RAIRO-Theor. Inf. Appl. 44 (2010) 37–58
DOI: 10.1051/ita/2010004

Available online at:
www.rairo-ita.org

ON THE COMPUTATION OF COVERT CHANNEL CAPACITY

Eugene Asarin[1] and Cătălin Dima[2]

Abstract. We address the problem of computing the capacity of a covert channel, modeled as a nondeterministic transducer. We give three possible statements of the notion of "covert channel capacity" and relate the different definitions. We then provide several methods allowing the computation of lower and upper bounds for the capacity of a channel. We show that, in some cases, including the case of input-deterministic channels, the capacity of the channel can be computed exactly (*e.g.* in the form of "the largest root of some polynomial").

Résumé. Nous proposons l'utilisation des transducteurs nondéterministes pour la modélisation des canaux cachés. Trois définitions possibles de la notion de "capacité du canal caché" sont proposées, et leurs relations sont étudiées. Nous donnons ensuite plusieurs méthodes permettant le calcul des bornes inférieures et supérieures pour la capacité d'un canal caché. Nous montrons que, dans certains cas, qui incluent le cas des transducteurs déterministes en entrée, la capacité peut être calculée exactement (sous forme de "la racine la plus grande d'un polynôme").

Mathematics Subject Classification. 94A17, 94A24, 68Q68.

1. Introduction

This paper is concerned with the computation of the capacity of covert channels, a problem that has been addressed for a long time and is of great interest in computer security or in the military domain (*e.g.* "The Orange Book" [5]).

Keywords and phrases. Covert channels, entropy, synchronous transducers.

[1] LIAFA, Université Denis Diderot and CNRS, Case 7014, 75205 Paris Cedex 13, France; asarin@liafa.jussieu.fr

[2] LACL, Université Paris-Est – Université Paris 12, 61 av. du Général de Gaulle, 94010 Créteil, France; dima@univ-paris12.fr

Article published by EDP Sciences

The first one to consider modeling covert channels as finite-state systems and to associate the computation of their capacity to the computation of the entropy of a finite-state system seems to have been [12]. The finite-state model in [12] assumes that one knows the amount of time needed for sending a bit of information, but does not argue on the complexity of the computation of this amount of time.

A different model, which seems to have been influenced by earlier work of Millen, is given in [11]. There, the possibility to transmit bits of information is related with the existence of sets of behaviors which do not satisfy (a variant of) the noninterference property [8]. However [11] concentrates mainly on logical properties modeling noninterference, and not on the computation of the covert channel capacity.

We investigate here a very basic setting of the problem of computing the capacity of a covert channel: in our setting, a covert channel is a system which allows unidirectional communication between two end users, High (the sender) and Low (the receiver). The channel is not controlled by either High nor Low, and, as such, incorporates some non-determinism, which abstracts from the channel activity that the end users are unable to control. Also, the channel itself may not identically transmit High's inputs to Low, but rather translate them into outputs for Low using some translation scheme that is also not controllable by High or Low. On the other hand, High and Low have a full knowledge of the channel, before High spying mission starts. Intuitively, the transmission of a bit succeeds if Low is able to tell apart a partial observation from any other observations, hence deducing that High has chosen a particular sequence of inputs. Also we assume that everything is synchronous, *i.e.* High and Low share a clock and each action from High is immediately followed by an observation by Low.

The attribute "covert" comes from the fact that High utilizes his communication capacities to send some message to Low, by encoding it into some sequence of inputs that is legally accepted as input by the channel. Hence, High and Low would like to agree on a "dictionary" that allows High to send n bits of information, as an input to the channel, such that Low, by observing the channel outputs, be able to distinguish the message High has sent. The main question that we are concerned with here is the asymptotics of the amount of time needed for sending n bits of information, function of n.

No probabilistic situations are modeled, and no feedback from Low to High is present. We do not model either any "game-like" situations, in which the system state might be controlled, *e.g.* in order to limit the bitrate through the channel.

The present paper is a first attempt to study the problem of computing the capacity of a covert channel, under the above simplifying assumptions, and, as such, investigates several possible ways to attack the problem.

We first show that the capacity of a covert channel can be modeled as a generalized form of entropy of a synchronous transducer. We actually give three variants of this definition, according to whether we work with ω-languages, regular ω-languages or with their finite prefixes. We then give some results relating the "prefix" definition of the bitrate with the ω-regular one, and show that the

definition based on general ω languages does not capture the intuition that the reception delay is independent of the number of bits in the message.

We then employ several different methods for under- and over-evaluating the bitrate, all related with spectral radius/entropy computation. An under-evaluation result based on Turán's theorem is given, and another one based on the computation of joint spectral radii.

The idea behind computing lower bounds as joint spectral radii can be summarized as follows: for each $n \in \mathbb{N}$, one constructs a set of matrices with "independent row uncertainties" [3], which represent sets of input words of length n that are translated into distinct output words. Then the joint spectral radius of such a set of matrices is a lower bound for the channel capacity. We also give an "almost monotonicity" result which shows that when passing from n to $n \cdot k$, the computed lower bound is closer to the channel capacity. We don't have yet a result showing that, in the limit, the joint spectral radius method "converges" to the channel capacity.

In the special case of "input-deterministic" channels, which are stronger than channels in which the input automaton is deterministic, we show that the covert channel capacity equals the entropy of the output language. This result is based on classical uniformization results for ω-rational relations. Also an ad-hoc method for computing the capacity in some particular cases is given, based on compositions of transducers.

The paper is organized as follows: the next section introduces our model and gives some basic properties. Section 3 relates the different notions of bitrate and shows that the ω-variant of bitrate has counter-intuitive properties. Section 4 gives a couple of underapproximation results of the bitrate, the first using Turán's Theorem and the second based on the computation of joint spectral radii. An example of computing lower bounds with the JSR method is given in this section. Section 5 gives some special cases in which the bitrate can be exactly computed and represented as the entropy of a regular language. We end with a section with conclusions and directions of future research.

2. Covert channels defined

We model covert channels as synchronous transducers [2]. Hence, High utilizes the input of the transducer to send his messages, and Low is supposed to receive the translations of those messages, as operated by the transducer, and decode them. The receiver is supposed to have complete information about the structure of the transducer, but does not have any supplementary information on the exact state of the transducer, except the information that he can deduce by observing the output history.

First, some notation: for a finite or infinite word w we denote $w[1..n]$ the prefix of length n of w – which is undefined when w has less than n letters. Recall that

the *entropy* of a language of finite words, $L \subseteq \Sigma^*$, or an ω-language $L \subseteq \Sigma^\omega$, is

$$\mathcal{H}(L) = \limsup_{n\to\infty} \frac{1}{n} \cdot \log_2 \operatorname{card}(L[1..n]).$$

Definition 2.1. A **covert channel** is a finite-state *synchronous transducer* $\mathcal{T} = (Q, \Sigma, \Gamma, \delta, q_0)$.

We denote $q \xrightarrow{a/b} r$ the transition in the channel $\mathcal{T}$ (from q to r with input a and output b). Also $\delta^* \subseteq Q \times \Sigma^* \times \Gamma^* \times Q$ denotes the usual reflexive-transitive closure of δ. We also assume that all states are *reachable* in the transducers utilized throughout this paper.

The covert channel $\mathcal{T}$ is supposed to accept finite or infinite words in Σ^∞ and generate, symbol by symbol, their translation in Γ^∞. Formally, the *finite transduction mapping* defined by $\mathcal{T}$ is the mapping $T_\mathcal{T} : Q\times\Sigma^* \to \mathcal{P}(\Gamma^*)$ defined by

$$T_\mathcal{T}(q, w) = \left\{w' \in \Gamma^* \mid \exists q' \in Q \text{ with } q \xrightarrow{w/w'} q' \in \delta^*\right\},$$

and the *infinite transduction mapping* is the mapping $T^\omega_\mathcal{T} : Q \times \Sigma^\omega \to \mathcal{P}(\Gamma^\omega)$ defined by

$$T^\omega_\mathcal{T}(q, w) = \left\{w' \in \Gamma^\omega \mid \forall n \in \mathbb{N}, \exists q' \in Q \text{ with } q \xrightarrow{w[1..n]/w'[1..n]} q' \in \delta^*\right\}.$$

In the sequel we will abuse notation and employ $\mathcal{T}$ as the mapping $T_\mathcal{T}$, and also $\mathcal{T}^\omega$ instead of $T^\omega_\mathcal{T}$.

The *underlying input automaton* of $\mathcal{T}$ is the automaton $In_\mathcal{T} = (Q, \Sigma, \delta_{in}, q_0)$, with $\delta_{in} = \delta\big|_{Q\times\Sigma\times Q}$. The *input language* for $\mathcal{T}$ is the language of $In_\mathcal{T}$, seen as a finite automaton working on finite words and having all states accepting. This language will be sometimes denoted as $\mathcal{L}_{in}(\mathcal{T}) = \mathcal{L}(In_\mathcal{T})$. On the other hand, the *output language* of $\mathcal{T}$ is $\mathcal{L}_{out}(\mathcal{T}) = \bigcup\left\{\mathcal{T}(w) \mid w \in \mathcal{L}_{in}(\mathcal{T})\right\}$. We may also define similarly the *input ω-language* and *output ω-language* for $\mathcal{T}$, denoted respectively $\mathcal{L}^\omega_{in}(\mathcal{T})$, and $\mathcal{L}^\omega_{out}(\mathcal{T})$.

The following definition is inspired from the nondeducibility on strategies of [16] (see also [11]), and models Low's observability of High activity:

Definition 2.2. Given a covert channel $\mathcal{T}$, a finite language $W \subseteq \Sigma^*$ is called **$\mathcal{T}$-distinguishable** if the following properties hold:

(1) All words in W are accepted by $\mathcal{T}$ as inputs and have the same length: $W \subseteq \mathcal{L}_{in}(\mathcal{T})[1..n]$ for some $n \in \mathbb{N}$.
(2) For any two words $w, w' \in W$, $\mathcal{T}(q_0, w) \cap \mathcal{T}(q_0, w') = \emptyset$.

An infinite language $W \subseteq \Sigma^\omega$ is called **$\mathcal{T}$-distinguishable** if $W \subseteq \mathcal{L}^\omega_{in}(\mathcal{T})$ and

$$\forall w, w' \in W, \mathcal{T}^\omega(q_0, w) \cap \mathcal{T}^\omega(q_0, w') = \emptyset.$$

Definition 2.3. The **bitrate** (or **capacity**) of $\mathcal{T}$ is:

$$\mathcal{B}(\mathcal{T}) = \limsup_{n\to\infty} \frac{1}{n} \cdot \log_2 \max\big\{ \operatorname{card}(W) \mid W \subseteq \Sigma^n, W \text{ is } \mathcal{T}\text{-distinguishable}\big\}.$$

We say that a language $W \subseteq \Sigma^*$ **realizes** the bitrate of $\mathcal{T}$ if there exists a strictly increasing sequence of integers $(k_n)_{n\in\mathbb{N}}$ satisfying the following properties:

(1) the language $W[1..k_n]$ is $\mathcal{T}$-distinguishable for all $n \in \mathbb{N}$; and
(2) $\mathcal{B}(\mathcal{T}) = \lim_{n\to\infty} \frac{1}{k_n} \cdot \log_2 \operatorname{card}(W[1..k_n])$.

The idea is that, if a channel has a bitrate of α, then, using some $\mathcal{T}$-distinguishable language $W \subseteq \Sigma^n$ as a "dictionary", High can transmit to Low one from $\approx 2^{\alpha n}$ messages, that is $\approx \alpha n$ bits of information.

The above definition has also the following "omega" and "omega-regular" variants:

Definition 2.4. The **ω-bitrate** (or **ω-capacity**) of $\mathcal{T}$ is:

$$\mathcal{B}^\omega(\mathcal{T}) = \sup\big\{\mathcal{H}(W) \mid W \subseteq \Sigma^\omega, W \text{ is } \mathcal{T}\text{-distinguishable}\big\}.$$

The **ω-regular bitrate** (or **ω-regular capacity**) of $\mathcal{T}$ is:

$$\mathcal{B}^\omega_r(\mathcal{T}) = \sup\big\{\mathcal{H}(W) \mid W \subseteq \Sigma^\omega, W \text{ is } \mathcal{T}\text{-distinguishable and } \omega\text{-regular}\big\}.$$

An ω-language $W \subseteq \Sigma^\omega$ **realizes** the ω-bitrate (ω-regular bitrate) of $\mathcal{T}$ if $\mathcal{H}(W) = \mathcal{B}^\omega(\mathcal{T})$ (resp. $\mathcal{H}(W) = \mathcal{B}^\omega_r(\mathcal{T})$).

The first problem which is addressed in this paper is the following:

Problem 2.5. Given a covert channel $\mathcal{T}$, compute the (simple, ω or ω-regular) bitrate of $\mathcal{T}$.

The term "compute" here can be understood as giving a representation for the real number $\mathcal{B}(\mathcal{T})$ (or $\mathcal{B}^\omega(\mathcal{T})$, or $\mathcal{B}^\omega_r(\mathcal{T})$, or their over- and under-approximations), *e.g.* in the form of "the largest root of some polynomial". A more precise exploration of computability of these reals in terms of recursive analysis will be a subject of future work.

A related problem is the following:

Problem 2.6. Given a covert channel $\mathcal{T}$, construct a (regular) realization W of $\mathcal{B}(\mathcal{T})$, and/or of $\mathcal{B}_r(\mathcal{T})$, and/or of $\mathcal{B}^\omega_r(\mathcal{T})$.

Solving the second problem would help High and Low in their choice of the largest dictionary: for any $n \in \mathbb{N}$, High can utilize words of $W[1..n]$ for transmitting $\approx \alpha n$ bits of information – where $\alpha = \mathcal{B}(\mathcal{T})$ or $\alpha = \mathcal{B}^\omega(\mathcal{T})$ or $\alpha = \mathcal{B}^\omega_r(\mathcal{T})$.

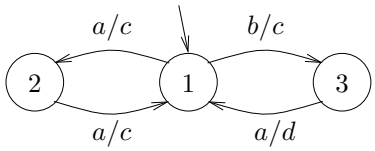

FIGURE 1. Channel $\mathcal{T}_1$.

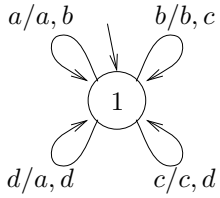

FIGURE 2. Channel $\mathcal{T}_2$.

2.1. Examples

The first example is from Figure 1. Here, both the input language $((a+b)a)^*$ and the input ω-language $((a+b)a)^\omega$ are $\mathcal{T}$-distinguishable, which means that

$$\mathcal{B}(\mathcal{T}_1) = \mathcal{B}^\omega(\mathcal{T}_1) = \mathcal{B}^\omega_r(\mathcal{T}_1) = \mathcal{H}(\mathcal{L}_{in}(\mathcal{T}_1)) = \log_2 \sqrt{2} = \frac{1}{2}.$$

The second example is from Figure 2. Note that in this four-leaf clover, the adjacent leaves induce "ambiguities" when translated. Intuitively, only the opposite leaves can be used for coding. This would mean that the bitrate of this channel is $\log_2 2 = 1$, and is realized by the complete automaton with one state and two loops. This will be proved in a subsequent section.

3. Basic properties of bitrates

We start with a couple of results concerning distinguishability. The first straightforward property that can be used for a rough over/underapproximation of the bitrate of a covert channel is the following:

Proposition 3.1. *For any covert channel $\mathcal{T}$ we have that $\mathcal{B}(\mathcal{T}) \le \mathcal{H}(\mathcal{L}_{out}(\mathcal{T})) \le \mathcal{H}(\mathcal{L}_{in}(\mathcal{T}))$.*

On the other side, for any $L \subseteq \Sigma^*$ having the property that $L[1..n]$ is $\mathcal{T}$-distinguishable for all $n \in \mathbb{N}$ (we also call such languages $\mathcal{T}$-distinguishable), we have that $\mathcal{B}(\mathcal{T}) \ge \mathcal{H}(L)$.

The upper bound above can be used directly. In order to use the lower bound, one can guess some distinguishable set W and utilize it for information transmission through the covert channel. The remaining problem of checking whether a

given (regular) language is distinguishable is addressed in the following proposition:

Proposition 3.2. *The problem of checking whether a given regular (ω-regular) language W is $\mathcal{T}$-distinguishable in a given covert channel $\mathcal{T}$ is decidable.*

Proof. The proof follows by a straightforward adaptation of the intersection construction: we construct a *synchronous composition* between two copies of $\mathcal{T}$ and two copies of the automaton which accepts W, denote it $\mathcal{A}$. The synchronous composition works on $\Sigma \times \Sigma$, and accepts exactly the set of pairs of non-distinguishable words (w, w'). Then W is distinguishable if and only if this synchronous composition has an empty language.

For the case when W is an ω-regular language, assume that the covert channel is $\mathcal{T} = (Q, \Sigma, \Gamma, \delta_{\mathcal{T}}, q_0)$, and the regular language W is accepted by the Büchi automaton $\mathcal{A} = (S, \Sigma, \delta_{\mathcal{A}}, s_0, F)$. The synchronous composition utilizes 7-tuples $(q_1, q_2, s_1, s_2, i_1, i_2, i_3)$ consisting of two states of $\mathcal{T}$, two states of $\mathcal{A}$ and three bits needed for bookkeeping: i_1 is used to record passages through the set of final states in the first copy of $\mathcal{A}$, i_2 records passages through the set of final states in the second copy of $\mathcal{A}$, while i_3 is needed for signaling that a pair of non-distinguishable *and distinct* words has been constructed so far. Note that the use of the indices i_1 and i_2 is reminiscent from the intersection construction for Büchi automata.

Formally, the synchronous composition is the Büchi automaton $\mathsf{Sync}(\mathcal{T}, \mathcal{A}) = (Q^2 \times S^2 \times \{0,1\}^3, \Sigma^2, \theta, \overline{s}_0, \overline{F})$ where

- $\overline{s}_0 = \big(q_0, q_0, s_0, s_0, 0, 0, 0\big)$;
- $\overline{F} = \big\{(q_1, q_2, s_1, s_2, 1, 1, 1) \mid q_1, q_2 \in Q, s_1, s_2 \in S\big\}$;
- θ is composed of tuples $(q_1, q_2, s_1, s_2, i_1, i_2, i_3) \xrightarrow{a,b} (r_1, r_2, t_1, t_2, j_1, j_2, j_3)$, where
  (1) $q_1 \xrightarrow{a/c} r_1, q_2, \xrightarrow{b/c} r_2 \in \delta_{\mathcal{T}}$, and $s_1 \xrightarrow{a} t_1, s_2 \xrightarrow{b} t_2$.
  (2) $j_1 = \begin{cases} 1 \text{ if } i_1 = 1 \text{ and } i_2 = 0, \text{ or } i_1 = 0 \text{ and } s_1 \in F, \\ 0 \text{ if } i_1 = i_2 = 1 \text{ or } i_1 = 0 \text{ and } s_1 \notin F \end{cases}$
  (3) $j_2 = \begin{cases} 1 \text{ if } i_2 = 1 \text{ and } i_1 = 0, \text{ or } i_2 = 0 \text{ and } s_2 \in F, \\ 0 \text{ if } i_1 = i_2 = 1 \text{ or } i_2 = 0 \text{ and } s_2 \notin F \end{cases}$
  (4) $j_3 = \begin{cases} 1 & \text{if } a \neq b \\ i_3 & \text{otherwise.} \end{cases}$

It is easy to see that

$$L^{\omega}(\mathsf{Sync}(\mathcal{T}, \mathcal{A})) = \big\{(w_1, w_2) \mid w_1, w_2 \in L^{\omega}(\mathcal{A}), w_1, w_2 \text{ are not } \mathcal{T}\text{-distinguishable}\big\}$$

A similar construction works for regular languages $W \subseteq \Sigma^*$. □

The simple and the ω-regular bitrates are related by the following result:

Proposition 3.3. *For any covert channel $\mathcal{T}$, $\mathcal{B}_r^{\omega}(\mathcal{T}) \leq \mathcal{B}(\mathcal{T})$.*

Proof. Note first that by definition of the ω-regular bitrate it is sufficient to prove that $\mathcal{H}(R) \leq \mathcal{B}(\mathcal{T})$ for any ω-regular R that is $\mathcal{T}$-distinguishable.

Unfortunately, there are $\mathcal{T}$-distinguishable ω-regular languages $R \subseteq \Sigma^\omega$ for which the n-prefix languages $R[1..n]$ might not all be $\mathcal{T}$-distinguishable. We will then need to prove that there exists some fixed integer N_R such that each prefix $R[1..n]$ can be extended to a distinguishable $R'_n \subseteq R[1..n+N_R]$, *i.e.* $R[1..n] \prec R'_n$, such that $\mathrm{card}(R'_n) = \mathrm{card}(R[1..n])$.

So take a Muller deterministic automaton $\mathcal{A} = (S, \Sigma, \theta, s_0, \mathcal{F})$, with $\mathcal{F} \subseteq 2^S$, accepting some $\mathcal{T}$-distinguishable ω-regular language. We construct first, from $\mathcal{A}$ and $\mathcal{T}$, the transducer whose transduction mapping is $\mathcal{T}\big|_{\mathcal{L}(\mathcal{A})}$, be it $\mathcal{T} \times \mathcal{A} = \big(Q \times S, \Sigma, \Gamma, \delta_\theta, (q_0, s_0), \mathcal{F}'\big)$ with

$$\mathcal{F}' = \big\{G \subseteq 2^{Q\times S} \mid \pi_S(G) \in \mathcal{F}\big\};$$
$$\delta_\theta = \big\{(q,s) \xrightarrow{a/b} (q',s') \mid q \xrightarrow{a/b} q' \in \delta, s \xrightarrow{a} s' \in \theta\big\}.$$

Denote then $Q_{\mathcal{T},\mathcal{A}}$ the set of states in $Q \times S$ that are accessible from (q_0, s_0) and *co-accessible* in $\mathcal{T} \times \mathcal{A}$, that is, for which there exists a path leading to some connected component in $\mathcal{F}'$.

In $\mathcal{T}_\mathcal{A}$, for each state $(q, s) \in Q \times S$ that is co-accessible we pick one of the connected components in $\mathcal{F}'$ that can be reached from (q, s) and label it as $F_{(q,s)}$. Define $w_{(q,s)}$ as the label of the shortest path that starts in (q, s) and reaches the first time one of the states in $F_{(q,s)}$.

We will then need to extend each $w_{(q,s)}$ into a word $w^1_{(q,s)}$ such that all such words have equal length. This can be easily done, since from each state of $F_{(q,s)}$ there exists an infinite path that remains in $F_{(q,s)}$.

With this assumption, define

$$N_1 = |w^1_{(q,s)}|,$$

where $|w|$ denotes the length of the word w. Note that N_1 is independent of the choice of (q, s), as all words $w^1_{(q,s)}$ have equal length.

Then, for each (q, s), choose $(q', s') \in F_{(q,s)}$ such that there exists a path from (q, s) to (q', s') whose input label is $w_{(q,s)}$. Furthermore, take a circuit in (q', s') that visits all the states of $F_{(q,s)}$, and denote its label as $w'_{(q,s)}$. Put

$$N_2 = \mathrm{lcm}\left\{|w'_{(q,s)}|\right\}$$

and

$$w^2_{(q,s)} = \big(w'_{(q,s)}\big)^{\frac{N_2}{|w'_{(q,s)}|}}.$$

Claim 3.4. In the above setting, given $w_1, w_2 \in \Sigma^*$ for which $(q_0, s_0) \xrightarrow{w_1} (q_1, s_1)$ and $(q_0, s_0) \xrightarrow{w_2} (q_2, s_2)$, we must have that

$$\mathcal{T}\big(q_0, w_1 w^1_{(q_1,r_1)} w^2_{(q_1,s_1)}\big) \cap \mathcal{T}\big(q_0, w_2 w^1_{(q_2,r_2)} w^2_{(q_2,s_2)}\big) = \emptyset.$$

For, if the claim does not hold, we would get a contradiction with the fact that R is $\mathcal{T}$-distinguishable: we would have the following sequences of extended transitions:

$$(q_0, s_0) \xrightarrow{w_1/z} (q_1, s_1) \xrightarrow{w^1_{(q_1,s_1)}/z'} (q'_1, s'_1) \xrightarrow{w^2_{(q_1,s_1)}/z''} (q'_1, s'_1); \tag{1}$$

$$(q_0, s_0) \xrightarrow{w_2/z} (q_2, s_2) \xrightarrow{w^1_{(q_2,s_2)}/z'} (q'_2, s'_2) \xrightarrow{w^2_{(q_2,s_2)}/z''} (q'_2, s'_2). \tag{2}$$

Note that the transductions in Γ^* of both input words are identical, the Γ-word $zz'z''$.

Since the last extended transition in (1) is a circuit, it can be iterated forever, and similarly for (2), which actually implies that

$$zz'(z'')^\omega \in \mathcal{T}\big(q_0, w_1 w^1_{(q_1,r_1)} (w^2_{(q_1,r_1)})^\omega\big) \cap \mathcal{T}\big(q_0, w_2 w^1_{(q_2,r_2)} (w^2_{(q_2,r_2)})^\omega\big).$$

But $w_1 w^1_{(q_1,r_1)} (w^2_{(q_1,r_1)})^\omega$ has as its set of repeatedly visited states exactly $F_{(q,s)}$. Hence we would have two words in R which cannot be distinguished by $\mathcal{T}$, contradiction.

Following the above claim, we can put the sum $N_R = N_1 + N_2$ as the searched-for bound. We might then construct the following extension of $\mathcal{L}(\mathcal{A})[1..n]$:

$$R_n = \big\{ w w^1_{(q,s)} w^2_{(q,s)} \mid (q,s) \text{ is the first state, in lexicographic order, for which } (q_0, s_0) \xrightarrow{w} (q, s) \text{ in } \mathcal{T} \times \mathcal{A} \big\}.$$

It then follows that $\operatorname{card}(R_n) = \operatorname{card}(\mathcal{L}(\mathcal{A})[1..n])$, $R_n \subseteq \Sigma^{n+N_R}$ and R_n is distinguishable. Therefore, for any automaton $\mathcal{A}$ accepting a $\mathcal{T}$-distinguishable language,

$$\mathcal{B}(\mathcal{T}) \geq \limsup_{n\to\infty} \frac{1}{n} \log_2 \operatorname{card} R_n = \limsup_{n\to\infty} \frac{1}{n} \log_2 \operatorname{card}\big(\mathcal{L}(\mathcal{A})[1..n]\big).$$

In conclusion, $\mathcal{B}^\omega_r(\mathcal{T}) \leq \mathcal{B}(\mathcal{T})$. □

3.1. ω-BITRATE IS TOO WEAK

In this subsection we prove that the ω-bitrate, in spite of having a nicer mathematical definition than the (simple) bitrate, cannot be really utilized as a measure of the amount of bits that the sender may transmit through the covert channel.

Consider the covert channel in Figure 3, and the following ω language (which is not ω-regular!)

$$L_\omega = \big\{ z a^{\overline{z}_2 + 1} b^\omega \mid z \in (0+1)^* \big\},$$

where $\overline{z}_2$ is the value of z as an integer written in base 2.

It is easy to see that L_ω is $\mathcal{T}$-distinguishable, since for each ω-word $w \in L$ there exists a unique integer $z \in \mathbb{N}$ such that $\mathcal{T}(w) = a^{\lfloor \log_2 z + 1 \rfloor} b^z a^\omega$. Hence,

$$\mathcal{B}^\omega(\mathcal{T}) \geq \mathcal{H}(L_\omega) = 1.$$

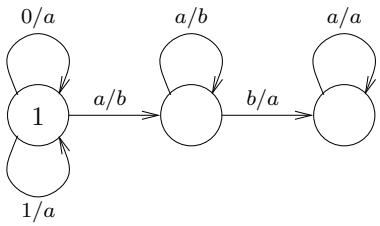

FIGURE 3. A covert channel whose ω-bitrate is 1.

Intuitively, to use $\mathcal{B}^\omega$ for receiving n bits of information, the receiver must wait for up to 2^n units of time to distinguish which bits have been sent.

This is counter-intuitive, since we want any message of length n to be distinguished by Low from other messages of length n after a delay which is small compared to n. As a consequence, in the sequel, we concentrate on the study of the "finitary" bitrate.

On the other hand, note that, by Proposition 3.1,

$$\mathcal{B}(\mathcal{T}) \leq \mathcal{H}\big(\mathcal{L}_{out}(\mathcal{T})\big) = 0,$$

which is consistent with our intuition that the channel in Figure 3 cannot be used to transmit one bit per time unit.

4. Approximation results

This section gives a couple of techniques that can be utilized for obtaining lower bounds for the bitrate.

4.1. Underapproximating the bitrate using Turán's theorem

In the sequel, the covert channel $\mathcal{T}$ is fixed, and we denote $L^n_{\mathcal{T}} = \mathcal{L}_{in}(\mathcal{T}) \cap \Sigma^n$. Consider the following *undirected graph* $G_n = (L^n_{\mathcal{T}}, E^n_{\mathcal{T}})$, where

$$E^n_{\mathcal{T}} = \big\{\{w, w'\} \subseteq L^n_{\mathcal{T}} \mid \mathcal{T}(w) \cap \mathcal{T}(w') = \emptyset\big\},$$

and denote $l = \mathrm{card}(L^n_{\mathcal{T}})$ and $e = \mathrm{card}(E^n_{\mathcal{T}})$.

Remark 4.1. Note that a $\mathcal{T}$-distinguishable set $W \subseteq L^n_{\mathcal{T}}$ forms a clique in G_n.

We will therefore use Turán's theorem [15] to underapproximate $\mathcal{B}(\mathcal{T})$ using the quantities l and e.

Theorem 4.2 Turán**.** *Assume that there is no clique of size larger than r in a graph $G = (V, E)$. Then*

$$\mathrm{card}(E) \leq \Big(1 - \frac{1}{r}\Big) \cdot \frac{\mathrm{card}(V)^2}{2}.$$

This means that, if W is a $\mathcal{T}$-distinguishable language of the largest size in Σ^n, then

$$\operatorname{card}(W) \geq \frac{l^2}{l^2 - 2e}. \tag{3}$$

We would like to use the inequality (3) to give an underapproximation of $\mathcal{B}(\mathcal{T})$. To this end we need to evaluate, function of n, the asymptotics of l and e function of n.

Clearly, l grows exponentially at the rate $2^{n\mathcal{H}(\mathcal{L}_{in}(\mathcal{T}))}$.

The asymptotics of e, function of n, can be derived from the entropy of a finite automaton which encodes pairs of words from Σ^* whose translation can be "ambiguous". Denote δ_{in} the input transition relation for $\mathcal{T}$, that is, $q \xrightarrow{a} r \in \delta_{in}$ if $q \xrightarrow{a/b} r \in \delta$ for some $b \in \Gamma$.

The automaton for the computation of e is $\mathsf{Pairs}_{\mathcal{T}} = (2^{Q\times Q}, \Sigma \times \Sigma, \theta, (q_0, q_0))$ with all states final and

$$\theta = \left\{\mathcal{R} \xrightarrow{a,b} \delta^2(\mathcal{R}) \mid \mathcal{R} \subseteq Q \times Q, \mathcal{R}, \delta^2(\mathcal{R}) \neq \emptyset\right\}$$

where the notation $\delta^2(\mathcal{R})$ stands for the application $\delta^2 : 2^{Q\times Q} \to 2^{Q\times Q}$ defined as:

$$\delta^2(\mathcal{R}) = \left\{(r_1, r_2) \in Q \times Q \mid \exists (q_1, q_2) \in \mathcal{R}_1, \exists x \in \Gamma \text{ with } q_1 \xrightarrow{a/x} r_1, q_2 \xrightarrow{a/x} r_2\right\}\Big\}.$$

Note that for each reachable state $\mathcal{R} \subseteq Q \times Q$ in $\mathsf{Pairs}(\mathcal{T})$ each pair $(q, q') \in \mathcal{R}$ encodes the existence of a pair of non-distinguishable words $w, w' \in \Sigma^*$ (with $|w| = |w'|$).

Remark 4.3. If we denote $P_n = \mathcal{L}(\mathsf{Pairs}_{\mathcal{T}}) \cap (\Sigma \times \Sigma)^n$, then the number of edges of G_n is $e = \frac{1}{2}\big(l^2 - \operatorname{card}(P_n)\big)$.

As a consequence, the asymptotics of e function of n can be computed by computing the entropy of the input automaton $In_{\mathcal{T}}$ and of the reachable part of $\mathsf{Pairs}(\mathcal{T})$.

Proposition 4.4. *Suppose that*

$$\limsup_{n\to\infty} \frac{\log_2 l}{n} = \alpha \text{ and } \limsup_{n\to\infty} \frac{\log_2 \operatorname{card}(P_n)}{n} = \beta.$$

Then an underapproximation for the bitrate of $\mathcal{T}$ is

$$\mathcal{B}(\mathcal{T}) \geq 2\alpha - \beta.$$

Proof. By straightforward calculations, and utilizing inequality (3) and the previous remark:

$$\begin{aligned}\mathcal{B}(\mathcal{T}) &\geq \limsup_{n\to\infty} \frac{1}{n} \log_2 \frac{l^2}{l^2 - 2e} = \limsup_{n\to\infty} \frac{1}{n} \log_2 \frac{l^2}{\mathrm{card}(P_n)} \\ &\geq \limsup_{n\to\infty} \frac{\log_2 l^2}{n} - \limsup_{n\to\infty} \frac{\log_2 \mathrm{card}(P_n)}{n} \\ &= 2\alpha - \beta. \end{aligned}$$

□

We may apply this proposition for the covert channel in Figure 2 as follows: first, note that

$$\limsup_{n\to\infty} \frac{\log_2 l}{n} = \mathcal{H}(\mathcal{L}_{in}(\mathcal{T})) = 2.$$

On the other hand, the entropy of $\mathsf{Pairs}(\mathcal{T}_2)$ is $\log_2 12$, as $\mathsf{Pairs}(\mathcal{T}_2)$ contains a single state with 12 loops. Therefore, $\mathcal{B}(\mathcal{T}_2) \geq 2 - \log_2 3$.

4.2. Underapproximating bitrates with joint spectral radii

In this section we suggest an alternative approach for searching underapproximations of the bitrates, in the case when the underlying input automaton for the given channel is deterministic. This approach involves the computation of the joint spectral radius of a particular set of matrices. Recall [3] that, for a given set of matrices $\mathcal{M}$, the *joint spectral radius* $\rho(\mathcal{M})$ is the following quantity:

$$\rho(\mathcal{M}) = \limsup_{k\to\infty} \max\left\{\|A_1 \cdot \ldots A_k\|^{1/k} \mid k \in \mathbb{N}, A_1, \ldots, A_k \in \mathcal{M}\right\},$$

where $\|\cdot\|$ is some matrix norm.

Sets of matrices which will be utilized in our approximation process are sets $\mathcal{M}$ bearing the property that, for each index i, there exists a set of rows $R_i \subseteq \mathbb{R}^n_{\geq 0}$ such that $\mathcal{M}$ consists of all the matrices whose i-th row belongs to R_i:

$$\mathcal{M} = \left\{(v_1^T, \ldots, v_n^T)^T \mid v_i \in R_i\right\}. \tag{4}$$

Following [3] we say that such sets of matrices have *independent row uncertainties* and refer them to as **IRU sets**.

In [3], a nice result on the *joint spectral radius* of IRU sets of matrices has been obtained:

Theorem 4.5 [3]**.** *If the IRU set $\mathcal{M}$ is composed of nonnegative matrices, then*

$$\rho(\mathcal{M}) = \max\{\rho(A) | A \in \mathcal{M}\}.$$

Moreover, $\rho(\mathcal{M})$ can be computed using a polynomial algorithm based on convex optimization techniques.

We may utilize this result for computing an underapproximation of the bitrate of a covert channel as follows: consider a channel $\mathcal{T} = (Q, \Sigma, \Gamma, \delta, q_0)$ with $Q = \{1, 2, \ldots, n\}$ and suppose that the input automaton for $\mathcal{T}$ is *deterministic*[1], that is, if $i \xrightarrow{a/b} j, i \xrightarrow{a/c} j' \in \delta$ then $j = j'$.

For an automaton $\mathcal{A}$, we denote $q \xrightarrow{w}_{\mathcal{A}} q'$ if there exists a run ρ in $\mathcal{A}$ which starts in q, ends in q' and is labeled with w. Also, given a set of words $W \subseteq \Sigma^k$ and a state $q \in Q$, we say that W is *$\mathcal{T}$-distinguishable in q* if $\mathcal{T}(q, w) \cap \mathcal{T}(q, w') = \emptyset$.

Suppose that, for each state $i \in Q$, we have a fixed set $W_i \subseteq \Sigma^k$ which is $\mathcal{T}$-distinguishable in i. We will then build a matrix of sets of words, $M = M(W_1, \ldots, W_n)$, with

$$M_{ij} = \left\{ w \in W_i \mid i \xrightarrow{w} j \in \delta_{in} \right\} \tag{5}$$

and, for each such matrix, associate the integer matrix $V = V(M)$ with

$$V(M)_{ij} = \operatorname{card}(M_{ij}). \tag{6}$$

Note first that the following set of matrices is IRU:

$$\mathcal{M}_k = \left\{ V(M(W_1, \ldots, W_n)) \mid W_i \subseteq \Sigma^k, W_i \text{ is } \mathcal{T}\text{-distinguishable in } i, 1 \leq i \leq n \right\}.$$

Proposition 4.6. *For any $k \in \mathbb{N}$, $k \geq 1$, the following lower bound for the bitrate holds:*

$$\frac{1}{k} \log_2 \rho(\mathcal{M}_k) \leq \mathcal{B}(\mathcal{T}).$$

This bound is almost monotonous in k in the following sense: for any $k, l \in \mathbb{N}$, $k, l \geq 1$,

$$\frac{1}{k} \log_2 \rho(\mathcal{M}_k) \leq \frac{1}{kl} \log_2 \rho(\mathcal{M}_{kl}). \tag{7}$$

Proof. Let us establish first that, for any sequence of matrices $M_1, \ldots, M_m \in \mathcal{M}_k$, the set $\bigcup_{j=1}^{n} (M_1 \cdot \ldots \cdot M_m)_{ij}$ is $\mathcal{T}$-distinguishable at i. This follows by induction on m: for $m = 1$ the property is satisfied by hypothesis on the construction of $\mathcal{M}_k$.

For the induction step, take two words $w, w' \in \bigcup_{j=1}^{n} (M_1 \cdot \ldots \cdot M_{m+1})_{ij}$, and consider their decomposition $w = w_1 \cdot w_2$, $w' = w_1' \cdot w_2'$ with $|w_1| = |w_1'| = km$ and $|w_2| = |w_2'| = k$.

Consider further the two states $j, j' \in Q$ with $i \xrightarrow{w_1} j, i \xrightarrow{w_1'} j' \in \delta_{in}$. Note that, by the assumption that $In_{\mathcal{T}}$ is deterministic, these two states are uniquely determined.

If $j \neq j'$, then, by induction, we must have that w_1 and w_1' are $\mathcal{T}$-distinguishable at i, which clearly implies that w and w' are $\mathcal{T}$-distinguishable at i.

[1]This requirement is weaker than input-determinism.

If $j = j'$, then, also by induction, we must have that $w_1 = w'_1$. But then w_2, w'_2 both belong to some set W_j which is $\mathcal{T}$-distinguishable at j by construction, and therefore:

$$\begin{aligned}\mathcal{T}(i, w_1 w_2) \cap \mathcal{T}(i, w_1 w'_2) &= \mathcal{T}(i, w_1) \cdot \mathcal{T}(j, w_2) \cap \mathcal{T}(i, w_1) \cdot \mathcal{T}(j, w'_2) \\ &= \mathcal{T}(i, w_1) \cdot \big(\mathcal{T}(j, w_2) \cap \mathcal{T}(j, w'_2)\big) \\ &= \emptyset,\end{aligned}$$

which proves the induction step.

Recall that each state in $\mathcal{T}$ is reachable, hence there exists a word w_0 (leading from q_0 to i) of length $n_0 \le \mathrm{card}(Q)$ for which the following word

$$w_0 \cdot \bigcup\nolimits_{j=1}^{n} (M_1 \cdot \ldots \cdot M_m)_{ij}$$

is $\mathcal{T}$-distinguishable (at q_0). Therefore,

$$\sum_{j=1}^{n} \big(V(M_1) \cdot \ldots \cdot V(M_m)\big)_{ij} \le \max\big\{\, \mathrm{card}(W) \mid W \subseteq \Sigma^{km+n_0} \text{ is } \mathcal{T}\text{-distinguishable}\big\},$$

which implies that

$$\begin{aligned}&\frac{1}{km} \log_2 \|V(M_1) \cdot \ldots \cdot V(M_m)\|_\infty \\ &\qquad \le \frac{1}{km} \log_2 \max\big\{\, \mathrm{card}(W) \mid W \subseteq \Sigma^{km+n_0}, W \text{ is } \mathcal{T}\text{-distinguishable}\big\}\end{aligned}$$

for any $M_1, \ldots, M_m \in \mathcal{M}_k$. (Here $\|\|_\infty$ is the ∞-norm on matrices.)

For $m \to \infty$, this gives

$$\frac{1}{k} \log_2 \rho(\mathcal{M}_k) \le \mathcal{B}(\mathcal{T}).$$

To prove the "monotonicity" inequality, note that, by our induction proof, we also have that $\big(M_1 \cdot \ldots \cdot M_m\big)_{ij}$ is some $\mathcal{T}$-distinguishable set at i consisting of words of length kl, and therefore there exists some matrix $M' \in \mathcal{M}_{kl}$ such that

$$\bigcup_{j=1}^{n} \big(M_1 \cdot \ldots \cdot M_m\big)_{ij} \subseteq \bigcup_{j=1}^{n} M'_{ij},$$

which straightforwardly implies the inequality (7). □

As an application, consider the channel from Figure 4. Intuitively, the loops in states 2 and 3 induce "ambiguities", hence they should not be both used forever.

It is easy to observe that the ω-regular language of the input Büchi automaton in which states 1 and 3 are final is $\mathcal{T}$-distinguishable. Then, by Proposition 3.3,

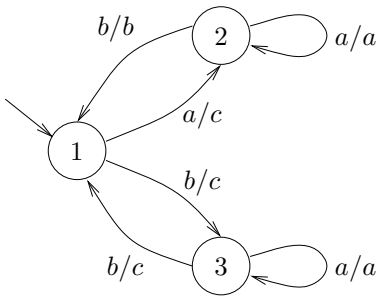

FIGURE 4. An example with a bitrate of $\log_2 \frac{1+\sqrt{13}}{2}$.

the bitrate of $\mathcal{T}$ equals the entropy of this Büchi automaton, which is the same as the entropy of $In_{\mathcal{T}}$, that is, $\log_2 \frac{1+\sqrt{13}}{2} \approx 1.203$. Note also that $\mathcal{T}$ is input-deterministic, hence the results from Section 5.2 below can be applied here.

Let us apply Proposition 4.6 with $k = 1, 2, 3$ and compute lower bounds for the bitrate of this covert channel. First, for every state we compute the k-step input/output function from $\{a, b\}^k$ to $\{a, b, c\}^k \times Q$. Next, using this function we find $\mathcal{T}$-distinguishable subsets W_i of $\{a, b\}^k$ at every state i (w.l.o.g. we consider only maximal W_i). For every W_i we find the corresponding i-th row of the matrix V, using identities (5) and (6). The results are presented in Tables 1-3.

The row sets generate then IRU sets of matrices $\mathcal{M}_k$. For $k = 1$, using Table 1, we get that the set $\mathcal{M}_1$ contains only 2 elements:

$$\begin{pmatrix} 0 & 1 & 0 \\ 1 & 1 & 0 \\ 1 & 0 & 1 \end{pmatrix}; \quad \begin{pmatrix} 0 & 0 & 1 \\ 1 & 1 & 0 \\ 1 & 0 & 1 \end{pmatrix}.$$

The set $\mathcal{M}_2$ contains 8 elements:

$$\begin{pmatrix} 2 & 1 & 0 \\ 1 & 2 & 0 \\ 1 & 1 & 1 \end{pmatrix}; \begin{pmatrix} 2 & 1 & 0 \\ 1 & 2 & 0 \\ 1 & 0 & 2 \end{pmatrix}; \begin{pmatrix} 2 & 1 & 0 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{pmatrix}; \begin{pmatrix} 2 & 1 & 0 \\ 1 & 1 & 1 \\ 1 & 0 & 2 \end{pmatrix};$$
$$\begin{pmatrix} 2 & 0 & 1 \\ 1 & 2 & 0 \\ 1 & 1 & 1 \end{pmatrix}; \begin{pmatrix} 2 & 0 & 1 \\ 1 & 2 & 0 \\ 1 & 0 & 2 \end{pmatrix}; \begin{pmatrix} 2 & 0 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{pmatrix}; \begin{pmatrix} 2 & 0 & 1 \\ 1 & 1 & 1 \\ 1 & 0 & 2 \end{pmatrix}.$$

Finally, for $k = 3$ the matrix set $\mathcal{M}_3$ contains $4 \cdot 3 \cdot 3 = 36$ matrices.

We conclude this example with the lower bounds $\frac{1}{k} \log_2 \rho(\mathcal{M}_k)$ of the bitrate presented in Table 4. The computations were done using the JAMA package, available from http://math.nist.gov/javanumerics/jama/.

TABLE 1. Computing W_i and row sets for $k = 1$. Underlined translations cannot be in the same W_i.

I/O function				
i	Translation	j	W_i	Variants for row i in $\mathcal{M}_1$
1	$\underline{a \to c}$	2	$\{a\}$	010
	$\underline{b \to c}$	3	$\{b\}$	001
2	$a \to a$	2	$\{a, b\}$	110
	$b \to b$	1		
3	$a \to a$	3	$\{a, b\}$	101
	$b \to c$	1		

TABLE 2. Computing W_i and row sets for $k = 2$. Underlined translations cannot be in the same W_i.

I/O function				
i	Translation	j	W_i	Variants for row i in $\mathcal{M}_2$
1	$\underline{aa \to ca}$	2	$\{aa, ab, bb\}$	210
	$ab \to cb$	1	$\{ba, ab, bb\}$	201
	$\underline{ba \to ca}$	3		
	$bb \to cc$	1		
2	$aa \to aa$	2	$\{aa, ab, ba\}$	120
	$ab \to ab$	1	$\{aa, ab, bb\}$	111
	$\underline{ba \to bc}$	2		
	$\underline{bb \to bc}$	1		
3	$aa \to aa$	3	$\{aa, ab, ba\}$	111
	$ab \to ac$	1	$\{aa, ab, bb\}$	102
	$\underline{ba \to cc}$	2		
	$\underline{bb \to cc}$	3		

TABLE 3. Sets W_i and row sets for $k = 3$.

i	W_i	Variants for row i in $\mathcal{M}_3$
1	$\{(aaa \text{ or } baa), aab, (aba \text{ or } abb), bab, (bba \text{ or } bbb)\}$	203, 212, 221, 230
2	$\{aaa, aab, (aba \text{ or } abb), (baa \text{ or } bba), bab, bbb\}$	330, 321, 312
3	$\{aaa, aab, (aba \text{ or } abb), (baa \text{ or } bba), bab, bbb\}$	303, 321, 312

5. Exact computation of the bitrate in special cases

Here we give an ad-hoc method for the exact computation of the bitrate in a special case. Then, for input-deterministic channels, we give a method based on uniformization results for ω-rational relations. Recall that a synchronous transducer (and hence a covert channel) is called *input-deterministic* if the underlying input automaton is deterministic and no two transitions between the same pair of states are labeled with the same input symbol.

TABLE 4. Lower bounds obtained with the JSR method; the exact bitrate is $\log_2 \frac{1+\sqrt{13}}{2} \approx 1.203$.

k	$\rho(\mathcal{M}_k)$	Bitrate bound $\frac{1}{k}\log_2 \rho(\mathcal{M}_k)$
1	1.618	0.694
2	3	0.792
3	5.541	0.823
4	11	0.865
5	21.511	0.885
6	43	0.904

5.1. SHRINKING THE INPUT LANGUAGE

Proposition 5.1. *Given a covert channel $\mathcal{T}$ and an input-deterministic transducer $\mathcal{D}$ whose input language is Σ^*, suppose that*

$$\forall w_1, w_2 \in \mathcal{L}_{in}(\mathcal{T}),\ \text{if } \mathcal{T}(w_1) \cap \mathcal{T}(w_2) = \emptyset \text{ then } \mathcal{T}(\mathcal{D}(w_1)) \cap \mathcal{T}(\mathcal{D}(w_2)) = \emptyset.$$

Then

$$\mathcal{B}(\mathcal{T}) \leq \mathcal{H}\big(\mathcal{D}(\mathcal{L}_{in}(\mathcal{T}))\big).$$

As a corollary, if $\mathcal{D}(\Sigma^n)$ is $\mathcal{T}$-distinguishable for an infinite sequence of integers n, then

$$\mathcal{B}(\mathcal{T}) = \mathcal{H}\big(\mathcal{D}(\mathcal{L}_{in}(\mathcal{T}))\big).$$

Proof. Note first that the hypothesis implies that $\mathcal{D}$ is injective on any language $L \subseteq \Sigma^n$ which is $\mathcal{T}$-distinguishable ($n \in \mathbb{N}$). Therefore, for any such L, $\mathcal{D}(L)$ is $\mathcal{T}$-distinguishable and

$$\operatorname{card}(L) = \operatorname{card}(\mathcal{D}(L)) \leq \operatorname{card}\big(\mathcal{D}(\mathcal{L}_{in}(\mathcal{T}) \cap \Sigma^n)\big).$$

This implies directly that $\mathcal{B}(\mathcal{T}) \leq \mathcal{H}\big(\mathcal{D}(\mathcal{L}_{in}(\mathcal{T}))\big)$, and then the corollary follows straightforwardly from Proposition 3.1. □

As an application, consider the covert channel in Figure 5a, and the input-deterministic transducer at (b). The transducer satisfies the conditions in Proposition 5.1, and proves that

$$\mathcal{B}(\mathcal{T}_4) = \mathcal{H}\big((a+c+e)^*\big) = \log_2 3.$$

On the other hand, the deterministic transducer in Figure 5c does not satisfy the conditions in Proposition 5.1, since $\mathcal{D}'(q,a) = \mathcal{D}'(q,c)$ but $\mathcal{T}(q,a) \cap \mathcal{T}(q,c) = \emptyset$. This confirms the intuition that, in Figure 5, a two-leaf clover whose leaves are the opposite loops labeled with x and y, in spite of being $\mathcal{T}$-distinguished, has a strictly smaller entropy that what could be used as bitrate.

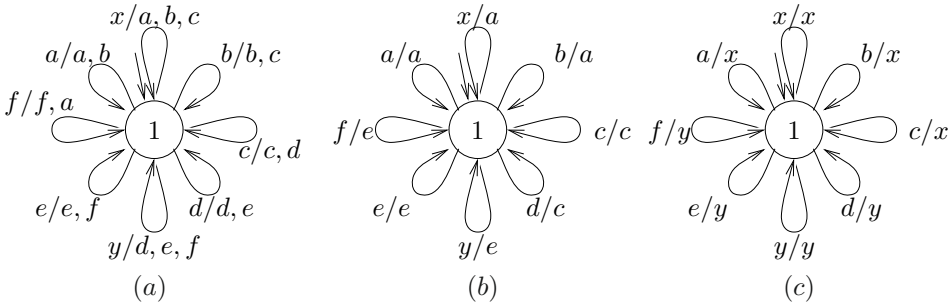

FIGURE 5. At (a), a covert channel of bitrate $\log_2 3$, and an input-deterministic transducer (b) that proves this, using Proposition 5.1. At (c), a deterministic transducer not satisfying the hypothesis in Proposition 5.1.

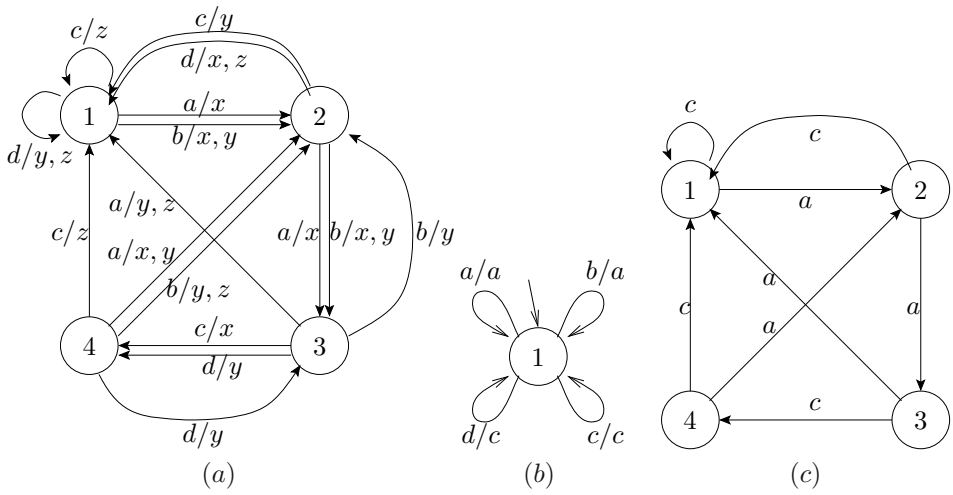

FIGURE 6. A second example of an application of Proposition 5.1.

As another application, consider the channel in Figure 6 at (a) with the deterministic transducer at (b), which fulfills the hypothesis in Proposition 5.1. The image $\mathcal{D}(\mathcal{L}(\mathcal{T}))$ is accepted by the automaton at (c), hence the (ω-regular or simple) bitrate of the channel at (a) is 1.

Unfortunately Proposition 5.1 cannot be applied in every situation. Consider the five-leaves clover in Figure 7. In this situation, there is no morphism $\mathcal{D}$ which satisfies the hypotheses in Proposition 5.1.

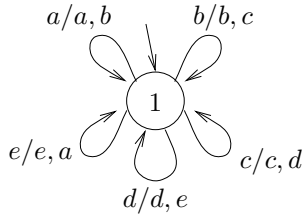

FIGURE 7. A covert channel whose bitrate cannot be computed using Proposition 5.1.

5.2. The case of input-deterministic covert channels

In this section we show that the two problems related with the computation of the bitrate are decidable for the case of input-deterministic covert channels.

Our results are based on two facts. The first one is a uniformization property for rational relations from [4] (here $\sharp(f)$ denotes the graph of f):

Proposition 5.2 [4]**.** *For any synchronous ω-rational relation $\mathcal{R} \subseteq \Sigma_1^\omega \times \Sigma_2^\omega$ there exists a synchronous ω-rational partial function $f : \Sigma_1^\omega \to \Sigma_2^\omega$ such that $\mathsf{dom}(f) = \mathsf{dom}(\mathcal{R})$ and $\sharp(f) \subseteq \mathcal{R}$.*

The second useful property is the following:

Lemma 5.3. *For any synchronous ω-rational partial function $f : \Sigma_1^\omega \to \Sigma_2^\omega$, if f is injective, then $\mathcal{H}(\mathsf{dom}(f)) = \mathcal{H}(\mathsf{range}(f))$.*

Proof. Clearly, for each $n \in \mathbb{N}$, $\mathrm{card}\big(\mathsf{dom}(f)[1..n]\big) \geq \mathrm{card}\big(\mathsf{range}(f)[1..n]\big)$. For the reverse inequality, take $\mathcal{U} = (Q, \Sigma_1, \Sigma_2, \delta, q_0, F)$ some transducer with Büchi acceptance condition, which realizes f, *i.e.*, $T_{\mathcal{U}} = f$. We may also assume that $\mathcal{U}$ is input-deterministic, due to the fact that f is a synchronous ω-regular partial function.

Remark 5.4. It is well-known that the entropy of a Büchi automaton equals the largest entropy of one of its subautomata containing a single SCC (strongly-connected component) which has a nonempty intersection with F. This remark can be seen as a corollary of a similar fact about spectral radii of matrices.

So take $\mathcal{U}' = (Q', \Sigma_1, \Sigma_2, \delta', q_0', F')$ a sub-transducer of $\mathcal{U}$ with $Q' \subseteq Q$ forming a SCC, $\delta' = \delta\big|_{Q \times \Sigma_1 \times \Sigma_2 \times Q'}$, $q_0' \in Q'$ and $F' = F \cap Q' \neq \emptyset$, and such that the input automaton $In_{\mathcal{U}'}$ has the same entropy as $\mathcal{L}_{in}(\mathcal{U})$.

Clearly $\mathcal{H}(\mathsf{dom}(f)) \geq \mathcal{H}(\mathcal{L}^\omega_{out}(\mathcal{U}'))$; also $\mathcal{H}(\mathcal{L}^\omega_{in}(\mathcal{U}')) = \mathcal{H}(\mathsf{dom}(f))$ by construction. It is also easy to see that $\mathcal{H}(\mathcal{L}^\omega_{out}(\mathcal{U}')) \leq \mathcal{H}(\mathsf{range}(f)) \leq \mathcal{H}(\mathsf{dom}(f))$. We will prove that for any $n \in \mathbb{N}$,

$$\mathrm{card}(\mathcal{L}^\omega_{in}(\mathcal{U}')[1..n]) \leq \mathrm{card}(\mathcal{L}^\omega_{out}(\mathcal{U}')[1..n + 3 \cdot \mathrm{card}(Q')]). \tag{8}$$

This would then imply that $\mathcal{H}(\mathcal{L}^\omega_{in}(\mathcal{U}')) \leq \mathcal{H}(\mathcal{L}^\omega_{out}(\mathcal{U}'))$ which would conclude our proof.

To this end, take two ω-words $w_1, w_2 \in \mathcal{L}^{\omega}_{in}(\mathcal{U}')$ and some $n \in \mathbb{N}$ such that $w_1[1..n] \neq w_2[1..n]$, and also $\mathcal{U}'(q'_0, w_1[1..n]) = \mathcal{U}'(q'_0, w_2[1..n])$. Denote further $q_1 = \delta'_{in}(q'_0, w_1[1..n])$ and $q_2 = \delta'_{in}(q'_0, w_2[1..n])$ (recall that $\mathcal{U}$ is synchronous and input-deterministic, and so is $\mathcal{U}'$).

We will actually prove the following claim:

Claim 5.5. (*) There exist $w'_1, w'_2 \in \Sigma_1^{3\cdot \mathrm{card}(Q')}$ with

$$T_{\mathcal{U}'}(q'_0, w_1[1..n]w'_1) \neq T_{\mathcal{U}'}(q'_0, w_2[1..n]w'_2).$$

Note that this claim implies the inequality (8), which would end our proof.

Assume then, for the sake of contradiction, that the desired property (*) does not hold. Choose $w_3, w_4, w_5 \in \Sigma_1^*$ and $q' \in F'$ such that

$$\delta'_{in}(q_1, w_3) = q_2, \qquad \delta'_{in}(q_2, w_4) = q', \qquad \delta'_{in}(q', w_5) = q_1.$$

Note that this is always possible, by the assumption that Q' is strongly connected, and that $F' \neq \emptyset$. Note also that the three words can be chosen such that $w_3w_4w_5 \leq 3 \cdot \mathrm{card}(Q')$.

The assumption that the property (*) does not hold implies that

$$T_{\mathcal{U}'}(q'_0, w_1[1..n]w_3w_4w_5) = T_{\mathcal{U}'}(q'_0, w_2[1..n]w_4w_5w_3).$$

But then it is easy to see that we also have, for any $m \in \mathbb{N}$,

$$T_{\mathcal{U}'}(q'_0, w_1[1..n](w_3w_4w_5)^m) = T_{\mathcal{U}'}(q'_0, w_2[1..n](w_4w_5w_3)^m),$$

which implies that

$$T_{\mathcal{U}'}(q'_0, w_1[1..n](w_3w_4w_5)^\omega) = T_{\mathcal{U}'}(q'_0, w_2[1..n](w_4w_5w_3)^\omega).$$

Note also that $w_1[1..n](w_3w_4w_5)^\omega \in \mathcal{L}^{\omega}_{in}(\mathcal{U}')$, due to the fact that:

$$\delta'_{in}\Big(q'_0, w_1[1..n](w_3w_4w_5)^m w_3w_4\Big) = q' \in F' \text{ for any } m \in \mathbb{N}$$

and the same holds for $w_2[1..n](w_4w_5w_3)^\omega$.

But these two facts are in contradiction with f being injective. □

Now we are ready to characterize the bitrate of input-deterministic channels.

Proposition 5.6. *For input-deterministic channels $\mathcal{T}$ the ω-regular bitrate equals the "simple" bitrate, and both are equal with the entropy of the output language of $\mathcal{T}$. Moreover one can effectively construct a regular realization of the bitrate.*

Proof. We may apply Proposition 5.2 to the synchronous ω-rational relation $T_{\mathcal{T}}^{-1} \subseteq \Gamma^\omega \times \Sigma^\omega$ – recall that $T_{\mathcal{T}}$ is the relation defined by $\mathcal{T}$ – to get a synchronous ω-rational partial function $f : \Gamma^\omega \to \Sigma^\omega$ with $\sharp(f) \subseteq T_{\mathcal{T}}^{-1}$. As a consequence of this latter fact, $\mathsf{range}(f)$ is a $\mathcal{T}$-distinguishable ω-regular language.

By virtue of Lemma 5.3, $\mathcal{H}(\mathsf{dom}(f)) = \mathcal{H}(\mathsf{range}(f))$. This result, combined with the fact that $\mathcal{H}\big(\mathcal{L}^{\omega}_{out}(\mathcal{T})\big) = \mathcal{H}\big(\mathcal{L}_{out}(\mathcal{T})\big) = \mathcal{H}\big(\mathsf{range}(T_{\mathcal{T}})\big) = \mathcal{H}\big(\mathsf{dom}(f)\big)$, and with the fact that $\mathsf{range}(f)$ is a $\mathcal{T}$-distinguishable ω-regular language, implies that $\mathcal{B}^{\omega}_{r}(\mathcal{T}) \geq \mathcal{H}\big(\mathcal{L}_{out}(\mathcal{T})\big)$.

Combining these with Proposition 3.1, which says that $\mathcal{B}(\mathcal{T}) \leq \mathcal{H}\big(\mathcal{L}_{out}(\mathcal{T})\big)$, and with Proposition 3.3, by which $\mathcal{B}^{\omega}_{r}(\mathcal{T}) \leq \mathcal{B}(\mathcal{T})$, we get the desired result. □

As an application, the channel in Figure 1 is an input-deterministic channel, therefore its bitrate equals the entropy of its output language, which is 1.

6. Conclusions

We have presented some results related to the computation of the capacity of a covert channel, modeled as a generalized form of entropy of a transducer. The results show that some underapproximations are possible using Turán's theorem or by computing some joint spectral radii, and in some special cases the exact computation can be done.

The main conjecture is that the covert channel capacity is computable in general, even for nondeterministic channels, by constructing a regular presentation of a realization of the bitrate.

Acknowledgements. The authors would like to thank the two anonymous reviewers whose remarks have helped to significantly improve our paper. Many thanks also to Dominique Perrin for pointing to us the uniformization result of Ch. Choffrut and S. Grigorieff from [4].

References

[1] M.-P. Béal, *Codage Symbolique*. Masson (1993).
[2] M.-P. Béal and O. Carton, Determinization of transducers over finite and infinite words. *Theoretical Computer Science* **289** (2002) 225–251.
[3] V. Blondel and Yu. Nesterov, Polynomial-time computation of the joint spectral radius for some sets of nonnegative matrices. *SIAM Journal of Matrix Analysis and Applications* **31** (2009) 865–876.
[4] C. Choffrut and S. Grigorieff, Uniformization of rational relations, in *Jewels are Forever, Contributions on Theoretical Computer Science in Honor of Arto Salomaa*, edited by J. Karhumäki, H.A. Maurer, Gh. Paun and G. Rozenberg. Springer (1999) 59–71.
[5] Department of defense trusted computer system evaluation criteria. DOD 5200.28-STD, National Computer Security Center (December 1985).
[6] C. Frougny and J. Sakarovitch, Synchronisation déterministe des automates à délai borné. *Theoretical Computer Science* **191** (1998) 61–77.
[7] F.R. Gantmacher, *The theory of matrices*. AMS Chelsea Publishing (1959).
[8] J.A. Goguen and J. Meseguer, Security policies and security models, in *Proceedings of the IEEE Symposium on Security and Privacy*, Oakland, CA, USA (1982) 11–20.
[9] R. Jungers, Vl. Protasov and V. Blondel, Efficient algorithms for deciding the type of growth of products of integer matrices. *Linear Algebra and its Applications* **428** (2008) 2296–2311.
[10] D. Lind and B. Marcus, *An Introduction to Symbolic Dynamics and Coding*. Cambridge University Press (1995).

[11] G. Lowe, Quantifying information flow, in *Proceedings of the 15th IEEE Computer Security Foundations Workshop (CSFW'02)*, IEEE Computer Society (2002) 18–31.
[12] J.K. Millen, Finite-state noiseless covert channels, in *Proceedings of the 2nd IEEE Computer Security Foundations Workshop (CSFW'89)*, IEEE Computer Society (1989) 81–86.
[13] V. Protasov, R. Jungers and V. Blondel, Joint spectral characteristics of matrices: a conic programming approach. http://www.inma.ucl.ac.be/~jungers/publis_dispo/conic.pdf (2009).
[14] W. Thomas, Languages, automata, and logic, in *Handbook of Formal Languages*, Vol. III. Springer Verlag (1997) 389–455.
[15] P. Turán, On an extremal problem in graph theory. *Matematicko Fizicki Lapok* **48** (1941) 436–452 (in Hungarian).
[16] J.T. Wittbold and D.M. Johnson, Information flow in nondeterministic systems, in *IEEE Symposium on Security and Privacy* (1990) 144–161.