Estimation of Congestion Price Using Probabilistic Packet Marking *

Micah Adler⋆ Jin-Yi Cai† Jonathan K. Shapiro⋆ Don Towsley⋆

⋆Department of Computer Science
University of Massachusetts at Amherst

†Computer Sciences Department
University of Wisconsin, Madison

July 23, 2002

Abstract

One key component of recent pricing-based congestion control schemes is an algorithm for probabilistically setting the Explicit Congestion Notification bit at routers so that a receiver can estimate the sum of link congestion prices along a path. We consider two such algorithms—a well-known algorithm called Random Early Marking (REM) and a novel algorithm called Self-Normalizing Additive Marking (SAM). We show that if link prices are unbounded, a class of REM-like algorithms are the only ones possible. Unfortunately, REM computes a biased estimate of total price and requires setting a parameter for which no uniformly good choice exists in a network setting. However, we show that if prices can be bounded and therefore normalized, then there is an alternate class of feasible algorithms, of which SAM is representative and furthermore, only the REM-like and SAM-like classes are possible. For properly normalized link prices, SAM returns an optimal price estimate (in terms of mean squared error), outperforming REM even if the REM parameter is chosen optimally. SAM does not require setting a parameter like REM, but does require a router to know its position along the path taken by a packet. We present an implementation of SAM for the Internet that exploits the existing semantics of the time-to-live field in IP to provide the necessary path position information. *Methods: statistics*.

*This work was supported in part by the National Science Foundation (NSF) under grant numbers ANI-9980552, EIA-0080119, CCR-0133664 and CCR-9634665. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

1 Introduction

Recent theoretical advances in optimization-based congestion control have led to the development of protocols in which congestion signals—or *prices* in the common terminology—are computed by links in the network and communicated to sessions. The prices represent Lagrange multipliers in a global optimization problem of maximizing the aggregate user utility in the network subject to a capacity constraint on each link. By knowing only the total price along its own path, each session can independently adapt its rate in a greedy fashion, optimizing its individual utility minus cost. When prices are set correctly by the network, the joint actions of all the users track the globally optimal rate allocation.

In considering the issues surrounding the deployment of such protocols in IP networks, the explicit congestion notification (ECN) bit in the IP header [9] has emerged as a key tool for practical implementations. The importance of ECN is three-fold. First, ECN decouples congestion signals from packet loss—a necessary condition for operating networks with low loss and low delay. Second, an ECN bit already exists in the standard IP header. As we will see, a single bit is sufficient to communicate prices. Thus the debate can focus on how to use the existing bit rather than on how many bits (if any) should be reserved[1]. Third (and most relevant to this paper), it has been demonstrated that routers can encode prices by probabilistically setting the ECN bit in such a way that the end-to-end marking probability encodes the sum of prices along a path. Thus receivers can estimate the total price along a session path, by recording the fraction of marked packets.

Optimization-based congestion control protocols consist of a component running at each link that sets the link's price and marks packets, and an component executed at end-hosts that estimates the total price and sets the transmission rate accordingly. Two classes of protocols have been proposed to date. The first, originally described by Gibbens and Kelly [6], employs an open-loop marking policy at links and adjusts rates iteratively at the end hosts. In the second class [8, 3] end-hosts set rates deterministically, and links combine an iterative algorithm for setting prices with probabilistic packet marking for encoding prices. We concern ourselves with this latter class of protocols where the link price computation and marking scheme are easily separable.

In this work, we assume link prices have converged to steady-state values and focus on the the problem of communicating the sum of fixed link prices along a path by means of packet marking, which we now formalize. Consider a set of links $1, \ldots, n$ forming an end-to-end path from a source to a receiver. Associated

[1] In actuality, two bits in the IP header have been reserved for the purposes of ECN. However, only one of these bits is used to carry congestion signals along the forward path.

with each link i is a non-negative price s_i. Let $z_n = \sum_{i=1}^{n} s_i$ denote the sum of prices along the path. As data packets traversing the path arrive at a receiver, the receiver must determine z_n and provide this quantity as feedback to the sender. We assume that a single bit in the packet header is available for the purpose of communicating this sum, as is the case in the current IP standard. The problem of path price estimation is to design a *marking algorithm*—that is, some strategy for computing the price bit X_i at each link i—to allow the receiver to estimate the total price z_n. To be practically implementable, a marking algorithm must obey the following design constraints: First, the algorithm must be fully distributed with each link making use of locally available information, namely, the price s_i and, if $i > 1$ the bit X_{i-1} computed at the previous step. In some cases, the step index i may also be considered available information. Second, the algorithm should not be required to maintain per-flow state, since this might impose prohibitive storage overheads on routers serving many simultaneous flows. This constraint is clearly satisfied if a link may not retain any memory of how previous packets were marked. Certainly, in this case, one must use randomization, as it is clear that information theoretically no deterministic algorithm can do the job.

In this work, we consider two probabilistic packet marking algorithms—one by Athuraliya and Low [3] called *REM*, and a novel algorithm we have developed called *SAM*.[2] We show that REM is essentially the only method possible when there are no further restrictions on s_i, except $s_i \geq 0$. However, this estimator involves setting a parameter that can be tricky and for which no uniformly good choice exists. Furthermore, the REM estimator is biased. When the additional information of the step index i is known at the i^{th} step and when we assume that each s_i is bounded by some fixed upper bound, say $0 \leq s_i \leq 1$, our SAM method becomes feasible. Moreover, when link prices are restricted to a finite interval, variations of SAM and REM are the only possible methods. We compare REM and SAM in terms of two common metrics. SAM is shown to be optimal in terms of mean squared error (M. S. E.) for uniform a priori distribution of the average price z_n/n. Finally, we present an Internet implementation of SAM, exploiting the existing semantics of the IP time-to-live field to provide the step index i (or an estimate thereof) to each link along a path.

The rest of this paper is organized as follows: In Section 2, we present the REM and SAM algorithms along with a generalized model of all possible marking algorithms. In Section 3 we identify key properties of all feasible protocols and establish the uniqueness of REM for unbounded prices and of REM and SAM when prices are bounded. Sections 4 and 4.2 compare REM and SAM in terms of the tail probability of their price estimates and considers the problem of setting a key parameter in REM. In Section 4.3 we compare REM and SAM in terms of mean squared error and establish the optimality of REM under this criterion. We

[2] We will define the acronyms REM and SAM below.

present an implementation for SAM for the Internet in Section 5.

2 Probabilistic Packet Marking

2.1 Random Early Marking

The Random Early Marking (REM) scheme proposed by Athuraliya and Low [3] is, as far as we are aware, the only existing marking algorithm for price estimation. In REM, the designer selects some base $\phi > 1$. The initial price bit X_0 is set to 0. The i^{th} link, where $i \geq 1$, sets the price bit to 1 with probability $1 - \phi^{-s_i}$. Thinking in terms of conditioning on the incoming price bit X_{i-1}, if $X_{i-1} = 1$ then $X_i = 1$ as well, and if $X_{i-1} = 0$ then with probability ϕ^{-s_i} set $X_i = 0$, and with probability $1 - \phi^{-s_i}$ set $X_i = 1$.

The bit arriving at the receiver is X_n. It is clear that $X_n = 0$ with probability $\prod_{i=1}^{n} \phi^{-s_i} = \phi^{-\sum_{i=1}^{n} s_i}$, and $X_n = 1$ otherwise. Hence the expectation $\mathbf{E}[X_n] = 1 - \phi^{-z_n}$. To estimate the total price z_n the receiver first collects N packets, obtaining N independent samples of the price bit $X_n^{(1)}, X_n^{(2)}, \ldots, X_n^{(N)}$. The receiver then takes $\overline{X} = (\sum_{j=1}^{N} X_n^{(j)})/N$, and estimates z_n to be approximately $-\log_\phi(1 - \overline{X})$.

Note that since $\log_\phi(x)$ is a non-linear function, the expectation $\mathbf{E}[-\log_\phi(1 - \overline{X})]$ is not equal to $-\log_\phi(1 - \mathbf{E}[\overline{X}]) = z_n$. By Jensen's inequality, since log is a strictly convex function, we have

$$\mathbf{E}[-\log(1 - \overline{X})] > z_n.$$

However, even though REM is a biased estimator, as $N \to \infty$ we do have almost everywhere convergence $-\log(1 - \overline{X}) \to z_n$, a.s.. Note also that in REM, the local computation at each step depends only on the local price s_i and the previous bit X_{i-1}, but does not depend on the step index i. Finally, observe that the base ϕ is a parameter that must be chosen by the designer. Athuraliya and Low give no prescription for setting ϕ, but do observe that it should be chosen so as to keep the end-to-end marking probability away from the extreme values of 0 and 1.

2.2 Self-Normalized Additive Marking

Suppose we restrict the range of each link price s_i to be $0 \leq s_i \leq 1$, and suppose the step index i is known for local computation at the i th step. Under these conditions, an alternative scheme is feasible. Again, we set $X_0 = 0$. At each step $i \geq 1$, link i leaves the price bit unchanged ($X_i = X_{i-1}$) with probability $(i-1)/i$. With probability s_i/i the link sets the bit to 1 and sets it to 0 otherwise. The resulting X_n is a 0-1 random variable with $\mathbf{E}[X_n] = \sum_{i=1}^{n} s_i/n$. We thus have an unbiased estimator for z_n/n; we simply collect N i.i.d. samples and compute the average $\overline{X}$. Since the step index is known at each step, the receiver can determine n and thus obtain z_n. We call this scheme *Self-normalizing Additive Marking* (SAM).

2.3 Generalized Protocol Model

The most general one-bit on-line assignment protocol can be described as follows. Without loss of generality, let $X_0 = 0$. Consider the i th step, where $i \geq 1$. The bit X_{i-1} is either 0 or 1. If $X_{i-1} = \varepsilon$ then we assign the bit X_i according to a 0-1 random variable Z_ε, where $\varepsilon = 0,1$. Thus all possible assignments of the bit at step i are defined by two 0-1 random variables Z_0 and Z_1, the distributions of which depend on i and s_i. Let $p_i = \Pr[X_i = 1]$. Then,

$$p_i = p_{i-1} f(i, s_i) + (1 - p_{i-1}) g(i, s_i), \tag{1}$$

where

$$f(i, s_i) = \Pr[Z_0 = 1], \quad \text{and} \quad g(i, s_i) = \Pr[Z_1 = 1].$$

3 Characterization of All Protocols

In this section we provide a characterization of all feasible protocols, such that, for all $(s_1, s_2, \ldots, s_n)$ the estimator converges to z_n, when sample size $N \to \infty$. In Subsection 3.1 we prove that the probability $p_n = \Pr[X_n = 1]$, as a function of $(s_1, s_2, \ldots, s_n)$ must be a function of $\sum_{i=1}^n s_i$, and must be continuous and strictly monotonic in this single argument. In Subsections 3.2 and 3.3 we give a complete analytic characterizations of all feasible protocols for the cases of unbounded and bounded link prices.

3.1 Strict monotonicity as a function of $\sum_{i=1}^n s_i$

No matter what it does at each step i, a marking algorithm ultimately produces a 0-1 random variable X_n. Thus looking at the problem externally any algorithm can be characterized by the probability that $X_n = 0$. This probability must be a function of $s_1, s_2, \ldots, s_n$; we will call it $p_n(s_1, s_2, \ldots, s_n)$.

Theorem 1. *If for all $(s_1, s_2, \ldots, s_n)$ the estimator converges to $z_n = \sum_{i=1}^n s_i$ asymptotically, as the number of sample points $N \to \infty$, p_n must be a function of the sum z_n, and be continuous and strictly monotonic in its single argument z_n.*

Proof. Let us fix the length n of the path. Suppose we take i.i.d. samples $Y_1, Y_2, \ldots, Y_N$, each a 0-1 random variable with $\Pr[Y_k = 1] = p_n(s_1, s_2, \ldots, s_n)$. Given the samples, the only quantity one can hope to infer is this value $p_n(s_1, s_2, \ldots, s_n)$, since it determines the distribution. In other words, if $(s_1, s_2, \ldots, s_n)$ and $(s'_1, s'_2, \ldots, s'_n)$ are such that $p_n(s_1, s_2, \ldots, s_n) = p_n(s'_1, s'_2, \ldots, s'_n)$, then for any $N, Y_1, Y_2, \ldots, Y_N$ are identically distributed for $(s_1, s_2, \ldots, s_n)$ and for $(s'_1, s'_2, \ldots, s'_n)$. In particular, if $\sum_{i=1}^n s_i \neq \sum_{i=1}^n s'_i$, then intuitively there would be no way to distinguish $\sum_{i=1}^n s_i$ from $\sum_{i=1}^n s'_i$.

The mean of samples $\overline{Y} = \sum_{k=1}^N Y_k/N$ is a *sufficient statistic* for $p_n(s_1, s_2, \ldots, s_n)$, thus $\overline{Y}$ contains all the information about $p_n(s_1, s_2, \ldots, s_n)$, which is the only quantity that we can hope to infer. Therefore we may assume that any inference rule from the sample is a function of the mean $\overline{Y}$, call it $G(\overline{Y})$. We want to show that, if any estimate $G(\overline{Y})$ for the sum $\sum_{i=1}^n s_i$ is to converge to the right value $\sum_{i=1}^n s_i$, as $N \to \infty$, then,
(I) $\sum_{i=1}^n s_i \neq \sum_{i=1}^n s'_i$ implies that $p_n(s_1, s_2, \ldots, s_n) \neq p_n(s'_1, s'_2, \ldots, s'_n)$; and
(II) $p_n(s_1, s_2, \ldots, s_n)$ must be a function of $\sum_{i=1}^n s_i$, i.e., $\sum_{i=1}^n s_i = \sum_{i=1}^n s'_i$ implies that $p_n(s_1, s_2, \ldots, s_n) = p_n(s'_1, s'_2, \ldots, s'_n)$.

By the strong Law of Large Numbers, we have an almost everywhere convergence, $\overline{Y} \to \mathbf{E}[Y_1] = p_n(s_1, \ldots, s_n), a.s.$. Since G is the result of a computable process, it is continuous, and thus we have $G(\overline{Y}) \to G(p_n(s_1, \ldots, s_n)), a.s.$. Hence the value $G(p_n(s_1, \ldots, s_n))$ must determine the sum $\sum_{j=1}^n s_j$ for all $(s_1, \ldots, s_n)$. (Note that the almost everywhere convergence refers to the measure space on Y_k's, for *any* fixed $(s_1, \ldots, s_n)$.) If $\sum_{j=1}^n s_j \neq \sum_{j=1}^n s'_j$, yet $p_n(s_1, \ldots, s_n) = p_n(s'_1, \ldots, s'_n)$, then for either $(s_1, \ldots, s_n)$ or $(s'_1, \ldots, s'_n)$ the convergence to $G(p_n(s_1, \ldots, s_n))$ is wrong. (I) is proved.

To show (II), we note that for on-line protocols, for each $1 \le i \le n$, $p_i(s_1, \ldots, s_i)$ represents the whole problem for the instance of a path of length i, and thus we may inductively assume that p_{i-1} is a function of $\sum_{j=1}^{i-1} s_j$. Denote by $s = \sum_{j=1}^{i-1} s_j$ and $t = s_i$. We want show that p_i is a function of $s+t$. By (1), since all the functions involved are the result of a computable process, and by the induction hypothesis for p_{i-1}, we see that p_i is a continuous function of s and t; we write it as $q_i(s,t)$ for the moment.

For a contradiction, suppose that for some $(s,t) \neq (s',t')$, $s+t = s'+t'$, yet we have $q_i(s,t) \neq q_i(s',t')$. Let $v = (q_i(s,t) + q_i(s',t'))/2$. By the intermediate value theorem, there must be a point (s^*, t^*) on the line segment between (s,t) and (s',t') but distinct from the two end points, such that $q_i(s^*, t^*) = v$. On the other hand, consider the rectilinear path from (s,t) to (s',t'), first along the s-axis, then along the t-axis. Again by the intermediate value theorem, there must be a point $(\overline{s}, \overline{t})$, distinct from the two end points (s,t) and (s',t'), such that $q_i(\overline{s}, \overline{t}) = v$. (To the careful reader we note that, in the case where we restricted the range of $s_i \in [0,1]$, all these points are in the domain of definition of q_i, namely $[0, i-1] \times [0,1]$.) However, geometrically it is clear that $\bar{s}+\bar{t} \neq s^* + t^*$. Thus the function value $v = q_i(s^*, t^*) = q_i(\bar{s}, \bar{t})$ does not uniquely determine the sum $s^* + t^* \neq \bar{s} + \bar{t}$. This contradiction to (I) proves that $q_i(s,t)$ is indeed a function of $s+t$. The induction is complete, and (II) is proved.

From now on we will denote this function as $p_n(\sum_{j=1}^{n} s_j)$ in place of $p_n(s_1, \ldots, s_n)$.

We next show that $p_n(\cdot)$ is strictly monotonic in its single argument. By (I) $p_n(\cdot)$ is 1-1. And as we noted before it is continuous. Hence by the intermediate value theorem again, it is strictly monotonic (either increasing or decreasing). Similarly we can conclude that G is also strictly monotonic (in the range of p_n). □

3.2 Solutions of functional equations over $[0, \infty)$

Now we fix $i \geq 1$. To simplify expressions, define $s = \sum_{j=1}^{i} s_j$ and $t = s_{i+1}$. Consider the functional equation transferring the probability from step i to $i+1$:

$$h(s+t) = p(s)f(t) + (1 - p(s))g(t). \tag{2}$$

Note that implicitly, all of these functions can depend on i, which is fixed.

Theorem 2. *Suppose h, p, f and g are real valued functions defined on $[0, \infty)$, and satisfy the functional equation (2) for all $s, t \geq 0$. Assume furthermore that p is strictly monotonic and bounded, h is non-constant, and f and g are continuous. Then there exists a constant $0 < \psi < 1$, such that each function h, p, f and g is of the form $c + c'\psi^x$ for some constants c and c'. More precisely, there exist constants $0 < \psi < 1$, and a, b, c and d, such that*

$$\begin{aligned} p(x) &= a + b\psi^x \\ f(x) &= c + (1-a)d\psi^x \\ g(x) &= c - ad\psi^x \\ h(x) &= c + bd\psi^x \end{aligned}$$

Remark: In this theorem and the one that follows, we explicitly assume that f and g are continuous. In fact, one may make this assumption without loss of generality since f and g are computable functions and, in the strict sense of computability, all computable real functions are continuous (See [13], Theorem 4.3.1, page 108).

Proof. From (2) if we take the difference at (s,t) and (s',t), we get

$$h(s+t) - h(s'+t) \quad = \quad (p(s) - p(s')) \cdot (f(t) - g(t)). \tag{3}$$

In (3) we set $t = 0$, then

$$h(s) - h(s') \quad = \quad (p(s) - p(s')) \cdot (f(0) - g(0)). \tag{4}$$

If $f(0) = g(0)$, then $h(s)$ is a constant function identically, which is a contradiction. Hence $f(0) \neq g(0)$. Denote the non-zero constant $f(0) - g(0)$ by d, and let $v(t) = \frac{f(t)-g(t)}{d}$, then (3) and (4) imply that

$$(p(s) - p(s')) \cdot (f(t) - g(t)) = h(s+t) - h(s'+t) = (p(s+t) - p(s'+t)) \cdot d,$$

hence

$$p(s+t) - p(s'+t) = (p(s) - p(s')) \cdot v(t). \tag{5}$$

Since p is bounded and monotonic, the limit $\lim_{x\to\infty} p(x)$ exists, which we will denote by a. Let $q(x) = p(x) - a$, and we take the limit $s' \to \infty$ in (5), then

$$q(s+t) = q(s)\ v(t). \tag{6}$$

In particular

$$q(t) = q(0)v(t). \tag{7}$$

Since p is strictly monotonic, $q(0) \neq 0$. Divide by $q(0)$ in (6) gives

$$v(s+t) = v(s) \cdot v(t). \tag{8}$$

We note that p is monotonic and non-constant, so is q, and by (7) and the fact that $q(0) \neq 0$, so is v. Hence Lemma 1 applies to v, so that there exists a constant $\psi > 0$, such that $v(x) = \psi^x, \forall x \geq 0$. And in fact, v is strictly monotonic and bounded, it follows that $\psi < 1$.

Let $b = q(0)$. Then

$$\begin{aligned} p(t) \quad &= \quad a + q(x) && (9)\\ &= \quad a + bv(t) && (10)\\ &= \quad a + b\psi^x. && (11) \end{aligned}$$

We complete the proof by deriving similar equations for the other functions. In (2) let $t = 0$ and $s \to \infty$, we see that the limit $\lim_{s\to\infty} h(s)$ exists. Denote this limit by c. Then by taking the limit $s \to \infty$ in (2) for any fixed t we get

$$af(t) + (1-a)g(t) = c. \tag{12}$$

Recall the definition of $v(t)$, the left-hand-side of (12) is

$$g(t) + a(f(t) - g(t)) = g(t) + adv(t). \tag{13}$$

Hence g is of the form

$$g(x) = c - ad\psi^x. \tag{14}$$

It follows then that

$$\begin{aligned} f(x) &= g(x) + (f(x) - g(x)) & (15) \\ &= g(x) + dv(x) & (16) \\ &= c + (1-c)d\psi^x. & (17) \end{aligned}$$

Finally in (4) we take the limit $s' \to \infty$, we get

$$h(x) = c + bd\psi^x. \tag{18}$$

□

The above proof uses the following well known lemma. For the sake of completeness we include a proof here.

Lemma 1. *Let $v(x)$ be a monotonic function defined on $[0,\infty)$, and is not identically zero, and satisfies the following functional equation*

$$v(s+t) = v(s)v(t) \tag{19}$$

for all $s,t \geq 0$, then there exists some constant $\psi > 0$, such that

$$v(x) = \psi^x. \tag{20}$$

Proof. Since v is not identically zero, and $v(t) = v(0)v(t)$, we have $v(0) \neq 0$. Then $v(0) = v(0)^2 \Longrightarrow v(0) = 1$.

Let $\psi = v(1)$, then by a simple induction $v(n) = \psi^n$, for all non-negative integers $n \geq 0$. Being monotonic, it follows that $\psi > 0$. Also by monotonicity, since all $v(n) > 0$, we must have $v(t) > 0$ for all real $t \geq 0$.

Now for any positive rational number $r = \frac{n}{m}$,

$$\psi^m = v(m) = v(nr) = (v(r))^n,$$

and $v(r)$ being positive, $v(r) = \psi^r$.

Finally by monotonicity, it follows that for all real $x \in [0, \infty)$,

$$v(x) = \psi^x.$$

□

In the following we will write $\phi = \psi^{-1}$, thus $\phi > 1$.

In order to be a probability and strictly monotonic, the constants a and b in the function p must also satisfy

$$0 \leq a, a+b \leq 1, \quad \text{and} \quad b \neq 0. \tag{21}$$

We note that given this complete characterization, it is easy to see that REM corresponds to the choice of constants $a = 1$ and $b = -1$ for $p(s)$. There is a dual choice of $a = 0$ and $b = 1$, which we will call REM*.

For all parameters (technically for all *computable* parameters) a and b satisfying (21), the function $p(\cdot)$ is realizable as the probability function of some one-bit on-line protocol as defined. In fact, if $n = 1$, we can just take X_1 such that $\Pr[X_1 = 1] = a + b\phi^{-s_1}$. This is legitimate since $a + b\phi^{-s_1}$ is always between a and $a+b$, and thus $0 \leq a + b\phi^{-s_1} \leq 1$, for all $s_1 \geq 0$. For $n > 1$, inductively we can assume $\Pr[X_{n-1} = 1] = a + b\phi^{-(z_{n-1})}$ as p_{n-1}, then we let $f(s_n) = a + (1-a)\phi^{-s_n}$ and $g(s_n) = a - a\phi^{-s_n}$. Again it is easy to see that both $0 \leq f(s_n), g(s_n) \leq 1$, for all $s_n \geq 0$. It follows that $p_n(z_n) = h(z_n) = a + b\phi^{-z_n}$. We will call all these feasible protocols REM-like.

For any fixed ϕ, the question of what choices of a and b are the best remains unanswered. We will subsequently show that, in terms of M.S.E., REM and REM* are the best choices of all these REM-like protocols with the same ϕ.

3.3 Solutions of functional equations over $[0, 1]$

In the previous subsection, we gave a complete characterization of probability functions of all admissible one-bit on-line protocols as defined before, provided that $s_i \in [0, \infty)$ $i = 1, \ldots, n$. When we have the further restriction that $s_i \in [0,1]$ $i = 1,\dots,n$, there are other solutions to the functional equations, which we turn to in this subsection.

Fix $i \geq 1$ and consider again the functional equation (2), except now f and g are only defined for $x \in [0,1]$, and p is defined for $x \in [0,i]$ and h is defined for $x \in [0,i+1]$. Note that implicitly, all these functions can depend on i, which is fixed.

Theorem 3. *Suppose h, p, f and g are real valued functions defined on $[0,i+1]$, $[0,i]$, $[0,1]$ and $[0,1]$, respectively, and satisfy the functional equation (2), for all $s \in [0,i]$ and $t \in [0,1]$. Assume furthermore that p is strictly monotonic and bounded, h is non-constant, and f and g are continuous. Then there are just two classes of solutions:*

1. *There exists a constant $\psi > 0$, $\psi \neq 1$, such that each function h, p, f and g is of the form $a + b\psi^x$ for some constants a and b. Or*

2. *Each function h, p, f and g is an affine linear function of x of the form $a + bx$ for some constants a and b.*

Here note that all these constants may depend on i.

Proof. From (2) if we take the difference at (s,t) and (s',t), where $0 \leq s, s' \leq i$ and $0 \leq t \leq 1$, we get

$$h(s+t) - h(s'+t) \quad = \quad (p(s) - p(s')) \cdot (f(t) - g(t)). \tag{22}$$

In (22) we set $t = 0$, then

$$h(s) - h(s') \quad = \quad (p(s) - p(s')) \cdot (f(0) - g(0)). \tag{23}$$

If $f(0) = g(0)$, then $h(s)$ is a constant function identically, which is a contradiction. Hence $f(0) \neq g(0)$. Let

$$v(t) = \frac{f(t) - g(t)}{f(0) - g(0)}, \tag{24}$$

then (22) and (23) imply that

$$(p(s) - p(s')) \cdot (f(t) - g(t)) = h(s+t) - h(s'+t) = (p(s+t) - p(s'+t)) \cdot (f(0) - g(0)),$$

provided that $0 \leq s+t, s'+t \leq i$, so that $p(\cdot)$ is defined there. Hence

$$p(s+t) - p(s'+t) = (p(s) - p(s')) \cdot v(t), \tag{25}$$

for $0 \leq t \leq 1, 0 \leq s, s', s+t, s'+t \leq i$.

Suppose in addition $0 \leq t' \leq 1$, $0 \leq t+t' \leq 1$, and $0 \leq s+t+t', s'+t+t' \leq i$, then we can apply (25) twice to get

$$\begin{aligned} p(s+t+t') - p(s'+t+t') &= (p(s)-p(s')) \cdot v(t+t') && (26) \\ &= (p(s+t)-p(s'+t)) \cdot v(t') && (27) \\ &= (p(s)-p(s')) \cdot v(t) \cdot v(t'). && (28) \end{aligned}$$

Since p is strictly monotonic, by taking sufficiently small but unequal $s \neq s'$, we conclude that for all $0 \leq t, t' < 1$, and $0 \leq t+t' < 1$,

$$v(t+t') = v(t) \cdot v(t'). \tag{29}$$

By continuity, this holds for all $0 \leq t, t' \leq 1$ and $0 \leq t+t' \leq 1$.

Since p is strictly monotonic, it follows from (25) that v is positive. By taking $V(t) = \log v(t)$, then V is additive

$$V(t+t') = V(t) + V(t'). \tag{30}$$

Hence, for $0 \leq x \leq 1$, and any integer $m \geq 1$, $V(x) = mV(x/m)$. Lemma 2 applies, and we conclude that there exists some constant $\psi > 0$, such that

$$v(x) = \psi^x. \tag{31}$$

If $\psi \neq 1$, then this leads to the first class of solutions as discussed previously. (The case $\psi > 1$ reduces to the case $\psi < 1$ by the reversal transformation of $x \mapsto 1-x$. Of course such reversal transformation is only feasible for finite intervals.)

If $\psi = 1$, then $v(t)$ is identically 1, and $f(t) - g(t)$ is a constant independent of t. In this case, repeated applications of (25) gives, for any integer $m \geq 1$, and any $0 \leq s \leq i$,

$$\begin{aligned} p(s) - p(0) &= \sum_{k=1}^{m} \left[p(\frac{k}{m}s) - p(\frac{k-1}{m}s) \right] && (32) \\ &= m \left[p(\frac{s}{m}) - p(0) \right]. && (33) \end{aligned}$$

Let $a = p(0)$ and $q(x) = [p(ix) - p(0)]/i$, then for all $0 \leq x \leq 1$,

$$q(x) = mq(x/m). \tag{34}$$

Lemma 2 applies again, and we conclude that $q(x) = bx$ for some constant $b = q(1)$. We can revert back to $p(s)$, for $0 \leq s \leq i$, and get

$$p(s) = a + bs. \tag{35}$$

It follows easily that h, f and g also take the form of affine linear functions. We omit the details. □

The following Lemma is also well known and is essentially the same as Lemma 1, except it is for a finite interval. For completeness we include a proof here.

Lemma 2. *If F is a continuous function defined on $[0,1]$, and satisfies the functional equation*

$$F(x) = mF\left(\frac{x}{m}\right) \tag{36}$$

for all $x \in [0,1]$ and integer $m \geq 1$, then there exists a constant c, such that $F(x) = cx$.

Proof. For any integers $1 \leq k, m \leq \ell$,

$$F(m/\ell) = mF(1/\ell) \quad \text{and} \quad F(k/\ell) = kF(1/\ell). \tag{37}$$

Hence,

$$F(m/\ell) = \frac{m}{k} F(k/\ell). \tag{38}$$

Then for any $0 \leq x \leq 1$, let $m/\ell \to x$, and $k/\ell \to 1$, while maintaining $1 \leq k, m \leq \ell$ at all time, (this is clearly possible to do), then $m/k \to x$ as well. Then by continuity of F,

$$F(x) = F(1)x. \tag{39}$$

□

The first class of solutions is essentially the exponential family discussed above.[3] As we showed before, if we want the functional equation to hold over functions defined over $[0,\infty)$, then there is only this first class of solutions with $0 < \psi < 1$; the second class of solutions is not possible. What makes it possible here is the restriction of the functional equation to a finite interval.

For a path of length n, denote by $\theta = \sum_{i=1}^{n} s_i/n$. Any admissible protocol of the second class must have $\Pr[X_i = 0] = a + b\theta$ for some constants a and b (which may depend on i.) Since $a + b\theta$ is a probability, $0 \leq a, a + b \leq 1$. SAM simply takes $a = 0$ and $b = 1$ and is thus an unbiased estimator of θ. There is a dual

[3]For finite interval $[0,i]$, $\psi > 1$ is possible; but it is easily transformed to the case with $\psi < 1$, by reversing the map $x \mapsto i - x$. For the infinite interval $[0,\infty)$, $\psi > 1$ is impossible, and we get $0 < \psi < 1$.

choice that corresponds to $a = 1$ and $b = -1$. In Section 4.3, we show that SAM and its dual, are uniquely optimal with respect to the criterion of Mean Square Error, among all solutions of the second class.

It is easy to verify that all feasible choices of (a,b) can be realized in a one-bit on-line protocol when each $s_j \in [0,1]$, and if at step i we know the index i. Assume we have the probability function $(\sum_{j=1}^{i-1} s_j)/(i-1)$ (as in SAM) for the $i-1$ step. Then let $f(s_i) = a + \frac{i-1}{i}b + \frac{b}{i}s_i$, and $g(s_i) = a + \frac{b}{i}s_i$. These choices are both legitimate since both $a + \frac{i-1}{i}b = \frac{1}{i}a + \frac{i-1}{i}(a+b)$, and $a + \frac{1}{i}b = \frac{i-1}{i}a + \frac{1}{i}(a+b)$, are convex combinations of a and $a+b$, and therefore, since both $a, a+b \in [0,1]$, it follows that all four numbers $f(0) = a + \frac{i-1}{i}b, f(1) = a+b, g(0) = a, g(1) = a + \frac{1}{i}b \in [0,1]$.

3.3.1 An extended REM family

When we considered the REM-like protocols over $[0,\infty)$, with probability function $p(s) = a + b\varphi^{-s}$ as in Theorem 2, the parameters a and b must satisfy $0 \le a, a+b \le 1$. However, if we restrict REM-like protocols to a finite interval, the solutions to our functional equation (2) may have more general coefficients a, b from Theorem 3.

Indeed we claim that the following "stretched" version of REM is feasible: At step i, let

$$f_i(s_i) = \frac{1-\varphi^{-(i-1)-s_i}}{1-\varphi^{-i}}, \quad \text{and} \quad g_i(s_i) = \frac{1-\varphi^{-s_i}}{1-\varphi^{-i}}.$$

Note that $0 \le f_i(s_i), g_i(s_i) \le 1$ for $0 \le s_i \le 1$, for all $i \ge 1$ and $\varphi > 1$. Then from the recurrence relation (2) and induction on n, it is easy to verify that this choice of f_i and g_i achieves

$$p_n(s) = \frac{1-\varphi^{-s}}{1-\varphi^{-n}},$$

where $s = \sum_{i=1}^{n} s_i$. If we denote $\theta = s/n$, and the above function as $F(\theta)$, then

$$F(\theta) = \frac{1-\psi^{-\theta}}{1-\psi^{-1}}, \tag{40}$$

where $\psi = \varphi^n$. Thus this is REM with a "stretch" so that it spans the total spectrum of $[0,1]$, for every base ψ. It is natural to expect this "stretched" version of REM to work better than REM when we know in advance that $s_i \in [0,1]$.

4 Evaluation

4.1 Comparison of Tail Probabilities

We next consider the receiver's problem of estimating the price of a path using either SAM or REM for marking packets. Suppose the receiver collects N packets, giving it N samples of the price bit. Let $B =$

$\sum_{j=1}^{N} X_n^{(j)}$ be the number of samples for which the price bit is set. The receiver can then estimate the path price by estimating the end-to-end marking probability. Let p denote the true end-to-end marking probability. The estimated marking probability is $\hat{p} = B/N$. For now, we assume the path length n is known to the receiver.

To simplify expressions, we will drop the superscript for the path price notation, thus

$$z = \sum_{i=1}^{n} s_i.$$

Let $\hat{z}$ be the price estimate provided by either algorithm. For SAM, we have

$$\hat{z} = \hat{p}n, \tag{41}$$

whereas for REM,

$$\hat{z} = -\log(1-\hat{p}). \tag{42}$$

The true path price z can also be expressed using equations (41) and (42) by substituting the true marking probability p for the estimated probability $\hat{p}$ on the right-hand side. Informally, we can think of the efficiency of a marking algorithm as the number of samples required to estimate the true price with high confidence. This notion is captured in the metric of *error probability*, denoted $err(\varepsilon)$ and defined as the probability that the price estimate falls outside of some range about the true price, where the range is determined by a parameter ε. Formally,

$$err(\varepsilon) = 1 - \Pr[(1-\varepsilon)z \leq \hat{z} \leq (1+\varepsilon)z] \tag{43}$$

It is natural to compare REM and SAM on the basis of efficiency, and the error probability provides one tractable metric for doing so.

Since both algorithms use the estimated marking probability $\hat{p}$ to estimate the price, it is also useful to relate the acceptable variation in $\hat{z}$ (as defined by the parameter ε) to an equivalent variation in $\hat{p}$,

$$err(\varepsilon) = 1 - \Pr[(1-\delta^-)p \leq \hat{p} \leq (1+\delta^+)p], \tag{44}$$

where δ^- and δ^+ depend on the value of ε, the marking probability p, and the choice of marking algorithm. As will become clear below, we must distinguish between the values of δ for the upper and lower tail since these may not be equal.

Noting that the price estimates (41) and (42) are increasing in $\hat{p}$, we may conclude that

$$\hat{z} = (1-\varepsilon)z \quad \Leftrightarrow \quad \hat{p} = (1-\delta^-)p \tag{45}$$

$$\hat{z} = (1+\varepsilon)z \quad \Leftrightarrow \quad \hat{p} = (1+\delta^+)p \tag{46}$$

Taking SAM as an example, let us fix a value of ε and require that

$$\Pr[\hat{z} \leq (1-\varepsilon)\,z] = \Pr[\hat{p} \leq (1-\delta^-)\,p] \tag{47}$$

$$\Pr[\hat{z} \geq (1+\varepsilon)\,z] = \Pr[\hat{p} \geq (1+\delta^+)\,p]. \tag{48}$$

We can now solve for the values of δ^- and δ^+ that make this requirement true. Using equation (41) and observation (45), we have

$$\hat{z} = (1-\varepsilon)\,z = (1-\delta^-)\,pN. \tag{49}$$

Using the fact that $z = pN$ we can rewrite (49)

$$(1-\delta^-)\,pN = (1-\varepsilon)\,pN$$

Thus, for a fixed $\varepsilon \in (0,1)$ we may set $\delta^- = \varepsilon$ for the case of SAM. Essentially identical reasoning establishes that $\delta^+ = \varepsilon$.

In the case of REM, a more complex relationship holds. Using equation (42) and observation (45), we can write

$$\hat{z} = (1-\varepsilon)\,z = -\log_\phi(1-(1-\delta^-)(1-\phi^{-z})),$$

Solving for δ^-, we have

$$\delta^- = \frac{\phi^{-(1-\varepsilon)z} - \phi^{-z}}{1-\phi^{-z}} \tag{50}$$

Applying the same reasoning for the upper tail gives

$$\delta^+ = \frac{\phi^{-z} - \phi^{-(1+\varepsilon)z}}{1-\phi^{-z}} \tag{51}$$

To facilitate the comparison of REM and SAM, we adopt a network model in which link prices are independent random variables uniformly distributed on the interval $[0,1]$. This model is perhaps not representative of the true distribution of congestion prices in a real network, where a relatively small fraction of links are highly congested and the majority of links are uncongested. The benefit of using this model is its simplicity; the expected path price $E[s]$ is proportional to path length.

To gain some understanding of how the error probability behaves as path length increases under our simple network model, we generated a set of n_{max} links with prices uniformly distributed on $[0,1]$. We then compute the end-to-end marking probability for a path consisting n links where $n = 1,2,\cdots,n_{max}$ for both REM and SAM. Since we expect the performance of REM to depend on the choice of parameter ϕ, we consider two different values, $\phi = 1.01$ and $\phi = e$, as a baseline for comparison. Finally, for fixed values for the error tolerance ε and the number of samples N, we can compute the resulting δ^+ and δ^-.

Since we know the true marking probability (given a set of link prices), we can compute the error probability (43) exactly. Treating each packet sent as a Bernoulli trial with probability of heads p, we have

$$\Pr[\hat{z} > (1+\varepsilon)z] = \sum_{B=\lceil n(1+\delta^+)p \rceil}^{n} r(n,B,p) \tag{52}$$

$$\Pr[\hat{z} < (1-\varepsilon)z] = \sum_{B=0}^{\lfloor n(1-\delta^-)p \rfloor} r(n,B,p), \tag{53}$$

where $r(n,B,p) = \binom{n}{b} p^B (1-p)^{(n-B)}$ is the probability mass function for a Bernoulli random variable. The error probability defined in (43) can also be written

$$err(\varepsilon) = \Pr\{\hat{z} \notin [(1-\varepsilon)s, (1+\varepsilon)s]\},$$

from which it is easily seen that $err(\varepsilon)$ is the sum of equations (52) and (53).

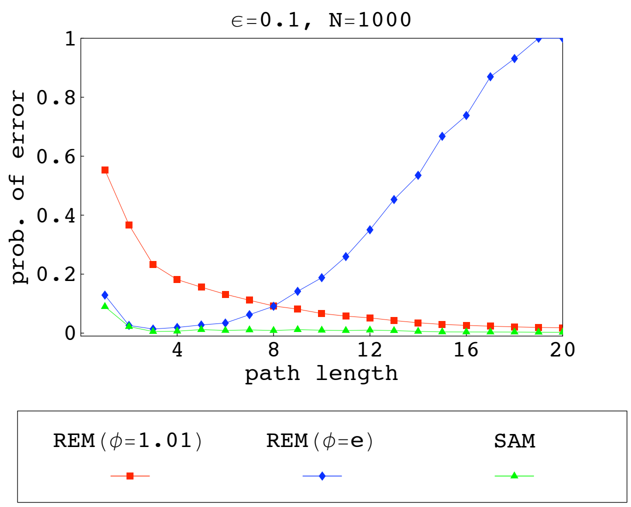

Figure 1: Error probability as a function of path length for SAM and for two parameterizations of REM. We observe that SAM yields and error probability that is largely independent of path length and that this error probability is matched by REM only at specific path lengths, which depend on the value of the parameter ϕ.

Figure 1 shows the dependence of error probability on path length for two parameterizations of REM ($\phi = 1.01$ and $\phi = e$) and for SAM. For this plot, we have fixed the number of samples at 1000 and the error tolerance parameter ε at 0.1. The data plotted are averaged over 10 independently generated sets of link prices. We observe several interesting features in these results. First, the error probability of the SAM price estimate is unaffected by path length. Second, the REM error probability does depend on path length, with the two different parameterizations yielding error probabilities comparable to SAM at different path lengths. This result implies that the appropriate choice of ϕ may be path dependent. We note also that the error probability for $\phi = e$ can become 1 at long path lengths. This situation corresponds to an extremely high marking probability for which no unmarked packets were seen within 1000 samples.

These results suggest that SAM is well-suited for marking in a network environment where sessions see varying path lengths and path prices. We have seen that REM can also perform comparably to SAM but that its performance depends on the choice of parameter ϕ. To compare the two algorithms fairly, we must investigate the issue of parameter setting in REM more thoroughly.

4.2 The Effect of Parameter ϕ in REM

Figure 1 shows that the REM algorithm with $\phi = e$ performs quite well at short path lengths but performs poorly for longer paths, whereas $\phi = 1.1$ performs well on longer paths but poorly on short paths. This result suggests that a version of REM in which ϕ is selected according to the path length[4] might have performance comparable to SAM.

In the case of either REM or SAM, one must collect a significant number of packets in order make a close estimate of path price. The number of marked packets B is a Bernoulli random variable. However, since the number of samples is large, we may approximate B as a normally distributed random variable with mean $\mu = Np$, variance $\sigma^2 = Np(1-p)$ and CDF $F(x;\mu,\sigma)$.[5] Under this approximation, the error probability can be written

$$err(\varepsilon) \approx 1 - \int_{(1-\delta^-)pn}^{(1+\delta^+)pn} dF(x;\mu(p),\sigma(p)), \tag{54}$$

where we have made explicit the functional dependence on p, the end-to-end marking probability.

Recall that the REM marking probability depends on the total path price s and the parameter ϕ. The limits of integration in (54) depend on p as does the pdf dN. Thus, the error probability is a continuous function of both s and ϕ.

Figure 2 shows the error probability for REM as a function of ϕ for values of the total path price s ranging from 0.1 to 10. Path prices ranging over three orders of magnitude is well within the realm of possibility for

[4] Recall that in the network model underlying Fig. 1, path price is proportional to path length.

[5] One rule of thumb for evaluating the validity of the normal approximation to the binomial is, for a binomial distribution with parameters N and p, that $Np \geq 5$ and $Np(1-p) \geq 5$ [11]. These conditions are satisfied in our model for $N > 200$ in the case of SAM and REM with $\phi = e$. For REM with $\phi = 1.1$, the conditions are satisfied for $N > 200$ for all path lengths greater than a single hop.

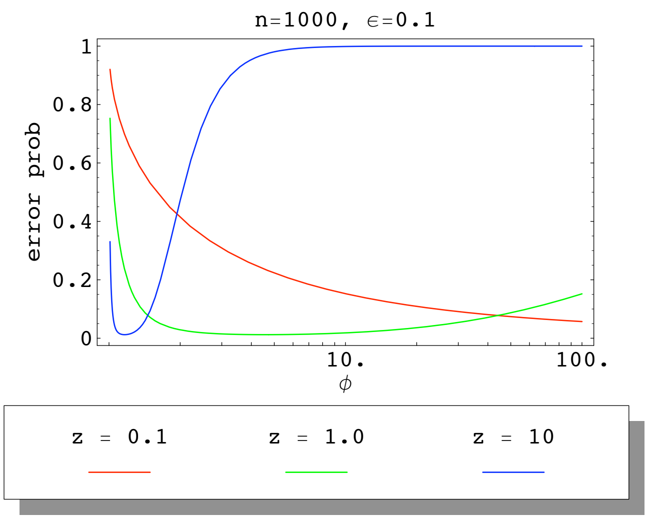

Figure 2: Error probability as a function of REM parameter ϕ for three values of total path cost $S = s^{n}$. Note that the x-axis is on a logarithmic scale.

REM due to the varying number of hops and levels of congestion seen along different paths, and due to the fact that link price is not actually constrained to $[0,1]$ in REM. The plots shown in the figure suggest that it is impossible to fix a single value for ϕ that will yield a low error probability for all paths. Rather, it appears that the appropriate choice of ϕ is indeed path dependent. In particular, for a given path with path price s, there is an optimal parameterization $\phi = \phi^*(z)$ for which error probability is minimized.

Although ϕ^* is path dependent, it is still necessary for each router along a path to use the same value when marking an individual flow. Incorporating such "path optimization" into REM would certainly add complexity to the implementation. For example, setting ϕ^* on a per-flow basis would require either per-flow storage at routers or including the value of ϕ^* in each packet header. More fundamentally, the value of ϕ^* depends on the end-to-end price, which is precisely the quantity to be estimated. Thus it would be necessary to jointly refine estimates of the price and ϕ^* as the protocol proceeds (and demonstrate the convergence of such an approach).

We next consider how well REM can perform in the best possible circumstances—if the path price (and, hence, ϕ^*) is known in advance. To address this question, we generated a sequence of links with prices uniformly distributed on $[0,1]$. For each experimental run, we computed ϕ^* for all paths starting at the first link and traversing a fixed number of hops in sequence. We obtained ϕ^* numerically by applying the normal approximation discussed above and then running a gradient descent algorithm on the resulting error probability function. For each path length, we collected 1000 samples of the marking bit with both SAM and a version of REM configured with ϕ^* for that path. We then compared the reduction in error probability

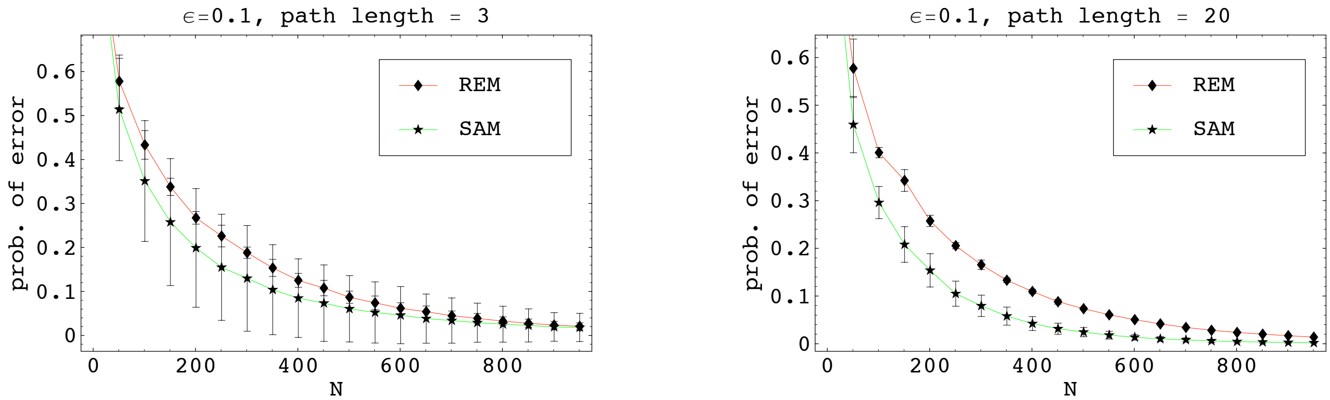

Figure 3: Error probability as a function of number of samples for and optimally parameterized REM and for SAM.

as samples are accumulated for SAM and the optimally parameterized REM.

Figure 3 shows results for path lengths of 3 and 20 averaged over 10 experiment iterations. We see that SAM still performs as well or better than REM even when ϕ^* is known in advance.

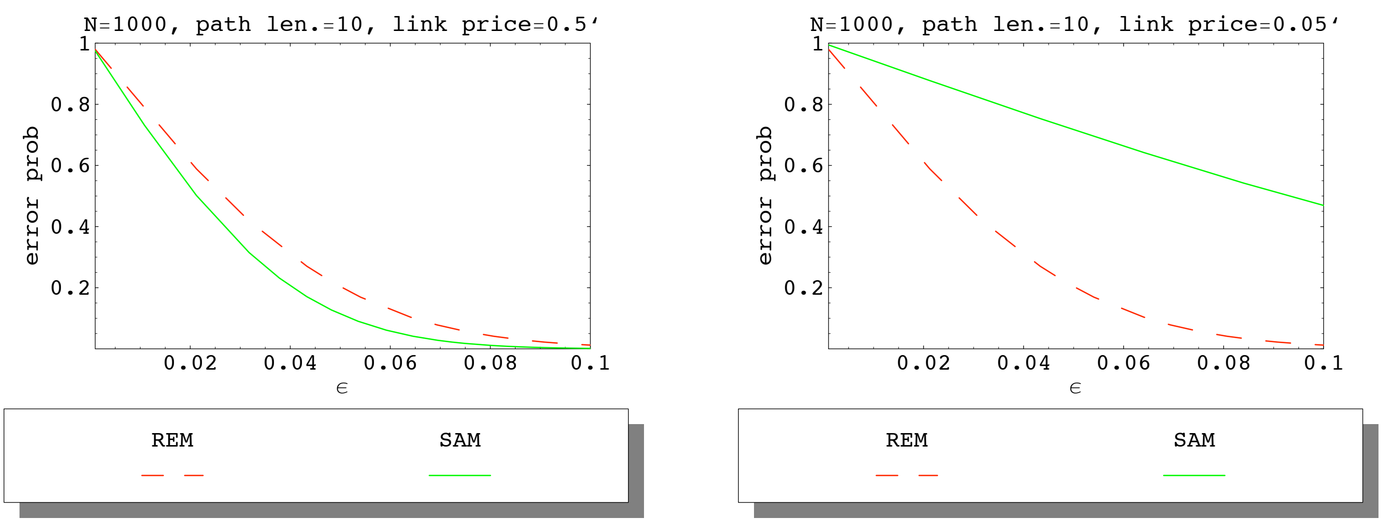

Figure 4: Sensitivity of error probability to the parameter ε for a path of 10 links with constant price of 0.5 on the left and 0.05 on the right. These figures show the importance of correctly normalizing link price for SAM.

Figure 4 shows the dependence of error probability on the error tolerance parameter ε for both SAM and optimally parameterized REM. For this analysis, we fix the price on each link to the same value and evaluate (43) for a range of ε. The figures presented use a path of 10 links and 1000 samples. In the left-hand figure the link price is set to 0.5. It turns out that SAM slightly outperforms optimized REM here. In the right-hand plot the link price is set to 0.05. Here optimized REM clearly outperforms SAM because optimized REM is able to maintain an end-to-end marking probability close to 0.5, which SAM cannot do. These results indicate that the performance of SAM relies on link prices being normalized "correctly"; at the very least, we require a mean link price close to 0.5.[6] We consider this issue in more detail in Section 4.3 In related results, omitted here for reasons of space, we find *unoptimized* REM can perform as badly as SAM (or worse) if ϕ is set far from its optimal value.

4.3 Optimality in terms of Mean Square Error

The comparison among different *statistics* $\hat{\theta}$ for the same quantity θ is, in general, a multi-faceted endeavor. Several factors enter into consideration. One can compare with respect to mean, variance or higher moments, or tail distributions (as we have done above). But in terms of tail distributions, there is the choice of the parameter ε, and the comparison based on the quantity $\Pr[|\hat{\theta}-\theta| > \varepsilon]$ can vary: one *statistic* could be better than another for one setting of ε, but worse for another. In terms of convergence when the sample number $N \to \infty$, one can also discuss the rate of convergence. Finally, there is the issue of prior distribution of the parameter θ itself. If we attempt to compare REM with SAM there is the additional difficulty that in REM the estimated parameter ranges over all $[0,\infty)$ while SAM makes some a priori assumption on the range. Taking into account of all these disparate considerations, one classical choice of a measure in such cases is called *Mean Square Error* (M. S. E.) with respect to an a priori distribution on θ.

For this subsection, let us define our parameter to be $\theta = \sum_{i=1}^{n} s_i/n$. To improve the prospects of REM in this comparison, we can allow the parameter ϕ to depend on n as

$$\phi(n) = \phi^{1/n}$$

. In the remainder of this section, we assume that ϕ implicitly contains this dependence.

The formulation of M. S. E. in our problem is as follows: Suppose θ has distribution $d\mu(\theta)$. For each parameter θ, the protocol constructs a 0-1 random variable Y with $\Pr[Y=1] = F(\theta)$. Then N i.i.d. samples are taken, and the mean $\overline{Y}$ is computed. Then we estimate θ by $G(\overline{Y})$, a function of the mean. The *M. S. E.* is

$$\int_{\Theta} \mathbf{E}_{F(\theta)}[(G(\overline{Y})-\theta)^2]d\mu(\theta).$$

[6] Whether link prices can be normalized correctly remains an open question. For now, however, we will assume that such a normalization is possible for practical implementations.

As we show in Section 3, REM corresponds to $F(\theta) = 1-\phi^{-\theta}$, and $G(Y) = -\log_\phi(1-Y)$ Unfortunately, REM has infinite expectation and mean square error, and consequently performs poorly in terms of M. S. E.,

$$\mathbf{E}_{F(\theta)}[G(\overline{Y})] = \infty$$

and

$$\mathbf{E}_{F(\theta)}[G(\overline{Y})-\theta]^2 = \infty.$$

This is because there is a non-zero probability that $\overline{Y} = 1$, and then $G(\overline{Y}) = \infty$. Note that this $G(Y)$ is the inverse function of $F(\theta) = 1-\phi^{-\theta}$ as defined over the *infinite interval* $[0,\infty)$. When we compare REM with SAM over $[0,1]$, a natural modification to REM is to "infer" $\theta = 1$ whenever the statistic $\overline{Y} > \max F(\theta) = F(1) = 1-\phi^{-1}$. For a more fair comparison between the two algorithms, it is reasonable to modify REM by taking its inference function G defined on $[0,1]$ to be

$$G(y) = \begin{cases} F^{-1}(y) & \text{if } 0 \leq y \leq 1-\phi^{-1} \\ 1 & \text{if } 1-\phi^{-1} < y \leq 1 \end{cases} \tag{55}$$

Note that $[0,1-\phi^{-1}] = [F(0),F(1)]$, so that G is still the inverse function of F on the range of F, and thus Theorem 4 (presented below) applies. With this modification, REM no longer has infinite expectation and square error.

Theorem 4. *Let F be a continuously differentiable and strictly monotonic function on $[0,1]$ and let $G = F^{-1}$ be its inverse function defined on the image interval of $F([0,1])$.*[7] *Let N be an integer ≥ 1. Let $\overline{Y} = \sum_{k=1}^{N} Y_k/N$ where $Y_1,\ldots,Y_N$ are i.i.d. 0-1 random variables with $\Pr[Y_i = 1] = F(\theta)$. Then the M.S.E. of $G(\overline{Y})-\theta$ has an approximate order of $\frac{1}{N}\cdot\int_0^1 \frac{F(\theta)(1-F(\theta))}{[F'(\theta)]^2}d\theta$. i.e.,*

$$N\cdot\int_0^1 \mathbf{E}_{N,F(\theta)}[(G(\overline{Y})-\theta)^2]d\theta \to \int_0^1 \frac{F(\theta)(1-F(\theta))}{[F'(\theta)]^2}d\theta,$$

when $N\to\infty$. Here $\mathbf{E}_{N,F(\theta)}$ denotes the expectation over the Binomial Distribution $B(N,F(\theta))$.

Proof. For any $\theta \in (0,1)$, we denote by $\theta' = F(\theta)$, then $\theta = G(\theta')$, Thus by the Mean Value Theorem, there exists some $\xi = \xi(\overline{Y},\theta)$ between $\overline{Y}$ and θ', such that

$$G(\overline{Y})-\theta = G'(\xi)(\overline{Y}-\theta'). \tag{56}$$

[7] We note that if F is monotonic increasing then this interval is $[F(0),F(1)]$, and if it is decreasing then it is $[F(1),F(0)]$. Moreover, where G is defined, G is also continuously differentiable and $G'(F(\theta)) = 1/F'(\theta)$.

In the following the expectation sign $\mathbf{E}$ will denote $\mathbf{E}_{N,\theta'}$ over the Binomial Distribution $B(N,\theta')$. Note that

$$\mathbf{E}[(\overline{Y}-\theta')^2] = \mathbf{Var}[\overline{Y}] = \frac{\theta'(1-\theta')}{N}. \tag{57}$$

Denote by $v(\theta) = \theta'(1-\theta')$. Denote by $\alpha = N \cdot (\overline{Y}-\theta')^2$, $\beta = (G'(\xi))^2$ and $\beta_0 = (G'(\theta'))^2 = \frac{1}{(F'(\theta))^2}$, then (57) can be written as

$$\mathbf{E}[\alpha] = v(\theta). \tag{58}$$

Now

$$N \cdot \int_0^1 \mathbf{E}[(G(\overline{Y})-\theta)^2] d\theta = \int_0^1 \mathbf{E}[\alpha \cdot \beta] d\theta. \tag{59}$$

We wish to exchange the order of the limit and the integration.

Since G is continuously differentiable, $(G'(y))^2$ is bounded on the compact set $F([0,1])$, thus, there exists some constant A, such that $(G'(y))^2 \leq A$, for all $y \in F([0,1])$. Therefore the quantity under the integral $\int_0^1$ is bounded

$$\mathbf{E}[\alpha \cdot \beta] \leq Av(\theta),$$

which is an integrable function of θ, independent of N.

Hence Lebesgue's Dominated Convergence Theorem [12, 7] applies, and we get

$$\lim_{N\to\infty} \int_0^1 \mathbf{E}[\alpha \cdot \beta] d\theta = \int_0^1 \lim_{N\to\infty} \mathbf{E}[\alpha \cdot \beta] d\theta. \tag{60}$$

Next we want to show that, for every fixed $\theta \in (0,1)$,

$$\lim_{N\to\infty} \mathbf{E}[\alpha \cdot \beta] = v(\theta) \cdot \beta_0. \tag{61}$$

We note that by (57) and (58), the RHS of (61) is exactly $\mathbf{E}[\alpha] \cdot \beta_0$, which is also $\mathbf{E}[\alpha \cdot \beta_0]$, since β_0 is a constant for a fixed θ. Note also that, the RHS of (61) is by definition the same as

$$\frac{F(\theta)(1-F(\theta))}{(F'(\theta))^2},$$

therefore the claim in (61) will complete the proof.

To show (61), note that $(G'(y))^2$ is continuous, and thus $\forall \varepsilon > 0, \exists \delta > 0$, such that whenever $|y-\theta'| < \delta$, $|(G'(y))^2 - (G'(\theta'))^2| < \varepsilon/(2v(\theta))$.

By Chernoff bound [2], for the given ε and δ, $\exists N_0$, such that if $N > N_0$, then

$$\mu[|\overline{Y}-\theta'| \geq \delta] < 2e^{-2\delta^2 N} < \varepsilon/(4AN), \tag{62}$$

where μ denotes the probability measure on the Binomial measure space $\Omega = B(N, \theta')$.

Denote by $B = \{|\overline{Y} - \theta'| \geq \delta\}$ the event in Ω. We note that the expectation $\mathbf{E}$, or more precisely $\mathbf{E}_{N,\theta'}$, is nothing but integration with respect to the measure μ, and thus we can write

$$\mathbf{E}[\alpha \cdot \beta] - \mathbf{E}[\alpha \cdot \beta_0] = \int_{\Omega} \alpha \cdot (\beta - \beta_0) d\mu = I_1 + I_2, \tag{63}$$

where

$$I_1 = \int_B \alpha \cdot (\beta - \beta_0) d\mu \tag{64}$$

$$I_2 = \int_{\Omega - B} \alpha \cdot (\beta - \beta_0) d\mu \tag{65}$$

On $\Omega - B$, $|\xi - \theta'| < \delta$, and so

$$|\beta - \beta_0| = \left|(G'(\xi))^2 - (G'(\theta'))^2\right| < \varepsilon/(2v(\theta)). \tag{66}$$

Therefore, since α is non-negative,

$$|I_2| \leq \int_{\Omega - B} \alpha \cdot |\beta - \beta_0| \, d\mu \tag{67}$$

$$\leq \frac{\varepsilon}{2v(\theta)} \int_{\Omega - B} \alpha d\mu \tag{68}$$

$$\leq \frac{\varepsilon}{2v(\theta)} \int_{\Omega} \alpha d\mu \tag{69}$$

$$= \frac{\varepsilon}{2}. \tag{70}$$

For I_1, we have $|\beta - \beta_0| \leq 2A$, and $\alpha \leq N$, thus by the Chernoff bound (62),

$$|I_1| \leq 2A \cdot \int_B \alpha d\mu \tag{71}$$

$$\leq 2AN\mu(B) \tag{72}$$

$$< \frac{\varepsilon}{2}. \tag{73}$$

Combining (71) and (67), we get

$$|\mathbf{E}[\alpha \cdot \beta] - \mathbf{E}[\alpha \cdot \beta_0]| \leq |I_1| + |I_2| < \varepsilon. \tag{74}$$

□

4.4 Comparison within the REM family

Consider a REM-like protocol over $[0,\infty)$, with probability function $p(s) = a + b\varphi^{-s}$ as in Theorem 2. If we take $F(\theta) = F_{a,b}(\theta) = a + b\varphi^{-\theta}$, we observe that the integration over the whole interval $[0,\infty)$

$$\int_0^\infty \frac{F(\theta)(1-F(\theta))}{[F'(\theta)]^2} d\theta = \infty,$$

for all choices of a and b, where $0 \leq a, a+b \leq 1$ and $b \neq 0$, as stipulated in Section (3). Thus, in order to compare REM or REM* with the general $F_{a,b}$, we will consider only any finite interval $[0,A]$.

In this case, let

$$\begin{aligned}
I_{a,b} &= \frac{1}{A}\int_0^A \frac{F(\theta)(1-F(\theta))}{[F'(\theta)]^2} d\theta && (75)\\
&= \frac{1}{Ab^2(\log\varphi)^3}\int_1^{\varphi^{-A}} \frac{(a+bz)(1-a-bz)}{z^3} dz && (76)\\
&= \frac{1}{Ab^2(\log\varphi)^3}\left[-Ab^2\log\varphi + b(1-2a)(\varphi^A - 1) + \frac{a(1-a)(\varphi^{2A}-1)}{2}\right] && (77)\\
&= \frac{-1}{(\log\varphi)^2} + \frac{(\varphi^A - 1)}{2Ab^2(\log\varphi)^3}\left[2b(1-2a) + a(1-a)(\varphi^A + 1)\right]. && (78)
\end{aligned}$$

Note that for REM and REM*

$$I_{1,-1} = I_{0,1} = \frac{-1}{(\log\varphi)^2} + \frac{(\varphi^A - 1)}{A(\log\varphi)^3},$$

we see that

$$I_{a,b} - I_{1,-1} = \frac{(\varphi^A - 1)}{2Ab^2(\log\varphi)^3}\left[2b(1-2a) + a(1-a)(\varphi^A + 1) - 2b^2\right].$$

Note that $\varphi^A > 1$, the quantity in the bracket is $\geq 2(b - 2ab + a(1-a) - b^2) = 2[(a+b)(1-(a+b))] \geq 0$, and furthermore the inequality is strict for all choices of (a,b) other than for $(1,-1)$ or $(0,1)$.

We conclude that among REM-like protocols, in terms of M.S.E., for every choice of the base $\varphi > 1$, and every finite interval $[0,A]$, the choice of parameters (a,b) is uniquely optimal with $(1,-1)$ for REM and $(0,1)$ for REM*.

4.4.1 Comparison within the SAM family

We now concentrate on the family of SAM-like estimators, where $F(\theta) = a + b\theta$, and G is the inverse function of F. Here, from Theorem 3, $0 \leq a, a+b \leq 1$ and $b \neq 0$, since F is strictly monotonic and represents a probability.

We first show that within the family of SAM-like protocols identified in Theorem 3, the SAM protocol presented in Section 2 is optimal in terms of M. S. E. To see this, we compute the difference in variance between an arbitrary feasible SAM-like protocol and SAM itself. Note that since the probability functions for SAM-like protocols are affine linear, $\mathbf{E}[G(\overline{Y}) - \theta]^2 = \mathbf{Var}[G(\overline{Y})]$. For any SAM-like protocol,

$$\begin{aligned}
\mathbf{Var}[G(\overline{Y})] &= \mathbf{Var}[\frac{\overline{Y}}{b}] \\
&= \frac{\mathbf{Var}[Y]}{b^2 N} \\
&= \frac{F(\theta)(1-F(\theta))}{b^2 N} \\
&= \frac{1}{N}\left(\theta + \frac{a}{b}\right)\left(\frac{1-a}{b} - \theta\right) \qquad (79)
\end{aligned}$$

SAM corresponds to $F(\theta) = \theta$, thus

$$\mathbf{Var}[\overline{Y}] = \frac{\theta(1-\theta)}{N}. \qquad (80)$$

Hence the difference

$$N \cdot [\mathbf{Var}[G(\overline{Y})] - \mathbf{Var}[\overline{Y}]] = \frac{a(1-a)}{b^2} + \left(\frac{1-b-2a}{b}\right)\theta.$$

Now we assume θ has a uniform distribution on $[0,1]$, then

$$\begin{aligned}
& N \cdot \int_0^1 [\mathbf{Var}[G(\overline{Y})] - \mathbf{Var}[\overline{Y}]] d\theta \\
&= \frac{1}{2b^2}[(a+b)(1-a-b) + a(1-a)] \\
&\geq 0, \qquad (81)
\end{aligned}$$

by elementary calculation, and since $0 \leq a, a+b \leq 1$.

We note that this result holds for *any* distribution on θ that has expectation $1/2$. In particular, this would apply if each s_i is independently distributed and symmetric about $1/2$.

The inequality (81) is strict, unless $(a+b)(1-a-b) = a(1-a) = 0$, which can happen in one of two ways (since $b = 0$ is not allowed). If $a = 0$ then $b = 1$ and we have SAM as presented in Section 2. If $a = 1$ then $b = -1$ and we have the dual of SAM with $F(\theta) = 1 - \theta$. We conclude that in terms of M.S.E. with respect to uniform distribution (or any other distribution with expectation 1/2) on θ, SAM (or its dual) is optimal.

4.4.2 Comparison of SAM and REM

For REM, we obtain a lower bound for the M. S. E. using the following theorem:

Theorem 5. *For every $\phi > 1$, with $F(\theta) = 1 - \phi^{-\theta}$ and G defined in (55), the M.S.E. of REM is asymptotically greater than $\frac{.54685578}{N}$. More precisely,*

1.

$$N \cdot \int_0^1 \mathbf{E}_{N,F(\theta)}[(G(\overline{Y}) - \theta)^2] d\theta \to \frac{\phi - 1 - \log_e \phi}{(\log_e \phi)^3} \tag{82}$$

2.

$$I(\phi) = \frac{\phi - 1 - \log_e \phi}{(\log_e \phi)^3}, \tag{83}$$

is strictly monotonic decreasing in $[1, \phi_0)$, and strictly monotonic increasing in (ϕ_0, ∞), and achieves an unique minimum at ϕ_0, with value $I(\phi_0) = \frac{(\phi_0+2)2}{27(\phi_0-1)}$. Here ϕ_0 is the unique solution to the equation $\frac{1}{\log \phi} - \frac{1}{\phi - 1} = \frac{1}{3}$, and $\phi_0 \approx 8.577356793$, and $I(\phi_0) \approx .54685578$.

(In the following log will stand for $\log_e$, namely log base e.)

Proof. By Theorem 4, we only need to evaluate the integral

$$\int_0^1 \frac{F(\theta)(1 - F(\theta))}{[F'(\theta)]^2} d\theta$$

which can be shown easily to be

$$I(\phi) = \frac{\phi - 1 - \log_e \phi}{(\log_e \phi)^3}.$$

The second claim takes a little more work. First we note that $\lim_{\phi \to \infty} I(\phi) = \infty$. When $\phi \to 1^+$, write $\phi = 1 + \xi$, then the numerator of $I(1 + \xi)$ has a zero of order 2 but the denominator has zero of order 3.

$$I(1 + \xi) = \frac{\xi - (\xi - \xi^2/2 + \dots)}{(\xi - \xi^2/2 + \dots)^3} \to \infty.$$

Thus $\lim_{\phi \to 1^+} I(\phi) = \infty$ as well.

We now wish to prove that $I(\phi)$ has a unique minimum for $\phi \in [1, \infty)$.

Let

$$h(\phi) = \phi - 1 - \log \phi.$$

Since $h'(\phi) = 1 - 1/\phi > 0$, for all $\phi > 1$, h is strictly monotonic increasing. Also $h(1) = 0$. Therefore

$$h(\phi) > 0, \quad \forall \phi > 1.$$

Let

$$g(\phi) = (\phi-1)^2 - \phi(\log\phi)^2.$$

Then $g'(\phi) = 2(\phi-1) - (\log\phi)^2 - 2\log\phi$, and

$$g''(\phi) = 2 - \frac{2\log\phi}{\phi} - \frac{2}{\phi} = \frac{2h(\phi)}{\phi}.$$

Since $h(\phi) > 0$, $g''(\phi) > 0$, for all $\phi > 1$. Hence, $g'(\phi)$ is strictly monotonic increasing. But $g'(1) = 0$. It follows that $g'(\phi) > 0$ for all $\phi > 1$. Hence, $g(\phi)$ is strictly monotonic increasing. Then $g(1) = 0$, we conclude that g is always positive

$$g(\phi) > 0, \quad \forall\phi > 1.$$

Let

$$f(\phi) = \frac{1}{\log\phi} - \frac{1}{\phi-1}.$$

$f'(\phi) = -\frac{1}{(\log\phi)^2\phi} + \frac{1}{(\phi-1)^2}$. Since $g(\phi) > 0$,

$$f'(\phi) < 0, \quad \forall\phi > 1.$$

So f is strictly monotonic decreasing for all $\phi > 1$.

Setting $\phi = 1+\xi$, we can expand f at $\xi = 0$ in a Laurent series,

$$\begin{aligned} \frac{1}{\log(1+\xi)} - \frac{1}{\xi} &= \frac{1}{\xi}\left[1 + (\frac{\xi}{2} - \frac{\xi^2}{3} + \dots) + (\frac{\xi}{2} - \frac{\xi^2}{3} + \dots)^2 + \dots\right] - \frac{1}{\xi} && (84) \\ &= \frac{1}{\xi} + \frac{1}{2} - \frac{1}{12}\xi + \dots - \frac{1}{\xi} && (85) \\ &= \frac{1}{2} - \frac{1}{12}\xi + \dots \to \frac{1}{2}. && (86) \end{aligned}$$

Also

$$\lim_{\phi\to\infty} f(\phi) = 0.$$

Hence

$$f(\phi) - \frac{1}{3}, \quad (87)$$

has a unique zero ϕ_0 in $[1,\infty)$. Numerically, $\phi_0 \approx 8.577356793$.

Finally we return to $I(\phi)$.

$$I'(\phi) = \frac{\left(1-\frac{1}{\phi}\right)(\log\phi)^3 - 3(\log\phi)^2\cdot\frac{1}{\phi}\cdot(\phi-1-\log\phi)}{(\log\phi)^6} \tag{88}$$

$$= \frac{3(\phi-1)}{\phi(\log\phi)^3}\cdot\left[\frac{1}{3}-\left(\frac{1}{\log\phi}-\frac{1}{\phi-1}\right)\right] \tag{89}$$

$$= \frac{3(\phi-1)}{\phi(\log\phi)^3}\cdot\left[\frac{1}{3}-f(\phi)\right]. \tag{90}$$

Thus, $I'(\phi) < 0$ for $\phi < \phi_0$, and $I'(\phi) > 0$ for $\phi > \phi_0$. So $I(\phi)$ achieves a unique minimum at ϕ_0. It follows by elementary calculation and (87) that

$$I(\phi_0) = \frac{(\phi_0+2)^2}{27(\phi_0-1)} \approx .54685578$$

□

The M. S. E. for SAM is readily computed as follows. Noting that $G(Y) = Y$, $E[\overline{Y}] = \theta$, and making use of (80) we obtain

$$E[(\overline{Y}-\theta)^2] = \frac{\theta(1-\theta)}{N}$$

If we assume θ has a uniform distribution on $[0,1]$, then the M. S. E. is $1/6N$ Comparing this result with the asymptotic M. S. E. for REM in Theorem 5, we conclude that, for a uniform a priori distribution of θ, no matter what choice we make for the base ϕ in REM, in terms of M.S.E. it is worse than SAM over a finite interval.

4.4.3 Comparison of SAM and extended REM

How does the "stretched" version of REM compare to SAM?

We apply Theorem 4. For $F(\theta)$ in (40) we get

$$\frac{F(\theta)(1-F(\theta))}{[F'(\theta)]^2} = \frac{(\psi^\theta-1)(1-\psi^{\theta-1})}{(\log\psi)^2}.$$

The integration can be computed in closed form

$$I(\psi) = \int_0^1 \frac{F(\theta)(1-F(\theta))}{[F'(\theta)]^2}d\theta = \frac{1}{2(\log\psi)^3}\left[\psi - 2\log\psi - \frac{1}{\psi}\right]. \tag{91}$$

It is easy to see that

$$\lim_{\psi\to\infty} I(\psi) = \infty.$$

For $\psi \to 1^+$, write $\phi = 1+\xi$, then both the numerator and denominator of $I(1+\xi)$ have a zero of order 3. By expanding in Taylor series,

$$\psi - 2\log\psi - \frac{1}{\psi} = \frac{\xi^3}{3} - \frac{\xi^4}{2} + \dots,$$

and

$$2(\log\psi)^3 = 2\xi^3 - 3\xi^4 + \dots,$$

and it follows that

$$\lim_{\psi\to 1^+} I(\psi) = \frac{1}{6}.$$

Recall that $1/6$ is exactly the same value for SAM in the corresponding integral! Of course this should be no surprise since in fact, both $F(\theta)$ and $F'(\theta)$, as a family of functions indexed by the base ψ, approaches to the identity function θ and its dirivative 1 *uniformly* over $[0,1]$, as $\psi \to 1^+$. Thus we could also have derived the limiting value by taking limit under the integral sign.

However, by having the closed form (91), we can show that $I(\psi) > 1/6$ for all $\psi > 1$. This means that the "streched" REM can get close to SAM but always strictly inferior.

The proof that $I(\psi) > 1/6$ for all $\psi > 1$ is enclosed below.

4.4.4 The function $I(\psi)$

We give the proof that $I(\psi) > 1/6$ for all $\psi > 1$.

Let $f(x) = \frac{x-2\log x - \frac{1}{x}}{(\log x)^3}$. We wish to show that $f(x) > 1/3$ for all $x > 1$. As we know $\lim_{x\to 1^+} f(x) = 1/3$. It is sufficient to show that $f(x)$ is strictly monotonic increasing for $x > 1$.

$$f'(x) = \frac{1}{(\log x)^4}\left[(1-\frac{2}{x}+\frac{1}{x^2})\log x - 3(1-\frac{2\log x}{x} - \frac{1}{x^2})\right].$$

Let g denote the nemerator,

$$g(x) = \log x(1+\frac{4}{x}+\frac{1}{x^2}) - 3(1-\frac{1}{x^2}).$$

To show $f'(x) > 0$, it suffices to show that $g(x) > 0$, for $x > 1$. Since $g(1) = 0$, it suffices to show that $g'(x) > 0$, for $x > 1$, thus g is strictly monotonic increasing, and in particular $g(x) > 0$, for all $x > 1$.

$$g'(x) = \frac{1}{x}\left[1+\frac{4}{x}-\frac{5}{x^2}-\frac{2\log x}{x}\left(2+\frac{1}{x}\right)\right].$$

Denote the expression inside the bracket as h, and change the variable to $y = 1/x$, we get

$$h(y) = 1 + 4y - 5y^2 + 2y\log y(2+y).$$

It is sufficient to show that $h(y) > 0$ for all $0 < y < 1$.

We have

$$h'(y) = 4 - 10y + 2((2+y) + \log y(2+2y)) = 4[2 - 2y + (1+y)\log y].$$

$$h''(y)/4 = -2 + \log y + (1+y)/y = 1/y - 1 + \log y = x - 1 - \log x.$$

As

$$(x - 1 - \log x)' = 1 - 1/x > 0,$$

for $x > 1$, and $[x - 1 - \log x]\mid_1 = 0$, it follows that $x - 1 - \log x > 0$ for all $x > 1$, i.e., $h''(y) > 0$ for all $0 < y < 1$.

So $h'(y)$ is strictly monotonic increasing for $0 < y < 1$. And $h'(1) = 0$ implies that $h'(y) < 0$ for $0 < y < 1$.

Hence h is strictly monotonic decreasing for $0 < y < 1$. And $h(1) = 0$ implies that $h(y) > 0$ for $0 < y < 1$. This completes the proof.

5 Implementing SAM in the Internet

Implementation of SAM on the current Internet is complicated by the fact that routers are typically not aware of their position along the path taken by a particular packet. Without this information, a router clearly cannot determine the correct marking probability for an incoming packet. One possible way to address this difficulty is to introduce a field in the IP header to be incremented at each hop which would contain the path length. However, requiring a change to a standard header would be a serious barrier to deployment. In addition, introducing a new field would effectively make additional bits available for packet marking and these bits might be used more profitably by some alternative marking scheme. Since we are interested in easily deployed single-bit schemes, we are motivated to explore other solutions.

The time-to-live (TTL) field in the IP header is an 8-bit field used to limit the maximum lifetime of a packet in the network. In addition to serving this intended purpose, the TTL field provides some information about path lengths and thus could plausibly be used by a marking algorithm. Unlike a path length field that is initialized to zero and incremented, TTL is initialized to some positive value and decremented. One problem with using TTL to perform marking within the network is that the routers along the path are unaware of the initial value placed in the TTL. Another problem is that the IP protocol allows routers to decrement the TTL value by more than one. Thus, even if a router knew the initial TTL value, it could not be sure of the number of intervening routers between it and the source on the basis of the observed value.

In the remainder of this section, we show how SAM can be implemented in the Internet using only the existing IP TTL field and a single ECN bit for marking. We do not require that routers know the initial TTL value. Instead, we will initially assume that the TTL field is always initialized to the maximum value of 255. We will show that in the case when the TTL is actually initialized to a lower value, the protocol still computes a correct estimate. Conceptually, assuming too high a value is equivalent to appending a chain of links with zero price to the beginning of the path, which collectively decrement the TTL to its actual initial value. However, such a mismatch between the guessed initial TTL value and the true value leads to slower convergence. We assume that each router knows the amount by which it will decrement the TTL, but require no knowledge about the behavior of other routers.

Consider the i^{th} link along a path. Assume that the link is initialized with its own price s_i, the maximum possible TTL value, denoted Ω, and the amount by which it will decrement the TTL of any packet passing through it, denoted k_i. In addition to these stored values, each arriving packet provides the router with an ECN bit, with expected value θ_{i-1}, and a TTL field with value τ_i. Note that we may write

$$\begin{aligned} \tau_i &= T - K_{i-1}, \\ K_{i-1} &= \sum_{1}^{i-1} k_i \end{aligned}$$

For each packet received, the router computes

$$\begin{aligned} t_i &= \Omega - \tau_i \\ &= \tilde{T} + K_{i-1}. \end{aligned}$$

The value t_i is the path position inferred by router i and has the property $t_i \geq i$ with equality holding in the case that the TTL field is actually initialized to Ω and each preceding router only decrements the TTL by one. Also note that necessarily $t_i > t_{i-1}$ for all i.

Theorem 6. The expected value of the marking bit emerging from a chain of n routers running Algorithm 1 is

$$a_n = \frac{z_n}{(\tilde{T} + K_n)}.$$

Algorithm 1 TTL-SAM algorithm

Given: s_i, k_i, Ω

Input: (τ_i, θ_{i-1})

With probability $\frac{t_i}{t_i+k_i}$, set $\theta_i = \theta_{i-1}$

With probability $\frac{s_i}{t_i+k_i}$, set $\theta_i = 1$

Otherwise, set $\theta_i = 0$

$\tau_{i+1} = \tau_i - k_i$

Output: (τ_i, θ_i);

Proof: The proof is by induction on n, the length of the router chain. The base case for $n = 1$ follows trivially from the Algorithm definition. We provide the inductive step. Consider the expected value of the bit emerging from router i

$$\theta_i = \frac{t_i}{t_i + k_i}\theta_{i-1} + \frac{1}{t_i + k_i}s_i$$

Using the substitution $t_i = \tilde{T} + K_{i-1}$ and the fact that $K_i = K_{i-1} + k_i$, we have

$$\theta_i = \frac{\tilde{T} + K_{i-1}}{\tilde{T} + K_{i-1} + k_i}\theta_{i-1} + \frac{1}{\tilde{T} + K_i}s_i$$

By hypothesis,

$$\theta_{i-1} = \frac{z_{i-1}}{\tilde{T} + K_{i-1}}.$$

The theorem follows. □

The receiver can recover the sum of path prices using an estimate of the marking probability$\hat{\theta}_n$ and the TTL value of arriving packets $\tau_{n+1} = T - K_n$. The path price estimate is simply

$$\hat{\theta}_n \cdot (\Omega - \tau_{n+1})$$

In practice, it is extremely rare for routers to decrement the TTL by more than one. We will therefore assume henceforward that $k_i = 1$ for all i. It is also rare for sources to initialize the TTL field to its maximum value. The IP standard simply states that the TTL should be at least as large as the (unknown) diameter of the Internet [5] with 64 being a recommended value [10]. The default values chosen by popular operating systems vary between 30 and 255 [1]. There is a motivation to choose as low a value as possible to limit the lifetime of misrouted packets. Unfortunately, we expect the effect on SAM of a source setting TTL to less than the guessed value to be slower convergence since the probability of any router overwriting the marking bit would be reduced.

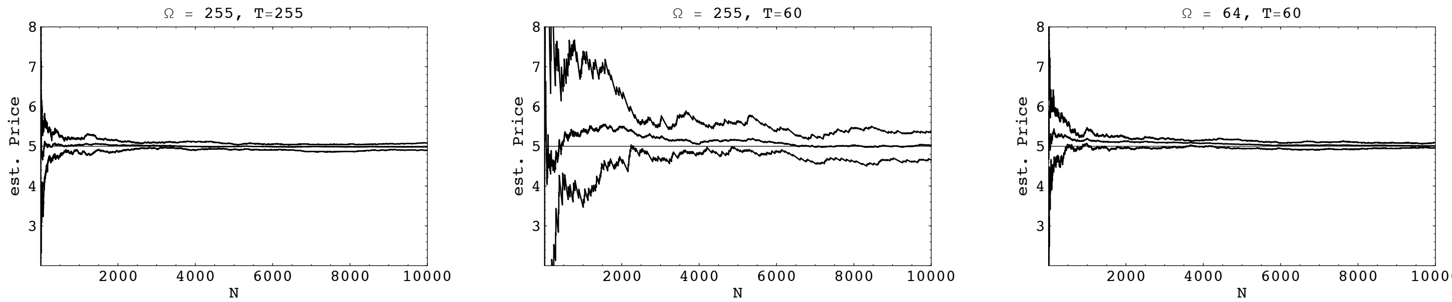

Figure 5: Convergence of SAM using the TTL field for different combinations of Ω and T.

Figure 5 shows the convergence of SAM for three different combinations of Ω and T on a 10 link path with a price of 0.5 on each link. For each setting of the parameters we executed 10 simulation runs, collecting 10^4 packets in each run. The plots in Fig. 5 show the evolution of the minimum, mean and maximum price estimates. We see that in all three cases, the mean price estimate quickly converges to the correct value, but that the mismatch between Ω and T introduces substantial variability in the estimate. If we can ensure a small difference between the initial TTL and the guessed value, SAM can achieve extremely good performance. If SAM were adopted in the Internet, this fact would serve as a strong incentive for users to set initial TTLs to 255 and for operating systems to standardize around 255 as a default.

However, despite the fact that initial TTL values are user-configurable parameters in most modern operating systems, users typically do not modify the default setting unless extremely long paths are encountered. Indeed it is likely that in many cases users do not know how to change these parameters or lack the authorization to do so. However, we note that the default values chosen in practice by operating systems tend to be equal to or slightly less (between 1 and 4) than some power of two and are never lower than 30 [1]. Furthermore, measurements put the average path length in the Internet somewhere around 16 hops with paths of more than 30 hops being exceedingly rare [4]. Based on these observations, it may be possible for a SAM router to make a better guess for Ω based on the observed TTL value of a packet. For example, SAM might guess that Ω is the smallest power of two greater than τ but at least 32. Specifically, define

$$\begin{aligned} \Omega(\tau) &= [2^{\lambda}]_{32}^{255} \\ \lambda &= \lfloor \log_2 \tau \rfloor, \end{aligned} \tag{92}$$

where $[x]_a^b = \min(\max(a,x),b)$. Using this rule, the guessed initial TTL will likely be very close (within 4) to the actual value. In extremely rare cases, a path may be so long that the guessed TLL will change at some point along the path. Consider, for example, a packet with initial TTL of T traversing a long path. For simplicity of explanation, assume T is a power of 2. The first $k = T/2$ routers along the path will correctly set $\Omega = T$. At router, $k+1$, however, $\Omega = T/2$. The expected value of the marking bit arriving at this router is $\theta_k = z_k/k$. The TTL-SAM algorithm at link $k+1$ will assume it is the first link along the path since $t_{k+1} = 0$ and with therefore overwrite the arriving bit with probability one, destroying all information about the path prior to itself. Unfortunately, link $k+1$ cannot distinguish between being the first link in the path and guessing a value of Ω lower than preceding routers.

It can be shown that the true path price cannot be recovered by means of local corrections at the links when Ω changes mid-path. However, this situation can be detected at the receiver if the initial TTL value T is sent end-to-end by the source. Specifically, if the receiver sees that $T - \tau_{n+1} \geq T/2$ then it knows that the value of Ω changed along the path and the SAM price estimate must be regarded as biased. We emphasize that such biased estimates are very rare events. A packet with an initial TTL of 32 (the value used in older Microsoft operating systems) would be discarded by the network before generating such an event. A packet with an initial TTL of 60 (a value used in several real-world operating systems) would have to traverse 28 hops before reaching a TTL of 32. A packet with an initial TTL of 128 (the default value for newer Microsoft operating systems) would have to traverse 64 hops.

6 Conclusion

In this paper we have considered the problem of estimating the sum of congestion prices along a path using a one-bit probabilistic packet marking algorithms. We showed that REM, the only previously proposed algorithm we are aware of, is, in fact, essentially unique if link prices are unbounded. By introducing a finite bound on link prices and allowing links to know their position along a packet's path, we found that an alternate class of algorithms became possible. We introduced SAM, a novel marking algorithm and showed that SAM together with the existing REM algorithm represent the only two possible classes of marking algorithms when link prices have finite bounds. By examining the tail probabilities of the two price estimates, we demonstrated the difficulty in setting the parameter φ in REM, which makes REM difficult to deploy in heterogeneous network environments. Furthermore, we showed that in terms of mean squared error, SAM out-performs even an optimally parameterized REM. Finally, we showed that path position information required by SAM is already available in the form of the IP TTL field

Future work to be done in this area includes accurately characterizing the prior distribution of link prices and path prices in large networks. SAM depends critically on the assumption that link prices can be effectively normalized to a finite range, symmetrically distributed about a mean value. This is a strong assumption, given the nature of congestion prices, which represent gradients and thus can, in principle, take on infinite values and given the fact that, along any given path, only a few links are likely to be congested. Another open question is how to relate the performance of price estimation algorithms to the performance of the congestion control schemes in which they are embedded.

References

[1] ”Swiss Academic and Research Network”. Default ttl values in tcp/ip. ”http://www.switch.ch/docs/ttl_default.html”, 1999.

[2] Noga Alon and Joel H. Spencer. *The Probabilistic Method, Second Edition*. John Wiley & Sons, 2000.

[3] Sanjeewa Athuraliya and Steven H. Low. Optimization flow control ii: Implementation. Technical report, Caltech, 2000.

[4] F. Begtasevic and P. Van Mieghen. Measurements of the hopcount in internet. In *Proc. Passive and Active Measurement*, 2001.

[5] R. Braden. Requirements for internet hosts – communication layers. Technical report, IETF, 1989. RFC 1122.

[6] R. J. Gibbens and F. P. Kelly. Resource pricing and the evolution of congestion control. *Automatica*, 35, 1999.

[7] Paul R. Halmos. *Measure Theory*. Van Nostrand, 1950.

[8] S. H. Low and D. E. Lapsley. Optimization flow control, i: Basic algorithm and convergence. *IEEE/ACM Transactions on Networking*, December 1999.

[9] K. Ramakrishnan, S. Floyd, and D. Black. The addition of explicit congestion notification (ecn) to ip. Technical report, IETF, 2001.

[10] J. Reynolds and J. Postel. Assigned numbers. Technical report, IETF, 1994. RFC1700.

[11] Walter A. Rosenkrantz. *Introduction to Probability and Statistics for Scientists and Engineers*. McGraw Hill, 1997.

[12] Walter Rudin. *Real and Complex Analysis, Third Edition*. McGraw-Hill, 1987.

[13] K. Weihrauch. *Computable Analysis. An Introduction*. Springer-Verlag, 2000.