

On Doubly-Efficient Interactive Proof Systems

Oded Goldreich¹

July 30, 2017

¹Department of Computer Science, Weizmann Institute of Science, Rehovot, ISRAEL.
oded.goldreich@weizmann.ac.il

Abstract

An interactive proof system is called doubly-efficient if the prescribed prover strategy can be implemented in polynomial-time and the verifier’s strategy can be implemented in almost-linear-time. Such proof systems, introduced by Goldwasser, Kalai, and Rothblum (*JACM*, 2015), make the benefits of interactive proof system available to real-life agents who are restricted to polynomial-time computation.

We survey the main two constructions of doubly-efficient interactive proof system: the system for sets having (uniform) bounded-depth circuits (due to Goldwasser, Kalai and Rothblum (*JACM*, 2015)), and the system for sets that are recognized in polynomial-time and small space (due to Reingold, Rothblum, and Rothblum (*STOC*, 2016)). Our presentation of the first construction is quite complete (except for some technical details), but for the second construction we only provide an overview. In addition, we present simpler constructions of doubly-efficient interactive proof system for “locally characterizable” sets like t -no-CLIQUE (due to Goldreich and Rothblum (*ECCC*, 2017)).

Contents

1	Introduction	2
1.1	The notion of doubly-efficient interactive proof systems	2
1.2	Doubly-efficient NP-proof systems	3
1.3	The power of doubly-efficient interactive proof systems	4
1.4	An upper bound on doubly-efficient interactive proof systems	5
1.5	On doubly-efficient argument systems	6
1.6	Preliminaries: The sum-check protocol	7
1.7	Organization	7
2	Simple doubly-efficient interactive proof systems	8
2.1	The case of t -CLIQUE	8
2.2	The general result	11
2.2.1	A natural class: locally-characterizable sets	11
2.2.2	Proof of Theorem 2.2	12
2.2.3	Generalization: round versus computation trade-off	14
2.2.4	Extension to a wider class	14
3	On the doubly-efficient interactive proof systems of GKR	16
3.1	Overview	16
3.2	The main module	17
3.3	Evaluating $\hat{\phi}_i$	20
3.4	Details	22
3.4.1	Applying the sum-check protocol to Eq. (3.3)	22
3.4.2	The highly uniform circuits in use	22
4	Overview of the doubly-efficient interactive proof systems of RRR	25
4.1	High level structure	25
4.2	Warm-up: Batch verification for NP	27
4.3	Batch verification for unambiguous IP	30
	Appendix: Defining Interactive Proofs and Arguments	33
A.1	The basic definition of interactive proofs	34
A.2	On computationally bounded provers: an overview	35
A.2.1	How powerful should the prover be?	36
A.2.2	Computational-soundness	37
	Acknowledgements	37
	References	38

Chapter 1

Introduction

The notion of interactive proof systems, put forward by Goldwasser, Micali, and Rackoff [26], and the demonstration of their power by Lund, Fortnow, Karloff, and Nisan [32] and Shamir [37] are among the most celebrated achievements of complexity theory. Recall that an interactive proof system for a set S is associated with an interactive verification procedure, V , that can be made to accept any input in S but no input outside of S . That is, there exists an interactive strategy for the prover that makes V accept any input in S , but no strategy can make V accept an input outside of S , except with negligible probability. (See the appendix for a formal definition of interactive proofs, and [19, Chap. 9] for a wider perspective.)

The original definition does not restrict the complexity of the strategy of the prescribed prover and the constructions of [32, 37] use prover strategies of high complexity. This fact limits the applicability of these proof systems in practice. (Nevertheless, such proof systems may be actually applied when the prover knows something that the verifier does not know, such as an NP-witness to an NP-claim, and when the proof system offers an advantage (over NP-proof systems) such as being zero-knowledge [26, 21].)

1.1 The notion of doubly-efficient interactive proof systems

Seeking to make interactive proof systems available for a wider range of applications, Goldwasser, Kalai and Rothblum put forward a notion of *doubly-efficient* interactive proof systems (also called *interactive proofs for muggles* [25] and *interactive proofs for delegating computation* [36]). In these proof systems the prescribed prover strategy can be implemented in polynomial-time and the verifier's strategy can be implemented in almost-linear-time. That is, doubly-efficient interactive proof systems are restricted by two additional efficiency requirements:

Prover's efficiency requirement: The prescribed prover strategy (referred to in the completeness condition) should be implemented in polynomial-time.

Verifier's efficiency requirement: The verifier strategy should be implemented in almost-linear time.

(We stress that unlike in *argument systems*, the soundness condition holds for all possible cheating strategies (not only for feasible ones).)¹ Restricting the prescribed prover to run in polynomial-time

¹See further discussion in Section 1.5.

implies that such systems may exist only for sets in \mathcal{BPP} , whereas a polynomial-time verifier can check membership in such sets by itself. However, restricting the verifier to run in almost-linear-time implies that something can be gained by interacting with a more powerful prover, even though the latter is restricted to polynomial-time.

The potential applicability of doubly-efficient interactive proof systems was demonstrated by Goldwasser, Kalai and Rothblum [25], who constructed such proof systems for any set that has log-space uniform circuits of small depth (e.g., log-space uniform \mathcal{NC}). A recent work of Reingold, Rothblum, and Rothblum [36] provided doubly-efficient (constant-round) proof systems for any set that can be decided in polynomial-time and small amount of space (e.g., for all sets in \mathcal{SC}). These two results are actually incomparable. The two constructions will be reviewed in Chapters 3 and 4, respectively, but before doing so we shall consider simpler constructions (see Section 1.2 and Chapter 2).

Terminology: We keep the term *almost linear* vague on purpose, but whenever appropriate we shall spell-out a specific interpretation of it. The most strict interpretation is that a function $f : \mathbb{N} \rightarrow \mathbb{N}$ is *almost linear* if $f(n) = \tilde{O}(n) = \text{poly}(\log n) \cdot n$. This interpretation is suitable for Chapter 3 and most of Chapter 2 (i.e., for the results of [25] and [23], resp). In contrast, Chapter 4 (following [36]) uses a much more liberal interpretation by which a sequence of functions of the form $f_\epsilon : \mathbb{N} \rightarrow \mathbb{N}$ (representing the complexities of a sequence of constructions) is *almost linear* if $f_\epsilon(n) \leq n^{1+\epsilon}$ for every $\epsilon > 0$. We mention that these interpretations have been used in the past in other settings (see discussion in [20, Sec. 13.3.3]).

1.2 Doubly-efficient NP-proof systems

For starters we mention that doubly-efficient interactive proof systems exist for a few sets in \mathcal{P} that are not believed to have almost-linear decision procedures. Examples include **perfect matching** and t -**CLIQUE**, for any constant $t \geq 3$. In all these cases, the proof systems are actually of the NP-type. The point is that in each of these cases, an easily verified NP-witness (i.e., one that can be verified in almost-linear time) can be found in polynomial-time.

The foregoing assertion is quite evident for **perfect matching**, by virtue of the known matching algorithms and the fact that checking whether a set of edges is a perfect matching in a given graph can be done in linear time. The same hold for the maximum flow problem, by virtue of the min-cut dual. More generally, morally, the same holds for linear programming (again by duality).² Turning to fixed-parameter problems, such as t -**CLIQUE**, let us now consider the (popular) sets

t -**CLIQUE**: The set of n -vertex graphs that contain a clique of size t .

t -**SUM**: The set of n -long sequences of integers, say, in $[-n^{t+3}, n^{t+3}]$, that contain t elements that sum-up to zero.

t -**OV**: The set of n -long sequences of vectors over $\text{GF}(2)^{\log^2 n}$ that contain t vectors such that their coordinate-wise multiplication yields the all-zero vector.

²Formally, issues may arise with respect to the length of the description of the solutions to the primal and dual programs, but these issues can be resolved by padding the input to that length (which seems quite natural in this context).

In all cases, the NP-witness is a set of t elements, which can be found in time $\binom{n}{t}$ and verified in $\text{poly}(t \cdot \log n)$ -time. Recall that t -CLIQUE is conjectured to require time $n^{c' \cdot t}$, where c' is any constant smaller than one third of the Boolean matrix multiplication exponent (see, e.g., [1]), and that 3-SUM is conjectured to require almost quadratic time (see, e.g., [34], which promoted it), whereas t -OV generalizes the Orthogonal Vectors problem (see [6]).³ Furthermore, t -CLIQUE is $\mathcal{W}[1]$ -complete [15], solving it in time $n^{o(t)}$ refutes the ETH [13], and lower bounds for t -SUM are known to follow from lower bounds for t -CLIQUE (see [2]).

In general, the class of sets having doubly-efficient NP-proof systems is a subclass of $\mathcal{P} \cap \text{Ntime}(\tilde{L})$, where \tilde{L} denotes the set of almost-linear functions. (Indeed, the class of sets having doubly-efficient NP-proof systems consists of all sets $S \in \text{Ntime}(\tilde{L})$ associated with a witness relation R such that given $x \in S$ we can find $y \in R(x) \stackrel{\text{def}}{=} \{w : (x, w) \in R\}$ in polynomial-time.)⁴ As in the case of \mathcal{IP} -versus- \mathcal{NP} , the question is what can be gained by allowing the verifier to be interactive, toss coins, and rule by statistical evidence. That is, moving beyond doubly-efficient NP-proof systems, we focus on the power of doubly-efficient *interactive* proof systems.

1.3 The power of doubly-efficient interactive proof systems

A quite simple doubly-efficient interactive proof system for t -no-CLIQUE (i.e., the complement of t -CLIQUE) is presented in Section 2.1 (following [23, Sec. 3]). This system is based on the observation that membership of an n -vertex graph in t -no-CLIQUE can be captured by an explicit low degree polynomial having n^t terms. Furthermore, each term of this polynomial can be evaluated in almost-linear time (in the size of the graph). Hence, applying the Sum-Check protocol (reviewed in Section 1.6) yields the desired proof system.

The same idea can be applied to a natural class of “locally characterizable” sets defined in Section 2.2 (following [23, Sec. 5]). This class, which contains t -no-CLIQUE, is a subclass of $\mathcal{NC} \cap \mathcal{SC}$. This means that doubly-efficient interactive proof systems for locally characterizable sets can be obtained from either Theorem 1.1 or Theorem 1.2, but the point of presenting the former proof systems is that they are considerably simpler than those obtained by either Theorem 1.1 or Theorem 1.2. Still, the latter theorems yield the most general results known (regarding doubly-efficient interactive proof systems).

Theorem 1.1 (doubly-efficient interactive proof systems for log-space uniform \mathcal{NC} [25]): *Every set that is decidable by a family of log-space uniform circuits of depth d such that $d(n) = n^{o(1)}$, has a doubly-efficient interactive proof system. Furthermore, the proof system uses $O(\log n) \cdot d(n)$ rounds, and the verifier runs in $\tilde{O}(n + d(n))$ -time.*

³Indeed, the Orthogonal Vectors problem corresponds to the case of $t = 2$. Our formulation of t -OV is different but equivalent to the one in [6], where the sequence of vectors is partitioned into t equal parts and a YES-instance has to take a single vector from each part.

⁴Formally, S has an doubly-efficient NP-proof systems if and only if there exists a relation R such that

1. if $(x, y) \in R$, then $|y|$ is almost linear in $|x|$;
2. $S = \{x : \exists w \text{ s.t. } (x, w) \in R\}$;
3. there exists a polynomial-time algorithm that, on input $x \in S$, outputs an element of $\{w : (x, w) \in R\}$;
4. there exists an almost-linear time algorithm that, on input (x, y) , decides whether or not $(x, y) \in R$.

Although circuit size and depth is related to time and space [10], this relation is not tight enough to relate \mathcal{NC} and \mathcal{SC} . Hence, Theorem 1.1 is incomparable to the following

Theorem 1.2 (doubly-efficient interactive proof systems for \mathcal{SC} [36]): *Every set that is decidable by an algorithm that runs in polynomial time and has space complexity s such that $s(n) < \sqrt{n}$, has a doubly-efficient interactive proof system. Furthermore, for any constant $\delta > 0$, the proof system uses $\exp(\tilde{O}(1/\delta))$ rounds, and verifier runs in $(\tilde{O}(n) + \tilde{O}(s(n)^2) \cdot n^\delta)$ -time.*

As noted in [36], Theorem 1.2 can be extended to randomized algorithms by first reducing their randomness complexity to linear (using adequate pseudorandom generators), and then letting the verifier toss coins for the derived algorithm and send the outcomes to the prover (asking it to prove membership in the corresponding residual set). A begging open problem is whether the upper bound of $s(n)^2$ can be replaced by $s(n)$; that is,

Problem 1.3 (a possible quantitative improvement of Theorem 1.2): *Does every set that is decidable by an algorithm that runs in polynomial time and linear space have a doubly-efficient interactive proof system?*

1.4 An upper bound on doubly-efficient interactive proof systems

As stated upfront, doubly-efficient interactive proof systems exists only for sets in \mathcal{BPP} . This is the case, since a decision procedure can just emulate the interaction between the prescribed polynomial-time prover (and the polynomial-time verifier). Using the hypothesis that the verifier runs in almost linear time, it follows that such sets are decidable in almost linear space. This is the case since the hypothesis implies that the communication and the verifier's randomness are (at most) almost-linear (in the length of the input), and the same holds for the space complexity of the verifier. Hence, a machine of almost linear space complexity can decide membership in the set by emulating all possible interactions. Let us state the conclusion for future reference.

Theorem 1.4 (upper bound): *Every set that has a doubly-efficient interactive proof system can be decided in $\mathcal{BPP} \cap \text{Dspace}(\tilde{\ell})$, where $\tilde{\ell}$ is an almost linear function.*

Note that even if Theorem 1.2 is improved as suggested in Problem 1.3, a gap will remain between it and Theorem 1.4, since $\text{TiSp}(T, s)$ is not necessarily equal $\text{Time}(T) \cap \text{Space}(s)$. Hence, another begging open problem is whether the power of doubly-efficient interactive proof systems is captured by $\text{TiSp}(\text{poly}, s)$ or by $\mathcal{BPP} \cap \text{Space}(s)$ (or by neither).

Problem 1.5 (on the gap between Theorems 1.2 and 1.4): *Prove or provide evidence against at least one of the following conjectures:*

1. *For $s(n) = \sqrt{n}$, every set in $\mathcal{P} \cap \text{Dspace}(s)$ has a doubly-efficient interactive proof system.*
2. *Every set that has a doubly-efficient interactive proof system can be decided by a probabilistic algorithm that runs in polynomial time and almost linear space.*

1.5 On doubly-efficient argument systems

Recall that argument systems are defined as interactive proof systems with the exception that the soundness condition is replaced by a computational soundness condition. That is, while the (standard) soundness condition requires that no cheating strategy can make the prover accept false assertions, except with negligible probability, the computational soundness condition only requires the infeasibility of cheating (i.e., that cheating strategies that can be implemented by polynomial size circuits may only fool the verifier with negligible probability).⁵

Doubly-efficient argument systems for any set in \mathcal{BPP} are implicit in Kilian’s argument system for sets in \mathcal{NP} , which relies on collision resistance hashing [30].⁶ This system uses a constant number of rounds. Furthermore, assuming the existence of computational PIR schemes [31], every set in \mathcal{P} has a two-message doubly-efficient argument system [28, 29, 11].⁷ This result also extends to \mathcal{BPP} , whereas sets having a doubly-efficient argument system are in \mathcal{BPP} (since we can decide membership in such a set by emulating the interaction between the prescribed prover and verifier strategies).⁸

The fact that argument systems are always asserted by relying on an intractability assumption is no coincidence, since these asserted systems do not satisfy the information theoretic soundness requirement. In fact, the existence of an argument system that is not an interactive proof system (i.e., does not satisfy standard soundness) implies a complexity separation (which is not known unconditionally). Specifically:

Theorem 1.6 (arguments that are not proofs imply separations): *Let V be a verifier strategy for an argument system for a set S , and suppose that V does not satisfy the (information theoretic) soundness requirement. Then, \mathcal{PSPACE} is not contained in \mathcal{P}/poly . In particular, $\mathcal{BPP} \neq \mathcal{PSPACE}$.*

In other words, Theorem 1.6 asserts that a gap between information theoretic soundness and computational soundness means a gap between computationally unbounded strategies and computationally bounded strategies. Recalling that, in the current setting (of fooling a probabilistic polynomial-time verifier), the former can be implemented in \mathcal{PSPACE} , whereas the computational restriction refers to \mathcal{P}/poly , the main claim follows (and $\mathcal{BPP} \neq \mathcal{PSPACE}$ follows, since $\mathcal{BPP} \subset \mathcal{P}/\text{poly}$ and $\mathcal{BPP} \subseteq \mathcal{PSPACE}$).

Proof: Assume, for simplicity and without loss of generality, that in the said argument system each message of the prover consists of a single bit, and let $f : \{0, 1\}^* \times \{0, 1\}^* \rightarrow \{0, 1\}$ denote an optimal prover strategy with respect to this system (i.e., $f(x, \gamma)$ is the message sent by the prover on common input x after receiving the sequence of verifier messages described in the communication transcript γ). Then, $f \in \mathcal{PSPACE}$, since an optimal prover strategy (w.r.t a probabilistic polynomial-time verifier) can be implemented in polynomial space. On the other hand, $f \notin \mathcal{P}/\text{poly}$, since otherwise a polynomial size circuit could implement the optimal strategy (for convincing V), and so there

⁵Polynomial size circuits are preferred over probabilistic polynomial-time algorithms in order to account for auxiliary information that may be available to the prover (esp., when used as a subroutine inside a higher-level application).

⁶The hypothesis also implies the existence of one-way functions, and hence of pseudorandom generators [27]. Using such a generator, we can reduce the randomness complexity of the decision procedures for \mathcal{BPP} to linear, and let the verifier send a random-tape for such a procedure (as part of its first message).

⁷The result of [29], which builds on [28], uses a computational PIR of quasi-polynomial security. The assumption was weakened to standard (polynomial security) by [11].

⁸The validity of this decision procedure refers only to the probability that the prescribed prover convinces the (prescribed) verifier.

will be no gap between the information theoretic soundness and the computational soundness of this proof system. ■

1.6 Preliminaries: The sum-check protocol

The sum-check protocol, designed by Lund, Fortnow, Karloff, and Nisan [32], is a key ingredient in some of the constructions that we present. In particular, it will be used in Chapters 2 and 3.

Fixing a finite field \mathcal{F} and a set $H \subset \mathcal{F}$ (e.g., H may consist of the 0 and 1 elements of \mathcal{F}), we consider an m -variate polynomial $P : \mathcal{F}^m \rightarrow \mathcal{F}$ of individual degree d . Given a value v , the sum-check protocol is used to prove that

$$\sum_{\sigma_1, \dots, \sigma_m \in H} P(\sigma_1, \dots, \sigma_m) = v. \quad (1.1)$$

The sum-check protocol proceeds in m iterations, starting with $v_0 = v$, such that in the i^{th} iteration the parties act as follows.

Prover's move: The prover computes a univariate polynomial of degree d

$$P_i(z) \stackrel{\text{def}}{=} \sum_{\sigma_{i+1}, \dots, \sigma_m \in H} P(r_1, \dots, r_{i-1}, z, \sigma_{i+1}, \dots, \sigma_m) \quad (1.2)$$

where r_1, \dots, r_{i-1} are as determined in prior iterations, and sends P_i to the verifier.

Verifier's move: The verifier checks that $\sum_{\sigma \in H} P_i(\sigma) = v_{i-1}$ and rejects if inequality holds.

Otherwise, it selects r_i uniformly in \mathcal{F} , and sends it to the prover, while setting $v_i \leftarrow P_i(r_i)$.

If all iterations are completed, the verifier conducts a final check. It computes the value of $P(r_1, \dots, r_m)$ and accepts if and only if this value equals v_m .

Clearly, if Eq. (1.1) holds (and the prover acts according to the protocol), then the verifier accepts with probability 1. Otherwise, no matter what the prover does, the verifier accepts with probability at most $m \cdot d/|\mathcal{F}|$. The complexity of verification is dominated by the complexity of evaluating P (on a single point). As for the prescribed prover, it may compute the relevant P_i 's by interpolation, which is based on computing the value of P at $(d+1) \cdot 2^{m-i}$ points, for $i \in [m]$.

1.7 Organization

In Chapter 2 we present simple constructions of doubly-efficient interactive proof systems, first for t -no-CLIQUE, and next for a natural class that contains it and is contained in uniform \mathcal{NC} (and also in \mathcal{SC}). These proof systems are due to Goldreich and Rothblum [23]. In Chapter 3 we present the proof systems of Goldwasser, Kalai and Rothblum [25], which are applicable to sets that are recognized by small depth circuits (e.g., uniform \mathcal{NC}). In Chapter 4 we provide an outline of the proof systems of Reingold, Rothblum, and Rothblum [36], which are applicable to sets recognized in polynomial-time and small space (e.g., sets in \mathcal{SC}).

Conventions: We assume that the verifier (resp., prover) has *direct access* to the common input; that is, each bit in the input can be read in unit cost. Unless explicitly stated differently, all logarithms are to base 2.

Chapter 2

Simple doubly-efficient interactive proof systems

In this chapter, we present simple constructions of doubly-efficient interactive proof systems for sets in \mathcal{P} that are believed to have relatively high complexity. Specifically, such a construction is presented for t -no-CLIQUE (i.e., the set of n -vertex graphs that do not contain a clique of size t). Next, we present a generic construction of such proof systems for a natural class that contains t -no-CLIQUE and is in \mathcal{NC} (and also in \mathcal{SC}). The proof systems presented here are significantly simpler than the proof systems presented by Goldwasser, Kalai and Rothblum [25], let alone those presented by Reingold, Rothblum, and Rothblum [36].

Section 2.1 presents the simplest proof system, which is for t -no-CLIQUE. The generic construction for any *locally-characterizable set* is presented in Section 2.2: The corresponding class is defined in Section 2.2.1, the basic construction appear in Section 2.2.2, and ramifications appear in Sections 2.2.3 and 2.2.4.

Note: Although the generic construction for the class of locally-characterizable sets, which is presented in Section 2.2.2, almost meets the parameters of construction tailored for t -no-CLIQUE, we present the latter construction first (in Sections 2.1). This is done since the construction tailored for t -no-CLIQUE is simpler and actually more efficient. We also mention that the construction for locally-characterizable sets (captured by Theorem 2.2) can be extended to a wider class of sets, as detailed in Section 2.2.4.

2.1 The case of t -CLIQUE

For a parameter $t \in \mathbb{N}$, given a graph $G = ([n], E)$, the problem we consider is determining whether there exist t vertices such that the subgraph induced by them is a clique; that is, does there exist distinct $v_1, \dots, v_t \in [n]$ such that $\bigwedge_{j,k \in [t]: j < k} \chi_{v_j, v_k}$, where $\chi_{u,v} = 1$ if and only if $\{u, v\} \in E$. (Our focus is on simple graphs, and so we assume that $\chi_{v,v} = 0$ for every $v \in [n]$.)

We actually consider the set of YES-instances (i.e., graphs having a t -clique), denoted t -CLIQUE, and the set of NO-instances, denoted t -no-CLIQUE. As outlined in Section 1.2, the set t -CLIQUE has an NP-proof system that uses proofs of length $t \log n$. We shall present a doubly-efficient interactive proof for the set t -no-CLIQUE.

Fixing a generic graph $G = ([n], E)$ (and recalling that $\chi_{u,v} = 1$ if and only if $\{u, v\} \in E$), we consider the expression

$$\sum_{v^{(1)}, \dots, v^{(t)} \in [n]} \prod_{j, k \in [t]: j < k} \chi_{v^{(j)}, v^{(k)}}. \quad (2.1)$$

Observe that Eq. (2.1) is zero if and only if G is in t -no-CLIQUE. Hence, proving that G is in t -no-CLIQUE reduces to proving that the sum of n^t terms evaluates to zero, and the sum-check protocol seems adequate for the task, except that the expression is over the integers (rather than over a finite field) and a generic term (in Eq. (2.1)) is not expressed as a polynomial in $v^{(1)}, \dots, v^{(t)}$.

The first problem is easy to handle: we can just use a finite field \mathcal{F} of prime size greater than n^t (and consider Eq. (2.1) as an expression in \mathcal{F}). The second problem is more acute: we wish to represent a choice from the sequence of values $(\chi_{u,w})_{u,w \in [n]}$, which may be viewed as a function $\chi : [n] \times [n] \rightarrow \{0, 1\}$, as a (low degree) polynomial (in the description of this choice). The way to do it is to identify $[n]$ with $\{0, 1\}^\ell$, where $\ell = \log_2 n$, and to consider a low-degree extension of the function $\chi : \{0, 1\}^\ell \times \{0, 1\}^\ell \rightarrow \{0, 1\}$. Specifically, we use the polynomial $\hat{\chi} : \mathcal{F}^\ell \times \mathcal{F}^\ell \rightarrow \mathcal{F}$ defined by

$$\hat{\chi}(y, z) = \sum_{\alpha, \beta \in \{0, 1\}^\ell} \prod_{i \in [\ell]} (y_i \alpha_i + (1 - y_i)(1 - \alpha_i)) \cdot \prod_{i \in [\ell]} (z_i \beta_i + (1 - z_i)(1 - \beta_i)) \cdot \chi_{\alpha, \beta}.$$

(There is some abuse of notation here: In the first two occurrences, α and β are viewed as elements of \mathcal{F}^ℓ , which happen to reside in $\{0, 1\}^\ell \subset \mathcal{F}^\ell$, whereas in the last occurrence α and β are viewed as elements of $[n] \equiv \{0, 1\}^\ell$.)

Hence, fixing a generic graph $G = ([n], E)$ and letting $\ell = \log n$, we consider a finite field \mathcal{F} of prime size greater than n^t , and identify $\{0, 1\}$ with the set H containing the zero and one elements of \mathcal{F} . Using this identification (and recalling that $\chi_{u,v} = 1$ if and only if $\{u, v\} \in E$), we define a polynomial $P : (\mathcal{F}^\ell)^t \rightarrow \mathcal{F}$ such that

$$P(z^{(1)}, \dots, z^{(t)}) = \prod_{j, k \in [t]: j < k} \sum_{\alpha, \beta \in H^\ell} \text{EQ}(\alpha\beta, z^{(j)} z^{(k)}) \cdot \chi_{\alpha, \beta} \quad (2.2)$$

$$\text{where } \text{EQ}(\bar{\gamma}, \bar{z}) = \prod_{i \in [2\ell]} (z_i \gamma_i + (1 - z_i)(1 - \gamma_i)). \quad (2.3)$$

Note that each $\bar{z} \in (H^\ell)^t$ it holds that $P(\bar{z}) \in \{0, 1\}$, and that $P(\bar{z}) = 1$ if and only if \bar{z} describes a sequence of t distinct vertices that induces a clique in $G = ([n], E)$. (This holds since, for distinct $v^{(1)}, \dots, v^{(t)} \in H^\ell$, it holds that $P(v^{(1)}, \dots, v^{(t)}) = 1$ if and only if $\{v^{(j)}, v^{(k)}\} \in E$ for every $j < k$.)¹

Note that P has individual degree $O(t)$, and that it is straightforward to evaluate P in time $O(t^2 \cdot 2^{2\ell} \cdot \ell) = \tilde{O}(t^2 \cdot n^2)$. Also, for $\bar{\gamma}, \bar{z} \in H^{2\ell}$, it holds that $\text{EQ}(\bar{\gamma}, \bar{z}) = 1$ if $\bar{\gamma} = \bar{z}$ and $\text{EQ}(\bar{\gamma}, \bar{z}) = 0$ otherwise. Hence, for $\bar{z} \in H^{t\ell}$, it holds that $P(\bar{z}) = \prod_{j, k \in [t]: j < k} \chi_{z^{(j)}, z^{(k)}}$.

The key observation is that the graph G is a NO-instance if and only if for all $\bar{z} \in (H^\ell)^t$ it holds that $P(\bar{z}) = 0$. Hence, we obtain an interactive proof system (for the set of NO-instances) by applying the sum-check protocol to the claim $\sum_{\bar{z} \in (H^\ell)^t} P(\bar{z}) = 0$.

The complexity of the verifier's strategy is dominated by the evaluation of P (as defined in Eq. (2.2)), which reduces to $\binom{t}{2}$ computations of sums having the form

$$\sum_{\alpha, \beta \in H^\ell} \text{EQ}(\alpha\beta, r^{(j)} r^{(k)}) \cdot \chi_{\alpha, \beta} \quad (2.4)$$

¹Also, if $v^{(j)} = v^{(k)}$ for some $j < k$, then $P(v^{(1)}, \dots, v^{(t)}) = 0$, since $\chi_{v^{(j)}, v^{(k)}} = 0$.

where $j < k \in [t]$ and $r^{(1)} \dots r^{(t)} \in (\mathcal{F}^\ell)^t$ are determined in the execution of the sum-check protocol. Writing Eq. (2.4) as $\sum_{(\alpha, \beta) \in H^{2\ell} \cap E} \mathbf{EQ}(\alpha\beta, r^{(j)}r^{(k)})$, we can evaluate this sum in time $O(|E| \cdot \ell) = \tilde{O}(|E| + n)$.

Turning our attention to the complexity of proving, we observe that the prover has to compute the polynomials that arise in each of the iterations of the sum-check protocol. The prover can do so by computing partial sums with at most $|H|^{t\ell} = n^t$ terms, where computing each such term amounts to $\text{poly}(t)$ evaluations of P . Hence, the prover's complexity is definitely bounded by $\tilde{O}(n^t \cdot |E|)$. A closer inspection reveals that we can do better, but before taking this closer look let us summarize what we have shown so far.

Proposition 2.1 (doubly-efficient interactive proof systems for t -no-CLIQUE): *For every constant t , the set t -no-CLIQUE has a doubly-efficient interactive proof system. Furthermore, the proof system uses logarithmically many rounds, and (on input G) the verifier runs in $\tilde{O}(|G|)$ -time.*

Indeed, the point of Proposition 2.1 is that is the simplicity of its proof; that is, the fact that it is proved by merely applying the sum-check protocol to the claim $\sum_{\bar{z} \in (H^\ell)^t} P(\bar{z}) = 0$, where P is defined in Eq. (2.2).

A finer analysis of the prover's running time. Recall that we already know that the prover's complexity is bounded by $\tilde{O}(n^t \cdot |E|)$. Our goal is to obtain an upper bound of $\tilde{O}(n^t)$. Towards this end, we note that in the i^{th} iteration of the sum-check, the prover has to provide the univariate polynomial (in z)

$$\sum_{\bar{s} \in H^{t\ell-i}} P(\bar{r}, z, \bar{s}) \quad (2.5)$$

where $\bar{r} \in \mathcal{F}^{i-1}$ was determined in the previous iteration. This univariate polynomial can be found by interpolation, and so the complexity of finding it is $\text{poly}(t) \cdot 2^{t\ell-i} \cdot O(|E| \cdot \ell)$, which is $\tilde{O}(n^t)$ for $i \geq 2\ell$. Thus, we focus on the case of $i \in [2\ell - 1]$. In that case, the complexity is $n^t/2^i$ times the complexity of evaluating P on $\bar{r}u\bar{s}$, where $(\bar{r}, u) \in \mathcal{F}^{i-1} \times \mathcal{F}$ and $\bar{s} \in H^{t\ell-i}$ (for $O(t^2)$ values of $u \in \mathcal{F}$). So the question is what is the complexity of evaluating P on such a $t\ell$ -element argument (which has a $(t\ell - i)$ -long suffix in $H^{t\ell-i}$).

Focusing on evaluating each of the inner sums (captured by Eq. (2.4)), we consider evaluating the sum $\sum_{\alpha, \beta \in H^\ell} \mathbf{EQ}(\alpha\beta, v) \cdot \chi_{\alpha, \beta}$, when given $v = uw \in \mathcal{F}^{2\ell}$ such that $w \in H^{2\ell-p}$, where p is determined by i and j, k (indeed, $p = i \in [2\ell]$ if $(j, k) = (1, 2)$, but $p < i$ otherwise). Letting $\chi_\gamma = \chi_{\alpha, \beta}$, where $\gamma = \alpha\beta$ such that $|\alpha| = |\gamma|/2$, we have

$$\begin{aligned} \sum_{\gamma \in H^{2\ell}} \mathbf{EQ}(\gamma, uw) \cdot \chi_\gamma &= \sum_{(\gamma', \gamma'') \in H^p \times H^{2\ell-p}} \mathbf{EQ}(\gamma', u) \cdot \mathbf{EQ}(\gamma'', w) \cdot \chi_{\gamma'\gamma''} \\ &= \sum_{\gamma' \in H^p} \mathbf{EQ}(\gamma', u) \cdot \chi_{\gamma'w} \end{aligned}$$

where the second equality holds because for $\gamma'', w \in H^{2\ell-p}$ it holds that $\mathbf{EQ}(\gamma'', w) = 1$ if $\gamma'' = w$ and $\mathbf{EQ}(\gamma'', w) = 0$ otherwise. Hence, evaluating this sum takes time $O(2^p \cdot \ell)$. The final observation is that, in our application, it holds that $p \leq i$, since the p values in $\mathcal{F} \setminus H$ can only arise from the values determined in the previous $i - 1$ iterations and the single value used for interpolation in the current iteration. It follows that the complexity of implementing the prover's strategy in the i^{th} iteration is $\text{poly}(t) \cdot 2^{t\ell-i} \cdot O(2^p \cdot \ell) = \tilde{O}(n^t)$.

Credits. The foregoing construction is reproduced from the work of Goldreich and Rothblum [23]. This proof system is comparable to system presented by Thaler [39]. In addition, several recent works [12, 42, 9] constructed *non-interactive* proof systems for t -no-CLIQUE and related problems. For a detailed discussion, the interested reader is referred to [23, Sec. 1.1]. We also mention that, as discussed in [23, Sec. 1.1], the foregoing proof system is not only simpler than those of [25, 36] (which will be reviewed in Chapters 3 and 4, respectively), but also yields better round complexity.

2.2 The general result

In this section we present simple constructions of proof systems for a natural sub-class of $\mathcal{NC} \cap \mathcal{SC}$, which is believed not to be in $\text{Dtime}(p)$ for any fixed polynomial p . Loosely speaking, this class consists of all sets that can be locally-characterized by the conjunction of polynomially many local conditions that can be expressed by Boolean formulae of polylogarithmic size (see definition in Section 2.2.1). The following result will become precise once the definition of locally-characterizable sets (i.e., Definition 2.3) is in place.

Theorem 2.2 (doubly-efficient interactive proof systems for locally-characterizable sets): *Every locally-characterizable set has a simple doubly-efficient interactive proof system. Specifically, on input of length n , the verifier runs in $\tilde{O}(n)$ -time and the strategy of the prescribed prover can be implemented in $\tilde{O}(n^{c+1})$ -time, where n^c denotes the number of local conditions in the characterization. Furthermore, the interactive proof system uses public coins, has logarithmic round-complexity, and uses polylogarithmic communication.*

We mention that that using different settings of parameters yields a general trade-off between the computational complexity and the number of communication rounds (see Theorem 2.6). The class of “locally-characterizable” sets contains two t -parameterized problems that are widely believed not to be in, say, $\text{Dtime}(n^{t/5})$:

t -no-CLIQUE: The set of n -vertex graphs that do not contain a clique of size t .

t -no-SUM: The set of n -long sequences of integers that contain no t elements that sum-up to zero.²

(Indeed, the aforementioned sets are the complements of the NP-sets called t -CLIQUE and t -SUM.) Recall that a direct construction of a proof system for t -no-CLIQUE was presented in Section 2.1, with a corresponding prover strategy that can be implemented in time $\tilde{O}(n^t)$. An analogous (but more cumbersome) construction for t -no-SUM appears in [23, Sec. 4].

2.2.1 A natural class: locally-characterizable sets

The following definition is related but different from the definition of “local characterization” that is often used in the property testing literature (see, Sudan’s survey [38]). Most importantly, the latter definitions do not specify the complexity of the functions ϕ_n and $\pi_{n,j}$, and typically take p to be a constant.³

²Alternatively, we may consider sequences of positive integers and ask if they contain a t -subset that sums-up to target integer, which is also given as part of the input.

³In addition, the notion used in property testing does not restrict the domain of Φ_x to have size $\text{poly}(|x|)$, although this can be assumed without loss of generality.

Definition 2.3 (locally-characterizable sets): *A set S is locally-characterizable if there exist a constant c , a polynomial p and a polynomial-time algorithm that on input n outputs $\text{poly}(\log n)$ -sized formulae $\phi_n : [n]^{p(\log n)} \times \{0, 1\}^{p(\log n)} \rightarrow \{0, 1\}$ and $\pi_{n,1}, \dots, \pi_{n,p(\log n)} : \{0, 1\}^{c \log n} \rightarrow [n]$ such that, for every $x \in \{0, 1\}^n$, it holds that $x \in S$ if and only if for all $w \in \{0, 1\}^{c \log n}$*

$$\Phi_x(w) \stackrel{\text{def}}{=} \phi_n(\pi_{n,1}(w), \dots, \pi_{n,p(\log n)}(w), x_{\pi_{n,1}(w)}, \dots, x_{\pi_{n,p(\log n)}(w)}) \quad (2.6)$$

equals 0.⁴

That is, each value of $w \in \{0, 1\}^{c \log n}$ yields a local condition that refers to polylogarithmically many locations in the input (i.e., the locations $\pi_{n,1}(w), \dots, \pi_{n,p(\log n)}(w) \in [n]$). This local condition is captured by ϕ_n , and in its general form it depends both on the selected locations and on the value of the input in these locations. A simplified form, which suffices in many case, uses a local condition that only depends on the values of the input in these locations (i.e., $\phi_n : [n]^{p(\log n)} \times \{0, 1\}^{p(\log n)} \rightarrow \{0, 1\}$ only depends on the $p(\log n)$ -bit suffix).

The simplified form (in which $\phi_n : \{0, 1\}^{p(\log n)} \rightarrow \{0, 1\}$) suffices for capturing the specific problems mentioned above. Specifically, for fixed $t \in \mathbb{N}$, when representing n -vertex graphs by their adjacency matrix, denoted $x = (x_{r,c})_{r,c \in [n]}$, the t -CLIQUE problem is captured by $\Phi_x(i_1, \dots, i_t) = \bigwedge_{j < k} x_{i_j, i_k}$; that is, we use $\phi_{n^2} : \{0, 1\}^{t^2} \rightarrow \{0, 1\}$ and $\pi_{n^2, (j,k)} : \{0, 1\}^{t \log n} \rightarrow [n] \times [n]$ (for $j, k \in [t]$) such that $\phi_{n^2}(\sigma_{1,1}, \dots, \sigma_{t,t}) = \bigwedge_{j < k} \sigma_{j,k}$ and $\pi_{n^2, (j,k)}(i_1, \dots, i_t) = (i_j, i_k)$. Likewise, considering n -long sequences over $[-m, m]$ (where $m = \text{poly}(n)$), with some abuse of notation, the t -SUM problem is captured by $\Phi_x(i_1, \dots, i_t) = 1$ if and only if $\sum_{j \in [t]} x_{i_j} = 0$; that is, we use $\phi_n : [m]^t \rightarrow \{0, 1\}$ such that $\phi_n(z_1, \dots, z_t) = \text{TruthValue}(\sum_{j \in [t]} z_j = 0)$ and $\pi_{n,j} : \{0, 1\}^{t \log n} \rightarrow [n]$ such that $\pi_{n,j}(i_1, \dots, i_t) = i_j$ (for each $j \in [t]$).

Note that the complement of every locally-characterizable set has a doubly-efficient interactive proof system. In this proof system, on input $x \in \{0, 1\}^n$, letting $\ell' = c\ell = c \log n$, the prover finds an adequate $w \in \{0, 1\}^{\ell'}$, sends it to the verifier, who retrieves the bits $x_{\pi_{n,1}(w)}, \dots, x_{\pi_{n,p(\ell)}(w)}$ and evaluates ϕ_n on them. Indeed, in this NP-proof system, the prover runs in time $2^{\ell'} \cdot \tilde{O}(|x|) = \text{poly}(|x|)$, whereas the verifier runs in $\text{poly}(\log |x|)$ -time (given direct access to the input). On the other hand, the set itself (i.e., every locally-characterizable set) has a doubly-efficient interactive proof systems (since computing the function $\sum_{w \in \{0,1\}^{\ell'}} \Phi_x(w)$, where Φ_x is as in Eq. (2.6), is in \mathcal{SC} , and so the proof system of [36] can be used).⁵ Here we present a much simpler construction.

2.2.2 Proof of Theorem 2.2

Letting $\ell = \log n$, we associated $[n]$ with $\{0, 1\}^\ell$, and derive from each $\pi_{n,j} : \{0, 1\}^{c\ell} \rightarrow [n]$ Boolean formulae $\pi_{n,j,1}, \dots, \pi_{n,j,\ell} : \{0, 1\}^{c\ell} \rightarrow \{0, 1\}$ such that $\pi_{n,j,k}(w)$ is the k^{th} bit of $\pi_{n,j}(w)$. We may assume, without loss of generality, that the depth of each of the formulae (i.e., ϕ_n and the $\pi_{n,j,k}$'s) is logarithmic in their size (which is $\text{poly}(\ell)$).⁶

⁴To simplify our exposition, we require that in case of inputs in S , the predicate ϕ_n evaluates to 0 (rather than to 1).

⁵Alternatively, if the algorithm guaranteed by Definition 2.3 uses linear space (i.e., $O(\log n)$ space), then computing $\sum_{w \in \{0,1\}^{\ell'}} \Phi_x(w)$ is in log-uniform \mathcal{NC} , and one can use the proof system of [25].

⁶Note that the transformation to this form can be performed in polynomial time, whereas the relevant formulae are of $\text{poly}(\log n)$ -size.

Next, for a finite field \mathcal{F} of size $\text{poly}(n)$, we construct arithmetic formulae $\widehat{\phi}_n : \mathcal{F}^{(\ell+1) \cdot p(\ell)} \rightarrow \mathcal{F}$ and $\widehat{\pi}_{n,j,k} : \mathcal{F}^{c\ell} \rightarrow \mathcal{F}$ such that $\widehat{\phi}_n$ (resp., $\widehat{\pi}_{n,j,k}$) agrees with ϕ_n on $H^{\ell \cdot p(\ell) + p(\ell)}$ (resp., with $\pi_{n,j,k}$ on $H^{c\ell}$). The crucial point is that these arithmetic formulae preserve the depth of the Boolean counterparts, and so the degrees of the functions that they compute is upper bounded by $D = \text{poly}(\log n) \ll |\mathcal{F}|$. (Note: \mathcal{F} is chosen so that the latter inequality holds.) Now, letting $\widehat{\pi}_{n,j} : \mathcal{F}^{c\ell} \rightarrow \mathcal{F}^\ell$ such that $\widehat{\pi}_{n,j}(\bar{z}) = (\widehat{\pi}_{n,j,1}(\bar{z}), \dots, \widehat{\pi}_{n,j,\ell}(\bar{z}))$, we consider the function $\widehat{\Phi}_x : \mathcal{F}^{c\ell} \rightarrow \mathcal{F}$ (i.e., an extension of Φ_x) such that

$$\widehat{\Phi}_x(\bar{z}) \stackrel{\text{def}}{=} \widehat{\phi}_n(\widehat{\pi}_{n,1}(\bar{z}), \dots, \widehat{\pi}_{n,p(\ell)}(\bar{z}), X_1, \dots, X_{p(\ell)}) \quad (2.7)$$

$$\text{where } X_j = \sum_{\alpha \in \{0,1\}^\ell} \text{EQ}(\widehat{\pi}_{n,j}(\bar{z}), \alpha) \cdot x_\alpha \quad (2.8)$$

and, as in Eq. (2.3), $\text{EQ}(\bar{y}, \alpha) = \prod_{i \in [\ell]} (y_i \alpha_i + (1 - y_i)(1 - \alpha_i))$. That is, the value of $\widehat{\Phi}_x(\bar{z})$ is obtained by feeding to $\widehat{\phi}_n : \mathcal{F}^{(\ell+1) \cdot p(\ell)} \rightarrow \mathcal{F}$ the $(p(\ell) \cdot \ell + p(\ell))$ -long sequence consisting of the $p(\ell) \cdot \ell$ -long sequence $(\widehat{\pi}_{n,1}(\bar{z}), \dots, \widehat{\pi}_{n,p(\ell)}(\bar{z})) \in (\mathcal{F}^\ell)^{p(\ell)}$ and the $p(\ell)$ -long sequence whose j^{th} location contains the field element $\sum_{\alpha \in \{0,1\}^\ell} \text{EQ}(\widehat{\pi}_{n,j}(\bar{z}), \alpha) \cdot x_\alpha$.

We invoke the sum-check protocol on the claim that the sum $\sum_{w \in \{0,1\}^{c\ell}} \widehat{\Phi}_x(w)$ equals 0, relying on the fact that \mathcal{F} is larger than $2^{c\ell}$. This protocol performs $c\ell$ iterations (i.e., it is applied only to the outer sum), and the verifier evaluates the residual expression, which amounts to evaluating the $\widehat{\pi}_{n,j}$'s and $\widehat{\phi}_n$ as well as computing $\text{poly}(\ell)$ sums of 2^ℓ terms each. The prover's computation is dominated by computing a sum of $2^{c\ell}$ terms, where each term requires a computation of the type conducted by the verifier. Recalling that $2^\ell = n$, it follows that the verifier runs in almost-linear-time (i.e., it runs in $\widetilde{O}(n)$ -time), whereas the prover runs in polynomial-time (i.e., it runs in $\widetilde{O}(n^{c+1})$ -time). This completes the proof of Theorem 2.2.

Comment: Applied to the t -no-SUM problem, the foregoing generic construction yields a proof system of complexity that is comparable to that of the tailored proof system presented in [23, Sec. 4]. In case of t -no-CLIQUE, the generic construction yields a verifier that runs in time that is almost linear in the size of the adjacency matrix, whereas the tailored proof system (presented in Section 2.1) yields a verifier that runs in time that is almost linear in the number of edges. Furthermore, the prover's complexity in this generic construction is n times slower than that of the tailored proof system.

Remark 2.4 (a relaxation of Definition 2.3): *Theorem 2.2 holds also when relaxing the notion of locally-characterizable sets such that the $\text{poly}(\log n)$ -sized formulae are required to be generated in $\widetilde{O}(n)$ -time, rather than in $\text{poly}(\log n)$ -time. Actually, the foregoing proof remains intact even if the said formulae may depend on x itself, but we consider the class as defined in Definition 2.3 more natural.*

Remark 2.5 (counting the number of violated condition): *The foregoing interactive proof system is applicable also to the task of verifying the number of violated conditions. The same applies also to the problem-tailored proof system presented in Section 2.1. Note that the number of violated conditions in t -no-CLIQUE is the number of t -cliques in the graph.*

2.2.3 Generalization: round versus computation trade-off

By using a set H of arbitrary size (rather than $H \equiv \{0, 1\}$), we obtain a general trade-off between the computational complexity and the number of communication rounds. Specifically, the computational complexity increases by a factor of $\tilde{O}(|H|)$, whereas the number of rounds is decreased by a factor of $\log |H|$.

Theorem 2.6 (Theorem 2.2, generalized): *Every locally-characterizable set has a simple doubly-efficient interactive proof system. Specifically, on input of length n and using a parameter $h \leq n$, we get public-coin interactive proof systems of round-complexity $O(\log_h n)$, verification time $\tilde{O}(h \cdot n)$, and proving time $\tilde{O}(h \cdot n^{c+1})$.*

In particular, setting $h = n^\epsilon$ for any constant $\epsilon > 0$, yields a constant-round interactive proof of verification time $\tilde{O}(n^{1+\epsilon})$. On the other hand, using $h = \log n$ maintains the computational complexity bounds of Theorem 2.2 (i.e., $\tilde{O}(n)$ -time verification and $\tilde{O}(n^{c+1})$ -time prover strategy) while using $o(\log n)$ rounds of communication.

Proof: We use $H = [h]$ but maintain $\ell = \log n$. Defining the $\pi_{n,j,k}$'s as in the proof of Theorem 2.2 and letting $d = \log h$, we associate $\{0, 1\}^{c\ell}$ with $H^{c\ell/d}$ and consider $\pi_{n,j,k} : H^{c\ell/d} \rightarrow \{0, 1\}$. Now, we let $\hat{\pi}_{n,j,k} : \mathcal{F}^{c\ell/d} \rightarrow \mathcal{F}$ be the low degree extension of $\pi_{n,j,k}$, and define $\hat{\pi}_{n,j} : \mathcal{F}^{c\ell/d} \rightarrow \mathcal{F}^\ell$ as before. The definition of $\hat{\Phi}_x : \mathcal{F}^{c\ell/d} \rightarrow \mathcal{F}$ is analogous to Eq. (2.7) except that Eq. (2.8) is replaced by

$$X_i = \sum_{\alpha \in H^{\ell/d}} \left(\prod_{i \in [\ell/d]} \prod_{\beta \in H \setminus \{\alpha_i\}} \frac{\beta - z_i}{\beta - \alpha_i} \right) \cdot x_\alpha \quad (2.9)$$

Invoking the sum-check protocol (w.r.t the non-binary H) on the claim that the sum $\sum_{w \in H^{c\ell/d}} \hat{\Phi}_x(w)$ equals 0, yields a protocol that performs $c\ell/d$ iterations. Again, the verifier evaluates the residual expression, which amounts to evaluating the $\hat{\pi}_{n,j}$'s and $\hat{\phi}_n$ as well as computing $\text{poly}(\ell)$ sums of $h^{\ell/d} = 2^\ell$ terms each, but here each term calls for evaluating a polynomial that has an arithmetic formula of size $O(\ell \cdot h/d)$. The prover's computation is dominated by computing a sum of $h^{c\ell/d} = 2^{c\ell}$ terms, where each term requires a computation of the type conducted by the verifier. ■

2.2.4 Extension to a wider class

As a motivation towards the following extension, we mention the problem of finding a dominating set of constant size t . This problem does not seem to be locally characterizable in the sense of Definition 2.3 (cf. [35]), but it definitely reside in the following class.

Definition 2.7 (locally $\forall\exists$ -characterizable sets): *A set S is locally $\forall\exists$ -characterizable if there exist constants c, c' , a polynomial p and an almost-linear time algorithm that on input 1^n outputs $\text{poly}(\log n)$ -sized formulae $\phi_n : [n]^{p(\log n)} \times \{0, 1\}^{p(\log n)} \rightarrow \{0, 1\}$ and $\pi_{n,1}, \dots, \pi_{n,p(\log n)} : \{0, 1\}^{(c+c') \log n} \rightarrow [n]$ such that, for every $x \in \{0, 1\}^n$, it holds that $x \in S$ if and only if for all $w \in \{0, 1\}^{c \log n}$ there exists $w' \in \{0, 1\}^{c' \log n}$ such that*

$$\Phi_x(w, w') \stackrel{\text{def}}{=} \phi_n(\pi_{n,1}(w, w'), \dots, \pi_{n,p(\log n)}(w, w'), x_{\pi_{n,1}(w, w')}, \dots, x_{\pi_{n,p(\log n)}(w, w')}) \quad (2.10)$$

equals 0.

Like in Definition 2.3, sometimes one may use a simplified form in which ϕ_n only depends on its $p(\log n)$ -bit long suffix. This simplification suffices for characterizing the set of graphs not having a dominating set of size $t = O(1)$. Specifically, when representing n -vertex graphs by their adjacency matrix x (augmented with 1's on the diagonal), the t -dominating set problem is captured by $\Phi_x(i_1, \dots, i_t, i_{t+1}) = \bigvee_{j \in [t]} x_{i_j, i_{t+1}}$ (i.e., x has no t -dominating set iff for every $w = (i_1, \dots, i_t) \in [n]^t$ there exists $w' = i_{t+1} \in [n]$ such that $\Phi_x(w, w') = 0$); that is, we used $\phi_{n^2} : \{0, 1\}^t \rightarrow \{0, 1\}$ and $\pi_{n^2, j} : \{0, 1\}^{(t+1) \log n} \rightarrow [n] \times [n]$ such that $\phi_{n^2}(\sigma_1, \dots, \sigma_t) = \bigvee_j \sigma_j$ and $\pi_{n^2, j}(i_1, \dots, i_t, i_{t+1}) = (i_j, i_{t+1})$ for $j \in [t]$.

Every locally $\forall\exists$ -characterizable set is in \mathcal{SC} (resp., in \mathcal{NC}) and so the existence of a doubly-efficient interactive proof system for it (and its complement) is guaranteed by [36] (resp., by [25] if the algorithm guaranteed by Definition 2.7 uses log-space). A direct construction that establishes the following result is provided in [23, Sec. 5.4].

Theorem 2.8 (Theorem 2.2, extended): *Every locally $\forall\exists$ -characterizable set has a simple doubly-efficient interactive proof system. Specifically, on input of length n , the verifier runs in $\tilde{O}(n)$ -time and the strategy of the prescribed prover can be implemented in $\tilde{O}(n^{c+c'+1})$ -time, where $n^{c+c'}$ denotes the number of local conditions in the characterization. Furthermore, the interactive proof system uses public coins and has logarithmic round complexity.*

Note that the complement of every locally $\forall\exists$ -characterizable set also has a doubly-efficient interactive proof system. In this proof system, on input $x \in \{0, 1\}^n$, letting $\ell = \log n$, the prover finds an adequate $w \in \{0, 1\}^{c\ell}$, sends it to the verifier, and the parties engage in a doubly-efficient interactive proof of the residual claim that *for every $w \in \{0, 1\}^{c'\ell}$ it holds that $\Phi_{x,w}(w') \stackrel{\text{def}}{=} \Phi_x(w, w')$ evaluates to 0.* (Such a proof system is provided by Theorem 2.2.)

Remark 2.9 (beyond $\forall\exists$ -characterization): *The notion of a locally $\forall\exists$ -characterizable set can be further extended to allow a constant number of (alternating) quantifiers; for example, a $\forall\exists\forall$ -characterization corresponds to the case that for all $w \in \{0, 1\}^{c \log n}$ there exists $w' \in \{0, 1\}^{c' \log n}$ such that for all $w'' \in \{0, 1\}^{c'' \log n}$ it holds that $\Phi_x(w, w', w'') = 0$. The proof of Theorem 2.8 extends naturally to that case (cf. the proof of Toda's Theorem [40]).*

Lastly, we note the correspondence between the foregoing local characterizations and the levels of a known hierarchy of *parameterized complexity classes* [16]. In particular, Definition 2.3 corresponds to a class denoted $\mathcal{W}[1]$, and Definition 2.7 corresponds to $\mathcal{W}[2]$. (In terms of the \mathcal{W} -hierarchy, our definitions are restricted in requiring that the “defining circuits” be more uniform.)

Chapter 3

On the doubly-efficient interactive proof systems of GKR

We present a somewhat simpler variant of the doubly-efficient interactive proof systems of Goldwasser, Kalai, and Rothblum [25]. Recall that these proof systems apply to log-space uniform sets in \mathcal{NC} (or, more generally, to inputs that are acceptable by log-space uniform bounded-depth circuits, where the number of rounds in the proof system is linearly related to the depth of the circuit). Our simplification is in the handling of the log-space uniformity condition. Rather than having the prover provide the verifier with bits of the encoding of the circuit and establish their correctness, we employ the proof system to a highly regular universal circuit that constructs and evaluates the log-space uniform circuit in question.

3.1 Overview

We provide an (alternative) exposition of the result of Goldwasser, Kalai, and Rothblum [25], which asserts the following –

Theorem 3.1 (doubly-efficient interactive proof systems for log-space uniform \mathcal{NC}): *Let $d : \mathbb{N} \rightarrow \mathbb{N}$ and A be an algorithm of logarithmic space complexity that on input 1^n outputs a Boolean circuit $C_n : \{0, 1\}^n \rightarrow \{0, 1\}$ of bounded fan-in and depth $d(n)$. Then, the set $S = \{x : C_{|x|}(x) = 1\}$ has an interactive proof system with $\text{poly}(\log n) \cdot d(n)$ rounds in which the prescribed prover strategy runs in polynomial-time while the verifier runs in $\tilde{O}(n + d(n))$ -time.*

(We note that the number of rounds is actually $O(\log n) \cdot d(n)$, but our simplified proof only yield a bound of $O(\log n)^2 \cdot d(n)$.)¹

The core of the proof system asserted in Theorem 3.1 is an iterative process in which a claim about the values of the gates that are at distance $i - 1$ from the output gate is reduced to a claim about the values of the gates at distance i . (All claims refer to the computation of $C_{|x|}(x)$.) Hence, in $d(n)$ iterations, the claim regarding the value of the output gate is reduced to a claim regarding the values of the bits of the inputs, whereas this claim can be verified in almost linear time.

¹The first (resp., second) expression omits an additive term of $O(\log n)^2$ (resp., $O(\log n)^2$), which is immaterial in the typical case where $d \geq \log n$.

The aforementioned claims refer to the values of specified locations in corresponding encodings (of the string describing all the gate-values at a certain layer of the circuit). Specifically, the encoding used is the low degree extension of the said string (viewed as a function), and the claims are claims about the evaluations of these polynomials at specific points.

The different codewords (or polynomials) are related via the structure of the circuit, and so (as usual) this structure is represented by a function that describes the adjacency relation (among the gates) of the circuit. We consider a low degree extension of this function, and the problem is allowing the verifier to evaluate this polynomial in almost-linear (in n) time. Here is where the uniformity condition comes into play.

Goldwasser, Kalai, and Rothblum used the uniformity condition in order to construct an interactive proof system by which the evaluation of this low degree polynomial is outsourced to the prover (who proves the validity of the provided answer) [25]. We present an alternative approach: We consider a “universal” circuit that, on input $x \in \{0, 1\}^n$, first constructs the log-space uniform circuit C_n , and next emulates the computation of $C_n(x)$. We apply the aforementioned interactive process to this universal circuit rather than to C_n , with the benefit being that this universal circuit is much more uniform than C_n . Consequently, a low degree extension² of the adjacency relation of that (universal) circuit can be evaluated in $\text{poly}(\log n)$ -time. We stress that the universal circuit preserves the complexities of C_n upto our level of interest; that is, the universal circuit has depth $\text{poly}(\log n) \cdot d(n)$ and size $\text{poly}(n)$.

Organization. The main module of the interactive proof system asserted in Theorem 3.1 is presented in Section 3.2, while avoiding the problem of evaluating the low degree extension of the function that represents the adjacency relation of the circuit. Coping with this computational problem is the contents of Section 3.3. Additional details regarding Sections 3.2 and 3.3 are presented in Section 3.4 (see Sections 3.4.1 and 3.4.2, resp.).

3.2 The main module

For simplicity (and w.l.o.g.), we assume that C_n has fan-in two and uses only NAND-gates. Viewing this gates as operating in $\text{GF}(2)$, we have $\text{NAND}(a, b) = 1 - (a \cdot b)$.

By augmenting the circuit with gates that are fed by no gate (and feed no gate), we can present the circuit as having $d(n) + 1$ layers of gates such that each layer has exactly $k(n) = \text{poly}(n)$ gates, where (by convention) gates that are fed nothing always evaluate to 0. As usual, the gates at layer i are only fed by gates at layer $i + 1$, and the leaves (at layer $d(n)$) are input-variables or constants. Furthermore, we may assume that, for $j \in [n]$, the j^{th} leaf is fed by the j^{th} input-variable, and all other leaves are fed by the constant 0. (The output is produced at the first gate of layer zero.)

The high level protocol. On input $x \in \{0, 1\}^n$, the prescribed prover computes the values of all layers. Letting $d = d(n)$ and $k = k(n)$, we denote the values at the i^{th} layer by $\alpha_i \in \{0, 1\}^k$; in particular, $\alpha_d = x0^{k-n}$ and $\alpha_0 = C_n(x)0^{k-1}$. For a sufficiently large finite field, denoted \mathcal{F} , we let $H \subset \mathcal{F}$ be a fixed set of size approximately $\log n$ and $m = \log_{|H|} k = (\log k) / \log |H|$ such that $k = |H|^m$. Viewing each α_i as a function from H^m to $\{0, 1\}$, the prover encodes α_i using its low degree extension $\hat{\alpha}_i : F^m \rightarrow \mathcal{F}$; that is,

²Actually, we use a low degree polynomial that agrees with the function; this low degree polynomial may have higher degree than the standard low degree extension.

$$\widehat{\alpha}_i(z_1, \dots, z_m) = \sum_{\sigma_1, \dots, \sigma_m \in H} \text{EQ}(z_1 \cdots z_m, \sigma_1 \cdots \sigma_m) \cdot \alpha_i(\sigma_1, \dots, \sigma_m) \quad (3.1)$$

where EQ is a low degree extension of the function that tests equality over H^m (e.g., $\text{EQ}(z_1 \cdots z_m, \sigma_1 \cdots \sigma_m) = \prod_{i \in [m]} \prod_{\beta \in H'} (z_i - \sigma_i + \beta) / \beta$, where $H' = \{\beta' - \beta'' : \beta', \beta'' \in H\} \setminus \{0\}$). Hence, proving that $x \in S$ is equivalent to proving that $\widehat{\alpha}_0(1^m) = 1$, where $1^m \in H^m$ corresponds to the fixed (e.g., first) location of the output gate in the zero layer.

This proof is conducted in $d(n)$ iterations, where in each iteration a multi-round interactive protocol is employed. Specifically, in i^{th} iteration, the correctness of the claim $\widehat{\alpha}_{i-1}(\bar{r}_{i-1}) = v_{i-1}$, where $\bar{r}_{i-1} \in \mathcal{F}^m$ and $v_{i-1} \in \mathcal{F}$ are known to both parties, is reduced (via the interactive protocol) to the claim $\widehat{\alpha}_i(\bar{r}_i) = v_i$, where $\bar{r}_i \in \mathcal{F}^m$ and $v_i \in \mathcal{F}$ are determined (by this protocol) such that both parties get these values. We stress that, with the exception of $i = d$, the $\widehat{\alpha}_i$'s are not known (or given) to the verifier; still, the claims made at the beginning (and at the end) of each iteration are well defined (i.e., they refer to the low degree extension of the values at certain layers of the circuit in a computation of the circuit on input $x \in \{0, 1\}^n$).

After the last iteration, the verifier is left with a claim of the form $\widehat{\alpha}_d(\bar{r}_d) = v_d$, which it can verify using Eq. (3.1) and its knowledge of $\alpha_d = x0^{k-n}$. Actually, some care is required here too, since the verifier needs to operate in time $\tilde{O}(n)$ (rather than in time $\tilde{O}(k)$, where $k = |H|^m = \text{poly}(n)$). For example, associating $[n]$ with $H^{m'} \equiv H^{m'} 1^{m-m'}$, we use the fact that for every $\sigma_1, \dots, \sigma_m \in H$ it holds that $\alpha_d(\sigma_1, \dots, \sigma_m) = x_{\sigma_1 \cdots \sigma_m}$ if $\sigma_{m-m'+1} \cdots \sigma_m = 1^{m-m'}$ and $\alpha_d(\sigma_1, \dots, \sigma_m) = 0$ otherwise. Hence,

$$\begin{aligned} \widehat{\alpha}_d(z_1, \dots, z_m) &= \sum_{\sigma_1, \dots, \sigma_m \in H} \text{EQ}(z_1 \cdots z_m, \sigma_1 \cdots \sigma_m) \cdot \alpha_d(\sigma_1, \dots, \sigma_m) \\ &= \sum_{\sigma_1, \dots, \sigma_{m'} \in H} \text{EQ}(z_1 \cdots z_m, \sigma_1 \cdots \sigma_{m'} 1^{m-m'}) \cdot \alpha_d(\sigma_1, \dots, \sigma_{m'} 1^{m-m'}). \end{aligned}$$

where $\alpha_d(\sigma_1, \dots, \sigma_{m'} 1^{m-m'})$ equals the bit of x in location $\sigma_1 \cdots \sigma_{m'} \in H^{m'} \equiv [n]$.

A single iteration. The core of the iterative proof is the interactive protocol that is performed in each iteration. This protocol is based on the relation between subsequent α_i 's, which is based on the structure of the circuit. Specifically, recall that the i^{th} iteration reduces a claim regarding $\widehat{\alpha}_{i-1}$ to a claim regarding $\widehat{\alpha}_i$, where these polynomials encode the values of neighboring layers in the circuit. To describe the relation between such neighboring layers, we consider the adjacency relation of the circuit. Actually, we let $\phi_i : H^m \times H^m \times H^m \rightarrow \{0, 1\}$ represent the feeding relation between the i^{th} and $(i-1)^{\text{st}}$ layer of the circuit such that $\phi_i(\bar{u}, \bar{v}, \bar{w}) = 1$ if and only if the gate in location $\bar{u} \in H^m$ in layer $i-1$ is fed by the gates in locations \bar{v} and \bar{w} in layer i . Hence, for every $\bar{u} \in H^m$, we have

$$\alpha_{i-1}(\bar{u}) = \sum_{\bar{v}, \bar{w} \in H^m} \phi_i(\bar{u}, \bar{v}, \bar{w}) \cdot (1 - \alpha_i(\bar{v})) \cdot \alpha_i(\bar{w}) \quad (3.2)$$

Recall that, by Eq. (3.1), we have $\widehat{\alpha}_{i-1}(\bar{z}) = \sum_{\bar{u} \in H^m} \text{EQ}(\bar{z}, \bar{u}) \cdot \alpha_{i-1}(\bar{u})$. Combining this with Eq. (3.2), while letting $\widehat{\phi}_i : \mathcal{F}^{3m} \rightarrow \mathcal{F}$ denote the low degree extension of ϕ_i , we have

$$\widehat{\alpha}_{i-1}(\bar{z}) = \sum_{\bar{u}, \bar{v}, \bar{w} \in H^m} \text{EQ}(\bar{z}, \bar{u}) \cdot \widehat{\phi}_i(\bar{u}, \bar{v}, \bar{w}) \cdot (1 - \widehat{\alpha}_i(\bar{v})) \cdot \widehat{\alpha}_i(\bar{w}). \quad (3.3)$$

For any fixed point $\bar{r}_{i-1} \in F^m$, the expression at the r.h.s of Eq. (3.3) can be written as $\sum_{\bar{u}, \bar{v}, \bar{w} \in H^m} p_{\bar{r}_{i-1}}(\bar{u}, \bar{v}, \bar{w})$, where $p_{\bar{r}_{i-1}}(\bar{z}', \bar{z}'', \bar{z}''') \stackrel{\text{def}}{=} \text{EQ}(\bar{r}_{i-1}, \bar{z}') \cdot \widehat{\phi}_i(\bar{z}', \bar{z}'', \bar{z}''') \cdot (1 - \widehat{\alpha}_i(\bar{z}'') \cdot \widehat{\alpha}_i(\bar{z}'''))$ is a low degree polynomial. Applying the sum-check protocol to Eq. (3.3) allows to reduce a claim regarding the value of $\widehat{\alpha}_{i-1}$ at a specific point $\bar{r}_{i-1} \in F^m$ to a claim regarding the value of the polynomial $p_{\bar{r}_{i-1}}$ at a random point $(\bar{r}', \bar{r}'', \bar{r}''')$ in F^{3m} (for details see Section 3.4.1). This is not quite what we aimed for but it is close enough, since the verifier can easily evaluate EQ at any point. However, there are two problems:

1. The verifier needs to be able to evaluate $\widehat{\phi}_i$ in $\widetilde{O}(n)$ -time. The definition of $\widehat{\phi}_i$ does imply that evaluating $\widehat{\phi}_i$ reduces to $|H|^{3m} = k^3$ evaluations of ϕ_i , but $n \leq k = \text{poly}(n)$ which may be much larger than n . Hence, evaluating $\widehat{\phi}_i$ in $\widetilde{O}(n)$ -time is not obvious.
2. After evaluating EQ and $\widehat{\phi}_i$, the verifier is left with a residual claim that refers to two evaluations of $\widehat{\alpha}_i$ rather than to a single one.

The second problem is handled by augmenting the protocol such that the prover provides the value of $\widehat{\alpha}_i$ at both points (i.e., at \bar{r}'' and \bar{r}''') as well as on the line that connects these points, and is asked to prove the correctness of the value at a random point on this line. That is, if the prover claims that $\widehat{\alpha}_i(\bar{r}'') = v''$ and $\widehat{\alpha}_i(\bar{r}''') = v'''$, then it also provides a low degree univariate polynomial p' such that $p'(z) = \widehat{\alpha}_i((1-z)\bar{r}'' + z\bar{r}''')$, which satisfies that $p'(0) = v''$ and $p'(1) = v'''$, and the verifier selects a random $r \in \mathcal{F}$ and sends it to the prover. The claim to be proved in the next iteration is that the value of $\widehat{\alpha}_i$ at $\bar{r}_i = (1-r)\bar{r}'' + r\bar{r}'''$ equals $v_r = p'(r)$. Hence, the full protocol that is run in iteration i is as follows.

Construction 3.2 (reducing a claim about $\widehat{\alpha}_{i-1}$ to a claim about $\widehat{\alpha}_i$): *For known $\bar{r}_{i-1} \in \mathcal{F}^m$ and $v_{i-1} \in \mathcal{F}$, the entry claim is $\widehat{\alpha}_{i-1}(\bar{r}_{i-1}) = v_{i-1}$. The parties proceed as follows.*

1. *Applying the sum-check protocol to the entry claim, when expanded according to Eq. (3.3), determines $\bar{r}', \bar{r}'', \bar{r}''' \in \mathcal{F}^m$ and a value $v \in \mathcal{F}$ such that the residual claim for verification is*

$$\text{EQ}(\bar{r}_{i-1}, \bar{r}') \cdot \widehat{\phi}_i(\bar{r}', \bar{r}'', \bar{r}''') \cdot (1 - \widehat{\alpha}_i(\bar{r}'') \cdot \widehat{\alpha}_i(\bar{r}''')) = v. \quad (3.4)$$

2. *The prover sends a univariate polynomial p' of degree $(|H| - 1) \cdot m$ such that $p'(z) = \widehat{\alpha}_i((1-z)\bar{r}'' + z\bar{r}''')$.*
3. *The verifier checks whether v equals*

$$\text{EQ}(\bar{r}_{i-1}, \bar{r}') \cdot \widehat{\phi}_i(\bar{r}', \bar{r}'', \bar{r}''') \cdot (1 - p'(0) \cdot p'(1)) \quad (3.5)$$

and continues only if equality holds (otherwise it rejects).

Note that this requires evaluating EQ and $\widehat{\phi}_i$.

4. *The verifier selects a random $r \in \mathcal{F}$, and sends it to the prover. Both parties set $\bar{r}_i = (1-r)\bar{r}'' + r\bar{r}'''$ and $v_r = p'(r)$.*

The exit claim is $\widehat{\alpha}_i(\bar{r}_i) = v_i$.

The prescribed prover strategy in Construction 3.2 can be implemented in $\text{poly}(n)$ -time. Assuming that $\widehat{\phi}_i$ can be evaluated in $\text{poly}(\log n)$ -time, the verifier's strategy in Construction 3.2 can be implemented in $\text{poly}(\log n)$ -time, provided that $\log |\mathcal{F}| \leq \text{poly}(\log n)$ (whereas we will use $|\mathcal{F}| = \text{poly}(\log n)$). One can readily verify that if the entry claim is correct, then the exit claim is correct, whereas if the entry claim is false, then with probability at least $1 - O(|H| \cdot m / |\mathcal{F}|)$ the exit claim is false.

We are still left with the problem of evaluating $\widehat{\phi}_i$ in $\text{poly}(\log n)$ -time. Actually, in Section 3.3 we present a solution to a relaxed version of this problem, in which $\widehat{\phi}_i$ is replaced by an arbitrary polynomial of degree at most D that agrees with ϕ_i on H^m . Before turning to Section 3.3, we note that the foregoing analysis remains valid, except that the error probability grows from $O(|H| \cdot m / |\mathcal{F}|)$ to $O(D + |H| \cdot m) / |\mathcal{F}|$.

3.3 Evaluating $\widehat{\phi}_i$

The general problem that we face is evaluating the low degree extension of a function $\phi : H^m \rightarrow \{0, 1\}$ that is computable in $O(\log n)$ space, where $|H|^m = \text{poly}(n)$. (In our setting we are interested in ϕ_i 's, which range over H^{3m} , which can be treated as slices of a single $\phi : H^{3m+1} \rightarrow \{0, 1\}$ such that $\phi_i(\bar{z}) = \phi(i, \bar{z})$, where $i \in [m]$ is viewed as an element of H .)³ This problem is solved in [25] by using an auxiliary multi-round interactive protocol, but here we bypass it by making several observations.

Recall that the function ϕ that we deal with here describes the structure of arbitrary log-space uniform circuits. The first observation is that we may use instead a function that describes the structure of a highly-uniform *universal circuit*. For starters, consider a universal circuit U_n that, given a description of a $\text{poly}(n)$ -sized circuit $C_n : \{0, 1\}^n \rightarrow \{0, 1\}$ of depth $d = d(n)$ and an input $x \in \{0, 1\}^n$, computes $C_n(x)$; that is, $U_n(\langle C_n \rangle, x) = C_n(x)$. By using a suitable description, we may get a highly uniform circuit U_n of $\text{poly}(n)$ -size and depth $O(\log n) \cdot d(n)$. The idea is applying the main module to U_n rather than to C_n .

The problem with the foregoing idea is that after the last iteration (i.e., the last application of Construction 3.2), the verifier needs to evaluate the low degree extension of the input to the circuit, whereas in our application the input to U_n consists of both $\langle C_n \rangle$ and x , which means that the verifier needs to compute $\langle C_n \rangle$. By the log-space uniformity condition, this can be done in $\text{poly}(n)$ -time, but we wish the verifier to perform this step in $\widetilde{O}(n)$ -time. So, instead, we consider a modified universal circuit U'_n , which is tailored to the log-space algorithm A that constructs the circuit family $\{C_n\}$. On input $x \in \{0, 1\}^n$, the circuit U'_n first constructs C_n (see details below), and then computes $U_n(\langle C_n \rangle, x)$. Now, applying the main module to U'_n rather than to U_n , the verifier only need to construct a low degree extension of x (as in the case that we use C_n itself).

Turning to the construction of C_n by U'_n , following [25], we use the observation that log-space uniform circuits (here $\{C_n\}_{n \in \mathbb{N}}$) can be constructed by an *highly uniform* \mathcal{NC} circuit. This observation relies on the standard construction of \mathcal{NC} circuits for log-space computation, a construction that is based on constructing the digraph of instantaneous configurations of the computation of a log-space machine (on a fixed input) and raising it to a power larger than its dimension (by repeated squaring). For details see Section 3.4.2. Recalling that U_n is also highly uniform (see Section 3.4.2 for details) and combining these two constructions, we obtain a highly uniform construction of U'_n .

³Recall that $m = O(\log n) / \log \log n < |H| \approx \log n$. For simplicity of notation, we write $\phi : H^m \rightarrow \{0, 1\}$ rather than $\phi : H^{3m+1} \rightarrow \{0, 1\}$.

So far, we avoided specifying what we meant by a highly uniform circuit. A notion that suffices for our purposes is presented now. Considering circuits of size $s(n) = \text{poly}(n)$ and letting $\ell = 3 \log s(n)$, we consider the predicate $\phi : \{0, 1\}^\ell \rightarrow \{0, 1\}$ that describes the adjacency relation of the circuit. We say that the circuit is **highly uniform** if ϕ can be computed by $\text{poly}(\ell)$ -sized formula that can be constructed in $\text{poly}(\ell)$ -time. The reader may verify that the \mathcal{NC} circuits that we used above (for constructing the log-space uniform circuits C_n) are highly uniform; in fact, they are even “more uniform” than that. Likewise for the circuit U_n , and consequently for U'_n . Hence, the predicate $\phi : \{0, 1\}^\ell \rightarrow \{0, 1\}$ that describes the adjacency relation of U'_n can be computed by a $\text{poly}(\log n)$ -size formula that can be constructed in $\text{poly}(\log n)$ -time. Viewing ϕ as ranging over H^m , where $m = \ell / \log |H|$, the following result implies that there exists a low degree polynomial over \mathcal{F}^m that agrees with ϕ on H^m and can be evaluated in $\text{poly}(\log n)$ -time.⁴

Theorem 3.3 (evaluating a low degree polynomial that agrees with a highly uniform Boolean circuit): *For $\phi : H^m \rightarrow \{0, 1\}$, where $H \equiv \{0, 1\}^{\ell'}$, let $\phi' : \{0, 1\}^{m\ell'} \rightarrow \{0, 1\}$ be the corresponding Boolean function that views elements of H as ℓ' -bit long strings. Suppose that ϕ' is highly uniform; that is, ϕ' can be computed by formula of size $s' = \text{poly}(m\ell')$ that can be constructed in $\text{poly}(s')$ -time. Then, there exists a polynomial $\widehat{\Phi} : \mathcal{F}^m \rightarrow \mathcal{F}$ of degree $\text{poly}(s') \cdot |H|$ that agrees with ϕ on H^m and can be evaluated in $\text{poly}(\ell)$ -time, where $\ell = \log |H|^m$, provided that $|H| = \text{poly}(m)$.*

We stress that the polynomial $\widehat{\Phi} : \mathcal{F}^m \rightarrow \mathcal{F}$ is *not* the (canonical) low degree extension of ϕ (or ϕ'); such a low degree extension has degree $m \cdot |H|$ (or $m\ell' \cdot |H|$) and exists regardless of the size of the formula ϕ . In contrast, the polynomial $\widehat{\Phi}$ is derived from the formula computing ϕ' , and its degree bound is derived from the depth of ϕ' . (As stated above, we apply Theorem 3.3 to the function $\phi : H^{3m+1} \rightarrow \{0, 1\}$ that describes the relation between gates in a circuit as in Section 3.2, where $\phi_i(\bar{z}) = \phi(i, \bar{z})$ for every $i \in [m] \subset H$.)

Proof: Associating H with $\{0, 1\}^{\ell'}$, we first consider the computation of the j^{th} bit in the binary representation of $z \in H$. Denote by H_j the set of elements in H that are represented by ℓ' -bit strings in which the j^{th} position holds the value 1. Observe that, for every $z \in H$, it holds that the expression $\sum_{a \in H_j} \prod_{b \in H \setminus \{a\}} (z - b) / (a - b)$ (with arithmetic of the field \mathcal{F}) equals 1 if $z \in H_j$ (i.e., the j^{th} bit of z is 1) and equals 0 otherwise (i.e., if the j^{th} bit of z is zero).

Next, consider the $\text{poly}(\ell)$ -size formula that computes $\phi' : \{0, 1\}^\ell \rightarrow \{0, 1\}$, where $\ell = m\ell'$. Then, for $z_1, \dots, z_m \in H$, it holds that

$$\phi(z_1, \dots, z_m) = \phi'(z_{i,1}, \dots, z_{i,\ell'}, \dots, z_{m,1}, \dots, z_{m,\ell'}), \quad (3.6)$$

where $z_{i,j} \leftarrow \sum_{a \in H_j} \text{EQ}(z_i, a) \in \{0, 1\}$ (and $\text{EQ} : \mathcal{F} \times \mathcal{F} \rightarrow \mathcal{F}$ is a low degree extension of the function that tests equality over H ; i.e., $\text{EQ}(z, a) = \prod_{b \in H \setminus \{a\}} (z - b) / (a - b)$). Now, construct an arithmetic circuit $\widehat{\Phi}$ (which is actually a formula) that computes a polynomial that agrees with ϕ on H^m by mimicing the Boolean formula ϕ' . Specifically, the input gate $z_{i,j}$ is replaced by the formula $\sum_{a \in H_j} \text{EQ}(z_i, a)$, and the NAND of the Boolean subformulae ϕ'_1 and ϕ'_2 is replaced by a gate that computes $1 - \widehat{\Phi}_1 \cdot \widehat{\Phi}_2$, where $\widehat{\Phi}_1$ and $\widehat{\Phi}_2$ are the corresponding arithmetic subformulae. (Indeed, we may assume, w.l.o.g., that ϕ' uses only NAND-gates, and that its depth is $O(\log s')$.)

Note that the degree of the resulting polynomial $\widehat{\Phi}$ is upper-bounded by $\text{poly}(s') \cdot |H|$, where the $|H|$ factor is due to the computation of the $z_{i,j}$'s and the $\text{poly}(s')$ is exponential in the depth of

⁴Recall that $|H| \approx \log n$.

the formula ϕ' (which, w.l.o.g., is logarithmic in its size). To see that $\widehat{\Phi}$ equals ϕ on H^m , note that when evaluated on input in H^m each of the $z_{i,j}$'s is assigned a Boolean value, whereas the residual subformula preserves Boolean values. Hence, the value of ϕ on a point in H^m is given by the value of ϕ' on the corresponding bit string (per Eq. (3.6)), which in turn equals the value of $\widehat{\Phi}$ on this bit string. ■

Wrapping-up. As stated above, we modify Construction 3.2 by replacing the polynomials $\widehat{\phi}_i$ by the polynomial $\widehat{\Phi}$ (i.e., replacing $\widehat{\phi}_i(\cdot)$ by $\widehat{\Phi}(i, \cdot)$ where $\widehat{\phi}_i : \mathcal{F}^{3m} \rightarrow \mathcal{F}$ and where $\widehat{\Phi}_i : \mathcal{F}^{1+3m} \rightarrow \mathcal{F}$). Doing so provides an efficient implementation of Step 3 of Construction 3.2, whereas the soundness error grows from $d(n) \cdot O(|H| \cdot m/|\mathcal{F}|)$ to $d(n) \cdot O((m + \text{poly}(s') \cdot |H|) \cdot |H|/|\mathcal{F}|)$, where s' is as in Theorem 3.3. Hence, we need to guarantee that $O((m + \text{poly}(s')) \cdot |H|^2/|\mathcal{F}|) = o(1/d(n))$, where $s' = \text{poly}(m \log |H|)$ is an upper bound on the size of the formula ϕ' that we use. (Note that $s' \geq 3m \log |H|$ definitely holds, whereas $s' = O(m \log |H|)$ may actually suffice.) Recalling that $|H| \approx \log n$ and $m < \log n$, it suffices to have $|\mathcal{F}| = \omega(d(n) \cdot \text{poly}(\log n))$, which implies that $|\mathcal{F}| = \text{poly}(|H|) \cdot d(n)$ suffices. This completes the proof of Theorem 3.1, modulo some details provided in Section 3.4.

3.4 Details

In this section we provide some tedious details, which some readers may find quite straightforward.

3.4.1 Applying the sum-check protocol to Eq. (3.3)

Here we detail how the sum-check protocol is used to reduce the claim regarding the value of $\widehat{\alpha}_i$ at a given point \bar{r}_{i-1} to a claim regarding $\widehat{\alpha}_{i-1}$. Towards this goal, we define a polynomial $p_{\bar{r}_{i-1}} : \mathcal{F}^{3m} \rightarrow \mathcal{F}$ such that

$$p_{\bar{r}_{i-1}}(\bar{z}', \bar{z}'', \bar{z}''') = \text{EQ}(\bar{r}_{i-1}, \bar{z}') \cdot \widehat{\phi}_i(\bar{z}', \bar{z}'', \bar{z}''') \cdot (1 - \widehat{\alpha}_i(\bar{z}'') \cdot \widehat{\alpha}_i(\bar{z}''')). \quad (3.7)$$

Recall that the claim that we wish to verify is $\widehat{\alpha}_{i-1}(\bar{r}_{i-1}) = v_{i-1}$. Expanding $\widehat{\alpha}_{i-1}$ according to Eq. (3.3) and using Eq. (3.7), we have

$$\widehat{\alpha}_{i-1}(\bar{r}_{i-1}) = \sum_{\bar{u}, \bar{v}, \bar{w} \in H^m} p_{\bar{r}_{i-1}}(\bar{u}, \bar{v}, \bar{w}). \quad (3.8)$$

Using $3m$ rounds of interaction, the sum-check protocol reduces the claim that the r.h.s of Eq. (3.8) equals v_{i-1} to a claim regarding the value of $p_{\bar{r}_{i-1}}$ at a random point $(\bar{r}', \bar{r}'', \bar{r}''') \in \mathcal{F}^{3m}$. This is captured by Step 1 of Construction 3.2.

3.4.2 The highly uniform circuits in use

Here we detail the construction of the universal circuit U'_n , which combines a construction of \mathcal{NC} circuits for constructing log-space uniform circuits and a construction of the universal circuit U_n . The point that we wish to establish here is that both these constructions are highly uniform (in the sense that the adjacency relation function describing each of these two $\text{poly}(n)$ -sized circuit can be computed by a $\text{poly}(\log n)$ -size formula that can be computed in $\text{poly}(\log n)$ -time).

We start with the highly uniform construction of \mathcal{NC} circuits for constructing log-space uniform circuits, denoted $\{C_n\}_{n \in \mathbb{N}}$. Recall that this construction relies on the standard construction of \mathcal{NC} circuits for log-space computations, a construction that is based on constructing the digraph of instantaneous configurations of the computation of a log-space machine (on a fixed input) and raising it to a power larger than its dimension (by repeated squaring). In our case, the input (i.e., 1^n) is actually used only to determine the space bound, and so we can ignore it. However, we should note that the said log-space computation produces many output bits (which describe the circuit C_n , rather than a single decision bit). Without loss of generality, we can incorporate the index of the next bit to be produced in the instantaneous configuration (since the log-space machine can maintain a corresponding counter).

Constructing the digraph of instantaneous configurations. Observe that the digraph that represents the consecutive configurations in the foregoing log-space computation can be constructed very easily (i.e., by a highly uniform \mathcal{NC} circuit); that is, the pair $(\gamma, \gamma') \in \{0, 1\}^{2 \cdot O(\log n)}$ is an edge in that digraph if and only if the log-space machine moves from the instantaneous configuration represented by γ to the configuration represented by γ' in a single step. Note that we actually consider an \mathcal{NC} circuit that has no input and always produces this digraph as its output. The output gate that corresponds to the entry (γ, γ') in the adjacency matrix representing this digraph is 1 if and only if the pair (γ, γ') satisfies a very simple relation that is only slightly more complex than testing equality.⁵ Hence, the value of the entry (γ, γ') in the matrix can be determined in $\text{poly}(\log n)$ -time as a function of γ and γ' . It follows that the circuit that produces this matrix, which consists merely of computing the constants 0 and 1 and feeding them to the $\text{poly}(n)$ -many output bits, is highly uniform; that is, its adjacency relation can be expressed a formula of size $O(\log n)$ that can be constructed in time $\text{poly}(\log n)$.

This yields a highly uniform $\text{poly}(n)$ -by- $\text{poly}(n)$ matrix M that when raised to the power $\text{poly}(n)$ yields the desired bits; that is, the value of the i^{th} output bit (i.e., the i^{th} bit in the description of C_n) equals the value of the entry $(\mathbf{s}_i, \mathbf{f}_i)$ in the resulting matrix (i.e., of $M^{\text{poly}(n)}$), where \mathbf{s}_i denotes the initial configuration that is used for producing the i^{th} bit and \mathbf{f}_i denotes the final configuration that actually produces this bit. As we show next, there exists a highly uniform \mathcal{NC} circuit for raising a given $\text{poly}(n)$ -by- $\text{poly}(n)$ matrix to the power $\text{poly}(n)$ (by repeated squaring), since the inner product of $\text{poly}(n)$ -bit long vectors has a highly uniform \mathcal{NC} circuit.

Constructing circuits for matrix multiplication. We now detail the highly uniform \mathcal{NC} circuit used to compute matrix multiplication. Let A and B be n -by- n matrices, then the $(i, j)^{\text{th}}$ bit in their product is given by $\sum_{k \in [n]} A_{i,k} B_{k,j}$. Hence, we first use n^3 gates such that the gate indexed by (i, j, k) computes $A_{i,k} B_{k,j}$, which means that this gate is fed by inputs with indices $(1, i, k)$ and $(2, k, j)$, and then we use n^2 trees of $n-1$ addition-gates for computing each n -way sum. Specifically, the sum that corresponds to index (i, j) is computed by gates that are indexed by (i, j, α) , where $\alpha \in \{0, 1\}^{\leq \log n}$, such that the gate indexed by (i, j, α) is fed by the gates indexed $(i, j, \alpha 0)$ and $(i, j, \alpha 1)$. Hence, the description of the adjacency relation of this circuit is very simple; for example,

⁵Specifically, recall that γ represents the contents of the work-space as well as the finite state of the machine and the location of its head (on the work-tape). The bits of γ' should equal the corresponding bits of γ , except for the bit at the head's location and the bits encoding the finite state that change according to a finite function. (The location of the head on the work-tape is indicated by the corresponding location of a special symbol in γ .) Hence, the value of the entry (γ, γ') can be determined as a conjunction of $O(\log n)$ conditions, where each of these conditions refers to a constant number of bits.

for gates that correspond to the n -way sums, we can use

$$\phi((i, j, \alpha), (i', j', \alpha'), (i'', j'', \alpha'')) = \text{EQ}(ij, i'j') \wedge \text{EQ}(ij, i''j'') \wedge \text{EQ}(\alpha 0, \alpha') \wedge \text{EQ}(\alpha 1, \alpha'').$$

Constructing the universal circuit U_n . Recall that, for fixed depth $d = d(n)$ and size $s = s(n)$, the circuit U_n is given a description of a circuit C_n (of depth d and size s) and a string $x \in \{0, 1\}^n$, and is supposed to output $C_n(x)$. The circuit U_n emulates the computation of C_n in a layer-by-layer manner. Specifically, U_n computes the value of the j^{th} gate in layer i in the emulated computation of C_n (on x), denoted $v_{i,j}$, as a function of the description of the circuit C_n and the two adequate values of gates in layer $i + 1$. In particular, $v_{d,j}$ is the j^{th} bit of x if $j \leq n$ and equals 0 otherwise, and for every $i \in [d]$ and $j \in [s]$,

$$v_{i-1,j} = \bigvee_{k',k'' \in [s]} (c_{(i-1,j),(i,k'),(i,k'')} \wedge \text{NAND}(v_{i,k'}, v_{i,k''})),$$

where $c_{(i-1,j),(i,k'),(i,k'')}$ is a bit in the description of C_n that indicates whether or not the $(i - 1, j)^{\text{th}}$ gate is fed by gates indexed the (i, k') and (i, k'') . Again, the description of the adjacency relation of this circuit is very simple.

Chapter 4

Overview of the doubly-efficient interactive proof systems of RRR

We provide an overview of the doubly-efficient interactive proof systems of Reingold, Rothblum, and Rothblum [36]. Recall that by their result, any set that is decidable in polynomial-time by an algorithm of space complexity $s(n) \leq n^{0.499}$, has a constant-round interactive proof system in which the prover runs polynomial time and the verifier runs in time $\tilde{O}(n)$.

Theorem 4.1 (doubly-efficient interactive proof systems for \mathcal{SC}): *Suppose that S is decidable by an algorithm that runs in polynomial time and has space complexity s . Then, for any constant $\delta > 0$, the set S has an interactive proof system with $\exp(\tilde{O}(1/\delta))$ rounds in which the prescribed prover strategy runs in polynomial-time while the verifier runs in $(\tilde{O}(n) + \text{poly}(s(n)) \cdot n^\delta)$ -time.*

Actually, the unspecified polynomial in [36] is almost quadratic (priv. comm. with the authors of [36], May 2017). The proof of Theorem 4.1 spans 70 pages in [36], and in this note we merely attempt to present an overview of it.

4.1 High level structure

Theorem 4.1 is proved by considering the instantaneous configurations of a space bounded algorithm. The interactive proof system will refer to claims of the form *on input x , machine M moves from configuration γ to configuration γ' in t steps*. The initial claim refers to the case that $\gamma \in \{0, 1\}^{s(|x|)}$ and $\gamma' \in \{0, 1\}^{s(|x|)}$ are the initial and accepting configurations of M , respectively, and that $t \leq \text{poly}(|x|)$.

The core of the proof of Theorem 4.1 is reducing the construction of an interactive proof system for a claim regarding t -step computations (of a space-bounded machine) to the construction of an interactive proof system for a claim regarding t/k -step computations (of such a machine), where k is a parameter (e.g., $k = n^\delta$ or so). This is done by having the prover send $\gamma_1, \dots, \gamma_{k-1} \in \{0, 1\}^s$ and prove that, *for every $i \in [k]$, on input x , machine M moves from configuration γ_{i-1} to configuration γ_i in t/k steps, where $\gamma_0 = \gamma$ and $\gamma_k = \gamma'$.*

Hence, we have reduced proving a single claim regarding a t -long computation of M to k claims regarding t/k -long computations of M . But have we gained anything? This depends on whether we can perform k verifications (regarding t/k -long computations) at a cost that is lower

than performing each of these verifications separately. In other words, we seek a *batch verification* procedure in which verifying k claims is cheaper than verifying each of the claims separately.

To see that batch verification is not a pipe dream, consider the task of proving k claims that are each in \mathcal{PSPACE} . That is, for a fixed set $S \in \mathcal{PSPACE}$, given $x_1, \dots, x_k \in \{0, 1\}^n$, we wish to prove $(\forall i \in [k]) x_i \in S$. Then, the complexity of verifying the latter claim by employing the proof system of [32, 38] is only moderately larger than the complexity of verifying membership in S using this generic construction. Specifically, suppose that S has space complexity $s(n)$, and recall that verification via the generic system of [32, 38] takes time that is proportional to $s(n)^4$. Then, the time complexity of batch verification k claims (each of length n) is proportional to $(s(n) + \log k)^4$, which is only moderately larger than $s(n)^4$.

The following result provides a batch verification procedure (or rather an interactive proof for batch verification) for any set that has an interactive proof system of a certain type, where the nature of this type will remain unspecified momentarily.¹ We shall denote the length of a single instance by n , and the number of instances by $k(n)$. For a set S and $k : \mathbb{N} \rightarrow \mathbb{N}$, we let

$$S^k = \cup_{n \in \mathbb{N}} \{(x_1, \dots, x_{k(n)}) \in \{0, 1\}^{k(n) \cdot n} : (\forall i \in [k(n)]) x_i \in S\}. \quad (4.1)$$

Theorem 4.2 (batch verification for certain interactive proof systems): *For $c : \mathbb{N} \rightarrow \mathbb{N}$ and a constant $\ell \in \mathbb{N}$, suppose that S has an ℓ -round interactive proof system of a particular type \mathcal{T} with total communication c , poly-logarithmic verification time, and polynomial-time prover strategy. Then, for every constant $\alpha > 0$ and $k : \mathbb{N} \rightarrow \mathbb{N}$, the set S^k has an $O(\ell/\alpha)$ -round interactive proof system of type \mathcal{T} with total communication $\tilde{O}(k^\alpha \cdot c + k)$, poly-logarithmic verification time, and polynomial-time prover strategy.*

Using Theorem 4.2, the construction of the proof system asserted in Theorem 4.1 can be recursively presented as follows.

Construction 4.3 (recursive description of the interactive proof system): *Let $n, s, t, k \in \mathbb{N}$, where $s = s(n)$ and $k < t$, and fixed M and $x \in \{0, 1\}^n$. To prove that on input x , machine M moves in t steps from configuration $\gamma_0 \in \{0, 1\}^s$ to configuration $\gamma_k \in \{0, 1\}^s$, the parties proceed as follows.*

1. *The prover sends $\gamma_1, \dots, \gamma_{k-1} \in \{0, 1\}^s$ such that, for every $i \in [k]$, it holds that on input x , machine M moves in t/k steps from configuration γ_{i-1} to configuration γ_i .*

In order to support the hypothesis made in the next step, each γ_i is sent using an error correcting code. In addition, we also assume that x is presented by such a code.

2. *The parties invoke an interactive proof system in order to verify the foregoing claim regarding the γ_i 's, where the i^{th} claim refers to moving from γ_{i-1} to γ_i (in t/k steps on input x).² This interactive proof system is obtained by applying Theorem 4.2 to the interactive proof system that refers to t/k step computation, where the latter is obtained by a recursive call. Recall that Theorem 4.2 requires that the latter system be of type \mathcal{T} .*

Hence, at the bottom of the recursion we have communication complexity $\text{poly}(s)$, at its j^{th} level (from the bottom) we have communication complexity $O(k^{j^\alpha} \cdot \text{poly}(s) + k^{(j-1)\alpha+1})$, where

¹Readers who are too curious to wait may note that the said type is restricted to public coin systems that are “unambiguous” (akin to the notion of unique solutions in NP-proof systems) and support local checking of the prover messages (akin to PCPs (or rather PCPPs)).

²Hence, the input to the i^{th} claim is $x_i = (x, \gamma_{i-1}, \gamma_i)$.

α is the constant in Theorem 4.2 and the recursion has $\log_k t$ levels. Using $k^{(\log_k t) \cdot \alpha} = t^\alpha$, at the top level of the recursion, this yields communication complexity $O(t^\alpha \cdot (\text{poly}(s) + k))$, whereas verification time is poly-logarithmic, proving time is polynomial, and the number of rounds is $O(1/\alpha)^{\log_k t}$.

In order to support the use of the derived system in subsequent recursive calls, we will show that the derived system is indeed of type \mathcal{T} .

Theorem 4.1 follows by combining Construction 4.3 with Theorem 4.2, while setting $\alpha = \delta/O(1)$ and $k = n^\delta$. (Actually, Construction 4.3 is invoked after letting the verifier encode x as well as the initial and accepting configurations under an error correcting code.)

In light of Construction 4.3, we focus on the proof of Theorem 4.2. We start with the simple case in which the set has an NP-proof system, which holds at the very bottom of the recursion but not at higher levels.

4.2 Warm-up: Batch verification for NP

We start by introducing two of the restrictions that make up the foregoing type \mathcal{T} of proof systems. The first restriction asserts that a convincing prover strategy is essentially unique; that is, if the prover wishes to convince the verifier of the fact that $x \in S$, then it has no choice but to follow a unique strategy. In the case of NP-proof systems, this means that the NP-witness is unique; that is, $S \in \mathcal{UP}$.

The second restriction is that verification can be performed in poly-logarithmic time provided that the input is presented in encoded form, under a suitable error correcting code. This condition is satisfied by a variety of PCPP systems, starting from the ones based on Reed-Muller encoding (cf., [5, 17], as interpreted by [8, 14]). Readers that are unfamiliar with PCPPs (i.e., PCPs of Proximity), may consider standard PCP systems and verification in almost linear time. Actually, for simplicity we shall just do that, while focusing on the communication complexity of the derived interactive proof system for batch verification.

Hence, our starting point is a set $S \in \mathcal{UP}$ and a corresponding PCP system. Actually, we shall assume that the PCP system uses **input-oblivious queries**; that is, its queries are determined non-adaptively based solely on its internal coin tosses, and only the final decision depends on its actual input.³ Such PCP systems are known and can be derived from any PCPP system.

Theorem 4.4 (batch verification for \mathcal{UP}): *For $c : \mathbb{N} \rightarrow \mathbb{N}$ and a constant $\ell \in \mathbb{N}$, suppose that $S \in \mathcal{UP}$ has a PCP system that uses input-oblivious queries and proofs of length c . Then, for every constant $\alpha > 0$ and $k : \mathbb{N} \rightarrow \mathbb{N}$, the set S^k has an $O(1/\alpha)$ -round interactive proof system with total communication $O(k^\alpha \cdot c + k)$.*

Proof Sketch: For a parameter $d \leq k$ to be determined (e.g., $d = \sqrt{k}$), consider a parity-check function $F : \{0, 1\}^k \rightarrow \{0, 1\}^{O(d \log k)}$ such that for every $y \in \{0, 1\}^k$ and $v \in \{0, 1\}^{O(d \log k)}$ the Hamming ball of radius d centered at y contains at most a single pre-image of v under F . (Indeed, in the case $d = 1$ we can let $F(y)$ be the XOR of the bits of y , and in the general case we can use a Reed-Solomon code.)

³Hence, these PCP systems are not truly input-oblivious in the sense studied in [22].

Construction 4.4.1 (basic construction): *On input $(x_1, \dots, x_k) \in \{0, 1\}^{k \cdot n}$, the proof system proceeds as follows.*

1. *Let w_i be the unique NP-witness associated with x_i , and $\pi_i \in \{0, 1\}^{c(n)}$ be the PCP-oracle derived from it. (We assume that the NP-witness is easy to extract from the PCP-oracle.)⁴ For each $j \in [c(n)]$, let $\pi_{i,j}$ denote the j^{th} bit of π_i . Then, for every $j \in [c(n)]$, the prover computes $v_j \leftarrow F(\pi_{1,j} \cdots \pi_{k,j})$, and sends $v_1, \dots, v_{c(n)}$ to the verifier. Let us denote the values actually sent by $\tilde{v}_1, \dots, \tilde{v}_{c(n)}$.*

Pictorially, the prover forms a k -by- $c(n)$ Boolean matrix such that the i^{th} row equals π_i , and applies the parity check function to each column. These values form a very partial commitment of the prover to the values of the matrix; once these values are determined, each column is pseudo-determined in the sense that a valid revealing may either equal the original column or must differ from it on more than d corruptions.

2. *The verifier generates a sequence of queries, $j_1, \dots, j_q \in [c(n)]$, for the PCP verifier, denoted V . It sends this sequence to the prover, who responds with the values of the corresponding columns. The verifier checks (i) whether the values of these columns match the parity-check values, and (ii) whether V would have accepted each of the inputs when given the corresponding answers. Specifically, suppose that the prover answered with $(\tilde{\pi}_{i,j_{i'}})_{i \in [k], i' \in [q]}$. Then, the verifier performs the following two checks:*

- (a) *For every $i' \in [q]$, it checks whether $F(\tilde{\pi}_{1,j_{i'}} \cdots \tilde{\pi}_{k,j_{i'}}) = \tilde{v}_{j_{i'}}$.*
- (b) *For each $i \in [k]$, it checks whether V would have accepted x_i when making the queries j_1, \dots, j_q and receiving the answers $\tilde{\pi}_{i,j_1}, \dots, \tilde{\pi}_{i,j_q}$.*

3. *The verifier selects uniformly a set R of $O(k/d)$ rows, sends their indices to the prover, who provides their contents. The verifier checks that (i) each of these rows equals the PCP-oracle derived from the unique NP-witness for the corresponding input, and (ii) that the values of these rows match the values of the columns provided in Step 2. Specifically, denoting the value of row i as sent in this step by r_i , the verifier performs the following two checks (for each such $i \in R$):*

- (a) *Letting w'_i denote the purported NP-witness extracted from r_i , the verifier checks that w'_i is a valid NP-witness for x_i and that r_i is the PCP-oracle derived from w'_i .*
- (b) *For every $i' \in [q]$, it checks whether $r_{i,j_{i'}} = \tilde{\pi}_{i,j_{i'}}$, where $r_i = r_{i,1} \cdots r_{i,c(n)}$*

The basic intuition is that Step 1 leaves the prover with the option of either cheating in Step 2 on more than d entries of one of the columns or providing the correct values for all columns. In the first case, the prover is likely to be caught cheating in Step 3, whereas in the second case it is likely that V will reject $x_i \notin S$ (when invoked in Step 2). Before presenting a more orderly analysis, let us consider the communication complexity of the above system: In Step 1 the prover sends $O(d \log k) \cdot c(n)$ bits, in Step 2 it sends $q \cdot k$ bits, and in Step 3 it sends $O(k/d) \cdot c(n)$ bits. Picking $d = \sqrt{k}$ and $q = O(1)$, we get total communication of $\tilde{O}(k^{1/2}) \cdot c(n) + O(k)$. We comment that the

⁴Indeed, the NP-witness may appear as a prefix of the PCP-oracle. By the *PCP-oracle derived from an NP-witness* we mean the PCP-oracle that is outlined in the description of the PCP system, while recalling that all known PCP systems are described in this matter.

verification time is vastly dominated by Step 3a, whereas the prover's strategy can be implemented in polynomial-time (given the NP-witnesses).

Claim 4.4.2 (soundness of Construction 4.4.1): *Construction 4.4.1 is an interactive proof system for S^k .*

Proof: Turning to the more orderly analysis of the soundness of Construction 4.4.1, we assume without loss of generality that the values provided by the prover satisfy the conditions stated in Steps 2a and 3b. We consider two cases when referring to the values $\tilde{v}_1, \dots, \tilde{v}_{c(n)}$ provided by the prover in Step 1.

The first case is that, with probability at least half, over the choice of the query-sequence (j_1, \dots, j_q) in Step 2, the $\tilde{\pi}_{i,j_{i'}}$'s provided by the prover satisfy $|\{i \in [k] : \tilde{\pi}_{i,j_{i'}} \neq \pi_{i,j_{i'}}\}| > d$ for some $i' \in [q]$, where π_i denotes the PCP-oracle derived for $x_i \in S$ and is defined as the all-zero string in case $x_i \notin S$. In this case, for each query-sequence (j_1, \dots, j_q) that belongs to the majority, with high probability, the set of rows chosen in Step 3 contains a row i such that $\tilde{\pi}_{i,j_{i'}} \neq \pi_{i,j_{i'}}$ for some $i' \in [q]$, and so $\tilde{\pi}_i$ does not satisfy the condition stated in Step 3a, which causes the verifier to reject. We stress that the uniqueness of NP-witnesses implies the uniqueness of the PCP-oracles derived from them, and so $\tilde{\pi}_i \neq \pi_i$ implies that the string extracted from $\tilde{\pi}_i$ (so that $\tilde{\pi}_i$ is derived from this string) is not a valid NP-witness for x_i .

Turning to the complimentary case, where with probability at least half (over the choice of (j_1, \dots, j_q) made in Step 2) it holds that $|\{i \in [k] : \tilde{\pi}_{i,j_{i'}} \neq \pi_{i,j_{i'}}\}| \leq d$ for every $i' \in [q]$, we let $(\pi'_{i,j})_{i \in [k], j \in [c(n)]}$ denote the unique matrix such that for every $j \in [c(n)]$ the string $\pi'_{1,j} \cdots \pi'_{k,j}$ is the unique string that resides in the intersection of the Hamming ball of radius d centered at $\pi_{1,j} \cdots \pi_{k,j}$ and the pre-image of \tilde{v}_j under F . Hence, with probability at least half (over the choice of (j_1, \dots, j_q)), for each $i \in [k]$, it holds that $\tilde{\pi}_{i,j_{i'}} = \pi'_{i,j_{i'}}$ for every $i' \in [q]$, since $\tilde{\pi}_{1,j_{i'}} \cdots \tilde{\pi}_{k,j_{i'}}$ resides in the said Hamming ball and is in the pre-image of $\tilde{v}_{j_{i'}}$ under F .⁵ By the soundness of the PCP verifier (which is invoked in Step 2b), it follows that $x_i \in S$.

Hence, either the verifier rejects with probability at least $1/2$ or all x_i 's are in S . ■

Conclusion. By setting $d = \sqrt{k}$, we have established the claim of the theorem for any $\alpha > 1/2$. To establish the claim for smaller constant values of $\alpha > 0$, we use recursion. Specifically, we set $d = k^\alpha$, and employ Construction 4.4.1, except that in Step 3 we invoke Construction 4.4.1 on the $O(k/d)$ instances that correspond to the selected rows. Hence, we reduce batch verification for k instances to batch verification for $O(k/d)$ instances. However, the claim to be verified is not only that $x_i \in S$ for each selected row i (i.e., $i \in R$), but rather that there exists an NP-witness for x_i such that the PCP-oracle derived from it match the values $\tilde{\pi}_{i,j_1}, \dots, \tilde{\pi}_{i,j_q}$ received in Step 2 (i.e., $\tilde{\pi}_{i,j_{i'}} = y_{i,j_{i'}}$ for every $i' \in [q]$, where $y_{i,j}$ is a variable representing the j^{th} bit of the oracle derived from the NP-witness for x_i).⁶ After $1/\alpha$ recursive calls, we just use the straightforward verification that is used originally in Step 3. ■

⁵Actually, it suffices to establish a weaker statement asserting that, for each $i \in [k]$, it holds that $\Pr_{(j_1, \dots, j_q)}[(\forall i' \in [q]) \tilde{\pi}_{i,j_{i'}} = \pi'_{i,j_{i'}}] \geq 1/2$.

⁶The corresponding NP-claim is that $(x_i, \tilde{\pi}_{i,j_1}, \dots, \tilde{\pi}_{i,j_q}) \in S'$, which holds if there exists $y_i = (y_{i,1}, \dots, y_{i,c(n)})$ such that y_i is the PCP-oracle derived from the NP-witness for $x_i \in S$ and $y_{i,j_{i'}} = \tilde{\pi}_{i,j_{i'}}$ holds for every $i' \in [q]$.

4.3 Batch verification for unambiguous IP

The proof of Theorem 4.4 makes essential use of the hypothesis that each YES-instance has a unique proof, which in the context of \mathcal{NP} means that it has a unique NP-witness. Seeking to extend Theorem 4.4 to sets that only have interactive proof systems, we seek an adequate notion of unique proving strategies, since an interactive strategy is the interactive analogue of a written proof (i.e., an NP-witness). Intuitively, we want to require that, at each round, there is at most one prover message that may lead the verifier to accept. This condition must be phrased while accounting for the fact that the verifier strategy is probabilistic and consequently the phrase “leading the verifier to accept” should be given a probabilistic interpretation. Indeed, we require that, at each round, there is at most one prover message that may lead the verifier to accept with probability at least $1/2$.

Definition 4.5 (unambiguous interactive proof systems): *Let (P, V) be an interactive proof system, where P is a deterministic prover strategy that satisfies the (perfect) completeness condition. The system (P, V) is called **unambiguous** if for every partial communication transcript (of V with P) that ends with a verifier message, there exists at most one prover message such that, in the residual communication, the verifier accepts with probability at least $1/2$. More formally:*

Let $\langle P', V \rangle_j(x)$ denote a random variable representing the distribution of the transcripts of the first j messages in an interaction between P' and V on common input x . Then, assuming that the verifier sends the first message, we require that for every i and every $(\alpha_1, \beta_1, \dots, \alpha_{i-1}, \beta_{i-1}, \alpha_i, \beta_i)$ in the support of $\langle P, V \rangle_{2i}(x)$ and every \tilde{P} , it holds that

$$\Pr \left[(\tilde{P}, V)(x) = 1 \mid \begin{array}{l} \langle \tilde{P}, V \rangle_{2i-1}(x) = (\alpha_1, \beta_1, \dots, \alpha_{i-1}, \beta_{i-1}, \alpha_i) \\ \langle \tilde{P}, V \rangle_{2i}(x) \neq (\alpha_1, \beta_1, \dots, \alpha_{i-1}, \beta_{i-1}, \alpha_i, \beta_i) \end{array} \right] < 1/2, \quad (4.2)$$

where $(\tilde{P}, V)(x)$ denotes the output of V after interacting with \tilde{P} on input x .

That is, Eq. (4.2) asserts that any message other than β_i leads the verifier to accept with probability that is smaller than $1/2$, whereas the message β_i may lead the verifier to accept with probability 1 .⁷

Note that in the special case of NP-proof systems, Definition 4.5 means that there is at most one prover-message (which corresponds to the NP-witness) that makes the (deterministic) verifier accept. Hence, the notion of unambiguous interactive proof systems is a natural extension of the notion of an NP-proof system with unique proofs (or unique NP-witnesses). We also mention that the celebrated interactive proof systems for \mathcal{PSPACE} (of [32, 38]) are unambiguous, since the sum-check protocol is unambiguous.⁸

When seeking batch verification for interactive proof systems, we shall restrict our attention to public-coin systems [4]. The essential feature of public-coin systems that we use is the fact that the next message of the verifier can be selected without looking at the previous prover messages. Furthermore, we use the fact that given the previous messages of the verifier (but not the previous

⁷In particular, if x is a YES-instance (i.e., $\Pr[\langle P, V \rangle_{2i}(x) = 1] = 1$), then the last message in $\langle P, V \rangle_{2i}(x)$ is uniquely determined by the first $2i - 1$ messages.

⁸At the beginning of each round, when one variable in the sum is stripped, there is a unique low-degree univariate polynomial that describes the residual function. If the prover sends any other polynomial, then it will be caught with high probability, no matter how it plays in the subsequent rounds.

messages of the prover), one can efficiently sample the distribution of the next verifier message. Needless to say, these features are trivially satisfied by public-coin interactive proof systems.

Lastly, we need a notion that extends the definition of PCP to the interactive context. Viewing PCP systems as NP-proof systems that support verification based on inspecting few randomly selected bits in the NP-witness, we consider (public-coin) interactive proofs systems in which the final verification is based on inspecting few randomly selected bits in the sequence of prover's messages. That is, these proof systems consists of two stages.

1. In the interaction stage, the verifier generates messages obliviously of the prover's prior messages, which are merely recorded for future use.
2. In the final verification stage, which takes place after the entire interaction is completed, the verifier decides based on inspecting few of the bits sent by the prover (and possibly based on its own coin tosses and input). Equivalently, the verifier final decision is based on all messages sent by the verifier and few probes made into the record of the prover's messages.

A proof system that satisfies the foregoing condition is called a **probabilistically checkable interactive proof** (PCIP). Note that the combination of the unambiguity condition and the probabilistically checkable condition is problematic; actually, a PCIP cannot be unambiguous (in the strict sense of Definition 4.5), since changing a single bit in a message is unlikely to be detected. Still, a natural relaxation of the unambiguity condition is applicable to PCIP: Rather than requiring that the successful proving strategy is unique, we require that all successful prover strategies generate messages that are close to one another (in Hamming distance). This notion is indeed akin to the definition of PCPP (see [8, 14]), where the soundness condition holds only with respect to inputs that are far from the predetermined set.

Observe that unambiguous public-coin interactive proof systems can be transformed into ones that satisfy the PCIP condition, by having the prover encode its messages under an error correcting code, and letting the verifier run a PCPP to verify that the original verifier would have accepted the original messages. That is, the new verifier checks that all messages of the new verifier are valid codewords, and that they encode messages that would have convinced the original verifier.

Having such a PCIP at our disposal, we employ the ideas that underlie Construction 4.4.1 to each round of interaction in the PCIP. Thus, the emulation of a typical round of interaction (of the PCIP) looks as follows, when referring to the input (x_1, \dots, x_k) .

Construction 4.5.1 (emulation of a single round):

1. Let β_i be the (unique) message that the original prover would have sent in the current round regarding the input x_i (when given the history of communication regarding this input).⁹ For each $j \in [c(n)]$, let $\beta_{i,j}$ denote the j^{th} bit of β_i . Then, for every $j \in [c(n)]$, the prover computes $v_j \leftarrow F(\beta_{1,j} \cdots \beta_{k,j})$, and sends $v_1, \dots, v_{c(n)}$ to the verifier. Recall that $F : \{0, 1\}^k \rightarrow \{0, 1\}^{O(d \log k)}$, where d is the commitment parameter.

(Indeed, note the analogy to Step 1 in Construction 4.4.1.)

We shall refer to the foregoing k messages as to the matrix of the current round.

⁹That is, we consider k executions of the PCIP, where the i^{th} execution refers to the input x_i .

2. *The verifier answers with a single random message, which will be used in all k copies of the PCIP emulation.*

(Using the same verifier message in all k copies may increase the soundness error by a factor of k , but actually the same may happen when k independently selected messages are used. In any case, before employing the construction, we should reduce the soundness error of the original PCIP so to compensate for this loss.)

Note that verification steps analogous to Steps 2 and 3 of Construction 4.4.1 are not performed at this point, but rather after all rounds of interaction of the PCIP are completed. Once this happens, the following verification steps take place, where $\ell = O(1)$ denotes the number of rounds in the PCIP.

Construction 4.5.2 (final verification):

1. *The verifier generate a sequence of queries, $j_1, \dots, j_q \in [c(n)]$, for the PCIP verifier, denoted V . Note that V may ask different queries to the messages of the different ℓ rounds, but for sake of simplicity (and at the cost of increasing q by a factor of ℓ), we assume that it asks the same queries in each round.*

As in Step 2 of Construction 4.4.1, the verifier sends j_1, \dots, j_q to the prover, who responds with the values of the corresponding columns. Note that the prover sends q columns per each of the ℓ matrices.

For each of these ℓ matrices, the verifier checks (i) whether the values of these columns match the parity-check values for the relevant matrix (as provided in Step 1 of Construction 4.5.1) and (ii) whether V would have accepted each of the k inputs when given the corresponding answers (as well as its own messages as sent in Step 2 of Construction 4.5.1). Note that the first condition refers to individual matrices, whereas the second condition refers to the sequence of ℓ matrices.

Specifically, let $\tilde{v}_1^{(r)}, \dots, \tilde{v}_{c(n)}^{(r)}$ denote the parity check values sent in round r , and suppose that the prover's current answers are $(\tilde{\pi}_{i,j_{i'}}^{(r)})_{r \in [\ell], i \in [k], i' \in [q]}$, where $(\tilde{\pi}_{i,j}^{(r)})_{i \in [k]}$ is claimed to be the j^{th} column of the matrix of round r . Then, the verifier performs the following two checks:

- (a) *For every $r \in [\ell]$ and $i' \in [q]$, it checks whether $F(\tilde{\pi}_{1,j_{i'}}^{(r)} \cdots \tilde{\pi}_{k,j_{i'}}^{(r)}) = \tilde{v}_{j_{i'}}^{(r)}$.*
- (b) *For each $i \in [k]$, it checks whether V would have accepted x_i when making the queries j_1, \dots, j_q and receiving the answers $\tilde{\pi}_{i,j_1}^{(1)}, \dots, \tilde{\pi}_{i,j_q}^{(1)}, \dots, \tilde{\pi}_{i,j_1}^{(\ell)}, \dots, \tilde{\pi}_{i,j_q}^{(\ell)}$ and when its own messages in the i^{th} copy are as recorded in the execution of Construction 4.5.1.*

(Indeed, the current step is analogous to Step 2 of Construction 4.4.1, and the next step is analogous to Step 3 in that construction.)

2. *The verifier selects uniformly a set of $k' = O(k/d)$ rows, and the prover is required to prove that these rows contain values (in the ℓ matrices) that (i) correspond to the unambiguous transcript that makes the PCIP verifier accept, and (ii) match the values of the columns provided in Step 1 of Construction 4.5.1.*

Unlike in Construction 4.4.1, the unambiguity condition cannot be verified by merely inspecting the transcripts of the relevant rows (in all ℓ matrices), because this condition refers to

possible random interactions that extend all possible prefixes of the transcript, where the crucial point is that in such random interactions the future messages of the verifier are not known. Letting $(\alpha_1, \beta_1, \dots, \alpha_\ell, \beta_\ell)$ denote a transcript that corresponds to a revealed row (in the ℓ matrices), we need to verify a condition analogous to Eq. (4.2), for each $i \in [\ell]$. This is done by invoking the original PCIP for ℓ times, such that in the i^{th} invocation we continue the execution at the end of the i^{th} round as if the first $2i$ messages were $(\alpha_1, \beta_1, \dots, \alpha_i, \beta_i)$. Note that in the i^{th} invocation (which refers to an i -round long prefix), $\ell - i$ rounds of (additional) interaction are performed, followed by the final verification step of the original PCIP. All these $k' \cdot \ell$ interactive processes are run in parallel.¹⁰

Combining Constructions 4.5.1 and 4.5.2, we obtain a PCIP of $O(\ell)$ rounds of query complexity $\text{poly}(q) \cdot k$, with communication complexity $O(c(n) \cdot d \log k) + \text{poly}(q) \cdot k + \frac{k}{d} \cdot c(n)$, where first term is due to Construction 4.5.1, the other terms (as well as the query complexity) are due to Construction 4.5.2.¹¹ Now, recalling that we wish to obtain communication complexity $c(n) \cdot k^\alpha$, for any $\alpha > 0$, as well as low query complexity, we face a problem since the first goal forces using $d = k^\alpha$ (just as in the proof of Theorem 4.4). The solution is to replace the sending of the k' rows in Step 2 of Construction 4.5.2 by having the prover employ the entire construction recursively so to allow for the verification of the corresponding conditions. We stress that this recursion is more complex than in the proof of Theorem 4.4, since we need to verify an unambiguity claim regarding a PCIP whose (constant) number of rounds varies throughout the recursive calls.

Another problem that we face is that the resulting query complexity is $\text{poly}(q) \cdot k$, whereas Theorem 4.2 asserts polylogarithmic query complexity. The solution is to use a *query reduction step*; that is, after deriving a PCIP of $O(\ell)$ rounds of query complexity $\text{poly}(q) \cdot k$ and communication complexity $O(c(n) \cdot k^\alpha) + \text{poly}(q) \cdot k$, we reduce the query complexity to $q = \text{poly}(\log n)$. Indeed, query reduction is a standard PCP technique, but one should verify that it preserves the unambiguity condition.

Conclusion: Deriving Theorem 4.2. The foregoing outline completes our high-level description of the proof of Theorem 4.2, and it implicitly specifies the type \mathcal{T} of interactive proof systems to which this theorem refers. To spell it out, type \mathcal{T} consists of public-coin unambiguous PCIPs (of proximity).

¹⁰The most straightforward interpretation of the above is that the j^{th} interaction that refers to an i -round prefix takes place in round j of the parallel execution, regardless of i . It is also natural to let this interaction take place in round $i + j$. Actually, the mapping of additional interactions to parallel rounds is immaterial as long as it is monotone and injective.

¹¹Specifically, the $\text{poly}(q) \cdot k$ term (resp., the $\frac{k}{d} \cdot c(n)$ term) is due to Step 1a (resp., Step 2) of Construction 4.5.2.

Appendix: Defining Interactive Proofs and Arguments

In this appendix we review the basic definition of interactive proof systems as well as the definition of argument systems. The latter definition is reviewed with the context of a general discussion regarding the power of the prover. (The text was adapted from [19, Sec. 9.1].)

A.1 The basic definition of interactive proofs

Loosely speaking, an interactive proof is a “game” between a computationally bounded verifier and a computationally unbounded prover whose goal is to convince the verifier of the validity of some assertion. Specifically, the verifier employs a probabilistic polynomial-time strategy (whereas no computational restrictions apply to the prover’s strategy). It is required that if the assertion holds then the verifier always accepts (i.e., when interacting with an appropriate prover strategy). On the other hand, if the assertion is false then the verifier must reject with probability at least $\frac{1}{2}$, no matter what strategy is being employed by the prover. (The error probability can be reduced by running such a proof system several times.)

We formalize the interaction between parties by referring to the *strategies* that the parties employ.¹² A **strategy** for a party is a *function mapping the party’s view of the interaction so far to a description of this party’s next move*; that is, such a strategy describes (or rather prescribes) the *party’s next move* (i.e., its next message or its final decision) *as a function of the common input* (i.e., the aforementioned assertion), *the party’s internal coin tosses, and all messages it has received so far*. Note that this formulation presumes (implicitly) that each party records the outcomes of its past coin tosses as well as all the messages it has received, and determines its moves based on these. Thus, an interaction between two parties, employing strategies A and B respectively, is determined by the common input, denoted x , and the randomness of both parties, denoted r_A and r_B . Assuming that A takes the first move (and B takes the last move), the corresponding (t -round) **interaction transcript** (on common input x and randomness r_A and r_B) is $\alpha_1, \beta_1, \dots, \alpha_t, \beta_t$, where $\alpha_i = A(x, r_A, \beta_1, \dots, \beta_{i-1})$ and $\beta_i = B(x, r_B, \alpha_1, \dots, \alpha_i)$. The corresponding final decision of A is defined as $A(x, r_A, \beta_1, \dots, \beta_t)$.

We say that a party employs a **probabilistic polynomial-time strategy** if its next move can be computed in a number of steps that is *polynomial in the length of the common input*. In particular, this means that, on input common input x , the strategy may only consider a polynomial in $|x|$

¹²An alternative formulation refers to the interactive machines that capture the behavior of each of the parties (see, e.g., [18, Sec. 4.2.1.1]). Such an interactive machine invokes the corresponding strategy, while handling the communication with the other party and keeping a record of all messages received so far.

many messages, which are each of $\text{poly}(|x|)$ length.¹³ Intuitively, if the other party exceeds an a priori (polynomial in $|x|$) bound on the total length of the messages that it is allowed to send, then the execution is suspended. Thus, referring to the aforementioned strategies, we say that A is a probabilistic polynomial-time strategy if, for every i and $r_A, \beta_1, \dots, \beta_i$, the value of $A(x, r_A, \beta_1, \dots, \beta_i)$ can be computed in time polynomial in $|x|$. Again, in proper use, it must hold that $|r_A|, t$ and the $|\beta_i|$'s are all polynomial in $|x|$.

Definition A.1 (Interactive Proof systems – IP):¹⁴ *An interactive proof system for a set S is a two-party game, between a verifier executing a probabilistic polynomial-time strategy, denoted V , and a prover that executes a (computationally unbounded) strategy, denoted P , satisfying the following two conditions:*

- **Completeness:** *For every $x \in S$, the verifier V always accepts after interacting with the prover P on common input x .*
- **Soundness:** *For every $x \notin S$ and every strategy P^* , the verifier V rejects with probability at least $\frac{1}{2}$ after interacting with P^* on common input x .*

We denote by \mathcal{IP} the class of sets having interactive proof systems.

The error probability (in the soundness condition) can be reduced by successive applications of the proof system. (This is easy to see in the case of sequential repetitions, but holds also for parallel repetitions; see [19, Exer. 9.1].) In particular, repeating the proving process for k times, reduces the probability that the verifier is fooled (i.e., accepts a false assertion) to 2^{-k} , and we can afford doing so for any $k = \text{poly}(|x|)$.

A.2 On computationally bounded provers: an overview

Recall that our definition of interactive proofs (i.e., Definition A.1) makes no reference to the computational abilities of the potential prover. This fact has two conflicting consequences:

1. The completeness condition does not provide any upper bound on the complexity of the corresponding proving strategy (which convinces the verifier to accept valid assertions).
2. The soundness condition guarantees that, regardless of the computational effort spend by a cheating prover, the verifier cannot be fooled to accept invalid assertions (with probability exceeding the soundness error).

Note that providing an upper-bound on the complexity of the (prescribed) prover strategy P of a specific interactive proof system (P, V) only strengthens the claim that (P, V) is a proof system for the corresponding set (of valid assertions). We stress that the prescribed prover strategy is referred to only in the completeness condition (and is irrelevant to the soundness condition). On the other hand, relaxing the definition of interactive proofs such that soundness holds only for a specific class of cheating prover strategies (rather than for all cheating prover strategies) weakens the corresponding claim. In this advanced section we consider both possibilities.

¹³Needless to say, the number of internal coin tosses fed to a polynomial-time strategy must also be bounded by a polynomial in the length of x .

¹⁴We follow the convention of specifying strategies for both the verifier and the prover. An alternative presentation only specifies the verifier's strategy, while rephrasing the completeness condition as follows: *There exists a prover strategy P such that, for every $x \in S$, the verifier V always accepts after interacting with P on common input x .*

A.2.1 How powerful should the prover be?

Suppose that a set S is in \mathcal{IP} . This means that there exists a verifier V that can be convinced to accept any input in S but cannot be fooled to accept any input not in S (except with small probability). One may ask how powerful should a prover be such that it can convince the verifier V to accept any input in S . It is known (see [19, Prop. 9.5]) that an optimal prover strategy (for convincing any fixed verifier V) can be implemented in polynomial-space, and that we cannot expect any better for a generic set in $\mathcal{PSPACE} = \mathcal{IP}$ (because the emulation of the interaction of V with any optimal prover strategy yields a decision procedure for the set). Still, we may seek better upper-bounds on the complexity of some prover strategy that convinces a *specific* verifier, which in turn corresponds to a specific set S . More interestingly, considering all possible verifiers that give rise to interactive proof systems for S , we wish to upper-bound the computational power that suffices for convincing any of these verifiers (to accept any input in S).

We stress that, unlike the case of computationally-sound proof systems (see Section A.2.2), we do not restrict the power of the prover in the soundness condition, but rather consider the minimum complexity of provers meeting the completeness condition. Specifically, we are interested in *relatively efficient* provers that meet the completeness condition. The term “relatively efficient prover” has been given three different interpretations, which are briefly surveyed next.

1. A prover is considered *relatively efficient* if, when given an auxiliary input (in addition to the common input in S), it works in (probabilistic) polynomial-time. Specifically, in case $S \in \mathcal{NP}$, the auxiliary input maybe an NP-proof that the common input is in the set. Still, even in this case the interactive proof need not consist of the prover sending the auxiliary input to the verifier; for example, an alternative procedure may allow the prover to be zero-knowledge. This interpretation is adequate and in fact crucial for applications in which such an auxiliary input is available to the otherwise polynomial-time parties. Typically, such auxiliary input is available in cryptographic applications in which parties wish to prove in (zero-knowledge) that they have correctly conducted some computation. In these cases, the NP-proof is just the transcript of the computation by which the claimed result has been generated, and thus the auxiliary input is available to the party that plays the role of the prover.
2. A prover is considered *relatively efficient* if it can be implemented by a probabilistic polynomial-time oracle machine with oracle access to the set S itself. Note that the prover in the standard proof system for Graph Non-Isomorphism has this property (see [19, Const. 9.3]). This interpretation generalizes the notion of self-reducibility of NP-proof systems. Recall that by self-reducibility of an NP-set (or rather of the corresponding NP-proof system) we mean that the search problem of finding an NP-witness is polynomial-time reducible to deciding membership in the set (cf. [19, Def. 2.14]). Here we require that implementing the prover strategy (in the relevant interactive proof) be polynomial-time reducible to deciding membership in the set.
3. A prover is considered *relatively efficient* if it can be implemented by a probabilistic machine that runs in time that is polynomial in the deterministic complexity of the set. This interpretation relates the time-complexity of convincing a “lazy person” (i.e., a verifier) to the time-complexity of determining the truth (i.e., deciding membership in the set).

Hence, in contrast to the first interpretation, which is adequate in settings where assertions are generated along with their NP-proofs, the current interpretation is adequate in settings

in which the prover is given only the assertion and has to find a proof to it by itself (before trying to convince a lazy verifier of its validity).

Indeed, this interpretation generalizes one of the two efficiency conditions that underly the definition of doubly-efficient interactive proof systems. Specifically, for set in \mathcal{P} (or even in \mathcal{BPP}), such a *relatively efficient* prover is required to run in polynomial-time.

A.2.2 Computational-soundness

Relaxing the soundness condition such that it only refers to relatively-efficient ways of trying to fool the verifier (rather than to all possible ways) yields a fundamentally different notion of a proof system. The verifier’s verdict in such a system is not absolutely sound, but is rather sound *provided that the potential cheating prover does not exceed the presumed complexity limits*. As in Section A.2.1, the notion of “relative efficiency” can be given different interpretations, the most popular one being that the cheating prover strategy can be implemented by a (non-uniform) family of polynomial-size circuits. The latter interpretation coincides with the first interpretation used in Section A.2.1 (i.e., a probabilistic polynomial-time strategy that is given an auxiliary input (of polynomial length)). Specifically, in this case, the soundness condition is replaced by the following **computational soundness** condition that asserts that it is infeasible to fool the verifier into accepting false statements. Formally:

For every prover strategy that is implementable by a family of polynomial-size circuits $\{C_n\}$, and every sufficiently long $x \in \{0, 1\}^ \setminus S$, the probability that V accepts x when interacting with $C_{|x|}$ is less than $1/2$.*

As in case of standard soundness, the computational-soundness error can be reduced by repetitions. We warn, however, that unlike in the case of standard soundness (where both sequential and parallel repetitions will do), the computational-soundness error cannot *always* be reduced by parallel repetitions (see [7]). Hence, it is often preferred to define computational soundness as having negligible error (i.e., error that decreases faster than the reciprocal of any polynomial function).¹⁵

It is common and natural to consider proof systems in which the prover strategies considered both in the completeness and soundness conditions satisfy the same notion of relative efficiency. Protocols that satisfy these conditions with respect to the foregoing interpretation are called **arguments**. We mention that argument systems may be more efficient (e.g., in terms of their communication complexity) than interactive proof systems (cf. [30] versus [24]).

¹⁵In such a case, the definition asserts that *for every prover strategy that is implementable by a family of polynomial-size circuits $\{C_n\}$, every positive polynomial p , and every sufficiently long $x \in \{0, 1\}^* \setminus S$, the probability that V accepts x when interacting with $C_{|x|}$ is less than $1/p(|x|)$.*

Acknowledgements

I am grateful to Guy Rothblum for many useful discussions regarding the subject. In particular, the expositions in Chapters 3 and 4 are based on his oral presentations of [25] and [36], respectively, and Chapter 2 reproduces text from our joint work [23].

I also wish to thank Ran Canetti and Ron Rothblum for discussion of various related topics. In particular, Theorem 1.6 was inspired by a discussion with Ran.

Bibliography

- [1] Amir Abboud, Arturs Backurs, and Virginia Vassilevska Williams. If the Current Clique Algorithms are Optimal, So is Valiant’s Parser. In *46th IEEE Symposium on Foundations of Computer Science*, pages 98–117, 2015.
- [2] Amir Abboud, Kevin Lewi, and Ryan Williams. Losing Weight by Gaining Edges. In *22nd ESA*, pages 1–12, 2014
- [3] Noga Alon, Oded Goldreich, Joahn Hastad, and Rene Peralta. Simple Constructions of Almost k -wise Independent Random Variables. *Journal of Random Structures and Algorithms*, Vol. 3, No. 3, pages 289–304, 1992. Preliminary version in *31st FOCS*, 1990.
- [4] Laszlo Babai. Trading Group Theory for Randomness. In *17th ACM Symposium on the Theory of Computing*, pages 421–429, 1985.
- [5] Laszlo Babai, Lance Fortnow, Leonid Levin, and Mario Szegedy. Checking Computations in Polylogarithmic Time. In *23rd ACM Symposium on the Theory of Computing*, pages 21–31, 1991.
- [6] Marshall Ball, Alon Rosen, Manuel Sabin, and Prashant N. Vasudevan. Average-Case Fine-Grained Hardness. In *48th ACM Symposium on the Theory of Computing*, pages 483–496, 2017.
- [7] Mihir Bellare, Russell Impagliazzo, and Moni Naor. Does Parallel Repetition Lower the Error in Computationally Sound Protocols? In *38th IEEE Symposium on Foundations of Computer Science*, pages 374–383, 1997.
- [8] Eli Ben-Sasson, Oded Goldreich, Prahladh Harsha, Madhu Sudan, and Salil Vadhan. Robust PCPs of Proximity, Shorter PCPs, and Applications to Coding. *SIAM Journal on Computing*, Vol. 36 (4), pages 889–974, 2006. Extended abstract in *36th STOC*, 2004.
- [9] Andreas Björklund and Petteri Kaski. How Proofs are Prepared at Camelot. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing*, pages 391–400, 2016.
- [10] Allan Borodin. On Relating Time and Space to Size and Depth. *SIAM Journal on Computing*, Vol. 6 (4), pages 733–744, 1977.
- [11] Zvika Brakerski, Justin Holmgren, and Yael Tauman Kalai. Non-interactive delegation and batch NP verification from standard computational assumptions. In *48th ACM Symposium on the Theory of Computing*, pages 474–482, 2017.

- [12] Marco L. Carmosino, Jiawei Gao, Russell Impagliazzo, Ivan Mihajlin, Ramamohan Paturi, and Stefan Schneider. Nondeterministic Extensions of the Strong Exponential Time Hypothesis and Consequences for Non-reducibility. In *2016 ACM Conference on Innovations in Theoretical Computer Science*, pages 261–270, 2016.
- [13] Jianer Chen, Benny Chor, Mike Fellows, Xiuzhen Huang, David W. Juedes, Iyad A. Kanj, Ge Xia. Tight lower bounds for certain parameterized NP-hard problems. *Inf. Comput.*, Vol. 201 (2), pages 216–231, 2005.
- [14] Irit Dinur and Omer Reingold. Assignment-testers: Towards a combinatorial proof of the PCP-Theorem. *SIAM Journal on Computing*, Vol. 36 (4), pages 975–1024, 2006. Extended abstract in *45th FOCS*, 2004.
- [15] Rodney G. Downey and Michael R. Fellows. Fixed-parameter tractability and completeness II: On completeness for $W[1]$. *Theoretical Computer Science A*, Vol. 141 (12), pages 109–131, 1995.
- [16] Rodney G. Downey and Michael R. Fellows. *Parameterized Complexity*. Springer-Verlag Monographs in Computer Science, 1999.
- [17] Uriel Feige, Shafi Goldwasser, Laszlo Lovász, Shmuel Safra, and Mario Szegedy. Approximating Clique is almost NP-complete. *Journal of the ACM*, Vol. 43, pages 268–292, 1996. Preliminary version in *32nd FOCS*, 1991.
- [18] Oded Goldreich. *Foundations of Cryptography: Basic Tools*. Cambridge University Press, 2001.
- [19] Oded Goldreich. *Computational Complexity: A Conceptual Perspective*. Cambridge University Press, 2008.
- [20] Oded Goldreich. *Introduction to Property Testing*. Cambridge University Press, to appear in 2017.
- [21] Oded Goldreich, Silvio Micali, and Avi Wigderson. Proofs that Yield Nothing but their Validity or All Languages in NP Have Zero-Knowledge Proof Systems. *Journal of the ACM*, Vol. 38, No. 3, pages 691–729, 1991. Preliminary version in *27th FOCS*, 1986.
- [22] Oded Goldreich and Or Meir. Input-Oblivious Proof Systems and a Uniform Complexity Perspective on P/poly. *TOCT*, Vol. 7 (4), pages 16:1–16:13, 2015.
- [23] Oded Goldreich and Guy N. Rothblum. Simple doubly-efficient interactive proof systems for locally-characterizable sets. *ECCC*, TR17-018, 2017.
- [24] Oded Goldreich, Salil Vadhan, and Avi Wigderson. On interactive proofs with a laconic provers. *Computational Complexity*, Vol. 11, pages 1–53, 2002.
- [25] Shafi Goldwasser, Yael Tauman Kalai, Guy N. Rothblum. Delegating Computation: Interactive Proofs for Muggles. *Journal of the ACM*, Vol. 62(4), Art. 27:1-27:64, 2015. Extended abstract in *40th STOC*, pages 113–122, 2008.

- [26] Shafi Goldwasser, Silvio Micali and Charles Rackoff. The Knowledge Complexity of Interactive Proof Systems. *SIAM Journal on Computing*, Vol. 18, pages 186–208, 1989. Preliminary version in *17th STOC*, 1985. Earlier versions date to 1982.
- [27] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A Pseudorandom Generator from any One-way Function. *SIAM Journal on Computing*, Volume 28, Number 4, pages 1364–1396, 1999.
- [28] Yael Tauman Kalai, Ran Raz, and Ron D. Rothblum. Delegation for bounded space. In *45th ACM Symposium on the Theory of Computing*, pages 565–574, 2013.
- [29] Yael Tauman Kalai, Ran Raz, and Ron D. Rothblum. How to delegate computations: the power of no-signaling proofs. In *46th ACM Symposium on the Theory of Computing*, pages 485–494, 2014.
- [30] Joe Kilian. A Note on Efficient Zero-Knowledge Proofs and Arguments. In *24th ACM Symposium on the Theory of Computing*, pages 723–732, 1992.
- [31] Eyal Kushilevitz and Rafail Ostrovsky. Replication is NOT Needed: SINGLE Database, Computationally-Private Information Retrieval. In *38th IEEE Symposium on Foundations of Computer Science*, pages 364–373, 1977.
- [32] Carsten Lund, Lance Fortnow, Howard Karloff, and Noam Nisan. Algebraic methods for interactive proof systems. *Journal of the ACM*, Vol. 39, No. 4, pages 859–868, 1992. Extended abstract in *31st FOCS*, 1990.
- [33] Joseph Naor and Moni Naor. Small-bias Probability Spaces: Efficient Constructions and Applications. *SIAM Journal on Computing*, Vol 22, pages 838–856, 1993. Preliminary version in *22nd STOC*, 1990.
- [34] Mihai Patrascu. Towards polynomial lower bounds for dynamic problems. In *42nd ACM Symposium on the Theory of Computing*, pages 603–610, 2010.
- [35] Mihai Patrascu and Ryan Williams. On the Possibility of Faster SAT Algorithms. In *21st SODA*, pages 1065–1075, 2010.
- [36] Omer Reingold, Guy N. Rothblum, Ron D. Rothblum. Constant-round interactive proofs for delegating computation. In *48th ACM Symposium on the Theory of Computing*, pages 49–62, 2016.
- [37] Adi Shamir. $IP = PSPACE$. *Journal of the ACM*, Vol. 39, No. 4, pages 869–877, 1992. Preliminary version in *31st FOCS*, 1990.
- [38] Madhu Sudan. Invariances in Property Testing. In *Property Testing: Current Research and Surveys*. Springer, Lecture Notes in Computer Science (Vol. 6390), pages 211–227, 2010.
- [39] Justin Thaler. Semi-Streaming Algorithms for Annotated Graph Streams. In *43rd International Colloquium on Automata, Languages, and Programming*, pages 59:1–59:14, 2016.
- [40] Seinosuke Toda. PP is as hard as the polynomial-time hierarchy. *SIAM Journal on Computing*, Vol. 20 (5), pages 865–877, 1991.

- [41] Virginia Vassilevska Williams. Hardness of Easy Problems: Basing Hardness on Popular Conjectures such as the Strong Exponential Time Hypothesis. In *10th International Symposium on Parameterized and Exact Computation*, pages 17–29, 2015.
- [42] Ryan Williams. Strong ETH Breaks With Merlin and Arthur: Short Non-Interactive Proofs of Batch Evaluation. In *31st Conference on Computational Complexity*, pages 2:1–2:17, 2016.