A Combinatorial Approach to Measuring Anonymity

Matthew Edman
Department of Computer Science
Rensselaer Polytechnic Institute
Troy, NY 12180, USA
edmanm2@cs.rpi.edu

Fikret Sivrikaya
Department of Computer Science
Rensselaer Polytechnic Institute
Troy, NY 12180, USA
sivrif@cs.rpi.edu

Bülent Yener
Department of Computer Science
Rensselaer Polytechnic Institute
Troy, NY 12180, USA
yener@cs.rpi.edu

***Abstract*— In this paper we define a new metric for quantifying the degree of anonymity collectively afforded to users of an anonymous communication system. We show how our metric, based on the permanent of a matrix, can be useful in evaluating the amount of information needed by an observer to reveal the communication pattern as a whole. We also show how our model can be extended to include probabilistic information learned by an attacker about possible sender-recipient relationships. Our work is intended to serve as a complementary tool to existing information-theoretic metrics, which typically consider the anonymity of the system from the perspective of a single user or message.**

I. Introduction

Starting with Chaum's work on mix-based anonymity systems [1], many research papers have been devoted to the subject of designing and evaluating systems for anonymous communication. Some designs exist only in the literature, while a few have been implemented and publicly deployed (*e.g.*, [2], [3]).

With each new system design or implementation, it becomes more important to be able to evaluate and compare the privacy afforded by a system to its users. Most previous work has focused on evaluating the anonymity of a system from the perspective of a single user or message in that system; however, it is not clear how to generalize such measurements to express the level of anonymity of the system as a whole (we elaborate on this point in Section II).

In this paper we define a new system-wide metric, based on the permanent of a matrix, which measures the amount of information needed by an observer to reveal the overall communication pattern between senders and recipients in an anonymity system. Our metric can be used alongside existing metrics, which typically measure the anonymity of a system from the perspective of a single user or message, in order to provide a more complete representation of the privacy provided by that system.

The rest of this paper is structured as follows. First, we will review some of the previous work done to establish measures of anonymity. In Section II, we introduce our model and formally define our proposed anonymity metric. We then show how our model can be generalized to include probabilistic information in Section III. We demonstrate how to apply our metric to some common types of mixes in Section IV. In Section V, we compare our metric to some existing metrics to illustrate the more important differences. Section VI describes how we can obtain a lower-bound on the anonymity level of a system while avoiding some computational complexity inherent in computing the permanent of a matrix. Finally, we conclude in Section VII.

This material is based upon work supported by the National Science Foundation under Grant No. 0540989.

A. Related Work

Chaum introduced the notion of an *anonymity set* in his work on DC-Networks [4]. An anonymity set is the set of participants who are likely to be the sender or recipient of a particular message. As the size of an anonymity set increases, so does the anonymity of the members of that set. Kesdogan, Egner, and Büschkes also use this metric for evaluating their design of Stop-and-Go MIXes (SG-MIXes) [5].

Serjantov and Danezis [6] showed that simply measuring the size of an anonymity set is inadequate for expressing instances where not all members of that set are equally likely to have sent a particular message. They go on to define an *effective anonymity set size*, based on the information theoretic concept of entropy, as

$$S = -\sum_{u=1}^{n} p_u \log_2(p_u),$$

where n is the number of users in the anonymity set, and p_u is the probability that a user u had a role $r \in \{sender, recipient\}$ for a particular message. The authors interpret the effective anonymity set size S as the amount of additional information the attacker needs in order to identify the user who was either the sender or recipient of a particular message.

Diaz et al. [7] independently proposed a similar entropy-based metric they refer to as the *degree of anonymity*, a term first introduced by Reiter and Rubin in the Crowds design [8]. They define the degree of anonymity d as

$$d = \frac{S}{S_{max}},$$

where S is as above and S_{max} is the maximum entropy of the system and is equal to $\log_2(n)$. The difference between the degree of anonymity and the previous entropy-based metric is that, by dividing S by the maximum entropy of the system, d is normalized to the range $[0, 1]$ and becomes a measure of anonymity independent of the number of users involved.

In Tóth and Hornák's analysis of non-adaptive real-time systems [9] they introduce the notions of *source-hiding* and *destination-hiding*. A system is source-hiding with parameter Θ if the observer cannot assign a sender to any delivered message with a probability greater than Θ. Similarly, a system is destination-hiding with parameter Ω if the observer cannot assign a recipient to any sent message with a probability greater than Ω.

Tóth and Hornák [10] later analyze the two entropy-based metrics and highlight some of the shortcomings of both by showing that a system can appear near-optimal according to the entropy-based metrics even though an attacker may still be able to guess the sender of some messages with high probability. They use this to argue for using, as a measure, the maximum probability that an attacker can assign to a sender or recipient with respect to a particular message. Such a measure focuses better on the local aspect of anonymity, which may be of more interest to individual users of the system.

Instead of measuring how much protection a system affords a single entity, whether that entity is a sender, recipient, or message in a system, Newman, Moskowitz, and Syverson [11] propose an entropy-based approach to evaluating how much protection a Traffic Analysis Prevention (TAP) system can provide to its users collectively. Specifically, the authors consider systems that perform actions such as padding and rerouting in order to increase the number of potential traffic matrices (TMs), thereby decreasing the probability of an observer being able to determine the actual TM based on her observations. The authors introduce an entropy-based approach to measure the amount of uncertainty the adversary has in determining the actual TM from the set of possible TMs.

II. A Permanent-based Anonymity Metric

The previous work on measuring anonymity mostly focuses on the level of anonymity from the perspective of a single user or message, with the most notable exception being the work done by Newman et al [11]. The anonymity set size [4], the entropy-based effective anonymity set size [6], and the normalized entropy-based metric [7] all follow this model.

Using such metrics, one can measure *sender anonymity* for a given recipient (or a message received). Or, one can measure *recipient anonymity* for a given sender (or a message sent). It is not clear how to generalize such a metric to clearly express a system-wide anonymity level. One can use the minimum degree of anonymity among all users as the anonymity degree of the whole system, but this only reveals the "weakest link" in the chain; it may not capture the overall system behavior. Another possibility is to add up the individual anonymity degrees for all users to obtain the anonymity level of the whole system. The problem with this approach is that it does not consider the interdependence between anonymity sets of different users.

We propose a new system-wide metric, based on the permanent of a matrix, which measures the amount of additional information needed to *reveal the whole communication pattern* between senders and recipients in the system. We stress that our approach is not intended to replace the existing entropy-based metrics. Rather, it can be used as a complementary tool, which we believe represents a reasonable and intuitive combination of the individual anonymity levels of the users of a system, to obtain a more complete, system-wide perspective.

A. Preliminaries

We consider an anonymity system, such as a network of mixes, which aims to provide its users (senders and recipients) with anonymous communication. As is common in the literature, we assume the existence of an observer who is able to see some or all of the messages entering and exiting the anonymity system. We also assume that every message the observer is able to see entering the system will be among the messages he is able to see exiting the system, and *vice versa*. That is, there exists a one-to-one relation between inputs to and outputs from the anonymity system.

In our model, an input may be a message or flow entering the anonymity system on one end, and an output may be a message or flow leaving the system at another end. In general, if the system provides perfect anonymity, then any input is equally likely to correspond to any output; however, due to the design of a system, or after a successful attack by an active adversary, some input-output pairings can be rendered infeasible, decreasing the level of anonymity of the system. For example, Danezis & Serjantov consider route-length restrctions in mix network to eliminate some possible sender-recipient pairings [6].

One might also consider a system providing anonymity to its users with some quality-of-service guarantees, implying an upper bound on the latency for each message going through the system. In addition, there is a lower bound on latency due to the time required for processing packets in the system, such as cryptographic operations and routing over several mixes. We note that these assumptions are similar to those made by Tóth and Hornák in their PROB-channel model [9].

B. A Permanent-based Approach

Let the inputs of an anonymity system be denoted by the set $S = \{s_i\}$ and the outputs by $T = \{t_i\}$. Given a set of possible associations between inputs and outputs, we construct a bipartite graph $G = (V_1, V_2, E)$ to represent the system, where $V_1 = S$, $V_2 = T$, and E is the set of edges representing all possible (s_i, t_j) mappings.

While initially defining our metric, we consider a real-time anonymity system that provides a minimum and maximum latency of messages through the system. More precisely, the delay Δ_i for some message m_i going through the anonymity system is bounded by

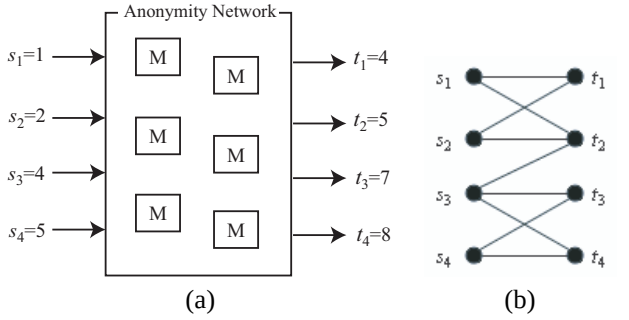

Fig. 1. (a) An example mix network, with the entry and exit times observed for four messages. (b) The corresponding bipartite graph, given that $\Delta_{min} = 1$ and $\Delta_{max} = 4$.

$$\Delta_{min} \leq \Delta_i \leq \Delta_{max},$$

for some given Δ_{min} and Δ_{max}. Even if we treat the internals of the system as a black-box, with the times of messages entering and exiting being the only observable information, we can still obtain a restricted set of possible mappings between inputs and outputs. We can then construct a bipartite graph G, with the entry and exit times of messages constituting the vertices (V_1 and V_2, respectively) of G. That is, each s_i and t_j represent not messages, but rather numeric *timestamps* associated with messages entering or exiting the anonymity system, respectively.

For any s_i and t_j, if $\Delta_{min} \leq t_j - s_i \leq \Delta_{max}$, then there is an edge in G connecting the vertices corresponding to s_i and t_j.

Fig. 1(a) presents an example with 4 messages, each entering the network at time s_i for some i and leaving at time t_j for some j, where $i, j \in \{1, 2, 3, 4\}$. Assuming that, for this system, $\Delta_{min} = 1$ and $\Delta_{max} = 4$ (the time unit used is not relevant), we conclude that a message entering at $s_1 = 1$ will leave the system sometime in the interval $[2, 5]$, resulting in two edges in G; (s_1, t_1) and (s_1, t_2). Continuing in this manner, we can obtain the bipartite graph in Fig. 1(b).

Note that the correct relationship between inputs and outputs corresponds to a perfect matching on the constructed bipartite graph. When the anonymity provided by the system is maximal, we obtain a complete bipartite graph, $G = K_{n,n}$, where n is the number of inputs (or, equivalently, outputs).

Considering all possible permutations of $\{1, 2, ..., n\}$, there are $n!$ ways of mapping a set of n inputs to a set of n outputs, hence $K_{n,n}$ has $n!$ different perfect matchings. Thus, intuitively, a passive adversary observing the system has a $1/n!$ probability of identifying the correct matching. If, on the other hand, G contains a single perfect matching, then no anonymity is provided at all.

Following this intuition, given a bipartite graph G representing the system, we note that the number of perfect matchings in G, combined with the normalization we present later in this section, can provide an indication as to the *strength* of the anonymity system.

A bipartite graph $G = (V_1, V_2, E)$ can be represented by its adjacency matrix A; a (0,1)-matrix of size $n \times n$, where $n = |V_1| = |V_2|$. For each $u \in V_1$ and $v \in V_2$, if the edge (u, v) exists in G, the entry $A(u, v)$ is set to 1, otherwise it is set to 0. It is known that counting the number of perfect matchings in G is equivalent to the *permanent* of A, which is given by

$$per(A) = \sum_{\pi} \prod_{i=1}^{n} A(i, \pi(i)), \tag{1}$$

where the summation is over all permutations of $\{1, 2, ..., n\}$. For a (0,1)-matrix, the summation terms in (1) are either 0 or 1. A term in the summation is 1 if and only if all entries $A(1, \pi(1)), A(2, \pi(2)), ..., A(n, \pi(n))$ are 1, which means that G has a perfect matching $\{(1, \pi(1)), (2, \pi(2)), ..., (n, \pi(n))\}$.

We can assume that there exists at least one perfect matching between inputs and outputs, indicating the true communication pattern. At most, every input potentially corresponds to any output. Thus, the number of perfect matchings in an $n \times n$ (0,1)-matrix A is bounded by $1 \leq per(A) \leq n!$.

C. Definition

We now precisely define our metric based on the matrix permanent. In doing so, we obtain a normalized value in the range $[0, 1]$ representing the anonymity level of the system. A value of 0 means no anonymity is provided, whereas a value of 1 means that the anonymity provided by the system is maximal. This is the same interpretation as used in the normalized entropy-based metric [7].

Given an $n \times n$ (0,1)-matrix A representing possible input-output correlations in an anonymity system, we define the system's anonymity level as

$$d(A) = \begin{cases} 0 & n = 1 \\ \frac{\log(per(A))}{\log(n!)} & n > 1 \end{cases} \tag{2}$$

where $n!$ is the permanent of the all-1 matrix J of size $n \times n$. Note that when $d(A) = 1$, the degree of anonymity of the system is maximal and the corresponding graph G is a complete bipartite graph $K_{n,n}$. On the other hand, when $d(A) = 0$, G has only a single perfect matching and the system provides no anonymity for any participant. For the system in Fig. 1, the level of anonymity can be computed as $\log(4)/\log(24) \approx 0.44$.

While the entropy-based metrics of Serjantov & Danezis [6] and Diaz et al. [7] have intuitive interpretations with and without normalization, we argue that the normalization we presented is necessary with our permanent-based approach.

Consider, again, the example system in Fig. 1(a). If the system had an additional input and output, $s_5 = 9$ and $t_5 = 10$, respectively, then the corresponding bipartite graph would have a single additional edge linking s_5 to t_5. Arguably, this modified system has weaker anonymity properties than the first since there is an edge linking s_5 to t_5 with absolute certainty; however, simply counting the number of perfect matchings

in the bipartite graph does not reflect the overall decrease in anonymity, since the number of perfect matchings is the same for both systems. If we apply our normalized measure, then the anonymity level of the modified system decreases from $\log(4)/\log(24) \approx 0.44$ to $\log(4)/\log(120) \approx 0.29$, correctly indicating that the latter system is less desirable than the former.

III. Generalized Permanent-based Anonymity Metric

We used the permanent of a (0,1)-matrix as a basis to define our permanent-based anonymity metric. Such a metric captures a scenario where one can model the feasibility between inputs and outputs as a (0,1) relation. In general, an observer may also possess probabilistic information about correlations between inputs and output, possibly obtained by observing the inputs and outputs of each individual mix in the system (instead of treating the system as a black box), or by some probabilistic attack. In such a case, one can construct a doubly stochastic matrix, where the sum of each row and column is 1, rather than a simple (0,1)-matrix. An entry (u, v) in this matrix represents the probability that input u is associated with output v.

In this section, we generalize our metric to capture scenarios involving probabilistic information. This generalization has a resemblance to the effective anonymity set size definition [6], except instead of counting the possible senders of a received message, the authors use the probabilities assigned to each possible sender to compute an entropy-based effective anonymity set size.

A. Including Probabilistic Information

When the matrix is no longer a (0,1)-matrix, the permanent does not have the same intuitive meaning as the number of perfect matchings in a bipartite graph. Still, we can show that the permanent of a doubly stochastic matrix can yield insight into the anonymity level of a system when the values of the matrix are probabilities rather than simply 0s and 1s.

Consider a *permutation matrix* P of size $n \times n$, which is a special doubly stochastic matrix, obtained by permuting the rows of the identity matrix I_n according to some permutation of the numbers 1 to n. Each row and column of P has a single nonzero entry, which must be 1. The permanent of P is 1 since there exists a single permutation that yields a non-zero summation term in the permanent. This is, in fact, the maximum value of the permanent of a doubly stochastic matrix.

As the total value becomes more evenly distributed to the other entries in a row or column, rather than on a single value, the value of the permanent decreases. Indeed, when the values are uniformly distributed (that is, when all entries of P are $1/n$) the permanent is minimum and is equal to $\frac{n!}{n^n}$. This lower bound on the permanent of a doubly stochastic matrix is known as the *Van der Waerden conjecture* and was proven by Egorychev [12] and also, independently, by Falikman [13]. Thus, the permanent of a doubly stochastic matrix P is bounded by the inequality $n!/n^n \leq per(P) \leq 1$.

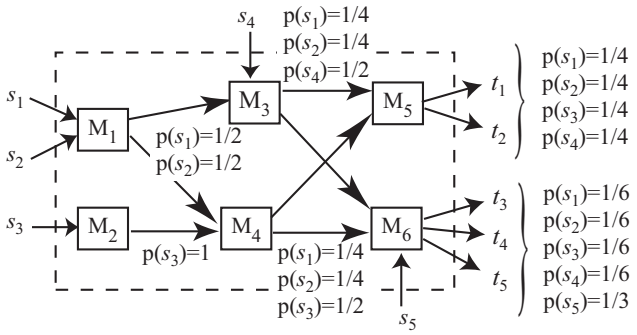

Fig. 2. A sample mix network observed globally. Five messages enter and exit the system, and each message entering a mix is equally likely to follow any outgoing link. The probabilities represent the likelihood of messages being on a particular link.

$$P = \begin{array}{c} \\ s_1 \\ s_2 \\ s_3 \\ s_4 \\ s_5 \end{array} \begin{array}{c} \begin{array}{ccccc} t_1 & t_2 & t_3 & t_4 & t_5 \end{array} \\ \begin{bmatrix} 1/4 & 1/4 & 1/4 & 1/4 & 0 \\ 1/4 & 1/4 & 1/4 & 1/4 & 0 \\ 1/6 & 1/6 & 1/6 & 1/6 & 1/3 \\ 1/6 & 1/6 & 1/6 & 1/6 & 1/3 \\ 1/6 & 1/6 & 1/6 & 1/6 & 1/3 \end{bmatrix} \end{array}$$

Fig. 3. The doubly stochastic matrix P corresponding to the system in Fig. 2. In this example, $per(P) = 0.0417$ and $d(P) = \frac{log(0.0417)}{log(\frac{5!}{5^5})} = 0.9747$

In our initial model using a (0,1)-matrix A, a greater value of $per(A)$ indicated a higher level of anonymity. In contrast, the level of anonymity of a system represented by a doubly stochastic matrix P is highest when its permanent is minimum (that is, when the probabilities in P are uniformly distributed). It is worth noting that in both the basic model using a (0,1)-matrix A and the probabilistic model using a doubly stochastic matrix P, the anonymity level of both systems system is minimum when $per(A) = per(P) = 1$, indicating a single perfect matching.

B. Definition

Given a doubly stochastic matrix P representing the probabilities of input-output relationships in an anonymity system, we define the degree of anonymity, d, as

$$d(P) = \begin{cases} 0 & n = 1 \\ \frac{\log(per(P))}{\log(\frac{n!}{n^n})} & n > 1 \end{cases} \quad (3)$$

where we recall that $n!/n^n$ is the minimum value of the permanent of an $n \times n$ doubly stochastic matrix, by the Van der Waerden conjecture.

Fig. 2 gives a simple example of how to obtain probabilities for mapping inputs to outputs in an anonymity network composed of several mixes. Fig. 3 presents the doubly stochastic matrix corresponding to the example system. Again, we use the same notation s_i for inputs and t_i for outputs, but note that the inputs and outputs do not represent the time instants in this scenario. Rather, they are simply used to label the incoming and outgoing messages to and from the anonymity system.

The anonymity system in Fig. 2 follows the scheme presented by Serjantov & Danezis [6], in that each mix in the system is equally likely to have sent an incoming message on any one of its outgoing links. A global observer can then form probability distributions on the links that represent the likelihood of finding a particular message on a specific link.

The probability distribution on an outgoing link of a mix is obtained by adding all distributions on the incoming links of that mix, and dividing by the number of outgoing links. For example, mix M_4 in Fig. 2 has two input links, with probability distributions $\{(s_1, 1/2), (s_2, 1/2)\}$ and $\{(s_3, 1)\}$, so the probability distribution on each of its outgoing links becomes $\{(s_1, 1/4), (s_2, 1/4), (s_3, 1/2)\}$. Following in this manner, we can obtain the probabilities of relating a particular input to a particular output and construct the doubly stochastic matrix P corresponding to this system, which is given in Fig. 3. Given the matrix P, the degree of anonymity $d(P)$ can be computed as in (3) and is approximately $\log(0.042)/\log(0.038) \approx 0.97$.

IV. Application to Some Common Mix Types

When defining our permanent-based anonymity metric in Section II, we used the example of a real-time anonymity network. We now extend our analysis and apply our permanent-based metric to some common types of high-latency mixes from the literature.

A. Threshold Mixes

A *threshold mix* is a mix that collects incoming messages until it has received some threshold N, applies a cryptographic transformation to each message received, and then forwards all N messages on to their next destination in a random order. We refer to each time the mix purges its store of messages as a mix *round*.

For an individual round, each input message is equally likely to correspond to any of the N output messages; however, messages within a particular round do not "blend" with messages from any previous or subsequent round. Fig. 4(a) shows the bipartite graph of possible relationships between inputs to and outputs from two rounds of a threshold mix with $N = 3$. After one round, the graph is a complete bipartite graph $G = K_{3,3}$. The permanent of the graph's corresponding adjacency matrix is $3! = 6$ and yields an optimal degree of anonymity of $d = \log(3!)/\log(3!) = 1.0$.

After the second round, each message is blending with a smaller fraction (one-half, to be precise) of the total messages observed exiting the mix. The corresponding bipartite graph G is composed of two smaller complete bipartite graphs where $G = K_{3,3} \cup K_{3,3}$. The permanent of G's adjacency matrix representation shown in Fig. 4(b), is thus $3! \times 3!$. The overall degree of anonymity decreases from $d = 1.0$ to $d = \log(3! \times 3!)/\log(6!) = 0.5447$.

Extending this example to any threshold mix with a threshold N, over r rounds we can compute the overall degree of anonymity of the mix as follows:

$$d(A) = \frac{\log(per(A))}{\log(n!)} = \frac{\log((N!)^r)}{\log((N \times r)!)}$$

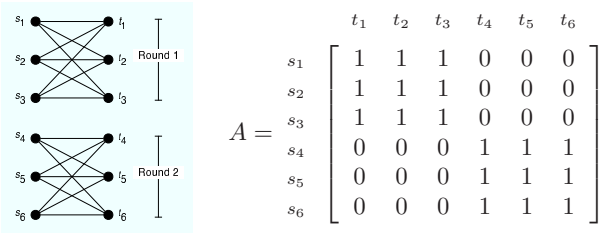

$$A = \begin{array}{c} \\ s_1 \\ s_2 \\ s_3 \\ s_4 \\ s_5 \\ s_6 \end{array} \begin{array}{c} \begin{array}{cccccc} t_1 & t_2 & t_3 & t_4 & t_5 & t_6 \end{array} \\ \begin{bmatrix} 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 1 & 1 & 1 \end{bmatrix} \end{array}$$

(a) (b)

Fig. 4. (a) Graph G of possible sender-recipient relationships over two rounds of a threshold mix with a threshold $N = 3$. (b) Adjacency matrix A corresponding to two rounds of the threshold mix. In this example, $per(A) = 36$ and $d(A) = 0.5447$.

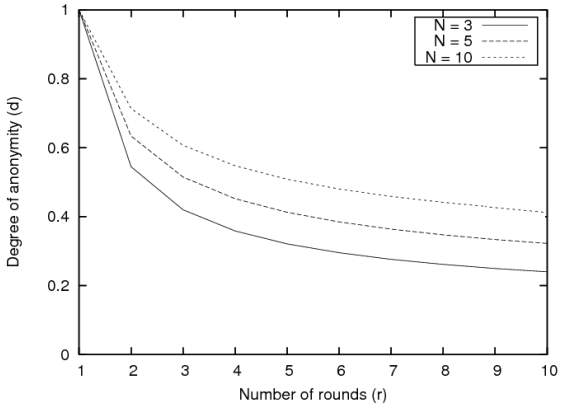

Fig. 5. Degree of anonymity for threshold mixes with $N = 3, 5$, and 10 over ten rounds.

The graph in Fig. 5 shows how the degree of anonymity changes over r rounds for various threshold values.

B. Timed Mixes

A *timed mix* is a mix that collects messages for t seconds, applies a cryptographic transformation to each message received, and then forwards all messages on to their next destination in a random order. Since there is no longer a constant number, or threshold, of messages that enter and exit the mix during each round, the degree of anonymity of a timed mix depends on both the *duration* of each mix round and the *arrival rate* of messages into the mix.

Let the arrival rate of messages into the mix be δ messages per second. The average number of messages arriving to and exiting from the mix during a single round of length t is then $\delta \times t$. Following in the same manner as we did for threshold mixes, we can express the degree of anonymity for a timed mix with parameters δ and t over r rounds as

$$d(A) = \frac{\log(per(A))}{\log(n!)} = \frac{\log(((\delta t)!)^r)}{\log((\delta t r)!)}.$$

From the previous equation, it is easy to see that a threshold mix with a threshold of N is equivalent to a timed mix with a round length of t seconds when $N = \delta \times t$. We can use this analysis to select parameters of threshold or timed mixes to provide a desired degree of anonymity.

Consider a timed mix that has an average message arrival rate of $\delta = 1/6$ messages per second. If we have a threshold mix with a threshold of $N = 10$ messages, we would require the timed mix to have a round length of at least $t = 60$ seconds to provide the same minimum degree of anonymity as the threshold mix. Similar arguments can be made for other values of N, δ, and t.

C. Pool Mixes

A *pool mix* is a mix that, at end of each round, randomly selects a subset of the messages inside the mix to cryptographically transform and forward to their next destination. The remaining messages will be retained internally for the next round. Pool mixes may determine the length of a single round according to a threshold or timed algorithm, as above, or some combination of the two.

We let N be the number of inputs into the mix in a single round. The mix will then have N outputs selected randomly from a pool of $N + n$ messages inside the mix, where n is the number of messages retained in the mix at every round. Initially, the pool mix is primed with n "dummy messages" created by the mix itself that are indistinguishable from authentic messages.

Pool mixes are able to blend messages across multiple rounds of the mix, unlike simple threshold or timed mixes that only blend together messages within the same round. Indeed, a message exiting a pool mix may correspond to *any* previous message that has ever entered the mix. There is a non-zero possibility that a message will remain in the mix indefinitely; however, the probability of a message remaining in the mix for r rounds decreases as r increases. Serjantov & Danezis showed in [6] that a message exiting the pool mix at round r corresponds to a message that previously entered the mix at round $0 < x \leq r$ with a probability of

$$p(r, x) = \frac{N}{N+n} \left(\frac{n}{N+n} \right)^{r-x},$$

where, again, N is the number of messages that enter the mix every round and n is the number of messages retained in the mix each round. Each individual output from the mix at round r then has a probability of $p(r, x)/N$ of corresponding to each input at some previous round x.

For the sake of example, we will consider a threshold pool mix with a threshold of $N = 2$ and pool size of $n = 2$. Initially, the pool mix creates two dummy messages and places them in its pool. After receiving two more input messages, the mix will randomly select two messages to forward to their next destinations and retain the other two messages in the mix.

Let us consider message s_1. At the end of the first round, s_1 has a $1/2$ chance of being forwarded to its next destination and a $1/2$ chance of being retained in the mix ($1/4$ for each output or position in the pool). After the second round, s_1 has a $p(2, 1)/N = 1/8$ chance of corresponding to each of the round's two outputs, t_3 and t_4.

Continuing in this manner through two mix rounds, we can populate a doubly stochastic matrix P as shown in Fig. 6.

$$P = \begin{array}{c} \\ s_1 \\ s_2 \\ s_3 \\ s_4 \\ p_1 \\ p_2 \end{array} \begin{array}{c} \begin{array}{cccccc} t_1 & t_2 & t_3 & t_4 & {p'}_1 & {p'}_2 \end{array} \\ \begin{bmatrix} 1/4 & 1/4 & 1/8 & 1/8 & 1/8 & 1/8 \\ 1/4 & 1/4 & 1/8 & 1/8 & 1/8 & 1/8 \\ 0 & 0 & 1/4 & 1/4 & 1/4 & 1/4 \\ 0 & 0 & 1/4 & 1/4 & 1/4 & 1/4 \\ 1/4 & 1/4 & 1/8 & 1/8 & 1/8 & 1/8 \\ 1/4 & 1/4 & 1/8 & 1/8 & 1/8 & 1/8 \end{bmatrix} \end{array}$$

Fig. 6. Doubly stochastic matrix corresponding to two rounds a pool mix with $N = 2$ and $n = 2$. Each p_i is a message in the pool at the start of the mix round and each ${p'}_i$ is a message in the pool after the mix fires.

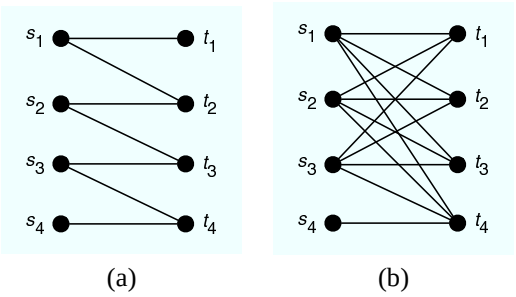

(a) (b)

Fig. 7. (a) A system that provides no anonymity for any messages, from a global perspective. (b) Another system that provides a reasonable degree of anonymity for all but one message sent.

From P, we can calculate the degree of anonymity of the example pool mix as

$$d(P) = \frac{\log(per(P))}{\log(\frac{m!}{m^m})} = \frac{\log(0.0176)}{\log(0.0154)} \approx 0.9688,$$

where $m = r \times N + n$ is the number of rows and columns in the $m \times m$ matrix P. A similar analysis can be made for other values of r, N, and n.

V. Comparison to Existing Measures of Anonymity

In this section we highlight a few of the more important differences between our proposed metric and some of the current metrics. Given that the more common existing metrics measure anonymity specific to a particular message or user, whereas our metric is a system-wide metric, direct quantitative comparisons between metrics are not especially meaningful. Instead, we will show how our permanent-based metric can be used to identify properties of a system that other metrics might not.

A. Information-theoretic Metrics

While individual users of an anonymity system are most likely interested in only their own level of anonymity, such a narrow focus can overlook important properties of a system as a whole. Consider if an adversary is able to determine through passive observation the possible links in the two systems represented as bipartite graphs in Fig. 7. Using the simple entropy-based metric, we can compute an effective anonymity set size for each input and output. In Fig. 7(a), s_4 and t_1 have an effective anonymity set size of $S = -(1 \times \log(1)) = 0$, while the rest have an effective anonymity set size of $S =$

$-(2 \times (1/2\log(1/2))) = 1$. The normalized entropy-based metric also yields a non-zero degree of anonymity for all but two nodes in the first system.

A clever adversary would be able to further his analysis in Fig. 7(a) and eliminate many additional links by working backwards from vertices with a degree of one. In our first example system, an adversary could eliminate s_1 from t_2's sender anonymity set, since t_1 is positively linked to s_1. Indeed, when we apply our system-wide metric, we see that the bipartite graph has only a single perfect matching, thus identifying the true communication pattern of the system.

While the entropy-based metrics indicate that, from the perspective of a single message sent or received, many messages have a non-zero degree of anonymity, the system in Fig. 7(a) in fact provides *no* anonymity for *any* message sent or received. Our system-wide metric, though, correctly identifies this weakness in the system.

B. Source- and Destination-hiding Metrics

The notions of *source-hiding* and *destination-hiding* are measures of the highest probability an adversary is able to assign to a link between a particular message and its sender or recipient. In that sense, they are a "worst-case" measure of a system.

In both example systems given in Fig. 7, the source- and destination-hiding metrics would identify that the given systems provide no anonymity for some communicants. That is, the two systems are both source-hiding with parameter $\Theta = 1.0$ and destination-hiding with parameter $\Omega = 1.0$, meaning there are one or more senders positively linked to their recipients and *vice versa*. The two example systems, however, clearly provide different overall levels of anonymity. Some of the messages in Fig. 7(b) even have a nearly optimal level of anonymity.

With our permanent-based metric, we are interested in the maximum probabilities an adversary is able to assign to *an entire set* of matchings between senders and recipients (a maximum weight perfect matching), instead of individual links, indicating the most likely true communication pattern in the system. The authors of [9] refer to this as *global back-tracing*, which they dismiss as inefficient as it is exponential in the number of sent messages.

With our permanent-based metric, global back-tracing would be equivalent to finding a maximum weight perfect matching in a bipartite graph with probabilities assigned to each edge in the graph. Using an algorithm such as the Hungarian method [14], we can more easily find a single maximum weight perfect matching in the weighted bipartite graph in $O(n^3)$. Enumerating all perfect matchings is still, of course, exponential.

VI. Bounding the Anonymity

The anonymity metric we defined in Sections II and III is based on computing the permanent of a matrix. While it is NP-hard to compute the permanent of a matrix [15], even if the entries are 0-1, there has been an on-going and promising research effort to develop efficient approximations and bounds on the permanent of 0-1 and real matrices [16], [17].

Jerrum et al. have given a fully polynomial randomized approximation algorithm, which provides an arbitrarily close approximation, in time polynomially dependent on n and the desired error [18]. Their algorithm is based on an almost uniform sampling of perfect matchings using a Markov chain Monte Carlo method. Bezakova et al. [19] improve this running time to $O(n^7 \log^4 n)$ by using a new "cooling schedule" for the simulated annealing algorithm running on top of the Markov chain. Still, these algorithms are impractical for large systems, since the degrees of the polynomial running times are too large.

We now describe an approach to obtaining bounds on the degree of anonymity of a given system by using known, easy-to-compute inequalities relating to the permanent of a matrix, which avoid the described computational complexity of computing the exact or approximated value of the permanent.

A. A Simple Bound

When defining our permanent-based metric in Section II, we argued that the greater the number of possible sets of matchings between inputs to and outputs from the anonymity system, the harder it is for an observer to correctly identify the correct relationship between senders and recipients. Subsequently, we can obtain a lower bound on the anonymity level of a system by determining the minimum number of possible perfect matchings in the system's corresponding bipartite graph.

The best known general lower bound for the permanent of a (0,1)-matrix was given by Ostrand [20] and is an improvement on an earlier lower bound due to Jurkat and Ryser [21]. Ostrand proved

$$per(A) \geq \prod_{i=1}^{n} max\{1, r_i - i + 1\}, \tag{4}$$

where A is a (0,1)-matrix with a nonzero permanent and $r_1, r_2, ..., r_n$ are row sums of A arranged such that $r_1 \leq r_2 \leq ... \leq r_n$.

In our model, we can assume that the permanent of a (0,1)-matrix representing possible sender-recipient relationships is greater than 0, since there must exist at least one perfect matching (the actual communication pattern). It is also possible to permute the rows of A such that the row sums $\{r_i\}$ satisfy $r_1 \leq r_2 \leq ... \leq r_n$, without distorting the sender-recipient relationships represented by the matrix. Further, the permanent of a matrix is invariant under a permutation of its rows [22]. Thus, we can apply Ostrand's lower bound to our model without loss of generality.

Let $p_{min}(A)$ be the minimum possible number of perfect matchings in a (0,1)-matrix A given by (4). Substituting $p_{min}(A)$ into our original definition of the degree of anonymity of a system in (2), we obtain the inequality

$$d(A) \geq \frac{\log(p_{min}(A))}{\log(n!)},$$

which we interpret as the minimum degree of anonymity of the system represented by A.

B. Generalizing the Bound

In Section III, we considered instances where an observer can assign probabilities to relationships between senders and recipients and then form a doubly stochastic matrix of those probabilities. When the entries of the matrix are probabilities instead of simply 0s and 1s, we can no longer use the bounds presented above.

Minc has shown a lower bound on the permanent of nonnegative matrices [23], an improvement of previous work done by Jurkat and Ryser [24]. Minc proved

$$per(A) \leq \prod_{i=1}^{n} \sum_{t=1}^{i} a_{it}^{*} + (r_1 - n a_{11}^{*}) \prod_{j=2}^{n} \sum_{s=1}^{j-1} a_{js}^{'}, \qquad (5)$$

where $(a_{i1}^{*}, ..., a_{in}^{*})$ is an n-tuple of A representing the ith row of A arranged in nonincreasing order, $(a_{i1}^{'}, ..., a_{in}^{'})$ is an n-tuple of A representing the ith row of A arranged in nondecreasing order, r_1 is the row sum of the first row of A, and $a_{11}^{'}$ and a_{11}^{*} are the least and greatest values, respectively, in the first row of A.

We showed in Section III that a higher value for the permanent of a doubly stochastic matrix indicates a lower degree of anonymity. Continuing with this logic, we let $p_{max}(P)$ be the maximum value of the permanent of a doubly stochastic matrix P, according to (5). Substituting $p_{max}(P)$ into (3), we have the inequality

$$d(P) \geq \frac{\log(p_{max}(P))}{\log(\frac{n!}{n^n})},$$

which we interpret as the minimum possible degree of anonymity yielded by a system in the presence of an observer who can assign probabilities to relationships between senders and recipients.

VII. Conclusion

We have described a new approach to quantifying the degree of anonymity of a system as a whole by modeling the possible associations between inputs to and outputs from the system as a bipartite graph and using the permanent of the graph's adjacency matrix to derive a degree of anonymity for that system. We have also shown how our permanent-based anonymity metric can be generalized to include probabilistic information determined by an attacker.

Our work is intended to complement the existing measures of anonymity that have been developed. Though users of a particular anonymity system are most likely interested in only their own level of anonymity, we have shown that our combinatorial approach can help identify properties of a system that other metrics might not.

References

[1] D. Chaum, "Untraceable electronic mail, return addresses, and digital pseudonyms," *Communications of the ACM*, vol. 4, no. 2, February 1981.

[2] G. Danezis, R. Dingledine, and N. Mathewson, "Mixminion: Design of a Type III Anonymous Remailer Protocol," in *Proceedings of the 2003 IEEE Symposium on Security and Privacy*, May 2003.

[3] R. Dingledine, N. Mathewson, and P. Syverson, "Tor: The second-generation onion router," in *Proceedings of the 13th USENIX Security Symposium*, August 2004.

[4] D. Chaum, "The dining cryptographers problem: Unconditional sender and recipient untraceability," *Journal of Cryptology*, vol. 1, pp. 65–75, 1988.

[5] D. Kesdogan, J. Egner, and R. Büschkes, "Stop-and-go MIXes: Providing probabilistic anonymity in an open system," in *Proceedings of Information Hiding Workshop (IH 1998)*. Springer-Verlag, LNCS 1525, 1998.

[6] A. Serjantov and G. Danezis, "Towards an information theoretic metric for anonymity," in *Proceedings of Privacy Enhancing Technologies Workshop (PET 2002)*, R. Dingledine and P. Syverson, Eds. Springer-Verlag, LNCS 2482, April 2002.

[7] C. Díaz, S. Seys, J. Claessens, and B. Preneel, "Towards measuring anonymity," in *Proceedings of Privacy Enhancing Technologies Workshop (PET 2002)*, R. Dingledine and P. Syverson, Eds. Springer-Verlag, LNCS 2482, April 2002.

[8] M. Reiter and A. Rubin, "Crowds: Anonymity for web transactions," *ACM Transactions on Information and System Security*, vol. 1, no. 1, June 1998.

[9] G. Tóth and Z. Hornák, "Measuring anonymity in a non-adaptive, real-time system," in *Proceedings of Privacy Enhancing Technologies workshop (PET 2004)*, ser. Springer-Verlag, LNCS, vol. 3424, 2004, pp. 226–241.

[10] G. Tóth, Z. Hornák, and F. Vajda, "Measuring anonymity revisited," in *Proceedings of the Ninth Nordic Workshop on Secure IT Systems*, S. Liimatainen and T. Virtanen, Eds., Espoo, Finland, November 2004, pp. 85–90.

[11] R. E. Newman, I. S. Moskowitz, P. Syverson, and A. Serjantov, "Metrics for traffic analysis prevention," in *Proceedings of Privacy Enhancing Technologies workshop (PET 2003)*, R. Dingledine, Ed. Springer-Verlag, LNCS 2760, March 2003.

[12] G. P. Egorychev, "The solution of van der waerden's problem for permanents," *Advances in Mathematics*, vol. 42, pp. 299–305, 1981.

[13] D. I. Falikman, "Proof of the van der waerden's conjecture on the permanent of a doubly stochastic matrix," *Mat. Zametki*, vol. 29, no. 6, pp. 931–938, 957, 1981, in Russian.

[14] H. W. Kuhn, "The Hungarian method for the assignment problem," *Naval Research Logistics Quarterly*, vol. 2, no. 1, pp. 83–98, 1955.

[15] L. G. Valiant, "The complexity of computing the permanent," *Theoretical Computer Science*, vol. 8, no. 2, pp. 189–201, April 1979.

[16] R. A. Servedio and A. Wan, "Computing sparse permanents faster," *Inf. Process. Lett.*, vol. 96, no. 3, pp. 89–92, 2005.

[17] P. Sankowski, "Alternative algorithms for counting all matchings in graphs," vol. 2607. Springer Berlin / Heidelberg, March 2003, pp. 427–438.

[18] M. Jerrum, A. Sinclair, and E. Vigoda, "A polynomial-time approximation algorithm for the permanent of a matrix with nonnegative entries," *J. ACM*, vol. 51, no. 4, pp. 671–697, July 2004.

[19] I. Bezakova, D. Stefankovic, V. V. Vazirani, and E. Vigoda, "Accelerating simulated annealing for the permanent and combinatorial counting problems," in *SODA '06: Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*. ACM Press, 2006, pp. 900–907.

[20] P. A. Ostrand, "Systems of distinct representatives ii," *Journal of Mathematics Analysis and Applications*, vol. 32, pp. 1–4, 1970.

[21] W. B. Jurkat and H. J. Ryser, "Matrix factorizations of determinants and permanents," *Journal of Algebra*, vol. 3, pp. 1–27, 1966.

[22] R. A. Brualdi and H. J. Ryser, *Combinatorial Matrix Theory*, ser. Encyclopedia of Mathematics, 1991.

[23] H. Minc, "Bounds for permanents of nonnegative matrices," in *Proceedings of the Edinburgh Mathematical Society*, vol. 16, 1968, pp. 233–237.

[24] W. B. Jurkat and H. J. Ryser, "Term ranks and permanents of nonnegative matrices," *Journal of Algebra*, vol. 5, pp. 342–357, 1967.