

Optimizing the Design Parameters of Threshold Pool Mixes for Anonymity and Delay

David Rebollo-Monedero^{*,a}, Javier Parra-Arnau^a, Jordi Forné^a, Claudia Diaz^b

^a*Department of Telematics Engineering, Universitat Politècnica de Catalunya (UPC)*

C. Jordi Girona 1-3, 08034 Barcelona, Spain

^b*Katholieke Universiteit Leuven, ESAT/SCD/IBBT/COSIC*

Kasteelpark Arenberg 10, B-3001 Leuven-Heverlee, Belgium

Abstract

The provision of content confidentiality via message encryption is by no means sufficient when facing the significant privacy risks present in online communications. Indeed, the privacy literature abounds with examples of traffic analysis techniques aimed to reveal a great deal of information, merely from the knowledge, even if probabilistic, of who is communicating with whom, when, and how frequently. Anonymous-communication systems emerge as a response against such traffic analysis. The technology underlying many of these systems consists in mixes, particularly threshold pool mixes, essentially trusted nodes that receive, store, reorder and delay messages in batches. However, the anonymity gained from the statistical difficulty to link incoming and outgoing messages comes at the expense of introducing a potentially costly delay in those messages.

In this paper we address the design of such mixes in a systematic fashion, by defining quantitative measures of both anonymity and delay, and by mathematically formalizing practical design decisions as a multiobjective optimization problem. Our extensive theoretical analysis finds the optimal mix parametrization and characterizes the optimal trade-off between the contrasting aspects of anonymity and delay, for two information-theoretic measures of anonymity. Experimental results show that mix optimization may lead to substantial delay reductions for a fixed anonymity.

Key words: Anonymous communications, threshold pool mix, optimal design, Shannon entropy, min-entropy, delay

1. Introduction

During the last two decades, the Internet and the Web have been gradually integrating into people's daily lives and have enabled new forms of communication such as e-mail, instant messaging and social networking. The so-called network of networks [1] not only has become an essential communication channel but also has transformed people's habits —online shopping, electronic voting and streaming media are just other examples of services and applications built upon this network.

Despite the myriad of benefits that the Internet has brought to users, it has also laid stress on the need for privacy protection. One of the reasons behind this is that the Internet, as many other data communication networks, requires that every user be identified by a unique address, in order for messages to be routed through the network. Internet service providers (ISPs) are precisely in charge of allocating addresses to users and keeping the correspondence between user identifiers and addresses. In this manner, users wishing to communicate through the Internet

just need to attach the source and destination addresses to the message to be sent. On the one hand, these addresses enable the intermediary entities (switches, routers, firewalls) involved in the communication process to forward these messages until the destination address is reached. But on the other hand, since the addresses are transmitted in clear, the entities themselves or any adversary capable of intercepting the messages may ascertain who is communicating with whom. In other words, the inherent operation of this network poses serious privacy concerns. Imagine, for example, a user who browses a Web forum on pregnancy tips from the office, during her lunch break. Even if the Web server is trustworthy and the contents of the information exchanged are encrypted, a privacy adversary, perhaps a human resources manager at her company, might still capture and inspect the packet headers to infer sensitive interests of the user, and take discriminative action against her possible promotion.

Clearly, the use of encryption techniques is not enough to mitigate these privacy risks. Concealing the content of messages hinders adversaries in their efforts to learn the information users exchange, but does not prevent those adversaries from unveiling who is communicating with whom, when, or how frequently. Motivated by this, privacy-enhancing technologies referred to as *anonymous communi-*

*Corresponding author. Tel.: +34 93 401 7027.

Email addresses: david.rebollo@entel.upc.edu (David Rebollo-Monedero), javier.parra@entel.upc.edu (Javier Parra-Arnau), jforne@entel.upc.edu (Jordi Forné), claudia.diaz@esat.kuleuven.be (Claudia Diaz)

ation systems have emerged [12]. The first proposal of an anonymous communication system was the *mix* proposed by Chaum in 1981 to provide anonymous email services [4]. Since then, many other anonymous communication systems have been proposed. Some of these are variations of Chaum’s mix design, meant to be used for message-based applications with toleration to latency, such as email [6, 23, 34, 9, 16, 13], while others focus on providing bidirectional anonymous channels for applications with low-latency requirements, such as web browsing [33, 22, 32, 17].

The goal of anonymous communication systems is to prevent an adversary from linking an outgoing message to its corresponding input message. To achieve this, mix-based systems delay and reorder messages, in addition to modifying cryptographically their contents, so that inputs and outputs cannot be correlated on the basis of content or timing. A popular family of mix types is known as *threshold pool mixes* [15]. These mixes collect a number of incoming messages, store them in the internal memory, and output a fraction of them once the number of collected messages kept reaches a certain threshold.

Certainly, delaying messages negatively affects the usability of these systems and hence imposes a cost on them. Yet, higher delays provide users with a higher degree of message unlinkability. Thus, mix systems pose an inherent trade-off between anonymity and delay. This compromise between anonymity and delay has not been studied in the anonymous communication literature in a mathematical, formal fashion. Given a maximum expected delay that users are willing to tolerate, there are no methods to optimize the mix parameters such that anonymity is maximized while respecting delay constraints. Similarly, given anonymity requirements, there are not methods that can assist mix designers in selecting parameters that meet those requirements while minimizing the expected message delay.

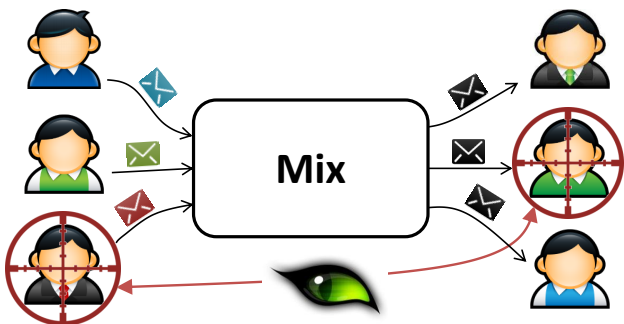


Figure 1: Many anonymous-communication systems are built upon the idea of Chaum’s mix. Essentially, a mix can be seen as a black box that forwards messages in such a way that prevents an adversary from linking an outgoing message to its corresponding input message.

1.1. Contribution and Organization

We have motivated the necessity of anonymous communication systems to face the important privacy risks derived

from online traffic analysis. Concordantly, the object of this work is to address the problem of designing threshold pool mixes, a key building block in the area of anonymous communications, in a manner that contemplates the optimal trade-off between two contrasting aspects: anonymity and delay. The first aspect, the anonymity of the peers involved in the communication, gained through the statistical difficulty to link incoming and outgoing messages, is something we would like to maximize. The second aspect, which we wish to minimize, is the potentially costly delay experienced by the messages sent through the mix. Thus we are faced with the inescapable compromise of maximizing anonymity at the cost of delay, or minimizing delay for a fixed anonymity.

We approach the issue in a systematic fashion, drawing upon the methodology of multiobjective optimization, long established in scientific fields such as data compression or economics, but perhaps less so in the area of information privacy. More precisely, we consider with a standard adversary model (global passive adversary) and adopt three quantifiable measures of anonymity in the literature, Shannon’s entropy, min-entropy and collision entropy. We then proceed to mathematically formalize practical decisions in the design of threshold pool mixes in terms of a multiobjective optimization problem. Our extensive theoretical analysis finds the optimal mix parametrization and characterizes the optimal trade-off between the aspects of anonymity and delay, for the two information-theoretic measures of anonymity adopted. Our results show that mix optimization may lead to substantial delay reductions for a given level of anonymity.

It must be pointed out that the choice of anonymity metrics in the formulation of our model is driven by our own recent work on the measurement of privacy [31]. In the cited work, we aim to offer a unified perspective on a number of privacy metrics, under their conceptual connection with the stochastic notion of an adversary’s estimation error in ascertaining confidential information, modeled as a random event, or the adversary’s effort in removing any residual uncertainty. For the sake of readability and self-containment, we provide here a short summary (Sec. 3.3) on the interpretation of various types of entropies as anonymity metrics, adapted from that work for the specific context of the current paper.

Finally, it should be noted that our contribution on the design of threshold pool mixes may be regarded from two perspectives. First, as a step towards the design of such mixes in keeping with the philosophy of *usable privacy*, that is, practical privacy enhancing technology that realistically takes into account the cost of its presence. Secondly, as an illustration of the applicability of formalized multiobjective optimization techniques to the still emerging field of privacy in information systems.

The remainder of this paper is organized as follows. Sec. 2 discusses the state of art in anonymous-communication systems and anonymity metrics. After some mathematical preliminaries and a concise exposition on the in-

terpretation of various types of entropies as an anonymity metrics, Sec. 3 presents a formulation of the problem of optimizing the anonymity-delay trade-off in threshold pool mixes. Sec. 4 investigates this problem theoretically, while Sec. 5 illustrates numerically the main results. Lastly, conclusions are drawn in Sec. 6.

2. State of the Art

In this section, we provide a review of anonymous-communication systems. We first examine systems based on variations of the original mix devised by Chaum, including threshold pool mixes. And secondly, we describe approaches that do not rely on the principles of mixes and are designed for applications with low-latency constraints, such as web browsing. Lastly, we touch upon the most widely used anonymity metrics in the literature.

2.1. Anonymous Communication Systems

As described in Sec. 1, mixes delay and reorder messages to preserve the *unlinkability* [28] of incoming and outgoing messages. The ultimate purpose is to protect the *anonymity* of users exchanging messages, or in other words, to conceal *who talks to whom*. The idea behind Chaum’s mix is conceptually simple. Users wishing to anonymously send messages to other peers encrypt these messages with the public key of the mix. The mix collects a number of these encrypted messages and stores them in its internal memory. When the number of collected messages reaches a certain threshold, the messages are decrypted, the information about senders is removed, and the mix forwards *all* the messages to their recipients in a random order.

In the literature, this process of collecting, storing and forwarding messages when a condition is satisfied is normally referred to as a *round*. An important group of mixes called *pool* mixes operate on this basis. Depending on the *flushing* condition, we may distinguish different types of pool mixes [15]. Possibly, the most relevant form of pool mixes are *threshold* pool mixes, where the condition is imposed on the number of messages stored, as in the case of Chaum’s mixes. The main difference is that threshold pool mixes do not flush all messages in each round, but only a fraction instead [16]. This strategy increases the level of anonymity as the set of possible incoming messages linkable to an outgoing target message includes all those messages that entered the mix before the target message was flushed. On the downside, incoming messages may be stored in the mix for an arbitrarily long period of time, which degrades the usability of the system.

Another important group of pool mixes outputs messages based on time [40]. Essentially, these *timed* mixes forward all messages kept in the memory every fixed interval of time called timeout. The major advantage of these mixes is that the delay experienced by messages is upper bounded, in contrast to the case of threshold pool

mixes. The flip side is that the unlinkability between incoming and outgoing messages may be seriously compromised when the number of messages arriving in that interval of time is small. Motivated by this, mix designs implement a combination of the strategies based on threshold and those based on time. Namely, these systems flush (all or a fraction of) messages when a timeout expires, provided that the number of messages stored meets a predefined threshold [27].

An alternative to pool mixes are the mixes based on the concept of *stop-and-go*, known as *continuous* mixes [23]. Specifically, this approach abandons the idea of rounds and gives the user the possibility of specifying the time that their messages will be stored in the mix. To this end, for each message to be sent the sender selects a random delay from an exponential distribution. This information is then attached to the message, which is encrypted with the mix’s public key and then sent to the mix. Once the mix decrypts the message, the mix keeps it for the time specified by the user before forwarding it to the message recipient.

The use of networks of mixes has also been thoroughly studied in the literature. The main reason to route over multiple mixes is to limit the trust that is placed on each single mix. In order to trace messages, an adversary must compromise all the mixes along the path. Depending on the network topology, we may classify the existent approaches into *cascade mixes*, *free-route networks* and *restricted-route networks*. The application of cascade mixes was already suggested by Chaum in his original work [4]. Fundamentally, this approach contemplates the concatenation of mixes in a chain. In contrast to this alternative where messages are routed through a fixed path, free-route networks recommend that users choose random paths to route their own messages [34]. In the end, restricted-route networks consider the case where every mix in the network is connected to a reduced number of neighboring mixes [9].

Active attacks on mixes refer to adversaries who are not only able to inspect incoming and outgoing traffic, but also able to insert, remove and delay messages. Specific attacks of this sort include *blending attacks* [39], in which the adversary ensures, by means of delaying (honest) messages and generating (malicious) messages, that a target message can be fully traceable when routed through the mix. The use of randomized mixing strategies together with dummy traffic mitigates, to some extent, this attack [16].

Apart from the systems based on mixes, there exist other approaches that also seek to provide anonymize communication channels for applications with low-latency requirements. Given these latency requirements, these systems do not delay messages and are thus vulnerable to timing attacks deployed by global passive adversaries. They thus provide anonymity only towards weaker adversaries, such as those who control a subset of the routers that form the anonymous network, but cannot observe the communications of nodes they do not control.

One such system is *Crowds*, which aims at protecting user anonymity when accessing the Web [33]. *Crowds* contemplates that a group of users wanting to anonymously browse the Web will collaborate to submit their requests, by forming a ‘crowd’. A user who decides to send a request to a Web server, selects first a member of the crowd at random and then forwards the request to that member. When this member receives the request, it flips a biased coin to determine whether to further send the request to another member or to submit it to the Web server. This process is repeated until the request is finally relayed to the intended destination. As a result, the Web server, as well as crowd members forwarding the request, cannot ascertain the identity of the true sender, that is, the member who initiated the request. Path lengths in *Crowds* follow a geometric distribution, which has been proven optimal in terms of anonymity for a given mean latency [11].

Another family of low-latency anonymity systems, that do not delay or reorder messages, are based on *onion routing* [22, 32, 17]. In these systems, users establish circuits that pass through multiple onion routers. Packets being sent through those circuits are encrypted in a layered (*onion*) fashion, so that each of the routers peels (in the upstream direction), or adds (in the downstream direction) a layer of encryption. The second-generation version of onion routing, Tor [17], is, with more than half a million daily users, the most widely used anonymous communication system. For performance reasons, Tor nodes do not delay the packets that they route, and thus it does not protect against adversaries that can observe both ends of the communication (the link between the sender and the first mix, and the link between the last mix and the final recipient). The security of Tor thus relies on geographic diversity, and it aims to protect against adversaries who can only observe some parts of the network.

2.2. Anonymity Metrics

This section assumes familiarity with certain types of entropies, formally defined and briefly commented later, in the mathematical review of Sec. 3.2.

A great effort has been devoted in prior work to the investigation of privacy metrics in general and anonymity metrics in particular. Some of the best-known privacy metrics come from statistical disclosure control (SDC), a research area that deals with the problem of disseminating data about individuals without compromising their privacy. Among these metrics, the most popular is *k-anonymity*, which was first proposed in [45, 36]. Later, numerous extensions and enhancements were introduced in an attempt to address the limitations of this proposal. *p*-Sensitive [49], *l*-diversity [25], *t*-closeness [24], an average version of *t*-closeness [30], δ -disclosure [3] and differential privacy [18] are some of these approaches.

In the specific field of anonymous communication systems, many proposals focus on measuring the extent to which these systems are effective. A key point is that the degree of anonymity achieved depends on the capabilities

of the adversary, and often anonymity metrics are tailored to the threat model assumptions. A study of adversarial models for these systems can be found in [29].

One of the earliest proposals for assessing the degree of anonymity provided by these systems appears in the *Crowds* protocol [33], explained above. Concisely, the cited work defines three anonymity requirements for each candidate sender of a message, in the form of probabilistic events from the point of view of the adversary. The strongest requirement, “beyond suspicion”, is fulfilled whenever all candidates are equally likely. Next in the list is the requirement of “probable innocence”, demanding that the given user’s likelihood be less than $1/2$. The weakest requirement, “possible innocence”, boils down to stipulating that this likelihood simply not be 1. One of the main theoretical results in that work states conditions under which the intermediate requirement is satisfied for all senders.

In the context of mix systems, [23] defined the *anonymity set* of users as the set of possible senders of a given message, or recipients, in the sense that the likelihood of them fulfilling the role in question is nonzero. A simple measure of anonymity was proposed by [2], namely the logarithm of the number of users involved in the communication, that is, the Hartley entropy of the anonymity set. The main drawback of this metric is that it does not contemplate the probabilistic information that an adversary may obtain about users when observing the system. In other words, this approach ignores the fact that certain users may be more likely to be the senders of a particular message.

Several approaches have considered the use of information-theoretic quantities to evaluate anonymous communication systems. The most commonly used are those proposed in [14, 38], in which the degree of anonymity observable by an adversary is measured essentially as the Shannon entropy of the probability distribution of possible senders of a given message. The use of entropy as a measure of privacy, however, is by no means new. As a matter of fact, Shannon’s work in the fifties introduced the concept of *equivocation* as the conditional entropy of a private message given an observed cryptogram [41], later used in the formulation of the problem of the wiretap channel [50, 8] as a measure of confidentiality.

Still in the case of information-theoretic measures, [44] formalizes the notion of unlinkability by using Shannon’s entropy.

On the other hand, [48, 47] argue that a worst-case metric should be considered instead of Shannon’s entropy, since the latter contemplates an average case. The authors refer to this worst-case metric as *local anonymity*, essentially equivalent to min-entropy, and concordantly define the *source hiding* property as the requirement that no sender probability exceed a given threshold. Another approach [42] proposes a method for quantifying the property of *relationship anonymity*, as defined in [28]. More specifically, the authors make use of Shannon’s entropy and min-entropy for measuring this property. Similarly, [5] evaluates Shannon’s entropy, min-entropy and Hart-

ley’s entropy as anonymity metrics, and proposes then to use Rényi’s entropy, which may be regarded as a generalization of those three metrics.

Alternatives include possibilistic —instead of probabilistic— approaches, such as those suggested in [46, 26, 20]. According to these metrics, subjects are considered anonymous if an adversary cannot determine their actions with absolute certainty. Further, [19] proposes a combinatorial anonymity metric that counts the number of possible one-to-one correspondences between a set of senders and a set of receivers, by means of the permanent of the matrix of adjacencies of the associated bipartite graph, consistent with message timing observations ruling out some of the permutations. It must be stressed that probability distributions weighting such possibilities are not considered, or from a mathematically equivalent perspective, that those probabilities are considered equally likely. Another difference with respect to most metrics based on probabilities is that this metric is directly defined on a group of consistent matchings between senders and receivers, rather than defined on the set of senders or receivers corresponding to one given message. Some limitations and extensions of this approach may be found in [21].

For the purposes of this paper, we may effectively conclude that in the scenario of anonymous-communication systems, the anonymity measures may be classified essentially into two groups, namely probabilistic and possibilistic metrics. In the latter group we include any metrics that do not probabilistically weight possibilities, even if they resort to sophisticated combinatorial methods. While most information-theoretic quantities are probabilistic, including Shannon’s entropy and min-entropy, Hartley’s entropy is clearly possibilistic. Because Hartley’s entropy is the logarithm of the anonymity set size, both metrics are fundamentally the same. Rényi’s entropies encompass these three entropies, and except for the singular case of Hartley’s, constitute a parametric family of probabilistic measures. The concepts of probable innocence in Crowds, local anonymity and source hiding, are essentially equivalent to min-entropy, as they all involve the most likely sender of a message. This preliminary classification is represented in Fig. 2.

Acknowledging the general relevance of information-theoretic concepts in the privacy literature, and particularly inspired by our own work on the specific subject of privacy metrics [31], whose relation to the present paper was explained in Sec. 1.1, we shall contemplate here Shannon’s entropy and min-entropy as anonymity metrics.

3. Formal Problem Statement

This section is devoted to the description of our adversarial model, and the mathematical formulation of the problem of optimizing the anonymity-delay trade-off of a threshold pool mix. Our formulation is prefaced by a quick, statistical and information-theoretic background, in which we revise the definition of Shannon’s entropy and min-entropy,

and by a succinct interpretation of these type of entropies, adapted from our own work [31] on the specific subject of privacy metrics, to the present context of anonymous-communication systems. The solution of the problem formulated is the object of the next section.

3.1. Adversarial Model

We adhere to an *adversarial model* commonly adopted in the literature of anonymous communication reviewed in Sec. 2. The specification of our adversary model is detailed enough for the purposes of the scope of the paper, but further technicalities on the topic may be found in the citations included in our state-of-the-art section. In our model, we assume that a privacy adversary, unable to compromise the threshold pool mix itself, but able to analyze its input and output traffic, wishes to ascertain the correspondence between its input and output messages, to ultimately unveil who is communicating with whom. Messages are not only forwarded in batches, but also encrypted with different session keys as well as padded, in order to hinder such adversaries in their efforts to establish message correspondence based on timing, bitwise or size comparisons. However, the adversary is assumed to know the principles of operation of the mix, and to be capable of deducing the probabilities of input and output message correspondence.

3.2. Notation and Mathematical Preliminaries

Throughout the paper, we shall follow the convention of using uppercase letters for *random variables* (r.v.’s), and lowercase letters for particular values they take on, assumed to belong to a *discrete alphabet*. *Probability mass functions* (PMFs) and *probability density functions* (PDFs) are denoted by p and subindexed by the corresponding r.v., or by an index on the alphabet. The *expectation* and the *variance* operators are denoted by E and Var , respectively. We occasionally prefer to denote the expectation and the variance of a r.v. X more compactly as μ_X and σ_X^2 . Expectations can model the special case of (weighted) averages over a countable set of data points $\{x_i\}_{i=1}^{\infty}$, simply by defining an r.v. X distributed over this set according to a PMF $p_X(x_i) = p_i$, so that $E X = \sum_{i=1}^{\infty} p_i x_i$.

Recall [7] that the *Shannon entropy* $H_1(X)$ of an r.v. X with PMF p_X is defined as

$$H_1(X) = -E \log p_X(X) = -\sum_{i=1}^{\infty} p_i \log p_i,$$

in terms of the above example of discrete alphabet, under the convention that $0 \log 0 = 0$. Recall also that this entropy is a measure of the uncertainty of the outcome of a random variable distributed according to such PMF. All theoretical results expressed in terms of logarithms are valid for any base, but the customary basis of 2 is chosen in all numerical computations, and concordantly entropy units are *bits*. Shannon entropy belongs in fact to a more general class of functionals known as *Rényi entropies* [35],

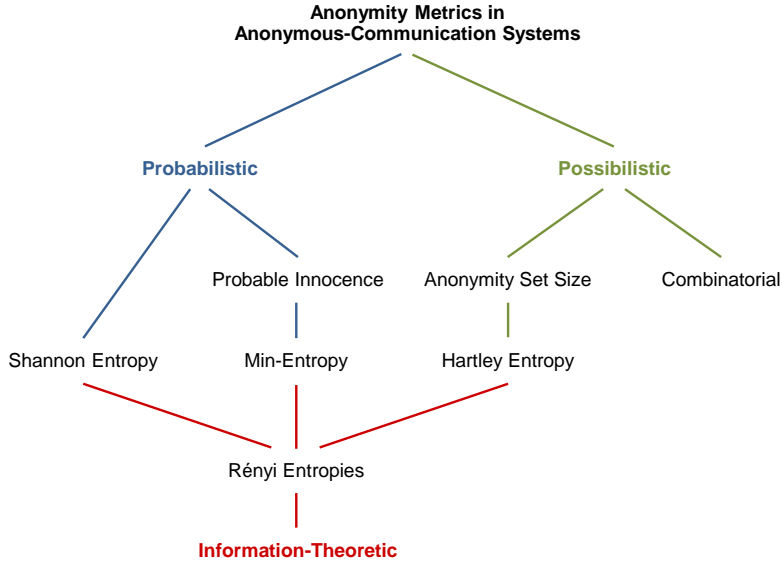


Figure 2: A preliminary classification of anonymity metrics in anonymous-communication systems.

all of which satisfy certain additivity axiom regarding statistically independent r.v.'s. Precisely, the Rényi entropy $H_\alpha(X)$ of order α of the same discrete r.v. X , with PMF $p_X(x_i) = p_i$ as above, is defined for any real $0 < \alpha \neq 1$, as

$$H_\alpha(X) = -\log \sqrt[\alpha-1]{\mathbb{E}[p_X(X)^{\alpha-1}]} = \frac{1}{1-\alpha} \log \sum_{i=1}^{\infty} p_i^\alpha,$$

definition that may be construed as the negative logarithm of the (generalized) *power mean* of order $\alpha - 1$ of the PMF, weighted by itself. The definition is commonly extended for the extreme values of the order $\alpha = 0, 1, \infty$, simply by taking limits of the above expression. In this manner, Shannon entropy is obtained when $\alpha = 1$, hence subindex in the notation $H_1(X)$ introduced earlier. It can be shown that for a given distribution, Rényi entropies are nonincreasing with α , and that for each α , they are maximized, among all distributions on $\{1, \dots, n\}$, by the *uniform distribution* $p_i = 1/n$ for all $i = 1, \dots, n$, for which $H_\alpha(X) = \log n$.

The extreme case when $\alpha = \infty$ yields a type of entropy known as *min-entropy*. In the previous example with PMF p_i , denote $p_{\max} = \max_{i=1,2,\dots} p_i$. Min-entropy is equivalently defined as

$$H_\infty(X) = \min_{i=1,2,\dots} -\log p_i = -\log p_{\max}.$$

A third type of entropy of particular significance within the Rényi family, mentioned in the state-of-the-art section and obtained for $\alpha = 0$, is *Hartley's entropy*, namely the logarithm of the cardinality of the alphabet (after removal without loss of generality of values with zero probability), that is, $H_0(X) = \log n$ in our finite-alphabet example. Finally, $\alpha = 2$ yields the *collision entropy*

$$H_2(X) = -\log \mathbb{E} p_X(X) = -\log \sum_{i=1}^{\infty} p_i^2,$$

which may be seen as the probability that two independent copies of X coincide.

For these special examples of entropy, the monotonicity and maximization properties of Rényi entropies already stated boil down to

$$H_0(X) \geq H_1(X) \geq H_2(X) \geq H_\infty(X), \quad (1)$$

with equality if and only if X is uniformly distributed on a finite alphabet, i.e., if and only if $p_i = 1/n$ for all $i = 1, \dots, n$.

Recall that commonly, message arrivals in communications systems are modeled by a *Poisson process* with *rate* λ , and equivalently interarrival times T are assumed to be independent, identically distributed according to an *exponential distribution* with parameter λ and PDF $p_T(t) = \lambda e^{-\lambda t}$, and concordantly, $\mu_T = 1/\lambda = \sigma_T$. An r.v. D is said to have an *Erlang distribution* when characterized by a PDF of the form

$$p_D(d) = \frac{\lambda^k d^{k-1}}{(k-1)!} e^{-\lambda d},$$

equivalently interpreted as the sum of k independent exponential r.v.'s with a common parameter λ .

Our theoretical analysis will exploit well-known properties of geometric r.v.'s, which model for example the number of tosses of a coin until the first head occurs. Rigorously, a discrete r.v. X *geometrically distributed* on $\{1, 2, \dots\}$ with parameter p , counts the number of independent Bernoulli trials needed to obtain one success, which occurs with probability p . Recall that $p_X(x) = (1-p)^{x-1} p$, $\mathbb{E} X = 1/p$, $H_1(X) = H_1(p)/p$ and $H_\infty(X) = p$, where $H_1(p)$ denotes the Shannon entropy of a binary r.v. of parameter p :

$$H_1(p) = -p \log p - (1-p) \log(1-p).$$

It is also common to define the alphabet starting with 0 in lieu of 1; in that case, $p_X = (1-p)^x p$, $\mathbb{E} X = 1/p - 1$ and the entropies remain the same.

3.3. Entropies as Anonymity Metrics for Pool Mixes

In the state-of-the-art subsection on anonymity metrics, Sec. 2.2, we briefly introduced three measures of uncertainty and information, namely Hartley’s, min- and Shannon’s entropies, whose precise definition we recalled among the mathematical preliminaries of Sec. 3.2. These measures of uncertainty of a random event are essentially scalar-valued functions of probability distributions across a set of possible outcomes. Their particular significance and wide application in the fields of information theory, statistics and engineering is unquestionable, but they are in fact found in a larger variety of fields, including for example demography and ecology, in the form of diversity indexes. In the context of anonymous-communication systems in general and pool mixes in particular, the knowledge of the privacy adversary may be modeled by a probability distribution on the possible senders of a given message, enabling the interpretation of these measures under the conceptual perspective of an adversary’s estimation error in ascertaining the outcome of a random event, or effort in removing any residual uncertainty. The following interpretations are based mainly on [31] and adapted to the subject of this paper.

- First, Hartley’s entropy is a possibilistic metric, in the sense that it disregards the likelihood of the values of an r.v., whereas both Shannon’s entropy and min-entropy are probabilistic. In principle, one could measure the attained degree of anonymity, merely by the cardinality of the set of candidate senders, or equivalently, by the logarithm of such cardinality. Loosely speaking, Hartley’s entropy may be regarded as a *best-case* metric from the point of view of users (worst for adversaries), in the sense that it represents a privacy adversary’s thorough effort in considering any and all possibilities, regardless of their likelihood. In pool mixes, however, the set of candidate output messages for a given input may be infinite, rendering Hartley’s entropy inappropriate.
- Secondly, min-entropy may be connected to the error in *maximum a posteriori* (MAP) estimation, where the adversary simply guesses the most likely outcome. More specifically, the maximum probability p_{\max} of a random event X , its min-entropy $H_{\infty}(X)$ and the probability of error $\epsilon_{\text{MAP}} = 1 - p_{\max}$ in MAP estimation, are bijectively related and thus essentially equivalent, via the relationship

$$\epsilon_{\text{MAP}} = 1 - 2^{-H_{\infty}(X)}.$$

This may be construed as a *worst-case* metric, in the sense that users are only concerned with the most vulnerable statistical link between senders and messages^(a). Although not specific to pool mixes, in our review of the state of the art on anonymity metrics,

Sec. 2.2, we mentioned the requirement of probable innocence. Clearly, probable innocence for all possible candidates of a given message is equivalent to $p_{\max} < 1/2$, in turn equivalent to $\epsilon_{\text{MAX}} > 1/2$, and finally equivalent to $H_{\infty}(X) > 1$ bit.

- Last but not least, a well-known interpretation of Shannon’s entropy refers to the game of 20 questions, in which one player must guess what the other is thinking through a series of yes/no questions, as quickly as possible. Informally, Shannon’s entropy is a lower bound on—and often good approximation to the minimum of—the average number of binary questions regarding the nature of possible outcomes of an event, to determine which one in fact has come to pass, intelligently exploiting their known probabilities. Inspired by the above interpretation of Shannon entropy as the effective uncertainty within a set endowed with a probability distribution, [38] proposed it as a measure of anonymity. This measure takes into account the underlying probability distribution in its entirety, between the extremes posed by the previous two, yielding a quantity bounded according to (1). For this reason, one may think of it as an *average-case* metric. An alternative interpretation of Shannon’s entropy in the context of privacy is offered in [31] on the basis of the *asymptotic equipartition property* (AEP) and the concept of *typical set* [7]. Precisely, for an adversary jointly estimating sequences of uncertain outcomes, rather than individually guessing single occurrences, Shannon’s entropy is a measure of the effective cardinality of the set of candidate sequences.

3.4. Formulation of the Threshold Pool Mix Operation and Performance Metrics

Consider a *threshold pool mix* in steady state, forwarding batches of k messages at a time, with a buffer of $n \geq k$ messages, thus keeping at least $m = n - k \geq 0$ messages at a given time. More precisely, the manner in which the mix operates is the following. At some point, the mix contains its minimum of m messages, and waits for k additional messages to come in, for a total of $n = m + k$. Once this threshold n is reached, k messages are drawn randomly, independently and uniformly among all stored messages, regardless of the order in which they arrived, and sent out simultaneously. This leaves the mix again at its minimum m , and the process is repeated from that point on. In the

est after a statistically related observation is available, akin to the Bayesian statistical concepts of posterior and prior probabilities. Under this interpretation, conditional entropy is understood as remaining uncertainty, reducing the initial uncertainty prior to the aforementioned observation, naturally terming the resulting difference as information leaked. The cited work discusses a possible definition of conditional min-entropy as the negative logarithm of the average across all observations of the posterior probability most vulnerable to guess.

^(a)In a broad overview on the information-theoretic measuring of the flow of information, [43] recalls the well-known interpretation of the conditional Shannon entropy of certain information of inter-

extreme case when $m = 0$, that is, $k = n$, the threshold pool mix forwards all of its messages in every batch, thus becoming a *Chaumian mix*. The trivial event of the *absence of a mix* may be readily represented by $k = 1 = n$.

Fig. 3 depicts a portion of the operation of a threshold pool mix with $k = 2$, $n = 5$ and $m = 3$, where an adversary strives to link the target message labeled as 0 to one of the outgoing messages. Note that our reference time is considered to be the arrival time of this target message. According to this, we assign the label 1 to the first message arriving after this target message, and -1 to the last message that arrived before it. The other messages are labeled analogously, from -2 to 3. As depicted in this figure, after forwarding several messages, the mix finally outputs the target message. In this case, message 0, which we choose as reference, is delayed one output batch and one message within this batch, thus experiencing a delay of three messages. As a result of the modification of the flow of messages, an adversary merely observing input and output timings cannot discern the destination of the targeted incoming message with absolute certainty. This figure will be used repeatedly in this section as a running example to clarify and illustrate the introduction of a number of formalisms exploited in our theoretical analysis, presented later in Sec. 4. Concordantly, the detailed explanation of the figure at hand will be provided as we delve into the remainder of the current section.

Define the r.v. Δ as the *delay in number of messages* incurred by a message from its arrival in the mix until it is flushed; zero whenever the message is forwarded immediately, a positive integer otherwise. In the example of Fig. 3, message 0 experiences a delay of $\Delta = 3$ messages, as it is flushed when message 3 arrives. To be precise, in any statistical analysis of Δ we shall assume that the mix is steady state, and that the corresponding message is chosen either uniformly at random within an input batch, or deterministically from its output. This latter technicality is necessary because of the cyclic nature of the mix. To see this, take the simple case of $k = n = 3$. The sequence of delays incurred for each of the input messages will be periodic and deterministic, specifically $\dots, 2, 1, 0, 2, 1, 0, \dots$. If Δ is not defined for a deterministically given input message, but for a message chosen uniformly at random within the input batch, then Δ becomes uniformly distributed, identically for all choices across all batches. From the perspective of output messages, because messages are shuffled before they are flushed, this technicality does not apply. Finally, because output shuffling is uniform, our definition for Δ from the perspective of uniformly randomly chosen inputs and deterministic outputs is one and the same. Under these convenient assumptions, the probability distribution of Δ , whether referring to input or output messages, will be identical from any other message (albeit not necessarily statistically independent). We would like to stress that $\Delta = 0$ represents the event according to which we randomly choose an input message causing the mix to flush immediately upon its arrival, and to select

it as any of its outputs. From the perspective of a given output message, $\Delta = 0$ indexes the input triggering the corresponding batch.

Define the *expected delay* (in number of messages)

$$\bar{\delta} = \mu_{\Delta} = \mathbb{E} \Delta. \quad (2)$$

Let the *delay in time units* be modeled by a real-valued r.v. D , related to the discrete r.v. Δ via the message inter-arrival times T_i of future input messages, also real-valued r.v.'s, simply by

$$D = \sum_{i=1}^{\Delta} T_i,$$

with $D = 0$ whenever $\Delta = 0$. The following proposition shows the existence of a simple proportionality relationship between delay in number of messages and delay in time units, owing to the fact that both quantities are measured as expectations.

Proposition 0 (Delay in Time Units vs. Messages). *Reasonably, we shall assume that the interarrival times T_i and the delay in messages Δ are pairwise statistically independent. Suppose further that the interarrival times have a common expectation $\mathbb{E} T_i = \mu_T = \bar{t}$, a common variance $\text{Var} T_i = \sigma_T^2$, and that they are uncorrelated.*

- (i) *Under those mild assumptions, the expected delay in time units μ_D is proportional to the expected delay in number of messages μ_{Δ} ; precisely,*

$$\mu_D = \mu_T \mu_{\Delta} = \bar{t} \bar{\delta}.$$

- (ii) *Similarly, the variance of the delay in time units σ_D^2 is a linear combination of the variance and the expectation of the delay in number of messages, according to*

$$\sigma_D^2 = \sigma_T^2 \mu_{\Delta} + \mu_T^2 \sigma_{\Delta}^2.$$

Assuming further that message arrivals follow a Poisson process with rate $\lambda = 1/\bar{t}$,

$$\sigma_D^2 = \bar{t}^2 (\mu_{\Delta} + \sigma_{\Delta}^2).$$

Proof: The proportionality of the expectations follows from a simple application of iterated expectation:

$$\mathbb{E} T = \mathbb{E} \mathbb{E}[T|\Delta] = \mathbb{E} \left[\sum_{i=1}^{\Delta} \mathbb{E} T_i \right] = \mathbb{E}[\Delta \mu_T] = \mu_{\Delta} \mu_T.$$

The statement on the variances is a consequence of the Pythagorean identity for nonlinear estimation, viewing

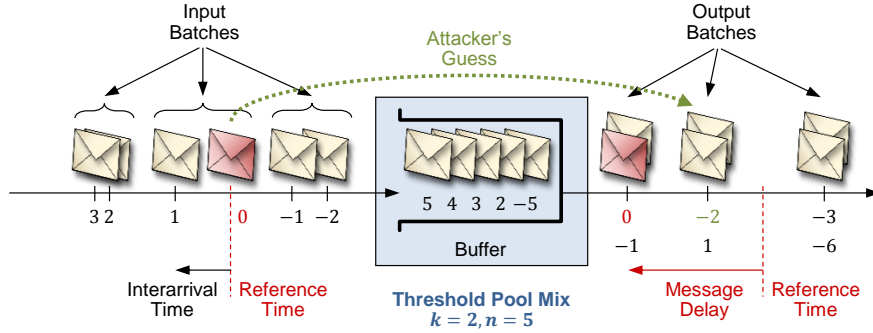
$$\mathbb{E}[D|\Delta] = \Delta \mu_T$$

as the estimate of the unknown D given the observation Δ , that is,

$$\text{Var} D = \text{Var} \mathbb{E}[D|\Delta] + \mathbb{E} \text{Var}[D|\Delta].$$

Observe that in the second term, because T_i are uncorrelated, the variance of the sum is the sum of variances, and therefore,

$$\text{Var}[D|\Delta] = \sum_{i=1}^{\Delta} \text{Var} T_i = \Delta \sigma_T^2.$$



(a) In this example, we show a k -input, k -output threshold pool mix, with $k = 2$ and with a buffer of $n = 5$ messages. This means that, when the number of messages stored by the mix reaches $n = 5$, the mix chooses $k = 2$ messages at random and forwards them simultaneously in a batch.

Incoming Message	Mix Buffer	Flushed Messages	Delay Δ [# of Messages]	Batch B	Index I
-2	<div style="border: 1px solid black; padding: 2px; display: inline-block;">-2 -3 -5 -6</div>		3	1	1
-1	<div style="border: 1px solid black; padding: 2px; display: inline-block;">-1 -2 0 -5 0</div> → <div style="border: 1px solid black; padding: 2px; display: inline-block;">-6 -3</div>	<div style="border: 1px solid black; padding: 2px; display: inline-block;">-6 -3</div>	4	2	0
0	<div style="border: 1px solid black; padding: 2px; display: inline-block;">0 -1 -2 -5</div>		3	1	1
1	<div style="border: 1px solid black; padding: 2px; display: inline-block;">1 0 -1 1 -5</div> → <div style="border: 1px solid black; padding: 2px; display: inline-block;">1 -2</div>	<div style="border: 1px solid black; padding: 2px; display: inline-block;">1 -2</div>	0	0	0
2	<div style="border: 1px solid black; padding: 2px; display: inline-block;">2 0 -1 -5</div>		≥ 3	≥ 1	0 - 1
3	<div style="border: 1px solid black; padding: 2px; display: inline-block;">3 2 0 1 -5</div> → <div style="border: 1px solid black; padding: 2px; display: inline-block;">-1 0</div>	<div style="border: 1px solid black; padding: 2px; display: inline-block;">-1 0</div>	≥ 2	≥ 1	0 - 1

(b) For each incoming message, we represent the resulting buffer state and the messages flushed, along with the corresponding delay Δ (in number of messages), batch B and index I .

Figure 3: Example of operation of a threshold pool mix. Messages are indexed in order of arrival, from -2 to 3 . Message 0 , which we choose as reference, is not flushed out instantaneously. On the one hand, it arrives before the buffer threshold is reached. On the other, it must wait for the second upcoming flush. In the end, message 0 is delayed until message 3 arrives, and accordingly, $\Delta = 3$, $B = 1$ and $I = 1$.

In the special case of Poisson arrivals, T_i are exponentially distributed with parameter $\lambda = 1/\bar{t}$, thus $\mu_T = \sigma_T = \bar{t}$. ■

Before proceeding, we characterize Δ by breaking it down into two components. It suffices to note that the delay Δ in number of messages an incoming message experiences until it is sent out, is completely determined by the *batch* B of k messages it belongs to, and the *index* I within that batch, which we both number starting from 0. Formally, B and I are r.v.'s defined as

$$\left. \begin{aligned} B &= \lfloor \Delta/k \rfloor \\ I &= \Delta \bmod k \end{aligned} \right\},$$

so that

$$\Delta = kB + I. \quad (3)$$

Clearly, there is a one-to-one correspondence between Δ and the pair (B, I) . We noticed in the example of Fig. 3 that message 0 was delayed by $\Delta = 3$ messages. In fact, it is not flushed in the immediately upcoming batch, when message 1 arrives, batch which we would index with 0, but in the following batch, when message 3 arrives, numbered with $B = 1$. Furthermore, because message 0 arrives precisely one message before the buffer is flush, indicated by $I = 1$. In other words, the reference message is delayed one output batch and one message within this batch.

It is important to note at this point that for a given *output* message, Δ , in addition to measuring delay, unequivocally indexes the corresponding input. Consequently, the statistical properties of this r.v. accurately portray the uncertainty regarding the input message corresponding to a given output. Take for instance message 0 at the output of the mix in Fig. 3, flushed when message 3 arrives. $\Delta = 3$ indicates that the matching input occurred three messages ago, thus its knowledge would dissipate its otherwise uncertain origin.

When an *input* message is given instead, Δ does not completely remove the uncertainty among all possible corresponding outputs; it merely narrows it down to an output batch. As a matter of fact, the value of the batch B suffices to pinpoint the output batch. However, viewing I as the index of the matching output within the batch, the pair (B, I) fully determines the operation of the mix on an input message, and its corresponding output. In our running example, message 0 at the input of the mix is not forwarded the first time the mix flushes, but in the second batch, numbered by $B = 1 > 0$. Index $I = 1 > 0$ indicates that the second among the two output messages in this batch matches the given input, message 0. But we mentioned that there is a one-to-one correspondence between Δ and the pair (B, I) . Consequently, Δ characterizes the statistical uncertainty of message correspondence, regardless of whether an input or an output message is given.

To sum up, the delay in number of messages Δ not only indexes message correspondence, thereby characterizing the anonymity in the communication, but also suffices to characterize the delay in time units, at least for

the purpose of averages, up to a proportionality constant. Based on the latter observation, we shall hereafter analyze delays only in terms of number of messages, in lieu of time units.

Having adopted message delay as our candidate to measure the cost incurred by the use of such a mix, we now turn to the matter of proposing adequate measures of the privacy gained in doing so. In keeping with the *anonymity metrics* adopted in the literature, explored in Sec. 2.2, in this paper we shall consider two information-theoretic quantities, namely Shannon's entropy $H_1(\Delta)$ and min-entropy $H_\infty(\Delta)$, of the (identically distributed) message correspondence index Δ . Both quantities are suitable, widely common measures of the uncertainty in the index of message correspondence Δ , in other words, measures of the dispersion of the probability distribution of candidate output messages for a given input. A succinct review and interpretation of both types of entropies was provided in Secs. 3.2 and 3.3, along with an argument against Hartley's entropy for our particular purposes.

As a final consideration, it must be pointed out that the Shannon entropy and the min-entropy of the r.v. Δ , effectively a measure of uncertainty on the correspondence between input and output messages, are by no means the only ways to quantitatively characterize the anonymity introduced by the pool. Concordantly with the principle established in [31], privacy may be construed as the error in the estimation of certain confidential information of interest to the privacy adversary, from any disclosed observation. In the context of our work, the sensitive information to be obfuscated consists in the communication patterns inferable via traffic analysis, beyond content protected via encryption, in a communications system, chiefly the identity of the senders and receivers of messages in a computer network. The observation consists in the timing of the messages entering and exiting the pool mix. Naturally, the algorithmic operation of the mix is considered known to the attacker. Within this general principle, the choice of a specific privacy metric remains entirely contingent upon the application and the adversarial model at hand. This is of course in addition to mathematical tractability, in the event that said metric is meant to be part of a theoretical model aiming to assess or improve the design of the anonymous-communication mechanism.

Our characterization of message unlinkability by means of the uncertainty between input and output messages is directly inspired by the extensive literature on the subject, particularly the work by Serjantov [37, 39]. Alternative privacy metrics resort to combinatorial approaches that count correspondences between senders and receivers consistent with observed message timings, ruling out some of the permutations [19].

It is relatively straightforward to conceptually formulate alternative anonymity models, but mathematical and computational implications may just as easily become overly complex. As an example suggestive of a measure of anonymity different from the one preferred in this work, suppose that

the anonymity of the communications between 3 senders and 2 receivers is to be protected by means of a threshold pool mix. Suppose further that the privacy adversary is interested in unveiling communication patterns representable by the 3-by-2 matrix of conditional probabilities modeling the transmission frequency of any given sender-receiver pair. From the observation of the message timing and knowing the operation parameters of the pool mix, privacy could be measured as the error in the estimation of said matrix, for instance as the Frobenius norm between the estimate and the correct matrix. The mathematical tractability of this model, however, is no trivial matter.

An enthralling approach to assessing the anonymity of a threshold pool mix, completely different in its fundamental model from the one followed here, appears in [51]. While this is a perfectly sound, mathematically tractable model, rather than measuring the uncertainty in the correspondence between input and output messages, or even sender and receiver identity, the privacy model in the cited work focuses on sender and receiver linkability. In their example on simple Chaumian mixes (Sec. 4), 3 senders and 3 receivers communicate via a single pool mix with threshold 2. Attacker observations are defined as pairs consisting of a subset of senders and a subset of receivers consistent with the anonymized communication. Say sender 2 sends a message to receiver 1, and 3 to 2. The only possible observation is that indicating that senders 2 and 3 and receivers 1 and 2 are involved; all other observations given this anonymous event are assigned zero probability. Although there is only one possible observation per anonymous communication and thus the conditional probabilities of observations given anonymous events are trivial, there are several equally likely anonymous events explaining each observation, among all possible anonymous events a priori. The conditional probabilities of all the anonymous events contemplated given all candidate observations are then used in the computation of a number of metrics, including the Shannon conditional entropy, additively normalized in the form of mutual information, and the conditional min-entropy described in [43].

4. Theoretical Characterization of the Anonymity-Delay Trade-Off

In the previous section, we specified the design parameters of threshold pool mixes and established measures of delay incurred and of anonymity gained. In this section, equipped with quantifiable design variables and objectives, we may proceed to tackle the problem of selection of mix parameters to attain an optimal anonymity-delay trade-off, in the sense that anonymity is maximized for a given expected delay, or expected delay is minimized for a given anonymity. For that purpose, we first express expected delay and anonymity as functions of the mix parameters, then show which choice of parameters is optimal, and finally characterize the optimal performance attained. We

make use of well-known properties of geometric r.v.'s, reviewed in the background subsection, Sec. 3.2. All theoretical results are expressed to be valid for any logarithmic base.

4.1. Delay in Messages and Time Units

The parameters k , m and n and the details of the operation of threshold pool mixes were specified in the previous section, Sec. 3.4. In the same section, the delay in number of messages Δ was defined and shown to be indicative of delay in time units, as far as averages are concerned, up to a proportionality constant. But Δ was also shown to be indicative of anonymity, as it indexes the correspondence between input and output messages. Finally, we broke down Δ into two characteristic components, the batch B a message is flushed in and the index I within the batch. This notation was illustrated in the example shown in Fig. 3, for a mix with $k = 2$, $n = 5$ and $m = 3$.

Our first proposition will complete the characterization of Δ by analyzing the statistical properties of B and I . The proposition in question will contemplate the extreme case of a Chaumian mix, in which the buffer is completely flushed in every round, represented by $m = 0$, or equivalently, $k = n$. Further, we shall also see that the proposition is consistent with the complete absence of a mix, modeled by $k = 1 = n$ and yielding $\Delta = 0$ with probability 1.

Proposition 1 (Batch, Index and Delay in Messages).

- (i) B has a geometric distribution, supported on the set $\{0, 1, \dots\}$, with parameter k/n .
- (ii) I is uniformly distributed on $\{0, 1, \dots, k-1\}$.
- (iii) B and I are statistically independent.
- (iv) Furthermore, the PMF of Δ is

$$p_{\Delta}(\delta) = \frac{1}{n} \left(1 - \frac{k}{n}\right)^{\lfloor \frac{\delta}{k} \rfloor} = \frac{1}{n} \left(\frac{m}{n}\right)^{\lfloor \frac{\delta}{k} \rfloor}.$$

- (v) Finally, in the trivial case when $m = 0$, we have $B = 0$ with probability 1, thus $\Delta = I$, uniformly distributed on $\{0, 1, \dots, n-1\}$.

Proof: The probability of a specific message being chosen from a pool of n , after k independent trials without replacement, is k/n . If this particular message is not chosen in a batch of size k , it may be chosen in the next independently, and with the same likelihood. The position of this message within the batch, by symmetry, is equally likely, regardless of the batch number. The expression for p_{Δ} is a straightforward consequence of the relationship between Δ and (B, I) , and involves routine manipulation of the formula for the PMF of the geometric r.v. B . ■

Fig. 4 plots the PMF $p_{\Delta}(\delta)$ characterized in Prop. 1(iv). The depiction illustrates the general case when $k < n$, for which the distribution decays geometrically in steps of size k , and has unbounded support. In the special case of the Chaumian mix, when $k = n$, the expression in the proposition reduces to a uniform distribution between 0

and $n - 1$. As intuition would suggest, $p_\Delta(\delta)$ reaches its maximum value $1/n$ at $\delta = 0$, regardless of k .




Figure 4: PMF $p_\Delta(\delta) = \text{P}\{\Delta = \delta\}$ of the delay Δ in number of messages, according to Prop. 1(iv).

The next proposition characterizes the PDF $p_D(d)$ of the delay in time units D , in terms of the PMF $p_\Delta(\delta)$ of the delay in number of messages Δ .

Proposition 2 (Delay in Time Units). *Assume that message arrivals follow a Poisson process with rate $\lambda = 1/\bar{t}$.*

- (i) *The PDF $p_D(d)$ of the delay in time units D is a mixture of a Dirac delta $\delta_{\text{Dirac}}(d)$ at $d = 0$, with weight equal to $\text{P}\{\Delta = 0\} = 1/n$, and of Erlang r.v.'s weighted according to the PMF $p_\Delta(\delta)$ of the delay in number of messages Δ , given by*

$$p_D(d) = \frac{1}{n} \delta_{\text{Dirac}}(d) + \lambda e^{-\lambda d} \sum_{\delta=0}^{\infty} p_\Delta(\delta + 1) \frac{(\lambda d)^\delta}{\delta!}.$$

- (ii) *In the special case when $k = 1$, $p_D(d)$ becomes a simple mixture of a Dirac delta at $d = 0$ and an exponential r.v. with parameter λ/n ; more specifically,*

$$p_D(d) = \frac{1}{n} \delta_{\text{Dirac}}(d) + \left(1 - \frac{1}{n}\right) \frac{\lambda}{n} e^{-\frac{\lambda}{n} d}.$$

Proof: We showed in Prop. 1 that $p_\Delta(0) = 1/n$. Let $p_{D|\Delta}(d|\delta)$ denote the conditional PDF of D given Δ . From the definition of D as the sum of Δ interarrival terms, $p_{D|\Delta}(d|0) = \delta_{\text{Dirac}}(d)$, and

$$p_D(d) = \text{E}_\Delta p_{D|\Delta}(d|\Delta) = \frac{1}{n} \delta_{\text{Dirac}}(d) + \sum_{\delta=1}^{\infty} p_\Delta(\delta) p_{D|\Delta}(d|\delta).$$

Under the Poisson assumption, for any $\delta > 0$, D given δ is the sum of δ independent, exponentially distributed r.v.'s with common parameter λ , i.e., $D|\delta$ has an Erlang distribution, precisely,

$$p_{D|\Delta}(d|\delta) = \frac{\lambda^\delta d^{\delta-1}}{(\delta-1)!} e^{-\lambda d}.$$

The first statement of the proposition follows.

The second assertion involves particularizing Prop. 1 for $k = 1$, which implies that $\Delta = B$ is geometrically

distributed with parameter $1/n$, starting at 0, which in turn particularizes $p_{D|\Delta}(d|\delta)$ above for

$$p_\Delta(\delta) = \frac{1}{n} \left(1 - \frac{1}{n}\right)^\delta,$$

resulting in

$$p_{D|\Delta}(d|\delta) = \left(1 - \frac{1}{n}\right) \frac{\lambda}{n} e^{-\lambda d} \sum_{\delta=0}^{\infty} \frac{\left(\left(1 - \frac{1}{n}\right) \lambda d\right)^\delta}{\delta!}.$$

To complete the proof, recall that $e^x = \sum_{n=0}^{\infty} \frac{x^n}{n!}$. ■

4.2. Expected Delay and Entropies

We are now ready to express the expected delay in number of messages (2), and the Shannon and min- entropies, in terms of the mix parameters.

Proposition 3 (Delay Expectation and Variance).

$$\bar{\delta} = \mu_\Delta = \text{E} \Delta = m + \frac{k-1}{2} = n - \frac{k+1}{2} \quad \text{and}$$

$$\sigma_\Delta^2 = \text{Var} \Delta = m n + \frac{k^2-1}{12}.$$

Proof: Prop. 1 asserts that B and I are geometric and uniform r.v.'s, respectively. Direct application of the well-known expressions for the mean and variance of said r.v.'s leads to $\text{E} B = \frac{n}{k} - 1$, $\text{E} I = \frac{k-1}{2}$, $\text{Var} B = \frac{m n}{k^2}$ and $\text{Var} I = \frac{k^2-1}{12}$. The linearity of the expectation operator applied to (3) implies that $\text{E} \Delta = k \text{E} B + \text{E} I$. On the other hand, because B and I are independent, they are also uncorrelated; consequently, $\text{Var} \Delta = k^2 \text{Var} B + \text{Var} I$. ■

Conforming with intuition, $\bar{\delta} = 0$ if, and only if, there is no mix, or equivalently, $k = 1$ and $m = 0$. We mentioned earlier that the Hartley entropy $H_0(\Delta)$ of Δ was not a suitable measure of anonymity. It is now clear, particularly from the expression for the PMF $p_\Delta(\delta)$ in Prop. 1(iv) and its depiction in Fig. 4, that as long as $k < n$, regardless of the specific values of k and n , $H_0(\Delta)$ remains uninformatively infinite. The following proposition and the rest of this paper concordantly focuses mainly on Shannon's entropy $H_1(\Delta)$ and min-entropy $H_\infty(\Delta)$. While the former uncertainty metric somewhat captures the shape of the PMF in its entirety, the latter metric merely exploits the fact that a MAP anonymity adversary would guess $\Delta = 0$ (or $\Delta = 1, \dots, \Delta = k-1$) as a maximally likely outcome, as represented in Fig. 3. We would like to stress that in the case $m = 0$, equivalent to require that $k = n$, Δ becomes uniformly distributed and consequently all entropies attain the (unconstrained) maximum $H_\alpha(\Delta) = \log n$. In the first statement of the following proposition, $m = 0$ is correctly represented under the usual convention $0 \log 0 = 0$, mentioned in Sec. 3.2.

Before proceeding, we must hasten to acknowledge that one of the results in Prop. 4, namely the expression for the Shannon entropy (i), already appeared in [38, 37], although its complete derivation is available in [10]. The

proof developed here, however, is substantially more compact. Further, these references do not consider min-entropy, the anonymity-delay trade-off, nor the optimality of the mix parameters to attain it, central aspects in this paper.

Proposition 4 (Entropy).

- (i) $H_1(\Delta) = \frac{H_1(k/n)}{k/n} + \log k = (n \log n - m \log m)/k$,
- (ii) $H_\infty(\Delta) = \log n$, and
- (iii) $H_2(\Delta) = \log(m+n)$.
- (iv) *More generally, for any $0 < \alpha \neq 1$,*

$$H_\alpha(\Delta) = \frac{1}{1-\alpha} \log \frac{k}{n^\alpha - m^\alpha}.$$

Proof: Assertion (i) concerning the Shannon entropy is a consequence of Prop. 1, characterizing Δ in terms of the pair (B, I) . Indeed, the one-to-one correspondence observed in Sec. 4.1 and the statistical independence shown in the proposition in that subsection enable us to write $H_1(\Delta) = H_1(B, I) = H_1(B) + H_1(I)$. Since B is geometric,

$$\begin{aligned} H_1(B) &= \frac{H_1(k/n)}{k/n} = \\ &= -\log(k/n) - \frac{1-k/n}{k/n} \log(1-k/n) = -\log \frac{k}{n} - \frac{m}{k} \log \frac{m}{n}, \end{aligned}$$

and since I is uniform, $H_1(I) = \log k$. The simplified expression for $H_1(\Delta)$ in the current proposition follows from routine logarithmic manipulation.

Part (ii) of the proposition is an immediate consequence of the definition of min-entropy and the PMF of Δ , the latter expressed in Prop. 1(iv).

The proof of statement (iii) on the collision entropy follows immediately from the general statement on Rényi entropies, which we proceed to show, simply by computing the summation in its definition. Specifically, on account of Prop. 1(iv),

$$\begin{aligned} \sum_{\delta=1}^{\infty} p_\Delta(\delta)^\alpha &= k \sum_{b=0}^{\infty} \left(\frac{1}{n} \left(1 - \frac{k}{n}\right)^b \right)^\alpha = \\ &= \frac{k}{n^\alpha} \sum_{b=0}^{\infty} \left(1 - \frac{k}{n}\right)^{\alpha b} = \frac{k/n^\alpha}{1 - \left(1 - \frac{k}{n}\right)^\alpha} = \frac{k}{n^\alpha - (n-k)^\alpha}. \end{aligned}$$

As one may expect, statement (i) in the theorem may alternatively be proven by taking the limit of (iv) as α approaches 1, concretely using l'Hôpital's rule. Conversely, statement (iv) may also be shown in a manner entirely analogous to that for (i). ■

4.3. Optimal Trade-Off

Equipped with the previous propositions, we are finally ready to solve the multiobjective optimization problem that characterizes the *optimal trade-off between expected delay and anonymity* in the design of threshold pool mixes. As stated in the beginning of this section, by optimal mix parameters we mean those parameters maximizing anonymity for a given expected delay, or equivalently, minimizing expected delay for a given anonymity.

It should be noted that Prop. 3 implies that only two disjoint types of expected delay are possible, namely integer and half-integer delays, and concordantly those cases will be distinguished. Recall that the set of *half-integers* is the set $\mathbb{Z} + 1/2$. More precisely, said proposition indicates that k being odd is equivalent to $\bar{\delta}$ being an integer, and k being even, to $\bar{\delta}$ being a half-integer.

The following theorem, Theor. 5, finds the best anonymity under a constraint on the expected delay, for the two possible cases of the latter quantity, integer and half-integer. We must stress that best here means highest entropy for a given expected delay $\bar{\delta}$. Consider, for example, $\bar{\delta} = 2$. Prop. 3 implies that the only possible mixes are those determined by (k, n) equal to $(1, 3)$, $(3, 4)$ and $(5, 5)$, with odd k in all three cases. Our theorem will demonstrate that the first, the only mix with $k = 1$, is the best choice in Shannon entropy, whereas the last, determined by the condition $k = n$, is optimal under the alternative min-entropy criterion. Take now the constraint $\bar{\delta} = 3/2$, encompassing the mixes parametrized by (k, n) equal to $(2, 3)$ and $(4, 4)$, with even k in both cases. While Shannon entropy favors the only choice with $k = 2$, namely $(2, 3)$, min-entropy advocates instead for $(4, 4)$. For such half-integer-valued $\bar{\delta}$, no mix with $k = 1$ is possible, let alone optimal.

On account of (3) and Prop. 3, it is also important to remark that the case $k = 1$ corresponds to $\Delta = B$, which means that the delay is geometrically distributed, precisely the solution for the case of Shannon entropy with integer expected delays.

Theorem 5 (Optimal Mix Parameters).

- (i) *For each integer expected delay $\bar{\delta}$, the mix has maximum Shannon entropy $H_1(\Delta)$ among all mixes with that delay if, and only if, $k = 1$. In the half-integer case, maximum Shannon entropy is then equivalent to $k = 2$.*
- (ii) *Regardless of whether the expected delay is an integer or a half-integer, a mix has maximum min-entropy among all mixes with the same delay if, and only if, $k = n$.*
- (iii) *Any mix parameters achieve maximum collision entropy for some delay.*

Proof: The case of integer delays and Shannon entropy is an immediate consequence of the fact that $k = 1$ is equivalent to requiring that Δ be geometrically distributed, and the fact that geometric r.v.'s maximize the Shannon entropy subject to a constraint on their expectation [7, §12]. Consider now the half-integer case, which corresponds to the case of even k . Define $B' = \lfloor \Delta/2 \rfloor$ and $I' = \Delta \bmod 2$. It should immediately be noted that $B' = B$, and thus geometrically distributed, provided that $k = 2$. For any even k , $\Delta = 2B' + I'$, and $I' = I \bmod 2$. This implies that the uniformity and independence of the batch index I , demonstrated in Prop. 1, carries over to the binary r.v. I' . Therefore,

$$H_1(\Delta) = H_1(B', I') = H_1(B') + H_1(I') = H_1(B') + \log 2$$

and

$$E \Delta = 2 E B' + E I' = 2 E B' + 1/2.$$

This reduces the problem of maximizing $H_1(\Delta)$ for a fixed $E \Delta$ to the problem of maximizing $H_1(B')$ for a fixed $E B'$, maximum attained whenever B' is geometrically distributed, i.e., whenever $k = 2$.

As far as the optimality condition for the min-entropy case is concerned, we proceed by seeking the minimum expected delay $\bar{\delta}$ for a fixed $H_\infty(\Delta)$, invoking the results of Props. 3 and 4. Precisely, regarding these two quantities as functions of k and n , $H_\infty(\Delta) = \log n$ depends only on n , whereas $\bar{\delta}$ decreases with $k \leq n$.

It is left to prove the optimality condition in the collision entropy case. Props. 3 and 4(iii) imply that

$$H_2(\Delta) = \log(2\bar{\delta} + 1).$$

This means that once $\bar{\delta}$ is fixed, $H_2(\Delta)$ is uniquely determined, in other words, any consistent combination of mix parameters will produce a single collision entropy value, formally both the maximum and the minimum for a given expected delay. ■

A numerical example should clarify the assertions in Theor. 5, and illustrate the formulas in Props. 3 and 4. A mix with buffer size $n = 3$ and batch size $k = 1$ imposes a delay $\bar{\delta} = 2$ while offering a level of anonymity quantified by a Shannon entropy of $H_1(\Delta) \simeq 2.755$ bit. No other combination of mix parameters n and k with $\bar{\delta} = 2$ —or lower—can attain higher $H_1(\Delta)$. This includes the combinations $n = 4$ and $k = 3$, $n = 5$ and $k = 5$, and so on. Making $n = 3$ and $k = 2$ yields $\bar{\delta} = 1.5$ and $H_1(\Delta) \simeq 2.377$, the maximum Shannon entropy among any mixes with that particular delay. The first entropy, approximately 2.755, is indeed higher than the second, 2.377, but 2 is also a higher delay than 1.5. As a matter of fact, both entropy-delay pairs define points lying on the optimal trade-off curve. None is better than the other in the sense of multiobjective optimization, that is, when anonymity is maximized for a given delay, or delay minimized for a given anonymity.

If anonymity is to be measured by min-entropy, then $k = n$ produces the best anonymity-delay trade-off. For $n = 5 = k$, $\bar{\delta} = 2$ and $H_\infty(\Delta) \simeq 2.322$, meaning that no other mix satisfying that delay constraint can offer better anonymity, when measured as min-entropy. For $n = 6 = k$, $\bar{\delta} = 2.5$ and $H_\infty(\Delta) \simeq 2.585$; anonymity is improved at the cost of delay. A more detailed numerical illustration is presented later in Sec. 5.

Our next and last theorem provides explicit closed-form expressions for the optimal anonymity-delay trade-off. Denote by $H_1(\bar{\delta})$ the values of the function $H_1 : \bar{\delta} \mapsto H_1(\Delta)$, and analogously for $H_\infty(\bar{\delta})$. Once again, recall that $0 \log 0 = 0$. The last note on decreasing increments, which may be regarded as a form of discrete concavity, means that, in practice, increasing delay yields diminishing returns in anonymity.

Theorem 6 (Optimal Anonymity-Delay Trade-Off).

- (i) *In the Shannon entropy case, the optimal trade-off between anonymity $H_1(\Delta)$ and expected delay $\bar{\delta}$ is attained at the points in the plane related by*

$$H_1(\bar{\delta}) = (\bar{\delta} + 1) H_1\left(\frac{1}{\bar{\delta} + 1}\right) = (\bar{\delta} + 1) \log(\bar{\delta} + 1) - \bar{\delta} \log \bar{\delta}$$

for $\bar{\delta} = 0, 1, \dots$, and

$$H_1(\bar{\delta}) = \frac{\bar{\delta} + 3/2}{2} H_1\left(\frac{2}{\bar{\delta} + 3/2}\right) + \log 2 = \frac{1}{2} ((\bar{\delta} + 3/2) \log(\bar{\delta} + 3/2) - (\bar{\delta} - 1/2) \log(\bar{\delta} - 1/2))$$

for $\bar{\delta} = 1/2, 3/2, \dots$. In addition, in the latter case,

$$H_1(\bar{\delta}) = \frac{1}{2} (H_1(\bar{\delta} - 1/2) + H_1(\bar{\delta} + 1/2)).$$

- (ii) *In the min-entropy case, the optimality curve is given by $H_\infty(\bar{\delta}) = \log(2\bar{\delta} + 1)$.*
 (iii) *The above expressions may be approximated for $\bar{\delta} \gg 1$, by $H_1(\bar{\delta}) \simeq \log \bar{\delta} + \log e$ (regardless of whether $\bar{\delta}$ is an integer or a half-integer), and*
 (iv) *$H_\infty(\bar{\delta}) \simeq \log \bar{\delta} + \log 2$.*
 (v) *Both entropies, as a function of the expected delay, are unbounded and strictly increasing. However, the integer increment functions $H_1(\bar{\delta} + 1) - H_1(\bar{\delta})$ and $H_\infty(\bar{\delta} + 1) - H_\infty(\bar{\delta})$ are strictly decreasing.*
 (vi) *Similarly to the min-entropy case, in the collision entropy case, $H_2(\bar{\delta}) = \log(2\bar{\delta} + 1)$, and*
 (vii) *$H_2(\bar{\delta}) \simeq \log \bar{\delta} + \log 2$.*

Proof: The exact formulas (i) and (ii) are routine manipulation of the results in Props. 3 and 4, restricted to the optimal values of k in Theor. 5.

The only nontrivial approximation is (iii), for the Shannon entropy case, which we show for the integer delay subcase, as the half-integer one is entirely analogous. Write

$$(\bar{\delta} + 1) \log(\bar{\delta} + 1) - \bar{\delta} \log \bar{\delta} = \log(\bar{\delta} + 1) + \bar{\delta} \log \frac{\bar{\delta} + 1}{\bar{\delta}}.$$

In order to complete the proof of this statement, it suffices to show that the last term approaches $\log e$ as $\bar{\delta} \rightarrow \infty$, or more simply in terms of natural logarithms, that it approaches 1. To see this, define $t = 1/\bar{\delta}$ and compute the limit as $t \rightarrow 0$, using l'Hôpital's rule:

$$\lim_{\bar{\delta} \rightarrow \infty} \bar{\delta} \ln \frac{\bar{\delta} + 1}{\bar{\delta}} = \lim_{t \rightarrow 0} \frac{\ln(1+t)}{t} = \lim_{t \rightarrow 0} \frac{1/(1+t)}{1} = 1.$$

The statements in (v) for $H_\infty(\bar{\delta})$ are similar albeit simpler to prove than those for $H_1(\bar{\delta})$, and hence omitted. The claim that $H_1(\bar{\delta})$ is unbounded may be seen from the logarithmic approximation (iii) to (i). To show that it is strictly increasing, we differentiate the first expression for integer $\bar{\delta}$, in terms of natural logarithms without loss of generality, and verify that it is strictly positive:

$$\frac{dH_1}{d\bar{\delta}} = \ln(\bar{\delta} + 1) - \ln \bar{\delta} > 0.$$

In light of the interpolation formula in (i), this also proves the half-integer case.

The statement regarding decreasing increments may be shown in a number of ways. A direct method for the integer case consists computing the derivative

$$\begin{aligned} \frac{d}{d\bar{\delta}} (H_1(\bar{\delta} + 1) - H_1(\bar{\delta})) &= \frac{dH_1}{d\bar{\delta}}(\bar{\delta} + 1) - \frac{dH_1}{d\bar{\delta}}(\bar{\delta}) = \\ &= \ln(\bar{\delta} + 2) + \ln(\bar{\delta}) - 2\ln(\bar{\delta} + 1), \end{aligned}$$

easily verified to be strictly negative from the concavity of the logarithm and Jensen's inequality. For the half-integer case, the interpolation (i) implies that

$$\begin{aligned} H_1(\bar{\delta} + 1) - H_1(\bar{\delta}) &= \frac{1}{2}(H_1(\bar{\delta} + 3/2) - H_1(\bar{\delta} + 1/2)) + \\ &+ H_1(\bar{\delta} + 1/2) - H_1(\bar{\delta} - 1/2), \end{aligned}$$

the sum of two integer increments at the integer values $\bar{\delta} + 1/2$ and $\bar{\delta} - 1/2$, just shown to be decreasing.

Alternative methods involve verifying that the second derivative of $H_1(\bar{\delta})$ is negative, which means that the first derivative is decreasing and that the function itself is concave, then applying the mean-value theorem, the fundamental theorem of calculus, or Jensen's.

Finally, the exact formula in statement (vi) concerning collision entropy was in fact shown and used at the end of the proof of Theor. 5, and being equal to that for min-entropy, so is approximation (vii). ■

5. Experimental Results

This section numerically illustrates the main theoretical results of Sec. 4, obtained for threshold pool mixes operating according to the model and the parameters specified in Sec. 3.4. Recall from the latter section that the r.v. Δ represents the delay from the moment a message enters the mix until it is forwarded, measured in terms of number of messages. The expected delay $\bar{\delta}$ in number of messages (2) is a measure proportional to the expected delay in time units, where the proportionality constant is the expected interarrival time (the inverse of the rate of message arrivals per time unit). We also argued that the delay Δ also indexes the correspondence between input and output messages, and that anonymity is measured as either the Shannon entropy $H_1(\Delta)$ or the min-entropy $H_\infty(\Delta)$, information-theoretic measures briefly reviewed in Sec. 3.2. All results are plotted in terms of those quantities; all entropies are given in bits.

First, we experimentally verify the formula of the PMF $p_\Delta(\delta)$ of the delay in number of messages Δ in Prop. 1(iv). To that end, we simulate a threshold mix with batch size $k = 2$ and buffer size $n = 5$, keeping track of the delay Δ of a total of 10^5 messages. The theoretical and experimental distributions are depicted in Fig. 5. According to Prop. 3, the mean and variance of the delay are $\bar{\delta} = 3.5$ and $\text{Var} \Delta = 15.25$, which our simulation estimated at 3.4999 and 15.272, respectively.

In the same simulation with $k = 2$ and $n = 5$, Poisson message arrivals are modeled simply by generating independent, exponentially distributed interarrival times, with

a common mean $\bar{t} = 2$ (arrival rate $\lambda = 1/\bar{t} = 1/2$). In order to verify the formula of the PDF $p_D(d)$ of the delay in time units D in Prop. 2(i), we estimate the weight of the Dirac delta component by counting the relative occurrence of the event $\{D = 0\}$, and we approximate the continuous portion with a simple Gaussian-kernel PDF estimate on the samples with $D > 0$. The resulting PDF estimate is plotted in Fig. 6, along with the theoretical counterpart. Additionally, we confirm the theoretical values for the expectation and variance of the delay in time units, $\mu_D = 7$ and $\sigma_D^2 = 75$, in Prop. 0, which the simulation accurately estimates as 6.9907 and 74.970, respectively.

Secondly, we consider all mixes parameterized by k and n with an expected delay $\bar{\delta}$ between 0 and 100, and compute their corresponding Shannon and min-entropies, according to the dependence established in Props. 3 and 4, which we plot in Fig. 7. Both integer and half-integer delays were included in the computation. The figure highlights, and numerically confirms, the optimal trade-off characterized in Theors. 5 and 6, that is, the frontier of maximum anonymity for a given expected delay, or minimum expected delay for a given anonymity, corresponding to $k = 1, 2$ in the Shannon entropy case, and $k = n$ in the min-entropy case. The figure equally highlights the effect of measuring anonymity according to one criterion or type of entropy, but optimizing for the other. Remarkably enough, our numerical computations suggest that optimizing for the wrong criterion yields the worst entropy for a given delay, not merely a suboptimal entropy.

Next, the two trade-offs are compared in a single plot provided as Fig. 8, which is numerically consistent with the exact formulae and the approximations found in Theor. 6. Incidentally, we observed the approximations to be fairly




Figure 8: Anonymity-delay trade-off for each entropy criterion.

accurate, and hardly distinguishable from the exact values in a plot, even for small $\bar{\delta}$.

Lastly, Fig. 9 plots the relative loss $(\bar{\delta} - \bar{\delta}_{\min})/\bar{\delta}_{\min}$ with respect to the optimal expected delay $\bar{\delta}_{\min}$, for a given entropy, for each of the mixes considered in Fig. 7. In the Shannon case, for each entropy value, the reference delay $\bar{\delta}_{\min}$ was interpolated from the first optimal trade-off for-




Figure 5: Experimental verification of the formula of the PMF $p_{\Delta}(\delta)$ of the delay Δ in number of messages in Prop. 1(iv). The simulation employed a threshold mix with batch size $k = 2$ and buffer size $n = 5$, and kept track of the delay Δ of a total of 10^5 messages.




Figure 6: Experimental verification of the formula of the PDF $p_D(d)$ of the delay D in time units in Prop. 2(i). The simulation employed a threshold mix with batch size $k = 2$ and buffer size $n = 5$, and kept track of the delay D of a total of 10^5 messages. Messages were generated according to a Poisson process, simply by drawing independent, exponentially distributed interarrival times, with mean $\bar{t} = 2$.

mula in Theor. 6. Observe that in the min-entropy case, a relative delay increment of up to 100% may be incurred by choosing the wrong mix parameters, and values of around 40% are common for the worst choice in the Shannon case. The upshot is that these two types of entropy constitute anonymity criteria leading to effectively different optimal parameters and mix performance.

6. Conclusion

This work addresses the problem of designing threshold pool mixes, a key building block in the area of anonymous communications, in a manner that contemplates the optimal trade-off between the two contrasting aspects of anonymity and delay.

We approach the issue of practical design of mixes in a systematic fashion, drawing upon the methodology of multiobjective optimization, long established in scientific

fields such as data compression or economics, but perhaps less so in the area of information privacy. More precisely, we depart from a standard adversary model and adopt three quantifiable measures of anonymity in the literature, Shannon's entropy, min-entropy, and collision entropy. A succinct review and interpretation of those types of entropies is provided in Sec. 3.2. For any message arrival process under the mild assumption of a common average in the sequence of interarrival times, we adopt a quantifiable measure of cost incurred, average delay. Having established a quantitative model of our design objectives enables us to formulate the issue of mix design as an optimization problem. The optimal solution to this problem is not a point, but a curve in the anonymity-delay plane representing the trade-off we seek.

Our extensive theoretical analysis expresses our optimization objectives, namely the expected delay $\bar{\delta}$ and entropies $H_1(\Delta)$, $H_\infty(\Delta)$, and $H_2(\Delta)$, in terms of the mix




Figure 7: Anonymity-delay region of all possible threshold pool mixes with $\bar{\delta} \in [0, 100]$. The optimal and worst frontiers are highlighted.

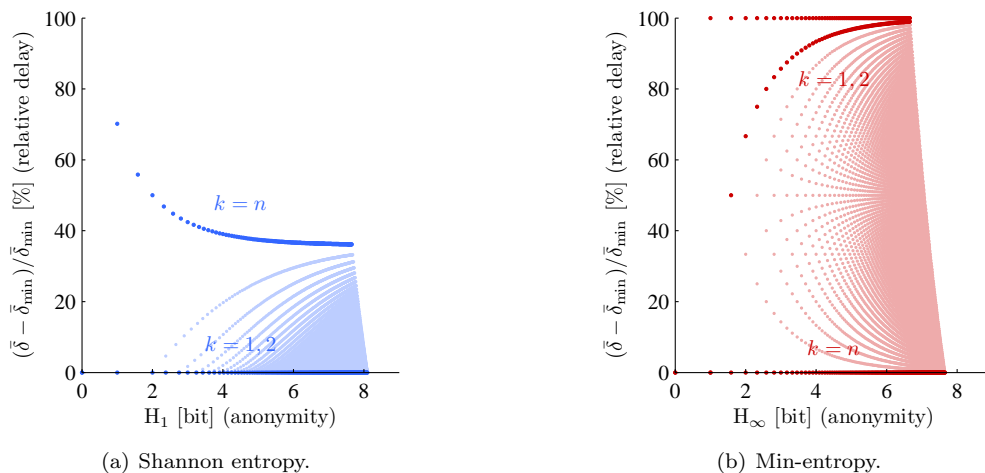


Figure 9: Relative loss in expected delay incurred by suboptimal choices of mix parameters.

parameters, batch size k and buffer size n . The theory continues by establishing the optimal parameters for each entropy criterion, along with exact and approximate expressions for the optimal anonymity-delay curve, the latter of the form $H(\Delta) \simeq \log \bar{\delta} + c$, with c a constant depending on the type of entropy. We also show that increments in expected delay yield positive albeit diminishing returns in anonymity.

A question that arises naturally from the conventional operation of threshold pool mixes, in the context of anonymity-delay optimality, is whether forwarded messages should be drawn uniformly from the pool's buffer. In other words, one may contemplate the possibility of a nonuniformly distributed choice, still random, but somehow biased towards messages that have remained inside the pool longer than others. The proof of the optimality results reached in Theor. 5 in the Shannon entropy case is based on the fact that geometric distributions, generated by independent Bernoulli trials and obtained for $k = 1$, maximize entropy subject to a constraint on expectation. The proof in question leads to the remarkable conclusion that uni-

form choices are optimal. However, this statement hinges on the measurement of cost as an average delay, instead of, say, a second-order moment, and on the measurement of anonymity as a Shannon entropy.

As for the case of min-entropy, only the highest likelihood of the message index correspondence Δ matters, which is $1/n$, attained in particular when $\Delta = 0$ regardless of k . In this case, the optimality condition $k = n$, equivalent to $m = 0$, means that the batch of k messages is forwarded as soon as possible, in order to minimize the delay. In other words, simple Chaumian mixes, with $m = 0$, are optimal threshold pool mixes when anonymity is measured as a min-entropy and cost as an average delay.

Somewhat surprisingly, any combination of mix parameters attains optimality in collision entropy for some delay, which means that the anonymity-delay region in this special case collapses into a curve.

Our main theoretical results, particularly the anonymity-delay region of possible mixes and the exact and approximate characterizations of its optimal frontier, are confirmed numerically. Our experiments report substantial

delay reductions by optimizing the mix parameters. In particular, we observe that a poor parametrization may lead to unnecessary average delay increments of up to 100% in the min-entropy case, and of up to 40% in the Shannon case. The experiments also delve into the effects of measuring anonymity according to one criterion, but optimizing according to the other, and conclude that these two types of entropy constitute anonymity criteria leading to effectively different optimal parameters and mix performance.

In closing, we would like to stress that our contribution on the design of threshold pool mixes may be regarded not only as a step towards a practical design that takes into account the factors of privacy and usability, but also as a methodological illustration of the applicability of formalized multiobjective optimization techniques to the still emerging field of privacy in information systems.

Acknowledgment

We would like to thank the anonymous referees for their extremely valuable comments. This work was partly supported by the Spanish Government through projects Consolider Ingenio 2010 CSD2007-00004 ARES, TEC2010-20572-C02-02 Consequence, and by the Government of Catalonia under grant 2009 SGR 1362. Additional supporting projects include IWT SBO SPION, FWO G.0360.11N, FWO G.0686.11N, and GOA TENSE (GOA/11/007). D. Rebollo-Monedero is the recipient of a Juan de la Cierva postdoctoral fellowship, JCI-2009-05259, from the Spanish Ministry of Science and Innovation.

References

- [1] T. Berners-Lee, J. Hendler, and O. Lassila, "The semantic Web," *Scient. Amer.*, pp. 35–43, May 2001.
- [2] O. Berthold, A. Pfitzmann, and R. Standtke, "The disadvantages of free MIX routes and how to overcome them," in *Proc. Design. Priv. Enhanc. Technol.: Workshop Design Issues Anon., Unobser.*, ser. Lecture Notes Comput. Sci. (LNCS). Berkeley, CA: Springer-Verlag, Jul. 2000, pp. 30–45.
- [3] J. Brickell and V. Shmatikov, "The cost of privacy: Destruction of data-mining utility in anonymized data publishing," in *Proc. ACM SIGKDD Int. Conf. Knowl. Disc., Data Min. (KDD)*, Las Vegas, NV, Aug. 2008.
- [4] D. Chaum, "Untraceable electronic mail, return addresses, and digital pseudonyms," *Commun. ACM*, vol. 24, no. 2, pp. 84–88, 1981.
- [5] S. Clauß and S. Schiffner, "Structuring anonymity metrics," in *Proc. ACM Workshop on Digit. Identity Manage.* Fairfax, VA: ACM, Nov. 2006, pp. 55–62.
- [6] L. Cottrell, "Mixmaster and remailer attacks," 1994. [Online]. Available: <http://obscura.com/~loki/remailer/remailer-essay.html>
- [7] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed. New York: Wiley, 2006.
- [8] I. Csiszár and J. Körner, "Broadcast channels with confidential messages," *IEEE Trans. Inform. Theory*, vol. 24, pp. 339–348, May 1978.
- [9] G. Danezis, "Mix-networks with restricted routes," in *Proc. Workshop Priv. Enhanc. Technol. (PET)*. Lecture Notes Comput. Sci. (LNCS), 2003, pp. 1–17.
- [10] G. Danezis, "Better anonymous communications," Ph.D. dissertation, Univ. of Cambridge, 2004.
- [11] G. Danezis, C. Diaz, E. Käsper, and C. Troncoso, "The wisdom of crowds: Attacks and optimal constructions," in *ESORICS*. Springer LNCS 5789, 2009, pp. 406–423.
- [12] G. Danezis, C. Diaz, and P. Syverson, "Systems for anonymous communication," in *CRC Handbook of Financial Cryptography and Security*. Chapman & Hall, 2010, pp. 341–390.
- [13] 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*, 2003, pp. 2–15.
- [14] C. Diaz, S. Seys, J. Claessens, and B. Preneel, "Towards measuring anonymity," in *Proc. Workshop Priv. Enhanc. Technol. (PET)*, ser. Lecture Notes Comput. Sci. (LNCS), vol. 2482. Springer-Verlag, Apr. 2002, pp. 54–68.
- [15] C. Diaz and B. Preneel, "Taxonomy of mixes and dummy traffic," in *Information Security Management, Education and Privacy (SEC'04)*, vol. 3. Kluwer, 2004, pp. 215–230.
- [16] C. Diaz and A. Serjantov, "Generalising mixes," in *Designing Privacy Enhancing Technologies, Proceedings of PET'03*. Springer-Verlag, LNCS 2760, 2003, pp. 18–31.
- [17] R. Dingledine, N. Mathewson, and P. Syverson, "Tor: The second-generation onion router," in *Proc. Conf. USENIX Secur. Symp.*, Berkeley, CA, 2004, pp. 21–21.
- [18] C. Dwork, "Differential privacy," in *Proc. Int. Colloq. Automata, Lang., Program.* Springer-Verlag, 2006, pp. 1–12.
- [19] M. Edman, F. Sivrikaya, and B. Yener, "A combinatorial approach to measuring anonymity," *IEEE J. Intell., Secur. Inform.*, pp. 356–363, 2007.
- [20] J. Feigenbaum, A. Johnson, and P. Syverson, "A model of onion routing with provable anonymity," in *Proc. Financ. Cryptogr., Data Secur. (FI)*. Springer-Verlag, 2007.
- [21] B. Gierlich, C. Troncoso, C. Diaz, B. Preneel, and I. Verbauwhede, "Revisiting a combinatorial approach toward measuring anonymity," in *Proc. Workshop Priv. Electron. Society*. ACM, 2008, pp. 111–116.
- [22] D. Goldschlag, M. Reed, and P. Syverson, "Hiding routing information," in *Proc. Inform. Hiding Workshop (IH)*, 1996, pp. 137–150.
- [23] D. Kesdogan, J. Egner, and R. Büschkes, "Stop-and-go mixes: Providing probabilistic anonymity in an open system," in *Proc. Inform. Hiding Workshop (IH)*. Springer-Verlag, 1998, pp. 83–98.
- [24] N. Li, T. Li, and S. Venkatasubramanian, "t-Closeness: Privacy beyond k-anonymity and l-diversity," in *Proc. IEEE Int. Conf. Data Eng. (ICDE)*, Istanbul, Turkey, Apr. 2007, pp. 106–115.
- [25] A. Machanavajjhala, J. Gehrke, D. Kiefer, and M. Venkatasubramanian, "l-Diversity: Privacy beyond k-anonymity," in *Proc. IEEE Int. Conf. Data Eng. (ICDE)*, Atlanta, GA, Apr. 2006, p. 24.
- [26] S. Mauw, J. Verschuren, and E. P. de Vink, "A formalization of anonymity and onion routing," in *Proc. European Symp. Res. Comput. Secur. (ESORICS)*, vol. 3193. Lecture Notes Comput. Sci. (LNCS), 2004, pp. 109–124.
- [27] U. Möller, L. Cottrell, P. Palfrader, and L. Sassaman, "Mixmaster protocol – Version 2," Internet Eng. Task Force, Internet Draft, Jul. 2003. [Online]. Available: <http://www.freehaven.net/anonbib/cache/mixmaster-spec.txt>
- [28] A. Pfitzmann and M. Hansen, "A terminology for talking about privacy by data minimization: Anonymity, unlinkability, undetectability, unobservability, pseudonymity, and identity management," Aug. 2010, v0.34. [Online]. Available: http://dud.inf.tu-dresden.de/literatur/Anon.Terminology_v0.34.pdf
- [29] J. F. Raymond, "Traffic analysis: Protocols, attacks, design issues and open problems," in *Proc. Design. Priv. Enhanc. Technol.: Workshop Design Issues Anon., Unobser.* Springer-Verlag, 2001, pp. 10–29.
- [30] D. Rebollo-Monedero, J. Forné, and J. Domingo-Ferrer, "From t-closeness-like privacy to postrandomization via information theory," *IEEE Trans. Knowl. Data Eng.*, vol. 22,

- no. 11, pp. 1623–1636, Nov. 2010. [Online]. Available: <http://doi.ieeecomputersociety.org/10.1109/TKDE.2009.190>
- [31] D. Rebollo-Monedero, J. Parra-Arnau, C. Diaz, and J. Forné, “On the measurement of privacy as an attacker’s estimation error,” *Int. J. Inform. Secur.*, vol. 12, no. 2, pp. 129–149, Apr. 2013. [Online]. Available: <http://dx.doi.org/10.1007/s10207-012-0182-5>
- [32] M. G. Reed, P. F. Syverson, and D. M. Goldschlag, “Proxies for anonymous routing,” in *Proc. Comput. Secur. Appl. Conf. (CSAC)*, San Diego, CA, Dec. 1996, pp. 9–13.
- [33] M. K. Reiter and A. D. Rubin, “Crowds: Anonymity for Web transactions,” *ACM Trans. Inform. Syst. Secur.*, vol. 1, no. 1, pp. 66–92, 1998.
- [34] M. Rennhard and B. Plattner, “Practical anonymity for the masses with mix-networks,” in *Proc. Int. Workshop Enabling Technol.: Infra. Col. Enterprises (WETICE)*. IEEE Comput. Soc., 2003, pp. 255–260.
- [35] A. Rényi, “On measures of entropy and information,” in *Proc. Berkeley Symp. Math. Stat., Prob.*, Berkeley, CA, Jun. 1961, pp. 547–561.
- [36] P. Samarati, “Protecting respondents’ identities in microdata release,” *IEEE Trans. Knowl. Data Eng.*, vol. 13, no. 6, pp. 1010–1027, 2001.
- [37] A. Serjantov, “On the anonymity of anonymity systems,” Univ. of Cambridge, Comput. Lab., Tech. Rep. UCAM-CL-TR-604, 2004.
- [38] A. Serjantov and G. Danezis, “Towards an information theoretic metric for anonymity,” in *Proc. Workshop Priv. Enhanc. Technol. (PET)*, ser. Lecture Notes Comput. Sci. (LNCS), vol. 2482. San Francisco, CA: Springer-Verlag, Apr. 2002, pp. 41–53.
- [39] A. Serjantov, R. Dingledine, and P. Syverson, “From a trickle to a flood: Active attacks on several mix types,” in *Proc. Inform. Hiding Workshop (IH)*, ser. Lecture Notes Comput. Sci. (LNCS), vol. 2578. Noordwijkerhout, The Netherlands: Springer-Verlag, Oct. 2002, pp. 36–52.
- [40] A. Serjantov and R. E. Newman, “On the anonymity of timed pool mixes,” in *Proc. Workshop Priv., Anon. Issues Netw., Distrib. Syst.* Kluwer, 2003, pp. 427–434.
- [41] C. E. Shannon, “Communication theory of secrecy systems,” *Bell Syst., Tech. J.*, 1949.
- [42] V. Shmatikov and M. H. Wang, “Measuring relationship anonymity in mix networks,” in *Proc. Workshop Priv. Electron. Society*. ACM, 2006, pp. 59–62.
- [43] G. Smith, “On the foundations of quantitative information flow,” in *Proc. Int. Conf. Found. Softw. Sci., Comput. Struct. (FoSSaCS)*, Mar. 2009, pp. 288–302.
- [44] S. Steinbrecher and S. Kopsell, “Modelling unlinkability,” in *Proc. Workshop Priv. Enhanc. Technol. (PET)*. Springer-Verlag, 2003, pp. 32–47.
- [45] L. Sweeney, “ k -Anonymity: A model for protecting privacy,” *Int. J. Uncertain., Fuzz., Knowl.-Based Syst.*, vol. 10, no. 5, pp. 557–570, 2002.
- [46] P. Syverson and S. Stubblebine, “Group principals and the formalization of anonymity,” in *Proc. World Congr. Formal Methods*, 1999, pp. 814–833.
- [47] G. Tóth and Z. Hornák, “Measuring anonymity in a non-adaptive, real-time system,” in *Proc. Workshop Priv. Enhanc. Technol. (PET)*, ser. Lecture Notes Comput. Sci. (LNCS), vol. 3424. Toronto, Canada: Springer-Verlag, May 2004, pp. 226–241.
- [48] G. Tóth, Z. Hornák, and F. Vajda, “Measuring anonymity revisited,” in *Proc. Nordic Workshop Secure IT Syst.*, Nov. 2004, pp. 85–90.
- [49] T. M. Truta and B. Vinay, “Privacy protection: p -Sensitive k -anonymity property,” in *Proc. Int. Workshop Priv. Data Manage. (PDM)*, Atlanta, GA, 2006, p. 94.
- [50] A. Wyner, “The wiretap channel,” *Bell Syst., Tech. J.* 54, 1975.
- [51] S. Zhioua, “Anonymity attacks on mix systems: A formal analysis,” in *Proc. Inform. Hiding Workshop (IH)*, Prague, Czech Republic, May 2011, pp. 133–147.



David Rebollo-Monedero received the M.S. and Ph.D. degrees in electrical engineering from Stanford University, in California, USA, in 2003 and 2007, respectively. His doctoral research at Stanford focused on data compression, more specifically, quantization and transforms for distributed source coding. Previously,

he was an information technology consultant for PricewaterhouseCoopers, in Barcelona, Spain, from 1997 to 2000, and was involved in the Retevisión startup venture. During the summer of 2003, still as a Ph.D. student at Stanford, he worked for Apple Computer with the QuickTime video codec team in California, USA. He is currently a postdoctoral researcher with the Information Security Group of the Department of Telematics Engineering at the Universitat Politècnica de Catalunya (UPC), also in Barcelona, where he investigates the application of data compression formalisms to privacy in information systems.



Javier Parra-Arnau was awarded the M.S. degree in electrical engineering by the Universitat Politècnica de Catalunya (UPC) in 2004. After finishing his degree, he gained a position as a project engineer in the communications department of a Spanish engineering company. Four years later he

joined the Information Security Group within the Department of Telematics Engineering at the UPC and continued to further develop his training. He was awarded the M.S. degree in Telematics Engineering in 2009 and decided to engage in research. He is currently a Ph.D. candidate at UPC, where he investigates mathematical models dealing with the inherent trade-off between privacy and data utility in information systems.



Jordi Forné received the M.S. degree in telecommunications engineering from the Universitat Politècnica de Catalunya (UPC) in 1992, and the Ph.D. degree in 1997. Currently, he is an associate professor of the Telecommunications Engineering School of Barcelona (ETSETB), and works with the Information Security

Group, both affiliated to the Department of Telematics Engineering of UPC in Barcelona. He is the coordinator of the Ph.D. program in Telematics Engineering and the director of the research M.S. program in Telematics Engineering. His research interests span a number of subfields within information security and privacy, including network security, electronic commerce and public-key infrastruc-

tures. He has been a member of the program committee of a number of security conferences, and he is editor of several international journals.



Claudia Diaz graduated in telecommunications engineering at the University of Vigo, Spain, in 2000. In 2005 she received the Ph.D. degree in engineering from KU Leuven, Belgium, with a Ph.D. dissertation entitled “Anonymity and Privacy in Electronic Services”.

She was a postdoctoral researcher between 2005 and 2010, and since 2010 she is an Assistant Professor at KU Leuven. Her research is focused on the modeling, design and analysis of privacy-enhancing technologies for a range of applications, including communications networks, social networks, vehicular networks, location-based services, and Web search.