

Characterizing the Adversarial Power in Uniform and Ergodic Peer Sampling

Emmanuelle Anceaume, Yann Busnel, Sébastien Gambs

► **To cite this version:**

Emmanuelle Anceaume, Yann Busnel, Sébastien Gambs. Characterizing the Adversarial Power in Uniform and Ergodic Peer Sampling. [Research Report] PI-1966, 2011, pp.15. <inria-00564293>

HAL Id: inria-00564293

<https://hal.inria.fr/inria-00564293>

Submitted on 8 Feb 2011

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Characterizing the Adversarial Power in Uniform and Ergodic Peer Sampling

Emmanuelle Anceaume^{*} Yann Busnel^{**} Sébastien Gambs^{***}

Abstract: We consider the problem of achieving uniform and ergodic peer sampling in large scale open systems under adversarial behavior. The main challenge is to guarantee that any honest peer is able to construct a uniform and non-definitive (ergodic) sample of the peers identifiers in the system, and this even in the presence of Byzantine peers controlled by the adversary. This sample is built out of a stream of peer identifiers received at each node. We consider and study two types of adversary; an omniscient adversary that has the capacity to eavesdrop on all the messages that are exchanged within the system, and a blind adversary that can only observe messages that have been sent or received by peers he controls. In both models, the adversary can disrupt the input stream by injecting new messages or dropping messages sent by honest peers. Given any sampling strategy, we quantify the minimum effort an adversary has to exert on any input stream to prevent the sampling strategy from outputting a uniform and ergodic sample. We derive lower bounds for both adversary models.

Key-words: Uniform and ergodic sampling, unstructured P2P, Byzantine adversary.

Caractérisation de la puissance d'un adversaire: application à un service d'échantillonnage uniforme et ergodique

Résumé : Nous considérons le problème de la mise en œuvre d'un service d'échantillonnage de nœuds uniforme et ergodique dans des systèmes distribués à très grande échelle soumis à l'influence d'un adversaire malveillant. Nous étudions deux types d'adversaire : un adversaire omniscient qui est capable d'espionner tous les messages échangés à travers les liens de communication du système, et un adversaire aveugle, qui ne peut observer que les messages transitants par les nœuds sous son contrôle. Dans chacun de ces modèles, l'adversaire peut perturber le flux entrant de chaque nœud par l'injection ou l'effacement de messages. Nous quantifions l'effort minimum qu'un adversaire doit exercer pour empêcher la stratégie d'échantillonnage de générer un échantillon uniforme et ergodique. Enfin, nous dérivons des bornes inférieures pour chacun des modèles d'adversaire considérés.

Mots clés : Echantillonnage uniforme et ergodique, P2P non structurés, adversaire byzantin

^{*} CNRS UMR 6074 IRISA, emmanuelle.anceaume@irisa.fr

^{**} LINA / Université de Nantes, Yann.Busnel@univ-nantes.fr

^{***} IRISA / Université de Rennes 1 - INRIA Rennes Bretagne Atlantique,sebastien.gambs@irisa.fr

1 Introduction

We investigate the problem of achieving uniform and ergodic peer sampling in large scale open systems in presence of adversarial nodes. Uniform sampling is a fundamental primitive ensuring that any individual in a population has the same probability to be selected as sample. Uniform sampling finds its root in many practical problems such as data collection, dissemination, load balancing and data-caching [5, 11, 13, 16] but achieving uniform sampling in large scale open systems has been shown to be difficult. One of the reasons for this is that the population of these systems is typically very large (*e.g.*, thousand or millions of peers) and displays a very high churn (recent studies on the eDonkey file-sharing network have shown that in average 500,000 peers connect and disconnect per day [14]). Moreover, openness makes unavoidable the presence of adversaries controlling a high number of peers. These adversaries strategize to isolate honest peers within the system by violating randomness assumptions that are at the core of these systems¹.

By relying on the topological properties of structured peer-to-peer systems, it has been shown that it is possible to guarantee that with high probability any peer is equally likely to appear in the local view of each other honest peer in a number of communication rounds polynomial in the size of the system. One way to achieve this is by imposing peers to frequently depart from their position and move to another random position in the system [2, 4]. In unstructured peer-to-peer systems, peers cannot rely on the topological nature of structured graphs to detect undesirable behaviors. To circumvent this issue, Bortnikov *et al.* [6] have relied on the properties of min-wise independent permutations, which are fed by the streams of gossiped peer ids and eventually converge towards uniform sampling on the peer ids. However, the resulting sample is definitive in the sense that no other peer id received in the input stream can ever appear in the random sample, which makes this sampling strategy uniform but not ergodic. Informally, the ergodicity property guarantees that each peer id has a non-zero probability to appear infinitely often in a sample. A preliminary step in determining conditions under which uniform and ergodic sampling is achievable in unstructured peer-to-peer systems potentially populated with a large proportion of Byzantine peers has been presented in a previous paper [3]. Briefly, it has been shown that imposing strict restrictions on the number of messages sent by malicious peers during a given period of time and providing each honest peer with a very large memory (proportional to the size of the system) are necessary and sufficient conditions to obtain uniform and ergodic sampling.

In this paper, we propose a characterization of the adversarial power towards biasing uniform and ergodic sampling. By adopting a statistical view of the input stream and by comparing distributions using metrics such as information divergence, we derive lower bounds on the work that the adversary has to exert to bias this input stream such that uniform and ergodic sampling does not hold. We consider and study two models of adversary; the omniscient adversary, that has the capacity to eavesdrop on all the messages that are exchanged within the system, and the blind one that can only observe messages that have been sent or received by peers he controls (*i.e.*, the malicious ones). To the best of our knowledge, we are not aware of any previous work that has specified and characterized the conditions for which uniform and ergodic sampling is achievable in the presence of adversarial behavior.

¹Scalability of structured peer-to-peer systems rely on the assumption that nodes are uniformly distributed over the structured communication graph, while connectivity of unstructured peer-to-peer systems result from the assumption that nodes choose their neighbors arbitrarily.

The outline of this paper is the following. In Section 2, we give an overview of the existing related work, whereas Section 3 describes the model of the system and the particular assumptions that we make. In Section 4, we describe the functionalities of a sampling component and the properties that it should guarantee while in Section 5, we present some background on information divergence of data-streams. The omniscient and blind adversary models, as well as the characterization of the minimum effort the adversary has to exert to bias the sampling properties, are respectively studied in Sections 6 and 7. Finally, Section 8 presents some concluding remarks.

2 Related Work

Different approaches have been proposed to deal with malicious behaviors in the peer sampling problem in unstructured overlays. Jesi *et al.* [12] propose a random sampling algorithm taking explicitly into account malicious peers. Their solution assumes that the ultimate goal of the malicious peers is to mutate the random graph into a hub-based graph, hub for which malicious peers gain the lead. This approach, also adopted in several structured based overlays [19] through auditing mechanisms, or in sensor networks [15], is effective only if the number of malicious peers is very small with respect to the size of the system (*i.e.*, typically of $\mathcal{O}(\log n)$). Recently, Bortnikov *et al.* [6] have proposed a uniform but not ergodic peer sampling algorithm that tolerates up to a linear number of malicious peers. Their sampling mechanism exploits the properties offered by min-wise permutations. Specifically, the sampling component is fed with the stream of peer ids periodically gossiped by peers, and outputs the peer id whose image value under the randomly chosen permutation is the smallest value ever encountered. Thus eventually, by the property of min-wise permutation, the sampler converges towards a random but permanent sample. In a previous work [3], we show that imposing strict restrictions on the number of messages sent by malicious peers during a given period of time and providing each honest peer with a very large memory (proportional to the size of the system) are necessary and sufficient conditions to obtain uniform and ergodic (non permanent) sampling. It is worth noting that our results complement two previous results [7, 10], in which an analysis of the class of uniform and ergodic sampling protocols is presented. Both previous work provide a complete analytical proof of a gossip-based protocol that reaches both uniformity and ergodicity, but in contrast to the present work, adversarial behaviors were not considered. Finally, taking a completely different approach from the previously mentioned papers, which are based on gossip algorithms or on distance function properties, the techniques presented in [20, 21] rely on social network topologies to guard against Sybil attacks. Both protocols take advantage of the fact that Sybil attacks try to alter the fast mixing property of social networks to defend against these attacks. However, in presence of malicious peers with a high degree, the performance of both protocols degrade drastically.

3 System Model

Model of the Network. We consider a system populated by a large collection of peers in which each peer is assigned a unique and permanent random identifier from an m -bit identifier space. Peer identifiers (simply denoted *ids* in the following) are derived by applying some standard strong cryptographic hash function on peers intrinsic characteristics. The value of m (128 for the standard MD5 hash function) is chosen to be large enough to make

the probability of identifiers collision negligible. The system is subject to churn, which is classically defined as the rate of turnover of peers in the system [9]. Each peer knows only a small set of peers existing within the system and this knowledge generally varies according to the activity of the system. The particular algorithm used by peers to choose their neighbors and to route messages induces the resulting overlay topology. In this work, we consider only an unstructured overlay. Unstructured overlays are assumed to conform with random graphs, in the sense that relationships among peers are mostly set according to a random process. Note that in the following we indifferently use both terms *node* and *peer*.

Adversary. We assume the presence of malicious (*i.e.*, Byzantine) peers that try to manipulate the system by exhibiting undesirable behaviors. In our context, this amounts to dropping chosen messages and injecting new ones to bias the sample list maintained by non malicious peers. We model malicious behaviors through an adversary that fully controls these malicious peers. The adversary cannot control more than a fraction k ($0 < k < 1$) of malicious peers in the whole system. A node that is not malicious is said *honest*. We distinguish between two models of adversary. The *omniscient* adversary model that represents an adversary that is able to eavesdrop all messages exchanged within the system, and the *blind* adversary model that represents an adversary that can only observe messages sent or received by malicious nodes. In both models, we assume that the adversary cannot tamper with the content a message exchanged between two honest nodes without being detected. This is achieved by the means of a signature scheme that ensures the authenticity and integrity of messages.

Sampling Assumptions. Similarly to Bortnikov *et al.* [6], we assume that there exists a time T_0 such that after that time, the churn of the system ceases. This assumption is necessary to make the notion of uniform sample meaningful. Thus from T_0 onwards, the population of the system \mathcal{S} is composed of $n \ll 2^m$ nodes, such that at least $(1 - k)n$ of them are honest and no more than kn of them are controlled by the adversary. The subset of honest nodes in the overlay is denoted by \mathcal{N} and we assume that all the nodes in \mathcal{N} are weakly connected from time T_0 onwards.

4 Sampling Component

Following the lines of [3], each node $u \in \mathcal{N}$ has locally access to a *sampling component*² as presented in Figure 1. The sampling component implements a *strategy* s and has uniquely access to a data structure Γ_u , referred to as the *sampling memory*. The size of the sampling memory Γ_u is bounded and is denoted by $|\Gamma_u|$. The sampling component S_u^s is fed with a stream $\langle v_i, v_j, \dots \rangle$ of (possibly non unique) node ids that correspond to the node ids periodically received by node $u \in \mathcal{N}$ (*e.g.*, through gossip algorithms, or random walks for instance). The fingerprint of an input stream is a collection of weighted points in which each node id is weighted by the number of times this node id appears in the stream. Specifically, a stream of node ids can be summarized by $\langle (v_1, m_1), \dots, (v_n, m_n) \rangle$, where v_i denotes the identifier of a node in \mathcal{S} and $m_i \in \mathbb{N}$ represents the number of times v_i appears in the stream.

At each time t , the following three steps are atomically executed: the first element of the stream, say node id v , is given as input to the sampler component. The sampling component S_u^s reads v , and removes it from the stream. According to its strategy s , S_u^s may store or not v in Γ_u and outputs at most one node id. For example, the strategy s may consist in storing

²Although malicious nodes have also access to a sampling component, we cannot impose any assumptions on how they feed it or use it as their behavior can be totally arbitrary.




Figure 1: Sampling component of node $u \in \mathcal{N}$.

v if Γ_u is not full, or in substituting v for a randomly chosen node id that belongs to Γ_u , or simply dropping v . The output at time t , denoted $S_u^s(t)$, is chosen among the node ids in Γ_u according to strategy s . For instance, strategy s may consist in choosing the smallest node id in Γ_u , or the smallest node id under a given min-wise permutation [6]. The maximum finite hitting time needed for the sampling component S_u^s to reach a uniform sample is denoted by T_s . Clearly, T_s depends on the strategy s implemented by the sampling component and also on the stream of node ids the sampling component is fed with. We assume that the sampling strategy is known by the adversary in the sense that the algorithm used is public knowledge. However, if the algorithm is a randomized one, the adversary has not access to the local random coins used by the honest peers.

Finally, δ represents the number of node ids injected by the adversary in the input stream of node u during the time interval T_s . Note that it does not matter whether the injected node ids correspond to the ids of malicious nodes or not as the unique goal of the adversary is to bias the input stream in such a way that whatever the strategy s of the sampler component, its output $S_u^s(t)$ cannot guarantee both the uniform and ergodic properties [3]. More precisely, these properties are defined as follows

Property 4.1 (Uniformity) *Let \mathcal{N} be a weakly connected graph from time T_0 onwards, then for any time $t \geq T_s$, for any node $v \in \mathcal{S}$, and for any node $u \in \mathcal{N}$,*

$$\mathbb{P}[v \in S_u^s(t)] = \frac{1}{|\mathcal{S}|}.$$

Property 4.2 (Ergodicity) *Let \mathcal{N} be a weakly connected graph from time T_0 onwards, then for any time $t \geq T_s$, for any node $v \in \mathcal{S}$, and for any node $u \in \mathcal{N}$,*

$$\mathbb{P}[\{t' | t' > t \wedge v \in S_u^s(t')\} = \emptyset] = 0,$$

where \emptyset represents the empty set.

Uniformity states that any node in the overlay should have the same probability to appear in the sample of honest nodes in the overlay, while ergodicity says that any node should have a non-zero probability to appear infinitely often in the sample of each honest node in the overlay. Note that uniformity by itself does not imply ergodicity, and *vice versa*. Indeed, the former does not impose any restriction on the freshness of output node ids, while the latter one does not provide any guarantee regarding the equiprobability of node ids to be chosen as samples. Moreover, as each node v in \mathcal{S} has a non-zero probability to be returned by $S_u^s(t)$

at time t , u must appear at least once in the input stream. Thus, $\forall v \in \mathcal{S}$, starting from time T_s , $m_v > 0$.

Note that the analysis presented in this paper is independent from the way the stream of node ids at each node u has been generated. That is, it may result from the propagation of node ids through gossip-based algorithms (namely through push, or pull or push-pull mechanisms initiated by u and its neighbors), from the node ids received during random walks initiated at u , or even from induced churn.

5 Information Divergence of Data Streams

A natural approach to detect changes on data streams is to model it as a distribution and to compute the distance between the observed stream and the ideal one. The distance we use in our context is the Kullback-Liebler (KL) divergence (sometimes called the relative entropy [8]).

Definition 5.1 (Kullback-Liebler divergence) *Given two probability distributions on events $p = \{p_1, \dots, p_n\}$ and $q = \{q_1, \dots, q_n\}$, the Kullback-Liebler divergence between p_i relative to q_i is defined as the expected value of the likelihood ratio with respect to q_i :*

$$\mathcal{D}(p||q) = \sum_{i=1}^n p_i \log_2 \frac{p_i}{q_i} \quad (1)$$

$$= H(p, q) - H(p), \quad (2)$$

where $H(p) = -\sum p_i \log_2 p_i$ is the (Shannon) entropy of p and $H(p, q) = -\sum p_i \log_2 q_i$ is the cross entropy of p and q (by convention, $0 \log_2 0 = 0$).

The KL-divergence is a member of a larger class of distances known as the Ali-Silvey distances [1]. For the sake of clarity, we will use the notation \log to denote the logarithm in base 2 for the rest of this paper.

Definition 5.2 (τ -closeness) *A stream of node ids σ is τ -close if the KL-divergence between the uniform probability distribution $p^{(U)}$ and the probability distribution q corresponding to σ is below or equal to a robustness threshold τ .*

Given two probability distributions, the Earth Mover's Distance (EMD) [18] measures the minimal amount of work needed to transform one distribution to another by moving the probability mass between events. We rely on this metric to quantify the effort that an adversary exerts to bias the input stream. In our context, a unit of work corresponds to dropping one id and to pushing another id instead in the input stream.

6 Omniscient Adversary Model

In this section, we study the behavior of an omniscient adversary, which has the capacity to eavesdrop on all the messages sent and received by all the nodes in \mathcal{S} . In the following, we demonstrate that the strategy that pushes all the probability mass over a single id is the one that maximizes the bias of the input stream so that it becomes far from the uniform distribution. We also describe an optimal strategy on how to achieve it.

In the following, we consider input streams of length T_s observed from time T_0 . Let $\bar{\sigma}$ be a stream such that the id of each node in \mathcal{S} appears exactly once in the stream, except for a unique id that appears in all the remaining slots. Therefore, we have $\exists v_i \in \mathcal{S}, m_{v_i} = T_s - (n-1)$ and $\forall v_j \neq v_i \in \mathcal{S}, m_{v_j} = 1$. The following theorem states that the probability distribution associated to this particular stream is the one that has the maximal divergence from the uniform distribution. In the following, notation $[1..n]$ denotes the set $\{1, 2, \dots, n\}$.

Theorem 6.1 (Maximal divergence from the uniform distribution) *Let $p^{(u)}$ be the uniform distribution corresponding to a uniform stream, i.e., $\forall i \in [1..n], p_i^{(u)} = \frac{1}{n}$, and \bar{q} be the probability distribution corresponding to $\bar{\sigma}$, i.e., $\exists v_i \in \mathcal{S}, \bar{q}_{v_i} = \frac{T_s - (n-1)}{T_s}$ and $\forall v_j \in \mathcal{S}, v_j \neq v_i, \bar{q}_{v_j} = \frac{1}{T_s}$. Then, for any possible probability distribution q ,*

$$\mathcal{D}(p^{(u)}||q) \leq \mathcal{D}(p^{(u)}||\bar{q}).$$

Proof. Let q be the probability distribution representing any valid input stream on $(T_0, T_s]$. We have $\forall v_i \in \mathcal{S}, q_{v_i} = \frac{m_{v_i}}{T_s}$, where m_{v_i} is the number of times v_i is present in the input stream. Therefore,

$$\begin{aligned} \mathcal{D}(p^{(u)}||q) &= H(p^{(u)}, q) - H(p^{(u)}) = \sum_{i=1}^n p_i^{(u)} \log(p_i^{(u)}) - \sum_{i=1}^n p_i^{(u)} \log\left(\frac{m_{v_i}}{T_s}\right) \\ &= \log\left(\frac{T_s}{n}\right) - \frac{1}{n} \log\left(\prod_{i=1}^n m_{v_i}\right). \end{aligned}$$

Therefore, maximizing $\mathcal{D}(p^{(u)}||q)$ amounts to minimizing $\log(\prod_{i=1}^n m_{v_i})$. We characterize the stream that minimizes $\prod_{i=1}^n m_{v_i}$ under the following constraints:

$$\begin{cases} 1 \leq m_{v_i} \leq T_s & \text{with } 1 \leq i \leq n, \\ \sum_{i=1}^n m_{v_i} = T_s. \end{cases} \quad (3)$$

From this set of constraints, we immediately have $1 \leq m_{v_i} \leq T_s - (n-1)$. To relax the second constraint, we pose $m_{v_n} = T_s - \sum_{i=1}^{n-1} m_{v_i}$. Let function f be such that

$$f(m_{v_1}, \dots, m_{v_{n-1}}) = \left(T_s - \sum_{i=1}^{n-1} m_{v_i}\right) \prod_{i=1}^{n-1} m_{v_i} = m_{v_j} \left(T_s - m_{v_j} - \sum_{i=1, i \neq j}^{n-1} m_{v_i}\right) \prod_{i=1, i \neq j}^{n-1} m_{v_i}.$$

Function f is differentiable on its domain $\mathcal{I}_s = [1..T_s - n + 1]^{n-1}$, thus we get

$$\frac{df}{dm_{v_j}}(m_{v_1}, \dots, m_{v_{n-1}}) = \left(T_s - 2m_{v_j} - \sum_{i=1, i \neq j}^{n-1} m_{v_i}\right) \prod_{i=1, i \neq j}^{n-1} m_{v_i}.$$

We have

$$\frac{df}{dm_{v_j}} \geq 0 \Leftrightarrow m_{v_j} \leq \frac{T_s - \sum_{i=1, i \neq j}^{n-1} m_{v_i}}{2} < T_s - (n-1).$$

Function f is then a paraboloid, which first increases then decreases, and then reaches its minimum on the domain's boundary. As f is symmetric and all m_{v_j} variables are independent, the minimum for each m_{v_j} can be determined separately. The minimum for a particular m_{v_j} is reached when $m_{v_j} = 1$ or $m_{v_j} = (T_s - n + 1)$. That is,

$$\begin{cases} f(m_{v_1}, \dots, 1, \dots, m_{v_{n-1}}) & = \prod_{i=1, i \neq j}^{n-1} m_{v_i} \left(T_s - 1 - \sum_{i=1, i \neq j}^{n-1} m_{v_i} \right) \\ f(m_{v_1}, \dots, T_s - n + 1, \dots, m_{v_{n-1}}) & = (T_s - n + 1) \prod_{i=1, i \neq j}^{n-1} m_{v_i} \left(n - 1 - \sum_{i=1, i \neq j}^{n-1} m_{v_i} \right) \end{cases}$$

$$\begin{aligned} f(m_{v_1}, \dots, T_s - n + 1, \dots, m_{v_{n-1}}) & = \prod_{i=1, i \neq j}^{n-1} m_{v_i} \left((T_s - n + 1) \left(n - 2 - \sum_{i=1, i \neq j}^{n-1} m_{v_i} \right) + T_s - n + 1 \right) \\ & = \prod_{i=1, i \neq j}^{n-1} m_{v_i} \left((T_s - n) \left(n - 2 - \sum_{i=1, i \neq j}^{n-1} m_{v_i} \right) + T_s - 1 - \sum_{i=1, i \neq j}^{n-1} m_{v_i} \right). \end{aligned}$$

As $\forall i, m_{v_i} \geq 1$, we have $n - 2 - \sum_{i=1, i \neq j}^{n-1} m_{v_i} \leq 0$. Moreover, as $T_s \geq n$, we have $f(m_{v_1}, \dots, T_s - n + 1, \dots, m_{v_{n-1}}) \leq f(m_{v_1}, \dots, 1, \dots, m_{v_{n-1}})$ and therefore the minimum is reached for $m_{v_j} = T_s - n + 1$. From the set of constraints (*cf.* Relation (3)), if the maximum of $\mathcal{D}(p^{(\mathcal{U})}||q)$ is reached for $m_{v_j} = T_s - n + 1$ then $\sum_{i=1, j \neq i}^n m_{v_i} = n - 1$ implies that $\forall i \in [1..n], i \neq j, m_{v_i} = 1$, which concludes the proof. \square

This allows us to formulate an upper-bound \mathcal{D}^{\max} on the KL-divergence between the uniform stream and any other stream:

$$\mathcal{D}^{\max} = \mathcal{D}(p^{(\mathcal{U})}||q) = \log \left(\frac{T_s}{n} \right) - \frac{1}{n} \log (T_s - n + 1). \quad (4)$$

Thus any input stream σ is \mathcal{D}^{\max} -close (*cf.* Definition 5.2).

To determine the minimal effort that the adversary has to exert to bias the input stream so that both uniformity and ergodicity properties do not hold, we use the Earth Mover's Distance (EMD) between the uniform distribution and the target one. In the following when we say that the adversary *replaces* the node id v_i by v_j , we mean that he drops v_i from the input stream and inject v_j instead.

Lemma 6.2 (Optimal strategy to maximize the divergence) *Given an input stream σ , replacing the less frequent id in σ with the most frequent one maximizes the gain in KL-divergence with respect to the uniform distribution for the same amount of work as measured by the EMD distance.*

Proof. Given an input stream σ represented by the probability distribution q , we construct the input stream σ' from σ by substituting one occurrence of node id v_i with node id v_j so that $\mathcal{D}(p^{(\mathcal{U})}||q')$ is maximized after this replacement (where q' denote the probability distribution representing σ'). This amounts to maximizing $[\mathcal{D}(p^{(\mathcal{U})}||q') - \mathcal{D}(p^{(\mathcal{U})}||q)]$. Recall that all node ids in \mathcal{S} must be present in σ' . Therefore, we search for the node id pair (v_i, v_j) such that

$$\begin{cases} m'_{v_j} & = m_{v_j} + 1 \\ m'_{v_i} & = m_{v_i} - 1 \\ v_j & = \arg \max_{v_j \in \mathcal{S}} \left(\log \left(\frac{1}{q'_{v_j}} \right) - \log \left(\frac{1}{q_{v_j}} \right) \right) \\ v_i & = \arg \max_{v_i \in \mathcal{S}} \left(\log \left(\frac{1}{q'_{v_i}} \right) - \log \left(\frac{1}{q_{v_i}} \right) \right) \end{cases}$$

leading to

$$\begin{cases} v_j = \arg \max_{v_j \in \mathcal{S}} \left(1 - \frac{1}{m_{v_j} + 1}\right) = \arg \min_{v_j \in \mathcal{S}} \frac{1}{m_{v_j} + 1} = \arg \max_{v_j \in \mathcal{S}} m_{v_j} \\ v_i = \arg \max_{v_i \in \mathcal{S}} \left(1 + \frac{1}{m_{v_i} - 1}\right) = \arg \max_{v_i \in \mathcal{S}} \frac{1}{m_{v_i} - 1} = \arg \min_{v_i \in \mathcal{S}} m_{v_i} \end{cases}$$

Thus the optimal node id replacement that maximizes the KL-divergence gain is obtained by replacing the less frequent node id v_i with the most frequent one v_j . \square

Algorithm 1 shows an optimal implementation of Lemma 6.2 with respect to the number of performed replacements. This algorithm is run by the adversary. Specifically, the inputs of the algorithm are τ_s and an input stream σ that will feed the sampler component S_u^s of node u . Recall that τ_s is the robustness threshold of the sampling strategy s implemented by S_u^s , *i.e.*, for any τ_s -close input stream σ , the sampling strategy s is capable of outputting a uniform and ergodic sample. The goal of the greedy Algorithm 1 is to tamper with the input stream σ in order to increase its KL-divergence above τ_s with a minimum effort.

Algorithm 1: Adversary biasing strategy

Data: an input stream σ , the robustness threshold τ_s

Result: the number of replacements ℓ if it exists

```

1 if  $\tau_s \geq \mathcal{D}^{\max}$  then
2   return "fail"
3 else
4    $\ell \leftarrow 0$ ;
5    $v_j \leftarrow \arg \max_{v_j \in \mathcal{S}} m_{v_j}$ ;
6   while  $(\mathcal{D}(p^{(\mathcal{U})} || q_\sigma) \leq \tau_s)$  do
7      $v_i \leftarrow \arg \min_{\{v \in \mathcal{S} : m_{v_i} \neq 1\}} m_{v_i}$ ;
8     let  $k$  be the index of an id such that  $\sigma[k] = v_i$  ;
9      $\sigma[k] \leftarrow v_j$  //one occurrence of  $v_i$  is dropped and  $v_j$  is injected instead;
10     $\ell \leftarrow \ell + 1$ ;
11  return  $\ell$ 

```

By assumption, the adversary is omniscient and therefore has the capacity to observe the entire input stream σ . From Section 4, the adversary knows the strategy s of the sampler, and thus can compute the value of τ_s . The value of the maximum divergence \mathcal{D}^{\max} is computed as Relation (4). If \mathcal{D}^{\max} is greater than or equal to the robustness threshold, the algorithm returns "fail". Otherwise at each iteration, the adversary performs the optimal id replacement until the KL-divergence exceeds the robustness threshold. Note that at lines (8) and (9) both m_{v_i} and m_{v_j} are updated. The counter ℓ returned by Algorithm 1 represents the number of replacements done by the adversary.

Consider s a sampling strategy, τ_s its robustness threshold, and σ an input stream. Let ℓ be the number of replacements executed in Algorithm 1. We denote $q_{\sigma(\ell)}$ the probability distribution derived from σ after these ℓ optimal replacements.

Corollary 6.3 (Lower bound on the effort exerted by an omniscient adversary) *The minimum number of replacements an omniscient adversary has to apply to exceed τ_s is*

$$\delta = \inf \left\{ \ell \in \mathbb{N} : D(p^{(\mathcal{U})} || q_{\sigma(\ell)}) > \tau_s \right\}. \quad (5)$$

7 Blind Adversary Model

In this section, we study the behavior of a blind adversary, that is an adversary that only has the capacity to observe messages sent or received by the nodes he controls. A strategy that the adversary might apply to bias the input stream is to choose a node id (possibly one that belongs to a malicious node but not necessarily) and to push it in the input stream as much as possible.

We show that this strategy is optimal and we give the lower bound on the expected minimum amount of work a blind adversary has to exert to bias the input stream.

Theorem 7.1 (Lower bound on the *expected* effort exerted by a blind adversary)

Let s be a sampling strategy, τ_s its robustness threshold and T_s the maximum convergence time of s . The minimum number of replacements a blind adversary has to apply to exceed τ_s in expectation is given when the input stream is the uniform one. We have

$$\begin{aligned} \tilde{\delta} = \inf \left\{ \ell \in \mathcal{I}_s : \frac{1}{n} \left(\left\lfloor \frac{\ell - 1}{\lfloor \frac{T_s}{n} - 1 \rfloor} \right\rfloor \log \left(\frac{T_s}{n} \right) + 2 \log (T_s) - \log (T_s + n\ell) \right. \right. \\ \left. \left. - \log \left(T_s - n \left(1 + (\ell - 1) \bmod \left\lfloor \frac{T_s}{n} - 1 \right\rfloor \right) \right) \right) > \tau_s \right\} \end{aligned} \quad (6)$$

where $\mathcal{I}_s = [0..T_s - n + 1 - \lfloor \frac{T_s}{n} \rfloor]$.

Proof. Let us consider the uniform node ids stream on a window of length T_s . For any $v_i \in \mathcal{S}$, v_i is present in the stream T_s/n in average. Thus the probability distribution $p^{(\mathcal{U})}$ is such that $\forall v_i \in \mathcal{S}, p_{v_i}^{(\mathcal{U})} = 1/n$. From the previous section, we have seen that the optimal strategy for the adversary to bias an input stream is to replace the less frequent node id in this stream with the most frequent one. By assumption, the adversary is blind and cannot observe all the node ids of the input stream. Thus the strategy the adversary applies consists choosing a specific node id v_j and repeatedly pushes v_j in the input stream. Let σ be an input stream and σ' be the stream obtained from σ after one step of this adversarial strategy (*i.e.*, replacing v_i by v_j for some $v_i \in \mathcal{S}$). We have

$$\mathcal{D}(p^{(\mathcal{U})} || q_{\sigma'}) - \mathcal{D}(p^{(\mathcal{U})} || q_{\sigma}) = \frac{1}{n} \left(\log \left(\frac{m_{v_j}}{m_{v_j} + 1} \right) + \log \left(\frac{m_{v_i}}{m_{v_i} - 1} \right) \right) \quad (7)$$

where q_{σ} and $q_{\sigma'}$ represent respectively the probability distributions of σ and σ' . In the following, $q_{\sigma(\ell)}$ denotes the probability distribution derived from σ after ℓ replacements. Given a sampling strategy s , we prove by induction on the number of optimal replacements ℓ that, starting from a uniform stream, the maximum KL-divergence after ℓ replacements is given by

$$\begin{aligned} \mathcal{D}(p^{(\mathcal{U})} || q_{\sigma(\ell)}) = \frac{1}{n} \left(\left\lfloor \frac{\ell - 1}{\lfloor \frac{T_s}{n} - 1 \rfloor} \right\rfloor \log \left(\frac{T_s}{n} \right) + 2 \log (T_s) - \log (T_s + n\ell) \right. \\ \left. - \log \left(T_s - n \left(1 + (\ell - 1) \bmod \left\lfloor \frac{T_s}{n} - 1 \right\rfloor \right) \right) \right). \end{aligned} \quad (8)$$

Note that ℓ cannot be greater than $(T_s - n + 1 - \lfloor \frac{T_s}{n} \rfloor)$. Indeed, all node ids in the initial uniform stream are present at least $\lfloor \frac{T_s}{n} \rfloor$ times and the maximum number of times a unique id can appear in the stream is $(T_s - n + 1)$.

For $\ell = 1$, the claim immediately holds from Equation 7. Now, assume that the claim also holds for all $1 \leq k \leq \ell$. We show that the claim holds for $k = \ell + 1$. The KL-divergence with respect to the uniform stream after $\ell + 1$ steps is

$$\mathcal{D}(p^{(\mathcal{U})} || q_{\sigma(\ell+1)}) = \mathcal{D}(p^{(\mathcal{U})} || q_{\sigma(\ell)}) + \mathcal{D}(p^{(\mathcal{U})} || q_{\sigma(\ell+1)}) - \mathcal{D}(p^{(\mathcal{U})} || q_{\sigma(\ell)}). \quad (9)$$

$\mathcal{D}(p^{(\mathcal{U})} || q_{\sigma(\ell)})$ is given by Equation 8, and $(\mathcal{D}(p^{(\mathcal{U})} || q_{\sigma(\ell+1)}) - \mathcal{D}(p^{(\mathcal{U})} || q_{\sigma(\ell)}))$ represents the gain of step $(\ell + 1)$. Two sub-cases need to be considered: (i) $\ell \bmod \lfloor \frac{T_s}{n} - 1 \rfloor \neq 0$ and (ii) $\ell \bmod \lfloor \frac{T_s}{n} - 1 \rfloor = 0$.

Case (i): the less frequent node id v_i in the stream at step $\ell + 1$ is the same as the one removed at step ℓ . After ℓ steps, $m_{v_j} = \frac{T_s}{n} + \ell$ and $m_{v_i} = \frac{T_s}{n} - (1 + (\ell - 1) \bmod \lfloor \frac{T_s}{n} - 1 \rfloor)$, thus the right part of Equation 7 is equal to

$$\begin{aligned} & \frac{1}{n} \left(\log \left(\frac{\frac{T_s}{n} + \ell}{\frac{T_s}{n} + \ell + 1} \right) + \log \left(\frac{\frac{T_s}{n} - (1 + (\ell - 1) \bmod \lfloor \frac{T_s}{n} - 1 \rfloor)}{\frac{T_s}{n} - (1 + (\ell - 1) \bmod \lfloor \frac{T_s}{n} - 1 \rfloor) - 1} \right) \right) \\ &= \frac{1}{n} \left(\log(T_s + n\ell) - \log(T_s + n(\ell + 1)) + \log \left(T_s - n \left(1 + (\ell - 1) \bmod \lfloor \frac{T_s}{n} - 1 \rfloor \right) \right) \right. \\ & \quad \left. - \log \left(T_s - n \left(2 + (\ell - 1) \bmod \lfloor \frac{T_s}{n} - 1 \rfloor \right) \right) \right). \end{aligned}$$

By assumption (i), $\lfloor \frac{\ell-1}{\lfloor \frac{T_s}{n} - 1 \rfloor} \rfloor = \lfloor \frac{\ell}{\lfloor \frac{T_s}{n} - 1 \rfloor} \rfloor$ and $(1 + (\ell - 1) \bmod \lfloor \frac{T_s}{n} - 1 \rfloor) = (\ell \bmod \lfloor \frac{T_s}{n} - 1 \rfloor)$.

From Equation 9, we get

$$\begin{aligned} \mathcal{D}(p^{(\mathcal{U})} || q_{\sigma(\ell+1)}) &= \frac{1}{n} \left(\left\lfloor \frac{\ell}{\lfloor \frac{T_s}{n} - 1 \rfloor} \right\rfloor \log \left(\frac{T_s}{n} \right) + 2 \log(T_s) \right. \\ & \quad \left. - \log(T_s + n(\ell + 1)) - \log \left(T_s - n \left(1 + \ell \bmod \lfloor \frac{T_s}{n} - 1 \rfloor \right) \right) \right), \end{aligned}$$

which ends Case (i).

Case (ii). The argumentation is the same as above. However, as $\ell \bmod \lfloor \frac{T_s}{n} - 1 \rfloor = 0$, the node id that has been previously replaced is now present exactly once in the stream. Thus the adversary needs to randomly choose another node id in the stream before processing the next step of his strategy. Thus applying Equation 7 at step $\ell + 1$ gives

$$\mathcal{D}(p^{(\mathcal{U})} || q_{\sigma(\ell+1)}) - \mathcal{D}(p^{(\mathcal{U})} || q_{\sigma(\ell)}) = \frac{1}{n} \left(\log \left(\frac{\frac{T_s}{n} + \ell}{\frac{T_s}{n} + \ell + 1} \right) + \log \left(\frac{\frac{T_s}{n}}{\frac{T_s}{n} - 1} \right) \right). \quad (10)$$

By assumption (($\ell - 1) \bmod \lfloor \frac{T_s}{n} - 1 \rfloor = \lfloor \frac{T_s}{n} - 1 \rfloor - 1$), and by combining the induction hypothesis 8 with the gain obtained at step $\ell + 1$ (Equation 10) we get

$$\mathcal{D}(p^{(\mathcal{U})} || q_{\sigma(\ell+1)}) = \frac{1}{n} \left(\left\lfloor \frac{\ell - 1}{\lfloor \frac{T_s}{n} - 1 \rfloor} \right\rfloor \log \left(\frac{T_s}{n} \right) + 3 \log(T_s) \right. \\ \left. - \log(T_s + n(\ell + 1)) - \log(T_s - n) - \log(n) \right).$$

By assumption of the case we have $\left\lfloor \frac{\ell}{\left\lfloor \frac{T_s}{n} - 1 \right\rfloor} \right\rfloor = \left\lfloor \frac{\ell-1}{\left\lfloor \frac{T_s}{n} - 1 \right\rfloor} \right\rfloor + 1$, which proves the induction:

$$\mathcal{D}(p^{(\ell)} || q_{\sigma^{(\ell+1)}}) = \frac{1}{n} \left(\left\lfloor \frac{\ell}{\left\lfloor \frac{T_s}{n} - 1 \right\rfloor} \right\rfloor \log \left(\frac{T_s}{n} \right) + 2 \log(T_s) - \log(T_s + n(\ell + 1)) - \log \left(T_s - n \left(1 + \ell \bmod \left\lfloor \frac{T_s}{n} - 1 \right\rfloor \right) \right) \right).$$

As a conclusion, any value of ℓ that allows the adversary to exceed the robustness threshold τ_s defeats the sampling strategy. Thus, the minimum number of replacement operations $\tilde{\delta}$ is the lower bound of this set of values. \square

We now evaluate the minimum amount of work a blind adversary has to exert, in the worst case, to bias the input stream. In the worst case, the node id v_i the adversary has chosen to blindly flood might be initially present only once in the input stream. In order to bias the input stream, the adversary needs to push id v_i sufficiently often so that the probability of appearance of id v_i reaches the uniform value, with respect to all the other node ids, and then to continue to push this id $\tilde{\delta}$ times so that the divergence between the resulting stream and the uniform one is maximum.

Theorem 7.2 (Lower bound on the effort exerted by a blind adversary) *Let s be a sampling strategy, τ_s its robustness threshold and T_s the maximum convergence time of s . The minimum number of replacements the adversary has to apply on a stream, in the worst case, to exceed τ_s is*

$$\tilde{\delta} + \left\lceil \frac{T_s}{n} \right\rceil - 1$$

Proof. The proof is immediate. First, the adversary has to raise the chosen id at least up to the uniform value. As in the worst case, this id is present only once in the initial stream, this costs $\left\lceil \frac{T_s}{n} \right\rceil - 1$ replacements to reach a number of occurrences equals to $\left\lceil \frac{T_s}{n} \right\rceil$. Moreover, once this id is present in the modified stream $\left\lceil \frac{T_s}{n} \right\rceil$ times, the adversarial follows the same strategy as before, which requires $\tilde{\delta}$ more steps to guarantee that the robustness threshold τ_s is exceeded. Note that this value is a worst-case bound and not the exact minimum value with respect to τ_s because after the first $(\left\lceil \frac{T_s}{n} \right\rceil - 1)$ steps, the modified stream could be different from the uniform one. In this situation, the KL-divergence to the uniform stream is strictly greater than 0, reducing accordingly the amount of work of the adversary to exceed τ_s . \square

8 Concluding Remarks

In this paper, we have focused on the problem of achieving uniform and ergodic peer sampling in large scale open systems potentially populated with malicious peers. This problem consists in guaranteeing that the knowledge of the system maintained by each honest peer is a uniform and non permanent sample of the whole population of the system. By modeling input streams as probability distributions, we have characterized the minimum effort (measured in terms of node ids' replacements) that an omniscient and a blind adversary have to exert on the input stream of node identifiers to exceed the robustness threshold that quantifies the power of a sampling strategy. Similarly to (pseudo)-random number generators that are considered as




Figure 2: Characterization of Sampling Strategies

the basic mathematical tool to generate complex probability distributions, we believe that uniform peer sampling should be regarded as a necessary building block to derive larger classes of sampling schemes. This building block is of utmost importance in systems in which the population is continuously evolving and thus, where it is impossible to capture the full complexity of the network through global snapshots. In our context, this dynamicity is exploited by the adversary to bias the peer sampling by over-representing targeted peers.

We conjecture that there exists a total order relationship on the power of ergodic (resp. non ergodic) uniform sampling strategies with respect to their robustness threshold τ . This can be depicted by using a planar representation as shown by Figure 2. Each τ -radius circle in this picture represents the class of τ -close uniform sampling strategies. For instance, in this representation the strategy proposed by Bortnikov *et al.* [6] belongs to the largest circle corresponding to $\tau = \infty$ on its non ergodic part (left side). As another illustrative example, the sampling strategy proposed in Busnel *et al.* [7] should be ranked as less powerful than the one proposed by Gurevich *et al.* [10] since the latter one achieves uniform and ergodic sampling despite message loss contrary to the former one. Both strategies would appear in the ergodic part of the representation (right side). Finally, we think that this classification can be used as a tool to precisely compare any two sampling strategies as different as they are (*e.g.*, [4] and [17]). As future work, we intend to rank the state-the-art sampling algorithms in the light of this new framework. Moreover, we plan to extend this work by integrating other dimensions, such as space and time resources, and deterministic *versus* probabilistic strategies.

References

- [1] S. M. Ali and S. D. Silvey. General Class of Coefficients of Divergence of One Distribution from Another. *Journal of the Royal Statistical Society. Series B (Methodological)*, 28(1):131–142, 1966.
- [2] E. Anceaume, F. V. Brasileiro, R. Ludinard, B. Sericola, and F. Tronel. Dependability Evaluation of Cluster-based Distributed Systems. *International Journal of Foundations*

- of *Computer Science (IJFCS)*, 2011. To appear.
- [3] E. Anceaume, Y. Busnel, and S. Gambs. Uniform and Ergodic Sampling in Unstructured Peer-to-Peer Systems with Malicious Nodes. In *Proceedings of the 14th International Conference On Principles Of Distributed Systems (OPODIS)*, volume 6490, pages 64–78, 2010.
 - [4] B. Awerbuch and C. Scheideler. Towards a Scalable and Robust Overlay Network. In *Proceedings of the 6th International Workshop on Peer-to-Peer Systems (IPTPS)*, 2007.
 - [5] M. Bertier, Y. Busnel, and A.-M. Kermarrec. On Gossip and Populations. In *Proceedings of the 16th International Colloquium on Structural Information and Communication Complexity (SIROCCO)*, 2009.
 - [6] E. Bortnikov, M. Gurevich, I. Keidar, G. Kliot, and A. Shraer. Brahms: Byzantine Resilient Random Membership Sampling. *Computer Networks*, 53:2340–2359, 2009. A former version appeared in the 27th ACM Symposium on Principles of Distributed Computing (PODC), 2008.
 - [7] Y. Busnel, R. Beraldi, and R. Baldoni. A Formal Characterization of Uniform Peer Sampling Based on View Shuffling. In *Proceedings of the International Conference on Parallel and Distributed Computing, Applications and Technologies (PDCAT)*, pages 360–365. IEEE Computer Society, 2009.
 - [8] T. Cover and J. Thomas. Elements of information theory. *Wiley New York*, 1991.
 - [9] P. B. Godfrey, S. Shenker, and I. Stoica. Minimizing churn in distributed systems. In *Proceedings of the ACM SIGCOMM*, 2006.
 - [10] M. Gurevich and I. Keidar. Correctness of Gossip-Based Membership under Message Loss. In *Proceedings of the 28th annual Symposium on Principles of distributed computing (PODC)*, Calgary, AL, Canada, 2009. ACM Press.
 - [11] M. Jelasity, S. Voulgaris, R. Guerraoui, A.-M. Kermarrec, and M. van Steen. Gossip-based Peer Sampling. *ACM Transaction on Computer System*, 25(3), 2007.
 - [12] G. P. Jesi, A. Montresor, and M. van Steen. Secure Peer Sampling. *Computer Networks*, 54(12):2086–2098, 2010.
 - [13] D. R. Karger and M. Ruhl. Simple Efficient Load Balancing Algorithms for Peer-to-Peer. In *Proceedings of the 3rd International Workshop on Peer-to-Peer Systems (IPTPS)*, 2004.
 - [14] S. Le Blond, F. Le Fessant, and E. Le Merrer. Finding Good Partners in Availability-Aware P2P Networks. In *Proceedings of the 11th International Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS)*, pages 472–484. Springer-Verlag, 2009.
 - [15] D. Liu, P. Ning, and W. Du. Detecting Malicious Beacon Nodes for Secure Location Discovery in Wireless Sensor Networks. In *Proceedings of the 25th IEEE International Conference on Distributed Computing Systems (ICDCS)*, 2005.

-
- [16] Q. Lv, P. Cao, E. Cohen, K. Li, and S. Shenker. Search and Replication in Unstructured Peer-to-Peer Networks. In *Proceedings of the International Conference on Supercomputing (ICS)*, pages 84–95, 2002.
- [17] L. Massoulié, E. L. Merrer, A.-M. Kermarrec, and A. Ganesh. Peer Counting and Sampling in Overlay Networks: Random Walk Methods. In *Proceedings of the 25th Annual Symposium on Principles of Distributed Computing (PODC)*, pages 123–132. ACM Press, 2006.
- [18] G. Monge. Mémoire sur la théorie des déblais et des remblais. *Histoire de l'Académie royale des sciences, avec les Mémoires de Mathématique et de Physique*, pages 666–704, 1781.
- [19] A. Singh, T.-W. Ngan, P. Druschel, and D. S. Wallach. Eclipse Attacks on Overlay Networks: Threats and Defenses. In *Proceedings of the 25th IEEE International Conference on Computer Communications (INFOCOM)*, 2006.
- [20] H. Yu, P. B. Gibbons, M. Kaminsky, and F. Xiao. SybilLimit: A Near-Optimal Social Network Defense against Sybil Attacks. In *Proceedings of the IEEE Symposium on Security and Privacy (SP)*, pages 3–17, 2008.
- [21] H. Yu, M. Kaminsky, P. B. Gibbons, and A. Flaxman. SybilGuard: Defending against Sybil Attacks via Social Networks. In *Proceedings of the ACM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM)*, pages 267–278, 2006.