

# BitMiner: Bits Mining in Internet Traffic Classification

Zhenlong Yuan<sup>\*‡</sup>, Yibo Xue<sup>†‡</sup> and Mihaela van der Schaar<sup>||</sup>

<sup>\*</sup>Department of Automation, Tsinghua University, Beijing, China

<sup>||</sup>Department of Electrical Engineering, UCLA, Los Angeles, CA, USA

<sup>†</sup>Tsinghua National Lab for Information Science and Technology, Beijing, China

<sup>‡</sup>Research Institute of Information Technology, Tsinghua University, Beijing, China  
yuanz11@mails.tsinghua.edu.cn, yiboxue@tsinghua.edu.cn, mihaela@ee.ucla.edu

## ABSTRACT

Traditionally, signatures used for traffic classification are constructed at the byte-level. However, as more and more data-transfer formats of network protocols and applications are encoded at the bit-level, byte-level signatures are losing their effectiveness in traffic classification. In this poster, we creatively construct bit-level signatures by associating the bit-values with their bit-positions in each traffic flow. Furthermore, we present BitMiner, an automated traffic mining tool that can mine application signatures at the most fine-grained bit-level granularity. Our preliminary test on popular peer-to-peer (P2P) applications, e.g. *Skype*, *Google Hangouts*, *PPTV*, *eMule*, *Xunlei* and *QQDownload*, reveals that although they all have no byte-level signatures, there are significant bit-level signatures hidden in their traffic.

## CCS Concepts

•Networks Network management;

## Keywords

Traffic classification, bit-level signatures, bits mining

## 1. INTRODUCTION

Signature based traffic classification has been playing an important role in a broad range of network operations and security management, such as quality-of-service control and intrusion detection. However, due to the increasing number of network applications and their frequent updates, it is becoming more challenging to keep track of the signatures. To address this challenge, a number of existing solutions have focused on automatically extracting signatures at the byte-level [4, 5], which first divide packet payloads into groups of consecutive bytes and then analyze to get the possible signatures. However, those solutions have two major limitations. Firstly, they are unable to discover signatures at the more fine-grained bit-level granularity. Note that previous work [1, 2] have revealed that bit-level characteristics

(group of 4 bits, less than 1 byte) are of great importance in identifying a few P2P applications. Secondly, they confine signatures to groups of consecutive bytes and thus are hard to discover the signatures that consist of inconsecutive bytes (e.g. 1 byte) in packet payloads. In this poster, we propose the novel bit-level signatures, and present an automated traffic mining tool (BitMiner) that can mine signatures at the most fine-grained bit-level granularity.

## 2. BITMINER

In this poster, we have two observations. The first is that an application signature should be robust enough to support per-flow identification due to the prevalence of asymmetric routing. For this reason, a favorable application signature should be one of the most frequent patterns in captured traffic after running an application for plenty of times. Therefore, our goal can turn into mining the most frequent patterns<sup>1</sup> in the application traffic. The second is that the bit-value of a bit-position in a flow often determines the bit-values of other bit-positions in this flow. Therefore, we are motivated to associate all the bit-values with their bit-positions in a flow for frequent pattern mining.

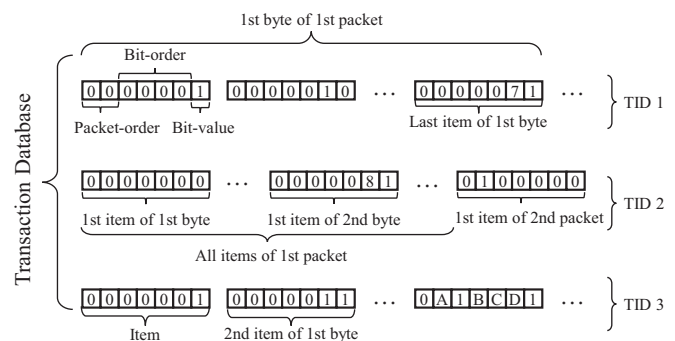


Figure 1: Format Traffic Flows to Transactions

As shown in Figure 1, we can take a bit-value with its position in a flow as an *item* and take all the bit-values with their individual positions in this flow as a *transaction*. Notice that we use only two hexadecimal characters to represent an *item*'s packet-order in a flow because application signatures are generally required to achieve early identification in practical use and the first 256 (0x00–0xFF) packets of a flow are sufficient enough. Similarly, we use four hexadecimal characters to represent one *item*'s bit-order in a

<sup>1</sup>From here, we start using some *terms* in *Data Mining*.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SIGCOMM '15 August 17–21, 2015, London, United Kingdom

© 2015 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-3542-3/15/08.

DOI: <http://dx.doi.org/10.1145/2785956.2789997>

Applications	Application Signatures	Support (Recall)
Skype	$\wedge(002\_0x02)+(002\_0\_0 \& 002\_4\_1 \& 002\_5\_1 \& 002\_6\_0 \& 002\_7\_1)*\$$	100.00%
Xunlei (Thunder)	$\wedge(001\_0\_0 \& 003\_0\_0 \& 003\_1\_1 \& 003\_2\_0)*\$$	100.00%
eMule	$\wedge(000\_0\_0 \& 000\_4\_0 \& 001\_0\_0 \& 001\_4\_0 \& 000\_6\_0 \& 001\_1\_0 \& 001\_5\_0)   (000\_0\_0 \& 000\_4\_0 \& 001\_0\_0 \& 001\_4\_0 \& 000\_6\_1 \& 001\_1\_1 \& 001\_5\_1)*\$$	100.00%
Google Hangouts	$\wedge(000\_0\_1 \& 000\_1\_0 \& 001\_1\_1 \& 001\_3\_0)*\$$	100.00%
PPTV (PPLive)	$\wedge(007\_0\_0 \& 007\_1\_0 \& 007\_2\_0 \& 007\_3\_0 \& 008\_0x00 \& 009\_0\_0 \& 009\_1\_0 \& 009\_2\_0 \& 009\_3\_0 \& 009\_6\_0 \& 00A\_0\_0 \& 00A\_1\_0 \& 00A\_2\_0 \& 00A\_3\_0)*\$$	100.00%
QQDownload	$\wedge(000\_1\_1 \& 000\_2\_1 \& 000\_5\_1 \& 000\_7\_0 \& 001\_0\_0 \& 001\_1\_0 \& 002\_0\_0 \& 002\_1\_0 \& 002\_7\_0 \& 003\_0\_0 \& 004\_0\_0 \& 005\_0\_0 \& 007\_5\_0 \& 009\_0\_0 \& 00A\_1\_0)*\$$	100.00%

Table 1: The Generated Bit-level Signatures

packet payload because the MTU of an IP packet over Ethernet networks is 1500-byte where 1 byte has 8 bit-orders.




Figure 2: An Example of How *BitMiner* Works

*BitMiner* consists of two parts: *Bit-table* and *Miner-tree*. Figure 2 shows an example of how *BitMiner* works. *Bit-table* is a hash table used for hashing and storing all the *items* read from a *transaction database*. In this process, *Bit-table* will read the *transaction database* twice. For the first time, *Bit-table* will count the *support* of every *item*. For the second time, *Bit-table* will remove the items whose *support* is below the initially set *support threshold* and sort the remaining *items* in every *transaction* by their *supports* (maximum to minimum). After that, all the sorted *transactions* will be entered into *Miner-tree* as a new *transaction database*.

*Miner-tree* is a prefix tree of the new *transactions*, which takes idea from the FP-tree [3] but is different. Note that there are probably multiple tasks running within an application and thus the signature could be a regular expression. Considering a *transaction* (flow) can only belong to one of the tasks, all the *transactions* are divided into multiple clusters to represent different tasks. Since the *items* in each *transaction* have been sorted by their *supports*, it is extremely fast to construct the *Miner-tree*.

After constructing the *Miner-tree*, there will be a pruning process controlled by two thresholds: *minimum support* and *minimum confidence*. Particularly, the *support* (defined as the proportion of *transactions* in a node from the whole *transaction database*) will be checked for every single node. Moreover, the *confidence* (defined as the proportion of *transactions* in all the child-nodes of a node from the node itself) will be checked for every parent node. In this way, it can be determined whether a branch should be removed or a parent node should stop splitting. Finally, the branches of the pruned *Miner-tree* are the target signature.

### 3. EVALUATION

*BitMiner* has been tested on the UDP traffic of six pop-

ular applications. As shown in Table 1, every signature is generated by *BitMiner* within a few seconds. The “(*p*)” represents a pattern (*p*) matching within one packet’s payload, the “ $\wedge(p)$ ” represents this matched packet is the first packet of a flow, the “(*p*)\$” represents this matched packet is the last packet of a flow, the “(*p*)<sup>+</sup>” represents this matched packet appears one or more times in succession within a flow, the “(*p*)<sup>\*</sup>” represents this matched packet appears zero or more times in succession within a flow, the “002\_0x02” represents the third byte value of a packet’s payload is 0x02, the “002\_4\_1” represents the fifth bit value of the third byte is 1, the “*p*&*p*” represents two patterns matching with one packet’s payload simultaneously and the “(*p*)/(*p*)” represents either one matched packet appears within a flow. For instance, the third byte values of the first one or more packets of a Skype flow are always 0x02 while five bit values of the third bytes of all the other packets are fixed. Specially, we also examine the other bits adjacent to the mined ones, such as the ‘second, third and fourth’ bits of the third bytes of Skype flows and the ‘fourth, fifth, sixth, seventh and eighth’ bits of the fourth bytes of Thunder flows. The results show that those bit-values are completely random (i.e. uniformly distributed). Also as shown in Table 1, the *support* represents the proportion of flows matched with the mined signature, which is equivalent to the *recall* in traffic classification. In addition, a longer signature generally means a better *precision*. For example, if we check the first 10 packets of a Thunder flow, the signature used for matching is totally 40 bits long, which may be robust enough to get a high precision in real-world situations.

### 4. ACKNOWLEDGEMENT

This work was supported by the National Key Technology R&D Program of China under Grant No.2012BAH46B04.

### 5. REFERENCES

- [1] D. Bonfiglio, M. Mellia, M. Meo, D. Rossi, and P. Tofanelli. Revealing skype traffic: when randomness plays with you. In *ACM SIGCOMM*, 2007.
- [2] A. Finamore, M. Mellia, M. Meo, and D. Rossi. Kiss: stochastic packet inspection classifier for udp traffic. *IEEE/ACM Transactions on Networking*, 2010.
- [3] J. Han, J. Pei, and Y. Yin. Mining frequent patterns without candidate generation. In *ACM SIGMOD*, 2000.
- [4] J. Ma, K. Levchenko, C. Kreibich, S. Savage, and G. M. Voelker. Unexpected means of protocol inference. In *ACM SIGCOMM IMC*, 2006.
- [5] Z. Zhang, Z. Zhang, P. P. Lee, Y. Liu, and G. Xie. Proword: an unsupervised approach to protocol feature word extraction. In *IEEE INFOCOM*, 2014.