Chapter 7

Complexity of Encoders

7.1 Complexity criteria

There are various criteria that are used to measure the performance and complexity of encoders, and their corresponding decoders. We list here the predominant factors that are usually taken into account while designing rate $p : q$ finite-state encoders.

The values of p and q. Typically, the rate p/q is chosen to be as close to $\mathsf{cap}(S)$ as possible, subject to having p and q small enough: The reason for the latter requirement is minimizing the number of outgoing edges, 2^p, from each state in the encoder and keeping to a minimum the number of input–output connections of the encoder.

Number of states in an encoder. In both hardware and software implementation of encoders $\mathcal{E}$, we will need $\lceil \log |V_{\mathcal{E}}| \rceil$ bits in order to represent the current state of $\mathcal{E}$. This motivates an encoder design with a relatively small number of states [Koh78, Ch. 9].

Gate complexity. In addition to representing the state of a finite-state encoder, we need, in hardware implementation, to realize the next-state function and the output function as a gate circuit. Hardware complexity is usually measured in terms of the number of required gates (e.g., NAND gates), and this number also includes the implementation of the memory bit cells that represent the encoder state (each memory bit cell can be realized by a fixed number of gates).

The number of states in hardware implementation becomes more significant in applications where we run several encoders *in parallel* with a common circuit for the next-state and output functions, but with duplicated hardware for representing the state of each encoder.

Time and space complexity of a RAM program. When the finite-state encoder is to be implemented as a computer program on a random-access machine (RAM) [AHU74, Ch. 1], the complexity is usually measured by the space requirements and the running time of the program.

Encoder anticipation. One way to implement a decoder for an encoder $\mathcal{E}$ with finite anticipation $\mathcal{A}(\mathcal{E})$ is by accumulating the past $\mathcal{A}(\mathcal{E})$ symbols (in $\Sigma(S^q)$) that were generated by $\mathcal{E}$; these symbols, with the current symbol, allow the decoder to simulate the state transitions of $\mathcal{E}$ and, hence, to reconstruct the sequence of input tags (see Section 4.1). The size of the required buffer thus depends on $\mathcal{A}(\mathcal{E})$.

Window length of sliding-block decodable encoders. A typical decoder of an (m, a)-sliding-block decodable encoder consists of a buffer that accumulates the past $\mathsf{m}+\mathsf{a}$ symbols (in $\Sigma(S^q)$) that were generated by the encoder. A decoding function $\mathcal{D}$: $(\Sigma(S^q))^{(\mathsf{m}+\mathsf{a}+1)} \to \{0, 1, \ldots, 2^p-1\}$ is then applied to the current symbol and to the contents of the buffer to reconstruct an input tag in $\{0, 1\}^p$ (see Section 4.3). From a complexity point-of-view, the window length, $\mathsf{m}+\mathsf{a}+1$, determines the size of the required buffer.

In order to establish a general framework for comparing the complexity of encoders generated by different methods of encoder synthesis, we need to set some canonical presentation of the constrained system S, in terms of which the complexity will be measured. We adopt the Shannon cover of S as such a distinguished presentation.

7.2 Number of states in the encoder

In this section we present upper and lower bounds on the smallest number of states in any (S, n)-encoder for a given constrained system S and integer n.

Let G be a deterministic presentation of S. As was described in Chapter 5, the state-splitting algorithm [ACH83] starts with an (A_G, n)-approximate eigenvector $\mathbf{x} = (x_v)_{v \in V_G}$, which guides the splitting of the states in G until we obtain an (S, n)-encoder with at most $\sum_{v \in V_G} x_v = \|\mathbf{x}\|_1$ states. Hence, we have the following.

Theorem 7.1 *Let S be a constrained system presented by a deterministic graph G and let n be a positive integer. Assume that $\mathsf{cap}(S) \geq \log n$. Then, there exists an (S, n)-encoder $\mathcal{E}$ such that*

$$|V_{\mathcal{E}}| \leq \min_{\mathbf{x} \in \mathcal{X}(A_G, n)} \|\mathbf{x}\|_1 \ .$$

On the other hand, the following lower bound on the number of states of any (S, n)-encoder was obtained in [MR91].

Theorem 7.2 *Let S be a constrained system presented by a deterministic graph G and let n be a positive integer. Assume that* $\mathsf{cap}(S) \ge \log n$. *Then, for any (S,n)-encoder $\mathcal{E}$,*

$$|V_{\mathcal{E}}| \ge \min_{\mathbf{x} \in \mathcal{X}(A_G, n)} \|\mathbf{x}\|_\infty .$$

Proof: Let $\mathcal{E}$ be an (S,n)-encoder and let $\Sigma = \Sigma(S)$. The following construction effectively provides an (A_G,n)-approximate eigenvector $\mathbf{x}$ which satisfies the inequality $|V_{\mathcal{E}}| \ge \|\mathbf{x}\|_\infty$.

(a) Construct a deterministic graph $H = H(\mathcal{E})$ which presents $S' = S(\mathcal{E})$.

This can be done using the determinizing graph of Section 2.2.1.

(b) For an irreducible sink H' of H, define a vector $\boldsymbol{\xi} \ne \mathbf{0}$ such that $A_{H'}\boldsymbol{\xi} = n\boldsymbol{\xi}$.

Let H' be an irreducible sink of H. Recall that each state $Z \in V_{H'} \left(\subseteq V_H\right)$ is a subset $T_{\mathcal{E}}(\mathbf{w}, v)$ of states of $\mathcal{E}$ that can be reached in $\mathcal{E}$ from a given state $v \in V_{\mathcal{E}}$ by paths that generate a given word $\mathbf{w}$. Let $\xi_Z = |Z|$ denote the number of states of $\mathcal{E}$ in Z and let $\boldsymbol{\xi}$ be the positive integer vector defined by $\boldsymbol{\xi} = (\xi_Z)_{Z \in V_{H'}}$. We now claim that

$$A_{H'}\boldsymbol{\xi} = n\boldsymbol{\xi} .$$

Consider a state $Z \in V_{H'}$. Since $\mathcal{E}$ has out-degree n, the number of edges in $\mathcal{E}$ outgoing from the set of states $Z \subseteq V_{\mathcal{E}}$ is $n|Z|$. Now, let E_a denote the set of edges in $\mathcal{E}$ labeled a that start at the states of $\mathcal{E}$ in Z and let Z_a denote the set of terminal states, in $\mathcal{E}$, of these edges. Note that the sets E_a, for $a \in \Sigma$, induce a partition on the edges of $\mathcal{E}$ outgoing from Z. Clearly, if $Z_a \ne \emptyset$, there is an edge $Z \xrightarrow{a} Z_a$ in H and, since H' is an irreducible sink, this edge is also contained in H'. We now claim that any state $u \in Z_a$ is accessible in $\mathcal{E}$ by exactly one edge labeled a whose initial state is in Z; otherwise, if $Z = T_{\mathcal{E}}(\mathbf{w}, v)$, the word $\mathbf{w}a$ could be generated in $\mathcal{E}$ by two distinct paths which start at v and terminate in u, contradicting the losslessness of $\mathcal{E}$. Hence, $|E_a| = |Z_a|$ and, so, the entry of $A_{H'}\boldsymbol{\xi}$ corresponding to the state Z in H' satisfies

$$\begin{aligned}(A_{H'}\boldsymbol{\xi})_Z &= \sum_{Y \in V_{H'}} (A_{H'})_{Z,Y}\xi_Y = \sum_{Y \in V_{H'}} (A_{H'})_{Z,Y}|Y| \\ &= \sum_{a \in \Sigma} |Z_a| = \sum_{a \in \Sigma} |E_a| = n|Z| = n\xi_Z ,\end{aligned}$$

as desired.

(c) Construct an (A_G,n)-approximate eigenvector $\mathbf{x} = \mathbf{x}(\mathcal{E})$ from $\boldsymbol{\xi}$.

As G and H' comply with the conditions of Lemma 2.13, each follower set of a state in H' is contained in a follower set of some state in G. Let $\mathbf{x} = (x_u)_{u \in V_G}$ be the nonnegative integer vector defined by

$$x_u = \max\left\{\, \xi_Z \;:\; Z \in V_{H'} \text{ and } \mathcal{F}_{H'}(Z) \subseteq \mathcal{F}_G(u) \,\right\} , \quad u \in V_G ,$$

and denote by $Z(u)$ some particular state Z in H' for which the maximum is attained. In case there is no state $Z \in V_{H'}$ such that $\mathcal{F}_{H'}(Z) \subseteq \mathcal{F}_G(u)$, define $x_u = 0$ and $Z(u) = \emptyset$. We claim that $\mathbf{x}$ is an (A_G, n)-approximate eigenvector. First, since $V_{H'}$ is nonempty, we have $\mathbf{x} \neq \mathbf{0}$. Now, let u be a state in G; if $x_u = 0$ then, trivially, $(A_G\mathbf{x})_u \geq nx_u$ and, so, we can assume that $x_u \neq 0$. Let $Z_a(u)$ be the terminal state in H' for an edge labeled a outgoing from $Z(u)$. Since $\mathcal{F}_{H'}\big(Z(u)\big) \subseteq \mathcal{F}_G(u)$, there exists an edge labeled a in G from u which terminates in some state u_a in G; and, since G and H' are both deterministic, we have $\mathcal{F}_{H'}\big(Z_a(u)\big) \subseteq \mathcal{F}_G(u_a)$. Furthermore, by the way $\mathbf{x}$ was defined, we have $x_{u_a} \geq \xi_{Z_a(u)}$ and, so, letting $\Sigma_{Z(u)}$ denote the set of labels of edges in H' outgoing from $Z(u)$, we have

$$(A_G\mathbf{x})_u \geq \sum_{a \in \Sigma_{Z(u)}} x_{u_a} \geq \sum_{a \in \Sigma_{Z(u)}} \xi_{Z_a(u)} = (A_{H'}\boldsymbol{\xi})_{Z(u)} = n\xi_{Z(u)} = nx_u \ ,$$

where we have used the equality $A_{H'}\boldsymbol{\xi} = n\boldsymbol{\xi}$. Hence, $A_G\mathbf{x} \geq n\mathbf{x}$.

The theorem now follows from the fact that each entry in $\mathbf{x}$ is a size of a subset of states of $V_{\mathcal{E}}$. □

The bound of Theorem 7.2 can be effectively computed by the Franaszek algorithm which was described in Section 5.2.2. The upper bound of Theorem 7.1 is at most $|V_G|$ times the lower bound of Theorem 7.2, which amounts to an additive term of $\log|V_G|$ in the number of bits required to represent the current state of the encoder.

There are examples of sequences of labeled graphs G for which the lower bound of Theorem 7.2 is *exponential* in the number of states of G [Ash88], [MR91]. We give here such an example, which appears in [AMR95] and [MR91].

Example 7.1 Let r be a positive integer and let Σ denote the alphabet of size $r^2{+}r{+}1$ given by $\{a\} \cup \{b_i\}_{i=1}^{r^2+r-1} \cup \{c\}$. Consider the constrained systems S_k that are presented by the graphs G_k of Figure 7.1 (from each state $u \leq k$ there are $r^2{+}r{-}1$ parallel outgoing edges labeled by the b_i's to state $k{+}u$). It is easy to verify that $\lambda(A_{G_k}) = \lambda = r{+}1$ and that every (A_{G_k}, λ)-approximate eigenvector is a multiple of $(\lambda\ \lambda^2\ \ldots\ \lambda^k\ 1\ \lambda\ \ldots\ \lambda^{k-1})^\top$. Hence, by Theorem 7.2, every $(S_k, r{+}1)$-encoder must have at least $(r{+}1)^k = \exp\{O(|V_{G_k}|)\}$ states.

On the other hand, note that the vector $\mathbf{x} = (x_u)_u$, whose nonzero components are $x_k = r$ and $x_{2k} = 1$, is an (A_{G_k}, r)-approximate eigenvector. Hence, if we can compromise on the rate and construct (S_k, r)-encoders instead, then the state-splitting algorithm provides such encoders with at most $r{+}1$ states. □

The bound of Theorem 7.2 is based on the existence of an approximate eigenvector $\mathbf{x} = \mathbf{x}(\mathcal{E})$, where each of the entries in $\mathbf{x}$ is a size of a subset of $V_{\mathcal{E}}$. The improvements on this bound, given in [MR91], are obtained by observing that some of these subsets might be disjoint. One such improvement (with proof left to the reader) is as follows.

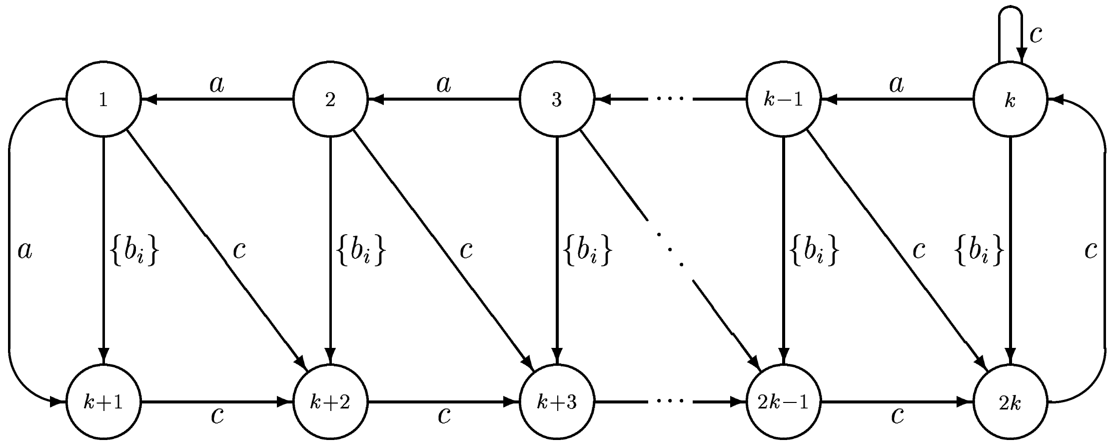

Figure 7.1: Labeled graph G_k for Example 7.1.

Theorem 7.3 *Let S be a constrained system presented by a deterministic graph G and let n be a positive integer. Assume that* $\mathsf{cap}(S) \geq \log n$*. Then, for any (S, n)-encoder $\mathcal{E}$,*

$$|V_{\mathcal{E}}| \geq \min_{\mathbf{x} \in \mathcal{X}(A_G, n)} \max_{U} \sum_{u \in U} x_u ,$$

where the maximum is taken over all subsets $U \subseteq V_G$ such that $\mathcal{F}_G(u) \cap \mathcal{F}_G(u') = \emptyset$ for every distinct states u and u' in U.

In fact, the preceding result can be generalized further to obtain the best general lower bound known on the number of states in any (S, n)-encoder, as it is stated in [MR91]. In order to state this result, we need the following definitions.

Let S be a constrained system presented by a deterministic graph G. For a state $u \in V_G$ and a word $\mathbf{w} \in \mathcal{F}_G(u)$, let $\tau_G(\mathbf{w}, u)$ be the terminal state of the path in G that starts at u and generates $\mathbf{w}$. (Using the notations of Section 2.2.1, we thus have $T_G(\mathbf{w}, u) = \{\tau_G(\mathbf{w}, u)\}$.) For a word $\mathbf{w} \notin \mathcal{F}_G(u)$, define $\tau_G(\mathbf{w}, u) = \emptyset$.

Let n be a positive integer and $\mathbf{x} = (x_u)_{u \in V_G}$ be an (A_G, n)-approximate eigenvector. For a word $\mathbf{w}$ and a subset $U \subseteq V_G$, let $I_G(\mathbf{x}, \mathbf{w}, U)$ denote a state $u \in U$ such that $x_{\tau_G(\mathbf{w},u)}$ is maximal (for the case where $\tau_G(\mathbf{w}, u) = \emptyset$, we define $x_\emptyset = 0$).

Let U be a subset of V_G. A list C of words is *U-complete* in G, if every word in $\bigcup_{u \in U} \mathcal{F}_G(u)$ either has a prefix in C or is a prefix of a word in C. Let $\mathcal{C}_G(U)$ denote the set of all finite U-complete lists in G. For example, the list $\mathcal{F}_G^m(U)$ of all words of length m that can be generated in G from states of U, belongs to $\mathcal{C}_G(U)$.

Finally, given an integer n, an (A_G, n)-approximate eigenvector $\mathbf{x}$, a subset U of V_G, and a list C of words, we define $\mu_G(\mathbf{x}, n, U, C)$ by

$$\mu_G(\mathbf{x}, n, U, C) = \sum_{u \in U} x_u - \sum_{\mathbf{w} \in C} n^{-\ell(\mathbf{w})} \sum_{u \in U - \{I_G(\mathbf{x}, \mathbf{w}, U)\}} x_{\tau_G(\mathbf{w}, u)}$$

(recall that $\ell(\mathbf{w})$ is the length of the word $\mathbf{w}$).

Theorem 7.4 *Let S be a constrained system presented by a deterministic graph G and let n be a positive integer. Assume that $\mathsf{cap}(S) \geq \log n$. Then, for any (S, n)-encoder $\mathcal{E}$,*

$$|V_\mathcal{E}| \geq \min_{\mathbf{y} \in \mathcal{X}(A_G, n)} \max_{U \subseteq V_G} \sup_{C \in \mathcal{C}_G(U)} \mu_G(\mathbf{y}, n, U, C) .$$

In particular, for all $U \subseteq V_G$ and m,

$$|V_\mathcal{E}| \geq \min_{\mathbf{y} \in \mathcal{X}(A_G, n)} \mu_G(\mathbf{y}, n, U, \mathcal{F}_G^m(U)) .$$

Example 7.2 Figure 4.6 depicts a rate 2 : 3 four-state encoder for the (1, 7)-RLL constrained system. The example therein is due to Weathers and Wolf [WW91], whereas the example of an encoder used in practice is due to Adler, Hassner, and Moussouris [AHM82] and has five states (see also [How89]). Using Theorem 7.4, a lower bound of 4 on the number of states of any such encoder is presented in [MR91]. Thus, the Weathers–Wolf encoder has the smallest possible number of encoder states. □

Example 7.3 Figure 4.4 depicts a rate 1 : 2 six-state encoder for the (2, 7)-RLL constrained system. This encoder is used in practice and is due to Franaszek [Fra72] (see also [EH78], [How89]). On the other hand, the encoder shown in Figure 4.5, which is due to Howell [How89], has only five states. Using Theorem 7.4, it can be shown that 5 is a lower bound on the number of encoder states for this system; thus, the Howell encoder has the smallest possible number of encoder states (note, however, that the anticipation of the Howell encoder is larger than Franaszek's). □

7.3 Values of p and q

When $\mathsf{cap}(S)$ is a rational number p/q, we can attain the bound of Theorem 4.2 by a rate $p : q$ finite-state encoder for S: Taking a deterministic presentation G of S, we have in this case an $(A_G^q, 2^p)$-approximate eigenvector which yields, by the state-splitting algorithm, an $(S^q, 2^p)$-encoder.

On the other hand, when $\mathsf{cap}(S)$ is not a rational number, we cannot attain the bound of Theorem 4.2 by a rate $p : q$ finite-state encoder for S. Still, we can approach the bound

$\mathsf{cap}(S)$ from below by a sequence of rate $p_m : q_m$ finite-state encoders $\mathcal{E}_m$. In fact, as stated in Theorem 4.3, we can approach capacity from below even by block encoders. It can be shown that in this way we obtain a sequence of rate $p_m : q_m$ block encoders $\mathcal{E}_m$ for S such that

$$\left| \frac{p_m}{q_m} - \mathsf{cap}(S) \right| \le \frac{\beta}{q_m}$$

for some constant $\beta = \beta(G)$. However, the constant β might be very large (e.g., exponential) in terms of the number of states of G. This means that q_m might need to be extremely large in order to have rates p_m/q_m close to capacity.

Obviously, for every sequence of rate $p_m : q_m$ finite-state encoders $\mathcal{E}_m$ for S, the number of edges in $\mathcal{E}_m$ is increasing exponentially with p_m. The question is whether convergence of p_m/q_m to $\mathsf{cap}(S)$ that is faster than $O(1/q_m)$ might force the number of states in $\mathcal{E}_m$ to blow up as well. The answer is given in the following result, which is proved in [MR91] using Theorem 7.2.

Theorem 7.5 *Let S be a constrained system with $\mathsf{cap}(S) = \log \lambda$.*

(a) If $\lambda = k^{s/t}$ for some positive integers k, s, and t, then there exists an integer N such that for any two positive integers p, q, where $p/q \le \log \lambda$ and t divides q, there is an $(S^q, 2^p)$-encoder with at most N states.

(b) If λ is not a rational power of an integer and, in addition,

$$\lim_{m\to\infty} \left(\frac{p_m}{q_m} - \log \lambda \right) \cdot q_m = 0 \ ,$$

then for any sequence of $(S^{q_m}, 2^{p_m})$-encoders $\mathcal{E}_m$,

$$\lim_{m\to\infty} |V_{\mathcal{E}_m}| = \infty \ .$$

Sketch of proof. Case (a): Let G be a deterministic presentation of S. Since $\lambda = k^{s/t}$, the matrix $(A_G)^t$ has an integer largest eigenvalue and an associated integer nonnegative right eigenvector $\mathbf{x}$. An (S^{tm}, k^{sm})-encoder $\mathcal{E}_m$ can therefore be obtained by the state-splitting algorithm for every m, with number of states which is at most $N = \|\mathbf{x}\|_1$. Write $q = tm$; we have $2^p \le k^{sm}$ and, so, an $(S^q, 2^p)$-encoder can be obtained by deleting excess edges from $\mathcal{E}_m$.

Case (b): It can be shown that if the values p_m/q_m approach $\log \lambda$ faster than $O(1/q_m)$, then the respective $((A_G)^{q_m}, 2^{p_m})$-approximate eigenvectors $\mathbf{x}$ (when scaled to have a fixed norm) approach a right eigenvector which must contain an irrational entry. Therefore, the largest components in such approximate eigenvectors tend to infinity. The result then follows from Theorem 7.2. □

If we choose p_m and q_m to be the continued fraction approximants of $\log \lambda$, we get

$$\left| \frac{p_m}{q_m} - \log \lambda \right| < \frac{\beta}{q_m^2}$$

for some constant β. So, in case (b), the fastest approach to capacity necessarily forces the number of states to grow without bound.

7.4 Encoder anticipation

7.4.1 Deciding upon existence of encoders with a given anticipation

We start with the following theorem, taken from [AMR96], which shows that checking whether there is an (S, n)-encoder with anticipation t is a decidable problem. A special case of this theorem, for $t = 0$, was alluded to in Section 4.4. Recall that $\mathcal{F}_G^t(u)$ stands for the set of words of length t that can be generated from a state u in a labeled graph G.

Theorem 7.6 *Let S be an irreducible constrained system with a Shannon cover G, let n and t be positive integers, and, for every state u in G, let $N(u,t) = |\mathcal{F}_G^t(u)|$. If there exists an (S, n)-encoder with anticipation t, then there exists an (S, n) encoder with anticipation $\leq t$ and at most $\sum_{u \in V_G} (2^{N(u,t)} - 1)$ states.*

By Lemma 2.9, we may assume that there is an *irreducible* (S, n)-encoder $\mathcal{E}$ with anticipation at most t. The proof of Theorem 7.6 is carried out by effectively constructing from $\mathcal{E}$ an (S, n)-encoder $\mathcal{E}'$ with anticipation $\leq t$ and with a number of states which is at most the bound stated in the theorem. We describe the construction of $\mathcal{E}'$ below, and the theorem will follow from the next two lemmas.

For a state $u \in V_G$ and a nonempty subset $\mathcal{F}$ of $\mathcal{F}_G^t(u)$, let $\Gamma(u, \mathcal{F})$ denote the set of all states v in $\mathcal{E}$ for which $\mathcal{F}_{\mathcal{E}}(v) \subseteq \mathcal{F}_G(u)$ and $\mathcal{F}_{\mathcal{E}}^t(v) = \mathcal{F}$. Whenever $\Gamma(u, \mathcal{F})$ is nonempty we designate a specific such state $v \in \Gamma(u, \mathcal{F})$ and call it $v(u, \mathcal{F})$. By Lemma 2.13, at least one $\Gamma(u, \mathcal{F})$ is nonempty.

We now define the labeled graph $\mathcal{E}'$ as follows. The states of $\mathcal{E}'$ are the pairs $(u, \mathcal{F})$ such that $\Gamma(u, \mathcal{F})$ is nonempty. We draw an edge $(u, \mathcal{F}) \xrightarrow{a} (\hat{u}, \widehat{\mathcal{F}})$ in $\mathcal{E}'$ if and only if there is an edge $u \xrightarrow{a} \hat{u}$ in G and an edge $v(u, \mathcal{F}) \xrightarrow{a} \hat{v}$ in $\mathcal{E}$ for some $\hat{v} \in \Gamma(\hat{u}, \widehat{\mathcal{F}})$.

Lemma 7.7 *For every $\ell \leq t+1$,*

$$\mathcal{F}_{\mathcal{E}'}^{\ell}\big((u, \mathcal{F})\big) = \mathcal{F}_{\mathcal{E}}^{\ell}\big(v(u, \mathcal{F})\big) .$$

Proof. We prove that $\mathcal{F}^{\ell}_{\mathcal{E}'}\big((u,\mathcal{F})\big) \subseteq \mathcal{F}^{\ell}_{\mathcal{E}}\big(v(u,\mathcal{F})\big)$ by induction on ℓ. We leave the reverse inclusion (which is not used here) to the reader.

The result is immediate for $\ell = 0$. Assume now that the result is true for some fixed $\ell \le t$. Let $w_0 w_1 \ldots w_\ell \in \mathcal{F}^{\ell+1}_{\mathcal{E}'}\big((u,\mathcal{F})\big)$, which implies that there is in $\mathcal{E}'$ a path of the form $(u,\mathcal{F}) \overset{w_0}{\to} (u_1,\mathcal{F}_1) \overset{w_1}{\to} (u_2,\mathcal{F}_2) \to \ldots \to (u_\ell,\mathcal{F}_\ell) \overset{w_\ell}{\to} (u_{\ell+1},\mathcal{F}_{\ell+1})$. By the inductive hypothesis, there is a path $v(u_1,\mathcal{F}_1) \overset{w_1}{\to} v_2 \overset{w_2}{\to} v_3 \to \ldots \to v_\ell \overset{w_\ell}{\to} v_{\ell+1}$ in $\mathcal{E}$. Therefore, the word $\mathbf{w} = w_1 w_2 \ldots w_\ell$ belongs to $\mathcal{F}_{\mathcal{E}}\big(v(u_1,\mathcal{F}_1)\big)$ and, since $\ell \le t$, we can extend $\mathbf{w}$ to form a word $\mathbf{w}\mathbf{w}'$ of length t that belongs to $\mathcal{F}_1$. Now, by definition of the edges in $\mathcal{E}'$, there is an edge $v(u,\mathcal{F}) \overset{w_0}{\to} \widehat{v}$ in $\mathcal{E}$ for some $\widehat{v} \in \Gamma(u_1,\mathcal{F}_1)$. Since $\mathbf{w}\mathbf{w}' \in \mathcal{F}_1$, there is a path labeled $\mathbf{w}$ outgoing from $\widehat{v}$ in $\mathcal{E}$ and, so, there is a path labeled $w_0 w_1 \ldots w_\ell$ outgoing from $v(u,\mathcal{F})$ in $\mathcal{E}$. Hence, $\mathcal{F}^{\ell+1}_{\mathcal{E}'}\big((u,\mathcal{F})\big) \subseteq \mathcal{F}^{\ell+1}_{\mathcal{E}}\big(v(u,\mathcal{F})\big)$, as desired. □

The next lemma shows that $\mathcal{E}'$ is an (S,n)-encoder with anticipation $\le t$.

Lemma 7.8 *The following three conditions hold:*

(a) The out-degree of each state in $\mathcal{E}'$ is n;

(b) $S(\mathcal{E}') \subseteq S$; and —

(c) $\mathcal{E}'$ has anticipation $\le t$.

Proof. *Part (a):* It suffices to show that there is a one-to-one correspondence between the outgoing edges of $(u,\mathcal{F})$ in $\mathcal{E}'$ and those of $v(u,\mathcal{F})$ in $\mathcal{E}$. Consider the mapping Φ from outgoing edges of $(u,\mathcal{F})$ to outgoing edges of $v(u,\mathcal{F})$ defined by

$$\Phi\Big((u,\mathcal{F}) \overset{a}{\to} (\widehat{u},\widehat{\mathcal{F}})\Big) = \Big(v(u,\mathcal{F}) \overset{a}{\to} \widehat{v}\Big)$$

where $\widehat{v} \in \Gamma(\widehat{u},\widehat{\mathcal{F}})$. To see that Φ is well-defined, observe that since $\mathcal{E}$ has anticipation at most t, there cannot be two distinct edges $v(u,\mathcal{F}) \overset{a}{\to} \widehat{v}$ and $v(u,\mathcal{F}) \overset{a}{\to} \widehat{v}'$ with $\widehat{v}$ and $\widehat{v}'$ both belonging to the same $\Gamma(\widehat{u},\widehat{\mathcal{F}})$. To see that Φ is onto, first consider an outgoing edge $v(u,\mathcal{F}) \overset{a}{\to} \widehat{v}$ from $v(u,\mathcal{F})$, and note that since $\mathcal{F} \subseteq \mathcal{F}_G(u)$, there is in G an outgoing edge $u \overset{a}{\to} \widehat{u}$ for some $\widehat{u}$. Let $\widehat{\mathcal{F}} = \mathcal{F}^t_{\mathcal{E}}(\widehat{v})$. We claim that $\widehat{v} \in \Gamma(\widehat{u},\widehat{\mathcal{F}})$. Of course $\mathcal{F}^t_{\mathcal{E}}(\widehat{v}) = \widehat{\mathcal{F}}$; and since $\mathcal{F}_{\mathcal{E}}(v(u,\mathcal{F})) \subseteq \mathcal{F}_G(u)$ and G is deterministic, $\mathcal{F}_{\mathcal{E}}(\widehat{v}) \subseteq \mathcal{F}_G(\widehat{u})$. Thus, by definition of $\mathcal{E}'$ there is an edge $(u,\mathcal{F}) \overset{a}{\to} (\widehat{u},\widehat{\mathcal{F}})$. We thus conclude that Φ is onto. Since u and a determine $\widehat{u}$ and since $\widehat{v}$ determines $\widehat{\mathcal{F}}$, it follows that Φ is 1–1. This completes the proof of (a).

Part (b): By definition of $\mathcal{E}'$, we see that whenever there is a path $(u_0,\mathcal{F}) \overset{w_0}{\to} (u_1,\mathcal{F}) \overset{w_1}{\to} (u_2,\mathcal{F}) \to \ldots \to (u_{\ell-1},\mathcal{F}) \overset{w_{\ell-1}}{\to} (u_\ell,\mathcal{F})$ in $\mathcal{E}'$, there is also a path $u_0 \overset{w_0}{\to} u_1 \overset{w_1}{\to} \ldots \to u_{\ell-1} \overset{w_{\ell-1}}{\to} u_\ell$ in G. Thus $S(\mathcal{E}') \subseteq S(G) = S$, as desired.

Part (c): We must show that the initial edge of any path γ of length $t+1$ in $\mathcal{E}'$ is determined by its label $w_0 w_1 \ldots w_t$ and its initial state $(u,\mathcal{F})$. Write the initial edge of γ as:

$(u, \mathcal{F}) \stackrel{w_0}{\to} (\hat{u}, \widehat{\mathcal{F}})$. By Lemma 7.7, there is a path in $\mathcal{E}$ with label $w_0 w_1 \ldots w_t$ that begins at state $v(u, \mathcal{F})$. Since $\mathcal{E}$ has anticipation at most t, the label sequence $w_0 w_1 \ldots w_t$ and $v(u, \mathcal{F})$ determine the initial edge $v(u, \mathcal{F}) \stackrel{w_0}{\to} \hat{v}$ of this path. So, it suffices to show that u, w_0, and $\hat{v}$ determine $\hat{u}$ and $\widehat{\mathcal{F}}$; for then $(u, \mathcal{F})$ and $w_0 w_1 \ldots w_t$ will determine the initial edge of γ.

Indeed, by definition of $\mathcal{E}'$, there must be an edge $u \stackrel{w_0}{\to} \hat{u}$ in G such that $\hat{v} \in \Gamma(\hat{u}, \widehat{\mathcal{F}})$. Since G is deterministic, u and w_0 determine $\hat{u}$. Furthermore, for any fixed $\hat{u}$, the sets $\Gamma(\hat{u}, \mathcal{G})$ are disjoint for distinct $\mathcal{G}$, and, so, $\hat{v}$ determines $\widehat{\mathcal{F}}$. It follows that u, w_0, and $\hat{v}$ determine $\hat{u}$ and $\widehat{\mathcal{F}}$, as desired, thus proving (c). □

Now, for every state $u \in V_G$, the number of distinct nonempty subsets $\Gamma(u, \mathcal{F})$ is bounded from above by $2^{N(u,t)} - 1$. This yields the desired upper bound of Theorem 7.6 on the number of states of $\mathcal{E}'$.

It follows by Theorem 7.6 that in order to verify whether there exists an (S, n)-encoder with anticipation t, we can exhaustively check all irreducible graphs $\mathcal{E}$ with labeling over $\Sigma(S)$, with out-degree n, and with number of states $|V_{\mathcal{E}}|$ which is at most the bound of Theorem 7.6. Checking that such a labeled graph $\mathcal{E}$ is an (S, n)-encoder can be done by the following finite procedure: Construct the determinizing graph H of $\mathcal{E}$ as in Section 2.2.1. Since $\mathcal{E}$ is irreducible, the states of any irreducible sink H' of H, as subsets of $V_{\mathcal{E}}$, must contain all the states of $\mathcal{E}$. Hence, we must have $S(\mathcal{E}) = S(H')$. Then, we verify that $S(G * H') = S(H')$; to this end, it suffices, by Lemma 2.9, to check that $S(H')$ is presented by an irreducible (deterministic) component G' of $G * H'$. The equality $S(H') = S(G')$, in turn, can be checked by Theorem 2.12, using the Moore algorithm of Section 2.6.2.

Finally, testing whether $\mathcal{E}$ has anticipation $\leq t$ can be done by the efficient algorithm described in Section 2.7.2.

By the Moore co-form construction of Section 2.2.7, the existence of such an encoder implies the existence of an (S, n)-encoder with anticipation *exactly* t.

7.4.2 Upper bounds on the anticipation

Continuing the discussion of Section 7.4.1, we now obtain more tractable upper and lower bounds on the smallest attainable anticipation of (S, n)-encoders in terms of n and a deterministic presentation of the constrained system S.

Let G be a deterministic presentation of S. The anticipation of an encoder obtained by the state-splitting algorithm [ACH83] is bounded from above by the number of splitting rounds. This, in turn, yields the following result, which is, so far, the best general upper bound known for the anticipation obtained by direct application of the state-splitting algorithm.

Theorem 7.9 *Let S be a constrained system presented by a deterministic graph G and let n be a positive integer. Assume that* $\mathsf{cap}(S) \geq \log n$. *Then, there exists an (S, n)-encoder*

$\mathcal{E}$, obtained by the state-splitting algorithm, such that,

$$\mathcal{A}(\mathcal{E}) \le \min_{\mathbf{x} \in \mathcal{X}(A_G,n)} \Big\{ \|\mathbf{x}\|_1 - \mathsf{w}(\mathbf{x}) \Big\} ,$$

where $\mathsf{w}(\mathbf{x})$ is the number of nonzero components in $\mathbf{x}$. Furthermore, if G has finite memory, then $\mathcal{E}$ is $(\mathcal{M}(G), \mathcal{A}(\mathcal{E}))$-definite.

This bound is quite poor, since it may be exponential in $|V_G|$, as is, indeed, the case for the constrained systems of Example 7.1. On the other hand, if G has finite memory, then the encoder $\mathcal{E}$ obtained by the state-splitting algorithm is guaranteed to be definite.

Now, suppose that G^t can be split fully in one round; that is, the splitting yields a labeled graph $\mathcal{E}_1$ with out-degree $\ge n^t$ at each state. By deleting excess edges, $\mathcal{E}_1$ can be made an (S^t, n^t)-encoder $\mathcal{E}_2$ with anticipation 1 over $\Sigma(S^t)$. Let $\mathcal{E}_3$ be the Moore co-form of $\mathcal{E}_2$ as in Section 2.2.7. Then $\mathcal{E}_3$ is an (S^t, n^t)-encoder with anticipation 2. If we replace the n^t outgoing edges from each state in $\mathcal{E}_3$ by an n-ary tree of depth t, we obtain an (S, n)-encoder $\mathcal{E}_4$ with anticipation $\le 3t-1$. Therefore, we have the following.

Theorem 7.10 *Let S be a constrained system presented by a deterministic graph G and let n and t be positive integers. Suppose that G^t can be split in one round, yielding a labeled graph with minimum out-degree at least n^t. Then, there is an (S, n)-encoder with anticipation $\le 3t-1$.*

In [Ash87b] and [Ash88], Ashley shows that for $t = O(|V_G|)$, G^t can be split in one round, yielding a labeled graph with minimum out-degree at least n^t; moreover, the splitting can be chosen to be $\mathbf{x}$-consistent with respect to *any* (A_G, n)-approximate eigenvector $\mathbf{x}$. This provides encoders with anticipation which is at most *linear* in $|V_G|$. The following theorem is a statement of Ashley's result.

Theorem 7.11 *Let S be a constrained system presented by a deterministic graph G and let n be a positive integer. Assume that $\mathsf{cap}(S) \ge \log n$. Then, there exists an (S, n)-encoder $\mathcal{E}$ such that,*

(a) when $n = \lambda(A_G)$,

$$\mathcal{A}(\mathcal{E}) \le 9|V_G| + 6\lceil \log_n |V_G| \rceil - 1 ;$$

(b) when $n < \lambda(A_G)$,

$$\mathcal{A}(\mathcal{E}) \le 15|V_G| + 3\lceil \log_n |V_G| \rceil - 1 .$$

Note, however, that the encoders obtained by splitting the tth power of G are typically not sliding-block decodable when $t > 1$, even when G has finite memory.

A further improvement on the upper bound of the smallest attainable anticipation is presented in [AMR95], using the stethering method which, in turn, is based on an earlier result by Adler, Goodwyn, and Weiss [AGW77] (see Chapter 6). The following result applies to the case where $n \le \lambda(A_G) - 1$.

Theorem 7.12 *Let S be a constrained system presented by a deterministic graph G and let n be a positive integer $\le \lambda(A_G) - 1$. Then, there is an (S, n)-encoder $\mathcal{E}$, obtained by the (punctured) stethering method, such that*

$$\mathcal{A}(\mathcal{E}) \le 1 + \min_{\mathbf{x} \in \mathcal{X}(A_G, n+1)} \left\{ \lceil \log_{n+1} \|\mathbf{x}\|_\infty \rceil \right\} .$$

Furthermore, if G has finite memory, then $\mathcal{E}$ is $(\mathcal{M}(G), \mathcal{A}(\mathcal{E}))$-definite, and hence any tagged (S, n)-encoder based on $\mathcal{E}$ is $(\mathcal{M}(G), \mathcal{A}(\mathcal{E}))$-sliding-block decodable.

In particular, when $n \le \lambda(A_G) - 1$, there always exists an $(A_G, n+1)$-approximate eigenvector $\mathbf{x}$ such that $\|\mathbf{x}\|_\infty \le (n+1)^{2|V_G|}$ [Ash87a], [Ash88]. Hence, we have the following.

Corollary 7.13 *Let S be a constrained system presented by a deterministic graph G and let n be a positive integer $\le \lambda(A_G) - 1$. Then, there is an (S, n)-encoder $\mathcal{E}$, obtained by the (punctured) stethering method, such that*

$$\mathcal{A}(\mathcal{E}) \le 2|V_G| + 1 .$$

Furthermore, if G has finite memory, then $\mathcal{E}$ is $(\mathcal{M}(G), 2|V_G|+1)$-definite, and hence any tagged (S, n)-encoder based on $\mathcal{E}$ is $(\mathcal{M}(G), 2|V_G|+1)$-sliding block decodable.

In terms of rate $p : q$ finite-state encoders, the requirement $n \le \lambda(A_G) - 1$ is implied by

$$\frac{p}{q} \le \mathsf{cap}(S) - \frac{1}{2^p q \log_e 2} ;$$

namely, we need a margin between the rate and capacity which decreases exponentially with p.

Applying the stethering method on a power of G, the following result is obtained in [AMR95].

Theorem 7.14 *Let S be a constrained system presented by a deterministic graph G and let n be a positive integer <u>smaller</u> than $\lambda(A_G)$. Then, there is an (S, n)-encoder $\mathcal{E}$ such that*

$$\mathcal{A}(\mathcal{E}) \le 12|V_G| - 1 .$$

Theorem 7.14 improves on Theorem 7.11, but it does not cover the case $n = \lambda(A_G)$. Also, the encoders guaranteed by Theorem 7.14 are typically not sliding-block decodable.

7.4.3 Lower bounds on the anticipation

The next theorem, taken from [MR91], provides a lower bound on the anticipation of any (S, n)-encoder. A special case of this bound appears in [Fra89].

Theorem 7.15 *Let S be a constrained system presented by a deterministic graph G and let n be a positive integer. Assume that $\mathsf{cap}(S) \geq \log n$. Then, for any (S, n)-encoder $\mathcal{E}$,*

$$\mathcal{A}(\mathcal{E}) \geq \min_{\mathbf{x} \in \mathcal{X}(A_G, n)} \left\{ \log_n \|\mathbf{x}\|_\infty \right\} .$$

Proof. The theorem trivially holds if $\mathcal{A}(\mathcal{E}) = \infty$, so we assume that $\mathcal{E}$ has finite anticipation $\mathcal{A}$. Let $\mathbf{x} = \mathbf{x}(\mathcal{E}) = (x_u)_{u \in V_G}$ be as in the proof of Theorem 7.2. We recall that by the way $\mathbf{x}$ was constructed, each nonzero component of $\mathbf{x}$ is a size of some subset $Z = T_{\mathcal{E}}(\mathbf{w}, v)$ of states in $\mathcal{E}$ which are accessible from $v \in V_{\mathcal{E}}$ by paths labeled $\mathbf{w}$.

Let $T_{\mathcal{E}}(\mathbf{w}, v)$ be such a subset whose size equals the largest component of $\mathbf{x}$ and let $\ell = \ell(\mathbf{w})$ (i.e., the length of $\mathbf{w}$). Since the out-degree of $\mathcal{E}$ is n, we have n^ℓ paths of length ℓ starting at v in $\mathcal{E}$ and, so,

$$n^\ell \geq |T_{\mathcal{E}}(\mathbf{w}, v)| = \max_{u \in V_G} x_u \ ,$$

implying

$$\ell \geq \min_{(y_u)_u \in \mathcal{X}(A_G, n)} \left\{ \log_n \left(\max\nolimits_{u \in V_G} y_u \right) \right\} .$$

Therefore, when $\mathcal{A} \geq \ell$, we are done.

Assume now that $\ell > \mathcal{A}$. Since $\mathcal{E}$ has finite anticipation $\mathcal{A}$, the first $\ell - \mathcal{A}$ edges of any path in $\mathcal{E}$ labeled $\mathbf{w}$ are uniquely determined, once we know the initial state v. It thus follows that the paths from v to $T_{\mathcal{E}}(\mathbf{w}, v)$ labeled $\mathbf{w}$ may differ only in their last $\mathcal{A}$ edges. Hence, we can have at most $n^{\mathcal{A}}$ such paths. Recalling that the number of such paths is $|T_{\mathcal{E}}(\mathbf{w}, v)|$, we have

$$n^{\mathcal{A}} \geq |T_{\mathcal{E}}(\mathbf{w}, v)| = \max_{u \in V_G} x_u \geq \min_{(y_u)_u \in \mathcal{X}(A_G, n)} \max_{u \in V_G} y_u \ ,$$

as claimed. □

There is some resemblance between the lower bound of Theorem 7.15 and the upper bound of Theorem 7.12. And, indeed, there are many cases where the difference between these bounds is at most 1. Note, however, that for the constrained systems $S_k = S(G_k)$ of Example 7.1, we obtain, by Theorem 7.12, an upper bound of $1 + \frac{1}{2}|V_{G_k}|$ on the smallest anticipation of any (S_k, r)-encoder, where the lower bound of Theorem 7.15 equals 2. In fact, this lower bound is tight [AMR95].

The following bound, proved in [AMR96] is, in a way, a converse of Theorem 7.10.

Theorem 7.16 *Let S be an irreducible constrained system presented by an irreducible deterministic graph G and let n and t be positive integers. If there is an (S, n)-encoder with anticipation t, then G^t can be split in one round, yielding a graph with minimum out-degree at least n^t.*

Proof. Let $\mathcal{E}$ be an (S, n)-encoder with anticipation t, let $\Sigma = \Sigma(S)$, and let H be the determinizing graph constructed from $\mathcal{E}$ as in Section 2.2.1. Recall that each state $Z \in V_H$ is a subset $T_{\mathcal{E}}(\mathbf{w}, v)$ of states of $\mathcal{E}$ that can be reached in $\mathcal{E}$ from a given state $v \in V_{\mathcal{E}}$ by paths that generate a given word $\mathbf{w}$. Let H' be an irreducible sink of H and let $\mathbf{x} = (x_u)_{u \in V_G}$ be the nonnegative integer vector defined in the proof of Theorem 7.2:

$$x_u = \max \{ |Z| \; : \; Z \in V_{H'} \text{ and } \mathcal{F}_{H'}(Z) \subseteq \mathcal{F}_G(u) \} \; , \quad u \in V_G \; ;$$

in case there is no state $Z \in V_{H'}$ such that $\mathcal{F}_{H'}(Z) \subseteq \mathcal{F}_G(u)$, define $x_u = 0$. Then, as in the proof of Theorem 7.2, $\mathbf{x}$ is an (A_G, n)-approximate eigenvector.

Let $Z = T_{\mathcal{E}}(\mathbf{w}, v)$ be a state in H' and suppose that Z contains two distinct states, z and z', of $\mathcal{E}$. First, we claim that there is no word $\mathbf{w}'$ of length t that can be generated in $\mathcal{E}$ from both z and z'. Otherwise, we would have in $\mathcal{E}$ two paths of length $\ell(\mathbf{w}) + t$, starting at the same state v, with the same labeling $\mathbf{w}\mathbf{w}'$, that do not agree in at least one of their first $\ell(\mathbf{w})$ edges. This, however, contradicts the fact that $\mathcal{E}$ has anticipation t.

For $\mathbf{w}' \in \mathcal{F}_{H'}^t(Z)$, denote by $Z_{\mathbf{w}'}$ the terminal state in H' of a path labeled $\mathbf{w}'$ starting at Z. As we have just shown, a word $\mathbf{w}' \in \mathcal{F}_{H'}^t(Z)$ can be generated in $\mathcal{E}$ from exactly one state $z \in Z$. Therefore, the sets $\mathcal{F}_{\mathcal{E}}^t(z)$, $z \in Z$, form a partition of $\mathcal{F}_{H'}^t(Z)$. Furthermore, by the losslessness of $\mathcal{E}$, the number of paths in $\mathcal{E}$ that start at $z \in Z$ and generate $\mathbf{w}' \in \mathcal{F}_{\mathcal{E}}^t(u)$ equals $|T_{\mathcal{E}}(\mathbf{w}', z)| = |Z_{\mathbf{w}'}|$. Since $\mathcal{E}$ is an (S, n)-encoder, we conclude:

$$\sum_{\mathbf{w}' \in \mathcal{F}_{\mathcal{E}}^t(z)} |Z_{\mathbf{w}'}| = n^t \quad \text{for every } z \in Z \; . \tag{7.1}$$

For each state $u \in V_G$ such that $x_u \neq 0$, select some $Z = Z(u) \in V_{H'}$ such that $|Z| = x_u$ and $\mathcal{F}_{H'}(Z) \subseteq \mathcal{F}_G(u)$. Now, the partition $\{\mathcal{F}_{\mathcal{E}}^t(z) : z \in Z\}$ of $\mathcal{F}_{H'}^t(Z)$ may be regarded as a partition of $\mathcal{F}_G^t(u)$ by appending the complement $\mathcal{F}_G^t(u) \setminus \mathcal{F}_{H'}^t(Z)$ to one of the atoms $\mathcal{F}_{\mathcal{E}}^t(z)$, $z \in Z$. Since G^t is deterministic, this defines a partition $P_{G^t}(u) = \{E_{G^t}(u, z)\}_{z \in Z(u)}$ of the outgoing edges from u in G^t into $|Z(u)| = x_u$ atoms. For $\mathbf{w}' \in \mathcal{F}_{\mathcal{E}}(z)$, let u' denote the terminal state of the edge in G^t that begins at u and is labeled $\mathbf{w}'$. Now, if $\mathbf{w}'' \in \mathcal{F}_{H'}(Z_{\mathbf{w}'})$, then $\mathbf{w}'\mathbf{w}'' \in \mathcal{F}_{H'}(Z) \subseteq \mathcal{F}_G(u)$. Since G is deterministic, this implies that $\mathbf{w}'' \in \mathcal{F}_G(u')$. Thus $\mathcal{F}_{H'}(Z_{\mathbf{w}'}) \subseteq \mathcal{F}_G(u')$ and, so, $|Z_{\mathbf{w}'}| \leq x_{u'}$. This, together with Equation (7.1), shows that the splitting of G^t defined by the partition $P_{G^t}(u)$ satisfies the following inequality:

$$\sum_{e \in E_{G^t}(u,z)} x_{\tau(e)} \geq n^t \quad \text{for every } u \in V_G \text{ and } z \in Z(u) \; .$$

Hence, the split graph has minimum out-degree at least n^t. □

Theorem 7.16 may be regarded as a lower bound on the anticipation of an encoder. This result, together with Theorem 7.10, shows that by one round of splitting of some power of G, one can obtain an encoder whose anticipation is within a constant factor from the smallest anticipation possible.

There are examples which show that neither of the lower bounds in Theorems 7.15, 7.16 implies the other. On the other hand, when $\mathsf{cap}(S) = \log n$, we claim that for irreducible constrained systems the lower bound of Theorem 7.16 implies that of Theorem 7.15. Indeed, let t denote the bound of Theorem 7.16. Then for each state $u \in V_G$ there is a partition $\{E_{G^t}(u,i)\}_{i=1}^{x_u}$ of the outgoing edges from u in G^t such that the vector $\mathbf{x} = (x_u)_u$ is a positive (A_G, n)-approximate eigenvector and

$$\sum_{e \in E_{G^t}(u,i)} x_{t(e)} \geq n^t \quad \text{for each } (u,i) \; . \tag{7.2}$$

We now claim that (7.2) holds in our case with equality for every (u,i). Otherwise, the corresponding splitting would yield an irreducible, lossless presentation of S^t with minimum out-degree at least n^t and at least one state with out-degree greater than n^t—contradicting the equality $\mathsf{cap}(S) = \log n$.

Let $u_{\max}$ be a state in G for which $x_{u_{\max}} = \|\mathbf{x}\|_\infty$. Also, let v be a state with an outgoing edge, in G^t, to $u_{\max}$. Then any edge e from v to $u_{\max}$ in G^t belongs to some $E_{G^t}(v,i)$ and so the equality $\sum_{e \in E_{G^t}(v,i)} x_{t(e)} = n^t$ implies

$$x_{u_{\max}} \leq n^t$$

—i.e., $t \geq \log_n \|\mathbf{x}\|_\infty \geq \min_{\mathbf{y} \in \mathcal{X}(A_G,n)} \Big\{ \log_n \|\mathbf{y}\|_\infty \Big\}$.

We end this section by mentioning without proof the improvements on Theorems 7.15 and 7.16 that have been obtained in [Ru96] and [RuR01].

Recall that Theorem 5.10 in Section 5.6.2 provides a necessary and sufficient condition for having (S,n)-encoders with anticipation t Such a characterization also implies a lower bound on the anticipation of (S,n)-encoders: given S and n, the anticipation any (S,n)-encoder is at least the smallest nonnegative integer t for which there exists a presentation G of S and an (A_G,n)-approximate eigenvector $\mathbf{x}$ that satisfy conditions (a)–(e) of Theorem 5.10.

The following is another result obtained in [Ru96] and [RuR01].

Theorem 7.17 *Let S be an irreducible constraint, let n be a positive integer where $\mathsf{cap}(S) \geq \log n$, and let G be any irreducible deterministic presentation of S. Suppose there exists some irreducible (S,n)-encoder with anticipation $t < \infty$. Then there exists an (A_G,n)-approximate eigenvector $\mathbf{x}$ such that the following holds:*

(a) $\|\mathbf{x}\|_\infty \leq n^t$.

(b) For every $k = 1, 2, \ldots, t$, the states of G^k can be split in one round consistently with the (A_G^k, n^k)-approximate eigenvector $\mathbf{x}$, such that the induced approximate eigenvector $\mathbf{x}'$ satisfies $\|\mathbf{x}'\|_\infty \le n^{t-k}$, and each of the states in G^k is split into no more than n^k states.

While Theorem 5.10 gives a necessary and sufficient condition on the existence of (S, n)-encoders with a given anticipation t, Theorem 7.17 gives only a *necessary* condition on the existence of such encoders. On the other hand, Theorem 7.17 allows to obtain a lower bound on the anticipation using *any* irreducible deterministic presentation of S—in particular the Shannon cover of S. Therefore, it will typically be easier to compute bounds using Theorem 7.17.

Note that Theorem 7.15 is equivalent to Theorem 7.17(a), while Theorem 7.16 is equivalent to Theorem 7.17(b) for the special case $k = t$. Examples in [RuR01] show that Theorem 7.17 (and hence Theorem 5.10) yields stronger bounds than these two former results.

The results in [RuR01] also imply tight lower bounds in certain practical cases. For example, it is shown therein that any rate $2 : 3$ finite-state encoder for the $(1, 7)$-RLL constraint must have anticipation at least 2, and the Weathers-Wolf encoder in 4.6 does attain this bound (and so does the encoder of Adler Coppersmith, and Hassner in [ACH83]). Similarly, any rate $1 : 2$ encoder for the $(2, 7)$-RLL constraint must have anticipation at least 3, and this bound is attained by the Franaszek encoder in Figure 4.4. A lower bound of 3 applies also to the anticipation of any rate $2 : 5$ encoder for the $(2, 18, 2)$-RLL constraint (see Figure 1.12); this bound is tight due to the constructions by Weigandt [Weig88] and Hollmann [Holl95].

7.5 Sliding-block decodability

The following is the analog of Theorem 7.6 for sliding-block decodable encoders. The special case of block decodable encoders was treated in Section 4.4.

Theorem 7.18 *Let S be an irreducible constrained system with a Shannon cover G, and let n be a positive integer and m and a be nonnegative integers. For every state u in G, let $N(u, \mathsf{a}) = |\mathcal{F}_G^{\mathsf{a}}(u)|$ and let $P(u, \mathsf{m})$ be the number of words of length m that can be generated in G by paths that terminate in state u. If there exists an (m, a)-sliding-block decodable (S, n)-encoder, then there exists such an encoder with at most $\sum_{u \in V_G} P(u, \mathsf{m})(2^{N(u,\mathsf{a})} - 1)$ states.*

Proof. The proof is similar to that of Theorem 7.6. In fact, that proof applies almost verbatim to the case of $(0, \mathsf{a})$-sliding-block decodable encoders, so we assume here that m is strictly positive. Let $\mathcal{E}$ be an irreducible (m, a)-sliding-block decodable (S, n)-encoder. Also, let H' be an irreducible sink of the determinizing graph of $\mathcal{E}$ obtained by the construction of Section 2.2.1. By construction of H', for every path from state v to state $\hat{v}$ in $\mathcal{E}$ that generates a word $\mathbf{w}$, there is a path in H' that generates $\mathbf{w}$, starting at a state Z and terminating in a state $\hat{Z}$, such that $v \in Z$ and $\hat{v} \in \hat{Z}$. Hence, by Lemma 2.13, there also exists a path in G that generates $\mathbf{w}$, starting at a state u and terminating in a state $\hat{u}$, such that $\mathcal{F}_{\mathcal{E}}(v) \subseteq \mathcal{F}_G(u)$ and $\mathcal{F}_{\mathcal{E}}(\hat{v}) \subseteq \mathcal{F}_G(\hat{u})$. It thus follows that for every state $\hat{v}$ in $\mathcal{E}$ and a word $\mathbf{w}$ that can be generated in $\mathcal{E}$ by a path terminating in $\hat{v}$, there is a path in G that generates $\mathbf{w}$ whose terminal state, $\hat{u}$, satisfies $\mathcal{F}_{\mathcal{E}}(\hat{v}) \subseteq \mathcal{F}_G(\hat{u})$.

For a state $u \in V_G$, a word $\mathbf{w}$ of length m that can be generated in G by a path terminating in u, and a nonempty subset $\mathcal{F} \subseteq \mathcal{F}_G^{\mathsf{a}}(u)$, we define $\Gamma(u, \mathbf{w}, \mathcal{F})$ to be the set of all states v in $\mathcal{E}$ which are terminal states of paths in $\mathcal{E}$ that generate $\mathbf{w}$ and such that $\mathcal{F}_{\mathcal{E}}(v) \subseteq \mathcal{F}_G(u)$ and $\mathcal{F}_{\mathcal{E}}^{\mathsf{a}}(v) = \mathcal{F}$. Note that each state of $\mathcal{E}$ is contained in some set $\Gamma(u, \mathbf{w}, \mathcal{F})$ and, so, at least one such set is nonempty.

A tagged (S, n)-encoder $\mathcal{E}'$ is now defined as follows. In each nonempty set $\Gamma(u, \mathbf{w}, \mathcal{F})$, we designate a state of $\mathcal{E}$ and call it $v(u, \mathbf{w}, \mathcal{F})$. The states of $\mathcal{E}'$ are triples $(u, \mathbf{w}, \mathcal{F})$ for which $\Gamma(u, \mathbf{w}, \mathcal{F})$ is nonempty.

Let u and $\hat{u}$ be states in G and let $\mathbf{w} = w_1 w_2 \ldots w_{\mathsf{m}}$ and $\hat{\mathbf{w}} = \hat{w}_1 \hat{w}_2 \ldots \hat{w}_{\mathsf{m}}$ be two words that can be generated by paths in G that terminate in u and $\hat{u}$, respectively. If $\Gamma(u, \mathbf{w}, \mathcal{F})$ and $\Gamma(\hat{u}, \hat{\mathbf{w}}, \hat{\mathcal{F}})$ are nonempty, then we draw a tagged edge $(u, \mathbf{w}, \mathcal{F}) \stackrel{s/b}{\to} (\hat{u}, \hat{\mathbf{w}}, \hat{\mathcal{F}})$ in $\mathcal{E}'$ if and only if the following four conditions hold:

(a) $\hat{w}_j = w_{j+1}$ for $j = 1, 2, \ldots, \mathsf{m}-1$;

(b) $b = \hat{w}_{\mathsf{m}}$;

(c) there is a tagged edge $v(u, \mathbf{w}, \mathcal{F}) \stackrel{s/b}{\to} \hat{v}$ in $\mathcal{E}$ for some $\hat{v} \in \Gamma(\hat{u}, \hat{\mathbf{w}}, \hat{\mathcal{F}})$;

(d) there is an edge $u \stackrel{b}{\to} \hat{u}$ in G.

By the proof of Theorem 7.6, it follows that $\mathcal{E}'$ is, indeed, an (S, n)-encoder and that $\mathcal{F}_{\mathcal{E}'}^{\mathsf{a}+1}\big((u, \mathbf{w}, \mathcal{F})\big) = \mathcal{F}_{\mathcal{E}}^{\mathsf{a}+1}\big(v(u, \mathbf{w}, \mathcal{F})\big)$. Furthermore, it can be shown by induction that, for every $\ell \le \mathsf{m}$, the paths of length ℓ in $\mathcal{E}'$ that terminate in $(u, w_1 w_2 \ldots w_{\mathsf{m}}, \mathcal{F})$, all have the same labeling $w_{\mathsf{m}-\ell+1} w_{\mathsf{m}-\ell+2} \ldots w_{\mathsf{m}}$. Since the 'outgoing picture'—including tagging—from state $(u, \mathbf{w}, \mathcal{F})$ in $\mathcal{E}'$ is the same as that from state $v(u, \mathbf{w}, \mathcal{F})$ in $\mathcal{E}$, it follows that $\mathcal{E}'$ is (m, a)-sliding-block decodable.

The upper bound on $|V_{\mathcal{E}'}|$ is now obtained by counting the number of distinct states $(u, \mathbf{w}, \mathcal{F})$. □

The upper bound on the number of states in Theorem 7.18 is doubly-exponential in the decoding look-ahead a. In [AM95], a stronger result is obtained where the upper bound on the number of states is singly-exponential. It is still open whether such an improvement is possible also for the doubly-exponential bound of Theorem 7.6.

The following bound is easily verified.

Proposition 7.19 *Let $\mathcal{E}$ be an irreducible (m, a)-sliding-block decodable encoder. Then,* $\mathsf{a} \geq \mathcal{A}(\mathcal{E})$.

Hence, we can apply the lower bounds on the anticipation which were presented in Section 7.4.3, to obtain lower bounds on the attainable look-ahead of sliding-block decodable encoders (but these do not give lower bounds on the decoding window length, $\mathsf{m} + \mathsf{a} + 1$, since m may be negative). On the other hand, Theorems 7.9 and 7.12 and Corollary 7.13 provide upper bounds on the look-ahead of encoders obtained by constructions that yield sliding-block decodable encoders for finite-type constrained systems.

We remark that Theorem 7.18 implies upper bounds on the the size of encoders which are sliding-block decodable also when m is negative: simply apply the theorem with $\mathsf{m} = 0$.

And finally we note that, at least for finite-type constrained systems, Hollmann [Holl96] has given a procedure for deciding if there exists a sliding block decodable (S, n)-encoder with a given window length; here, the window length $L = \mathsf{m} + \mathsf{a} + 1$, rather than m and a, is specified. Even for $L = 1$, this is a non-trivial problem, because one must consider the possibility that $\mathsf{a} = -\mathsf{m}$ may be arbitrarily large.

7.6 Gate complexity and time–space complexity

In this section, we discuss the gate complexity and time-space complexity of some of the encoding schemes that were mentioned in the previous sections. We start with the time-space complexity criterion, assuming that the encoders are to be implemented as a program on a random-access machine (RAM) [AHU74, Ch. 1]. The results on gate complexity will then follow by known results in complexity theory.

We define an *encoding scheme* as a function $(G, q, n) \mapsto \mathcal{E}(G, q, n)$, that maps a deterministic graph G and integers q and n into an $(S(G^q), n)$-encoder $\mathcal{E}(G, q, n)$. The state-splitting algorithm of [ACH83], the method described by Ashley in [Ash88], and the stethering method of [AMR95] are examples of encoding schemes.

For a given encoding scheme $(G, q, n) \mapsto \mathcal{E}(G, q, n)$, we can formalize the *encoding problem* as follows: We are to write an encoding program P on a RAM; an input instance to P consists of the following entries:

- a deterministic graph G over an alphabet Σ,
- an integer q,

- an integer $n \leq \lambda(A_G^q)$,
- a state u of $\mathcal{E}(G, q, n)$,
- an input tag $s \in \{0, 1, \ldots, n-1\}$.

For any input instance, the program P computes an output q-block over Σ and the next state of the tagged $(S(G^q), n)$-encoder $\mathcal{E}(G, q, n)$, given we are at state u in $\mathcal{E}(G, q, n)$ and the current input tag is s. Note that in order to perform its function, the program P does not necessarily have to generate the whole graph presentation of $\mathcal{E}(G, q, n)$.

We denote by $\mathrm{Poly}(\cdot)$ a fixed arbitrary multinomial, whose coefficients are absolute constants, independent of its arguments.

The following was proved in [AMR95] for the stethering coding scheme and for a variation of Ashley's construction [Ash88].

Theorem 7.20 *There exists an encoding scheme $(G, q, n) \mapsto \mathcal{E}(G, q, n)$ (such as the one presented in [Ash88] or [AMR95]), for which there is an encoding program P on a RAM that solves the encoding problem in time complexity which is at most* $\mathrm{Poly}\Big(|V_G|, q, \log|\Sigma|\Big)$.

In particular, if we now fix G, q, and n, we obtain an encoding program that simulates $\mathcal{E}(G, q, n)$ with a polynomial time and space complexity.

Theorem 7.20 applies to the constructions covered in Theorems 7.11, 7.12, and 7.14. In contrast, it is not known yet whether a polynomial encoder can be obtained by a direct application of the state-splitting algorithm.

For a positive integer ℓ, denote by I_ℓ the set of all possible inputs to P of size ℓ, according to some standard representation of the input. Now, if the time complexity of P on each element of I_ℓ is polynomial in ℓ, then for any input size ℓ, there exists a circuit C_ℓ with $\mathrm{Poly}(\ell) = \mathrm{Poly}\Big(|V_G|, q, \log|\Sigma|\Big)$ gates that implements P for inputs in I_ℓ. Furthermore, such circuits C_ℓ are 'uniform' in the sense that there is a program on a RAM that generates the layouts of C_ℓ in time complexity which is $\mathrm{Poly}(\ell)$. This is a consequence of a known result on the equivalence between polynomial circuit complexity and polynomial RAM complexity of decision problems [ST93, Theorem 2.3].

By Theorem 7.20, it thus follows that we can have such a polynomial circuit at hand for both Ashley's construction and the stethering construction. We summarize this in the following theorem.

Theorem 7.21 *For every constrained system S over an alphabet Σ, presented by a deterministic graph G, and for any positive integers q and $n \leq \lambda(A_G^q)$, there exists an (S^q, n)-encoder that can be implemented by a circuit consisting of* $\mathrm{Poly}\Big(|V_G|, q, \log|\Sigma|\Big)$ *gates and*

$O(|V_G| \log n)$ memory bit-cells. Furthermore, there exists a program on a RAM that generates the layout of such an implementation in polynomial-time.

Theorems 7.20 and 7.21 apply also to the decoding complexity of the corresponding encoders.

Problems

Problem 7.1 Prove Theorem 7.3 by modifying the end of the proof of Theorem 7.2.

Problem 7.2 Verify the assertion of Example 7.2.

Problem 7.3 Verify the assertion of Example 7.3.

Problem 7.4 Let S be the constrained system presented by the graph G in Figure 2.24. Is there a positive integer ℓ for which there exists a deterministic $(S^{2\ell}, 2^\ell)$-encoder? If yes, construct such an encoder; otherwise, explain.

Problem 7.5 Let S be the constrained system presented by the graph G in Figure 7.2.

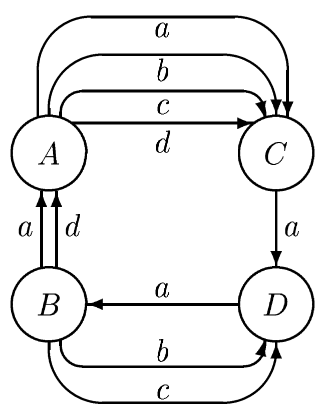

Figure 7.2: Graph G for Problem 7.5.

1. Compute the capacity of S.
2. Compute an $(A_G, 2)$-approximate eigenvector in which the largest entry is the smallest possible.
3. For every positive integer ℓ, determine the smallest anticipation of any $(S^\ell, 2^\ell)$-encoder.

Problem 7.6 Let S be the constrained system presented by the graph G in Figure 5.27.

1. Find the smallest anticipation possible of any $(S, 2)$-encoder.
2. Find the smallest number of states of any $(S^2, 4)$-encoder.