can we do next? If the authors could analyze the characteristics of address misconfigured P2P traffic and give some suggestions to solve the problem, it may be more interesting.

In Section I, the paper pointed out "In addition, from 2004 to 2007, the address misconfigured P2P traffic has increased more than 100% each year as shown in Figure 1". How to get this conclusion? Is there any reference or collected data to support this result?

In Section IV, the result analysis section, the authors should use more graphs and tables to describe the results. It would be more intuitive.

There are some writing errors here and there. For example, in paragraph 1 of Section VI, "Since such traffic is harmful for both end users and ISPs", "since" should be omitted. The authors must consistently use the technical words throughout the paper, such as the upper case "S" of "P2PScope" and "P2Pscope". The reviewer considers that the paper requires being sophisticated so that the readability of the paper is enhanced.

Your confidential comments to the Guest Editors:

* What is your confidence in your review of this paper?

2 (High)

1 (Medium)

0 (Low)

* Overall recommendation

3 (Accept: high quality, without changes)

2 (Conditional Accept: minor revision)

1 (Weak Reject: re-submit)

0 (Reject: hopeless)

Justification of your recommendation (a few words, e.g., strength and weakness):

The significance of the work is limited.

Paper ID (the number in the name of the file): 13

Paper title: Measurement and Diagnosis of Address Misconfigured P2P Traffic

Authors (or the first author): Zhichun Li, Anup Goyal, Yan Chen, Aleksandar Kuzmanovic

Please thoroughly analyze the paper and provide **detailed** comments on the following:

*1: Is the paper relevant to the topic of the special issue?

- No
- Yes

More info if any:

I'm not sure this paper really fits the CFP.

*2: How innovative is the paper?

- 5 (Very innovative)
- 4 (Innovative)
- 3 (Marginally)
- 2 (Not very much)
- 1 (Not)
- 0 (Not at all)

More info:

The paper is basically showing that byte ordering bugs in P2P clients can cause problems. A previous version of this paper appeared in INFOCOM 2010.

It is upsetting that-

- Authors DO NOT MENTION THIS
- This version is basically a shorter version of that paper (from 9 down to 6 pages)

*3: How would you rate the technical quality of the paper?

- 5 (Very high)
- 4 (High)
- 3 (Good)
- 2 (Needs improvement)
- 1 (Low)
- 0 (Very low)

More info:

The paper is technically sound even if simple methodologies are presented.

*4: How is the presentation?

- 5 (Excellent)
- 4 (Good)
- 3 (Above average)
- 2 (Below average)
- 1 (Fair)
- 0 (Poor)

More info if any:

The paper is badly written. It clearly suffer from being a shortened version of the infocom paper. The reader must already be an expert in P2P systems. Most of the terms are not defined. There are long sentences that give lot of details, but provide little intuition and insights.

*5: How is the quality of the reported results?

- 5 (Excellent)
- 4 (Good)
- 3 (Above average)
- 2 (Below average)
- 1 (Fair)
- 0 (Poor)

More info:

Bottom line: a software bug can generate some problems.

The takeover message of the paper is weak. Ok, the authors pinpointed some problems. But how general can be the approach? It sounds a lot like a custom method to identify the root cause.

The proposed methodology is not very innovative. It leverages on protocol knowledge and sort of brute force analysis to find the root cause of problems. The same methodology will fail in case i) proprietary ii) encrypted, iii) unknown protocols are used by other applications.

*6: How is the appropriateness of references and of related work?

- 5 (Excellent)
- 4 (Good)
- 3 (Above average)
- 2 (Below average)
- 1 (Fair)
- 0 (Poor)

More info if any:

*7: Is the paper of interest to readers of IEEE Network?

- 3 (Yes)
- 2 (May be)
- 1 (No)
- 0 (Not applicable)

More info if any:

Not really. It's about a minor problem, which is overstated by authors.

*8: Detailed comments to the authors (how the authors can improve the paper?)

Please, see the attached pdf in which I embedded some detailed comments.

There are several problems with this paper.

First, you do not mention that the same paper appeared in infocom 2010. Not really fair.

This version in particular is a cut of the infocom paper, which was originally 9 pages, which clearly makes this version weaker.

Major weakness

- You fail to convince me that the problem you face is important. You claim that there is a clear waste of internet resource. Then you state that approximately 7.9Gb/s of traffic is due to this. This is NOTHING compared to internet traffic and other waste of resource. I'm

thinking about spam, viruses, worms, or even the download of content using P2P that results fake at the end.

- The takeover message is basically that i) programmer must pay attention to byte ordering issues and ii) anti-p2p companies are doing nasty thing. Not really an exciting message.
- Your P2Pscope design is very simple. It relies on the knowledge of the P2P protocol. It will however not work for encrypted/obfuscated/proprietary/unknown protocols. E.g., can you handle skype traffic?
- The paper writing and organization is poor, especially for IEEE network audience. You assume the reader is an expert in P2P problems, protocols, architecture, etc. You assume everyone knows what a honey net is and how it works. You keep giving a lot of details, but fail to give the high level picture and message.
- There are some technical aspects that are not faced and explained. For example, you claim to use DPI techniques to identify P2P traffic. But what about obfuscated/encrypted traffic. eMule and BitTorrent indeed both support this... Second, you claim that you have identified two root causes. But all those the only possible one? How many may you miss? From a technical and scientific point of view, this paper is weak.

Your confidential comments to the Guest Editors:

I found it upsetting to find that the same (actually longer and more complete) paper was published at infocom 2010.

* What is your confidence in your review of this paper?

- 2 (High)
- 1 (Medium)
- 0 (Low)

* Overall recommendation

- 3 (Accept: high quality, without changes)
- 2 (Conditional Accept: minor revision)
- 1 (Weak Reject: re-submit)
- 0 (Reject: hopeless)

Justification of your recommendation (a few words, e.g., strength and weakness):

The paper deals with a problem which is not really interesting. Authors claim this is a major problem, but it just a very marginal problem. Indeed, claiming that 7Gb/s worldwide of traffic is due to this problem makes it really weak... I mean, compared to the TB of spam, or viruses, of worms, this is really marginal.

The takeover message is that byte ordering can be a problem for some P2P implementation. Ok, not really scientific. Second message, anti-p2p companies can play bad games. Ok, no news.

For the rest, the P2Pscope design is very simple and straight forward. It can deal with all known protocols, but it will fail in case of encrypted/obfuscated /unknown protocol. Note that authors never mention this.

Paper writing is definitively below the average. The reader must be already an expert of P2P systems and he must already know all the technical details. The paper will be mostly obscure to the average reader.

Paper ID (the number in the name of the file): 13

Paper title: Measurement and Diagnosis of Address Misconfigured P2P Traffic

Authors (or the first author): Zhichun Li, Anup Goyal, Yan Chen, Aleksander Kuzmanovic

Please thoroughly analyze the paper and provide **detailed** comments on the following:

*1: Is the paper relevant to the topic of the special issue?

- No
- Yes

More info if any:

*2: How innovative is the paper?

- 5 (Very innovative)
- 4 (Innovative)
- 3 (Marginally)
- 2 (Not very much)
- 1 (Not)
- 0 (Not at all)

More info:

Marginally to Innovative, with more weight towards Innovative.

*3: How would you rate the technical quality of the paper?

- 5 (Very high)
- 4 (High)
- 3 (Good)
- 2 (Needs improvement)
- 1 (Low)
- 0 (Very low)

More info:

*4: How is the presentation?

- 5 (Excellent)
- 4 (Good)
- 3 (Above average)
- 2 (Below average)
- 1 (Fair)
- 0 (Poor)

More info if any:

*5: How is the quality of the reported results?

- 5 (Excellent)
- 4 (Good)
- 3 (Above average)
- 2 (Below average)
- 1 (Fair)
- 0 (Poor)

More info:

I would rate the quality of reported results as “Average”, but that is not an option in the list above.

*6: How is the appropriateness of references and of related work?

- 5 (Excellent)
- 4 (Good)
- 3 (Above average)
- 2 (Below average)
- 1 (Fair)
- 0 (Poor)

More info if any:

*7: Is the paper of interest to readers of IEEE Network?

- 3 (Yes)
- 2 (May be)
- 1 (No)
- 0 (Not applicable)

More info if any:

*8: Detailed comments to the authors (how the authors can improve the paper?)

I think you have good data, but the presentation is lacking. The english and grammar need to be substantively tightened up because as the paper now stands, it is hard to make sense of some of the sentence constructs (e.g., “*the anti-P2P companies themselves somehow get misconfigured solely.*” Huh?) and the text does not really flow well.

But beyond the language and grammatical problems, the paper needs to be presented in a more rigorous fashion. As I said before, you have good data but the presentation is lacking. Here are some specific comments:

- 1) Please explain what you mean by “The NU HoneyNet sensor in Northwestern University has 10 discontinuous /24 IP blocks within three different /8 IP prefixes.” Does it matter to the reader that NU has 3 different /8 IP prefixes? Is it not enough to know that the university has 10 /24 IP blocks?
- 2) As far as I can tell, Table I is not referenced anywhere in the text. The three datasets in Table I have different start and end times --- is this important to making sense of the data? Clearly, the variations are huge --- LBL spans almost 4 years while GQ spans only 25 days!
- 3) Please spend some real-estate describing HoneyNet and HoneyFarm. Apparently the former is used by LBL and NU and the latter by GQ. Does HoneyFarm use Honeyd as well? Why is Honeyd important? I suspect that I can **guess** what the answers to my questions are, but a technical paper should not be written to allow the reader to guess.
- 4) Why is GQ missing from Table II? Furthermore, I don't think GQ dataset is used anywhere in the paper. Why introduce it at all?
- 5) In Section II-B, I do not follow your definition of “events”. I read it, but I am not sure what you are trying to say. You need to spend some more text describing your “events”.
- 6) In Section II-E, why muddle the discussion with Blaster and Welchia worms? They have nothing to do with your primary thesis under study. You can simply normalize all of the connections due to p2p peers and present these in a trendline.
- 7) In Section III-C, you may want to discuss when you perform the routability tests. Clearly, if the routability test is performed many minutes --- or even hours --- after your event

occurred, the responsible IP address may be been reassigned to other subscribers who may be running firewalls to block it, etc. Also, the newly assigned subscriber may not be running the same p2p client, so you need to distinguish between unroutable addresses as well as addresses that were routable but did not elicit a response.

- 8) Section IV: Is p2pscope deployed **only** at NU? If so, then the results are representative only for the NU dataset, no? Please clarify.
- 9) Section IV-B, bullet 3: You write, “Since we have not observed any misconfigured DHT traffic, we believe that the DHT does not have the byte ordering problem.” I do trust you, but please provide some raw numbers. How many unique peers did you see in the NU dataset (since section IV appears to be for the NU dataset)? Of these, how many used eMule (again, I can see from Table II that 76.52% of peers used eMule in NU, but 76.52% of what number?) Same problem with the rest of that bullet item and bullet item 4: you report percentages but I am not sure what was your population size?
- 10) Section IV-B, bullet 3: You state that you do a return routability test, but do not state when you do it (see comment 7 above).
- 11) Section IV-C, “... the peers in normal peer events depart gradually.” Maybe you mean *randomly* instead of *gradually*?
- 12) Section IV-C: “Moreover, the peers in anti-P2P peer events are from a small number of networks.” Can you quantify these networks? Are they tier 3 ISP networks, tier 1? or tier 2?

Your confidential comments to the Guest Editors:

None. All comments are public.

* What is your confidence in your review of this paper?

- 2 (High)
 1 (Medium)
 0 (Low)

* Overall recommendation

- 3 (Accept: high quality, without changes)
 2 (Conditional Accept: minor revision)
 1 (Weak Reject: re-submit)
 0 (Reject: hopeless)

Justification of your recommendation (a few words, e.g., strength and weakness):

Given a choice, I would rate this paper as “Conditional Accept: major revision”, but that is not a choice above. I think that the paper has good data, but the presentation needs to be spruced up heavily and resubmitted for another review cycle. In its current form, I would reject the paper. For reasons for my justification, please see review comments I made above.