Identification Systems
Lecture 5

Piotr Syga

2021-04-13

Tracking prevention

M. Klonowski, M. Kutyłowski, PS: *Chameleon RFID and Tracking Prevention*

- Constrained RFID tags without secrets
- Adversary may capture and temper with the tags
- Honest tags security is not compromised by capturing others
- Adversary is not ubiquitous
- Even compromised tags, when not observed *for a while* become untrackable again
- **Idea:** evolution of the ID.

Model

- We consider a system consisting of an external identification authority (database), a set of readers connected to the database and a set of Class 2 RFID Tags with unique IDs sampled from the the subspace $\mathcal{E} = \{x \in \{0,1\}^{2n} : 2|n \wedge 2|W_H(x)\}$.
- Each tag stores exactly 2 ID strings – one is the *previousID*, the other *currentID*. The protocol will evolve the IDs.
- The database stores *permanentID* of each tag, coupled with its last *presentedID*.
- The adversary is not omnipresent it has limited number of readers, each of those can eavesdrop or initiate the communication with a single tag during each round.
- The legitimate readers have read and write permissions on the *presentedID* column and read permission on *permanentID* – we assume that this communication is secure and is immediately repeated after failure (during a single round).

The Protocol

	RFID		Database
1.	$z :=$ ***currentID***		
2.		$\rightarrow^{z}$	
3.			**find row** (***presentedID***, ***permanentID***) **such that** $D_H(z,$ ***presentedID***$) == n$
4.			**select** n **indices that differ between** z **and** ***presentedID***, **and store it in** L
5.		$\leftarrow^{L}$	
6.	**check if** b_i **for** $i \in L$ ***currentID***$_i \neq$ ***previousID***$_i$		
7.	**in case of difference – abort**		
8.	***previousID*** $=$ ***currentID***		**update:** ***presentedID*** $= z$
9.	***currentID*** $=$ UPDATE(***currentID***)		

Figure: A single round of Chameleon protocol.

Updating the IDs

After each successful communication, the Tag changes its ID by selecting at random exactly n positions and reversing bits on the given index. The idea is to prevent the adversary that skipped several rounds from recognizing the tag.

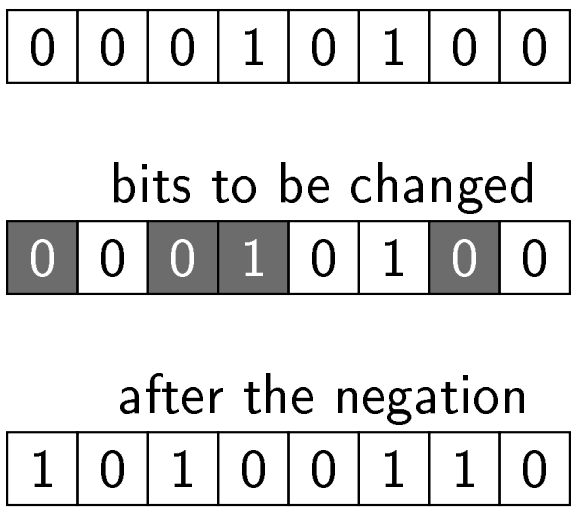

Figure: First round of the evolution ($ID_0 \rightarrow ID_1$).

Updating the IDs

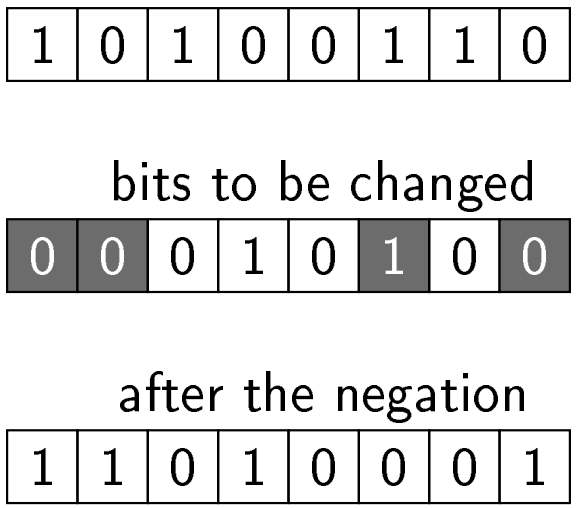

Figure: Second round of the evolution ($ID_1 \rightarrow ID_2$).

Updating the IDs

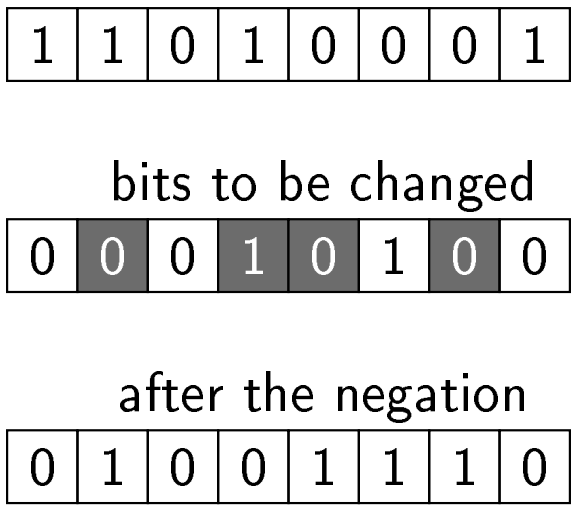

Figure: Third round of the evolution ($ID_2 \rightarrow ID_3$).

Updating the IDs

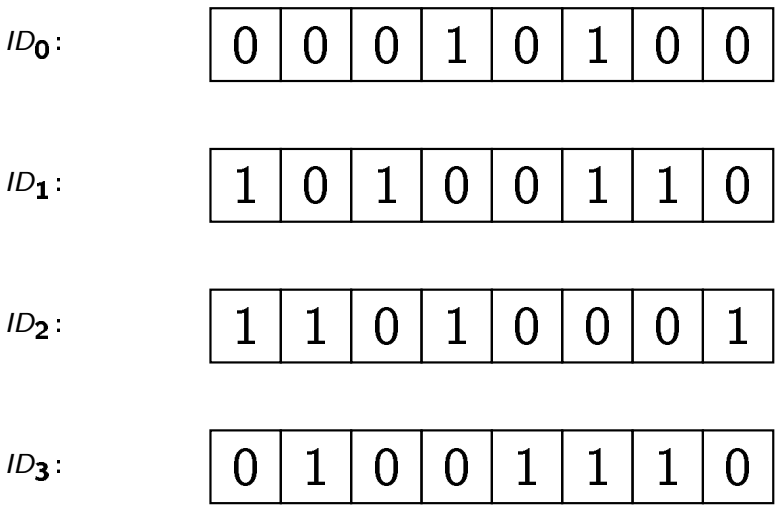

Figure: IDs of the same tag during the evolution.

Updating the IDs

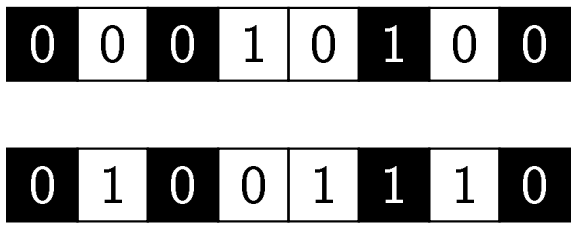

Figure: ID of a tag before the evolution and after 3 rounds of changes.

Getting a proper ID

- each evolution round preserves the the properties of the ID (i.e. new ID still is in $\mathcal{E}$)
- any proper ID may be reached within 2 rounds with positive probability
- evolution in each round is independent – modeling with Markov Chain
- sufficiently large set of possible new IDs

Number of possible new IDs

Let us consider a graph G with vertices $V = \{0, 1\}^{2n}$. Vertices a and b are connected **iff** $D_H(a, b) = n$.
Due to the Update procedure size of the neighborhood of each vertex is $\binom{2n}{n}$, hence $(n+1)C_n$.
Asymptotic of Catalan numbers C_n – $\frac{4^n}{\sqrt{n^3\pi}}$, hence $\frac{1}{\sqrt{n\pi}}$ of the ID space is connected with each node – for $n = 10$ it's 0.174 of the space, for $n = 20$ – 0.125, for $n = 40$ – 0.089, ...

Reachability of any state

- any ID with even Hamming weight may be reached within 2 rounds starting with an arbitrary ID $I \in \mathcal{E}$
- lets fix two identifiers ID, ID' from $\mathcal{E}$
- let $L \neq \emptyset$ be the set of all positions on which the two IDs differ – obviously $2| |L|$
- let us divide L into two disjoint sets: L_1, L_2 of an equal size
- let us select a set of indices A, s.t. $A \cap L = \emptyset$ and $|A| = n - |L_1|$
- in Update1 we negate all bits from L_1 and A, in Update2 we negate all bits from A and L_2.

Histogram of frequencies

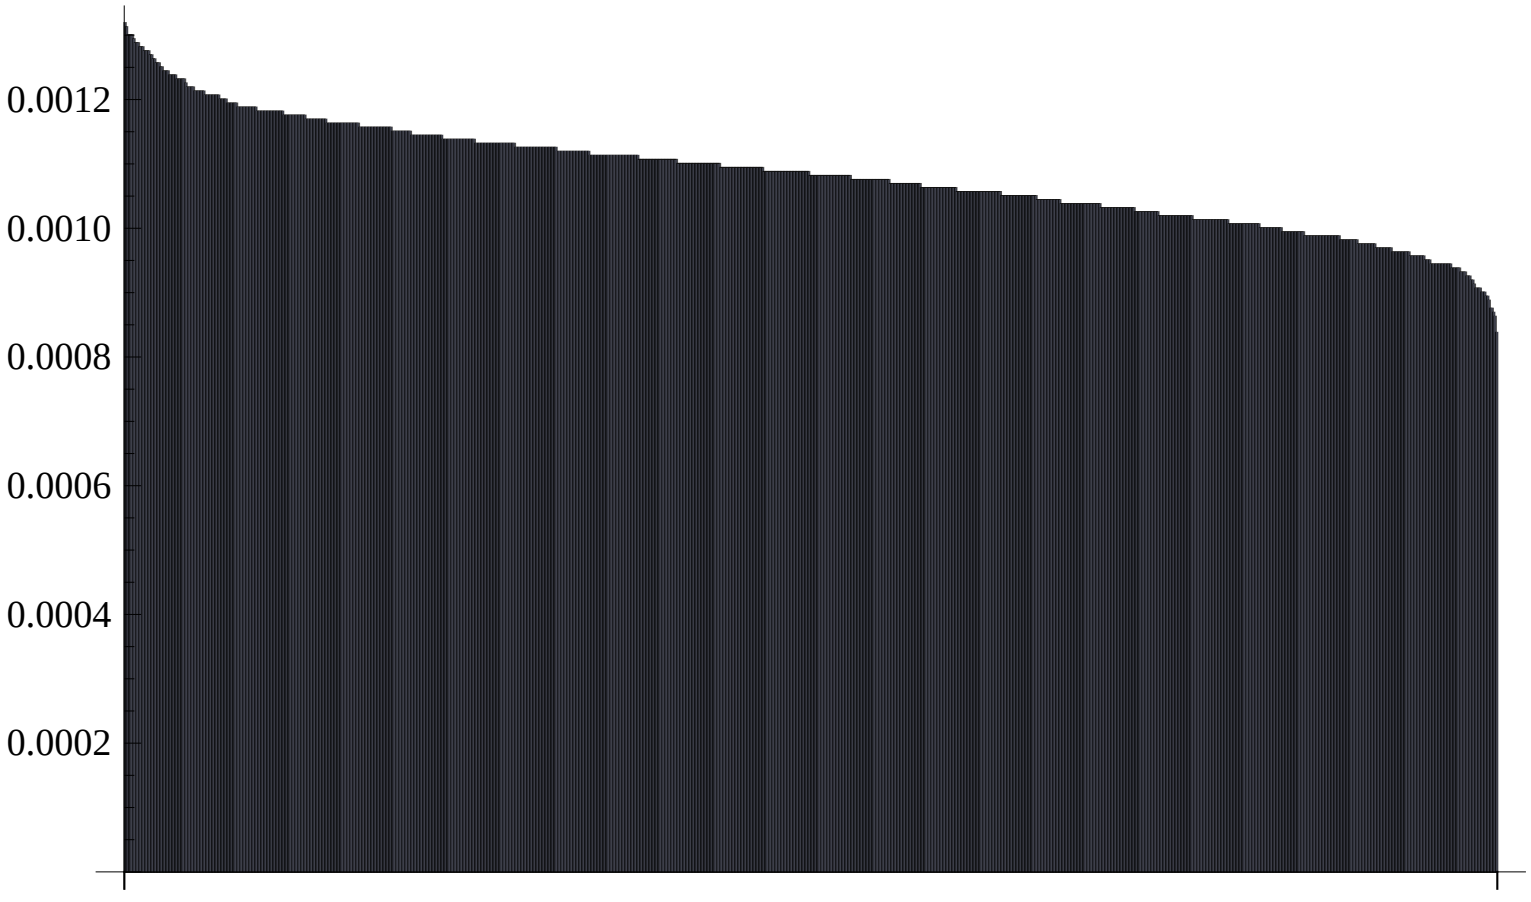

Figure: A histogram of IDs reached during 8 UPDATE iterations for 20000 trials, with $ID_0 = 11111111111$.

Markov Chain

- let us interpret IDs as a binary numbers and denote the ID after t-th Update by ID_t (by ID_0 we denote the starting ID)
- all the binary numbers are a proper subset of $\{0, 2^{2n} - 1\}$
- we represent the Update by transition matrix $U = [u_{i,j}]$ of size $2^{2n} \times 2^{2n}$, with $u_{i,j}$ denoting the probability of evolution from $ID_t = i$ to $ID_{t+1} = j$
- one can easily note that the transition is independent of t, and for each i, j from our set we get:

$$u_{i,j} = \begin{cases} \frac{1}{\binom{2n}{n}} & \text{if } D_H(i,j) = n \ , \\ 0 & \text{oth.} \end{cases}$$

- hence the distribution of IDs after t steps, with the starting distribution v_0, can be calculated as $v_t = v_0 U^t$

Security – idea

- if the probability of reaching a fixed ID is *close enough* independently of the starting ID after T Updates, then an adversary who missed the T Updates cannot determine which tag presents the ID
- since the transition matrix is symmetric, we can switch the conditional probability – we aim to find such T that after T Updates the distribution of the IDs is *close enough* to the uniform distribution
- we analyze a Markov chain ID_t with transition probabilities:

$$\Pr[ID_{t+1} = s' | ID_t = s] = \begin{cases} \frac{1}{\binom{2n}{n}}, & \text{if } D_H(s, s') = n \ , \\ 0, & \text{oth.} \end{cases}$$

Ergodicity

Claim

$\{ID_t\}_{t\geq 0}$ is an ergodic Markov chain convergent to the uniform distribution over $\mathcal{E}$.

- obviously the chain is finite
- we can achieve any state after 2 steps, hence it is irreducible
- since the return to a given state may be obtained e.g., after 2 or 3 steps it is aperiodic
- hence it is ergodic – and it has a stationary distribution
- since $\Pr[ID_{t+1} = s | ID_t = s'] = \Pr[ID_{t+1} = s' | ID_t = s]$, uniform distribution is the only stationary distribution of the process

Coupling

Coupling

Coupling for a Markov Chain M_t is a pair of process (X_t, X_t^*) such that $X_t \sim M_t$ and $X_t^* \sim M_t$, i.e., the mariginal distributions of the process are equal.

Coupling Lemma

Let M_t be an ergodic Markov chain convergent to U. Let (X_t, X_t^*) be any coupling for M_t and X_0 is distributed according to U, with X_0^* distributed in an arbitrary way. Then,

$$\mathrm{TVD}(X_t, X_t^*) \leq \Pr[X_t \neq X_t^*] \ .$$

Hence, if (X_t, X_t^*) is a coupling for M_t such that $\Pr[X_T \neq X_T^*] \leq \varepsilon$ for any $X_0 = s$ i $X_0^* = s^*$, then the mixing time $\tau(\varepsilon) \leq T$.

Mixing time for IDs

Let (ID_t, ID_t^*) be a coupling for the Update procedure, and let ID_t select it's next state first, followed by ID_t^* that aims to reduce the Hamming distance between them (with keeping the mariginal distribution).
Let $Z_t = \{i \in [2n] : ID_{t-1}(i) = ID_{t-1}^*(i)\}$ and ${Z^C}_t = [2n] \setminus Z_t$. Let $K_t = |{Z^C}_t| = 2n - |Z_t|$ denote the number of non-matched bits in round t.
In t-th round, the process ID_t chooses A_t of size n (with appropriate distribution). Let $l_t = |A_t \cap {Z^C}_t|$ denote the number of bits that were not matched between the couple in $t-1$-th round and were selected by the free process in t-th round to be negated. For ID_t^* we define a subset to be negated as $A_t^* = (A_t \cap Z_t) \cup B_t$, where

$$B_t = \begin{cases} B_t' \subset_R^{l_t} {Z^C}_t \setminus A_t & \text{if } l_t \leq \frac{K_t}{2}, \\ B_t' = ({Z^C}_t \setminus A_t) \cup \hat{B} & \text{if } l_t > \frac{K_t}{2}, \end{cases}$$

and $\hat{B} \subset_R^{(2l_t - K_t)} \left(A_t \cap {Z^C}_t\right)$.

Mixing time for IDs

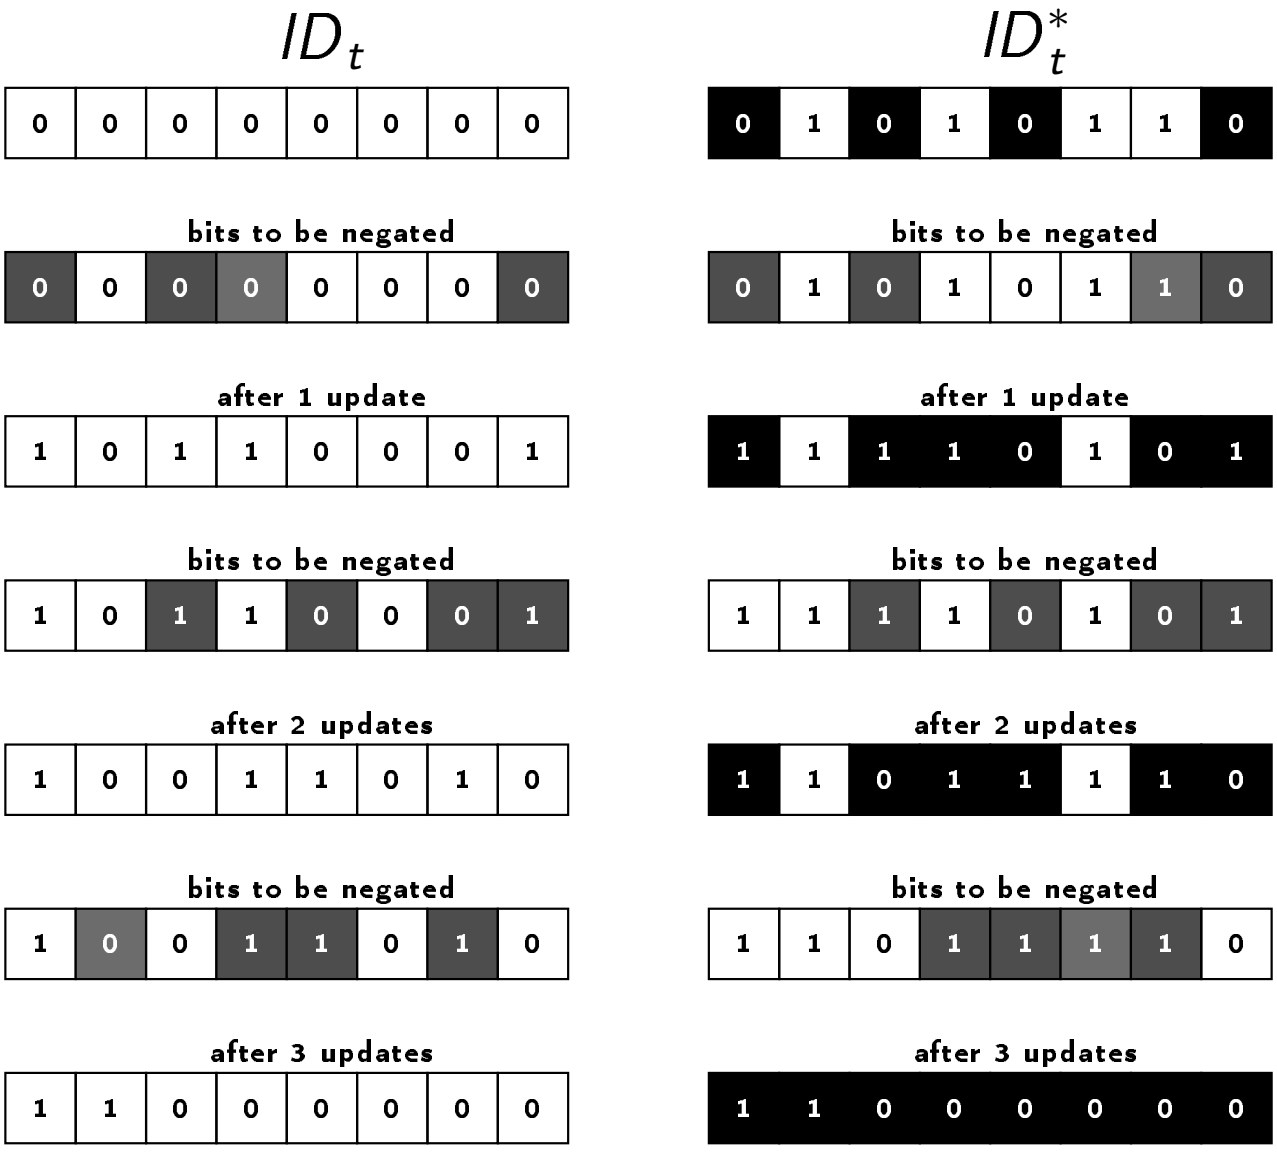

Mixing time for IDs

- let us recall that random variables with hypergeometric distribution $\mathcal{H}(N, n, m)$ describe the number of black balls randomly chosen from a bin containing m black balls and $N - m$ white balls within n selections (without returning the selected ball)
- let us observe that $K_{t+1} \sim 2|\mathcal{H}(2n, n, K_t) - \frac{K_t}{2}|$
- let h denote the numbet of bits chosen during update on which the processes did not match
- if $h < \frac{K_t}{2}$, then $K_{t+1} = 2(\frac{K_t}{2} - h)$, since we can reduce the gap by $2h$, otherwise, each position is negated by at least one process and $2h - K_t = 2(h - \frac{K_t}{2})$ positions are negated by both of them, hence they remain non-matched

Mixing time for IDs

- Since $E[X] = \frac{k}{2}$ and $\sigma^2 = \frac{k(2n-k)}{4(2n-1)} \leq \frac{k}{4}$, directly from Chernoff bound we get that for $X \sim \mathcal{H}(2n, n, k)$ and an even k:

$$\Pr\left[|X - \frac{k}{2}| \geq \frac{k}{4}\right] \leq 2\exp\left(-\frac{k}{8}\right) .$$

- Following the above for $k \geq 11$ we have $2\exp\left(-\frac{k}{8}\right) < \frac{1}{2}$, hence for each $K_t > 11$ we get $\Pr[K_{t+1} \leq \frac{K_t}{2}] > \frac{1}{2}$
- So after manual verification of the cases for $k < 11$ we get that for each $2n \geq K_t \geq 2$ the following is true:

$$\Pr\left[K_{t+1} \leq \frac{K_t}{2}\right] > \frac{1}{2} .$$

- Recalling that K_t is always even, we call a round successful if $K_{t+1} \leq \frac{K_t}{2}$.

Mixing time for IDs

- Since $K_0 \leq 2n$, we need at most $\log n + 1$ successful rounds (that are independent and occurs with probability $p = \frac{1}{2}$).
- Hence, the number of succesful rounds in m subsequent rounds stochastically dominates the distribution $B(m, \frac{1}{2})$.
- Using Chernoff bound we obtain that after $T > 3.6 \log n + 1.6$ updates $ID_T^* = ID_T$ with probability at least $1 - \frac{1}{2n}$.
- Hence, for any $ID_0 \in \mathcal{E}$ and $t \geq 3.6 \log n + 2$ we have $\mathrm{TVD}(ID_t, \mathcal{U}) \leq \frac{1}{2n}$.

Small n

The above proof utilizes Chernoff bound that is not very strict, esp. for small values. As an example, one can show that for $n = 4$, the obtained result in $3.6 \log(4) + 2 = 9.2$, by a simple enumeration one can show more precise bound of $\tau = \frac{32}{17}$.

Ambiguities – modification

Since the database does not know the current ID of the tag, the presented ID may be incorrectly attributed to some different tag (if presented ID belongs to the neighborhood of two different IDs). Since we shown that each ID is connected with rather large subset of the IDs, we slightly modify the protocol:

- We assume that the $|ID| = 8n$, with the same conditions on n and Hamming weight.
- In each update round we still negate only n bits.

The general properties of the protocol remain unchanged, yet the exact values differ.

Ambiguities – modification

Thm.

Let ID_t be tag's ID after t updates and by $\mathcal{U}$ let us denote the uniform distribution over $\mathcal{E}_{8n}$. Then $\mathrm{TVD}(ID_t, \mathcal{U}) \leq \frac{1}{8n}$, for $t \geq 12.44 \log n + \Theta\left(\sqrt{\log n}\right)$ and any starting identifier $ID_0 \in \mathcal{E}_{8n}$.

Hypergeometric distribution

Let $X \sim \mathcal{H}(8n, n, k)$ and k be even. Then

$$\Pr\left[|X - \frac{k}{8}| \geq \frac{k}{16}\right] \leq 2\exp\left(-\frac{k}{75}\right) .$$

Non-matched bits

For each $8n \geq K_t \geq 2$ we have

$$\Pr\left[K_{t+1} \leq \frac{7K_t}{8}\right] > \frac{1}{2} .$$

The probability of ambiguities

Probability bound

Let C_T^N denote that at least 1 ambiguity occurred during T rounds with N tags present. Then

$$\Pr\left[C_T^N\right] \leq N^2 T e^{-\frac{3n}{4}} .$$

The probability of ambiguities

Proof

Let $C_T^{(a,b)}$ denote ambiguity for tags a and b during T rounds. Since ID_t is stationary chain that is uniform in time we obtain

$$\Pr\left[C_T^N\right] \leq T \Pr\left[C_1^{(a,b)}\right] ,$$

for any a, b. Moreover, the update selects n positions at random, hence

$$\Pr\left[C_1^{(a,b)}\right] \leq N^2 \Pr\left[C_1^{(id_0,y)}\right] ,$$

where id_0 is an identifier with $W_H(id_0) = 0$ and a random ID y. Note that, since Update changes n bits, collision is not possible when $W_H(y) > 2n$. Hence,

$$\Pr\left[C_1^{(id_0,y)}\right] \leq \Pr\left[W_H(y) \leq 2n\right] \leq e^{-\frac{3n}{4}} .$$