

Fishing in Poisson Streams: Focusing on the Whales, Ignoring the Minnows

Maxim Raginsky
ECE, Duke University
Durham, NC 27708, USA

Sina Jafarpour
CS, Princeton University
Princeton, NJ 08540, USA

Rebecca Willett
ECE, Duke University
Durham, NC 27708, USA

Robert Calderbank
EE, Princeton University
Princeton, NJ 08540, USA

Abstract—This paper describes a low-complexity approach for reconstructing average packet arrival rates and instantaneous packet counts at a router in a communication network, where the arrivals of packets in each flow follow a Poisson process. Assuming that the rate vector of this Poisson process is sparse or approximately sparse, the goal is to maintain a compressed summary of the process sample paths using a small number of counters, such that at any time it is possible to reconstruct both the total number of packets in each flow and the underlying rate vector. We show that these tasks can be accomplished efficiently and accurately using compressed sensing with expander graphs. In particular, the compressive counts are a linear transformation of the underlying counting process by the adjacency matrix of an unbalanced expander. Such a matrix is binary and sparse, which allows for efficient incrementing when new packets arrive. We describe, analyze, and compare two methods that can be used to estimate both the current vector of total packet counts and the underlying vector of arrival rates.

I. INTRODUCTION

Successful management of large-scale communication networks rests crucially on the availability of accurate traffic measurements. From the viewpoint of such tasks as billing/accounting or intrusion detection, network traffic is composed of packet flows (or streams) arriving at or departing from routers in the network. As both the number of users and the data rates continue growing, there is increasing emphasis on traffic measurement architectures that are accurate, fast, and cheap. Naturally, some trade-offs between these three desiderata are inevitable. For instance, one could keep a dedicated counter for each flow, but, depending on the type of memory used, one could end up with an implementation that is either fast but expensive and unable to keep track of a large number of flows (e.g., using SRAMs, which have low access times, but are expensive and physically large) or cheap and high-density but slow (e.g., using DRAMs, which are cheap and small, but have longer access times).

Recent work has shown that a reasonable compromise between accuracy, speed and cost can be found if one takes into account certain prior knowledge about the relative flow sizes in a typical network. In particular, there is empirical evidence [1], [2] that flow sizes in IP networks follow a *heavy-tail* pattern: just a few flows (say, 10%) carry most of the traffic (say, 90%). Based on this observation, Estan and Varghese [3] proposed two methodologies (“sample-and-hold” and “multistage filters”) that use a small number of counters to keep track only of the flows whose sizes exceed a given

fraction of the total bandwidth. More recently, Lu et al. [4] developed a new technique, termed “Counter Braids,” which uses sparse random graphs to aggregate (or “braid”) the raw packet counts into a small number of counters. The total size of *each* flow can then be recovered at the end of a measurement epoch using a message passing decoder.

The approach of Estan and Varghese [3] allows one to keep track only of the few heavy flows while ignoring the rest (“focusing on the elephants, ignoring the mice,” as they put it), while Lu et al. [4] can recover the *entire* vector of flow sizes. Moreover, these two approaches rely on different modeling assumptions. Specifically, in [3] the flow sizes are assumed to be deterministic and subject to the heavy-tail behavior, while in [4] the flow sizes are i.i.d. realizations of a random variable with a heavy-tail distribution.

A. Our contribution

The present paper considers a more realistic setting where each flow (or stream) is modeled as a Poisson process with an unknown rate (measured in packets per unit time), and it is the *rates* corresponding to the streams at a given router that possess the heavy-tail property. This modeling assumption combines certain aspects of [3] and [4]: the heavy-tail property is present both on the level of coarse-grained, time-averaged behavior of the flows and on the level of actual traffic patterns, which are stochastic. Moreover, our model goes beyond the i.i.d. assumption of [4] and can account for the heterogeneous nature of the different flows entering a particular router.

The main goal is to reconstruct the underlying vector of rates while maintaining a small number of counters with low access times. To accomplish this goal, we exploit our recent work [5] on compressed sensing (CS) with Poisson-distributed observations. Mathematically, the heavy-tail property can be restated as follows: the vector of rates is, to a good approximation, *sparse*. This sparsity interpretation strongly suggests that CS can be used to accurately recover the underlying vector of rates from a small number of judiciously designed linear transformations of the observed flows. Building on the results from [5], we show that the raw packet counts can be mapped into a small number of “compressed” counts using the adjacency matrix of a properly constructed unbalanced expander. Such an adjacency matrix has binary entries and is sparse (i.e., each column has a small constant number of ones), which ensures that the counts can be updated using a small

number of operations as new packets arrive. The resulting architecture can be used to recover the raw packet counts as well. Since we are dealing here with Poisson streams, we would like to push the metaphor further and say that we are “focusing on the whales, ignoring the minnows.”

We analyze the performance of our scheme theoretically, describe an efficient implementation, and present preliminary experimental results.

B. Notation

Given a vector $\mathbf{u} \in \mathbb{R}^m$ and a set $S \subseteq \{1, \dots, m\}$, we will denote by \mathbf{u}^S the vector obtained by setting to zero all coordinates of \mathbf{u} that are in S^c , the complement of S : $\forall 1 \leq i \leq m, u_i^S = u_i 1_{\{i \in S\}}$. Given some $1 \leq k \leq m$, let S be the set of positions of the k largest (in magnitude) coordinates of \mathbf{u} . Then $\mathbf{u}^{(k)} \triangleq \mathbf{u}^S$ will denote the *best k -term approximation* of \mathbf{u} (in any norm on \mathbb{R}^m), and

$$\sigma_k(\mathbf{u}) \triangleq \|\mathbf{u} - \mathbf{u}^{(k)}\|_1 = \sum_{i \in S^c} |u_i|$$

will denote the resulting ℓ_1 approximation error. The ℓ_0 quasinnorm measures the number of nonzero coordinates of \mathbf{u} : $\|\mathbf{u}\|_0 \triangleq \sum_{i=1}^m 1_{\{u_i \neq 0\}}$. Given a vector \mathbf{u} , we will denote by \mathbf{u}^+ the vector obtained by setting to zero all negative components of \mathbf{u} : for all $1 \leq i \leq m$, $u_i^+ = \max\{0, u_i\}$.

II. PROBLEM FORMULATION

We wish to monitor a large number N of packet flows using a much smaller number M of counters. Each flow is a homogeneous Poisson process. Specifically, let $\boldsymbol{\lambda}^* \in \mathbb{R}_+^N$ denote the vector of rates, and let \mathbf{U} denote the random process $\mathbf{U} = \{\mathbf{U}_t\}_{t \in \mathbb{R}_+}$ with sample paths in \mathbb{Z}_+^N , where for any $t \in \mathbb{R}_+$ and any $\mathbf{k} \in \mathbb{Z}_+^N$ we have

$$\mathbb{P}_{\boldsymbol{\lambda}^*}(\mathbf{U}_t = \mathbf{k}) = \prod_{i=1}^N \frac{(\lambda_i^* t)^{k_i}}{k_i!} e^{-t\lambda_i^*}.$$

In other words, for each $i \in \{1, \dots, N\}$, the i th component of \mathbf{U} , which we will denote by $U^{(i)}$, is a homogeneous Poisson process with the rate of λ_i^* arrivals per unit time, and all the $U^{(i)}$'s are mutually conditionally independent given $\boldsymbol{\lambda}^*$.

The counters are updated in discrete time, every τ time units. Let $\mathbf{X} = \{\mathbf{X}_n\}_{n \in \mathbb{Z}_+}$ denote the sampled version of \mathbf{U} , where $\mathbf{X}_n \triangleq \mathbf{U}_{n\tau}$. The update takes place as follows. We have a binary matrix $\mathbf{A} \in \{0, 1\}^{M \times N}$, and at each time n let $\mathbf{Y}_n = \mathbf{A}\mathbf{X}_n$. The probabilistic law governing the evolution of the counter contents is now

$$\mathbb{P}_{\mathbf{A}, \boldsymbol{\lambda}^*}(\mathbf{Y}_n = \boldsymbol{\ell}) = \prod_{j=1}^M \frac{(n\tau \mathbf{A}\boldsymbol{\lambda}^*)_j^{\ell_j}}{\ell_j!} e^{-(n\tau \mathbf{A}\boldsymbol{\lambda}^*)_j},$$

where the case $(\mathbf{A}\boldsymbol{\lambda}^*)_j = 0$ for any $1 \leq j \leq M$ is handled using the fact that $\lambda^\ell e^{-\lambda}/\ell! \rightarrow 1_{\{\ell=0\}}$ as $\lambda \rightarrow 0$. In other words, \mathbf{Y} is a sampled version of an M -dimensional homogeneous Poisson process with the rate vector $\mathbf{A}\boldsymbol{\lambda}^*$.

The goal is to estimate the unknown rate vector $\boldsymbol{\lambda}^*$ after n time steps given $\mathbf{Y}^n \triangleq (\mathbf{Y}_1, \dots, \mathbf{Y}_n)$ using an estimator $\hat{\boldsymbol{\lambda}}_n$

based on \mathbf{Y}^n : $\hat{\boldsymbol{\lambda}}_n = \hat{\boldsymbol{\lambda}}_n(\mathbf{Y}^n)$. We measure the quality of such an estimator by the expected ℓ_1 risk:

$$R(\hat{\boldsymbol{\lambda}}_n, \boldsymbol{\lambda}^*) \triangleq \mathbb{E}_{\boldsymbol{\lambda}^*} \|\hat{\boldsymbol{\lambda}}_n - \boldsymbol{\lambda}^*\|_1 = \mathbb{E}_{\boldsymbol{\lambda}^*} \left[\sum_{i=1}^N |\hat{\lambda}_{n,i} - \lambda_i^*| \right],$$

where the expectation is taken w.r.t. the underlying flow process $\mathbf{X} \sim \mathbb{P}_{\boldsymbol{\lambda}^*}$. Assuming that the unknown rate vector $\boldsymbol{\lambda}^*$ is a member of a given class Λ^* , we would like to design the counter update matrix \mathbf{A} and an accompanying sequence of estimators $\{\hat{\boldsymbol{\lambda}}_n\}$ to attain low risk $R(\hat{\boldsymbol{\lambda}}_n, \boldsymbol{\lambda}^*)$ over Λ^* . One particular class of interest, which pertains to the heavy-tail behavior of network traffic, is defined by

$$\Sigma_{\alpha, L_0} \triangleq \{\boldsymbol{\lambda} \in \mathbb{R}_+^N : \|\boldsymbol{\lambda}\|_1 \leq L_0; \sigma_k(\boldsymbol{\lambda}) \leq L_0 k^{-\alpha}, \forall k\} \quad (1)$$

for some $L_0 > 0$ and $\alpha \geq 1$. Here, α is the power-law exponent that controls the tail behavior; in particular, the extreme regime $\alpha \rightarrow +\infty$ describes the fully sparse setting.

III. PRELIMINARIES

As we show in the sequel, a good choice for the counter update matrix \mathbf{A} is the adjacency matrix of a suitably constructed expander. Adjacency matrices of high-quality expanders have been proposed as an alternative to dense, random measurement matrices for sparse signal recovery [6]–[10]. This section summarizes the key results on expanders, as well as the results from our earlier work [5] on the use of expanders for sparse recovery under the Poisson observation model.

A. Expanders and sparse recovery

Definition 1. A (k, ε) -unbalanced expander, or simply a (k, ε) -expander, is a bipartite simple graph $G = (A, B, E)$ with left degree d , such that for any $S \subset A$ with $|S| \leq k$, the set of neighbors $\mathcal{N}(S)$ of S has size $|\mathcal{N}(S)| \geq (1 - \varepsilon)d|S|$.

Here, A (resp., B) corresponds to the components of the original signal (resp., its compressed representation). Hence, for a given $|A|$, a “high-quality” expander should have $|B|$, d , and ε as small as possible, while k should be as close as possible to $|B|$. The following proposition (cf. [5], [6]) tells us what we can expect:

Proposition 1. For any $1 \leq k \leq N/2$ and any $\varepsilon \in (0, 1)$, there exists a (k, ε) -expander with left set size N , left degree $d = O\left(\frac{\log(N/k)}{\varepsilon}\right)$ and right set size $M = O\left(\frac{k \log(N/k)}{\varepsilon^2}\right)$.

From this point on, given N and $1 \leq k \leq N/2$, we will denote by $G_{k,N}$ a fixed expander with $\varepsilon = 1/16$ whose existence is guaranteed by the above proposition (the value of ε is fixed for convenience). The following proposition is key to the use of expanders for sparse recovery:

Proposition 2. Let $\Phi = \mathbf{A}/d$ be the normalized adjacency matrix of $G_{k,N}$, and let \mathbf{u}, \mathbf{v} be two vectors in \mathbb{R}^N , such that $\|\mathbf{u}\|_1 \geq \|\mathbf{v}\|_1 - \Delta$ for some $\Delta > 0$. Then

$$\|\mathbf{u} - \mathbf{v}\|_1 \leq 4\sigma_k(\mathbf{u}) + 4\|\Phi\mathbf{u} - \Phi\mathbf{v}\|_1 + 2\Delta.$$

For future reference, we note that, since our expander is regular, there exists a minimal set $\Omega \subset A$ of size M , such that its neighborhood covers all of B , i.e., $\mathcal{N}(\Omega) = B$. Let $\mathbf{I}_\Omega \in \mathbb{R}^N$ be the vector with components $I_{\Omega,i} = 1_{\{i \in \Omega\}}$. In that case, note that $\Phi \mathbf{I}_\Omega \succeq I_{M \times 1/d}$.

B. Expander-based CS under the Poisson model

In [5], we have considered the following problem: Let $\theta^* \in \mathbb{R}_+^N$ be an unknown vector of Poisson intensities with known ℓ_1 norm $\|\theta^*\|_1 = L$ (in general, L may be a known upper bound on $\|\theta^*\|_1$). Given a fixed $1 \leq k \leq N/2$, let Φ be the normalized adjacency matrix of $G_{k,N}$. We observe a random vector $\mathbf{Z} \in \mathbb{Z}_+^M$ distributed according to $\mathbf{Z} \sim \text{Poisson}(\Phi \theta^*)$.

Let $\Theta \subset \mathbb{R}_+^N$ be a finite or countable set of candidate estimators of θ^* such that $\|\theta\|_1 \leq L, \forall \theta \in \Theta$, and for a given $c > 0$ define the set

$$\mathcal{F}_c \triangleq \left\{ \mathbf{f} \triangleq \theta + cL\mathbf{I}_\Omega : \theta \in \Theta \right\}.$$

Moreover, let $\text{pen}(\cdot) : \Theta \rightarrow \mathbb{R}_+$ be a *penalty* (or *regularization*) functional satisfying the *Kraft inequality*,

$$\sum_{\theta \in \Theta} e^{-\text{pen}(\theta)} \leq 1.$$

Since there is a one-to-one correspondence between Θ and \mathcal{F}_c , we will overload our notation and let $\text{pen}(\mathbf{f})$ denote $\text{pen}(\theta)$ whenever $\mathbf{f} = \theta + cL\mathbf{I}_\Omega$. In [5], we have shown the following:

Proposition 3. Consider the penalized maximum likelihood estimator (pMLE)

$$\hat{\mathbf{f}} \triangleq \underset{\mathbf{f} \in \mathcal{F}_c}{\text{argmin}} [-\log \mathbb{P}_{\Phi \mathbf{f}}(\mathbf{Z}) + 2\text{pen}(\mathbf{f})] \quad (2a)$$

$$\hat{\theta} \triangleq \hat{\mathbf{f}} - cL\mathbf{I}_\Omega. \quad (2b)$$

Then

$$\begin{aligned} \mathbb{E}\|\theta^* - \hat{\theta}\|_1^2 &= O\left(\sigma_k(\theta^*)^2 + c(ML)^2(2d + Mc)\right) \\ &+ O\left(\left(\frac{d}{c} + M\right) \min_{\theta \in \Theta} \left[\|\theta^* - \theta\|_1^2 + \frac{L\text{pen}(\theta)}{d}\right]\right). \end{aligned} \quad (3)$$

Prop. 3 effectively states that the squared ℓ_1 error of $\hat{\theta}$ scales with M times the best penalized ℓ_1 approximation error plus the k -term approximation error of θ^* . The first term in (3) is smaller for sparser θ^* , and the second term is smaller when there is a θ which is simultaneously a good approximation to θ^* and has a low penalty.

IV. TWO ESTIMATION STRATEGIES

We consider two estimation strategies. In both cases, we let our measurement matrix \mathbf{A} be the adjacency matrix of a $G_{k,N}$ for a fixed $k \leq N/2$. The first strategy, which we call the *direct method*, uses expander-based CS to first recover an estimate of \mathbf{X}_n from \mathbf{Y}_n , then constructs an estimate of λ^* . The second strategy, which we call the *penalized MLE* strategy (or pMLE), relies on the Poisson CS machinery presented in Section III-B and can be used when only the rates are of

interest. One benefit of pMLE compared to the direct method is its low complexity, which is derived from a preprocessing step based on the structure of the underlying expander $G_{k,N}$.

A. The direct method

The first approach is to use expander-based CS to obtain an estimate $\widehat{\mathbf{X}}_n$ of \mathbf{X}_n from $\mathbf{Y}_n = \mathbf{A}\mathbf{X}_n$, followed by letting

$$\widehat{\lambda}_n^{\text{dir}} = \frac{\widehat{\mathbf{X}}_n^+}{n\tau}. \quad (4)$$

This strategy is based on the observation that $\mathbf{X}_n/(n\tau)$ is the maximum-likelihood estimator of λ^* , and will serve as a “baseline” against which the penalized MLE will be compared. To obtain $\widehat{\lambda}_n^{\text{dir}}$, we need to solve the convex program

$$\text{minimize } \|\mathbf{u}\|_1 \quad \text{subject to } \mathbf{A}\mathbf{u} = \mathbf{Y}_n$$

which can be cast as a linear program [6]. The resulting solution $\widehat{\mathbf{X}}_n$ may have negative coordinates¹, hence the use of the $(\cdot)^+$ operation in (4). We then have the following result:

Theorem 4.

$$R\left(\widehat{\lambda}_n^{\text{dir}}, \lambda^*\right) \leq 4\sigma_k(\lambda^*) + \frac{\|(\lambda^*)^{1/2}\|_1}{\sqrt{n\tau}}, \quad (5)$$

where $(\lambda^*)^{1/2}$ is the vector with components $\sqrt{\lambda_i^*}, \forall i$.

Proof: We first observe that, by construction, $\widehat{\mathbf{X}}_n$ satisfies the relations $\mathbf{A}\widehat{\mathbf{X}}_n = \mathbf{A}\mathbf{X}_n$ and $\|\widehat{\mathbf{X}}_n\|_1 \leq \|\mathbf{X}_n\|_1$. Hence,

$$\begin{aligned} \mathbb{E}\|\widehat{\mathbf{X}}_n - n\tau\lambda^*\|_1 &\leq \mathbb{E}\|\widehat{\mathbf{X}}_n - \mathbf{X}_n\|_1 + \mathbb{E}\|\mathbf{X}_n - n\tau\lambda^*\|_1 \\ &\leq 4\mathbb{E}\sigma_k(\mathbf{X}_n) + \mathbb{E}\|\mathbf{X}_n - n\tau\lambda^*\|_1 \end{aligned} \quad (6)$$

where the first step uses the triangle inequality, while the second step uses Proposition 2 with $\Delta = 0$. To bound the first term in (6), let $S \subset \{1, \dots, N\}$ denote the positions of the k largest entries of λ^* . Then, by definition of the best k -term representation,

$$\sigma_k(\mathbf{X}_n) \leq \|\mathbf{X}_n - \mathbf{X}_n^S\|_1 = \sum_{i \in S^c} |X_{n,i}| = \sum_{i \in S^c} X_{n,i}.$$

Therefore,

$$\mathbb{E}\sigma_k(\mathbf{X}_n) \leq \mathbb{E}\left[\sum_{i \in S^c} X_{n,i}\right] = n\tau \sum_{i \in S^c} \lambda_i^* \equiv n\tau\sigma_k(\lambda^*).$$

To bound the second term, we can use concavity of the square root, as well as the fact that each $X_{n,i} \sim \text{Poisson}(n\tau\lambda_i^*)$, to write

$$\mathbb{E}\|\mathbf{X}_n - n\tau\lambda^*\|_1 \leq \sum_{i=1}^N \sqrt{\mathbb{E}(X_{n,i} - n\tau\lambda_i^*)^2} = \sum_{i=1}^N \sqrt{n\tau\lambda_i^*}.$$

Now, it is not hard to show that $\|\widehat{\mathbf{X}}_n^+ - n\tau\lambda^*\|_1 \leq \|\widehat{\mathbf{X}}_n - n\tau\lambda^*\|_1$. Therefore,

$$R\left(\widehat{\lambda}_n^{\text{dir}}, \lambda^*\right) \leq \frac{\mathbb{E}\|\widehat{\mathbf{X}}_n - n\tau\lambda^*\|_1}{n\tau} \leq 4\sigma_k(\lambda^*) + \frac{\|(\lambda^*)^{1/2}\|_1}{\sqrt{n\tau}},$$

which proves the theorem. \blacksquare

¹Khajehnejad et al. [9] have recently proposed the use of perturbed adjacency matrices of expanders to recover nonnegative sparse signals.

B. The penalized MLE approach

The second approach is based on the penalized MLE framework. Assume that we know a good upper bound L_0 on the total average arrival rate $\|\boldsymbol{\lambda}^*\|_1$. Let Λ be a sufficiently large finite set of candidate estimators with $\|\boldsymbol{\lambda}\|_1 \leq L_0$ for all $\boldsymbol{\lambda} \in \Lambda$, and let $\text{pen}(\cdot)$ be a penalty functional satisfying the Kraft inequality over Λ . Given n and τ , let $\Lambda_{n,\tau} \triangleq n\tau d\Lambda$ with the same penalty function.

We can now apply the results of Section III-B with $\mathbf{Z} = \mathbf{Y}_n$ and $\boldsymbol{\theta}^* = n\tau d\boldsymbol{\lambda}^*$. With this notation, define

$$\hat{\boldsymbol{\lambda}}_n^{\text{pMLE}} \triangleq \frac{\hat{\boldsymbol{\theta}}}{n\tau d},$$

where $\hat{\boldsymbol{\theta}}$ is the corresponding pMLE estimator. Then we have the following risk bound:

Theorem 5. Let $c = \frac{\gamma}{k \log(N/k)}$, where $\gamma > 0$ is chosen so that $c \ll 1$. Then

$$\begin{aligned} R\left(\hat{\boldsymbol{\lambda}}_n^{\text{pMLE}}, \boldsymbol{\lambda}^*\right) &= O\left(\sigma_k(\boldsymbol{\lambda}^*) + \sqrt{\gamma k} \log(N/k)\right) \\ &+ O\left(\log(N/k) \sqrt{\frac{k}{\gamma} \min_{\boldsymbol{\lambda} \in \Lambda} \left[\|\boldsymbol{\lambda}^* - \boldsymbol{\lambda}\|_1^2 + \frac{\text{pen}(\boldsymbol{\lambda})}{n\tau}\right]}\right). \end{aligned} \quad (7)$$

We now develop risk bounds under the heavy-tail condition. To this end, let us suppose that $\boldsymbol{\lambda}^*$ is a member of the heavy-tail class $\Sigma_{L_0, \alpha}$ defined in (1). Fix a small positive number δ , such that $L_0/\sqrt{\delta}$ is an integer, and define the set

$$\Lambda \triangleq \left\{ \boldsymbol{\lambda} \in \mathbb{R}_+^N : \|\boldsymbol{\lambda}\|_1 \leq L_0; \lambda_i \in \{m\sqrt{\delta}\}_{m=0}^{L_0/\sqrt{\delta}}, \forall i \right\}$$

These will be our candidate estimators of $\boldsymbol{\lambda}^*$. We can define the penalty function $\text{pen}(\boldsymbol{\lambda}) \asymp \|\boldsymbol{\lambda}\|_0 \log(\delta^{-1})$ so that it satisfies Kraft's inequality. Moreover, if δ is small enough, for any $\boldsymbol{\lambda} \in \Sigma_{\alpha, L_0}$ and any $1 \leq m \leq N$ we will be able to find some $\boldsymbol{\lambda}^{(m)} \in \Lambda$, such that $\|\boldsymbol{\lambda}\|_0 \asymp m$ and

$$\|\boldsymbol{\lambda} - \boldsymbol{\lambda}^{(m)}\|_1^2 \asymp m^{-2\alpha} + m\delta.$$

We will also assume that δ is sufficiently small, so that the penalty term $\frac{m \log(\delta^{-1})}{n\tau}$ dominates the quantization error $m\delta$. Thus, we can bound the minimum over $\boldsymbol{\lambda} \in \Lambda$ in (7) from above by

$$\min_{1 \leq m \leq N} \left[m^{-2\alpha} + \frac{m \log(\delta^{-1})}{n\tau} \right] \approx \left(\frac{\log(\delta^{-1})}{n\tau} \right)^{\frac{2\alpha}{2\alpha+1}}.$$

Using $\tilde{O}(\cdot)$ notation to hide factors that are logarithmic in N and k , we can particularize Theorem 5 to the heavy-tail case:

Theorem 6.

$$\begin{aligned} \sup_{\boldsymbol{\lambda}^* \in \Sigma_{\alpha, L_0}} R\left(\hat{\boldsymbol{\lambda}}_n^{\text{pMLE}}, \boldsymbol{\lambda}^*\right) \\ = O(k^{-\alpha}) + \tilde{O}\left(\sqrt{\gamma k}\right) + \tilde{O}\left(\sqrt{\frac{k}{\gamma}} \left(\frac{\log(\delta^{-1})}{n\tau}\right)^{\frac{\alpha}{2\alpha+1}}\right). \end{aligned}$$

Note that the risk bound here is worse than the benchmark bound of Theorem 4. However, in order to compute the direct

estimator one has to solve a linear program, whereas, as we show next, the pMLE can be approximated very efficiently with proper preprocessing of the observed counts \mathbf{Y}_n based on the structure of $G_{k,N}$.

V. EFFICIENT pMLE APPROXIMATION

In this section we present an efficient algorithm for approximating the pMLE estimate. The algorithm consists of two phases: (1) first, we preprocess \mathbf{Y}_n to isolate a subset A_1 of $A = \{1, \dots, N\}$ which is sufficiently small and is guaranteed to contain the locations of the k largest entries of $\boldsymbol{\lambda}^*$ (the whales); (2) then we construct a set Λ of candidate estimators whose support sets lie in A_1 , together with an appropriate penalty, and perform pMLE over this reduced set.

The success of this approach hinges on the assumption that the magnitude of the smallest whale is much larger compared to the total contribution of the minnows. Specifically, we make the following assumption: Let $S \subset A$ contain the locations of the k largest coordinates of $\boldsymbol{\lambda}^*$. Then we require that

$$\frac{\min_{i \in S} \lambda_i^*}{d} > \sigma_k(\boldsymbol{\lambda}^*). \quad (8)$$

One way to think about (8) is in terms of a signal-to-noise ratio, which must be strictly larger than the left degree d of the underlying expander [recall that $d = O(\log(N/k))$]. We also perturb our expander a bit as follows: choose an integer $k' > 0$ so that

$$\frac{15k'd}{16} \geq kd + 1. \quad (9)$$

Then we replace our original $(k, 1/16)$ -expander with left-degree d with a $(k', 1/16)$ -expander with the same left degree. The resulting procedure, displayed below as Algorithm 1, has the following guarantees:

Algorithm 1 Efficient pMLE approximation algorithm

Input: Measurement vector \mathbf{Y}_n , and the sensing matrix \mathbf{A} .
Output: An approximation $\hat{\boldsymbol{\lambda}}$

Let B_1 consist of the locations of the kd largest elements of \mathbf{Y}_n and let $B_2 = B \setminus B_1$.

Let A_2 contain the set of all variable nodes that have at least one neighbor in B_2 and let $A_1 = A \setminus A_2$.

Construct a candidate set of estimators Λ with support in A_1 and a penalty $\text{pen}(\cdot)$ over Λ .

Output $\arg \min_{\boldsymbol{\lambda} \in \Lambda} [-\log \mathbb{P}_{n\tau \mathbf{A} \boldsymbol{\lambda}}(\mathbf{Y}_n) + 2\text{pen}(\boldsymbol{\lambda})]$.

Theorem 7. The set A_1 constructed by Algorithm 1 has the following properties: (1) $S \subset A_1$; (2) $|A_1| \leq kd$; (3) A_1 can be found in time $O(Nd) = O(N \log(N/k))$.

Proof: (1) If we decompose \mathbf{X}_n as $\mathbf{X}_n^S + \mathbf{e}$, then $\mathbf{Y}_n = \mathbf{A} \mathbf{X}_n^S + \mathbf{A} \mathbf{e}$. Since each column of \mathbf{A} is d -sparse and \mathbf{X}_n^S is k -sparse, $\mathbf{A} \mathbf{X}_n^S$ is kd -sparse. On the other hand, $\mathbf{Y}_n = \mathbf{Y}_n^{B_1} + \mathbf{Y}_n^{B_2}$, where, by construction, $\mathbf{Y}_n^{B_1}$ is the best kd -term approximation of \mathbf{Y}_n . Hence,

$$\|\mathbf{Y}_n^{B_2}\|_1 = \|\mathbf{Y}_n - \mathbf{Y}_n^{B_1}\|_1 \leq \|\mathbf{A} \mathbf{e}\|_1 \leq d \|\mathbf{e}\|_1, \quad (10)$$

where the last inequality follows from the properties of A . Now, since only the nodes in A_2 have neighbors in B_2 ,

$$\|\mathbf{Y}_n^{B_2}\|_1 = \sum_{j \in B_2} \sum_{i \in A_2} A_{ji} X_{n,i} \geq \sum_{i \in A_2} X_{n,i} = \|\mathbf{X}_n^{A_2}\|_1. \quad (11)$$

Combining (10) and (11), we get the bound $\|\mathbf{X}_n^{A_2}\|_1 \leq d\|e\|_1$. Taking expectation of both sides, we obtain

$$\mathbb{E}\|\mathbf{X}_n^{A_2}\|_1 \leq d \cdot \mathbb{E}\|e\|_1 = d \cdot \mathbb{E}\sigma_k(\mathbf{X}_n) \leq dn\tau\sigma_k(\boldsymbol{\lambda}^*), \quad (12)$$

where the last step follows the same reasoning as in the proof of Theorem 4. Now suppose that $S \cap A_2 \neq \emptyset$. Then

$$\mathbb{E}\|\mathbf{X}_n^{A_2}\|_1 \geq \mathbb{E}\|\mathbf{X}_n^{S \cap A_2}\|_1 \geq n\tau \min_{i \in S} \lambda_i^* > dn\tau\sigma_k(\boldsymbol{\lambda}^*),$$

where the last step follows from (8). Since (12) must also hold, we arrive at a contradiction, and therefore $S \subset A_1$.

(2) Suppose, to the contrary, that $|A_1| > kd$. Let $A'_1 \subseteq A_1$ be any subset of size $kd + 1$. Now, Lemma 3.6 in [9] states that, provided $\varepsilon \leq 1 - 1/d$, then every (ℓ, ε) -expander with left degree d is also a $(\ell(1 - \varepsilon)d, 1 - 1/d)$ -expander with left degree d . We apply this result to our $(k', 1/16)$ -expander, where k' satisfies (9), to see that it is also a $(kd + 1, 1 - 1/d)$ -expander. Therefore, for the set A'_1 we must have $|\mathcal{N}(A'_1)| \geq |A'_1| = kd + 1$. On the other hand, $\mathcal{N}(A'_1) \subset B_1$, so $|\mathcal{N}(A'_1)| \leq kd$. This is a contradiction, hence we must have $|A_1| \leq kd$.

(3) Finding the sets B_1 and B_2 can be done in $O(M \log M)$ time by sorting \mathbf{Y}_n . The set A_1 can then be found in time $O(Nd)$, by sequentially eliminating all nodes connected to each node in A_2 . ■

Having identified the set A_1 , we can reduce the pMLE optimization only to those candidates whose support sets lie in A_1 . More precisely, if we originally start with a sufficiently rich class of estimators $\tilde{\Lambda}$, then the new feasible set can be reduced to

$$\Lambda \triangleq \left\{ \boldsymbol{\lambda} \in \tilde{\Lambda} : \text{Supp}(\boldsymbol{\lambda}) \subset A_1 \right\}.$$

Hence, by extracting the set A_1 , we can significantly reduce the complexity of finding the pMLE estimate. If $|\Lambda|$ is small, the optimization can be performed by brute-force search in $O(|\Lambda|)$ time. Otherwise, since $|A_1| \leq kd$, we can use the quantization technique from the preceding section with quantizer resolution $\sqrt{\delta}$ to construct a Λ of size at most $(L_0/\sqrt{\delta})^{kd}$. In this case, we can even assign the uniform penalty

$$\text{pen}(\boldsymbol{\lambda}) = \log |\Lambda| = O(k \log(N/k) \log(\delta^{-1})),$$

which amounts to a vanilla MLE over Λ .

VI. EXPERIMENTAL RESULTS

Here we compare penalized MLE with ℓ_1 -magic [11], a universal ℓ_1 minimization method, and with SSMP [10], an alternative method that employs combinatorial optimization. ℓ_1 -magic and SSMP both compute the “direct” estimator by solving a convex program. The pMLE estimate is computed using Algorithm 1 above and the Sparse Poisson Intensity Reconstruction ALgorithm (SPIRAL) [12] for reconstruction of sparse signals from indirect Poisson measurements.

Figures 1(a) through 2(c) report the result of numerical experiments, where the goal is to identify the k largest entries in the rate vector from the measured data. The set of k largest entries (the whales) is chosen at random. Since a random graph is, with overwhelming probability, an expander graph, each experiment was repeated 30 times².

Given a particular relative sizing of whales and minnows, Figure 1(a) reports values of k where recovery is possible with generic ℓ_1 algorithms (ℓ_1 -magic) but not with SSMP or pMLE. As k increases the first algorithm to fail is SSMP, and the probability of successful recovery falls more sharply than for pMLE. We also report the relative ℓ_1 error ($\|\boldsymbol{\lambda} - \hat{\boldsymbol{\lambda}}_n\|_1 / \|\boldsymbol{\lambda} - \boldsymbol{\lambda}^{(k)}\|_1$) as a function of k in Figure 1(b). However the complexity of ℓ_1 -magic is 2–3 orders of magnitude greater than pMLE, as shown in Figure 1(c).

The effect of increased variability in the size of whales is to reduce the value of k at which pMLE fails. The size of minnows in Figure 2 is the same as in Figure 1, but the variation in the size of whales is determined by an $\mathcal{N}(0, 1)$ Gaussian random variable. Here we see that still Algorithm 1 combined with SPIRAL is two order of magnitudes faster, but the probability of success drops substantially for $k > 80$.


VII. CONCLUSIONS

The compressed sensing algorithms based on Poisson observations and expander-graph sensing matrices provide a useful mechanism for accurately and efficiently estimating a collection of flow rates with relatively few counters. These techniques have the potential to significantly reduce the cost of hardware required for flow rate estimation. While previous approaches assumed packet counts matched the flow rates exactly or that flow rates were i.i.d., the approach in this paper accounts for the Poisson nature of packet counts with relatively mild assumptions about the underlying flow rates (i.e., that only a small fraction of them are large).

The “direct” estimation method (in which first the vector of flow counts is estimated using a linear program, and then the underlying flow rates are estimated using Poisson maximum likelihood) is juxtaposed with an “indirect” method (in which the flow rates are estimated in one pass from the compressive Poisson measurements using penalized likelihood estimation). The direct method can yield smaller error bounds, but this comes at a high computational cost relative to the efficient algorithms associated with the indirect method. These theoretical results are verified in our simulations.


The methods in this paper, along with related results in this area, are designed for settings in which the flow rates are sufficiently stationary, so that they can be accurately estimated in a fixed time window. Future directions include extending these approaches to a more realistic setting in which the flow rates evolve over time. In this case, the time window over which packets should be counted may be relatively short, but this can be mitigated by exploiting estimates of the flow rates in earlier time windows. Another direction for future research

²We observed similar results for experiments with larger number of trials.



(a) Probability of successful support recovery as a function of number of whales k . (b) Relative ℓ_1 error as a function of number of whales k . (c) Average recovery time as a function of number of whales k .

Fig. 1. Performance - Complexity tradeoff for ℓ_1 -magic, SSMP and pMLE. The number of flows $N = 5000$, the number of counters $M = 800$, and the number of updates $n = 40$. There are k whales (peaks with magnitude 1), and the remaining entries are minnows with magnitudes determined by a $\mathcal{N}(0, 10^{-6})$ random variable.



(a) Probability of successful support recovery as a function of number of whales k . (b) Relative ℓ_1 error as a function of number of whales k . (c) Average recovery time as a function of number of whales k .

Fig. 2. Performance - Complexity tradeoff for ℓ_1 -magic, SSMP and pMLE. The number of flows $N = 5000$, the number of counters $M = 800$, and the number of updates $n = 40$. There are k whales (peaks with magnitude determined by a $\mathcal{N}(0, 1)$ random variable), and the remaining entries are minnows with magnitudes determined by a $\mathcal{N}(0, 10^{-6})$ random variable.

will be to tighten the bounds for the indirect method using oracle inequalities based on the Kullback–Leibler divergence.

ACKNOWLEDGMENT

The work of M. Raginsky and R. Willett is supported by NSF CAREER Award No. CCF-06-43947, DARPA Grant No. HR0011-07-1-003, and NSF Grant DMS-08-11062. The work of R. Calderbank and S. Jafarpour is supported in part by NSF under grant DMS 0701226, by ONR under grant N00173-06-1-G006, and by AFOSR under grant FA9550-05-1-0443.

REFERENCES

- [1] W. Fang and L. Peterson, "Inter-AS traffic patterns and their implications," in *Proc. IEEE GLOBECOM*, 1999.
- [2] A. Feldmann, A. Greenberg, C. Lund, N. Reingold, J. Rexford, and F. True, "Deriving traffic demands for operational IP networks: methodology and experience," in *Proc. ACM SIGCOMM*, 2000.
- [3] C. Estan and G. Varghese, "New directions in traffic measurement and accounting: focusing on the elephants, ignoring the mice," *ACM Trans. Computer Sys.*, vol. 21, no. 3, pp. 270–313, 2003.
- [4] Y. Lu, A. Montanari, B. Prabhakar, S. Dharmapurikar, and A. Kabbani, "Counter Braids: a novel counter architecture for per-flow measurement," in *Proc. ACM SIGMETRICS*, 2008.
- [5] S. Jafarpour, R. Willett, M. Raginsky, and R. Calderbank, "Performance bounds for expander-based compressed sensing in the presence of Poisson noise," in *Proc. 43rd Asilomar Conf. on Signals, Systems, and Computers*, 2009, to appear.
- [6] R. Berinde, A. Gilbert, P. Indyk, H. Karloff, and M. Strauss, "Combining geometry and combinatorics: a unified approach to sparse signal recovery," in *Proc. 46th Allerton Conf. on Comm., Control, and Computing*, 2008.
- [7] R. Berinde, P. Indyk, and M. Ruzic, "Practical near-optimal sparse recovery in the ℓ_1 norm," in *Proc. 46th Allerton Conf. on Comm., Control, and Computing*, 2008.
- [8] S. Jafarpour, W. Xu, B. Hassibi, and R. Calderbank, "Efficient and robust compressed sensing using optimized expander graphs," *IEEE Trans. Inform. Theory*, vol. 55, no. 9, pp. 4299–4308, September 2009.
- [9] M. A. Khajehnejad, A. G. Dimakis, W. Xu, and B. Hassibi, "Sparse recovery of positive signals with minimal expansion," submitted, 2009.
- [10] R. Berinde and P. Indyk, "Sequential Sparse Matching Pursuit," in *Proc. 47th Allerton Conf. on Comm., Control, and Computing*, 2009.
- [11] E. Candès and J. Romberg, " ℓ_1 -MAGIC: recovery of sparse signals via convex programming," available at <http://www.acm.caltech.edu/l1magic>, 2005.
- [12] Z. Harmany, R. Marcia, and R. Willett, "Sparse Poisson intensity reconstruction algorithms," in *Proc. IEEE Statist. Signal Proc. Workshop*, 2009.