Anonymity and Covert Channels in Simple Timed Mix-firewalls*

Richard E. Newman[1], Vipan R. Nalla[1], and Ira S. Moskowitz[2]

[1] CISE Department
University of Florida
Gainesville, FL 32611-6120
{nemo,vreddy}@cise.ufl.edu
&
[2] Center for High Assurance Computer Systems, Code 5540
Naval Research Laboratory
Washington, DC 20375
moskowitz@itd.nrl.navy.mil

Abstract. Traditional methods for evaluating the amount of anonymity afforded by various Mix configurations have depended on either measuring the size of the set of possible senders of a particular message (the anonymity set size), or by measuring the entropy associated with the probability distribution of the messages possible senders. This paper explores further an alternative way of assessing the anonymity of a Mix system by considering the capacity of a covert channel from a sender behind the Mix to an observer of the Mix's output.

Initial work considered a simple model [5], with an observer (Eve) restricted to counting the number of messages leaving a Mix configured as a firewall guarding an enclave with one malicious sender (Alice) and some other naive senders (Clueless$_i$'s). Here, we consider the case where Eve can distinguish between multiple destinations, and the senders can select to which destination their message (if any) is sent each clock tick.

1 Introduction

In [5] the idea of measuring the lack of perfect anonymity (quasi-anonymity) via a covert channel was initiated. This idea was formalized in [6]. Our concern in this paper is to identify, and to calculate the capacity of, the covert channels that arise from the use of a Mix [1, 8] as an exit firewall from a private enclave (as briefly addressed in [5, Sec. 4].) In general, we refer to a covert channel that arises, due to a state of quasi-anonymity, as a quasi-anonymous channel [6]. The quasi-anonymous channel also serves the dual role of being a measure of the lack of perfect anonymity. [2] uses a similar model for statistical attacks in which Eve correlates senders' actions with observed output.

* Research supported by the Office of Naval Research.

2 Exit Mix-firewall Model

There are $N+1$ senders in a private enclave. Messages pass one way from the private enclave to a set of M receivers. The private enclave is behind a firewall which also functions as a timed Mix [8] that fires every tick, t, hence we call it a simple timed Mix-firewall. For the sake of simplicity we will refer to a simple timed Mix-firewall as a Mix-firewall in this paper. One of the $N+1$ senders, called Alice, is malicious. The other N clueless senders, $\text{Clueless}_i, i = 1, \ldots, N$, are benign. Each sender may send at most one message per unit time t to the set of receivers. All messages from the private enclave to the set of receivers pass through public lines that are subject to eavesdropping by an eavesdropper called Eve. The only action that Eve can take is to count the number of messages per t going from the Mix-firewall to each receiver, since the messages are otherwise indistinguishable. Eve knows that there are $N+1$ possible senders. The N clueless senders act in an independent and identical manner (i.i.d.) according to a fixed distribution $C_i, i = 1, \ldots, N$. Alice, by sending or not sending a message each t to at most one receiver, affects Eve's message counts. This is how Alice covertly communicates with Eve via a quasi-anonymous channel [6].

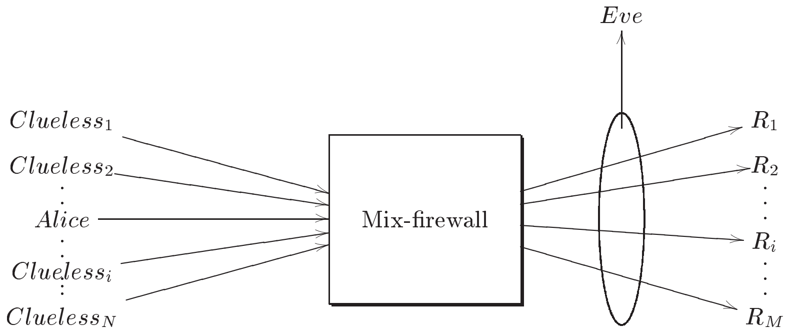

Fig. 1. Exit Mix-firewall model with N clueless senders and M distinguishable receivers

Alice acts independently (through ignorance of the clueless senders) when deciding to send a message; we call this the *ignorance assumption*. Alice has the same distribution each t. Between Alice and the N clueless senders, there are $N+1$ possible senders per t, and there are $M+1$ possible actions per sender (each sender may or may not transmit, and if it does transmit, it transmits to exactly one of M receivers).

We consider Alice to be the input to the quasi-anonymous channel, which is a proper communications channel [9]. Alice can send to one of the M receivers or not send a message. Thus, we represent the inputs to the quasi-anonymous channel by the $M+1$ input symbols $0, 1, \ldots, M$, where $i = 0$ represents Alice not sending a message, and $i \in \{1, \ldots, M\}$ represents Alice sending a message to the ith receiver R_i. The "receiver" in the quasi-anonymous channel is Eve. Eve receives the output symbols $e_j, j = 1, \ldots, K$. Eve receives e_1 if no sender sends a message. The other output symbols correspond to all the different ways the $N+1$ senders can send or not send, at most one message each, out of the private enclave, provided at least one sender does send a message.

For the sake of simplicity we introduce a dummy receiver R_0 (not shown above). If a sender does not send a message we consider that to be a "message" to R_0. For $N+1$ senders and M receivers, the output symbol e_j observed by Eve is an $M+1$ vector $\langle a_0^j, a_1^j, ..., a_M^j \rangle$, where a_i^j is how many messages the Mix-firewall sends to R_i. Of course it follows that $\sum_{i=0}^{M} a_i^j = N+1$.

The quasi-anonymous channel that we have been describing is a discrete memoryless channel (DMC). We define the channel matrix M as an $(M+1) \times K$ matrix, where M$[i,j]$ represents the conditional probability that Eve observes the output symbol e_j given that Alice input i. We model the clueless senders according to the i.i.d. C_i for each period of possible action t:

$$P(Clueless_i\ doesn't\ send\ a\ message) = p$$
$$P(Clueless_i\ sends\ a\ message\ to\ any\ receiver) = \frac{q}{M} = \frac{1-p}{M}$$

where in keeping with previous papers, $q = 1-p$ is the probability that Clueless$_i$ sends a message to any one of the M receivers. When Clueless$_i$ *does* send a message, the destination is uniformly distributed over the receivers $R_1, \ldots, R_M$. We call this the **semi-uniformity assumption**. Again, keep in mind that each clueless sender has the same distribution each t, but they all act independently of each other.

We model Alice according to the following distribution each t:

$$P(Alice\ sends\ a\ message\ to\ R_i) = x_i$$

Of course, this tells us that

$$x_0 = P(Alice\ doesn't\ send\ a\ message) = 1 - \sum_{i=1}^{M} x_i\ .$$

We let A represent the distribution for Alice's input behavior, and we denote by E the distribution of the output that Eve receives. Thus, the channel matrix M along with the distribution A totally determine the quasi-anonymous channel. This is because the elements of M take the distributions C_i into account, and M and A let one determine the distribution describing the outputs that Eve receives, $P(\text{Eve receives } e_j)$.

Now that we have our set-up behind our exit Mix-firewall model, we may now go on to analyze various cases in detail. Additional cases and more detail are available in [7].

3 Capacity Analyses of the Exit Mix-firewall Model

The mathematics of the problem gets quite complex. Therefore, we start with some simple special cases before attempting to analyze the problem in general.

The mutual information between A and E is given by

$$I(A,E) = H(A) - H(A|E) = H(E) - H(E|A) = I(E,A).$$

The capacity of the quasi-anonymous channel is given by [9]

$$C = \max_A I(A,E)\ ,$$

where the maximization is over the different possible values that the x_i may take (of course, the x_i are still constrained to represent a probability distribution). Recall $\mathrm{M}[i,j] = P(E = e_j|A = i)$, where $\mathrm{M}[i,j]$ is the entry in the i^{th} row and j^{th} column of the channel matrix, M. To distinguish the various channel matrices, we will adopt the notation that $\mathrm{M}_{N.M}$ is the channel matrix for N clueless senders and M receivers.

3.1 One Receiver ($M = 1$)

Case 1 — No clueless senders and one receiver ($N = 0,\ M = 1$)
Alice is the only sender, and there is only one receiver R_1. Alice sends either 0 (by not sending a message) or 1 (by sending a message). Eve receives either $e_1 = \langle 1,0\rangle$ (Alice did nothing) or $e_2 = \langle 0,1\rangle$ (Alice sent a message to the receiver). Since there is no noise (there are no clueless senders) the channel matrix M is the 2×2 identity matrix and it trivially follows that $P(E = e_1) = x_0$, and that $P(E = e_2) = x_1$.

$$\mathrm{M}_{0.1} = \begin{array}{c} \\ 0 \\ 1 \end{array} \begin{array}{c} \begin{array}{cc} e_1 & e_2 \end{array} \\ \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \end{array}$$

Since $x_0 = 1 - x_1$, we see that[3] $H(E) = -x_0 \log x_0 - (1-x_0)\log(1-x_0)$. The channel matrix is an identity matrix, so the conditional probability distribution $P(E|A)$ is made up of zeroes and ones, therefore $H(E|A)$ is identically zero. Hence, the capacity is the maximum over x_0 of $H(E)$, which is easily seen to be unity[4] (and occurs when $x_0 = 1/2$). Of course, we could have obtained this capacity[5] without appealing to mutual information since we can noiselessly send one bit per tick, but we wish to study the non-trivial cases and use this as a starting point.

Case 2 — N clueless senders and one receiver ($M = 1$)
This case reduces to the *indistinguishable receivers* case with N senders. This is the situation analyzed in [5] with both an exit Mix-firewall that we have been discussing and an entry Mix-firewall, with the receivers behind the latter. Alice can either send or not send a message, so the input alphabet again has two symbols. Eve observes $N + 2$ possible output symbols. That is, Eve sees $e_1 = \langle N+1, 0\rangle$, $e_2 = \langle N, 1\rangle$, $e_3 = \langle N-1, 2\rangle$, $\cdots$, $e_{N+2} = \langle 0, N+1\rangle$. A detailed discussion of this case can be found in [5].

[3] All logarithms are base 2.
[4] The units of capacity are bits per tick t, but we will take the units as being understood for the rest of the paper. Note that all symbols take one t to pass through the channel.
[5] This uses Shannon's [9] asymptotic definition of capacity, which is equivalent for noiseless channels (in units of bits per symbol).

3.2 Some Special Cases for Two Receivers ($M = 2$)

There are two possible receivers. Eve has knowledge of the network traffic, so Alice can signal Eve with an alphabet of three symbols: 1 or 2, if Alice transmits to R_1 or R_2, respectively, or the symbol 0 for not sending a message. Let us analyze the channel matrices and the entropies for different cases of senders.

The symbol e_j that Eve receives is an 3-tuple of the form $\langle a_0^j, a_1^j, a_2^j \rangle$, where a_i^j is the number of messages received by i^{th} receiver.[6] The index $i = 0$ stands for Alice not sending any message. The elements of the 3-tuple must sum to the total number of senders, $N + 1$,

$$\sum_{i=0}^{2} a_i = N + 1 \ .$$

Case 3 — No clueless senders and two receivers ($N = 0,\ M = 2$)
Alice is the only sender and can send messages to two possible receivers. The channel matrix is trivial and there is no anonymity in the channel.

$$\mathrm{M}_{0.2} = \begin{array}{c} \\ 0 \\ 1 \\ 2 \end{array} \begin{array}{c} \begin{array}{ccc} \langle 1,0,0 \rangle & \langle 0,1,0 \rangle & \langle 0,0,1 \rangle \end{array} \\ \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \end{array}$$

The subscript 0.2 represents one sender (Alice alone) and two receivers. The 3×3 channel matrix $\mathrm{M}_{0.2}[i, j]$ represents the conditional probability of Eve receiving the symbol e_j, when Alice sends to the Receiver i. '0' stands for not sending a message.

The mutual information I is given by the entropy $H(E)$ describing Eve

$$I(E, A) = H(E) = -x_1 \log x_1 - x_2 \log x_2 - (1 - x_1 - x_2) \log(1 - x_1 - x_2).$$

The capacity of this noiseless covert channel is $\log 3 \approx 1.58$ (at x_i=1/3, $i = 0, 1, 2$). This is the maximum capacity, which we note corresponds to zero anonymity.

Case 4 — $N = 1$ clueless sender and $M = 2$ receivers

There are only six symbols that Eve may receive since there are six ways to put two indistinguishable balls into three distinct urns.

Let us consider the channel matrix.

$$\mathrm{M}_{1.2} = \begin{array}{c} \\ 0 \\ 1 \\ 2 \end{array} \begin{array}{c} \begin{array}{cccccc} \langle 2,0,0 \rangle & \langle 1,1,0 \rangle & \langle 1,0,1 \rangle & \langle 0,2,0 \rangle & \langle 0,1,1 \rangle & \langle 0,0,2 \rangle \end{array} \\ \begin{pmatrix} p & q/2 & q/2 & 0 & 0 & 0 \\ 0 & p & 0 & q/2 & q/2 & 0 \\ 0 & 0 & p & 0 & q/2 & q/2 \end{pmatrix} \end{array}$$

[6] Recall that the a_i^j's of the output symbol are not directly related to A, which denotes the distribution of Alice.

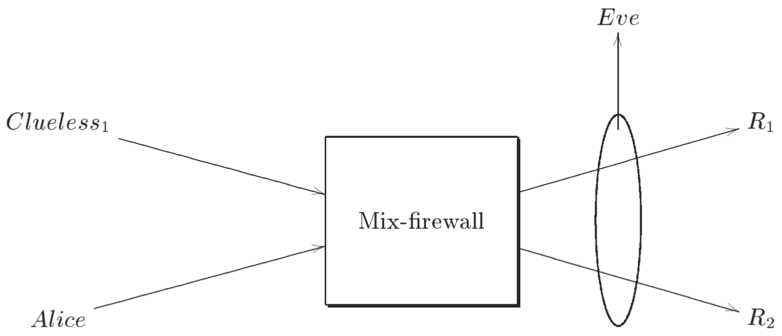

Fig. 2. Case 4: system with $N = 1$ clueless sender and $M = 2$ receivers

The 3×6 channel matrix $\mathrm{M}_{1.2}[i, j]$ represents the conditional probability of Eve receiving the symbol e_j when Alice sends to R_i. As noted, the dummy receiver R_0 corresponds to Alice not sending to any receiver (however this is still a transmission to Eve via the quasi-anonymous channel).

Given the above channel matrix we have:

$$
\begin{aligned}
H(E) = -\{&px_0 \log[px_0] \\
&+[qx_0/2 + px_1] \log[qx_0/2 + px_1] \\
&+[qx_0/2 + px_2] \log[qx_0/2 + px_2] \\
&+[qx_1/2] \log[qx_1/2] + [qx_1/2 + qx_2/2] \log[qx_1/2 + qx_2/2] \\
&+[qx_2/2] \log[qx_2/2]\}.
\end{aligned}
$$

The conditional entropy is given by

$$
H(E|A) = -\sum_{i=0}^{2} \left[p(x_i) \sum_{j=1}^{6} p(e_j|x_i) \log p(e_j|x_i) \right] = h_2(p) \ ,
$$

where $h_2(p)$ denotes the function

$$
\begin{aligned}
h_2(p) \ = \ & -(1-p)/2 \log((1-p)/2) - (1-p)/2 \log((1-p)/2) - p \log p \\
& = -(1-p) \log((1-p)/2) - p \log p \ .
\end{aligned}
$$

The mutual information between Alice and Eve is given by:

$$
I(A, E) = H(E) - H(E|A) \ .
$$

and the channel capacity is given by :

$$
\begin{aligned}
C &= \max_{A} I(A, E) \\
&= \max_{x_1, x_2} -\{p x_0 \log[p x_0] \\
&\quad + [q x_0/2 + p x_1] \log[q x_0/2 + p x_1] \\
&\quad + [q x_0/2 + p x_2] \log[q x_0/2 + p x_2] \\
&\quad + [q x_1/2] \log[q x_1/2] + [q x_1/2 + q x_2/2] \log[q x_1/2 + q x_2/2] \\
&\quad + [q x_2/2] \log[q x_2/2]\} - h_2(p).
\end{aligned}
$$

Note that the maximization is over x_1 and x_2, since x_0 is determined by these two probabilities (holds for any N). This equation is very difficult to solve analytically and requires numerical techniques. Figure 4 shows the capacity for this case with the curve labeled $N = 1$. From the plot the minimum capacity is approximately 0.92, when $p = 1/3$.

Case 5 — $N = 2$ clueless senders and $M = 2$ receivers

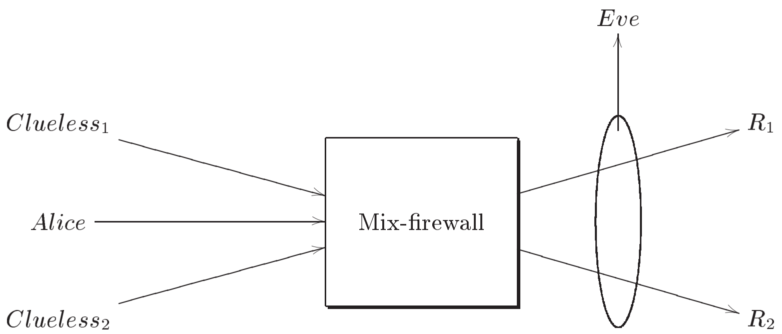

Fig. 3. Case 5: system with $N = 2$ clueless senders and $M = 2$ receivers

With two clueless senders and two receivers, Eve may receive ten symbols since there are ten different ways to put three indistinguishable balls into three distinct urns.

$$
M_{2.2} = \begin{array}{cc}
 & \begin{array}{cccccccccc} \langle 3,0,0\rangle & \langle 2,1,0\rangle & \langle 2,0,1\rangle & \langle 1,2,0\rangle & \langle 1,1,1\rangle & \langle 1,0,2\rangle & \langle 0,1,2\rangle & \langle 0,3,0\rangle & \langle 0,2,1\rangle & \langle 0,0,3\rangle \end{array} \\
\begin{array}{c} 0 \\ 1 \\ 2 \end{array} & \left(\begin{array}{cccccccccc}
p^2 & pq & pq & q^2/4 & q^2/2 & q^2/4 & 0 & 0 & 0 & 0 \\
0 & p^2 & 0 & pq & pq & 0 & q^2/4 & q^2/4 & q^2/2 & 0 \\
0 & 0 & p^2 & 0 & pq & pq & q^2/2 & 0 & q^2/4 & q^2/4
\end{array}\right)
\end{array}
$$

The 3×10 channel matrix $M_{2.2}[i, j]$ represents the conditional probability of Eve receiving e_j when Alice sends a message to receiver R_i.

Figure 4 shows the capacity for this case in the curve labeled $N = 2$. Again, the minimum capacity is found at $p = 1/3 = 1/(M + 1)$. From the plot the minimum capacity is approximately 0.62, when $p = 1/3$.

Case 6 — General Case: N clueless senders and M receivers

We now generalize the problem to N clueless senders and M receivers (refer again to Figure 1). There are $N + 1$ indistinguishable transmissions (including null transmissions) and they are sent into $M + 1$ distinct receivers (urns) (this also includes the null transmission, which by convention goes to R_0, not shown in the figure). Combinatorics tells us then that there are $K = \binom{N+M+1}{N+1}$ possible symbols e_j.

The rows of our channel matrix correspond to the actions of Alice. The ith row of $\mathrm{M}_{N.M}$ describes the conditional probabilities $p(e_j|x_i)$. By convention e_1 always corresponds to every sender not sending a message (which is equivalent to all senders sending to R_0). Therefore e_1 is the $M+1$ tuple $\langle N+1, 0, \ldots, 0\rangle$. Given our simplifying semi-uniformity assumption for the clueless senders' distribution, this term must be handled differently.

The first row of the channel matrix is made up of the terms $\mathrm{M}_{N.M}$[0,j]. (We will not always explicitly note that $j = 1, \ldots, \binom{N+M+1}{N+1}$.) Here, Alice is not sending any message (i.e., she is "sending" to R_0), so Alice contributes one to the term a_0^j in the $M+1$ tuple $\langle a_0^j, a_1^j, a_2^j, \ldots, a_M^j\rangle$ associated with e_j. In fact, this tuple is the "long hand" representation of e_j. Therefore the contributions to the $M+1$ tuple $\langle a_0^j - 1, a_1^j, a_2^j, \ldots, a_M^j\rangle$ describe what the N clueless senders are doing. That is, $a_0^j - 1$ clueless senders are not sending a message, a_1^j clueless senders are sending to R_1, etc. Hence, the multinomial coefficient $\binom{N}{a_0^j-1, a_1^j, \ldots, a_M^j}$ tells us how many ways this may occur.[7] For each such occurrence we see that the transmissions to R_0 affect the probability by $p^{a_0^j-1}$, and the transmissions to $R_i, i > 0$, due to the semi-uniformity assumption, contribute $(q/M)^{a_i^j}$. Since the actions are independent, the probabilities multiply, and since $a_0^j - 1 + a_1^j + \cdots + a_M^j = N$, we have a probability term of $p^{a_0^j-1}(q/M)^{N+1-a_0^j}$. Multiplying that term by the total number of ways of arriving at that arrangement we have that:
$\mathrm{M}_{N.M}[0, j] = \binom{N}{a_0^j-1, a_1^j, \ldots, a_M^j} p^{a_0^j-1}(q/M)^{N+1-a_0^j}$.

The other rows of the channel matrix are $\mathrm{M}_{N.M}[i, j],\ i > 0$. For row $i > 0$, we have a combinatorial term $\binom{N}{a_0^j, a_1^j, \ldots, a_{i-1}^j, a_i^j-1, a_{i+1}^j, \ldots, a_M^j}$ for the N clueless senders, a_0^j of which are sending to R_0 and $N - a_0^j$ of which are sending to the $R_i, i > 0$. Therefore, we see that under the uniformity assumption,
$\mathrm{M}_{N.M}[i, j] = \binom{N}{a_0^j, a_1^j, \ldots, a_{i-1}^j, a_i^j-1, a_{i+1}^j, \ldots, a_M^j} p^{a_0^j}(q/M)^{N-a_0^j}\ , i > 0$.

We show the plots of the mutual information when the clueless senders act (assumed throughout the paper) in a semi-uniform manner *and* when Alice also sends in a semi-uniform manner (i.e., $x_i = (1 - x_0)/M,\ i = 1, 2, ..., M$). We **conjecture** based upon our intuition, but do not prove, that Alice having a semi-uniform distribution of destinations $R_1, ..., R_M$ when the clueless senders act in a semi-uniform manner maximizes mutual information (achieves capacity). This has been supported by all of our numeric computations for capacity. With this conjecture, we can reduce the degrees of freedom for Alice from M to 1 (her distribution A is described entirely by x_0), which allows greater experimental and analytical exploration.

[7] The multinomial coefficient is taken to be zero, if any of the "bottom" entries are negative.

The channel matrix greatly simplifies when both the clueless senders and Alice act in a *totally uniform manner.* That is, when $x_0 = 1/(M+1)$, then $x_i = (1-x_0)/M = 1/(M+1)$ for all x_i, and $p = 1/(M+1)$. We have
$\mathrm{M}_{N.M}[0,j] = \binom{N}{a_0^j-1, a_1^j, \ldots, a_M^j} p^{a_0^j-1}(q/M)^{N+1-a_0^j}$, which simplifies to
$\mathrm{M}_{N.M}[0,j] = \binom{N}{a_0^j-1, a_1^j, \ldots, a_M^j} (\frac{1}{M+1})^N$. We also have
$\mathrm{M}_{N.M}[i,j] = \binom{N}{a_0^j, a_1^j, \ldots, a_{i-1}^j, a_i^j-1, a_{i+1}^j, \ldots, a_M^j} p^{a_0^j}(q/M)^{N-a_0^j}\ , i > 0,$
which simplifies to $\mathrm{M}_{N.M}[i,j] = \binom{N}{a_0^j, a_1^j, \ldots, a_{i-1}^j, a_i^j-1, a_{i+1}^j, \ldots, a_M^j} (\frac{1}{M+1})^N\ , i > 0$. Note that this form holds for $i = 0$ also, due to the total uniformity of the C_i.

To determine the distribution E describing Eve we need to sum over the columns of the channel matrix and use the total uniformity of A.
$P(E = e_j) = \sum_i P(E = e_j | A = i) P(A = i)\ i = 0, \ldots, M$. This gives us
$P(E = e_j) = (\frac{1}{M+1})^N \sum_{i=0}^{M} \binom{N}{a_0^j, \ldots, a_{i-1}^j, a_i^j-1, a_{i+1}^j, \ldots, a_M^j} = (\frac{1}{M+1})^N \binom{N+1}{a_0^j, \ldots, a_M^j}$.
From this we can compute the entropy $H(E)$ without too much trouble
$H(E) = (\frac{1}{M+1})^N \sum_j \binom{N+1}{a_0^j, \ldots, a_M^j} \left(N \log(M+1) - \log \binom{N+1}{a_0^j, \ldots, a_M^j} \right)$. However, the conditional entropy is more complicated, but is expressible. Therefore, we wrote Matlab code to calculate the mutual information, which is conjectured to be the capacity, when both the clueless senders and Alice act in a semi-uniform manner. Local exploration of nearby points all yield lower mutual information values.

4 Discussion of Results

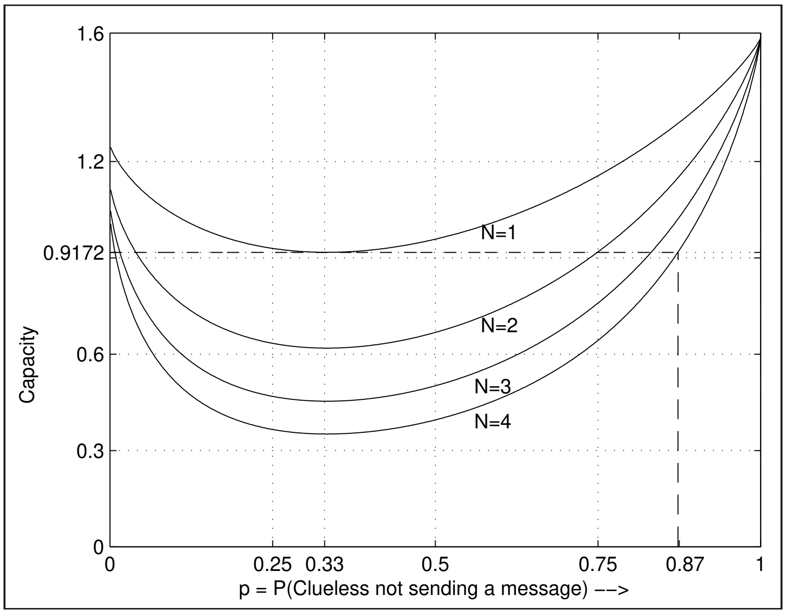

Fig. 4. Capacity for $M = 2$ receivers and $N = 1$ to 4 clueless senders

Figure 4 shows the capacity as a function of p with $M = 2$ receivers, for $N = 1, 2, 3, 4$ clueless senders. In all cases, the minimum capacity is realized at $p = 1/3$, and the capacity at $p = 1$ is identical to $\log 3$. As N increases, the capacity decreases, with the most marked effects at $p = 1/3$.

In Figure 4, the capacity (of course under the semi-uniformity assumption for C_i)) was determined numerically for any choice of A. However, for the remaining plots, we applied the semi-uniformity conjecture (that Alice is better off behaving semi-uniformly if that is what the clueless senders do). Thus, x_0 is the only free variable for Alice's distribution in what follows.

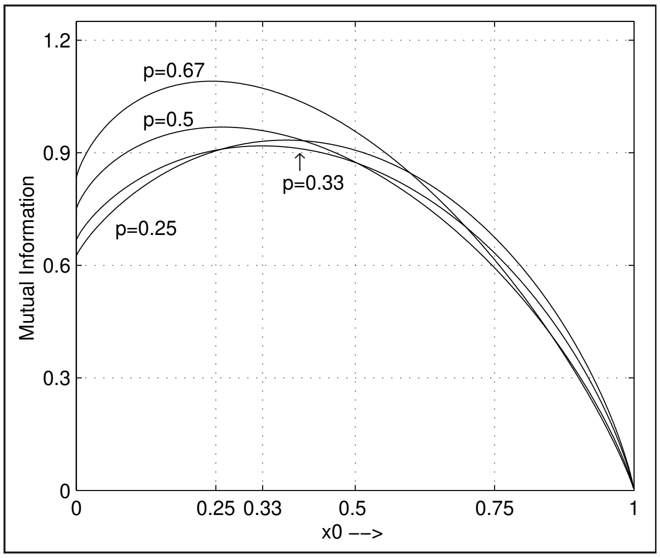

Fig. 5. Mutual information vs. x_0 for $M = 2$ receivers and $N = 1$ clueless sender, for $p = 0.25, 0.33, 0.5, 0.67$

The mutual information as a function of x_0 is shown in Figure 5 for $M = 2$ receivers and $N = 1$ clueless sender for $p = 0.25, 0.33, 0.5, 0.67$. Here, note that the curve with $p = 0.33$ has the smallest maximum value (capacity), and that the value of x_0 at which that maximum occurs is $x_0 = 0.33$. The x_0 value that maximizes the mutual information (i.e., for which capacity is reached) for the other curves is not 0.33, but the mutual information at $x_0 = 0.33$ is not much less than the capacity for any of the curves.

Figure 6 shows the mutual information curves for various values of x_0 as a function of p, with $N = 2$ clueless senders and $M = 2$ receivers. Note that the curve for $x_0 = 1/(M + 1) = 1/3$ has the largest minimum mutual information, and also has the greatest mutual information at the point where $p = 1$, i.e., when there is no noise since Clueless_1 is not sending any messages. The capacity for various values of p is, in essence, the curve that is the maximum at each p over all of the x_0 curves, and the lower bound on capacity occurs at $p = 1/3 = 1/(M+1)$.

Also observe that the $x_0 = 0.33$ curve has the highest value for $p = .33$, but for other values of p, other values of x_0 have higher mutual information

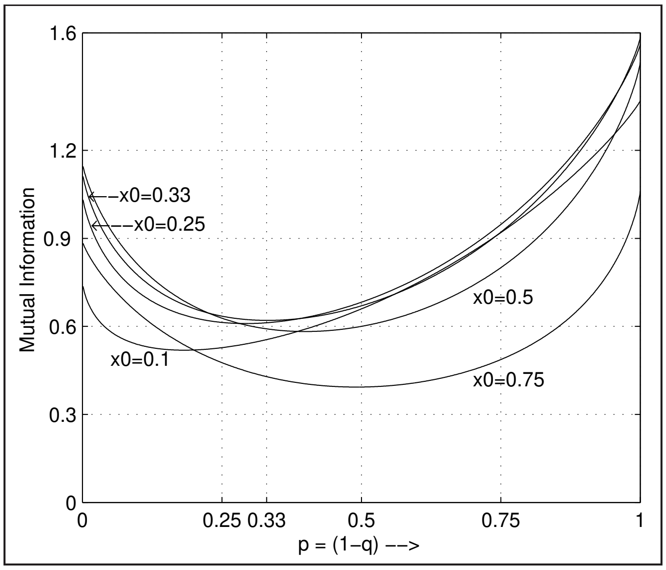

Fig. 6. Mutual information vs. p for $N = 2$ clueless senders and $M = 2$ receivers

(i.e., Alice has a strategy better than using $x_0 = 0.33$). However, the mutual information when $x_0 = 0.33$ is never much less than the capacity at any value of p, so in the absence of information about the behavior of the clueless senders, a good strategy for Alice is to just use $x_0 = 1/(M + 1)$. These observations are illustrated and expanded in the next two figures. Note the differences in concavity between Figure 5 and Figure 6. We will discuss concavity again later in the paper.

Figure 7 shows the optimal value for x_0, i.e., the one that maximizes mutual information and hence, achieves channel capacity, for $N = 1, 2, 3, 4$ clueless senders and $M = 3$ receivers as a function of p. A similar graph in [5] for $M = 1$ receiver is symmetric about $x_0 = 0.5$, but for $M > 1$ the symmetry is multidimensional, and the graph projected to the (p, x_0)-plane where the destinations are uniformly distributed is not symmetric. However, note that the optimum choice of x_0 is $1/(M + 1)$ both at $p = 1/(M + 1)$ and at $p = 1$, that is, when the clueless senders either create maximum noise or when they do not transmit at all (no noise). As N increases, the optimum x_0 for other values of p is further from $1/(M + 1)$. Also observe that Alice's best strategy is to do the opposite of what the clueless senders do, up to a point. If they are less likely to send messages ($p > 1/(M + 1)$), then Alice should be more likely to send messages ($x_0 < 1/(M + 1)$), whereas if Clueless_i is more likely to send messages (($p < 1/(M + 1)$), then Alice should be less likely to send messages ($x_0 > 1/(M + 1)$).

Figure 8 shows the degree to which the choice of $x_0 = 1/(M+1)$ can be suboptimal, for $N = 1, 2, 3, 4$ clueless senders and $M = 3$ receivers. The plot shows the mutual information for the given p and $x_0 = 1/(M + 1)$, normalized by dividing by the capacity (maximum mutual information) at that same p. Hence, it

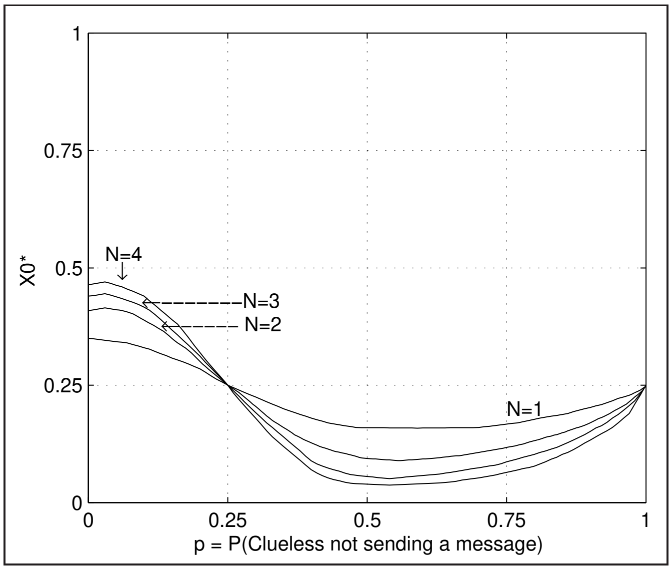

Fig. 7. Value of x_0 that maximizes mutual information for $N = 1, 2, 3, 4$ clueless senders and $M = 3$ receivers as a function of p

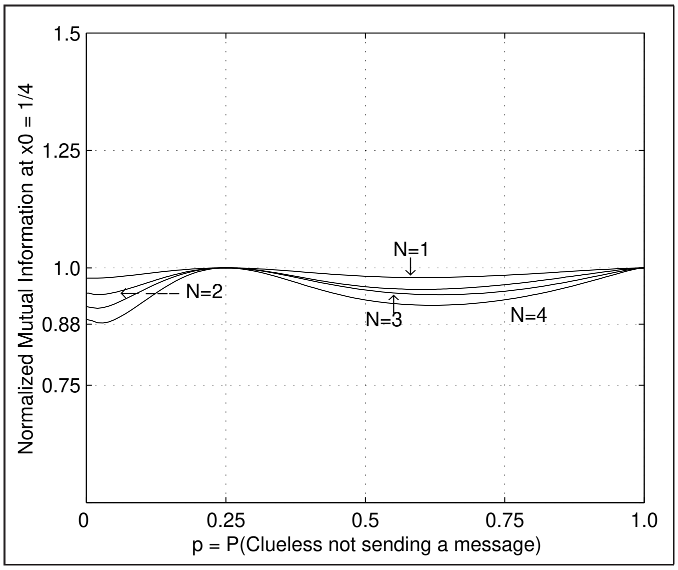

Fig. 8. Normalized mutual information when $x_0 = 1/4$ for $N = 1, 2, 3, 4$ clueless senders and $M = 3$ receivers

shows the degree to which a choice of $x_0 = 1/(M+1)$ fails to achieve the maximum mutual information. For $N = 2$, it is never worse than 0.94 (numerically), but for $N = 4$, its minimum is 0.88. The relationship of suboptimality for other choices of M and N, or for other distributions is not known.

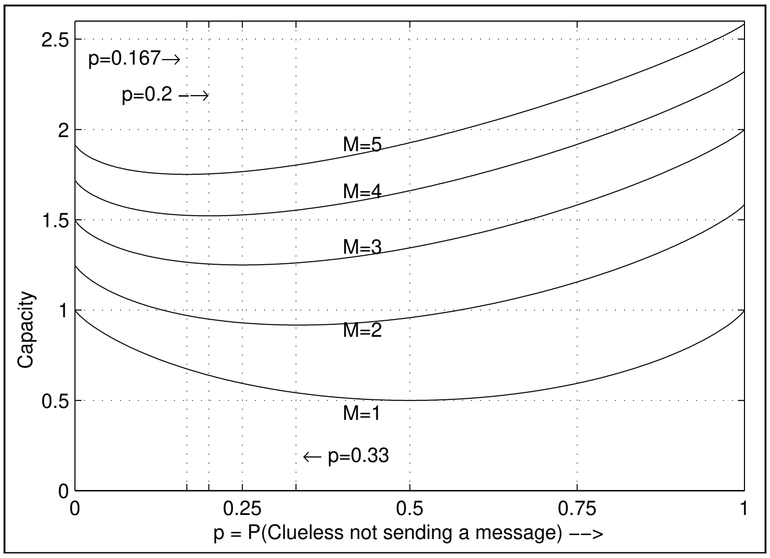

Fig. 9. Lower bound on capacity for $N = 1$ clueless sender and $M = 1$ to 5 receivers

In Figure 9, we show the lower bound on capacity of the channel as a function of p for $N = 1$ clueless sender and various values of M receivers. Numerical results show that this lower bound increases for all p as M increases, and the lower bound on the capacity for a given M occurs at $p = 1/(M+1)$, which is indicated by the dotted lines in the figure.

For Figure 10, we take the capacity at $p = 1/(M+1)$, which we found numerically to minimize the capacity of the covert channel, and plot this lower bound for capacity for many values of N and M. We retain the assumption that $x_i = (1-x_0)/(M+1)$ for $i = 1, 2, ..., M$, that is, given the semi-uniform distribution of transmissions to the receivers by the clueless senders, it is best for Alice to do likewise. Along the surface where $N = 0$, we have the noiseless channel, and the capacity is $\log(M+1)$, which is also the upper bound for capacity for all N and M. The values along the surface when $M = 1$ give us the same values we derived in [5].

Equations and curves for additional values and ranges of N and M may be found in a forthcoming technical report [7].

5 Comments and Generalizations

We first note that the maximum capacity of this (covert) quasi-anonymous channel is $\log(M+1)$ for M distinguishable receivers, and is achievable only if there

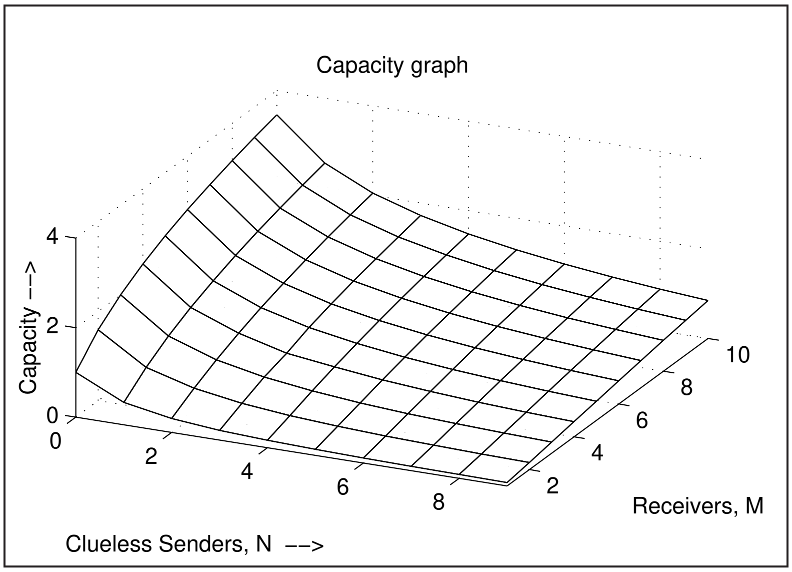

Fig. 10. Capacity lower bound for $N = 0$ to 9 clueless senders and $M = 1$ to 10.

are no other senders ($N = 0$) or if none of them ever send ($p = 1$), i.e., when the channel is noiseless.

Here are some of the observations from the different cases considered, under the semi-uniform assumption for the clueless senders and the semi-uniform conjecture for Alice, followed by some generalizations.

1. The capacity $C(p, N, M)$, as a function of the probability p that a clueless sender remains silent, with N clueless senders and M receivers, is strictly bounded below by $C(\frac{1}{M+1}, N, M)$, and is achieved with $x_0 = 1/(M + 1)$.
2. The lower bound for capacity for a given number M of receivers decreases as the number N of clueless senders increases,
$C(\frac{1}{M+1}, N, M) \ > \ C(\frac{1}{M+1}, N + 1, M)$.
3. The lower bound for capacity for a given number N of clueless senders increases as the number M of distinguishable receivers increases,
$C(\frac{1}{M+2}, N, M + 1) \ > \ C(\frac{1}{M+1}, N, M)$.

These observations are intuitive, but we have not shown them to be true numerically in the general case (we did for the case that $M = 1$ in [5]). It is interesting to note that increasing the number of distinguishable receivers increases the covert channel capacity, which in some sense *decreases* the (sender) anonymity in the system (Alice has more room in which to express herself). This is a bit contrary to the conventional view of anonymity in Mix networks, where more receivers tends to provide "greater anonymity." In this light, we note that Danezis and Serjantov investigated the effects of multiple receivers in statistical attacks on anonymity networks [3]. They found that Alice having multiple receivers greatly lowered a statistical attacker's certainty of Alice's receiver set.

While the graphs and numerical tests support that the "worst" thing the clueless senders can do is to send (or not) with uniform probability distribution over the R_i, $i = 0, 1, 2, ..., M$, we have not proven this mathematically. Nor have we proven that, under these conditions, the best Alice can do is to send (or not) to each receiver R_i with uniform probability, $x_i = 1/(M+1)$ for $i = 0, 1, 2, ..., M$, although the numerical computations support this. The proof in [5] of these conjectures for the case where $M = 1$ relied, in part, on the symmetry about $x_0 = 0.5$, which is not the case when $M > 1$, so another approach must be used. However, we should still be able to use the concavity/convexity results from [5]. Note that our conjecture that the best that Alice can do is to send in a semi-uniform manner, and the results illustrated in Figure 8, seem to be an extension of the interesting results of [4].

6 Conclusions and Future Work

This paper has taken a step towards tying the notion of capacity of a quasi-anonymous channel associated with an anonymity network to the amount of anonymity that the network provides. It explores the particular situation of a simple type of timed Mix (it fires every tick) that also acts as an exit firewall. Cases for varying numbers of distinguishable receivers and varying numbers of senders were considered, resulting in the observations that more senders (not surprisingly) decreases the covert channel capacity, while more receivers increases it. The latter observation is intuitive to communication engineers, but may not have occurred to many in the anonymity community, since the focus there is often on sender anonymity.

As the entropy H of the probability distribution associated with a message output from a Mix gives the effective size, 2^H, of the anonymity set, we wonder if the capacity of the residual quasi-anonymous channel in an anonymity system provides some measure of the effective size of the anonymity set for the system as a whole. That is, using the covert channel capacity as a standard yardstick, can we take the capacity of the covert channel for the observed transmission characteristics of clueless senders, equate it with the capacity for a (possibly smaller) set of clueless senders with maximum entropy (i.e., who introduce the maximum amount of noise into the channel for Alice), and use the size of this latter set as the effective number of clueless senders in the system. This is illustrated in Figure 4, with the vertical dashed line showing that $N = 4$ clueless senders that remain silent with probability $p = 0.87$ are in some sense equivalent to one clueless sender that sends with $p = 0.33$.

The case in which the Mix itself injects dummy messages into the stream randomly is not distinguishable from having an additional clueless sender. However, if the Mix predicates its injection of dummy messages upon the activity of the senders, then it can affect the channel matrix greatly, to the point of eliminating the covert channel entirely. We are also interested in the degree to which the Mix can reduce the covert channel capacity (increase anonymity) with a limited ability to inject dummy messages.

In future work we will analyze the situation where we have different (and more realistic) distributions for the clueless senders. We are also interested in different kinds of exit point Mix-firewalls, such as threshold Mixes, timed Mixes (where the time quantum is long enough to allow more than one message per sender to be sent before the Mix fires), timed-pool Mixes, and systems of Mixes.

7 Acknowledgements

We thank the reviewers and Andrei Serjantov for their helpful comments.

References

1. David Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24(2):84–88, 1981.
2. George Danezis and Andrei Serjantov. Statistical disclosure or intersection attacks on anonymity systems. In *Information Hiding, Sixth International Workshop*. Springer-Verlag, LNCS, 2004.
3. George Danezis and Andrei Serjantov. Statistical disclosure or intersection attacks on anonymity systems. In *IH 2004, Springer LNCS TBD*, page TBD, Toronto, Canada, May 2004.
4. E.E. Majani and H. Rumsey. Two results on binary input discrete memoryless channels. In *IEEE International Symposium on Information Theory*, page 104, June 1991.
5. Ira S. Moskowitz, Richard E. Newman, Daniel P. Crepeau, and Allen R. Miller. Covert channels and anonymizing networks. In *ACM WPES*, pages 79–88, Washington, October 2003.
6. Ira S. Moskowitz, Richard E. Newman, and Paul F. Syverson. Quasi-anonymous channels. In *IASTED CNIS*, pages 126–131, New York, December 2003.
7. Richard E. Newman, Vipan R. Nalla, and Ira S. Moskowitz. Covert channels and simple timed mix-firewalls. NRL Memorandum Report to appear, NRL, 2004.
8. Andrei Serjantov, Roger Dingledine, and Paul Syverson. From a trickle to a flood: Active attacks on several mix types. In F.A.P. Petitcolas, editor, *IH 2002, Springer LNCS 2578*, pages 36–52, Noordwijkerhout, the Netherlands, October 2002.
9. Claude E. Shannon. The mathematical theory of communication. *Bell Systems Technical Journal*, 30:50–64, 1948.