Small Hazard-Free Transducers

Johannes Bund ✉ 🆔
CISPA Helmholtz Center for Information Security, Saarbrücken, Germany

Christoph Lenzen ✉
CISPA Helmholtz Center for Information Security, Saarbrücken, Germany

Moti Medina ✉ 🆔
Faculty of Engineering, Bar-Ilan University, Ramat Gan, Israel

─── **Abstract** ───

Ikenmeyer et al. (JACM'19) proved an unconditional exponential separation between the hazard-free complexity and (standard) circuit complexity of explicit functions. This raises the question: which classes of functions permit efficient hazard-free circuits?

In this work, we prove that circuit implementations of *transducers* with small state space are such a class. A transducer is a finite state machine that transcribes, symbol by symbol, an input string of length n into an output string of length n. We present a construction that transforms any function arising from a transducer into an efficient circuit of size $\mathcal{O}(n)$ computing the *hazard-free extension* of the function. More precisely, given a transducer with s states, receiving n input symbols encoded by l bits, and computing n output symbols encoded by m bits, the transducer has a hazard-free circuit of size $2^{\mathcal{O}(s+\ell)}mn$ and depth $\mathcal{O}(s\log n + \ell)$; in particular, if $s, \ell, m \in \mathcal{O}(1)$, size and depth are asymptotically optimal. In light of the strong hardness results by Ikenmeyer et al. (JACM'19), we consider this a surprising result.

1 Introduction

Digital circuit design relies on a fundamental abstraction of the physical world. Electric voltages transmitted by wires are mapped to Boolean values, where high voltages correspond to 1 (true) and low voltages to 0 (false). By virtue of this abstraction, the behavior of digital circuitry can be described by Boolean formulas. However, this description does not account for the behavior of digital circuits in all cases: it offers no way of representing signals that are unstable, transitioning, oscillating, etc.

In this work, we study a classic extension of Boolean logic due to Kleene [?, §64], which allows for the presence of unspecified signals. In the following, we refer to the Boolean values $\mathbb{B} := \{0, 1\}$ as *stable,* while the additional third logical value $\mathfrak{u}$ is the *unstable* value. The resulting ternary set of logic values is denoted by $\mathbb{T} := \{0, 1, \mathfrak{u}\}$. Intuitively, $\mathfrak{u}$ may evaluate to any stable state at any point in the circuit, regardless of previous evaluations. Hence, we regard $\mathfrak{u}$ as a "superposition" of 0 and 1.

◼ **Table 1** Behaviour of basic gates `and`, `or`, and `not`.

and	0	1	ԉ		or	0	1	ԉ		not	
0	0	0	0		0	0	1	ԉ		0	1
1	0	1	ԉ		1	1	1	1		1	0
ԉ	0	ԉ	ԉ		ԉ	ԉ	1	ԉ		ԉ	ԉ

In Kleene logic, the basic gates[1] output a stable value if and only if the stable inputs already determine this output. The natural extension of the basic gates `and`, `or`, and `not` is given in Table 1:

By induction over the circuit structure, this defines for any circuit C consisting of such gates the function $C\colon \mathbb{T}^n \to \mathbb{T}^m$ it implements.

A key difference of Kleene logic to Boolean logic is that the laws of excluded middle and non-contradiction do not hold:

$$\mathtt{or}(x, \mathtt{not}(x)) \neq 1 \text{ and } \mathtt{and}(x, \mathtt{not}(x)) \neq 0,$$

as $\mathtt{not}(ԉ) = ԉ$ and $\mathtt{and}(ԉ, ԉ) = \mathtt{or}(ԉ, ԉ) = ԉ$. This drastically distinguishes $ԉ$ from an unknown Boolean value. Accounting for this limitation and the fact that constant stable inputs can easily be provided, we allow for constant-0 (and constant-1) gates in addition to the basic gates `or`, `and`, and `not`.

In general, digital logic cannot detect or prevent the propagation of unstable signals [?], matching the above limitation present in Kleene logic. On the other hand, it has been shown that CMOS logic gates implement the stated specification [?]. Furthermore, [?] extends the above model for combinational circuits to general clocked circuits. Thus, understanding combinational circuits in the model we use in this paper is equivalent to understanding worst-case propagation of unstable signals in general-purpose digital logic in a precise sense.

1.1 Hazards and Hazard-free Circuits

We strive for circuits that behave similarly to basic gates when receiving unstable inputs. That is, given all inputs, if changing an input bit from 0 to 1 has no effect on the output, then setting this input bit to $ԉ$ should also not affect the output. To formalize this concept, we make use of two operations. The first is the *superposition,* which results in the unstable value $ԉ$ whenever its inputs do not agree on a stable input value.

▶ **Definition 1** (Superposition). *Denote the* superposition *of two bits by the operator* $*\colon \mathbb{T} \times \mathbb{T} \to \mathbb{T}$. *For* $x, y \in \mathbb{T}$, $x * y = x$ *if* $x = y$ *and* $x * y = ԉ$ *otherwise. We extend the* $*$-operation *to* $x, y \in \mathbb{T}^n$ *by applying it at each bit* $i \in \{1, \ldots, n\}$, *such that*

$$(x * y)_i = \begin{cases} x_i & \text{if } x_i = y_i, \\ ԉ & \text{otherwise.} \end{cases}$$

The $*$-operation is associative and commutative, hence for $X \subseteq \mathbb{T}^n$ we define $*X \coloneqq x_1 * \ldots * x_n$, where $x_1, \ldots, x_n$ is an arbitrary enumeration of the elements of X. The second operation, called *resolution*, maps from ternary strings to sets of Boolean strings by replacing each $ԉ$ with both stable values.

---

[1] The specific choice of basic gates does not matter, see [?]; hence, we stick to `and`, `or`, and `not`.

▶ **Definition 2** (Resolution). *Denote the* resolution *by* $\mathrm{res}\colon \mathbb{T} \to \mathcal{P}(\mathbb{B})$. *For* $x \in \mathbb{T}$, $\mathrm{res}(x) = \{0, 1\}$ *if* $x = \mathfrak{u}$ *and* $\mathrm{res}(x) = \{x\}$ *otherwise. We extend this to bit strings of length n in the natural way, by setting for* $x \in \mathbb{T}^n$

$$\mathrm{res}(x) := \{y \in \mathbb{B}^n \mid \forall i \in \{1, \ldots, n\}\colon x_i \neq \mathfrak{u} \Rightarrow y_i = x_i\}.$$

For notational convenience, we extend all functions $f\colon X \to Y$ to sets of inputs $X' \subseteq X$, by applying them to each element of the input, i.e., $f(X') := \{f(x) \mid x \in X'\}$. For instance, $\mathrm{res}(\{0\mathfrak{u}0, 1\mathfrak{u}0\}) = \{000, 010, 100, 110\}$. First, we observe that superposition and resolution are not inverse functions.

▶ **Observation 3.** *Let* $X \subseteq \mathbb{T}^n$, *then taking the resolution of the superposition of all strings in X may add further strings, i.e.,* $X \subseteq \mathrm{res}(*\, X)$.

Note that a strict superset relation is possible. For instance, if there are two strings $x, y \in X$ that disagree on more than $\log(|X|)$ positions, then $|\mathrm{res}(*\, X)| > 2^{\log(|X|)} = |X|$ and hence $X \subset \mathrm{res}(*\, X)$. For example, if $X = \{101, 110\}$, then $\mathrm{res}(*\, X) = \mathrm{res}(1\mathfrak{u}\mathfrak{u}) = \{100, 101, 110, 111\}$.

Recall that `and` guarantees a stable output of 0 if at least one of its inputs is 0 – even if the other input is $\mathfrak{u}$. As a toy example, consider two circuits implementing a conjunction of x and y. First (needlessly involved) $\mathtt{and}(\mathtt{or}(x, \mathtt{not}(\mathtt{or}(y, \mathtt{not}(y)))), y)$ and second $\mathtt{and}(x, y)$. Under Boolean inputs, these expressions are equivalent, and so are the circuits. In contrast, for inputs $x = 0$ and $y = \mathfrak{u}$, we get that

$$\mathtt{and}(\mathtt{or}(0, \mathtt{not}(\mathtt{or}(\mathfrak{u}, \mathtt{not}(\mathfrak{u})))), \mathfrak{u}) = \mathfrak{u} \neq 0 = \mathtt{and}(0, \mathfrak{u}).$$

The first implementation has an unstable output, even though the stable input of 0 already determines that the output should be 0. This is referred to as a *hazard*.

A convenient formalization of this concept is as follows. For a circuit C with n inputs and m outputs, denote by $C(x) \in \mathbb{T}^m$ the output it computes on input $x \in \mathbb{T}^n$. We say that C *implements* the Boolean function $f\colon \mathbb{B}^n \to \mathbb{B}^m$, iff $C(x) = f(x)$ for all $x \in \mathbb{B}^n$. The desired behavior of the circuit is given by the *hazard-free extension* of f.

▶ **Definition 4** (Hazard-free Extensions). *For function* $f\colon \mathbb{B}^n \to \mathbb{B}^m$, *denote by* $f_{\mathfrak{u}}\colon \mathbb{T}^n \to \mathbb{T}^m$ *its* hazard-free extension, *which is defined by* $f_{\mathfrak{u}}(x) := *_{y \in \mathrm{res}(x)} f(y)$.

For any Boolean function f, there is a circuit implementing the hazard-free extension of f [**?**]. The hazard-free extension is the "most precise" extension of f computable by combinational logic. It is easy to show that $(f_{\mathfrak{u}})_i(x) = \mathfrak{u}$ entails that $C_i(x) = \mathfrak{u}$ for any circuit C implementing f: Restricted to Boolean inputs C and $f_{\mathfrak{u}}$ are identical. Changing input bits to $\mathfrak{u}$ can change the output bits to $\mathfrak{u}$ only, and if $(f_{\mathfrak{u}})(x)$ is $\mathfrak{u}$ at position i then $C(x)$ is $\mathfrak{u}$ at position i. In contrast, C has a *hazard* at $x \in \mathbb{T}^n$ iff it deviates from the desired behavior.

▶ **Definition 5** (k-Bit Hazards). *Circuit C implementing* $f\colon \mathbb{B}^n \to \mathbb{B}^n$ *has a* hazard *at* $x \in \mathbb{T}^n$ *iff* $C(x) \neq f_{\mathfrak{u}}(x)$. *If, for some* $k \in \mathbb{N}$, *x contains at most k many $\mathfrak{u}$s and C has a hazard on x, it is called a* k-bit *hazard. C is* k-bit hazard-free *if it has no k-bit hazards, i.e., if for every $x \in \mathbb{T}^n$ where $\mathfrak{u}$ appears at most k times,* $C(x) = f_{\mathfrak{u}}(x)$.

Note that C having a hazard at x is equivalent to $C(x)$ containing more $\mathfrak{u}$s than necessary. The smallest non-trivial example for a hazardous circuit is given by a naive implementation of a multiplexer circuit. A multiplexer has the Boolean specification $\mathtt{MUX}(a, b, s) = a$ if $s = 0$ and $\mathtt{MUX}(a, b, s) = b$ if $s = 1$. It can be implemented by the circuit corresponding to the Boolean formula $\mathtt{or}(\mathtt{and}(a, \mathtt{not}(s)), \mathtt{and}(b, s))$, resulting in a hazard at $(1, 1, \mathfrak{u})$, cf. [**?**].

Ikenmeyer et al. proved unconditional lower bounds on the complexity of hazard-free circuits implementing explicit functions [**?**]. More precisely, they show exponential gaps between the size of several Boolean circuits and their hazard-free counterparts. Furthermore, they show that hazard-free verification circuits for NP-hard problems cannot be of polynomial size unless the circuit equivalent of $P = NP$ holds. On the other hand, there are efficient implementations of sorting networks that avoid certain hazards [**?**], showing that hazard-free implementations do not always come at a high cost. This leads to the following question:

Which classes of Boolean functions allow for an efficient hazard-free implementation?

The key ingredient to the result from [**?**] is a circuit implementation of a finite-state transducer parsing the inputs, which leads to a parallel prefix computation (PPC) task. Ladner and Fischer [**?**] presented a general framework providing an efficient circuit implementation of arbitrary (small) transducers, giving rise to the most efficient adder circuits known to date. While the Ladner and Fischer framework fails to yield hazard-free circuits, the result from [**?**] suggests the possibility of a general hazard-free construction. Despite being a somewhat specialized class of circuits, the fact that addition can be phrased as a PPC problem renders hazard-free transducer circuits a key stepping stone towards hazard-free arithmetics.

1.2   Transducers

A deterministic finite-state transducer is a finite state machine that outputs a symbol on each state transition. We phrase our results for *Mealy machines* [**?**], but our techniques are not specific to this type of transducer.

▶ **Definition 6** (Mealy Machine). *A Mealy machine $T = (S, s_0, \Sigma, \Lambda, t, o)$ is a 6-tuple, where*
  (i) *S is the finite set of states,*
  (ii) *$s_0 \in S$ is the starting state,*
  (iii) *Σ is the finite input alphabet,*
  (iv) *Λ is the finite output alphabet,*
  (v) *$t\colon S \times \Sigma \to S$ is the state transition function, and*
  (vi) *$o\colon S \times \Sigma \to \Lambda$ is the output function.*

Each Mealy machine induces a *transcription function* mapping a string of input symbols to a string of output symbol of the same length.

▶ **Definition 7** (Transcription Function τ). *For Mealy machine $T = (S, s_0, \Sigma, \Lambda, t, o)$ and $n \in \mathbb{N}$, the transcription function $\tau_{T,n}\colon \Sigma^n \to \Lambda^n$ is given in the following way. Define for $i \in \{1, \ldots, n\}$ and $x \in \Sigma^n$ the state s_i after i steps inductively via $s_i := t(s_{i-1}, x_i)$. Then $\tau_{T,n}(x)_i := o(s_{i-1}, x_i)$.*

Note that every Boolean function $f\colon \mathbb{B}^n \to \mathbb{B}^m$ can (essentially) be realized by a deterministic finite-state transducer. A simple implementation could read the entire input string x and output $f(x)$ on reception of the last input symbol. This approach, however, requires an exponential number of states $|S| \in \mathcal{O}(2^n)$ to memorize the input. Accordingly, it is of interest to consider *small* transducers. In particular, important basic operations, like addition, max, and min, can be implemented by constant-size transducers.

1.3   Our Contribution

In this work, we establish that constant-size transducers allow for an efficient hazard-free circuit implementation. Denoting by ℓ and m the (constant) number of bits encoding an

input symbol and an output symbol, respectively, by n the length of the input string, by $|S|$ the number of states of the transducer, and by k an upper bound on the number of metastable bits in the input, our main result is as follows.

▶ **Theorem 27.** *For any integers $k \in \mathbb{N}$, $\ell, m, n \in \mathbb{N}_{>0}$ (with $k \leq n$) and Mealy machine $T = (S, s_0, \Sigma = \mathbb{B}^\ell, \Lambda \subseteq \mathbb{B}^m, t, o)$, there is a k-bit hazard-free circuit implementing $\tau_{T,n}$. For $\kappa := \sum_{i=0}^{\min\{|S|, 2^k\}} \binom{|S|}{i}$ and $\lambda := \min\{m, 2^{|S| \cdot |\Sigma|}\}$ the circuit has*

$$\text{size} \quad \mathcal{O}\left((\kappa^3 + (2^\ell/\ell)\kappa^2 + 2^\ell \kappa \lambda)n\right)$$
$$\text{and depth} \quad \mathcal{O}\left(\log \kappa \log n + \ell\right).$$

We remark that the proof of Theorem 27 shows that we can save a factor of κ in the third term, provided that the preimage of 1 under $o(\cdot, \sigma)_j$ (i.e., bit j of the output function with the second input fixed to σ) has size at most 2^k for each $\sigma \in \Sigma$ and $j \in [m]$. In this case, there is a k-bit hazard-free circuit implementing $\tau_{T,n}$ of size $\mathcal{O}\left((\kappa^3 + 2^\ell \kappa^2 + 2^\ell \lambda)n\right)$.

The asymptotic complexity depends on k, the upper bound on the numbers of ʮ's. For $k \in \mathbb{N}$ we consider two cases: $2^k \geq |S|$ and $2^k < |S|$. Let in both cases $\lambda := \min\{m, 2^{|S| \cdot |\Sigma|}\}$. If $2^k \geq |S|$, we apply the trivial bound of $2^{|S|}$ for the sum over the binomial coefficients κ. Note that in this case, trivially the preimage of 1 under $o(\cdot, \sigma)$ has size at most 2^k. Hence, as discussed above, the factor of κ in the third term of the size bound can be removed. This gives us the following size and depth bounds for a fully hazard-free implementation.

▶ **Corollary 8.** *For any integers $\ell, n \in \mathbb{N}$ and Mealy machine $T = (S, s_0, \Sigma = \mathbb{B}^\ell, \Lambda, t, o)$, the transcription function $\tau_{T,n}$ can be implemented by a hazard-free circuit*

$$\text{of size} \quad \mathcal{O}((2^{3|S|} + 2^{2|S|+\ell}/\ell + 2^\ell \lambda)n)$$
$$\text{and depth} \quad \mathcal{O}(|S| \log n + \ell).$$

We stress that this result stands out against the lower bound from [**?**], which proves an exponential dependence of the circuit size on n, for any general construction of hazard-free circuits. While the above theorem incurs exponential overheads in terms of the size of the transducer, the dependence on n is asymptotically optimal. Thus, for constant-size transducers, we obtain asymptotically optimal hazard-free implementations of their transcription functions, both with respect to size and depth. More generally, Theorem 27 shows that the task of implementing transcription functions is fixed-parameter tractable with respect to $\max\{\ell, |S|\}$.

If $2^k < |S|$, the Binomial Theorem [**?**] provides a stronger bound for κ, the sum over the binomial coefficients. Note that the respective factor in the third term of the size bound can still be removed if the output function satisfies the above requirement, but this does not hold true in general.

▶ **Corollary 9.** *Given integers $k, \ell, n \in \mathbb{N}$ and Mealy machine $T = (S, s_0, \Sigma = \mathbb{B}^\ell, \Lambda, t, o)$, such that $2^k < |S|$, the transcription function $\tau_{T,n}$ can be implemented by a k-bit hazard-free circuit*

$$\text{of size} \quad \mathcal{O}((|S|^{3 \cdot 2^k} + (2^\ell/\ell)|S|^{2 \cdot 2^k} + 2^\ell |S|^{2^k} \lambda)n)$$
$$\text{and depth} \quad \mathcal{O}(2^k \log(|S|) \log n + \ell).$$

The main insight underlying the proof of Theorem 27 is an understanding of how the encoding of a piece of information (such as an input) affects the ability of the circuit to keep track of this information. Due to the ambiguity presented by ʮ signals, naive encodings may

lose information crucial for determining a stable output, which cannot be recovered later. We tackle this problem by introducing a "universal" encoding that explicitly stores for each $A \subseteq S$ (of size at most 2^k) whether the state machine is currently in *some* state from A. This redundancy is sufficient to completely eliminate k-bit hazards, yet is affordable when $|S|$ or k are small.

Organization of this article

We discuss related work in Section 2. With the help of a toy example, Section 3 builds intuition and presents the key ideas needed to obtain Theorem 27. That is, Section 3.1 explains why the Ladner and Fischer framework fails, and Section 3.2 introduces the encoding used to resolve the main shortcoming of their approach. Finally, we prove Theorem 27 in Section 3.3.

2   Related Work

Applications of Hazard-Free Circuits

If timing constraints for accessing bistable elements, such as flip-flops or latches, are violated, they may become metastable. That is, their output signal exhibits an intermediate voltage between high (1) and low (0) for an unknown amount of time, before it resolves to either one of them. Downstream circuit components may respond as if subjected to an input of 1 or 0, or produce an intermediate output voltage themselves, where the responses of different components may be in conflict.

By the late 70's, there was an intense debate among hardware developers whether the problem of metastability can be dealt with deterministically, by suitable design of circuits [**?**, **?**, **?**, **?**]. Marino [**?**] proved by a topological argument that no circuit (with non-constant output) can avoid, resolve, or detect metastability in all cases. Ever since, the standard approach to evading metastability in applications where timing constraints cannot be guaranteed have been synchronizers, see, e.g., [**?**, Chap. 2]. Synchronizers trade time for decreased probability of ongoing metastability (and thus resulting errors).

As mentioned earlier, modeling propagation of metastability in a worst-case fashion matches Stephen Cole Kleene's "strong logic of indeterminacy" [**?**, §64]. Characterizing the complexity of hazard-free circuits thus is of immediate relevance for avoiding use of synchronizers, eliminating both incurred delays and the (remaining) probability of error due to deterministic guarantees. A non-trivial example of this is given by [**?**], where Gray code inputs that may present some metastability are sorted deterministically, with only constant-factor overheads compared to optimal sorting networks in Boolean logic.

To our knowledge, the first hazard-free multiplexer was published by Goto [**?**], but remained unnoticed by the western world for decades. Huffman [**?**] provided the first general construction of metastability-containing circuits. Goto and Huffman make no mention of Kleene logic, developing their own terms. Similarly, related work on cybersecurity [**?**, **?**] appears to derive from coming up with the concept independently again. See [**?**] for a survey covering some of these articles and discussing different logics. Further applications are discussed in [**?**]. In our opinion, all of this goes to show that the questions we study in this paper are fundamental and of widespread interest.

Complexity of Hazard-free Circuits

In [**?**] it was shown that for monotone functions, their *hazard-free complexity* (i.e., the size of the smallest hazard-free implementation) equals their monotone complexity (i.e., the size of the smallest implementation without negation gates). This yields a number of unconditional lower bounds as corollaries of results on monotone circuits. In particular, an exponential separation between hazard-free and standard circuit complexity follows from [**?**, **?**], and the naive monotone circuit of cubic size for Boolean matrix multiplication is optimal [**?**, **?**]. These lower bounds are complemented by a general construction yielding circuits of size $n^{\mathcal{O}(k)}|C|$ without k-bit hazards, where C is an arbitrary circuit implementing the desired function. Thus, for constant k, the overhead for removing k-bit hazards is polynomial in n. The above separation result implies that an overhead of $2^{k^{\Omega(1)}}$ is necessary, but it remains open whether the task is fixed-parameter tractable w.r.t. k.

Jukna [**?**] strengthens the results from Ikenmeyer et al. [**?**] on the gap between unconstrained and hazard-free circuit complexity. Moreover, Jukna shows that in general any Boolean function $f: \mathbb{B}^n \to \mathbb{B}$ can be implemented by a hazard-free circuit of size $\mathcal{O}(2^n/n)$. We remark that applying this to the transcription function $\tau_{T,n}: \mathbb{B}^{\ell n} \to \mathbb{B}^{mn}$ results in a circuit of size $\mathcal{O}((2^{\ell n}/(\ell n))mn) = \mathcal{O}((2^{\ell n}/\ell)m)$, which is much larger than the circuit presented in this work.

In contrast, some functions and specific hazards admit much more efficient solutions, e.g. the optimal sorting networks in [**?**]. If the possible positions of $\mathfrak{u}$ inputs are restricted to index set I, a construction based on hazard-free multiplexers avoids the respective hazards with a circuit of size $\mathcal{O}(2^{|I|}|C|)$ [**?**, Lemma 5.2], where C is as above. This can be seen as combining speculative computing [**?**, **?**] with hazard-free multiplexers.

Furthermore, we remark that the lower bound can be circumvented by using non-standard non-combinational logic [**?**]. Using clocked circuits and so-called masking registers, k-bit hazards can be eliminated with factor $\mathcal{O}(k)$-overhead. Masking registers also strictly increase the computational power of the system with each clock cycle. However, in this paper we consider combinational logic only.

Transducers

Our approach can be seen as an extension of the work of Ladner and Fischer [**?**]. This celebrated result yields the only asymptotically optimal adder constructions known to date, cf. [**?**]. Alongside the result for binary addition the authors point out general applicability of their parallel prefix computation (PPC) framework: for any transcription function, it allows constructing a circuit implementing it. However, as we discuss in detail in Section 3, their approach cannot be applied to our setting, as it does not take into account the uncertainty imposed by unstable inputs.

Our approach might also remind the reader of the power set construction [**?**, Thm. 1.39], which translates a non-deterministic finite-state automation into a deterministic one operating on the power set of the state space. This analogy is correct to the extent that we seek to maintain information on the set of states that are reachable by resolutions of the input. However, Kleene logic has the fundamentally different characteristic that the choice of encoding (e.g. of states) affects to what extent the circuit can keep track of the encoded information. In a nutshell, we prove that it is sufficient to maintain a bit vector indicating for each element $A \subseteq S$ of the power set whether, given the input, all states that could have been reached by the state machine are a subset of A. This resolves an issue that has no connection to the original power set construction.
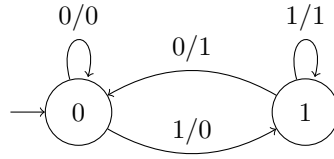
Figure 1 The shift transducer delays the input by one symbol. It serves as a running example.

3 Extending the PPC Framework to Hazard-free Circuits

In this section, we walk the reader through the main ideas underlying our framework, at hand of a simple running example, and then prove our main result. In Section 3.1 we demonstrate that a naive application of the parallel prefix framework [**?**] results in circuits that are not hazard-free. We then use the running example to illustrate how to overcome this hurdle and to obtain a hazard-free circuit by making use of the universal encoding introduced in Section 3.2. Finally, we prove Theorem 27 in Section 3.3.

The running example we use throughout this section is an extremely simple transducer: it simply shifts the input sequence by one bit, outputting a 0 on reception of the first symbol; see Figure 1 for an illustration. It has two states (referred to as 0 and 1), which are used to keep track of the most recently processed input bit. Hence, the transition and output functions are obvious: the automaton transitions to state $s \in \{0, 1\}$ on reception of input s, and outputs s when leaving state s. Thus, the transducer is formally specified by the 6-tuple $(S := \{0, 1\}, s_0 := 0, \Sigma := \{0, 1\}, \Lambda := \{0, 1\}, t(s, i) = i, o(s, i) = s)$.

Clearly, this transducer is a toy example, and it is pointless to construct a circuit implementing its transcription function – this is easily achieved by suitable rewiring the inputs instead. However, the shift transducer serves as a minimal example for illustrating both the obstacle we need to overcome and the general solution we provide for doing so.

3.1 The Classic PPC Framework

In their work, Ladner and Fischer observe that any transcription function $\tau_{T,n}$ on inputs $x \in \mathbb{B}^n$ can be efficiently implemented as a circuit by following four steps. Given an encoding of functions $S \to S$, each step takes the output of the previous step as input. For each $i \in \{1, \ldots, n\}$:

(Step 1)   compute the encoding of t_{x_i}, where $t_\sigma := t(\cdot, \sigma) \colon S \to S$ is the restricted transition function for symbol $\sigma \in \Sigma$,

(Step 2)   compute the composition $\pi_i := t_{x_i} \circ \ldots \circ t_{x_1}$ of restricted transition functions,

(Step 3)   compute the i-th state, i.e., evaluate $s_i = \pi_i(s_0)$, and

(Step 4)   compute the i-th output $o(s_{i-1}, x_i) = \tau_{T,n}(x)_i$.

Steps 1, 3, and 4 can be performed independently and hence in parallel for each i, based on the output of previous steps. This means that each of them can be performed by n copies of a circuit whose size (and thus depth) depends only on the transducer. In contrast, Step 2, the computation of all prefixes, inherently relies on information across all i's. To achieve small depth without blowing up the circuit size, Ladner and Fischer exploit the associativity of function composition.

For a constant-size Mealy machine, Steps 1, 3, and 4 can be performed by circuits of size $\mathcal{O}(n)$ and depth $\mathcal{O}(1)$, and Step 2 can be done by a circuit of size $\mathcal{O}(n)$ and depth $\mathcal{O}(\log n)$. In their argument showing this, Ladner and Fischer encode the space of functions $S \to S$

as Boolean $|S| \times |S|$ matrices.[2] Functional composition (Step 2), hence, becomes Boolean matrix multiplication, while evaluation of functions (Step 3) becomes Boolean matrix-vector multiplication.

Applying this to our example, states 0 and 1 of our transducer are represented by the column unit vector $e^{(0)} = \binom{1}{0}$ and $e^{(1)} = \binom{0}{1}$, respectively. Representing $t_\sigma \colon S \to S$ as a Boolean matrix in the natural way, the column corresponding to state s is the unit vector $e^{(t(s,\sigma))}$. Assuming that we make a best effort and use a hazard-free circuit for computing the encodings of t_σ, the hazard-free extension determines what our circuit will compute when receiving ʉ as an input symbol. Denote by $\mathcal{M}_{t_\sigma}$ the matrix computed by this hazard-free circuit for (the encoding of) the transition function restricted to input symbol σ. We thus obtain

$$\mathcal{M}_{t_0} = \begin{pmatrix} 1 & 1 \\ 0 & 0 \end{pmatrix}, \quad \mathcal{M}_{t_1} = \begin{pmatrix} 0 & 0 \\ 1 & 1 \end{pmatrix}, \quad \mathcal{M}_{t_ʉ} = \begin{pmatrix} 1 & 1 \\ 0 & 0 \end{pmatrix} * \begin{pmatrix} 0 & 0 \\ 1 & 1 \end{pmatrix} = \begin{pmatrix} ʉ & ʉ \\ ʉ & ʉ \end{pmatrix},$$

where the $*$ operator is applied component-wise: as each entry of the computed matrix depends on whether the input symbol was 0 or 1, ʉ must result in the all-ʉ matrix.

Function composition corresponds to Boolean matrix multiplication, i.e., for restricted functions t_σ and $t_{\sigma'}$ $(\sigma, \sigma' \in \Sigma)$, $\mathcal{M}_{t_\sigma \circ t_{\sigma'}} = \mathcal{M}_{t_\sigma} \cdot \mathcal{M}_{t_{\sigma'}}$, where $\cdot$ denotes the Boolean matrix multiplication operator. Similarly, function evaluation corresponds to matrix-vector multiplication, meaning that the framework stipulates to compute the encoding of π_i as $\mathcal{M}_{t_{x_i}} \cdot \ldots \cdot \mathcal{M}_{t_{x_1}}$ and hence s_i as $\mathcal{M}_{t_{x_i}} \cdot \ldots \cdot \mathcal{M}_{t_{x_1}} \cdot e^{(s_0)}$. Again, we make a best effort, i.e., assume that hazard-free circuits are used. Therefore, the circuit will compute $\mathcal{M}_{t_{x_{i-1}}} \cdot_ʉ \ldots \cdot_ʉ \mathcal{M}_{t_{x_1}} \cdot_ʉ e^{(s_0)}$.

Finally, the i-th output bit is computed according to the output function by mapping $e^{(0)}$ to output 0 and $e^{(1)}$ to output 1. Note that the redundant representation allows for some freedom: we can choose for $\binom{0}{0}$ and $\binom{1}{1}$ whether to map them to 0 or 1, respectively. As these state vectors cannot occur anyway, this choice has no impact on stable inputs. It might, however, affect the behavior of a (hazard-free) circuit confronted with unstable inputs.

Now consider Table 2, which breaks down the computation for two input strings, first the stable input 0010 and then an input containing a single unstable bit, 0ʉ10. A hazard-free circuit should output 0001 in the first case and 00ʉ1 in the latter case. However, multiplication of any function encoding with the all-ʉ matrix results again in the all-ʉ matrix, such that any further step of function composition will return $\mathcal{M}_{t_ʉ}$. Hence, for the second input, the hazard-free extension of the established approach will compute ʉ as the last symbol.

To identify the key issue, examine the sequence of matrices determined from the input symbols, which represent the transition functions restricted to the respective input bit. $\mathcal{M}_{t_0}$ will map any vector corresponding to a stable state, i.e., each unit vector, to $e^{(0)}$. This reflects the fact that a 0 is guaranteed to result in state 0. Accordingly, multiplying $\mathcal{M}_{t_0}$ with any matrix representing the transition function restricted to a stable input symbol will result in $\mathcal{M}_{t_0}$: no matter what happened to the state machine before, the state after receiving input symbol 0 is 0 (represented by $e^{(0)}$).

On the other hand, $\mathcal{M}_{t_ʉ}$ is the "correct" representation for input symbol ʉ: regardless of the previous state, the resolutions 0 and 1 of input symbol ʉ reach state 0 or 1 respectively, and $\binom{1}{0} * \binom{0}{1} = \binom{ʉ}{ʉ}$. Unfortunately, (the hazard-free extension of) Boolean matrix multiplication

---

[2] We remark that it would be more efficient to encode these functions by listing their values, reducing the size of the encoding from $|S|^2$ to $|S|\lceil \log |S| \rceil$. However, for $|S| \in \mathcal{O}(1)$ this does not affect the asymptotics.

■ **Table 2** Application of the Ladner and Fischer approach to the example transducer for two different input strings. Top: stable input word 0010, Bottom: unstable input word 0u10. Gray area: these values do not match the results from a hazard-free computation.

	i	0	1	2	3	4
input	x_i	-	0	0	1	0
Step 1, function encoding	$\mathcal{M}_{t_{x_i}}$	-	$\mathcal{M}_{t_0}$	$\mathcal{M}_{t_0}$	$\mathcal{M}_{t_1}$	$\mathcal{M}_{t_0}$
Step 2, function composition	$\mathcal{M}_{\pi_i}$	-	$\mathcal{M}_{t_0}$	$\mathcal{M}_{t_0}$	$\mathcal{M}_{t_1}$	$\mathcal{M}_{t_0}$
Step 3, function evaluation	$s_i = \pi_i(s_0)$	$e^{(0)}$	$e^{(0)}$	$e^{(0)}$	$e^{(1)}$	$e^{(0)}$
Step 4, output	$o(s_{i-1}, x_i)$	-	0	0	0	1

	i	0	1	2	3	4
input	x_i	-	0	u	1	0
Step 1, function encoding	$\mathcal{M}_{t_{x_i}}$	-	$\mathcal{M}_{t_0}$	$\mathcal{M}_{t_u}$	$\mathcal{M}_{t_1}$	$\mathcal{M}_{t_0}$
Step 2, function composition	$\mathcal{M}_{\pi_i}$	-	$\mathcal{M}_{t_0}$	$\mathcal{M}_{t_u}$	$\mathcal{M}_{t_u}$	$\mathcal{M}_{t_u}$
Step 3, function evaluation	$s_i = \pi_i(s_0)$	$e^{(0)}$	$e^{(0)}$	$\binom{u}{u}$	$\binom{0}{u}$	$\binom{0}{u}$
Step 4, output	$o(s_{i-1}, x_i)$	-	0	0	u	u

of any matrix with the all-u matrix $\mathcal{M}_{t_u}$ can never yield a matrix that is composed of column unit vectors. The circuit computes

$$\mathcal{M}_{t_1} \cdot_u \mathcal{M}_{t_u} \cdot_u \mathcal{M}_{t_0} \cdot_u e^{(s_0)} = \begin{pmatrix} 0 & 0 \\ 1 & 1 \end{pmatrix} \cdot_u \begin{pmatrix} u & u \\ u & u \end{pmatrix} \cdot_u \begin{pmatrix} 1 & 1 \\ 0 & 0 \end{pmatrix} \cdot_u \begin{pmatrix} 1 \\ 0 \end{pmatrix}$$

$$= \begin{pmatrix} 0 & 0 \\ u & u \end{pmatrix} \cdot_u \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} 0 \\ u \end{pmatrix}$$

as the (encoding of) state s_3 in Step 3. Thus, in Step 4 the circuit at its best effort can only output

$$o_u\left(\binom{0}{u}, 0\right) = o\left(\binom{0}{0}, 0\right) * o\left(\binom{0}{1}, 0\right) = o\left(\binom{0}{0}, 0\right) * 1 = \begin{cases} 1 & \text{if } o\left(\binom{0}{0}, 0\right) = 1 \\ u & \text{if } o\left(\binom{0}{0}, 0\right) = 0. \end{cases}$$

This might give the false hope of escaping the problem by leveraging our aforementioned freedom to choose $o\left(\binom{0}{0}, 0\right)$ by setting it to 1, but this is a red herring. If $o\left(\binom{0}{0}, 0\right) = 1$, then the input string 0u00 forces the circuit to incorrectly output $o_u\left(\binom{u}{0}, 0\right) = o\left(\binom{0}{0}, 0\right) * 0 = u$ as the final bit.

Intuitively, the main take-away from this example is that encoding the transition function as an $|S| \times |S|$ matrix is insufficient to keep track of the set of reachable states. The crucial problem is that uncertainty about the transducer's state can be removed (or reduced) by later input symbols. In our example, we have a very simple such case: any stable input symbol fully determines the attained state, regardless of the previous state.

In our approach we keep track of a strict subset of states $A \subset S$ the state machine could have reached when facing some inputs with some uncertainty, such that we can infer the set of states $B \subset S$ (ideally a singleton, if the uncertainty has been completely masked) that can be reached by the current state transition.

3.2   Suitable Encoding of Transition Functions

Before formalizing our encoding, we provide some intuition, by using, again, our toy example, the shift transducer. As we kept the example tiny, the number of subsets of the statespace, i.e., the power set of S, is small: the four possible subsets are $\emptyset$, $\{0\}$, $\{1\}$ and $\{0,1\}$. Already a single u input leads to the largest possible uncertainty about the state of the transducer, hence we choose $k = 1$ throughout the example.

Intuition

To avoid the pitfall discussed in Section 3.1, we now choose a highly redundant matrix representation. Fix an input symbol $\sigma \in \{0,1\}$. The corresponding $2^{|S|} \times 2^{|S|}$ Boolean matrix encodes for each pair of sets $A, B \subseteq S$ whether for each state in A receiving σ as next input symbol will result in a state from B. Again, assuming that a hazard-free circuit is used to compute the matrix representation, this choice fully determines the matrix $\mathcal{M}_{t_{\mathrm{u}}} = \mathcal{M}_{t_0} * \mathcal{M}_{t_1}$ resulting from input symbol u. Labeling the rows by subsets B and columns by the subsets A, the resulting matrices $\mathcal{M}_{t_0}$, $\mathcal{M}_{t_1}$, and $\mathcal{M}_{t_{\mathrm{u}}}$ are

$$
\begin{array}{c c}
\begin{array}{c}
\phantom{\{0,1\}} \\
\emptyset \\
\{0\} \\
\{1\} \\
\{0,1\}
\end{array}
\begin{array}{c}
\begin{array}{cccc}
\emptyset & \{0\} & \{1\} & \{0,1\}
\end{array} \\
\left(\begin{array}{cccc}
1 & 0 & 0 & 0 \\
1 & 1 & 1 & 1 \\
1 & 0 & 0 & 0 \\
1 & 1 & 1 & 1
\end{array}\right)
\end{array}
,
&
\begin{array}{c}
\phantom{\{0,1\}} \\
\emptyset \\
\{0\} \\
\{1\} \\
\{0,1\}
\end{array}
\begin{array}{c}
\begin{array}{cccc}
\emptyset & \{0\} & \{1\} & \{0,1\}
\end{array} \\
\left(\begin{array}{cccc}
1 & 0 & 0 & 0 \\
1 & 0 & 0 & 0 \\
1 & 1 & 1 & 1 \\
1 & 1 & 1 & 1
\end{array}\right)
\end{array}
, \text{ and}
\end{array}
$$

$$
\begin{array}{c}
\phantom{\{0,1\}} \\
\emptyset \\
\{0\} \\
\{1\} \\
\{0,1\}
\end{array}
\begin{array}{c}
\begin{array}{cccc}
\emptyset & \{0\} & \{1\} & \{0,1\}
\end{array} \\
\left(\begin{array}{cccc}
1 & 0 & 0 & 0 \\
1 & \mathrm{u} & \mathrm{u} & \mathrm{u} \\
1 & \mathrm{u} & \mathrm{u} & \mathrm{u} \\
1 & 1 & 1 & 1
\end{array}\right)
\end{array}
,
$$

respectively. Consider, for example, input symbol 0 and the corresponding matrix $\mathcal{M}_{t_0}$. Each set of states $\{0\}$, $\{1\}$ and $\{0,1\}$ will transition to state 0. As 0 is a subset of $\{0\}$ and $\{0,1\}$ the respective column vectors of the matrix are $\begin{pmatrix} 0 & 1 & 0 & 1 \end{pmatrix}^{\mathsf{T}}$. Note that each matrix maintains the trivialities that the empty set will always be mapped to a subset of any set (leftmost column), no non-empty set is mapped to a subset of the empty set (top row), and any set will be mapped to a subset of $S = \{0,1\}$ (bottom row).[3] Crucial to us is the point that the encoding now "takes note" of the fact that even when an input symbol is u, it remains certain that for any resolution of the input the transducer must end up in *some* state, reflected by the bottom row of $\mathcal{M}_{t_{\mathrm{u}}}$.

---

[3] Note that the submatrices induced by the rows and columns of singleton sets equal those we got in Section 3.1. Since we opted for a minimal example with only two states, none of the additional entries depend on the specific transition function. Note that this changes for $|S| > 2$.

Application to the example

Applying the framework of Section 3.1 with the new encoding to the input string $0\mathfrak{u}10$, this time Step 2 yields for π_3:

$$\mathcal{M}_{t_1} \cdot_{\mathfrak{u}} \mathcal{M}_{t_{\mathfrak{u}}} \cdot_{\mathfrak{u}} \mathcal{M}_{t_0} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{pmatrix} \cdot_{\mathfrak{u}} \begin{pmatrix} 1 & 0 & 0 & 0 \\ 1 & \mathfrak{u} & \mathfrak{u} & \mathfrak{u} \\ 1 & \mathfrak{u} & \mathfrak{u} & \mathfrak{u} \\ 1 & 1 & 1 & 1 \end{pmatrix} \cdot_{\mathfrak{u}} \begin{pmatrix} 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 \\ 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 \end{pmatrix}$$

$$= \begin{pmatrix} 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{pmatrix} \cdot_{\mathfrak{u}} \begin{pmatrix} 1 & 0 & 0 & 0 \\ 1 & \mathfrak{u} & \mathfrak{u} & \mathfrak{u} \\ 1 & \mathfrak{u} & \mathfrak{u} & \mathfrak{u} \\ 1 & 1 & 1 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{pmatrix} = \mathcal{M}_{t_1}.$$

As we can see, multiplying with $\mathcal{M}_{t_1}$ from the left now correctly recovers $\mathcal{M}_{t_1}$, i.e., regardless of previous possibly unstable input symbols, the computed matrix reflects that reading input symbol 1 results in state 1. The above behavior is not a corner case due to the small size of our example, a larger example is given in Appendix B.

A fundamental problem in hazard-free circuits is that the resolution of the superposition may add undesired values (Observation 3). Recall that for a Boolean function $f\colon \mathbb{B}^n \to \mathbb{B}^m$, we obtain $f_{\mathfrak{u}}(x)$ by mapping each $y \in \mathrm{res}(x)$ using f and then taking the $*$ operation over the resulting set. The latter might, depending on x and the encoding, lose information, as $\mathrm{res}(f_{\mathfrak{u}}(x))$ might be a strict superset of $f(\mathrm{res}(x))$. This becomes problematic when we subsequently apply some function, e.g., $g\colon \mathbb{B}^m \to \mathbb{B}$, that is constant on $f(\mathrm{res}(x))$, but not on $\mathrm{res}(f_{\mathfrak{u}}(x))$ for some $x \in \mathbb{T}^n$; we then get that

$$\mathfrak{u} = g_{\mathfrak{u}}(f_{\mathfrak{u}}(x)) = *(g(\mathrm{res}(f_{\mathfrak{u}}(x)))) \neq *(g(f(\mathrm{res}(x))) = (g \circ f)_{\mathfrak{u}}(x) \in \mathbb{B}\,.$$

A universal encoding for functions

The key idea underlying our solution to this problem is to maintain the information that f maps $\mathrm{res}(x)$ to $f(\mathrm{res}(x))$. As illustrated by the example, the encoding stores for each $A \subseteq \mathbb{B}^n$ and $B \subseteq \mathbb{B}^m$ whether $f(A) \subseteq B$. When composing functions, we then can retrieve the information that $g \circ f$ is constant on $\mathrm{res}(x)$.

The size of the encoding can be reduced if we only regard k-bit hazards. If the number of $\mathfrak{u}$'s in the input x can be bounded by an integer k, then there is also an upper bound on $|f(\mathrm{res}(x))|$. As each $\mathfrak{u}$ has two stable resolutions, we can readily bound $|f(\mathrm{res}(x))| \leq |\mathrm{res}(x)| \leq 2^k$. Hence, the encoding can be reduced to sets $A \subseteq \mathbb{B}^n$ and $B \subseteq \mathbb{B}^m$, where $|A| \leq 2^k$ and $|B| \leq 2^k$.

This leads to the following encoding, which is universal in the sense that it gives rise to k-bit hazard-free implementations of arbitrary transducers.

▶ **Definition 10** (Universal Function Encoding). *Denote by $\mathcal{P}_t(A)$ the set of all subsets of A with cardinality smaller equal to $t \in \{0,\ldots,|A|\}$, i.e., $\mathcal{P}_t(A) := \{A' \subseteq A \mid |A'| \leq t\}$. Given a function $f\colon S \to T$ and $k \in \mathbb{N}$, define*

$$\forall A \in \mathcal{P}_{2^k}(S), B \in \mathcal{P}_{2^k}(T)\colon \ (\mathcal{M}_f)_{BA} := \begin{cases} 1 & \text{if } f(A) \subseteq B \\ 0 & \text{else.} \end{cases}$$

Thus, the Boolean matrix $\mathcal{M}_f$ has dimension

$$|\mathcal{P}_{2^k}(T)| \times |\mathcal{P}_{2^k}(S)| = \sum_{i=0}^{\min\{|T|,2^k\}} \binom{|T|}{i} \times \sum_{i=0}^{\min\{|S|,2^k\}} \binom{|S|}{i}\,.$$

Moreover, for $s \in S$ and $A \in \mathcal{P}_{2^k}(S)$, define $e^{(s)}$ via $e_A^{(s)} := 1$ if $s \in A$ and $e_A^{(s)} := 0$ otherwise. Hence, for all $B \in \mathcal{P}_{2^k}(T)$ we have that $\left(\mathcal{M}_f \cdot e^{(s)}\right)_B = 1$ if $f(s) \in B$ and $\left(\mathcal{M}_f \cdot e^{(s)}\right)_B = 0$ otherwise.

We remark that we are mostly interested in the case where $T = S$, since for the restricted transition functions computed in Step 1 we only need to represent functions from S to S.

3.3   Proving the Main Result

Our goal in this subsection is to show Theorem 27. To this end, we first establish that the above encoding indeed keeps track of all required information to remove uncertainty in case the input allows it. We show that the above matrix representation is a suitable encoding, i.e., we show that the representation is capable of encoding the transition function without dropping information, here the transition function is restricted to a single input symbol.

The PPC framework of Ladner and Fischer

Recall that the PPC framework computes states s_i by (Step 1) translating input symbol x_i into the matrix representation of t_{x_i}, (Step 2) determining by matrix multiplication the transition function π_i resulting from a sequence of input symbols, and (Step 3) evaluating the transition function π_i on $e^{(s_0)}$ via matrix-vector multiplication. Given state s_i the output at position i can be determined by application of o_{x_i} (Step 4). In the PPC framework we replace the original matrix represenentation by the new universal function encoding. We show that the universal encoding overcomes the issue of information loss during function composition. If function composition does not lose information, then repeated application of function composition gives hazard-free transition functions π_i, which will be formalized in Corollary 20.

Main stepping stone

To show that for the universal encoding the strategy also succeeds in face of unstable inputs, we need to prove that composing functions and translating the composed function into its matrix representation is equivalent to first translating each function to its matrix representation and then multiplying these matrices. This is captured by the following theorem, which is our main stepping stone towards Theorem 27.

▶ **Theorem 11.** *Let $k \in \mathbb{N}$, $f_j\colon S \to T$ for all $j \in J$, $g_i\colon T \to U$ for all $i \in I$, $A \in \mathcal{P}_{2^k}(S)$, and $C \in \mathcal{P}_{2^k}(U)$. If $|J| \cdot |A| \leq 2^k$, then*

$$\left(\left(\underset{i \in I}{\text{\Large$*$}} \mathcal{M}_{g_i}\right) \cdot_{\mathfrak{u}} \left(\underset{j \in J}{\text{\Large$*$}} \mathcal{M}_{f_j}\right)\right)_{CA} = \left(\underset{(i,j) \in I \times J}{\text{\Large$*$}} \mathcal{M}_{g_i \circ f_j}\right)_{CA}.$$

The condition $|J| \cdot |A| \leq 2^k$ may seem non-intuitive at first. The product $|J| \cdot |A|$ corresponds to the number of resolutions of the input that has been processed so far. The product is bounded by 2^k, i.e., the number of resolutions of k many $\mathfrak{u}$'s. In the application of Theorem 11, set J corresponds to the resolutions of respective parts of the input. Set A corresponds to the current state of the transducer, and its size depends on the uncertainty of previous transitions. More intuition is provided in the example in Appendix B.

Before we prove the key stepping stone we discuss tools that are used in the proof of the main result and the key stepping stone. First, we define hazard-free multiplexers which

are used in Step 4 of the PPC framework. Second, we show that there is an efficient implementation of hazard-free matrix multiplication. Third, we introduce monotone resolutions, a technique used in the proofs. Last, we state a recent result on the complexity of hazard-free circuits for general functions, which is applied in the proof of the main theorem.

Hazard-free multiplexer

For later use we define the ℓ-input multiplexer $\texttt{MUX}_\ell$. A multiplexer is a circuit that selects one of its inputs according to a dedicated select input. The select input encodes index $i \in [2^\ell]$, where $[t] = \{0, \ldots, t-1\}$, for $t \in \mathbb{N}_{>0}$.

▶ **Definition 12.** *Let $\ell, b \in \mathbb{N}_{>0}$. An ℓ-input multiplexer $\texttt{MUX}_\ell$ receives inputs $x_i \in \mathbb{B}^b$ for $i \in [2^\ell]$, and select input $s \in \mathbb{B}^\ell$. Let $\langle \cdot \rangle \colon \mathbb{B}^\ell \to \mathbb{N}$ be the standard binary decoding function. Interpreting the select input s as an index, $\texttt{MUX}_\ell$ outputs $x_{\langle s \rangle}$, i.e., $\texttt{MUX}_\ell(x_0, \ldots, x_{2^\ell-1}, s) \coloneqq x_{\langle s \rangle}$.*

Hazard-free multiplexers are an efficient tool to select from a set of functions encoded by the universal function encoding. For ease of readability we name these multiplexers $\mathcal{M}\text{-}\texttt{MUX}$. Given select input $s \in \mathbb{B}^\ell$ and a set of functions $\{f_j : S \to T | j \in \mathbb{B}^\ell\}$, interpret j as binary numbers, set all $\mathcal{M}_{f_j}$ as the multiplexer inputs, and s as the select input. Then $\mathcal{M}\text{-}\texttt{MUX}$ selects the universal encoding of f_j, where $\langle j \rangle = \langle s \rangle$;

$$\mathcal{M}\text{-}\texttt{MUX}_\ell(\{f_j | j \in \mathbb{B}^\ell\}, s) \coloneqq \mathcal{M}_{f_s} \,.$$

The ℓ-input multiplexer can be implemented by a hazard-free circuit [**?**].

▶ **Corollary 13** (of [**?**, Lemma 5.1]). *Let $\ell, b \in \mathbb{N}_{>0}$, there is a hazard-free implementation of $\texttt{MUX}_\ell$ computing $(\texttt{MUX}_\ell)_{\texttt{u}}(x_0, \ldots, x_{2^\ell-1}, s)$ with $x_i \in \mathbb{B}^b$ for $i \in [2^\ell]$ and $s \in \mathbb{B}^\ell$. The implementation has*

$$\begin{aligned} \text{size} \quad & \mathcal{O}\left(2^\ell b\right) \\ \text{and depth} \quad & \mathcal{O}\left(\ell\right) \,. \end{aligned}$$

Proof. We only sketch the proof as this is a standard construction. The hazard-free implementation of a multiplexer receiving two input bits and a single select bit is given in [**?**]; it has constant size and and depth. The multiplexer can be extended to ℓ select bits by building a tree of multiplexers, where each layer is controlled by a bit of the select input. Extension to inputs of width b is simply done by copying the tree of multiplexers b times.   ◀

Hazard-free matrix multiplication

For Theorem 11 to be of use, we need a circuit implementing $\cdot_{\texttt{u}}$, i.e., hazard-free matrix multiplication. The standard Boolean matrix multiplication algorithm is known to be appropriate.

▶ **Corollary 14** (of [**?**, Lemma 4.2]). *There is a circuit of size $(2\beta-1)\alpha\gamma$ and depth $\lceil \log \beta \rceil + 1$ that computes $\mathcal{A} \cdot_{\texttt{u}} \mathcal{B}$ for matrices $\mathcal{A} \in \mathbb{T}^{\alpha \times \beta}$ and $\mathcal{B} \in \mathbb{T}^{\beta \times \gamma}$.*

Proof. The standard algorithm for Boolean matrix multiplication is monotone, i.e., does not use negations, and requires for each of the $\alpha\gamma$ entries of $\mathcal{A} \cdot_{\texttt{u}} \mathcal{B}$ a binary tree of β and gates (the leaves) and $\beta - 1$ or gates (internal nodes); monotone circuits are hazard-free.   ◀

We observe that hazard-free Boolean matrix multiplication is associative.

▶ **Observation 15** ($\cdot_{\mathfrak{u}}$ is associative). *For all $A \in \mathbb{T}^{\alpha \times \beta}$, $B \in \mathbb{T}^{\beta \times \gamma}$, and $C \in \mathbb{T}^{\gamma \times \delta}$, we have that $(A \cdot_{\mathfrak{u}} B) \cdot_{\mathfrak{u}} C = A \cdot_{\mathfrak{u}} (B \cdot_{\mathfrak{u}} C)$.*

Proof. As `or` and `and` are associative also on $\mathbb{T}$, this follows by the same straightforward calculation as for matrices over arbitrary (semi)rings.[4] ◀

To prove Theorem 11, we first need to establish that matrix multiplication indeed is equivalent to function composition for stable inputs.

▶ **Lemma 16.** *Let f and g be functions $f \colon S \to T$ and $g \colon T \to U$. For all $A \in \mathcal{P}_{2^k}(S)$ and $C \in \mathcal{P}_{2^k}(U)$, it holds that $(\mathcal{M}_g \cdot \mathcal{M}_f)_{CA} = (\mathcal{M}_{g \circ f})_{CA}$.*

Proof. Suppose $(\mathcal{M}_{g \circ f})_{CA} = 1$, i.e., $(g \circ f)(A) = g(f(A)) \subseteq C$. Note that $A \in \mathcal{P}_{2^k}(S) \Rightarrow f(A) \in \mathcal{P}_{2^k}(T)$. Then,

$$((\mathcal{M}_g) \cdot (\mathcal{M}_f))_{CA} = \sum_{B \in \mathcal{P}_{2^k}(T)} (\mathcal{M}_g)_{CB} (\mathcal{M}_f)_{BA} \geq (\mathcal{M}_g)_{Cf(A)} (\mathcal{M}_f)_{f(A)A} = 1 \cdot 1 = 1,$$

where we use that $(\mathcal{M}_f)_{f(A)A} = 1$ by Definition 10 and $(\mathcal{M}_g)_{Cf(A)} = 1$ because $g(f(A)) \subseteq C$.

Now consider the case that $(\mathcal{M}_{g \circ f})_{CA} = 0$, i.e., there exists $a \in A$ so that $g(f(a)) \not\subseteq C$. Accordingly,

$$\forall B \in \mathcal{P}_{2^k}(T) \colon (\mathcal{M}_f)_{BA} = 1 \Rightarrow (\mathcal{M}_g)_{CB} = 0,$$

because

$$(\mathcal{M}_f)_{BA} = 1 \Leftrightarrow f(A) \subseteq B \Rightarrow g(B) \not\subseteq C \Leftrightarrow (\mathcal{M}_g)_{CB} = 0.$$

Thus, for all $B \in \mathcal{P}_{2^k}(T)$, we have that $(\mathcal{M}_f)_{BA} (\mathcal{M}_g)_{CB} = 0$, leading to

$$(\mathcal{M}_g \cdot \mathcal{M}_f)_{CA} = \sum_{B \subseteq T} (\mathcal{M}_g)_{CB} (\mathcal{M}_f)_{BA} = 0.$$ ◀

Monotone resolutions

To understand how multiplying matrices plays out when the multiplicands are not stable, we exploit the monotonicity of matrix multiplication. Because flipping matrix entries of multiplicands from 0 to 1 can only flip entries from 0 to 1 in the product, we can restrict our attention to only two resolutions of each matrix: we simultaneously replace all $\mathfrak{u}$ entries by either 0 or 1, respectively.

▶ **Definition 17.** *For $A \in \mathbb{T}^{\alpha \times \beta}$ and $b \in \mathbb{B}$, define $A^{(b)} \in \mathbb{B}^{\alpha \times \beta}$ via*

$$\forall (i,j) \in \{1, \ldots, \alpha\} \times \{1, \ldots, \beta\} \colon A_{ij}^{(b)} := \begin{cases} b & \text{if } A_{ij} = \mathfrak{u} \\ A_{ij} & \text{else}. \end{cases}$$

With this definition, the above intuition is formalized by the following lemma.

▶ **Lemma 18.** *For all $G \in \mathbb{T}^{\alpha \times \beta}$, $F \in \mathbb{T}^{\beta \times \gamma}$ and all $i \in \{1, \ldots, \alpha\}$, $j \in \{1, \ldots, \gamma\}$, we have*

$$(G \cdot_{\mathfrak{u}} F)_{ij} = \mathfrak{u} \Leftrightarrow \left(G^{(0)} \cdot F^{(0)} \right)_{ij} = 0 \wedge \left(G^{(1)} \cdot F^{(1)} \right)_{ij} = 1.$$

---

[4] Note that $(\mathbb{T}, \texttt{or}, \texttt{and})$ is only a (commutative) semiring, as its "addition," i.e., `or`, has no inverses.

Proof. Fix $i \in \{1, \ldots, \alpha\}$ and $j \in \{1, \ldots, \gamma\}$. If $(G \cdot_{\mathfrak{u}} F)_{ij} = b \in \mathbb{B}$, i.e., for all G', F' such that $G' \in \text{res}(G)$ and $F' \in \text{res}(F)$, $(G' \cdot F')_{ij} = b$, then since $G^{(0)}, G^{(1)} \in \text{res}(G)$ and $F^{(0)}, F^{(1)} \in \text{res}(F)$ it holds that that

$$\left(G^{(0)} \cdot F^{(0)} \right)_{ij} = \left(G^{(1)} \cdot F^{(1)} \right)_{ij} = b.$$

Now consider the case that $(G \cdot_{\mathfrak{u}} F)_{ij} = \mathfrak{u}$. Thus, there are $G', G'' \in \text{res}(G)$ and $F', F'' \in \text{res}(F)$ satisfying that $(G' \cdot F')_{ij} = 0$ and $(G'' \cdot F'')_{ij} = 1$, respectively. It follows that

$$\left(G^{(0)} \cdot F^{(0)} \right)_{ij} = \sum_{k=1}^{\beta} G_{ik}^{(0)} F_{kj}^{(0)} \leq \sum_{k=1}^{\beta} G_{ik}' F_{kj}' = (G' \cdot F')_{ij} = 0$$

and, analogously,

$$\left(G^{(1)} \cdot F^{(1)} \right)_{ij} \geq (G'' \cdot F'')_{ij} = 1 \,. \hspace{3cm} \blacktriangleleft$$

Hazard-free circuit complexity

Jukna [**?**] presents an upper bound on the complexity of the hazard-free implementation of a Boolean function f.

▶ **Theorem 19** ([**?**]). *Given a Boolean function $f \colon \mathbb{B}^n \to \mathbb{B}$, there is a hazard-free circuit implementing f that has size $\mathcal{O}(2^n / n)$ and depth $\mathcal{O}(n)$.*

The size bound of the implementation is shown in [**?**]. The depth bound follows by the following argument. The hazard-free construction consists of two consecutive parts that are recursively defined. The first part uses at most m recursions and the second part uses $n - m$ recursions of the same recursion step. As each recursion step has constant size, the construction has depth $\mathcal{O}(n)$.

Proving the key stepping stone

Using Lemma 18, proving Theorem 11 is reduced to showing correct behavior for matrices $\left(\ast_{j \in J} \mathcal{M}_{f_j} \right)^{(0)}$ and $\left(\ast_{j \in J} \mathcal{M}_{f_j} \right)^{(1)}$ instead of all resolutions of $\ast_{j \in J} \mathcal{M}_{f_j}$ and $\ast_{i \in I} \mathcal{M}_{g_i}$.

Proof of Theorem 11. Define $\preceq$ as the partial order $b \prec \mathfrak{u}$ for $b \in \mathbb{B}$ and observe that $\ast X \preceq \ast Y$ for $X \subseteq Y \subseteq \mathbb{B}$. By Observation 3, we obtain for $\mathcal{M}_{g_i}$ (and accordingly $\mathcal{M}_{f_j}$) that

$$\{\mathcal{M}_{g_i} | i \in I\} \subseteq \text{res} \left(\underset{i \in I}{\ast} \mathcal{M}_{g_i} \right) \,.$$

Thus, by the definition of the hazard-free extension (Definition 4) and the resolution (Definition 2),

$$\left(\left(\underset{i \in I}{\ast} \mathcal{M}_{g_i} \right) \cdot_{\mathfrak{u}} \left(\underset{j \in J}{\ast} \mathcal{M}_{f_j} \right) \right)_{CA} = \ast \left(\text{res} \left(\underset{i \in I}{\ast} \mathcal{M}_{g_i} \right) \cdot \text{res} \left(\underset{j \in J}{\ast} \mathcal{M}_{f_j} \right) \right)_{CA}$$

$$\succeq \ast \left(\{\mathcal{M}_{g_i} | i \in I\} \cdot \{\mathcal{M}_{f_j} | j \in J\} \right)_{CA}$$

$$= \left(\underset{(i,j) \in I \times J}{\ast} \mathcal{M}_{g_i} \cdot \mathcal{M}_{f_j} \right)_{CA}$$

$$= \left(\underset{(i,j) \in I \times J}{\ast} \mathcal{M}_{g_i \circ f_j} \right)_{CA} \,,$$

where the last equality follows from Lemma 16. The claimed equality follows if the l.h.s. equals $b \in \{0, 1\}$.

It remains to show the claimed equality assuming that the l.h.s. equals $\mathfrak{u}$. By application of Lemma 18 with $G = *_{i \in I} \mathcal{M}_{g_i}$ and $F = *_{j \in J} \mathcal{M}_{f_j}$, we obtain that

$$\left(\left(G^{(0)} \cdot F^{(0)} \right)_{CA} = 0 \right) \wedge \left(\left(G^{(1)} \cdot F^{(1)} \right)_{CA} = 1 \right) .$$

By the definition of matrix multiplication, this is equivalent to

$$\forall B \in \mathcal{P}_{2^k}(T) \colon G^{(0)}_{CB} = 0 \vee F^{(0)}_{BA} = 0 , \tag{1}$$

$$\exists B \in \mathcal{P}_{2^k}(T) \colon G^{(1)}_{CB} = 1 \wedge F^{(1)}_{BA} = 1 . \tag{2}$$

We observe from Definition 17 that for $b \in \mathbb{B}$ we have that

$$\left(\underset{i \in I}{*} \mathcal{M}_{g_i} \right)^{(b)}_{CB} = b \Leftrightarrow \exists i \in I \colon (\mathcal{M}_{g_i})_{CB} = b ;$$

an analogous statement holds for $\mathcal{M}_{f_j}$. Plugging this observation into equations (1) and (2), we get that

$$\forall B \in \mathcal{P}_{2^k}(T) \, \exists (i, j) \in I \times J \colon (\mathcal{M}_{g_i})_{CB} = 0 \vee \left(\mathcal{M}_{f_j} \right)_{BA} = 0 \tag{3}$$

$$\exists B \in \mathcal{P}_{2^k}(T) \, \exists (i, j) \in I \times J \colon (\mathcal{M}_{g_i})_{CB} = 1 \wedge \left(\mathcal{M}_{f_j} \right)_{BA} = 1 . \tag{4}$$

Let $B_0 = \bigcup_{j \in J} f_j(A)$ be the subset of T, to which states in A are mapped to by any f_j. As $|f_j(A)| \leq |A|$, cardinality $|B_0|$ is at most $|J| \cdot |A|$. By assumption $|J| \cdot |A| \leq 2^k$, we obtain $B_0 \in \mathcal{P}_{2^k}(T)$. Since $f_j(A) \subseteq B_0$ by construction, it holds that $\left(\mathcal{M}_{f_j} \right)_{B_0 A} = 1$ for all $j \in J$. Equation (3) thus entails that

$$\exists i \in I \colon (\mathcal{M}_{g_i})_{CB_0} = 0 \Leftrightarrow \exists i \in I \colon g_i(B_0) \not\subseteq C .$$

Hence, there are $i_0 \in I$ and $x \in B_0$ such that $g_{i_0}(x) \notin C$. By construction, $x \in f_{j_0}(A)$ for some $j_0 \in J$, yielding that $(g_{i_0} \circ f_{j_0})(A) = g_{i_0}(f_{j_0}(A)) \not\subseteq C$. We conclude that $\left(\mathcal{M}_{g_{i_0} \circ f_{j_0}} \right)_{CA} = 0$.

Now consider equation (4), which says that there are indices $i_1 \in I$ and $j_1 \in J$ such that $g_{i_1}(B_1) \subseteq C$ and $f_{j_1}(A) \subseteq B_1$. This immediately yields that $(g_{i_1} \circ f_{j_1})(A) \subseteq C$ and thus $\left(\mathcal{M}_{g_{i_1} \circ f_{j_1}} \right)_{CA} = 1$.

The desired equality now follows, because

$$\begin{aligned} \left(\underset{(i,j) \in I \times J}{*} \mathcal{M}_{g_i \circ f_j} \right)_{CA} &\succeq \left(*\{ \mathcal{M}_{g_{i_0} \circ f_{j_0}}, \mathcal{M}_{g_{i_1} \circ f_{j_1}} \} \right)_{CA} \\ &= *\{ \left(\mathcal{M}_{g_{i_0} \circ f_{j_0}} \right)_{CA}, \left(\mathcal{M}_{g_{i_1} \circ f_{j_1}} \right)_{CA} \} \\ &= *\{0, 1\} = \mathfrak{u} \end{aligned}$$

and $b \succeq \mathfrak{u}$ only holds if $b = \mathfrak{u}$. ◀

With Theorem 11 at our disposal, we are ready to prove our main result, Theorem 27. Following Step 1 and Step 2 of the parallel prefix framework given in Section 3.1, we need to compute $\pi_i = t_{x_i} \circ \ldots \circ t_{x_1}$ for all prefixes $x_i \ldots x_1$ of the input string. This computation can be phrased in terms of matrix multiplications, which is shown by the following corollary. It readily follows by inductive application of Theorem 11.

▶ **Corollary 20.** *Suppose that for $i \in [n]$, we are given mappings $E_i \colon \mathbb{B}^\ell \to F_i$ from input symbols $\mathbb{B}^\ell$ to function spaces F_i. Moreover, for all $i \in [n-1]$ the codomain of functions from F_i equals the domain of functions from F_{i+1}. Let $E \colon \mathbb{B}^{n\ell} \to (F_0 \to F_{n-1})$ denote a function that maps a binary string $x \in \mathbb{B}^{n\ell}$ to the composition of the corresponding functions, $E(x) := \circ_{i=0}^{n-1} E_i(x_i)$. Then, for all $x \in \mathbb{T}^{n\ell}$,*

$$\left(\mathcal{M}_{E(\cdot)}\right)_{\mathfrak{u}}(x) = \left(\mathcal{M}_{E_{n-1}(\cdot)}\right)_{\mathfrak{u}}(x_{n-1}) \cdot_{\mathfrak{u}} \left(\mathcal{M}_{E_{n-2}(\cdot)}\right)_{\mathfrak{u}}(x_{n-2}) \cdot_{\mathfrak{u}} \ldots \cdot_{\mathfrak{u}} \left(\mathcal{M}_{E_0(\cdot)}\right)_{\mathfrak{u}}(x_0).$$

Following this insight we observe that evaluation of the functions corresponds to hazard-free matrix-vector multiplication.

▶ **Corollary 21.** *For $j \in J$, let $f_j \colon S \to T$. Assume that $A \in \mathcal{P}_{2^k}(T)$, $S' \in \mathcal{P}_{2^k}(S)$, and $|J| \cdot |S'| \le k$. Then*

$$\left(\left(\underset{j \in J}{*} \mathcal{M}_{f_j}\right) \cdot_{\mathfrak{u}} \left(\underset{s \in S'}{*} e^{(s)}\right)\right)_A = \left(\underset{(j,s) \in J \times S'}{*} e^{(f_j(s))}\right)_A.$$

Proof. Define $g_s \colon \{\bullet\} \to S$ by $g_s(\bullet) := s$ for $s \in S$, such that $(\mathcal{M}_{g_s})_{\{\bullet\}A} = \left(e^{(s)}\right)_A$. By Theorem 11, we thus get that

$$\left(\left(\underset{j \in J}{*} \mathcal{M}_{f_j}\right) \cdot_{\mathfrak{u}} \left(\underset{s \in S'}{*} e^{(s)}\right)\right)_A = \left(\left(\underset{j \in J}{*} \mathcal{M}_{f_j}\right) \cdot_{\mathfrak{u}} \left(\underset{s \in S'}{*} \mathcal{M}_{g_s}\right)\right)_{\{\bullet\}A}$$

$$= \left(\underset{(j,s) \in J \times S'}{*} \mathcal{M}_{f_j \circ g_s}\right)_{\{\bullet\}A}$$

$$= \left(\underset{(j,s) \in J \times S'}{*} e^{(f_j(s))}\right)_A. \qquad \blacktriangleleft$$

Multiplication of all prefixes

Corollary 20 uses the hazard-free matrix product of all input prefixes. This can be efficiently implemented, similarly to the parallel prefix computation approach of Ladner and Fischer.

▶ **Corollary 22** (of [?, Section 2]). *For input matrices $\mathcal{A}_{n-1}, \ldots, \mathcal{A}_0$ of size $\alpha \times \alpha$, there is a circuit computing the hazard-free Boolean matrix multiplication of all prefixes; $\mathcal{A}_i \cdot_{\mathfrak{u}} \ldots \cdot_{\mathfrak{u}} \mathcal{A}_0$, for each $i \in [n]$. The circuit*

has size $\mathcal{O}\left(\alpha^3 n\right)$
and depth $\mathcal{O}\left(\log \alpha \log n\right).$

Proof. By Observation 15, $\cdot_{\mathfrak{u}}$ is associative. For an associative operator, Ladner and Fischer [?] present a family of circuits computing the application of all prefixes of the input. Let c be the size and d the depth of a circuit implementing the operator. The family has asymptotically optimal size $\mathcal{O}(cn)$ and depth $\mathcal{O}(d \log n)$. Corollary 14 offers an implementation of hazard-free Boolean $\alpha \times \alpha$ matrix multiplication of size $c = \alpha^3$ and depth $d = \log \alpha$. ◀

Output step

Before putting the above pieces together to derive our main results, we need to address how the final output is computed, i.e., Step 4. Here, we can exploit that (i) the output does not need to be represented in the universal encoding, and (ii) the universal encoding used in the

previous computations holds additional information that simplifies determining the output in a hazard-free way. We leverage these points in a case analysis to minimize the cost of the output stage.

The input to Step 4 is the vector encoding state s_{i-1} and input x_i for each $i \in \{1 \ldots n\}$; it computes output $o(s_{i-1}, x_i)$. We restrict the output function to an input symbol, as we did for the transition function, i.e., $o_\sigma : S \to \Lambda$ for $\sigma \in \Sigma$ is defined by $o_\sigma(s) := o(s, \sigma)$. In what follows we describe the computation of output bit $o_\sigma(s)_j$ $(= o(s, \sigma)_j)$, for $j \in [m]$. We remark that the preimage of 1 under $o_\sigma(s)_j$ is a set of states. Recall that by Definition 10, the state vector $e^{(s_{i-1})}$ encodes not only state s_{i-1}, but indicates for each subset of states $S' \subseteq S$ (with cardinality less or equal to 2^k) whether $s_{i-1} \in S'$. Thus, we can simply check whether s_{i-1} lies in the set that of states which $o_\sigma(s_{i-1})_j$ maps to 1 (provided its cardinality is at most 2^k).

▶ **Definition 23.** *For an input symbol $\sigma \in \Sigma$ and $j \in [m]$, we define*

$$A_{\sigma,j} := \{s \in S | o(s, \sigma)_j = 1\} .$$

Note that if $2^k < |S|$, we reduce the size of the universal encoding by encoding only sets of size less or equal to 2^k. Hence, we might not have computed a bit indicating whether $s_{i-1} \in A_{\sigma,j}$ as an entry of the vector $e^{(s_{i-1})}$. However, essentially all output bits are conveniently available if $\max_{\sigma \in \Sigma, j \in [m]} |A_{\sigma,j}| \leq 2^k$, which is captured by the following lemma.

▶ **Lemma 24.** *For $k \in \mathbb{N}$ and $S' \subseteq S$, $\Sigma' \subseteq \Sigma$, if $\max_{\sigma \in \Sigma, j \in [m]} |A_{\sigma,j}| \leq 2^k$ we have that*

$$\underset{s \in S', \sigma \in \Sigma'}{*} o(s, \sigma)_j = \underset{s \in S', \sigma \in \Sigma'}{*} e^{(s)}_{A_{\sigma,j}} . \tag{5}$$

Proof. First, as $2^k \geq \max_{\sigma \in \Sigma', j \in [m]} |A_{\sigma,j}|$, by Definition 10 every $A_{\sigma,j}$ is encoded in $e^{(s)}$, i.e., every entry $e^{(s)}_{A_{\sigma,j}}$ exists.

Next, we distinguish three cases for every evaluation of the l.h.s. of (5): 1, 0, and $\mathfrak{u}$. In the first case, $*_{s \in S', \sigma \in \Sigma} o(s, \sigma)_j = 1$, we apply Definition 23 and Definition 10 to show the claim, as follows.

$$\underset{s \in S', \sigma \in \Sigma'}{*} o(s, \sigma)_j = 1$$

$$\Leftrightarrow \forall s \in S', \sigma \in \Sigma' : o(s, \sigma)_j = 1$$

$$\overset{Def\ 23}{\Leftrightarrow} \forall s \in S', \sigma \in \Sigma' : s \in A_{\sigma,j}$$

$$\overset{Def\ 10}{\Leftrightarrow} \forall s \in S', \sigma \in \Sigma' : e^{(s)}_{A_{\sigma,j}} = 1$$

$$\Leftrightarrow \underset{s \in S', \sigma \in \Sigma'}{*} e^{(s)}_{A_{\sigma,j}} = 1$$

The second case, $*_{s \in S', \sigma \in \Sigma} o(s, \sigma)_j = 0$, is treated analogously.

$$\underset{s \in S', \sigma \in \Sigma'}{*} o(s, \sigma)_j = 0$$

$$\Leftrightarrow \forall s \in S', \sigma \in \Sigma' : o(s, \sigma)_j = 0$$

$$\overset{Def\ 23}{\Leftrightarrow} \forall s \in S', \sigma \in \Sigma' : s \notin A_{\sigma,j}$$

$$\overset{Def\ 10}{\Leftrightarrow} \forall s \in S', \sigma \in \Sigma' : e^{(s)}_{A_{\sigma,j}} = 0$$

$$\Leftrightarrow \underset{s \in S', \sigma \in \Sigma'}{*} e^{(s)}_{A_{\sigma,j}} = 0$$

In case the l.h.s. of (5) evaluates to $\mathfrak{u}$ the statement follows from the first two cases. As we showed equivalence in case 1 and 0, we deduce, that in case the l.h.s. of (5) evaluates to $\mathfrak{u}$, the r.h.s. of (5) also evaluates to $\mathfrak{u}$.   ◀

If the size of the largest preimage $(\max_{\sigma \in \Sigma, j \in [m]} |A_{\sigma,j}|)$ exceeds 2^k not all preimages have a corresponding entry in $e^{(s_{i-1})}$. We need to compute whether s_{i-1} is in the preimage. To this purpose, we define a cover of $A_{\sigma,j}$ with sets of size at most 2^k.

▶ **Definition 25.** *For an input symbol $\sigma \in \Sigma$, $j \in [m]$, and $k \in \mathbb{N}$ we define $\mathcal{A}^k_{\sigma,j}$, the cover of $A_{\sigma,j}$ contatining sets of cardinality smaller or equal to 2^k:*

$$\mathcal{A}^k_{\sigma,j} := \begin{cases} \{A_{\sigma,j}\} & \text{if } |A_{\sigma,j}| \le 2^k \,, \\ \{A \subseteq A_{\sigma,j} \,|\, |A| = 2^k\} & \text{else.} \end{cases}$$

If the state of the transducer is in one of the sets in $\mathcal{A}^k_{\sigma,j}$, then it is also in $A_{\sigma,j}$. All sets in $\mathcal{A}^k_{\sigma,j}$ have a corresponding entry in the state vector. We can take the or over all entries to see whether the transducer is in a state of $A_{\sigma,j}$ and hence whether it outputs 1.

▶ **Lemma 26.** *For $k \in \mathbb{N}$, $j \in [m]$, and $S' \subseteq S$, $\Sigma' \subseteq \Sigma$ such that $2^k < \max_{\sigma \in \Sigma, j \in [m]} |A_{\sigma,j}|$, if $|S'| \le 2^k$ we have that*

$$\mathop{\text{\Large $*$}}_{s \in S', \sigma \in \Sigma'} o(s,\sigma)_j = \mathop{\text{\Large $*$}}_{\sigma \in \Sigma'} \bigvee_{A \in \mathcal{A}^k_{\sigma,j}} \mathop{\text{\Large $*$}}_{s \in S'} e^{(s)}_A \tag{6}$$

Proof. Every $A \in \mathcal{A}^k_{\sigma,j}$ has cardinality smaller or equal to 2^k. Hence, there is an entry in vector $e^{(s)}$ corresponding to A, i.e., entry $e^{(s)}_A$ exists in the encoding.

From its definition we observe that $\mathcal{A}^k_{\sigma,j}$ is indeed a cover of $A_{\sigma,j}$, i.e.,

$$\bigcup_{A \in \mathcal{A}^k_{\sigma,j}} A = A_{\sigma,j} \,. \tag{7}$$

Moreover, Definition 25 ensures that every subset of $A_{\sigma,j}$ of size at most 2^k is contained in at least one set from $\mathcal{A}^k_{\sigma,j}$. On the other hand each $A \in \mathcal{A}^k_{\sigma,j}$ is a subset of $A_{\sigma,j}$. Hence, the assumption that $|S'| \le 2^k$ implies that

$$S' \subseteq A_{\sigma,j} \Leftrightarrow \exists A \in \mathcal{A}^k_{\sigma,j} : S' \subseteq A \,. \tag{8}$$

We make a case distinction on every possible evaluation of the l.h.s. of (6) and show the

equality for each case. The first case is $*_{s \in S', \sigma \in \Sigma'} e^{(s)}_{A_{\sigma,j}} = 1$. We get that

$$\mathop{*}_{s \in S', \sigma \in \Sigma'} o(s,\sigma)_j = 1$$

$$\Leftrightarrow \forall s \in S', \sigma \in \Sigma' : o(s,\sigma)_j = 1$$

$$\Leftrightarrow \forall s \in S', \sigma \in \Sigma' : s \in A_{\sigma,j}$$

$$\Leftrightarrow \forall \sigma \in \Sigma' : S' \subseteq A_{\sigma,j}$$

$$\overset{(8)}{\Leftrightarrow} \forall \sigma \in \Sigma' : \exists A \in A_{\sigma,j} : S' \subseteq A$$

$$\Leftrightarrow \forall \sigma \in \Sigma' : \exists A \in A_{\sigma,j} : \forall s \in S' : s \in A$$

$$\Leftrightarrow \forall \sigma \in \Sigma' : \exists A \in A_{\sigma,j} : \forall s \in S' : e^{(s)}_A = 1$$

$$\Leftrightarrow \forall \sigma \in \Sigma' : \exists A \in A_{\sigma,j} : \mathop{*}_{s \in S'} e^{(s)}_A = 1$$

$$\Leftrightarrow \forall \sigma \in \Sigma' : \bigvee_{A \in \mathcal{A}^k_{\sigma,j}} \mathop{*}_{s \in S'} e^{(s)}_A = 1$$

$$\Leftrightarrow \mathop{*}_{\sigma \in \Sigma'} \bigvee_{A \in \mathcal{A}^k_{\sigma,j}} \mathop{*}_{s \in S'} e^{(s)}_A = 1.$$

The second case, $*_{s \in S', \sigma \in \Sigma'} e^{(s)}_{A_{\sigma,j}} = 0$, is treated similarly, where now $S' \cap A_{\sigma,j} = \emptyset$ for each $\sigma \in \Sigma'$.

$$\mathop{*}_{s \in S', \sigma \in \Sigma} o(s,\sigma)_j = 0$$

$$\Leftrightarrow \forall s \in S', \sigma \in \Sigma' : o(s,\sigma)_j = 0$$

$$\Leftrightarrow \forall s \in S', \sigma \in \Sigma' : s \notin A_{\sigma,j}$$

$$\Leftrightarrow \forall s \in S', \sigma \in \Sigma', A \in \mathcal{A}^k_{\sigma,j} : s \notin A$$

$$\Leftrightarrow \forall s \in S', \sigma \in \Sigma', A \in \mathcal{A}^k_{\sigma,j} : e^{(s)}_A = 0$$

$$\Leftrightarrow \forall \sigma \in \Sigma', A \in \mathcal{A}^k_{\sigma,j} : \mathop{*}_{s \in S'} e^{(s)}_A = 0$$

$$\Leftrightarrow \forall \sigma \in \Sigma' : \bigvee_{A \in \mathcal{A}^k_{\sigma,j}} \mathop{*}_{s \in S'} e^{(s)}_A = 0$$

$$\Leftrightarrow \mathop{*}_{\sigma \in \Sigma'} \bigvee_{A \in \mathcal{A}^k_{\sigma,j}} \mathop{*}_{s \in S'} e^{(s)}_A = 0$$

In the final case, i.e., that the l.h.s. of (6) evaluates to $\mathtt{u}$, equality follows from the equivalence established in the previous two cases.   ◀

Main theorem

We are left with the task of showing that indeed the obtained circuit is correct, i.e., prove Theorem 27. The correctness of the construction is proven foremost by application of Corollary 20 and Corollary 21. The proof is mainly concerned with establishing the size and depth bound for the obtained circuit. Without going into detail, for constant $|S|$ the reader should be convinced that the depth of the circuit is logarithmic in n, because all operations except for Step 2 can be computed in parallel, while Step 2 exploits the associativity of matrix multiplication to obtain a circuit of depth logarithmic in n. The size of the circuit is
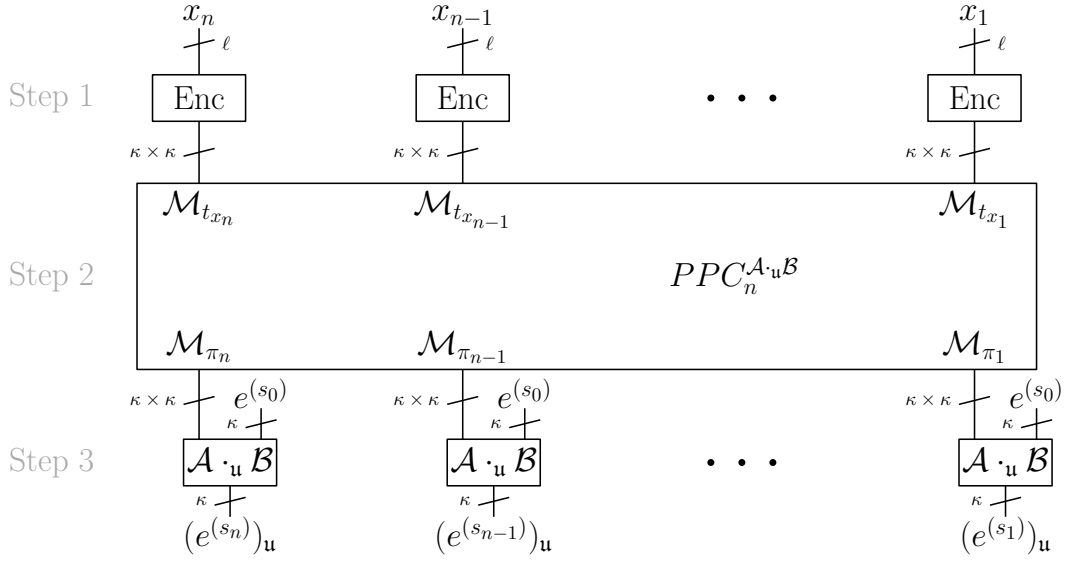
Figure 2 Steps 1 to 3 of the circuit implementing the transcription function. Enc denotes the computation of the universal encoding by the generic construction of [**?**]. Hazard-free Boolean matrix multiplication is denoted by $\mathcal{A} \cdot_{\mathfrak{u}} \mathcal{B}$.

linear in n, as Step 1, Step 3 and Step 4 each use a constant number of operations for each input symbol and Step 2 can be performed asymptotically optimally with a linear number of operations.

▶ **Theorem 27.** *For any integers $k \in \mathbb{N}$, $\ell, m, n \in \mathbb{N}_{>0}$ (with $k \leq n$) and Mealy machine $T = (S, s_0, \Sigma = \mathbb{B}^\ell, \Lambda \subseteq \mathbb{B}^m, t, o)$, there is a k-bit hazard-free circuit implementing $\tau_{T,n}$. For $\kappa \coloneqq \sum_{i=0}^{\min\{|S|, 2^k\}} \binom{|S|}{i}$ and $\lambda \coloneqq \min\{m, 2^{|S| \cdot |\Sigma|}\}$ the circuit has*

$$\text{size} \quad \mathcal{O}\left((\kappa^3 + (2^\ell/\ell)\kappa^2 + 2^\ell \kappa \lambda)n\right)$$
$$\text{and depth} \quad \mathcal{O}\left(\log \kappa \log n + \ell\right).$$
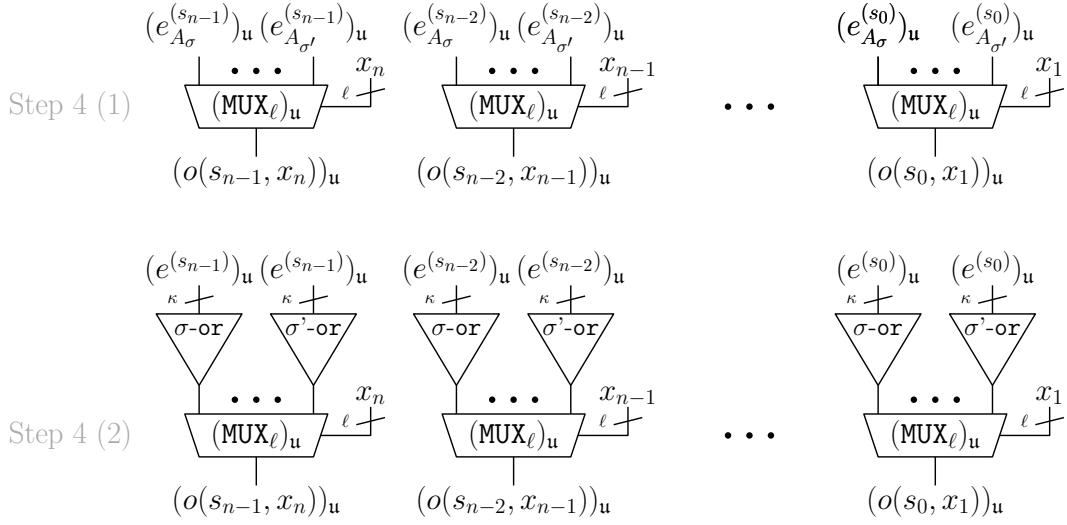
Proof. We show that there is a circuit computing the hazard-free extension of the transcription function $(\tau_{T,n})_{\mathfrak{u}}(x)$ for every $x \in \Sigma^n$, where we replace at most k many bits with $\mathfrak{u}$'s. We follow the steps of the PPC framework presented in Section 3.1 to compute output $((\tau_{T,n})_{\mathfrak{u}}(x))_i = o_{\mathfrak{u}}(s_{i-1}, x_i)$ at each position $i \in \{1 \ldots n\}$.

The block diagram of steps 1 to 3 of the circuit is depicted in Figure 2. The circuit mostly consists of hazard-free matrix multiplication blocks. In particular, the n-input parallel prefix circuit is an arrangement of hazard-free matrix multiplication blocks.

Step 1 computes the universal encoding of the restricted transition function t_{x_i} (a $\kappa \times \kappa$ matrix) from input x_i. Noting that the computation in Step 1 evaluates a function from $\mathbb{B}^\ell$ to $\mathbb{B}^{\kappa^2}$, we can directly apply Theorem 19 for each output bit seperately. Hence, there is a hazard-free encoding circuit of size $\mathcal{O}((2^\ell/\ell)\kappa^2)$ and depth $\mathcal{O}(\ell)$ implementing this step. Thus, the computation of the universal encoding of t_{x_i} for each i in parallel has size $\mathcal{O}((2^\ell/\ell)\kappa^2 n)$ and depth $\mathcal{O}(\ell)$.

Step 2 computes the encoding $\mathcal{M}_{\pi_i}$ of the composition $\pi_i = t_{x_i} \circ \ldots \circ t_{x_1}$. We define $E_j(x_j) = t(\cdot, x_j)$ for $j \in \{1, \ldots, i\}$ such that for $E(x) = \circ_{j=1}^i E_j(x_j)$ we have

$$\left(\mathcal{M}_{\pi_i}\right)_{\mathfrak{u}} = \left(\mathcal{M}_{E(x)}\right)_{\mathfrak{u}} = \left(\mathcal{M}_{E(\cdot)}\right)_{\mathfrak{u}}(x).$$

■ **Figure 3** Step 4 of the circuit implementing the transcription function, assuming $m = 1$, for the case every preimage is encoded in the state vector (1) and the case that there is at least one preimages not encoded in the state vector (2). To increase readability we do not enumerate all elements of Σ, but use σ, σ' and dots to denote that the step is repeated for every element of Σ.

Application of Corollary 20 then yields

$$\left(\mathcal{M}_{E(\cdot)}\right)_{\mathtt{u}}(x) = \left(\mathcal{M}_{E_j(\cdot)}\right)_{\mathtt{u}}(x_j) \cdot_{\mathtt{u}} \left(\mathcal{M}_{E_{j-1}(\cdot)}\right)_{\mathtt{u}}(x_{j-1}) \cdot_{\mathtt{u}} \ldots \cdot_{\mathtt{u}} \left(\mathcal{M}_{E_1(\cdot)}\right)_{\mathtt{u}}(x_1). \tag{9}$$

By Corollary 22, there is an efficient circuit computing the r.h.s. of (9) for each i. The Corollary gives also size $\mathcal{O}(\kappa^3 n)$ and depth $\mathcal{O}(\log \kappa \log n)$ for Step 2.

Step 3 computes the column unit vector corresponding to the i-th state, i.e., evaluation of the composition π_i on the initial state s_0. By Corollary 21, we can compute $\left(e^{(s_i)}\right)_{\mathtt{u}}$ by matrix-vector multiplication of $(\mathcal{M}_{\pi_i})_{\mathtt{u}}$ and $\left(e^{(s_0)}\right)_{\mathtt{u}} = e^{(s_0)}$. Hence, by Corollary 14 there is a k-bit hazard-free circuit that computes s_i. Evaluation of π_i in parallel for each i yields size $\mathcal{O}(\kappa^2 n)$ and depth $\mathcal{O}(\log \kappa)$ for Step 3.

Finally, Step 4 computes the i-th output $o(s_{i-1}, x_i)$ for each $i \in \{1 \ldots n\}$. W.l.o.g., assume that the width of an output symbol is 1, i.e., $m = 1$ and hence $\Lambda \subseteq \mathbb{B}$ (otherwise repeat the computation for each output bit separately). Viewing the computaion in Step 4 as a function from $\mathbb{B}^{\kappa+\ell}$ to $\mathbb{B}$, we could apply Theorem 19 to obtain a circuit of size $\mathcal{O}(2^{\kappa+\ell}/(\kappa + \ell))$ and depth $\mathcal{O}(\kappa + \ell)$. We now show how to obtain better results, bounding the size by $\mathcal{O}(2^\ell \kappa)$ with a circuit of depth $\mathcal{O}(\log \kappa + \ell)$.

Recall Definition 23. As $j = 0$ we omit j from the notation and write A_σ and $\mathcal{A}_\sigma^k$ instead of $A_{\sigma,j}$ and $\mathcal{A}_{\sigma,j}^k$. First, consider the case that $\max_{\sigma \in \Sigma} |A_\sigma| \leq 2^k$. A depiction of this case is given in Figure 3 (1). By directly using the outputs of the gates computing the respective bits of (the universal encoding of) the state vector, we readily obtain $o(s_{i-1}, \sigma)$ for each $\sigma \in \Sigma$. Thus, we are left with the task to choose the output corresponding to input x_i. We can do so by using a MUX. By Corollary 13, the implementation $\mathtt{MUX}(e_{A_\sigma}^{(s_{i-1})}, \ldots, e_{A_{\sigma'}}^{(s_{i-1})}, x_i)$ has size $\mathcal{O}(2^\ell)$ and depth $\mathcal{O}(\ell)$, where σ, σ' are representatives for every input symbol in Σ. If $m > 1$, the multiplexer is copied m times and wired accordingly. Step 4 can be performed in parallel for each i, hence the resulting size and depth bounds are $\mathcal{O}(2^\ell m n)$ and $\mathcal{O}(\ell)$, respectively.

The other case is that $\max_{\sigma \in \Sigma} |A_\sigma| > 2^k$. A depiction of the corresponding circuit is given in Figure 3 (2), where σ-$\mathtt{or}$ denotes the $\mathtt{or}$-tree over all $e_A^{(s_i)}$ for $A \in \mathcal{A}_\sigma^k$. Here, we

compute the output $o(s_{i-1}, \sigma)$ as the $\mathtt{or}$ over all entries corresponding to $\mathcal{A}_\sigma^k$ in the state vector. Correctness readily follows from Lemma 26. To bound size and depth of the resulting circuit, observe that the cardinality of $\mathcal{A}_\sigma^k$ is bounded by $\binom{|S|}{2^k}$, hence each $\mathtt{or}$-tree has size $\mathcal{O}(\binom{|S|}{2^k})$ and depth $\mathcal{O}(\log \binom{|S|}{2^k})$. As in the previous case, the i-th output is selected by a multiplexer. Hence, applying Corollary 13, we get that in this case Step 4 requires size $\mathcal{O}(\binom{|S|}{2^k} 2^\ell mn)$ and depth $\mathcal{O}(\log \binom{|S|}{2^k} + \ell)$.

Furthermore, there is a natural cap on m. Similar to the previous paragraph, consider each bit of the output seperately, i.e., for each position of an output symbol consider the function $o\colon S \times \Sigma \to \mathbb{B}$. For inputs from S and Σ there are $2^{|S| \cdot |\Sigma|}$ different one bit functions. Hence, we can enumerate all possible output functions. If $m > 2^{|S| \cdot |\Sigma|}$, we simply compute and reuse as often as needed the output of each of these possible functions, rather than replicating computations for identical output bits.

Finally, we can derive the asymptotic size and depth of the presented k-bit hazard-free implementation of the transcription function $\tau_{T,n}$. We distinguish two cases depending on the size of the largest preimage of o. In both cases m is capped at $2^{|S| \cdot |\Sigma|}$ as discussed above. Assume $\max_{\sigma \in \Sigma} |A_\sigma| \leq 2^k$, then the presented circuit has

> size: $\mathcal{O}\left((\kappa^3 + (2^\ell/\ell)\kappa^2 + 2^\ell m)n\right)$, and
>
> depth: $\mathcal{O}\left(\log \kappa \log n + \ell\right)$.

In case $\max_{\sigma \in \Sigma} |A_\sigma| > 2^k$ (and hence $|S| > 2^k$) we bound $\binom{|S|}{2^k}$ by κ, such that the circuit has

> size: $\mathcal{O}\left((\kappa^3 + (2^\ell/\ell)\kappa^2 + 2^\ell \kappa m)n\right)$, and
>
> depth: $\mathcal{O}\left(\log \kappa \log n + \ell\right)$. ◄

We remark that the main result holds also for $\Sigma \subseteq \mathbb{B}^\ell$, when choosing an (arbitrary) extension for the transition function to domain $S \times \mathbb{B}^\ell$. However, the choice how to extend the transition function t matters for the behavior of the hazard-free state machine. The desicion how to treat non-input symbols is important, because it is possible that an unstable input resolves to such a non-input symbol. If this happens, the choice of where t maps such resolutions to affects the value the hazard-free extension takes. A "bad" extension may result in unstable output without need, decreasing the utility of the constructed circuit. The task of finding a useful extension is nontrivial and depends on the application. A detailed discussion at hand of an example is given in Appendix A.

4   Conclusion

Any generic construction for hazard-free circuits incurs an exponential blow-up in circuit complexity [**?**]. In this work, we present a generic construction for transducers that is asymptotically optimal in the size of the inputs n, but yields an exponential overhead in the size of the transducer. By Theorem 27, for a transducer with $|S|$ states we obtain a circuit of size $\mathcal{O}(\kappa^3 n + 2^\ell \kappa^2 n + 2^\ell \kappa \lambda n)$ and depth $\mathcal{O}(\log \kappa \log n + \ell)$, where ℓ is the maximum number of bits of an input symbol, k is the maximum number of uncertain bits in the input, and κ^2 is the size of the universal function encoding. The universal encoding of functions is a key ingredient to the construction of the circuit. Together with Theorem 11, we show correctness of the construction. The Theorem states that, for the chosen encoding, the superposition of function composition is the matrix product of the superpositions of both functions. We specify our findings in Corollary 8 and Corollary 9.

This opens the stage for a k-bit hazard-free implementation of addition. We remark that, to make hazard-freeness meaningful in the context of addition, one has to limit the uncertainty of the input with respect to the encoded sum and choose appropriate encodings. The relevant definitions exceed the scope of this submission. In a preliminary publication [**?**], we show that it is possible to apply this construction to obtain efficient circuits for addition that avoid certain hazards. This demonstrates that the construction we present here is of interest and, in our view, has surprising consequences.

A   Extension of the input encoding

The main result, Theorem 27, assumes that each binary string of length ℓ is part of the input alphabet, i.e., $\Sigma = \mathbb{B}^\ell$. In this section we discuss the case where Σ is only a strict subset of $\mathbb{B}^\ell$. We choose extensions of the transition function $t\colon \Sigma \times S$ and ouput function $o\colon \Sigma \times S$ to domain $\mathbb{B}^\ell \times S$. The choice of extension is arbitrary, for any transition function and output function defined on $\mathbb{B}^\ell$ the construction described in the proof of the main theorem computes the hazard-free extension of the transcription function. An arbitrary extension may lead to undesirable results, however. We discuss this issue at hand of an example. Note that both, t and o have to be extended and the same problems may arise in both functions. We construct an example, that provides intuition on the problem, by extending the output function o. The transition function t is not discussed in further detail.

Example

Consider a simple finite-state transducer with a single state ($S = \{0\}$) with single bit outputs and output alphabet $\Lambda = \mathbb{B}$. Input symbols have three bits ($\ell = 3$) and the output function computes the `and` over the three input bits[5]. See Table 3 for the output table. For the sake of the example, assume that input 000 never occurs, such that the input alphabet is defined by $\Sigma = \mathbb{B}^3 \setminus \{000\}$. Even though the example is artificial, in the sense that it could be solved by simpler circuits, it is worth the discussion as it provides intuition to the problem.

Unknown inputs (which we model by $\mathfrak{u}$) may arise from transitioning input signals. Due to the analog behavior of signal transitions it might be the case that a bit is neither digital 0 nor 1 for some time during the transition. The signal can float near the threshold for 0 or 1 and the circuit designer can not be sure wheter the electronics interpret the signal as digital 0 or 1.

Extension of the output function

For an application of the framework we have to define an extension to o on input 000. As there is only a single output bit, we are left with exactly two choices:

$$o(0, 000) = 0\,, \text{ or} \tag{1}$$

$$o(0, 000) = 1\,. \tag{2}$$

Assume the input switches from 001 to 010, such that the input transitions via $0\mathfrak{u}\mathfrak{u}$ for some time. We choose extension (2), then the framework computes the hazard-free extension

---

[5] The `and` over the three bits is defined by two sequential `and` gates (as given in Table 1).

■ **Table 3** Output function of the example transducer and extensions (1) and (2).

σ	$o(\sigma, 0)$	(1)	(2)
000	-	0	1
001	0	0	0
010	0	0	0
100	0	0	0
011	0	0	0
101	0	0	0
110	0	0	0
111	1	1	1

of the output function

$$o_{\mathfrak{u}}(0, 0\mathfrak{u}\mathfrak{u})$$
$$= *\{o(0, 000), o(0, 001), o(0, 010), o(0, 011)\}$$
$$= *\{1, 0\}$$
$$= \mathfrak{u}\,.$$

In the sense of hazard-freeness this is perfectly fine, because we compute the most precise output for all specified inputs and the extension. However, input 001 and 010 both produce output 0. Hence, the output has no transition when switching from 001 to 010. But, due to the chosen extension the output is $\mathfrak{u}$ during the transition.

Contrary to that, if we choose extension (1) we obtain

$$o_{\mathfrak{u}}(0, 0\mathfrak{u}\mathfrak{u})$$
$$= *\{o(0, 000), o(0, 001), o(0, 010), o(0, 011)\}$$
$$= *\{0\}$$
$$= 0\,.$$

Hence, output signal remains stable during the transition of the inputs. Also one can verify that for each input transition the output remains unless there is a transition in the output (for any input transition to or from 111). Hence, the advisable choice of the extension for the specific application is extension (1).

To conclude, the example shows that finding an extension of the transition function, in case Σ is a strict subset of $\mathbb{B}^{\ell}$, is not a straightforward task. We refrain from defining a notion of the optimum choice and generating the according extension, as this is beyond the scope of this work.

B   Bound on k

We present an extension of our earlier example, the shift transducer. Here, we switch from output 0 to output 1 only if we see two consecutive 1's in the input. Vice versa, we switch from 1's to 0's only we see two consecutive 0 inputs. The transducer has three states $S = \{0, 1, 2\}$. It is depicted in Figure 4.

What we seek to demonstrate is the impact of bounding k. Accordingly, we choose $k = 1$, such that $2^k = 2 < 3 = |S|$. Thus, the encoding takes into account $\mathcal{P}_2(S) = \mathcal{P}(S) \setminus S$, i.e.,
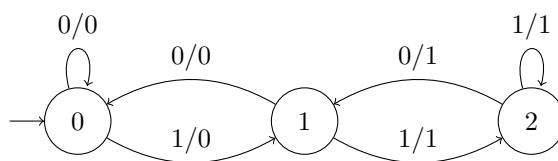
Figure 4 Extension of the first example transducer, that shifts to outputting ones only after reception of two consecutive ones. Vice versa the transducer shifts to outputting zeros upon reception of two consecutive zeros.

all proper subsets of S. For restricted transition functions t_0, t_1, and $t_{\mathfrak{u}}$, the encoding yields the following matrices.

$$
\mathcal{M}_{t_0} =
\begin{array}{c@{\,}c}
& \begin{array}{ccccccc} \emptyset & \{0\} & \{1\} & \{2\} & \{0,1\} & \{1,2\} & \{0,2\} \end{array} \\
\begin{array}{c} \emptyset \\ \{0\} \\ \{1\} \\ \{2\} \\ \{0,1\} \\ \{1,2\} \\ \{0,2\} \end{array} &
\left(\begin{array}{ccccccc}
1 & 0 & 0 & 0 & 0 & 0 & 0 \\
1 & 1 & 1 & 0 & 1 & 0 & 0 \\
1 & 0 & 0 & 1 & 0 & 0 & 0 \\
1 & 0 & 0 & 0 & 0 & 0 & 0 \\
1 & 1 & 1 & 1 & 1 & 1 & 1 \\
1 & 0 & 0 & 1 & 0 & 0 & 0 \\
1 & 1 & 1 & 0 & 1 & 0 & 0
\end{array}\right)
\end{array}
$$

$$
\mathcal{M}_{t_1} =
\begin{array}{c@{\,}c}
& \begin{array}{ccccccc} \emptyset & \{0\} & \{1\} & \{2\} & \{0,1\} & \{1,2\} & \{0,2\} \end{array} \\
\begin{array}{c} \emptyset \\ \{0\} \\ \{1\} \\ \{2\} \\ \{0,1\} \\ \{1,2\} \\ \{0,2\} \end{array} &
\left(\begin{array}{ccccccc}
1 & 0 & 0 & 0 & 0 & 0 & 0 \\
1 & 0 & 0 & 0 & 0 & 0 & 0 \\
1 & 1 & 0 & 0 & 0 & 0 & 0 \\
1 & 0 & 1 & 1 & 0 & 1 & 0 \\
1 & 1 & 0 & 0 & 0 & 0 & 0 \\
1 & 1 & 1 & 1 & 1 & 1 & 1 \\
1 & 0 & 1 & 1 & 0 & 1 & 0
\end{array}\right)
\end{array}
$$

$$
\mathcal{M}_{t_{\mathfrak{u}}} =
\begin{array}{c@{\,}c}
& \begin{array}{ccccccc} \emptyset & \{0\} & \{1\} & \{2\} & \{0,1\} & \{1,2\} & \{0,2\} \end{array} \\
\begin{array}{c} \emptyset \\ \{0\} \\ \{1\} \\ \{2\} \\ \{0,1\} \\ \{1,2\} \\ \{0,2\} \end{array} &
\left(\begin{array}{ccccccc}
1 & 0 & 0 & 0 & 0 & 0 & 0 \\
1 & \mathfrak{u} & \mathfrak{u} & 0 & \mathfrak{u} & 0 & 0 \\
1 & \mathfrak{u} & 0 & \mathfrak{u} & 0 & 0 & 0 \\
1 & 0 & \mathfrak{u} & \mathfrak{u} & 0 & \mathfrak{u} & 0 \\
1 & 1 & \mathfrak{u} & \mathfrak{u} & \mathfrak{u} & \mathfrak{u} & \mathfrak{u} \\
1 & \mathfrak{u} & \mathfrak{u} & 1 & \mathfrak{u} & \mathfrak{u} & \mathfrak{u} \\
1 & \mathfrak{u} & 1 & \mathfrak{u} & \mathfrak{u} & \mathfrak{u} & 0
\end{array}\right)
\end{array}
$$

In matrix $\mathcal{M}_{t_{\mathfrak{u}}}$, we see again the same behavior as in the naive encoding of the shift transducer. Columns corresponding to $\{0,1\}$, $\{1,2\}$, and $\{0,2\}$ contain only 0 and $\mathfrak{u}$ entries. These columns hence cannot prevent the resolution from containing the all 0's vector, which cannot be mapped to a single state in a way that guarantees correct output.

However, columns corresponding to $\{0\}$, $\{1\}$, and $\{2\}$ can be mapped to states without uncertainty, presuming the next input symbol allows it. Taking a look at Figure 4, we can see that the transducer goes back to state 0 after reception of two consecutive 0 inputs, regardless of the previous state. For the sake of this example, consider the input sequence

$\mathfrak{u}00$. Hence, we need to calculate matrix corresponding to $t_0 \circ t_0 \circ t_\mathfrak{u}$:

$$\mathcal{M}_{t_0 \circ t_0 \circ t_\mathfrak{u}} = \mathcal{M}_{t_0} \cdot \mathcal{M}_{t_0} \cdot \mathcal{M}_{t_\mathfrak{u}} =$$

$$\mathcal{M}_{t_0} \cdot \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & \mathfrak{u} & \mathfrak{u} & \mathfrak{u} & \mathfrak{u} & \mathfrak{u} \\ 1 & 0 & \mathfrak{u} & \mathfrak{u} & 0 & \mathfrak{u} & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & \mathfrak{u} & \mathfrak{u} & \mathfrak{u} \\ 1 & 0 & \mathfrak{u} & \mathfrak{u} & 0 & \mathfrak{u} & 0 \\ 1 & 1 & \mathfrak{u} & \mathfrak{u} & \mathfrak{u} & \mathfrak{u} & \mathfrak{u} \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & \mathfrak{u} & \mathfrak{u} & \mathfrak{u} \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & \mathfrak{u} & \mathfrak{u} & \mathfrak{u} \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & \mathfrak{u} & \mathfrak{u} & \mathfrak{u} \end{pmatrix}.$$

Here it becomes apparent why we need the assumption $|A| \cdot |J| \leq 2^k$ in Theorem 27. In matrix $\mathcal{M}_{t_0 \circ t_0 \circ t_\mathfrak{u}}$ only column unit vectors corresponding to singletons can be mapped to stable states. Column vectors corresponding to sets $\{0,1\}$, $\{1,2\}$ and $\{0,2\}$ contain only 0's and $\mathfrak{u}'s$, no 1's. Hence, they cannot be "properly" mapped to any state. We apply Theorem 27 to the last step, i.e., multiplication of $\mathcal{M}_{t_0}$ and $\mathcal{M}_{t_0 \circ t_\mathfrak{u}}$. Here, set J is mapped to $\mathrm{res}(0\mathfrak{u})$ yielding $|J| = 2$. As $k = 1$, condition $|A| \cdot |J| \leq 2^k$ is only satisfied for $|A| = 1$ (A cannot be the empty set). Thus, the condition is only satisfied for singleton sets, which is what we observed from the example.