HIERARCHICAL HEAVY HITTER DETECTION UNDER UNKNOWN MODELS

Sattar Vakili[⋆] *Qing Zhao*[†] *Chang Liu*[‡] *Chen-Nee Chuah*[‡]

[⋆] Electrical Engineering Department, Princeton University, Princeton, svakili@princeton.edu

[†]School of Electrical and Computer Engineering, Cornell University, qz16@cornell.edu

[‡]Electrical and Computer Engineering Department, University of California, Davis, {cchliu,chuah}@ucdavis.edu

ABSTRACT

We consider the problem of detecting heavy hitters and hierarchical heavy hitters among a large number of traffic flows modeled as random processes with unknown and potentially heavy-tailed distributions. The objective is an active inference strategy that determines, sequentially, which aggregated flow on the IP-prefix tree to probe in order to minimize the sample complexity under a reliability constraint. We propose an active inference strategy that induces a biased random walk on the flow aggregation tree based on confidence bounds of sample statistics. We then establish its order optimality in terms of both the size of the search space (i.e., the number of traffic flows) and the reliability requirement. The result also finds applications in noisy group testing and adaptive sampling with noisy response.

***Index Terms*—** Hierarchical Heavy Hitter, Anomaly Detection, Active Non-Parametric Composite Hypothesis Testing

1. INTRODUCTION

In Internet and other communication and financial networks, it is a common observation that a small number of flows, referred to as heavy hitters (HH), account for the most of the total traffic [1]. Quickly identifying the heavy hitters is thus crucial to network stability and security. With limited sampling resources at the router, however, maintaining a packet count of each individual flow is highly inefficient, if not infeasible. The key to an efficient solution is to consider prefix aggregation based on the source or destination IP addresses. This naturally leads to a binary tree structure (as illustrated in Fig. 1) with each node representing an aggregated flow with a specific IP prefix. The packet count of each node on the tree equals to the sum of the packet counts of its children.

A more complex version of the problem is hierarchical heavy hitter (HHH) detection, in which the search for flows with abnormal volume extends to aggregated flows in the upper levels of the IP prefix tree. In other words, there exist HHHes defined recursively in an ascending order of the level of the tree. Specifically, an upper-level node is an HHH if its mean remains above a given threshold after excluding all its abnormal descendants (if any). Otherwise, this upper-level node is only a reflecting point for merely being an ancestor of an HHH (see Fig. 1). HHH detection is of particular interest in detecting *distributed* denial-of-service attacks [2].

Each traffic flow is modeled as a stochastic process with an unknown and potentially heavy-tailed distribution. The objective of the problem is to detect all HHHes quickly and reliably by fully exploiting the hierarchical structure of the flow aggregation. Specifically, we seek an active inference strategy that determines, sequentially, which node on the tree to probe and when to terminate the search in order to minimize the sample complexity for a given level of detection reliability. We are particularly interested in strategies that achieve a sublinear scaling of the sample complexity with respect to the number of traffic flows. In other words, accurate detection can be achieved by examining only a diminishing fraction of the search space as the search space grows.

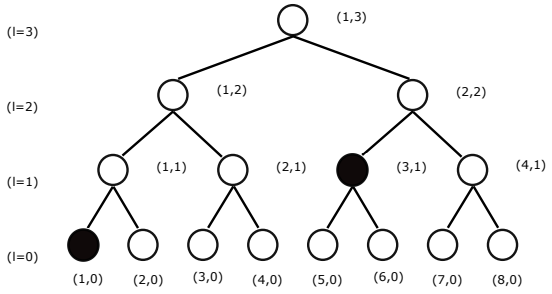

Fig. 1. A binary tree representing the flow aggregation model with l denoting the level of the tree and (k, l) the kth node on the l-th level ($(1, 0)$ is an HH, $(3, 1)$ is an HHH. Nodes $(1, 1)$, $(1, 2)$, $(1, 3)$, and $(2, 2)$ are reflecting points).

1.1. Main Results

We develop an active inference strategy for detecting an unknown number of HHHes among a large number N of aggregated data points with unknown distributions. The performance measure is the number of samples (i.e., detection delay) required for achieving a confidence level of $1 - \epsilon$ (i.e.,

[0]This work was supported by the U.S. Army Research Office under Grant W911NF-17-1-0464.

the probability that the declared HHH set does not equal to the true set is bounded by ϵ). By fully exploiting the flow aggregation model, the proposed active inference strategy has a sample complexity that is order optimal in both the size N of the search space and the reliability constraint ϵ.

Referred to as Confidence Bounds based Random Walk (CBRW), the proposed strategy consists of a global random walk on the tree interwoven with a local confidence-bound based test. Specifically, it induces a biased random walk that initiates at the root of the tree and eventually arrives and terminates at an HHH with the required reliability. Each move in the random walk is guided by the output of a local confidence-bound based sequential test carried on each child of the node currently being visited by the random walk. This local sequential test module ensures that the global random walk is more likely to move toward the HHH than move away from it and that the random walk terminates at a true HHH with a sufficiently high probability.

The sample complexity of CBRW is analyzed using properties of biased random walk on a tree and large deviation results on the concentration of the sample mean statistic. We show that the sample complexity of CBRW is in the order of $O(\log N + \log \frac{1}{\epsilon})$ provided that the gap between the mean value of each flow count and the given threshold is bounded away from 0. It is thus order optimal in both N and ϵ as determined by information-theoretic lower bounds. Of particular significance is that the effect on the sample complexity from an enlarged search space (increasing N) and an enhanced reliability (decreasing ϵ) is additive rather than multiplicative. This results from the random walk structure which effectively separates two objectives of moving to the HHHes with $O(\log N)$ samples and declaring the HHH at the desired confidence level with $O(\log \frac{1}{\epsilon})$ samples. The proposed strategy applies to unknown heavy-tailed distribution models and preserves its order optimality in both N and ϵ. Comprising of calculating confidence bounds of the mean and performing simple comparisons, the proposed strategy is computationally efficient.

The result also finds applications in noisy group testing and adaptive sampling with noisy response as detailed in the full version in [3].

1.2. Related Work

Prior solutions for online detection of HHHes typically involve adjusting which prefixes to monitor either at the arrival of each packet [4, 5, 6], or at periodic intervals [7, 8]. A particularly relevant work is the adaptive monitoring algorithm proposed by Jose *et al.* [7], where a fixed number of measurement rules are adjusted at periodic intervals based on the aggregate packet counts matching to each of these rules. At each time interval, the aggregate count is compared to a heuristically chosen threshold (e.g., a fraction of link capacity), to determine whether it is an HHH, and whether the rules need to be kept in the next interval, or expanded to monitor the children of the prefix, or collapsed and combined with upstream nodes. While the proposed CBRW has a similar flavor of moving among parent and children, which is very much inherent to the HHH detection problem, the decision criteria used to adjust the prefix is different. Instead of comparing with a fixed threshold, our decision is based on statistical metric determined by the desired detection error. Different from the heuristic studies in the literature, the proposed strategy offers performance grantee and order optimality. We provide a rigorous framework that succinctly captures the trade-off between detection time and overall detection performance. In [9], a quantitative group testing strategy was developed for HH detection based on IP prefix aggregation. It does not address HHH detection or the issue of unknown stochastic models of traffic flows.

2. PROBLEM FORMULATION

Consider a set of N stochastic processes conforming to a binary-tree structure with K leaf nodes as illustrated in Fig. 1. Let (l,k) ($l = 0, 1, \ldots, L, k = 1, \ldots, 2^{L-l}$) denote the kth node at level l of the tree. Let $\{X_{k,l}(t)\}_{t=1}^{\infty}$ denote the corresponding random process which is independent and identically distributed with an unknown distribution $f_{k,l}$ and an unknown mean $\mu_{k,l}$. For any node (k,l), $\mu_{k,l}$ is the sum of the mean values of the children of (k,l).

Associated with each level l of the tree is a given threshold η_l. An HHH is defined recursively in terms of l. Specifically, an HHH, at level l, is a node whose mean value remains above the threshold η_l after excluding all its decedents with mean value above the threshold at their respective level.

An active inference strategy $\pi = (\{a_t\}_{t\geq 1}, T_\pi, \mathcal{S}_\pi)$ consists of a sampling strategy $\{a_t\}_{t\geq 1}$, a stopping rule T_π, and a terminal decision rule $\mathcal{S}_\pi$. The sampling strategy $\{a_t\}_{t\geq 1}$ is a sequence of functions mapping from past actions and observations to an aggregated data point to be sampled at the current time t. The stopping rule T_π determines when to terminate the search, and the decision rule $\mathcal{S}_\pi$ declares the detected set of HHHes at the time of stopping. Let $\mathbb{E}_{\mathcal{F}}$ and $\mathbb{P}_{\mathcal{F}}$ denote, respectively, the expectation and the probability measure under distribution model $\mathcal{F} = \{f_{(k,l)}\}_{\substack{l=0,1,\ldots,L \\ k=1,\ldots,2^{L-l}}}$. The objective is as follows: minimize$_\pi$ $\mathbb{E}_{\mathcal{F}} T_\pi$, *s.t.* $\mathbb{P}_{\mathcal{F}}[\mathcal{S}_\pi \neq \mathcal{S}] \leq \epsilon$, where $\mathcal{S}$ is the true set of HHHes.

We consider a general distribution $f_{k,l}$ for each process. Due to different concentration behaviors, sub-Gaussian and heavy-tailed distributions are treated separately.

3. AN ACTIVE INFERENCE STRATEGY: CBRW

In this section, we present the Confidence Bounds based Random Walk (CBRW) policy and establish its order-optimal performance. We focus on the case of a single HHH and sub-Gaussian distributions. Extensions to multiple HHHes detection and heavy-tailed distributions are discussed in Sec. 4.

3.1. The CBRW Policy

The basic structure of CBRW consists of a global random-walk module interwoven with a local CB-based sequential test module at each step of the random walk. Specifically, the CBRW policy performs a biased random walk on the tree that eventually arrives and terminates at the HHH with the required reliability. Each move in the random walk (i.e., which neighboring node to visit next) is guided by the output of the local CB-based sequential test module. This module ensures that the random walk is more likely to move toward the HHH than to move away from the HHH and that the random walk terminates at the true HHH with high probability.

Consider first the local CB-based sequential test module. This local sequential test is carried out on a specific node (random process) $\{X(t)\}_{t=1}^{\infty}$, where we have omitted the node index (k,l) for simplicity. The goal is to determine whether the mean value of $\{X(t)\}_{t=1}^{\infty}$ is below a given threshold η at a confidence level of $1-\beta$ or above the threshold at a confidence level of $1-\alpha$. If the former is true, the test module outputs 0, indicating this node is unlikely to be an ancestor of an HHH or the HHH itself. If the latter is true, the output is 1. Let $\mathcal{L}(\alpha,\beta,\eta)$ denote this local sequential test with given parameters $\{\alpha,\beta,\eta\}$. It sequentially collects samples from $\{X(t)\}_{t=1}^{\infty}$. After collecting each sample, it determines whether to terminate the test and if yes, which value to output based on the following rule:

- If $\overline{X}(s)-\sqrt{\frac{2\xi\log\frac{2s^3}{\alpha}}{s}}>\eta$, terminate and output 1.
- If $\overline{X}(s)+\sqrt{\frac{2\xi\log\frac{2s^3}{\beta}}{s}}<\eta$, terminate and output 0.
- Otherwise, continue taking samples,

where $\overline{X}(s)$ denotes the sample mean obtained form s observations and ξ is the distribution parameter specified in Sec. 3.2.

We now specify the random walk on the tree based on the outputs of the local CB-based tests. Let (k,l) denote the current location of the random walk (which is initially set at the root node). Consider first (k,l) is a non-leaf node with $l>0$. The node (k,l) is first probed by the local module $\mathcal{L}(\alpha,\beta,\eta)$ with parameters set to $\alpha=\beta=p_0$ where $p_0\in(0,1-\frac{1}{\sqrt[3]{2}})$ and $\eta=\eta_l$. If the output is 0, the random walk moves to the parent of (k,l). If the output is 1, then the left child of (k,l) is tested by the local module $\mathcal{L}(\alpha,\beta,\eta)$ with parameters set to $\alpha=\beta=p_0$ and $\eta=\eta_{l-1}$. If the output is 1, the random walk moves to the left child. Otherwise, the right child of (k,l) is tested with the same set of parameters, and the random walk moves to the right child if this test outputs 1. If the outputs of the tests at (k,l) and its children are 1, 0, and 0, respectively, then (k,l) is likely to be an HHH and the random walk stays at (k,l). When the random walk stays at the same node (k,l), the same tests are repeated on (k,l) and its children with an increased confidence level. We increase the confidence level by dividing α and β by 2 iteratively. When the current value of α and β becomes smaller than $\frac{\epsilon}{3LC_{p_0}^H}$, the random walk stops and declares (k,l) as the HHH. The value of

$$C_{p_0}^H=\frac{1}{\left(1-\exp(-2(1-2(1-p_0)^3)^2)\right)^2} \tag{1}$$

ensures the desired confidence level of $1-\epsilon$ at detection of the HHH. If the random walk moves to a new location the values of α and β is reset to p_0. When the random walk arrives at a leaf node (k,l) with $l=0$, the leaf node is tested by the local module $\mathcal{L}(\alpha,\beta,\eta)$ with parameters set to $\alpha=\frac{\epsilon}{3LC_{p_0}^H}$, $\beta=p_0$ and $\eta=\eta_0$. If the output is 1, the random walk stops and declares (k,l) as the HHH. Otherwise, the random walk moves to the parent of (k,l). A pseudo-code for CBRW is given in [3].

3.2. Performance Analysis

We first analyze the sample complexity of the local CB-based sequential test module $\mathcal{L}(\alpha,\beta,\eta)$ in the lemma below. We then analyze the behavior of the random walk to establish the number of times that the local sequential test is carried out.

Recall that a real-valued random variable X is called sub-Gaussian [10] if, for all $\lambda\in(-\infty,\infty)$, $\mathbb{E}[e^{\lambda(X-\mathbb{E}[X])}]\le e^{\xi\lambda^2/2}$ for some constant $\xi>0$. We assume (an upper bound on) ξ is known. For sub-Gaussian random variables, Chernoff-Hoeffding concentration inequalities hold.

Lemma 1 *Let μ denote the expected value of an i.i.d. sub-Gaussian random process $\{X(t)\}_{t=1}^{\infty}$. Let $T_{\mathcal{L}}$ be the stopping time of the CB-based sequential test $\mathcal{L}(\alpha,\beta,\eta)$ applied to $\{X(t)\}_{t=1}^{\infty}$. We have, in the case of $\mu>\eta$,*

$$\mathbb{P}[\overline{X}(T_{\mathcal{L}})+\sqrt{\frac{2\xi\log\frac{2T_{\mathcal{L}}^3}{\beta}}{T}}<\eta]\le\beta, \tag{2}$$

$$\mathbb{E}[T_{\mathcal{L}}]\le\frac{48}{(\mu-\eta)^2}\log\frac{24\sqrt[3]{\frac{2}{\alpha}}}{(\mu-\eta)^2}+2. \tag{3}$$

In the case of $\mu<\eta$,

$$\mathbb{P}[\overline{X}(T_{\mathcal{L}})-\sqrt{\frac{2\xi\log\frac{2T_{\mathcal{L}}^3}{\alpha}}{2T}}>\eta]\le\alpha, \tag{4}$$

$$\mathbb{E}[T_{\mathcal{L}}]\le\frac{48}{(\mu-\eta)^2}\log\frac{24\sqrt[3]{\frac{2}{\beta}}}{(\mu-\eta)^2}+2. \tag{5}$$

Proof: Omitted due to space limit. See a detailed proof in [3].

The inequalities (2) and (4) establish the confidence levels for the local sequential CB-based test. Both results on the confidence levels and the sample complexity are based on the Chernoff-Hoeffding concentration inequalities.

The sample complexity of CBRW is order optimal in both N and $\frac{1}{\epsilon}$ as stated in Theorem 1 below.

Theorem 1 *Assume that there exists $\delta>0$ such that $\mu_{k,l}-\eta_l\ge\delta$ for all (k,l) ($l=0,1,\ldots,L,k=1,\ldots,2^{L-l}$). We have*

$$\mathbb{E}_{\mathcal{F}}[T_{CBRW}]=O(\log_2 N+\log\frac{1}{\epsilon}), \tag{6}$$

and

$$\mathbb{P}_{\mathcal{F}}[\mathcal{S}_{CBRW} \neq \mathcal{S}] \leq \epsilon. \tag{7}$$

Proof: Omitted due to space limit. See a detailed and finite-time analysis in [3].

4. EXTENSIONS

Detecting $|\mathcal{S}| > 1$ HHHes with $|\mathcal{S}|$ known can be easily implemented by sequentially locating the HHHes one by one. We assume that each HHH can be removed after it is located by CBRW[1]. To ensure that the reliability constraint holds, we replace ϵ with $\frac{\epsilon}{|\mathcal{S}|}$ in each search of a single HHH. The reliability constraint holds by union bound on the error probabilities of the searches for a single HHH.

When the number of HHHes is unknown, but an upper bound $S_{\max} \geq |\mathcal{S}|$ on the number of HHHHes is known, we can similarly detect the HHHes one by one. To ensure that the reliability constraint holds, we replace ϵ with $\frac{\epsilon}{2S_{\max}}$ in each search of a single HHH. The stopping rule for the overall search can be implemented by testing the root node. Specifically, the root node is tested by $\mathcal{L}(\epsilon_0, \epsilon_0, \eta_L)$ with $\epsilon_0 = \frac{\epsilon}{2S_{\max}}$ every $LC_{p_0}^H$ steps in the random walk. The reliability constraint holds by union bound on the error probabilities of the searches for a single HHH and error in stopping the overall search before finding all HHHes. The sample complexities of finding single HHHes simply add up to $O(|\mathcal{S}| \log K + |\mathcal{S}| \log \frac{1}{\epsilon})$ overall sample complexity.

The extension to more general distribution models can be implemented by modifying the local CB-based test $\mathcal{L}$ in a way that the confidence levels remain the same. As a result, the behavior of the random walk on the tree remains the same.

Specifically, for heavy-tailed distributions with a finite b'th $(1 < b < 2)$ moment we modify the test $\mathcal{L}$ as follows.

- If $\widehat{X}(s,\alpha) - 4u^{1/b}(\frac{\log \frac{2s^3}{\alpha}}{s})^{\frac{b-1}{b}} > \eta$, terminate and output 1.
- If $\widehat{X}(s,\beta) + 4u^{1/b}(\frac{\log \frac{2s^3}{\beta}}{s})^{\frac{b-1}{b}} < \eta$, terminate and output 0.
- Otherwise, continue taking samples,

where u is an upper bound on $\mathbb{E}[X^b]$ and $\widehat{X}(s,p)$ is the truncated sample mean defined as

$$\widehat{X}(s,p) = \frac{1}{s}\sum_{t=1}^{s} X(t)\mathbb{1}\left\{|X(t)| \leq (\frac{ut}{\log \frac{1}{p}})^{1/b}\right\}. \tag{8}$$

The resulting CBRW achieves the same $O(\log N + \log \frac{1}{\epsilon})$ sample complexity. See [3] for the detailed analysis.

5. SIMULATION

We compare the performance of CBRW with the strategy proposed in [7]. We consider a number of randomly chosen HHHes and a threshold set to 5% of the expected total traffic at leaf nodes. The traffic at the leaf nodes is randomly generated according to Poisson distributions with randomly chosen parameters. The results are obtained from 1000 Monte-Carlo runs. As shown in Fig. 2, while the existing strategy provides no reliability guarantee, CBRW satisfies the required error rate. This also shows that CBRW can be conservative and the algorithm parameters such as ϵ can be adjusted to further reduce the detection time for a given reliability constraint.

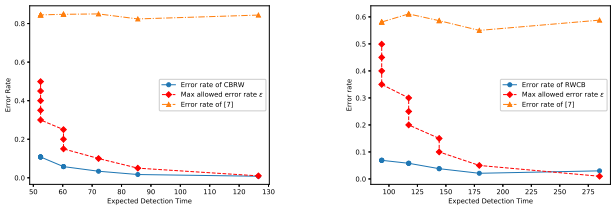

Fig. 2. The comparison between the detection error of CBRW and the strategy proposed in [7]on a graph with $L = 8$, 12 HHHes and 24 counters (on the left) and on a graph with $L = 9$, 13 HHHes and 26 counters (on the right).

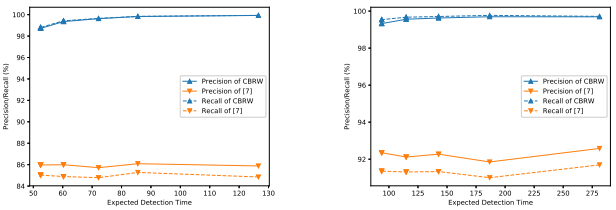

Fig. 3. The comparison between the Precision/Recall of CBRW and the strategy proposed in [7] on a graph with $L = 8$, 12 HHHes and 24 counters (on the left) and on a graph with $L = 9$, 13 HHHes and 26 counters (on the right).

In order to show a fair comparison, we also consider the performance measures used in [7]: Precision defined as the percentage of true HHHes in the detected set ($\frac{|\mathcal{S}\cap\mathcal{S}_\pi|}{|\mathcal{S}_\pi|} \times 100$), and Recall defined as the percentage of detected HHHes ($\frac{|\mathcal{S}\cap\mathcal{S}_\pi|}{|\mathcal{S}|} \times 100$). As shown in Fig. 3, CBRW shows a higher Precision and Recall in comparison to the existing strategy.

6. CONCLUSION

In this paper, we studied the problem of detecting HHHes among a large number of traffic flows modeled as random processes with unknown distributions. The proposed strategy detects the HHHes at the required confidence level with an order-optimal logarithmic sample complexity in both the problem size and the reliability constraint. The result also finds applications in noisy group testing and adaptive sampling with noisy response as detailed in the full version in [3].

[1]For example, the packet count of each detected HHH can be subtracted from the packet count of the parents.

7. REFERENCES

[1] K. Thompson, G. Miller, and R. Wilder, "Wide-area internet traffic patterns and characteristics," *IEEE Network*, vol. 11, pp. 1017, Nov 1997.

[2] P. E. Ayres, H. Sun, H. J. Chao, and W. C. Lau "Alpi: A ddos defense system for high-speed networks," *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 10, pp. 18641776, 2006. 24(10):18641776, 2006.

[3] S. Vakili, Q. Zhao, C. Liu, C.-N. Chuah, "Anomaly Detection in Hierarchical Data Streams under Unknown Models", available at: *arXiv:1709.03573 [cs.LG]*

[4] G. Cormode, F. Korn, S. Muthukrishnan, and D. Srivastava, "Finding hierarchical heavy hitters in streaming data" *ACM Trans. Knowl. Discov. Data*, vol. 1, no.4, pp 1-48, 2008.

[5] L. Yuan, C. Chuah, and P. Mohapatra, "Progme: Towards programmable network measurement" *In Proceedings of the 2007 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications, SIGCOMM, ACM* 2007.

[6] Y. Zhang, S. Singh, S. Sen, N. Duffield, and C. Lund, "Online identification of hierarchical heavy hitters: Algorithms, evaluation, and applications," *In Proceedings of the 4th ACM SIGCOMM Conference on Internet Measurement, IMC, ACM*, 2004.

[7] L. Jose, M. Yu, and J. Rexford, "Online measurement of large traffic aggregates on commodity switches" *In Proceedings of the 11th USENIX Conference on Hot Topics in Management of Internet, Cloud, and Enterprise Networks and Services, Hot-ICE, USENIX Association*, 2011.

[8] M. Mitzenmacher, T. Steinke, and J. Thaler, "Hierarchical heavy hitters with the space saving algorithm," *In Proceedings of the Meeting on Algorithm Engineering and Experiments, ALENEX*, 2012.

[9] C. Wang, Q. Zhao, CN Chuah, "Optimal Nested Test Plan for Combinatorial Quantitative Group Testing," https://arxiv.org/abs/1407.2283

[10] P. Chareka, O. Chareka, S. Kennendy, "Locally Sub-Gaussian Random Variable and the Stong Law of Large Numbers," *Atlantic Electronic Journal of Mathematics*, vol. 1, no. 1, pp. 75-81, 2006.