A Three-Dimensional Approach Towards Measuring Sender Anonymity *

Neeraj Jaggi Umesh MarappaReddy Rajiv Bagai
Department of Electrical Engineering and Computer Science
Wichita State University, Wichita, KS 67260
Email: {neeraj.jaggi, uxmarappareddy, rajiv.bagai}@wichita.edu

***Abstract*—We first illustrate using examples that existing measures in literature are not sufficient to fully characterize the anonymity provided by an anonymous system. We then propose a new *isolation* measure, based upon presence of outliers in a distribution, and show that this measure is critical towards quantifying the overall anonymity provided by the system. We provide justification for three distinct aspects of anonymity, important from the perspectives of a user, a system designer, and an attacker, leading to a three-dimensional approach towards measuring sender anonymity. We further show how two anonymous systems can be compared in terms of the degree of anonymity provided, using the proposed 3-tuple metric and appropriate weights reflecting the attributes desired in the system. Finally, we apply the proposed metric to an existing anonymous system and discuss the insights gained.**

***Index Terms*—Measuring Anonymity, Isolation, Privacy**

I. Introduction

The ability to maintain one's privacy on the Web has always been an important concern [1], [2]. Anonymity on the Web is critical to enable applications such as e-voting, auctions, payments, and to provide necessary cyber-security measures [3]. With renewed emphasis on identity protection, privacy awareness and selective information sharing on social networks [4], the ability to protect online activities from possible misuse is gaining importance [5].

Various approaches have been proposed towards providing anonymous transactions on the Web. A widely employed technique for providing anonymity is by introducing proxy server(s) between a sender-receiver pair, for instance as in Crowds [6]. Crowds [6] attempts to conceal user's identity by forwarding user's message via a number of intermediate nodes before it is delivered to the end server. Other techniques used to provide anonymity and information hiding include introducing delays, mixing [7] and encrypting the messages. More sophisticated systems providing anonymous communications include Onion routing [8] and Tor [9], which use a combination of above approaches to provide unlinkability of a message from its sender and/or receiver.

The goal of an anonymous system includes providing one or more of – *sender* anonymity, *receiver* anonymity, and unlinkability of sender and receiver. In this paper, we focus solely on sender anonymity, and provide a metric to quantify the level of anonymity provided by the system. Potential eavesdroppers or attackers could reside anywhere in the network, and protecting the identity of the sender of a message from malicious nodes is the key concern in most anonymous systems. Various approaches consider specific attack models such as local eavesdroppers or collaborating intermediate nodes to analyze the properties of the system. However, the approach presented in [10], [11] is slightly different, and considers only the a posteriori probabilities assigned by an attacker to various users who are suspected of having sent a particular message, after the information has been gained from the attack. Our model is also based upon the latter approach, and is thus suitable to analyze anonymous systems under any general attack model.

In case of sender anonymity, the attacker attempts to identify the originator or sender of a particular message, from possible senders belonging to the anonymity set $\mathcal{A}$. The objective of the anonymous system design is to prevent such identification under different models of attacks possible in the system. Once an attacker has gained sufficient information using a particular attack, he/she attempts to identify the sender in a probabilistic manner. The attacker assigns a probability p_i to each user i in the anonymity set $\mathcal{A}$, where p_i is a measure of suspicion with which the attacker considers user i to be the actual sender of the message. Let $\mathcal{N}$ denote the size of anonymity set $\mathcal{A}$ ($\mathcal{N} = |\mathcal{A}|$). A probability distribution is valid if $\forall i \in \{1, \ldots, \mathcal{N}\} : 0 \leq p_i \leq 1$ and $\sum_{i=1}^{\mathcal{N}} p_i = 1$. Given the probability distribution $\mathbf{p}$, the anonymity set $\mathcal{A}$ and the sender $j \in \mathcal{A}$, quantifying the lack of sender identification information available to the attacker amounts to measuring the level or degree of anonymity provided by the system.

Various approaches have been presented to quantify the degree of anonymity provided by an anonymous system. One of the approaches considers just the probability assigned to the actual sender j [6], [12]. Information theoretic measures towards quantifying the degree of anonymity are discussed in [13], [10]. A simple entropy based measure is given by [13]:

$$S = -\sum_{i=1}^{\mathcal{N}} p_i \log_2(p_i), \tag{1}$$

where S could be interpreted as the number of bits of additional information that the attacker needs in order to completely identify the sender j. The measure S equals zero when $p_j = 1$, and equals $\log_2(\mathcal{N})$ (which is the maximum possible entropy) when $p_i = \frac{1}{\mathcal{N}}$, $\forall i \in \{1, \ldots, \mathcal{N}\}$. Let us

* This work was supported in part by the US Navy Engineering Logistics Office Contract No. N41756-08-C-3077.

denote the latter probability distribution as $\mathcal{D}$. A normalized entropy based measure is given by [10]:

$$d = \frac{S}{\log_2 \mathcal{N}} = -\frac{\sum_{i=1}^{\mathcal{N}} p_i \log_2(p_i)}{\log_2 \mathcal{N}}. \quad (2)$$

An interesting property of this measure is that, when the sender j is completely identified ($p_j = 1$), d equals zero. And $d = 1$ under distribution $\mathcal{D}$, since the sender j is as likely to have sent the message as any other user in $\mathcal{A}$ in the eyes of the attacker, and is thus relatively (completely) unidentifiable.

Shortcomings of both simple and normalized entropy based measures are outlined in [11], where counterexamples are constructed to show that two distributions could have the same measure but provide practically different anonymities. Authors in [11] introduce *local* anonymity $\theta = \max_{1 \le i \le \mathcal{N}} \ p_i$, argue that θ is more important than d from the perspective of a user, and derive appropriate relations between θ, S and d. Since the probability assigned to the actual sender is upper bounded by θ, a low value of θ is sufficient to hide the actual sender's identity. However, considering θ alone is not sufficient to characterize the anonymity of the distribution, since the probability assigned to the sender must be viewed in relation with the probabilities assigned to other users. Thus, the normalized entropy d is important as well.

In this paper, we first show (in Section II) that the existing measures of S, d and θ are not sufficient to completely characterize the (sender) anonymity provided by the system. Next, we argue that a new measure is needed to quantify the isolation of a user (or a set of users) in the distribution $\mathbf{p}$. We propose a new measure to quantify this isolation, based upon presence of outliers in a distribution, and discuss its properties (in Section III). We then propose a three-dimensional approach towards measuring sender anonymity (in Section IV), and provide justification for the three distinct aspects in this approach, considering the perspectives of a user, system designer and attacker respectively. We discuss comparison of anonymous systems using the proposed 3-tuple metric and suitable choice of weights (in Section V). Finally, we apply the proposed metric to an existing anonymous system (in Section VI) and summarize our conclusions (in Section VII).

II. Drawbacks of existing measures

In this section, we consider examples of anonymous systems with different anonymity properties, and show that even though an attacker assigns quite different probability distributions in these systems, the existing anonymity measures declare the two systems to be equally anonymous. Let $\mathcal{D}_1$ and $\mathcal{D}_2$ denote the a posteriori probability distributions assigned in the two anonymous systems, and let the sizes of the anonymity set be denoted $\mathcal{N}_1$ and $\mathcal{N}_2$ respectively.

Example 1**:** Consider two systems with $\mathcal{N}_1 = 7$ and $\mathcal{N}_2 = N + 1$, where the distributions are given by:

$$\mathcal{D}_1 : \ p_1 = 0.4, \ p_2 = 0.4, \ p_k = 0.04, \ \forall k \in \{3, \ldots, 7\}.$$

$$\mathcal{D}_2 : \ p_1 = 0.4, \ p_k = \frac{0.6}{N}, \ \forall k \in \{2, \ldots, N+1\}.$$

Here, it is apparent that the attacker has more information about the actual sender in System 2, compared with System 1 (particularly if N is large). We have,

$$d_1 = -\frac{\sum_{i=1}^{\mathcal{N}_1} p_i \log_2(p_i)}{\log_2 \mathcal{N}_1} = 0.708.$$

$$d_2 = -\frac{\sum_{i=1}^{\mathcal{N}_2} p_i \log_2(p_i)}{\log_2 \mathcal{N}_2} = -\frac{0.4\log_2(0.4) + 0.6\log_2\left(\frac{0.6}{N}\right)}{\log_2(N+1)}.$$

Equating $d_1 = d_2$, we get,

$$0.708 = \frac{0.971 + 0.6\log_2 N}{\log_2(N+1)}.$$

For large values of N, using $\log_2(N+1) \approx \log_2 N$, we get,

$$\log_2 N = 8.99 \quad \Rightarrow \quad N \approx 508.$$

Choosing $N = 508$, $d_2 = 0.708 = d_1$ (up to 3 places of decimal). The local anonymity measure for both the systems is the same, i.e. $\theta_1 = \theta_2 = 0.4$. Note that the simple entropy measure (given by (1)) for the two systems is not the same, and is in fact counterintuitive. Computing, we get, $S_1 = 1.986$ and $S_2 = 6.364$. Thus, $S_1 < S_2$, which seems to suggest that System 2 has higher degree of anonymity than System 1.[1] □

In Example 1, we considered two systems with different sizes of anonymity sets. Next, we show that existing measures fail to differentiate systems with same anonymity set size.

Example 2**:** Consider two systems with $\mathcal{N}_1 = \mathcal{N}_2 = 100$, where the distributions are given by (for some $N \in \{3, \ldots, 97\}$):

$$\mathcal{D}_1 : \ p_1 = p_2 = 0.2, \ p_k = \frac{0.6}{98}, \ \forall k \in \{3, \ldots, 100\}.$$

$$\mathcal{D}_2 : \ p_1 = 0.2, \ p_k = \frac{0.5}{N}, \ \forall k \in \{2, \ldots, N+1\},$$

$$p_l = \frac{0.3}{(99 - N)}, \ \forall l \in \{N+2, \ldots, 100\}.$$

We have,

$$d_1 = -\frac{\sum_{i=1}^{100} p_i \log_2(p_i)}{\log_2 100} = 0.804.$$

$$d_2 = -\frac{0.2\log_2(0.2) + 0.5\log_2\left(\frac{0.5}{N}\right) + 0.3\log_2\left(\frac{0.3}{99-N}\right)}{\log_2 100}.$$

Equating $d_1 = d_2$, we get,

$$0.5\log_2 N + 0.3\log_2(99 - N) = 3.856.$$

The integer value of N which satisfies the above equation with minimum error is given by $N = 15$. Choosing $N = 15$, $d_2 = 0.806 \approx d_1$. The local anonymity measure for both the systems is the same, i.e. $\theta_1 = \theta_2 = 0.2$. In this example, since the attacker is able to isolate two users in System 1 and one user in System 2, the degree of anonymity provided by the two systems appears to be different. □

[1] Therefore, we do not consider simple entropy measure in other examples.

Next, we provide an example to show that normalized entropy measure could sometimes be misleading.

Example 3: Consider two systems with $\mathcal{N}_1 = \mathcal{N}_2 = N$, where the distributions are given by:

$$\mathcal{D}_1 : \ p_1 = 0.3, \ p_k = \frac{0.7}{N-1}, \ \forall k \in \{2, \ldots, N\}.$$

$$\mathcal{D}_2 : \ p_1 = p_2 = 0.3, \ p_k = \frac{0.4}{N-2}, \ \forall k \in \{3, \ldots, N\}.$$

Here, for all $N > 3$, $\theta_1 = \theta_2 = 0.3$ and $d_1 > d_2$. For e.g., for $N = 10$, $d_1 = 0.933$ and $d_2 = 0.834$. However, System 2 seems to provide a higher degree of anonymity than the System 1, since the attacker is able to isolate a single user in the anonymity set in System 1 (particularly for large N). □

III. Measuring Isolation in Anonymous Systems

In all the examples considered in the previous section, one observation is apparently common. In one of the distributions, exactly one user is being isolated by the attacker i.e. the probability assigned to exactly one user is substantially higher than that assigned to other users in the anonymity set. As a result, the degree of anonymity provided by the corresponding anonymous system seems to be lower than the other system. However, the measures d and θ are unable to satisfactorily quantify this *isolation* of the user in the anonymity set. In this Section, we propose a new measure, *Isolation factor* (denoted $\mathcal{I}$) to measure the degree of isolation in the system.

A user in the anonymity set is considered isolated, if the probability assigned to that user is significantly larger than that assigned to other users in the set. Viewing these probabilities as a set of observations, the probability assigned to the isolated user corresponds to an *outlier* in this set. An outlier is a statistical observation which appears to deviate considerably from the rest of the observations. Multiple approaches exist to detect an outlier in a sample of data [14]. Model-based methods identify observations which are deemed *unlikely* based upon the mean and standard deviation of the distribution, and are commonly used. One such method, which is able to identify multiple outliers, is the Peirce's criterion [15], [16], [17]. Here, observations deviating from the mean beyond an appropriately computed threshold are declared outliers. When the number of observations is relatively small (< 60), a simplified procedure to detect outliers based upon Peirce's criterion could be used [17]. However, we use the detailed procedure proposed in [16] which is valid for up to 150 observations. When the number of observations is larger, the algorithm is applied to multiple sets of observations, obtained by dividing the original set into equal sized pieces of < 150 observations each, as suggested in [16]. Outlier detection is discussed in the Appendix.

A. Isolation Factor

Using the above approach, we detect all outliers in the probability distribution which are significantly larger than the mean. Let $\mathcal{L}$ denote the number of outliers detected. A new measure, Isolation factor ($\mathcal{I}$) is defined in order to measure the extent of additional information the attacker would gain due to the presence of these outliers. If $\mathcal{L} = 0$, $\mathcal{I}$ is defined to be zero. The desired properties of this measure include:

- Isolation factor should decrease as number of outliers increase, since the information gained by the attacker decreases as the number of isolated users increase.
- Isolation factor should be proportional to the extent of deviation of the outliers, since the larger the deviation, the higher is the perceived suspicion of the attacker towards the corresponding users.
- For the same number and value(s) of outlier(s), the Isolation factor should increase with an increase in the size of the anonymity set, as the attacker's suspicion set becomes a smaller fraction of the total number of users.

For a probability distribution $\mathbf{p}$ with $\mathcal{N}$ users and $\mathcal{L}$ outliers, let us assume without loss of generality, that the first $\mathcal{L}$ probabilities in $\mathbf{p}$ correspond to the outliers. The Isolation factor $\mathcal{I}$ is defined as:

$$\mathcal{I} = \frac{\sqrt{\sum_{i=1}^{\mathcal{N}} (p_i - \bar{p})^2} - \sqrt{\sum_{i=\mathcal{L}+1}^{\mathcal{N}} (p_i - \bar{q})^2}}{\max\{1, \mathcal{L}\}}, \qquad (3)$$

where $\bar{p}$ is the sample mean with outliers (and equals $\frac{1}{\mathcal{N}}$), and $\bar{q}$ is the sample mean without outliers, i.e. $\bar{q} = \frac{\sum_{i=\mathcal{L}+1}^{\mathcal{N}} p_i}{\mathcal{N}-\mathcal{L}}$. Note that the definition of the measure satisfies the desired properties mentioned above. The numerator in (3) measures the impact of outliers in increasing the standard deviation of the distribution, while the denominator adjusts this impact based upon the number of outliers. Next, we discuss some interesting properties of this new measure.

B. Properties of Isolation Factor

Proposition 1: Isolation factor $\mathcal{I}$ lies in the range $[0, \ 1]$.

Proof: When there are no outliers, $\bar{q} = \bar{p}$ and $\mathcal{I} = 0$, indicating no isolation in the distribution. When user j is completely isolated i.e. $p_j = 1$ and $p_i = 0, \ \forall i \neq j$,

$$\mathcal{I} = \sqrt{\left(1 - \frac{1}{\mathcal{N}}\right)^2 + \frac{\mathcal{N}-1}{\mathcal{N}^2}} = \sqrt{\frac{\mathcal{N}-1}{\mathcal{N}}}. \qquad (4)$$

In this case, $\mathcal{I} \to 1$ as $\mathcal{N} \to \infty$. For all other scenarios, including more than one outlier, $0 < \mathcal{I} < 1$, as shown below.

$$\begin{aligned} \mathcal{I} &= \frac{\sqrt{\sum_{i=1}^{\mathcal{N}} (p_i - \bar{p})^2} - \sqrt{\sum_{i=\mathcal{L}+1}^{\mathcal{N}} (p_i - \bar{q})^2}}{\max\{1, \mathcal{L}\}} \\ &\leq \sqrt{\sum_{i=1}^{\mathcal{N}} (p_i - \bar{p})^2} - \sqrt{\sum_{i=\mathcal{L}+1}^{\mathcal{N}} (p_i - \bar{q})^2} \leq \sqrt{\sum_{i=1}^{\mathcal{N}} (p_i - \bar{p})^2}. \end{aligned}$$

Therefore, $\mathcal{I}^2 \leq \sum_{i=1}^{\mathcal{N}} (p_i - \bar{p})^2 = \sum_{i=1}^{\mathcal{N}} p_i^2 - \frac{1}{\mathcal{N}} \leq \sum_{i=1}^{\mathcal{N}} p_i^2$. Since $0 \leq p_i \leq 1, \ \forall i \in \{1, \ldots, \mathcal{N}\}, \ \sum_{i=1}^{\mathcal{N}} p_i^2 \leq \left(\sum_{i=1}^{\mathcal{N}} p_i\right)^2 = 1$, as $\sum_{i=1}^{\mathcal{N}} p_i = 1$. Thus $\mathcal{I}^2 \leq 1$ and hence $\mathcal{I} \leq 1$. ■

$\mathcal{I} = 0$ when there is no isolation in the system, and $\mathcal{I} = 1$ when the degree of isolation in the system is the maximum. Thus, the Isolation factor provides a measure of the extent of

additional information the attacker gains due to the presence of the outliers in the distribution. Let the outlier probabilities be denoted $p_1, \ldots, p_{\mathcal{L}}$. Let $\bar{l}$ denote the mean of outlier probabilities, i.e. $\bar{l} = \frac{\sum_{i=1}^{\mathcal{L}} p_i}{\mathcal{L}}$.

Proposition 2: Isolation factor $\mathcal{I}$ satisfies the inequality:

$$\mathcal{I} \leq \bar{l}\left(1 + \frac{1}{\sqrt{\mathcal{N}}}\right). \quad (5)$$

Proof: Follows from (3), using $\sqrt{A+B} \leq \sqrt{A} + \sqrt{B}$, for $A, B \geq 0$, and is omitted due to space constraints. Detailed proof is provided in [18]. ∎

IV. Three Dimensional Measure of Anonymity

We believe that the three measures namely, normalized entropy d, local anonymity θ and Isolation factor $\mathcal{I}$, are all essential towards measuring the degree of anonymity of an anonymous system. Note that all the three measures lie in the range $[0,\ 1]$. Although a higher value of d is desirable from the system's perspective, a lower value of θ and $\mathcal{I}$ would be preferred from an end user's perspective. Let us denote the degree of anonymity of a system by the 3-tuple $(d, \theta, \mathcal{I})$.

A. Motivation and Perspective

The objective of an anonymous system is to guarantee that the sender of a message be non-identifiable within the anonymity set. In the a posteriori probability distribution, this would be achieved when $d = 1$, i.e. $p_i = \frac{1}{\mathcal{N}},\ \forall i \in \{1, \ldots, \mathcal{N}\}$. Thus, the normalized entropy d is an important measure from the system designer's perspective.

However, as pointed out in [11], a higher value of d may not be sufficient to convince the end user. For a user trying to utilize the services provided by an anonymous system, she would like to have some guarantees on her maximum exposure for any message she sends in the system. This would be achieved when θ is minimized. Thus, the local anonymity θ is an important measure from the end user's perspective.

Now let us consider the attacker's point of view. Once the attacker has performed the attack and has assigned the probabilities based upon the information gained from the attack, the attacker would like to make a guess (his best bet) as to which user(s), from among those in the anonymity set, is most likely to have sent the message. The attacker's task gets easier when there is a single user (or a set of users) which is clearly isolated in the distribution. The larger the value of $\mathcal{I}$, the more successful is the attack from the attacker's perspective. Thus, the Isolation factor $\mathcal{I}$ is an important measure from the attacker's perspective, and including this measure helps better understand the overall anonymity of the system. Thus, all the three measures in the 3-tuple metric $(d, \theta, \mathcal{I})$ are important measures of sender anonymity provided by a system, albeit one measure may be rendered more important than the others, depending upon the perspective under consideration.

Let us consider an anonymous system with 3-tuple metric $(d, \theta, \mathcal{I})$. For a given $\mathcal{N}$, maximum anonymity is achieved under the probability distribution $\mathcal{D}$, defined in Section I. In this case, $d = 1$, $\theta = \frac{1}{\mathcal{N}}$ and $\mathcal{I} = 0$. As $\mathcal{N} \to \infty$, $\theta \to 0$, and the maximum anonymity corresponds to 3-tuple $(1, 0, 0)$.

B. Metric Interpretation and Evaluation

In order to characterize the overall anonymity of the system, we propose that the desired attributes of the system be represented using weights assigned to each of the three dimensions of the 3-tuple metric, and the overall anonymity of the system should be viewed in relation to these weights. Let W_d, W_θ and $W_\mathcal{I}$ denote the weights associated with d, θ and $\mathcal{I}$ respectively, such that,

$$0 \leq W_d,\ W_\theta,\ W_\mathcal{I} \leq 1, \text{ and } \quad W_d + W_\theta + W_\mathcal{I} = 1. \quad (6)$$

These weights are designed to reflect the attributes desired in the system. For instance, if an end user is only interested in local anonymity, she would view the system (and the 3-tuple metric) using weights $W_d = W_\mathcal{I} = 0$, and $W_\theta = 1$.

Once the weights have been assigned, representing $\sqrt{W_d} \cdot d$, $\sqrt{W_\theta} \cdot (1-\theta)$ and $\sqrt{W_\mathcal{I}} \cdot (1-\mathcal{I})$ on the $x-$, $y-$ and $z-$ coordinates respectively, the 3-tuple metric corresponds to a point in the 3-dimensional unit sphere[2]. The distance of this point from origin is interpreted as the overall anonymity of the system, with the maximum anonymity corresponding to a distance of 1. Specifically, this distance (denoted $\mathcal{R}$) equals

$$\mathcal{R} = \sqrt{W_d \cdot d^2 + W_\theta \cdot (1-\theta)^2 + W_\mathcal{I} \cdot (1-\mathcal{I})^2}. \quad (7)$$

Maximum possible distance, $\mathcal{R}_{\max} = \sqrt{W_d + W_\theta + W_\mathcal{I}} = 1$. The distribution $\mathcal{D}$ with $\mathcal{N} \to \infty$ results in maximum overall anonymity of $\mathcal{R} = 1$, regardless of the weights assigned. Similarly, the distribution with least anonymity, given by $p_1 = 1$, $p_k = 0,\ \forall k \in \{2, \ldots, \mathcal{N}\}$, results in minimum overall anonymity of $\mathcal{R} = 0$, as $\mathcal{N} \to \infty$, regardless of the weights assigned $\left(\text{since } d = 0,\ \theta = 1, \text{ and } \mathcal{I} = \sqrt{\frac{\mathcal{N}-1}{\mathcal{N}}} \text{ from } (4)\right)$.

Proposition 3: For distributions $\mathcal{D}_1$ and $\mathcal{D}_2$ with $d_1 > d_2$, $\theta_1 < \theta_2$ and $\mathcal{I}_1 < \mathcal{I}_2$, overall anonymity $\mathcal{R}_1 > \mathcal{R}_2$, for all choices of weights W_d, W_θ and $W_\mathcal{I}$.

Proof: Follows from (7). ∎

The degree of anonymity provided by a system is accurately characterized by 3-tuple $(d, \theta, \mathcal{I})$. The overall anonymity $\mathcal{R}$ should only be considered to compare different anonymous systems. For such comparisons, the choice of weights becomes important. However, from Proposition 3, this choice is not important for obvious comparisons, where one system is evidently better than the other. (Note that Proposition 3 can be extended to scenarios such as $d_1 > d_2$, $\theta_1 = \theta_2$ and $\mathcal{I}_1 = \mathcal{I}_2$.) Only in cases where this comparison is non-trivial, the weights should be chosen carefully to reflect user preferences during the comparison. Indeed, with a different choice of weights, the results of such comparisons are expected to differ. We recommend considering equal weights $W_d = W_\theta = W_\mathcal{I} = \frac{1}{3}$, in general. However, more sophisticated choices are also possible.

Consider anonymous systems $\mathcal{S}_1$, $\mathcal{S}_2$ and $\mathcal{S}_3$ with overall anonymity $\mathcal{R}_1$, $\mathcal{R}_2$ and $\mathcal{R}_3$ respectively. If $\mathcal{R}_1 > \mathcal{R}_2$, System

[2]Note that square root of weights is considered in order to assign the maximum overall anonymity of 1 to a system with distribution $\mathcal{D}$ and $\mathcal{N} \to \infty$, regardless of the weights assigned. Other ways to interpret the 3-tuple metric may also be plausible.

$\mathcal{S}_1$ is considered to be $\frac{\mathcal{R}_1}{\mathcal{R}_2}$ times more anonymous than System $\mathcal{S}_2$. Also, if $\mathcal{S}_1$ is u times more anonymous than $\mathcal{S}_2$, and $\mathcal{S}_2$ is v times more anonymous than $\mathcal{S}_3$, then $\mathcal{S}_1$ is $u \cdot v$ times more anonymous than $\mathcal{S}_3$.

V. Comparisons of Systems with New Metric

Example 1 revisited: Consider two systems with $\mathcal{N}_1 = 7$ and $\mathcal{N}_2 = 509$, where the distributions are given in Section II. We have $d_1 = d_2 = 0.708$, and $\theta_1 = \theta_2 = 0.4$. For $\mathcal{D}_1$, the number of outliers detected $\mathcal{L}_1 = 0$, and the Isolation factor $\mathcal{I}_1 = 0$. For $\mathcal{D}_2$, the number of outliers detected $\mathcal{L}_2 = 1$, and the Isolation factor is computed as $\mathcal{I}_2 = 0.398$. Considering weights $W_d = W_\theta = W_\mathcal{I} = \frac{1}{3}$, the overall anonymity of System 1 equals $\mathcal{R}_1 = 0.788$, and for System 2 equals $\mathcal{R}_2 = 0.639$. Thus, System 1 is more anonymous than System 2, under the chosen set of weights. Thus, using the 3-tuple metric allows us to distinguish the two systems in terms of the degree of anonymity provided. □

Example 2 revisited: Consider two systems with $\mathcal{N}_1 = \mathcal{N}_2 = 100$, where the distributions are given in Section II, and $N = 15$. We have $d_1 = 0.804$, $d_2 = 0.806$, and $\theta_1 = \theta_2 = 0.2$. For $\mathcal{D}_1$, the number of outliers detected $\mathcal{L}_1 = 2$, and the Isolation factor is computed as $\mathcal{I}_1 = 0.136$. For $\mathcal{D}_2$, the number of outliers detected $\mathcal{L}_2 = 1$, and the Isolation factor is computed as $\mathcal{I}_2 = 0.112$. Considering weights $W_d = W_\theta = W_\mathcal{I} = \frac{1}{3}$, the overall anonymity of System 1, $\mathcal{R}_1 = 0.823$, and for System 2, $\mathcal{R}_2 = 0.832$. Thus, System 2 is slightly more anonymous than System 1, under the chosen set of weights. In System 2 even though the first user is being isolated, the attacker does not suspect this user a lot more than the next 15 users. On the other hand, in System 1, the attacker suspects first 2 users very highly. Therefore, System 2 turns out to be more anonymous than System 1. □

Let us examine the impact of choice of weights upon comparison of systems, using Example 3 in Section II.

Example 3 revisited: Consider two systems with $\mathcal{N}_1 = \mathcal{N}_2 = 11$, where the distributions are as given in Section II. The 3-tuple metric for these systems are given by $(0.927, 0.3, 0.219)$ and $(0.821, 0.3, 0.163)$ respectively. Clearly, $d_1 > d_2$ and $(1 - \mathcal{I}_1) < (1 - \mathcal{I}_2)$.

Case I: Let us consider a system designer's perspective and choose $W_d = 0.8$, $W_\theta = 0.1$ and $W_\mathcal{I} = 0.1$. This results in $\mathcal{R}_1 = 0.893$ and $\mathcal{R}_2 = 0.811$, and System 1 evaluates to be more anonymous than System 2.

Case II: Let us consider an attacker's perspective and choose $W_d = 0.1$, $W_\theta = 0.1$ and $W_\mathcal{I} = 0.8$. This results in $\mathcal{R}_1 = 0.789$ and $\mathcal{R}_2 = 0.823$, and System 2 evaluates to be more anonymous than System 1.

Case III: Let us consider an end user's perspective and choose $W_d = 0.1$, $W_\theta = 0.8$ and $W_\mathcal{I} = 0.1$. This results in $\mathcal{R}_1 = 0.734$ and $\mathcal{R}_2 = 0.728$, and System 1 evaluates to be slightly more anonymous than System 2 (similar to Case I). Since $\theta_1 = \theta_2$, the degree of anonymity provided by the two systems is similar from the end user's perspective. □ Thus, an anonymous system may be more anonymous from the perspective of system designer, but could turn out to be less anonymous for an attacker. Anonymity of a system should be viewed keeping the desired attributes of system in mind.

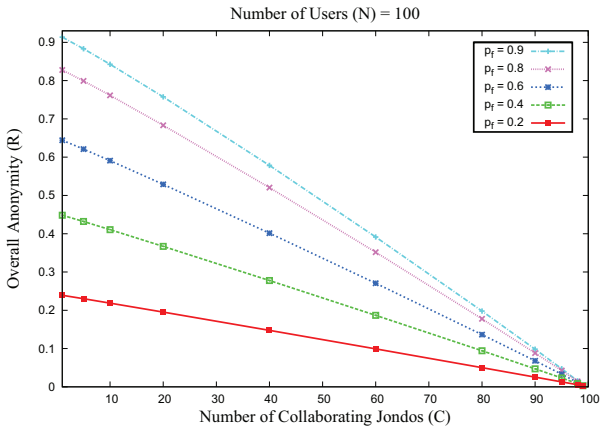

Fig. 1. Overall anonymity for crowds.

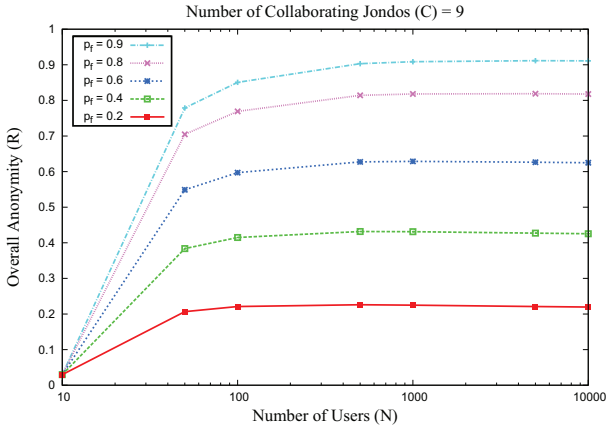

Fig. 2. Anonymity saturates as N increases.

VI. Metric Application to Crowds

Consider the Crowds [6] system with N users, where each user forwards the request to a randomly chosen user with probability p_f, and contacts the server directly with probability $(1 - p_f)$. [6], [10] consider an attack model with C collaborating jondos (or users) in the system. The size of anonymity set equals N, and the probability assigned to each collaborating jondo equals 0. Given that a collaborator is on the path between sender and receiver, the probability that the sender is the first collaborator's immediate predecessor is given by [6]:

$$p_1 = 1 - p_f \frac{N - C - 1}{N}. \tag{8}$$

Thus, one of the users (predecessor of first collaborating jondo on the path) is suspected to be the sender with probability p_1. Assuming, all other non-collaborative users are equally suspected,

$$p_i = \frac{1 - p_1}{N - C - 1} = \frac{p_f}{N}, \ \forall i \in \{2, \ldots, N - C\}. \tag{9}$$

Using this probability distribution, we evaluate the 3-tuple metric for Crowds for various values of p_f, N and C, and

compare these systems under equal weights, $W_d = W_\theta = W_\mathcal{I} = \frac{1}{3}$. Figure 1 depicts the overall anonymity ($\mathcal{R}$) for $N = 100$. As the number of collaborating jondos (C) increases, the anonymity of the system decreases. Further, the anonymity increases as the probability of forwarding (p_f) increases for all values of C, with the increase being higher when C is small. Figure 2 depicts the anonymity of the system under a fixed number of collaborating jondos ($C = 9$). We observe that the anonymity increases with an increase in N for all values of p_f. However, the anonymity of the system saturates as N is increased further, with the saturated anonymity value being quite close to (and greater than) the probability of forwarding (p_f). Thus, in order to achieve an overall anonymity of > 0.75, we should either have $N = 50$ and $p_f = 0.9$ or $N = 100$ and $p_f = 0.8$. When C is small, increasing the probability of forwarding (p_f) would substantially increase the anonymity in Crowds. However, when C is large, increasing the number of users (N) instead (such that C is a smaller fraction of N) would lead to a higher degree of anonymity.

VII. Summary and Conclusions

We proposed a new *isolation* measure, based upon presence of outliers in a distribution, which is critical towards quantifying the overall anonymity of a system. We proposed a three-dimensional approach towards measuring anonymity, and have applied our metric to an existing anonymous system, Crowds. We provided justification for three distinct aspects of the proposed 3-tuple metric, and provided interpretation of the metric in terms of attributes desired in the system.

References

[1] S. Garfinkel and G. Spafford, *Web Security, Privacy & Commerce, Second Ed.* O'Reilly Media, 2001.

[2] R. Geambasu, T. Kohno, A. Levy, and H. M. Levy, "Vanish: Increasing data privacy with self-destructing data," in *Proc.* 18^{th} *USENIX Security Symposium*, Montreal, Canada, Aug. 2009, pp. 299–315.

[3] L. A. Gordon, M. P. Loeb, and T. Sohail, "A framework for using insurance for cyber-risk management," *Communications of the ACM*, vol. 46, pp. 81–85, Mar. 2003.

[4] A. Narayanan and V. Shmatikov, "De-anonymizing social networks," in *Proc.* 30^{th} *IEEE Symposium on Security and Privacy*, Oakland, CA, USA, May 2009, pp. 173–187.

[5] C. Diaz, C. Troncoso, and A. Serjantov, "On the impact of social network profiling on anonymity," in *Proc.* 8^{th} *International Symposium on Privacy Enhancing Technologies (PETS)*, Leuven, Belgium, Jul. 2008, pp. 44–62.

[6] M. K. Reiter and A. D. Rubin, "Crowds: anonymity for web transactions," *ACM Transactions on Information and System Security*, vol. 1, pp. 66–92, Nov. 1998.

[7] D. L. Chaum, "Untraceable electronic mail, return addresses, and digital pseudonyms," *Communications of the ACM*, vol. 24, pp. 84–90, Feb. 1981.

[8] M. G. Reed, P. F. Syverson, and D. M. Goldschlag, "Anonymous connections and onion routing," *IEEE Journal on Selected Areas in Communications*, vol. 16, pp. 482–494, May 1998.

[9] R. Dingledine, N. Mathewson, and P. Syverson, "Tor: The second-generation onion router," in *Proc.* 13^{th} *USENIX Security Symposium*, Aug. 2004.

[10] C. Diaz, S. Seys, J. Claessens, and B. Preneel, "Towards measuring anonymity," in *Proc.* 2^{nd} *International Conference on Privacy Enhancing Technologies (PETS)*, San Francisco, CA, USA, 2002, pp. 54–68.

[11] G. Tóth, Z. Hornák, and F. Vajda, "Measuring anonymity revisited," in *Proc.* 9^{th} *Nordic Workshop on Secure IT Systems*, Espoo, Finland, Nov. 2004, pp. 85–90.

[12] J. Feigenbaum, A. Johnson, and P. Syverson, "Probabilistic analysis of onion routing in a black-box model," in *Proc. ACM workshop on Privacy in electronic society*, Alexandria, VA, USA, 2007, pp. 1–10.

[13] A. Serjantov and G. Danezis, "Towards an information theoretic metric for anonymity," in *Proc.* 2^{nd} *International Conference on Privacy Enhancing Technologies (PETS)*, San Francisco, CA, USA, 2002, pp. 41–53.

[14] P. J. Rousseeuw and A. M. Leroy, *Robust regression and outlier detection*. New York, NY, USA: John Wiley & Sons, Inc., 1987.

[15] B. Peirce, "Criterion for the rejection of doubtful observations," *Astronomical Journal*, vol. II, pp. 161–163, Jul. 1852.

[16] B. A. Gould, "On peirce's criterion for the rejection of doubtful observations, with tables for facilitating its application," *Astronomical Journal*, vol. IV, pp. 81–87, Apr. 1855.

[17] S. Ross, "Peirce's criterion for the elimination of suspect experimental data," *Journal of Engineering Technology*, vol. 20, 2003.

[18] N. Jaggi, U. M. Reddy, and R. Bagai, "A three-dimensional approach towards measuring sender anonymity," Technical Report, available online at http://www.cs.wichita.edu/~neeraj/techrep-anonymity.pdf, Nov. 2010.

Appendix

Outlier Detection: We briefly discuss the algorithm used to detect outliers in a set of observations. The algorithm starts by assuming that there is at least one outlier in the set. Using some calculations, this hypothesis is either accepted or rejected. If accepted, the algorithm assumes that there are at least two outliers in the set, and so on. The algorithm stops when a hypothesis gets rejected, and the number of outliers is declared to be the value corresponding to the previous iteration. Let N denote the total number of observations, and n denote the total number of suspected observations during an iteration of the algorithm. The sample standard deviation is denoted ε, and the error threshold used in outlier detection is given by $x \cdot \varepsilon$. $\lambda \cdot \varepsilon$ is the mean error after n observations have been declared outliers. To declare n observations as outliers, the following inequality should be satisfied [16]

$$\lambda^{N-n}\ e^{\frac{1}{2}\,n\,(x^2-1)}\ (\psi\, x)^y < Q^N, \tag{10}$$

where $\psi\, x = \frac{2}{\sqrt{2\pi}} \int_x^\infty e^{-\frac{1}{2}\,x^2}$, and

$$Q^N = \frac{n^n\,(N-n)^{N-n}}{N^N}, \quad \lambda^{N-n}\,R^n = Q^N. \tag{11}$$

$$x^2 = 1 + \frac{N-n}{n}(1-\lambda^2), \quad R = e^{\frac{1}{2}\,(x^2-1)}\,\psi\, x. \tag{12}$$

Assuming $n = 1$, Q is computed using (11). Now, assuming $\log_{10}R = 9.2$ (infact, $\log_{10}R = 10 - 9.2$, and we are just following the same notation as in [15], [16]), λ is computed using (11) and x is computed using (12). Then, R is obtained using (12) (an appropriate table can also be used [17]), and this value of R is used to recompute λ and x. This procedure is followed until the value of x converges. The converged value of x is used to define the error threshold. The observation i is declared outlier if $|p_i - \bar{p}| > x \cdot \varepsilon$. In the next iteration, we assume $n = 2$, compute x and check for outliers. The algorithm detects all outliers in the set, deviating in either direction from the mean. Since, we are interested only in observations whose value is significantly larger, we trim the output of the algorithm, and choose only the outliers which are greater than the mean.