

Scalable, transparent, and post-quantum secure computational integrity

Eli Ben-Sasson* Iddo Bentov[†] Yinon Horesh* Michael Riabzev*

March 6, 2018

Abstract

Human dignity demands that personal information, like medical and forensic data, be hidden from the public. But veils of secrecy designed to preserve privacy may also be abused to cover up lies and deceit by institutions entrusted with Data, unjustly harming citizens and eroding trust in central institutions.

Zero knowledge (ZK) proof systems are an ingenious cryptographic solution to this tension between the ideals of personal *privacy* and institutional *integrity*, enforcing the latter in a way that does not compromise the former. Public trust demands *transparency* from ZK systems, meaning they be set up with no reliance on any trusted party, and have no trapdoors that could be exploited by powerful parties to bear false witness. For ZK systems to be used with Big Data, it is imperative that the public verification process scale sublinearly in data size. Transparent ZK proofs that can be verified *exponentially* faster than data size were first described in the 1990s but early constructions were impractical, and no ZK system realized thus far in code (including that used by crypto-currencies like ZcashTM) has achieved *both* transparency and exponential verification speedup, simultaneously, for general computations.

Here we report the first realization of a transparent ZK system (ZK-STARK) in which verification scales exponentially faster than database size, and moreover, this exponential speedup in verification is observed concretely for meaningful and sequential computations, described next. Our system uses several recent advances on interactive oracle proofs (IOP), such as a “fast” (linear time) IOP system for error correcting codes.

Our proof-of-concept system allows the Police to prove to the public that the DNA profile of a Presidential Candidate does not appear in the forensic DNA profile database maintained by the Police. The proof, which is generated by the Police, relies on no external trusted party, and reveals no further information about the contents of the database, nor about the candidate’s profile. In particular, no DNA information is disclosed to any party outside the Police. The proof is shorter than the size of the DNA database, and verified faster than the time needed to examine that database naïvely.

*Technion — Israel Institute of Technology, Haifa, Israel; supported by the Israel Science Foundation (grant # 1501/14), the US–Israel Binational Science Foundation, and the European Union’s Horizon 2020 Research And Innovation Programme under grant agreement no. 693423.

[†]Cornell University, Ithaca, NY, USA.

Contents

1	Introduction	4
1.1	Main contribution	6
1.2	Discussion — Applications to decentralized societal functions	8
1.3	Comparison to other realized universal ZK systems	9
1.3.1	Theoretical discussion	9
1.3.2	Concrete performance	11
2	Methods	14
2.1	Fast Reed-Solomon Interactive Oracle Proof of Proximity (FRI3rd)	14
2.2	Arithmetization I — Algebraic Intermediate Representation (AIR)	14
2.3	Arithmetization II — Algebraic Linking Interactive Oracle Proof (ALI)	15
2.4	Low degree extension and composition degree	17
2.5	Minimizing Authentication Path Complexity (APC) and Communication Complexity (CC)	17
2.6	Organization of the remaining sections	18
3	On STIKs and STARKs — formal definitions and prior constructions	19
3.1	Scalable Transparent IOP of Knowledge (STIK)	19
3.2	Main Theorems	21
3.3	Scalable Transparent ARGument of Knowledge (STARK) as a realization of STIK	22
3.4	Prior STIK and STARK constructions	23
A	Measurements of the ZK-STARK for the DNA profile match	31
A.0.1	Prover	31
A.0.2	Verifier	32
B	From AIR to ZK-STARK	33
B.1	Preliminaries and notation	34
B.2	Algebraic Intermediate Representation (AIR)	35
B.3	Algebraic placement and routing (APR)	39
B.4	APR reduction	40
B.5	Algebraic linking IOP (ALI)	42
B.5.1	On the soundness error of the ALI protocol	44
B.6	Fast Reed-Solomon (RS) IOP of Proximity (IOPP) (FRI)	45
B.7	Proof of main theorems	45
B.7.1	Proof of Main Theorem 3.4	46
B.7.2	Proof of Main Lemma B.6	46
B.7.3	Proof of Lemma B.7	48
B.7.4	Proof of ZK-STIK Theorem 3.5	50
B.8	Realization considerations	51
C	Algebraic placement and routing (APR) reduction	52
C.1	The APR reduction for space bounded computation	52
C.1.1	Common definitions	52
C.1.2	Instance reduction	52

C.1.3	Witness reduction	53
C.2	Proof of Theorem B.12 for space bounded computation	53
C.2.1	Proof of completeness (Item 1)	53
C.2.2	Proof of soundness (Item 2)	54
C.2.3	Knowledge extraction (Item 3)	55
C.2.4	Instance properties (Item 4h)	55
C.3	The APR reduction for general computation	56
C.3.1	Common definitions	56
C.3.2	Instance reduction	57
C.3.3	Witness reduction	59
C.4	Proof of Theorem B.13 for general computation	59
C.4.1	Proof of completeness (Item 1)	59
C.4.2	Proof of soundness (Item 2)	61
C.4.3	Knowledge extraction (Item 3)	62
C.4.4	Instance properties (Item 4h)	62
D	Algebraic linking IOP (ALI)	63
D.1	The Algebraic Linking IOP (ALI) protocol	63
D.2	Proof of the ALI reduction Theorem B.15	64
D.2.1	Completeness — Part 2	64
D.2.2	Soundness — Part 3	64
D.2.3	Knowledge Extraction — Part 4	67
D.2.4	Perfect Zero-Knowledge — Part 5	68
D.2.5	Arithmetic complexity — Part 6	70
D.3	Conjectured soundness	70
E	An algebraic intermediate representation of the DNA profile match	71
E.1	Algebraic description of the Rijndael cipher	72
E.2	Implementation technique of the Rijndael cipher	72
E.3	State machine of the Rijndael cipher	73
E.4	From encryption to hash function: Davies-Meyer	74
E.5	DNA profile match (DPM)	76
F	The AIR of the SHA2 hash function	80
G	Affine rearrangeable networks	80
G.1	Combinatorial representation of back-to-back De Bruijn routing	80
G.2	Affine embedding of back-to-back De Bruijn routing	81

1 Introduction

Scalable verification of computational integrity over confidential datasets The problem addressed here is best illustrated by a hypothetical example: Suppose the Police (P), that is in charge of the national forensic DNA profile database (D), claims that the DNA profile (p) of a soon-to-be-appointed and alleged-to-be-corrupt Presidential Candidate, does not appear in D. Can cryptographic protocols convince the doubtful public to believe this claim, without compromising D or p , without relying on any external trusted party (e.g., the Chief Justice), and with “reasonable” computational resources?

The *DNA profile match (DPM)* example is a special case of a more general problem. A party (P) executing a computation (C) on a dataset (D) may have incentive to misreport the correct output ($C(D)$), raising the problem of *computational integrity (CI)*¹ — ensuring that P indeed reports $C(D)$ rather than an output more favorable to P. When the dataset D is public, any party (V) interested in *verifying* CI can naïvely re-execute C on D and compare its output to that reported by P, as a customer might inspect a restaurant bill, or as a new Bitcoin node will verify its blockchain [86]. This naïve solution does not *scale* because the time spent by the verifier (T_V) is as large as the time required to execute the program (T_C) and V must read the full dataset D. Commitment schemes based on cryptographic hash functions [33] are commonly used to compute a short immutable “fingerprint” cm_t for the state at time t of a large dataset D_t [33]. Typically cm_t is negligible in length² compared to D_t , and may be easily posted on a block-chain to serve as a public notice³. Thus, the CI solution we seek should have *scalable verification*, one in which verification time and communication complexity scale roughly like $\log T_C$ and $|cm_t|$ (the bit-length of cm_t), rather than like T_C and $|D_t|$; at the very least verification time/communication should be strictly less than T_C and $|D_t|$.

When the dataset D contains confidential data, the naïve solution can no longer be implemented and the party P in charge of D may conceal violations of computational integrity under the veil of secrecy. Prevailing methods for enforcing CI over confidential data rely on a “trusted party”, like an auditor or accountant to naïvely verify the computation on behalf of the public. This solution still offers no scaling, much like when the data is public. Worse still, it requires the public to trust a third party, which creates a potential single point of failure in the protocol, as this third party — to the extent it can be agreed upon — can be breached, bribed, or coerced by malicious parties.

Zero knowledge (ZK) proof and argument systems are automated protocols that replace human auditors as a means of guaranteeing computational integrity over confidential data for any efficient computation⁴, eliminating corruptibility and reducing costs [59]. A ZK system S for a computation C is a pair of randomized algorithms, $S = (P, V)$; the prover P is the algorithm used to prove computational integrity and the verifier V checks such proofs. The completeness and soundness of S imply that P can efficiently prove all truisms but will fail to convince V of any falsities (with all but negligible probability). The very first theoretical constructions of ZK systems with scalable verifiers for general computations⁵, discussed in the early 1990s, were based on *Probabilistically Checkable Proofs (PCP)*. (See Section 1.3 for recent alternative ZK constructions.) The celebrated PCP Theorem [7, 6, 3, 2] offered a surprising trade-off between the running time spent by the prover constructing the proof (T_P), and the running time consumed by the verifier

¹This problem is also known as *delegation of computation* [58], *certified computation* [41] and *verifiable computation* [52].

²Commonly, cm_t is the SHA2 hash of D_t which is 256 bits long for any dataset length.

³A recent report by the World Economic Forum mentions several use cases, among them monitoring blood diamonds and curbing human trafficking [68].

⁴In the interactive oracle proof model that we consider, as in the model of multi-prover interactive proofs, ZK proof systems exist for any language in nondeterministic exponential time (NEXP) [11, 15].

⁵Special cases for ZK, like proving membership/non-membership in a hidden-and-committed set — the “ZK-set” problem — are efficiently solved by other cryptographic means [84]

checking it (T_V); this trade-off means proving time increases *polynomially* compared to naïve computation time ($T_P = T_C^{O(1)}$) whereas verification time decreases *exponentially* with respect to it ($T_V = \log^{O(1)} T_C$).

A ZK system based on the PCP Theorem (ZK-PCP) [74, 85, 49, 75, 71] has three additional advantages that are essential for ongoing public trust in computational integrity. First, the assumptions on which the security of these constructions is founded — the existence of collision-resistant hash functions [74] for interactive solutions, and common access to a random function⁶ (the “random oracle model” [50]) for non-interactive ones [85] — are not known to be susceptible to attacks by large-scale quantum computers; we call such solutions *post-quantum secure*. The anticipated increase in scale of quantum computers [43] and the call for post-quantum cryptographic protocols, e.g., by the USA National Institute of Standards and Technology (NIST) [37], highlight the importance of a post-quantum secure ZK solution.

Second, ZK-PCPs are *proof of knowledge* (POK) systems, or, when realized as described above, *argument of knowledge* (ARK) systems [33, 9]. Informally, in the context of the DPM example, a ZK-ARK is a proof that convinces the public that the Police has used “the true” dataset D_t and Presidential Candidate DNA profile p whose commitments were previously announced (see Definition 3.3).

Third, and most important, ZK-PCPs are *transparent* (or “public randomness”⁷), which means that the randomness⁸ used by the verifier is *public*; in particular, setting up a ZK-PCP requires no external trusted setup phase, in contrast to newer ZK solutions, including the one used by the ZcashTM cryptocurrency (see Section 1.3). Transparency is essential for ongoing public trust because it severely limits the ability of even the most powerful of parties P to abuse the system, and thus transparent systems are ones which the public may reliably trust as long as there exists something unpredictable in the observable universe.

Summarizing, ZK-PCPs are an excellent method for ensuring public trust in CI over confidential data, and possess six core virtues: (i) *transparency*, (ii) *universality* — apply to any efficient computation C , even if it requires auxiliary (and possibly confidential) input like D_t above, (iii) *confidentiality* (ZK) — do not compromise auxiliary inputs like D_t , (iv) *post-quantum security*, (v) *proof/argument of knowledge* and (vi) *scalable verification*. Although ZK-PCPs have been known since the mid-1990’s, none have been realized in code thus far because, in the words of a recent survey [110], “*the proofs arising from the PCP theorem (despite asymptotic improvements) were so long and complicated that it would have taken thousands of years to generate and check them, and would have needed more storage bits than there are atoms in the universe.*” Consequently, recent realization efforts of ZK systems for general computations (surveyed in Section 1.3) focused on alternative techniques that do not achieve all of (i)–(vi), though some are extremely efficient in practice for concrete circuit sizes and for amortized computations.

Interactive Oracle Proofs (IOP) with scalable proofs To improve prover scalability without sacrificing properties (i)–(vi), a new model was recently suggested [22, 94], called an *interactive oracle proof (IOP)*⁹, a common generalization of the IP, PCP, and interactive PCP (IPCP) models [72]. As in the PCP setting, the IOP verifier need not read prover messages in entirety but rather may query them at random locations; as in the IP setting, prover and verifier interact over several rounds. As was the case for ZK-PCPs, a ZK-IOP system can be converted into an interactive ARK assuming a family of collision-resistant hash functions, and can be turned into a non-interactive argument in the random oracle model [22], which is typically realized using a standard hash function. As a strict generalization of IP/PCP/IPCP, the IOP model offers several

⁶Even though the random oracle model, per se, is unattainable, its use is prevalent in cryptography and the theoretical justification for it discussed, e.g., in [10] and following works.

⁷Transparent systems are also known as *Arthur-Merlin* protocols [4].

⁸Randomness is necessary for ZK proof systems for non-trivial computations [57, Section 3.2].

⁹Reingold et al. [94] use the name “Probabilistically Checkable Interactive Proofs” (PCIP).

advantages. Most relevant to this work is the improved *prover scalability* of IOPs, described below; this advantage holds both asymptotically — as input size $n \rightarrow \infty$ (cf. [15, 16]) — and for concrete input lengths that arise in practice. Based on this efficiency, a proof-of-concept implementation of an IOP, codenamed SCI, was recently reported¹⁰ [13]; however, SCI does not have ZK and its concrete argument length and proving time are still quite large. IOPs with (perfect) ZK and scalable verifiers were recently described, first for NP [17], then for NEXP [15]. In both works, prover running time (T_P) is bounded by $T_C \cdot \log^{O(1)} T_C$; we refer to this as *scalable proving time* (also known as *quasi-linear* proving time).

Henceforth, we shall call a (universal) ZK system (vi’) *fully scalable*, or, simply *scalable*, if both prover and verifier running times are scalable; this is justified because both running times are nearly-optimal, up to poly-logarithmic factors. A ZK-IOP system satisfying properties (i)–(v) and full scalability (vi’) will be called a *Scalable Transparent IOP of Knowledge* (ZK-STIK); see Section 3 for formal definitions. Summarizing, *theoretical* constructions of ZK-STIK systems were recently presented, but their concrete efficiency and applicability to “practical” computations have not been demonstrated thus far.

1.1 Main contribution

We present a new construction of a (doubly) scalable and transparent ZK system in the IOP model (a ZK-STIK); see Theorems 3.4 and 3.5 for details. We realize this system as a ZK-STARK and apply it to a proof-of-concept “meaningful” computation that is highly sequential in nature — the DPM problem presented earlier. Our realization achieves (i) verification time that is strictly smaller than naïve running time ($T_V < T_C$) and (ii) communication complexity that is strictly smaller than witness size. The core innovation and main source of improved performance in this system is the extended reliance on the IOP model, including the *Fast Reed-Solomon (RS) IOP of Proximity (IOPP)* (FRI) protocol discussed in Section 2 (cf. [14]) and a new arithmetization procedure (see Section 2.3). We stress that the exponential speedup in verification time and witness-size described next (and displayed in Figure 1) apply to any computation that is defined for arbitrarily large witness size, though the particular point at which this speedup materializes depends on the complexity of the computation (as defined in Section 2.2)¹¹.

DNA profile match computation As a proof-of-concept “meaningful” computation we construct a ZK-STARK for the *DNA profile match (DPM)* problem, which we describe informally next (see Appendix E for details). This computation addresses the following hypothetical scenario: Suppose that the Police (acting as the prover P) is in charge of the national forensic DNA profile database (D), and at previous time t has posted (say, on a block-chain) a hiding commitment cm_t to the state D_t of the database at that point in time. The Police now claims that the DNA profile p of the soon-to-be-appointed and alleged-to-be-corrupt Presidential Candidate, does not appear in D_t and thus wishes to create, in a scalable manner, a proof that will convince the public that the DPM computation was carried out correctly, and the output reported by the Police is correct (with respect to p and D_t).

The prevailing standard for DNA profiles, used in over 50 countries, is the *Combined DNA Index System* (CODIS) format; according to this standard an individual is represented by the Short Tandem Repeat (STR) count of his/her DNA, measured for a set of 20 “core loci” [87] (the number of core loci increased from 13 to 20 starting January 2017). The commitment cm_t to the state D_t of a CODIS database is assumed to be public information (say, published at time t on a blockchain), as is a commitment cm_p to the profile

¹⁰<https://github.com/elibensasson/SCI-POC>

¹¹In particular, a computation with parameters similar to the last row of Figure 4 will behave similarly to the DPM computation displayed on Figure 1.

p of the Presidential Candidate; we assume p was extracted by an independent laboratory that handed it (confidentially) to the Police while publishing cm_p publicly. Assume that the Police declares

“The value α is the result of the match search for the profile with commitment cm_p in the database with commitment cm_t ” (*)

The answer α is one of three possibilities: “no match”, “partial match”, or “full match”. The public (open source) computation C is the one that would have been executed by a trusted third party verifying the claim above. This computation requires three public inputs — cm_t, cm_p and A — and two confidential inputs: (i) a DNA profile database D' and individual DNA profile p' . The computation C terminates successfully if and only if the public inputs (cm_t, cm_p, A) and the confidential ones (D', p') satisfy three conditions: (i) the commitment cm' of the confidential input D' equals the public input cm_t ; (ii) the commitment cm'_p of the confidential input p' equals the public input cm_p ; and (iii) the output of the match search for the confidential input p' in the confidential dataset D' leads to the publicly announced outcome α ; see Appendix E.5 for details.

Let $|D(n)|$ denote the bit-length of a dataset $D(n)$ that contains n profiles (each profile is 40 bytes long); let $CC(n)$ denote the communication complexity of the ZK-STARK for $D(n)$, i.e., the total number of bits communicated between prover and verifier; similarly, let $T_C(n)$ denote the time needed to naïvely verify C by executing it on D with n entries, and let $T_V(n)$ denote the time required by V to verify it, (both measured on a fixed physical computer.)

Realizing time and witness-size compression Consider a computation C which requires auxiliary confidential input D that varies in size, like the DPM example. Any ZK-system $S = (P, V)$ for C induces a pair of *rate measures* for time and witness-size, respectively:

$$\rho_{\text{time}}(n) = \frac{T_V(n)}{T_C(n)}; \quad \rho_{\text{size}}(n) = \frac{CC(n)}{|D(n)|} \quad (1)$$

The rate measures (and thresholds defined next) depend on C and the system S , so the notation $\rho_{\text{time}}^{(S;C)}$ would be more precise, but we prefer notational simplicity and assume C and S are known.

A rate value smaller than 1 indicates *compression*, meaning verification in S is more efficient than naïve verification. In fully scalable ZK systems verifier complexity is poly-logarithmic in prover complexity. Therefore eventually, for large enough n , the system achieves compression. Our main claim here is that we exhibit, for the first time, time and witness-size compression for a ZK-STARK for a large-scale sequential computation. Define the *compression threshold* to be the smallest value n_0 such that for all $n \geq n_0$ the rate is less than 1,

$$\theta_{\text{time}} = \min \{n_0 \mid \forall n \geq n_0 \rho_{\text{time}}(n) < 1\}; \quad \theta_{\text{size}} = \min \{n_0 \mid \forall n \geq n_0 \rho_{\text{size}}(n) < 1\} \quad (2)$$

Figure 1 shows the rate measures for the DPM problem on a double logarithmic scale. The time compression threshold is at $\theta_{\text{time}} = 2.8 \times 10^5$ and the witness-size threshold is $\theta_{\text{size}} = 9 \times 10^3$. The largest database for which we could generate a proof during our tests is $n_{\text{max}} = 2^{20} \approx 1.14 \times 10^6$ DNA profiles; larger databases require more disk space and RAM than was available to us. Each profile occupies 40 bytes so $|D_{n_{\text{max}}}| \approx 43$ megabytes. The time-rate for n_{max} is $\rho_{\text{time}}(n_{\text{max}}) = 1/6$ and the witness-size rate is $\rho_{\text{size}}(n_{\text{max}}) = 1/100$. This figure also demonstrates that compression will improve if supported by stronger hardware than that on which our tests were executed. (see Appendix A for more measurements.)

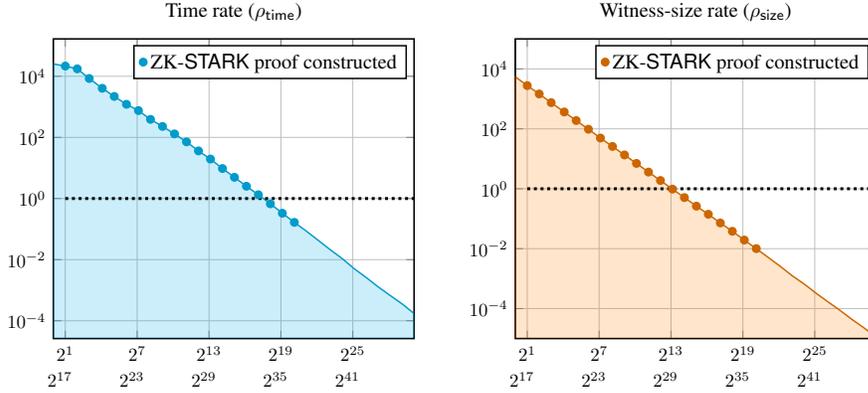


Figure 1: The time (left) and witness-size (right) rate functions of the DPM benchmark ZK-STARK as a function of (i) number of entries (n) in the database (upper horizontal axis) and (ii) number of multiplication gates (lower horizontal axis). Database size, in bytes, is $40 \cdot n$ and processing a single profile corresponds to a circuit with $\approx 2^{16}$ multiplication gates (bottom right entry of Figure 4, explained in Sections 1.3.2 and 2.2). Verifier complexity is independent of prover complexity, so it was measured (on a “standard” laptop; cf. Appendix A.0.2 for machine specifications) even for values of n that are larger than those for which a proof was generated; the values of n for which a proof was generated are marked by full circles.

1.2 Discussion — Applications to decentralized societal functions

Cryptocurrencies, led by Bitcoin, are disrupting established financial systems by suggesting a fully decentralized monetary system to replace fiat currency. Money is but one of the *societal functions* that could be decentralized, and legal contracts are already being replaced by automated *smart contracts* [103] in the Ethereum blockchain. We end this section by discussing the two expected impacts of ZK-STARK systems on decentralized public ledgers.

Scalability A heated discussion is taking place in blockchains today, surrounding the proper way to *scale* the transaction throughput without over-taxing the time and space of nodes participating in the network. As first pointed out by one of the co-authors [12] and embraced recently by several crypto-currency initiatives [66, 34, 76], fully scalable proof systems (even without zero-knowledge) could solve the scalability problem by exponentially decreasing verification time. In more detail, a single “prover node” can generate in quasilinear time a proof that will convince all other nodes to accept the validity of the current state of the ledger, without requiring those nodes to naïvely re-execute the computation, nor to store the entire blockchain’s state, which would be required for such a naïve verification.

Privacy The confidentiality of ZK proofs is already being used to enhance coin fungibility and financial privacy in cryptocurrencies. The Zerocash protocol [18] — recently implemented in the Zcash™ cryptocurrency [89, 67] — uses a particular kind of ZK proofs called Succinct Non-interactive ARGuments of Knowledge (ZK-SNARK) based on cryptographic *knowledge of exponent (KOE)* assumptions [53, 21] to maintain with integrity a decentralized registry whose entries are hiding commitments of unspent funds. These ZK-SNARKs are non-transparent as they require a “setup phase” which uses non-public randomness that, if compromised, could be used to compromise the system’s security (see Section 1.3). Looking forward, ZK-STARKs could replace ZK-SNARKs and achieve the fungibility and confidentiality of Zcash™, transparently. Currently, ZK-SNARKs are roughly $1000\times$ shorter than ZK-STARK proofs so replacing ZK-SNARKs with STARKs calls for more research to either shorten proof length, or aggregate and compress

several ZK-STARK proofs using incrementally verifiable computation [105] (cf. [29]).

1.3 Comparison to other realized universal ZK systems

Recent years have seen a dramatic effort to realize in code zero knowledge proof systems using various theoretical approaches that differ from that of our ZK-STARK. Many of these systems outperform our ZK-STARK for sufficiently small-size computations, for low-depth parallel computations, and/or for batched and amortized computations; all of these cases are extremely useful in practice. But for large scale computations, especially *sequential ones*, the improved full scalability of our IOP-based approach is, eventually, noticeable.

Next, we briefly survey the different *implemented* approaches that are *universal*, i.e., apply to general computations and languages in NP; the interested reader is referred to [110] and [13] for more information on computational integrity solutions, including ones that are non-universal and/or without zero-knowledge. We start by an “asymptotic” discussion in Section 1.3.1 and continue with a comparison of concrete parameters for published and realized systems (Section 1.3).

1.3.1 Theoretical discussion

Within the vast (and growing) literature on realizations of ZK systems, we must limit the scope of our discussion and do so somewhat arbitrarily, by considering only systems that are ZK, Turing complete, and which have been realized in code. We compare these for the most general class of computational integrity statements (see Definition 3.1 for a formal definition) and consider four properties: asymptotic (i) prover scalability (quasilinear running time), (ii) asymptotic verifier scalability (poly-logarithmic verification), (iii) transparency, and (iv) post-quantum security. The first three terms are formally defined in Definition 3.3, and the last one is informal, but could be replaced with the property of *reliance only on collision resistant hash functions*¹². Figure 2 summarizes our discussion, and we provide details next.

- **Homomorphic public-key cryptography (hPKC):** This approach, initiated by Ishai et al. [69] (for the “designated verifier” case) and Groth [60] (for the “publicly verifiable” case), uses an efficient information-theoretic model called a “linear PCP” that is then “compiled” into a cryptographic system using hPKC. An extremely efficient instantiation, based on Quadratic Span Programs, was introduced by Gennaro et. al [53] (see [64, 52, 81, 30, 62, 63] for related work and further improvements). It serves, e.g., as the proof system behind Zerocash and Zcash™. The first implementation of a QSP based system is called Pinocchio [88], with subsequent implementations including lib-SNARK [21, 96] (discussed in the next section) which is used in the Zerocash and Zcash™ implementations; additional implementations appear in [98, 101, 100, 99, 24, 108, 46].

The theoretical differences between hPKC and ZK-STARK are that of transparency and post-quantum security — hPKC lacks both. Verification time in hPKC is scalable (i.e., poly-logarithmic in T_C) only for computations that are repeated many times, because the hPKC “setup phase” requires time $\geq T_C$.

- **Discrete logarithm problem (DLP):** An approach initiated by Groth [61] (cf. [97]) and implemented in [31], relies on the hardness of the DLP to construct a system that is transparent. Shor’s quantum factoring algorithm solves the DLP efficiently, rendering this approach quantum-susceptible. Additionally, verifier complexity in the DLP approach requires time $\geq T_C$ hence it is non-scalable (according

¹²This assumption covers only the interactive setting; see discussion in Section 3.3.

to our definition of the term), although communication complexity in the DLP approach is logarithmic. We refer to the initial implementation of this system as BCCGP [31], and a recent improved version is called BulletProofs [35].

- **Interactive Proofs (IP) based:** IP protocols can be performed with zero knowledge [11] but only recently have IP protocols been efficiently “scaled down” to small depth (non-sequential) computations via so-called “proofs for muggles” of Goldwasser et al. [58, 94]. This led to a line of realizations in code, early works lacked ZK [42, 41, 104, 107], but the state-of-the-art ones, like [112] and Hyrax [109], do have it.

Like ZK-STARK, these recent IP-based proofs are transparent and have a scalable prover, but are quantum-susceptible and their verifier is not scalable, as it scales linearly with computation time for “standard” (i.e., sequential) computations (like other approaches, it is quite efficient for batched and amortized computations and for small circuits).

- **Secure multi-party computation (MPC):** This approach, suggested by Ishai et al. [70] and implemented first in the ZKBoo [55] system, and more recently, in Ligerio [1], “compiles” secure MPC protocols into ZK-PCP systems, by requiring the prover to commit to the transcript of a secure MPC protocol, and then reveal the view of one of the parties.

Like ZK-STARK, the MPC-based proofs are transparent, post-quantum secure and have scalable (quasilinear) proving time. However, MPC based systems have a non-scalable verifier, one that runs in time $\geq T_C$ and communication complexity is non-scalable, it is $\sqrt{T_C}$ in the state of the art system [1]; for concrete circuits and amortized computations it is, nevertheless, extremely efficient.

- **Incrementally Verifiable Computation (IVC):** This approach, suggested by Valiant [105] (cf. [39, 29]) reduces prover space consumption by relying on knowledge extraction assumptions; this approach can be applied on top of other proof systems with succinct (sub-linear) verifiers, including ZK-STARK, but thus far has been realized only for a single hPKC system [23].

Compared with ZK-STARK, systems built this way inherit most properties from the underlying proof system. In particular, the hPKC-based IVC is non-transparent and quantum-susceptible; however the verifier is scalable even for a computation executed only once, because the setup phase runs in poly-logarithmic time.

	prover scalability (quasilinear time)	verifier scalability (polylogarithmic time)	Transparency (public randomness)	Post-quantum security
hPKC	Yes	Only repeated computation	No	No
DLP	Yes	No	Yes	No
IP	Yes	No	Yes	No
MPC	Yes	No	Yes	Yes
IVC+hPKC	Yes	Yes	No	No
ZK-STARK	Yes	Yes	Yes	Yes

Figure 2: Theoretical comparison of universal (NP complete) realized ZK systems.

1.3.2 Concrete performance

Different ZK proof systems are based on different cryptographic assumptions and are designed for different computational problems. Their realizations are written in different programming languages and tested on varied hardware. Therefore, exact “apples to apples” comparisons are difficult, if not impossible, to perform. Having said that, in this section we *attempt* to qualitatively compare the realized proof systems reported in the previous section in terms of verifier, prover, and communication complexity for computational problems that are similar in nature to the DPM.

Arithmetic circuit complexity as standard measuring yard All realized proof systems surveyed here (including our ZK-STARK) use arithmetization to reduce computational integrity (CI) statements to statements about systems of low-degree polynomials over finite fields (see Section 2). All other surveyed systems use prime fields \mathbb{F}_p , though some (like MPC- and IP-based) could operate also over binary fields, like ZK-STARK; we stress that ZK-STARK could also operate over prime fields¹³ but we have not realized this in code. Most systems (ZK-STARK not included) reduce CI statements to *arithmetic circuits*, i.e., ones that correspond to constraints that are *quadratic* polynomials; ZK-STARK reduces to systems of higher-degree polynomials, e.g., for our DPM benchmark this degree is 8.

Arithmetic circuit complexity is a reasonable metric to use in order to compare various proof-systems. The main parameters that influence proof-system complexity (and are mentioned in prior works) are *circuit depth*, *circuit width* (number of gates in each “level” of the circuit), *nondeterministic witness size*, and *multiplication complexity*, i.e., the number of multiplication gates. (Addition complexity is also relevant, but most proof systems are less affected by it.)

Our DPM computation corresponds to an arithmetic circuit with the following parameters, when applied to a database with n entries (see Figure 4):

- circuit depth is $\text{depth}_n = 62 \cdot n$;
- circuit width is $w = 81$;
- witness complexity is $\text{wit}_n = 40 \cdot n$ bytes;
- multiplication complexity is $\text{mult}_n = 1467 \cdot 62 \cdot n = 90954 \cdot n \approx 2^{16.4} \cdot n$.

As discussed at length in Appendix A, we measured the full ZK-STARK system (prover+verifier) with 60 bits of security for $n = 2^k$, $k = 1, \dots, 20$, i.e., for arithmetic circuits with depth up to $\text{depth}_n \approx 2^{25.9}$, and with up to $\approx 2^{36.4}$ multiplication gates over $\mathbb{F}_{2^{64}}$; the (nonadaptive) verifier alone was measured even for larger inputs, up to $n = 2^{36}$ (see Appendix A).

To attempt an “apples to apples” comparison with other systems, we ran several of them on a single machine — one that is different than that used to measure the DPM code¹⁴ — using the same benchmark computation based on the “exhaustive subset-sum” computation measured in prior work [13]; the results are summarized by Figure 3. The systems measured thus far this way are

- libSNARK (commit dc78fd, September 7, 2017) with 80-bit security
- SCI (same measurements used in [13]) with 80-bit security

¹³the FRI system requires p to contain a sufficiently large multiplicative subgroup of order $2^{t+O(1)}$; such prime fields abound, as implied by Linnik’s Theorem [80].

¹⁴The machine used to measure the DPM code was kindly offered to us by Intel™ for a limited time, whereas for the “apples-to-apple” comparison we needed to provide other teams with access to a machine, for long periods of time.

- BCCGP with logarithmic communication complexity and 128-bit security, single threaded (same system used in [31])
- Ligerio with 60 bits of security (same system as reported in [1]);
- our ZK-STARK, with 60-bit security ZK-STARK; we estimate prover prover time for 80 bits to be at most 5% longer; cf. Appendix B.5.1.

We now briefly discuss the performance of these (and other prior reported works), focusing on the following complexity parameters: prover time, verifier time, and communication complexity.

Prover complexity Nearly all systems surveyed earlier have prover complexity that scales either linearly or nearly-linearly in computation size. As shown in Figure 3, our ZK-STARK prover is at least $10\times$ faster than the other measured systems across the full range of compared computations (all systems were tested up to maximal proving time of 12 hours). We hope to perform similar “apples-to-apples” comparisons (i.e., same machine, circuit depth, width and size) with other systems like Hyrax and BulletProofs in future work.

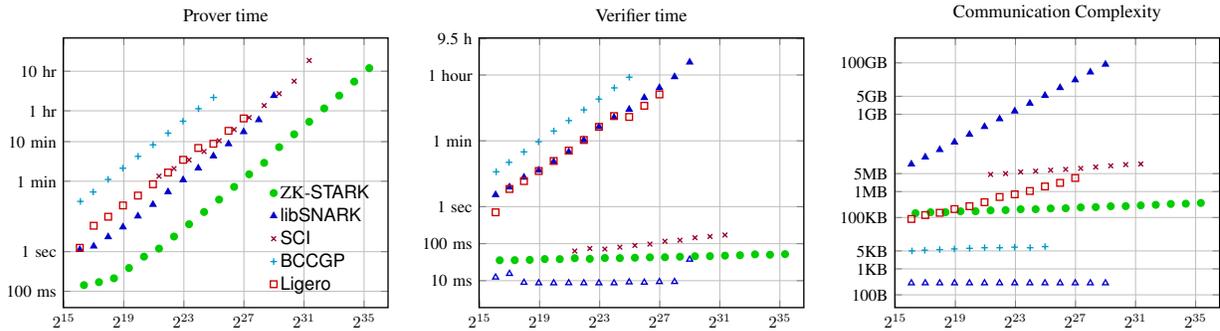


Figure 3: An “apples-to-apples” comparison of different realized proof systems as function of computation size, measured by number of multiplication gates. All systems were tested on the same server (specs below) and executed a computation of size and structure corresponding to the “exhaustive subset-sum” program from [13, Section 3]. The compared systems are SCI (purple x-marks), which lacks ZK, libSNARK (blue triangles), BCCGP (cyan +-marks), executed in single-thread mode, Ligerio (red squares) and ZK-STARK (green circles). From left to right, we measure prover time, verifier time and communication complexity. For libSNARK, the hollow marks in the middle and right plots measure *only post-processing* verification time and CC, respectively; the full marks measure *total* verification time and CC, and this includes the (non-transparent) key-generation phase. Server specification: 32 AMD cores at clock speed of 3.2GHz, with 512GB of DDR3 RAM. (Each pair of cores shares memory; this roughly corresponds to a machine with 16 cores and hyper-threading.)

Verifier complexity Different proof systems excel on different circuit topologies. For example, Ligerio achieves best performance for circuits of size s that are iterated s times (i.e., when $\text{depth} \approx w \approx \sqrt{\text{mult}}$), and Hyrax works best on small depth, massively parallel, circuits ($\text{depth} = O(1)$ and $w, \text{mult} \gg \text{depth}$). The concrete performance of IOP-based systems on such circuit topologies is an interesting question, left for future work.

For “deep” and “narrow” circuits, like the ones arising from the DPM, verifier arithmetic complexity of prior works scales at least like $\sqrt{\text{mult}}$ (and, often, like mult), whereas our ZK-STARK scales like $w + \log \text{mult}$ (see Theorems 3.4 and 3.5). Consequently, for medium- and large-scale sequential computations our ZK-STARK verifier time is better than other solutions, as shown by the middle plot of Figure 3. We

expect the comparison with other works, like **Hyrax** and **BulletProofs**, to behave similarly; in particular, the **Hyrax** prover reaches ≈ 10 seconds for a circuit with $\approx 2^{28}$ gates (but measured on a different machine than ours); **BCCGP** and **BulletProofs** require even greater running time [109, Figure 4.(i)]. For comparison, the **ZK-STARK** verifier for the DPM computation requires less than 50 ms (on a different machine), even for huge circuits¹⁵, with $n = 2^{36}$ entries (profiles), $wit_n \approx 2.5$ -terabyte size witnesses and arithmetic circuits with $\text{mult}_n = 2^{52}$ multiplication gates and depth $\approx 2^{40}$ (cf. Figure 7).

The hPKC systems like **Pinocchio** and **libSNARK**, and IVC+hPKC systems like that of [23], are different in this respect. They have a pre-processing phase that is performed only once per circuit. For **Pinocchio** and **libSNARK** pre-processing time grows *linearly* with circuit size. E.g., the **libSNARK** system requires ≈ 16 seconds for a computation with 2^{20} gates. For the IVC+hPKC system, pre-processing time is *constant* and does not depend on circuit size; however, this constant is quite large compared to our verifier time, it is ≈ 10 seconds for a computation similar to our DPM.

Communication complexity (CC) The use of a pre-processing phase in the hPKC and IVC+hPKC systems leads to extremely small post-processing CC; the **BCCGP** system also enjoys extremely short CC and, because its pre-processing is transparent, can be effectively replaced with a short seed to a pseudo-random generator. Concretely, for all computations measured in practice, post-processing CC of **Pinocchio**, **libSNARK** and the IVC+hPKC system are less than 300 bytes, and that of **BCCGP** is less than 7KB [31] (see also Figure 3). However, pre-processing key length scales linearly with circuit size for hPKC; the IVC+hPKC system is different in this respect, it has succinct pre-processing length even for large computation size, but once again, this length is concretely large — more than 40 MB for a computation like our DPM.

For **Ligero**, communication complexity scales like $70\sqrt{\text{mult}_n}$ field elements [1, Section 5.3], and for **Hyrax** it scales like $wit^{1/k} + 10 \cdot \text{depth} \cdot \log w$ field elements for arbitrary integer k [109, Section 1]; increasing k decreases CC but also increases verification time (which is at least $wit/(wit^{1/k})$). Using the estimate for **Hyrax**, a quick calculation shows that for a circuit arising from our DPM computation with, say $n = 2^{13}$ profiles, the CC of **Hyrax** would reach several megabytes, compared with **ZK-STARK** CC that is less than 1 megabyte even for $n = 2^{36}$ profiles.

Summary Among all ZK systems tested in the “apples-to-apples” manner described above, our **ZK-STARK** has the fastest prover for all circuit-sizes we were able to measure; in particular, it is $\approx 10\times$ faster than the second fastest measured system — **libSNARK**. Other systems perform better (shorter communication, faster verification) on small circuits (**ZKBoo**, **Ligero**), small-depth circuits (**Hyrax**), and on computations repeated many times with the same fixed circuit (**BulletProofs**, **Pinocchio**, **libSNARK**). However, for general large scale computations our **ZK-STARK** has verification time and communication complexity outperform all other *transparent* systems published thus far for this range of parameters. In other words, our particular **ZK-STARK** realization shows that the asymptotic benefits of full scalability and transparency are manifested already for concrete computations that are practically relevant, like the DPM, and suggest that our type of system is potentially useful for constructing scalability solutions, e.g., for decentralized crypto-currencies (as discussed in Section 1.2).

¹⁵We stress that current hardware does not support generating proofs for such large instances, as discussed later.

2 Methods

This section highlights the main innovative components that underlie the (double) scalability and concrete efficiency of our ZK-STARK; the exposition is short and informal. Later, in Section 3 we shall formally define the theoretical model which our ZK-STARK uses, and state the main theorems for this model (Theorems 3.4 and 3.5); then, in Appendix B we formally present the steps of the reduction (proved in subsequent sections).

Overview Many ZK systems (including ours) use *arithmetization*, a technique first¹⁶ used to prove circuit lower bounds [93, 102], then adopted to interactive proof systems [5, 82]. Arithmetization is the reduction of *computational* problems to *algebraic* problems, that involve “low degree” polynomials over a finite field \mathbb{F} ; in this context, “low degree” means degree is significantly smaller than field size.

The start point for arithmetization in all proof systems is a computational integrity statement which the prover wishes to prove, like

“ α is the result of executing C for T steps on (public) input x ” (**)

Notice the DPM statement (*) is a special case of (**). For our ZK-STARK, and for related prior systems [27, 25, 13], the end point of arithmetization is a pair of *Reed-Solomon (RS) proximity testing (RPT)* problems¹⁷, and the scalability of our ZK-STARK relies on a new solution to the RPT problem, discussed first; later we explain the arithmetization process in more detail.

2.1 Fast Reed-Solomon Interactive Oracle Proof of Proximity (FRI3rd)

For $S \subset \mathbb{F}$ and rate parameter $\rho \in (0, 1)$, the Reed-Solomon code $\text{RS}[\mathbb{F}, S, \rho]$ is the family of functions $f : S \rightarrow \mathbb{F}$ that are evaluations of polynomials of degree $< \rho|S|$. The RPT problem assumes a verifier is given oracle access to f , and to auxiliary information like a probabilistically checkable proof of proximity (PCPP) [25, 48] or an interactive oracle proof of proximity (IOPP) [22, 94, 15]; the verifier’s task is to distinguish with high probability and with a small number of queries to f and the auxiliary PCPP/IOPP oracle(s), between the case that $f \in \text{RS}[\mathbb{F}, S, \rho]$ and the case that f is 0.1-far from (all members of) $\text{RS}[\mathbb{F}, S, \rho]$ in relative Hamming distance. Finding solutions to the RPT problem (a special case of the “low-degree testing” problem) is a major bottleneck for transparent systems.

Our ZK-STARK uses a new protocol to solve RPT, called the *Fast RS IOPP (FRI)*. FRI is the first RPT solution to achieve prover arithmetic complexity that is *strictly linear* — $6 \cdot |S|$ arithmetic operations in \mathbb{F} — and verifier arithmetic complexity that is *strictly logarithmic*: $21 \cdot \log |S|$ arithmetic operations; additionally, the proof can be constructed in $\log |S|$ cycles on a parallel machine, and is structured as to lead to short arguments (see Section 2.5). FRI improves significantly, both asymptotically and concretely, on the previous RPT solutions which required *quasilinear* prover arithmetic complexity ($\theta(|S| \cdot \log^{O(1)} |S|)$). See [14] for a detailed description.

2.2 Arithmetization I — Algebraic Intermediate Representation (AIR)

Having discussed its end point, we return to describe the innovative components of our ZK-STARK within the arithmetization process itself. The arithmetization is comprised of several phases that are similar to other

¹⁶Earlier reductions, such as the one used in Gödel’s Incompleteness Theorem, involved *infinite* algebraic domains, in particular the natural numbers [56].

¹⁷The other solutions described in Section 1.3 have different end points.

program and circuit compilation processes, so we borrow terminology used there and adapt it to our process.

The first phase of arithmetization is that of constructing an *algebraic intermediate representation (AIR)* of the program C . Informally, the AIR is a set

$$\mathcal{P} = \{P_1(\vec{X}, \vec{Y}), \dots, P_s(\vec{X}, \vec{Y})\}$$

of low degree polynomials with coefficients in \mathbb{F} over a pair of variable sets $\vec{X} = (X_1, \dots, X_w)$ and $\vec{Y} = (Y_1, \dots, Y_w)$ that represent respectively the current and next state of the computation¹⁸ (see Appendix C and Definition B.3 for more details). The AIR defines the *transition relation* of the computation C in the sense that a pair $(\vec{x}, \vec{y}) \in \mathbb{F}^w \times \mathbb{F}^w$ corresponds to a single valid transition (or “cycle”) of C if and only if

$$P_1(\vec{x}, \vec{y}) = \dots = P_s(\vec{x}, \vec{y}) = 0,$$

i.e., if and only if (\vec{x}, \vec{y}) is a common solution of the AIR system \mathcal{P} . The following parameters of \mathcal{P} determine prover and verifier complexity, so minimizing them is a major goal of this phase. The *degree* of the AIR is $\deg(\mathcal{P}) = \max_{i=1}^s \deg(P_i)$; the *(state) width* is the number of variables (w) needed to represent a state; the *(AIR) size* is the number of constraints (s), and the *cycle count* is the number of machine cycles needed to execute¹⁹ C ; when the program processes a large number (n) of data elements, as is the case for the DPM benchmark, we are interested in the number of cycles per element, denoted c ; the total cycle count for n elements is $c \cdot n$. If the computation is “expanded” to a circuit (as commonly done in the other solutions described in Section 1.3), the cycle count is a lower bound on circuit depth; for the sake of comparison with those other systems, we compute in the rightmost column of Figure 4 the total number of multiplication gates for this expanded circuit, as this measure along with circuit depth, are the complexity measures that dictate prover and verifier complexity.

A major contributor to prover complexity in our benchmarks is the cost of proving computational integrity of repeated invocations of a cryptographic hash function; other computations are negligible compared to this cost. Thus, choice of the particular hash function (H) is of great importance, as is its definition in terms of \mathcal{P} . Our ZK-STARK uses the binary (characteristic 2) field $\mathbb{F}_{2^{64}}$ because (i) it has efficient arithmetic operations (e.g., addition is equivalent to exclusive-or) and (ii) its algebraic structure is needed for the FRI3rd protocol. Therefore, the cryptographic hash function we seek is one that is “*binary field friendly*”, meaning, informally, its AIR has small complexity parameters when defined over binary fields. Figure 4 summarizes the main AIR complexity parameters for the DPM benchmark described in Section 1 and for three hash functions: the Secure Hash Algorithm 2 (SHA2) family [92] and the Davies–Meyer [111] hash based on the Rijndael block cipher [44] with 128 bits (AES128+DM) and with 160 bits (Rij160+DM). See Appendices E and F for details.

2.3 Arithmetization II — Algebraic Linking Interactive Oracle Proof (ALI)

The main bottleneck for prover time and space complexity is the cost of performing *polynomial interpolation* and its inverse operation — multi-point *polynomial evaluation*; we discuss both in Section 2.4. The complexity measure that dominates this bottleneck is the *maximal degree* of a polynomial which the prover must interpolate and/or evaluate; for a computation on a dataset of size n denote this degree by $d^{\max}(n)$. Prior state-of-the-art [27, 20, 38, 13] gave

$$d_{\text{old}}^{\max}(n) = n \cdot c \cdot w \cdot d + n \cdot c \cdot s. \quad (3)$$

¹⁸This informal description omits, for simplicity, the *boundary conditions*, like public inputs and outputs of the computation.

¹⁹In general, this number may depend arbitrarily on the particular input, however, in all our benchmarks it depends linearly on the size (n) of the input dataset.

	State size w	Cycles c	Degree d	System size s	#× gates	#+ gates	total #× gates
SHA2	56	3762	11	1065	720	3194	2708640 ≈ 2 ^{21.3}
AES128+DM	62	48	8	327	486	1033	23328 ≈ 2 ^{14.5}
Rij160+DM	68	58	8	318	891	1390	51678 ≈ 2 ^{15.7}
DPM	81	62	8	626	1467	1520	90954 ≈ 2 ^{16.4}

Figure 4: Main complexity parameters of basic cryptographic primitives and the benchmark DNA profile match search program. The first four measures are explained in Section 2.2. The 5th and 6th columns measure total number of multiplication and addition gates in the AIR; the last column measures total number of multiplication gates (product of second and 5th columns) for the sake of comparison with other ZK systems that use this measure (cf. Section 1.3.2).

which leads to concretely large values (see first column of Figure 5). Our ZK-STARK reduces d^{\max} to

$$d_{\text{ZK-STARK}}^{\max}(n) = n \cdot c \cdot d \quad (4)$$

which results in a multiplicative savings factor of $6.5 \times 10^4 - 1.8 \times 10^5$ over prior works (see the last two columns of Figure 5). The improved efficiency of our ZK-STARK is due to two reasons, explained next. The first one completely removes the second summand of (3) and the second one removes w from its first summand.

Algebraic linking IOP (ALI) The second summand of (3) arises because our prover needs to apply a “local map” induced by the AIR system \mathcal{P} (see [17] for a discussion of “local maps”). Prior state-of-the-art systems, like [13], used a local map that checks each constraint of the AIR separately, leading to this second summand. Instead, our ZK-STARK uses a single round of interaction to reduce all s constraints to a *single constraint* that is a *random linear combination* of P_1, \dots, P_s . This round of interaction completely removes the second summand of (3).

Register-based encoding The naïve computation performed by the prover can be recorded by an *execution trace*, a two-dimensional array with $c \cdot n$ rows and w columns, in which each row represents the state of the computation at a single point in time¹⁹ and each column corresponds to an *algebraic register* tracked over all $c \cdot n$ cycles. Prior systems, like [13], encoded the *full* execution trace by a *single* Reed-Solomon codeword, leading to degree $n \cdot c \cdot w$; this degree is then multiplied by d to account for application of the afore-mentioned “local map” to the codeword, resulting in the first summand of (3). Our ZK-STARK uses a *separate* Reed-Solomon codeword for each register²⁰, leading to w many codewords, each of lower degree $n \cdot c$. At first glance this tradeoff may seem wasteful, because we now have to solve an RPT problem for each of these w codewords. However, the interaction and use of randomness allowed by the IOP model once again come to our aid: it suffices to solve a *single* RPT problem, applied to a *random* linear combination of all w codewords. The use of a single codeword per register also helps with reducing communication complexity, as explained in Section 2.5 below.

Figure 5 compares the d^{\max} value of our ZK-STARK to that of the prior state of the art [27, 20, 38, 13] and shows a multiplicative reduction factor of $6.5 \times 10^4 - 1.8 \times 10^5$ for the computations discussed in Section 2.2 and Figure 4.

²⁰For simplicity, the current description discusses the case of space bounded computations; the case of computations with large space also uses multiple codewords but the reduction is more complicated, see Appendix C.3.

	$d_{\text{old}}^{\max}(1)$	$d_{\text{ZK-STARK}}^{\max}(1)$	$d_{\text{old}}^{\max}/d_{\text{ZK-STARK}}^{\max}$
SHA2	6323922	41382	18491
AES128+DM	39504	384	6584
Rij160+DM	49996	464	6896
DPM	78988	496	10192

Figure 5: The maximal degree (d^{\max}) as given by formulas (3) and (4), respectively, for the computations discussed in Section 2.2. The last column gives the multiplicative improvement factor of $d_{\text{ZK-STARK}}^{\max}$ over the prior state of the art.

2.4 Low degree extension and composition degree

Decreasing d^{\max} affords our ZK-STARK greater scalability. But, eventually, as the input size n grows, so does d^{\max} . The main bottleneck for prover time and space is the computation of the *low degree extension (LDE)* of the execution trace, defined next.

Definition 2.1 (Low degree extension (LDE)). *Given finite subsets S, S' of a field \mathbb{F} satisfying $|S'| > |S|$ and a function $f : S \rightarrow \mathbb{F}$, the low degree extension (LDE) of f to S' is the function $f' : S' \rightarrow \mathbb{F}$ that has the same interpolating polynomial as that of f .*

The LDE is typically computed by polynomial interpolation, followed by a polynomial multi-point evaluation step. State of the art algorithms for interpolating and evaluating polynomials over binary fields are known as *additive FFT* algorithms because they resemble the fast Fourier transform (FFT) [40]. To improve prover scalability, our ZK-STARK uses the recent and novel additive FFT of Lin et al. [79], inspection of this algorithm shows it leads to arithmetic complexity of $3 \cdot |S'| \log |S'|$ for the sets S' that our system requires (see Theorem B.2). Prior additive FFTs [36, 106, 51] required $\theta(N \log N \log \log N)$ operations; moreover, the memory-access pattern of this encoding algorithm leads to favorable running times compared with prior implementations [13].

Using this additive FFT, which has *strictly* quasi-linear arithmetic complexity²¹, we also obtain the first ZK-STARK in which prover complexity is *strictly* quasi-linear, *and* verifier complexity is *strictly* logarithmic, in the size of the execution trace $c \cdot n \cdot w$; see Lemma B.6.

2.5 Minimizing Authentication Path Complexity (APC) and Communication Complexity (CC)

The largest contributor to communication complexity, and to verifier time and space complexity in ZK-STARK (and prior related works [27, 20, 38, 13]) is the cost of *realizing* the IOP model via Merkle trees. We now discuss the way our ZK-STARK reduces this cost.

The *commit-reveal* scheme of Kilian [74] (which uses the “cut-and-choose” method of Brassard et al. [33]) has the prover *commit* to each oracle by sending the root of a Merkle tree whose leaves are labeled by oracle entries. Recall an IOP involves *several* oracles, hence also several Merkle trees and several roots/commitments. After the prover has committed to all oracles, the verifier queries these oracles at randomly chosen positions. When the prover *reveals* the oracle answers to these queries, each answer must be appended with an *authentication path* proving the query answers are consistent with previously committed Merkle tree roots. Let λ denote the number of output bits of the cryptographic hash function used to construct a Merkle tree in our system; let AP_{total} denote the total number of authentication path nodes in all

²¹A function $g(n)$ is called *strictly* quasi-linear if $g(n) = O(n \log n)$, and called *strictly* logarithmic if $g(n) = O(\log n)$.

subtrees of Merkle trees whose leaves are query answers, and let q_{total} denote the total number of queries, made to all proof oracles. The total communication complexity (CC) of the proof system is

$$CC = q_{total} \cdot \log |\mathbb{F}| + AP_{total} \cdot \lambda \quad (5)$$

In addition to reducing the first summand above by improved soundness analysis, our ZK-STARK also reduces the second summand in two separate ways:

1. The ZK-STARK verifier queries *rows* of the (LDE of the) execution trace, each row comprised of w field elements that represent the state at some point in the computation (or its LDE). To reduce communication complexity, the ZK-STARK prover places each such row in a *single* sub-tree of the Merkle tree, and therefore only *one* authentication path is required per row (as opposed to w many paths in prior solutions).
2. The QUERY phase of the FRI protocol queries affine cosets of a fixed subspace. Accordingly, the ZK-STARK prover places each such coset in a *single* sub-tree of the Merkle tree, thereby reducing the number of authentication paths to one-per-coset (as opposed to one per field element).

Finally, to improve running time and further reduce communication complexity, we use the Davies–Meyer hash composed with AES as the hash function for our ZK-STARK commit-reveal scheme (recall AES is part of the instruction set of many modern processors).

2.6 Organization of the remaining sections

The following sections give a full and formal description of our construction. Section 3 formally defines the notion of a ZK-STIK and its realization as a ZK-STARK, and presents the main asymptotic results (Section 3.2); along the way we recall the formal definitions of the IOP model. Appendix B describes the main components used in our construction, and uses these to prove our main results (Appendix B.7). The compenets are then described in more detail in the remaining sections.

3 On STIKs and STARKs — formal definitions and prior constructions

In this section we state the main theorems that our ZK-STARK realizes (Section 3.2). Along the way we explain what constitutes a ZK-STARK (Section 3.3) and point to earlier relevant works that are variants of it (Section 3.4). We assume familiarity with standard definitions of zero knowledge (ZK) interactive proof (IP) and argument systems [59, 33], probabilistically checkable proofs (PCP) [2] and PCPs of proximity (PCPP) [26, 48], as well as interactive oracle proofs (IOP) [22, 94] and interactive oracle proofs of proximity (IOPP) [16].

A nondeterministic machine M that decides a language $L \in \text{NTIME}(T(n))$ in time $T(n)$ (n denotes instance size) *induces* a binary relation R_M consisting of all pairs $(\mathfrak{x}, \mathfrak{w})$ where $\mathfrak{x} \in L$ and \mathfrak{w} is a sequence of nondeterministic choices of $M(\mathfrak{x})$ that lead to an accepting state. In this case we say $R = R_M$ is *induced* by L and implicitly assume M is fixed and known. The language that we shall be most interested in, is the NEXP-complete *computational integrity* language²².

Definition 3.1 (Computational Integrity). *The binary relation R_{CI} is the set of pairs $(\mathfrak{x}, \mathfrak{w})$ where*

- $\mathfrak{x} = (M, x, y, T, S)$ with M a nondeterministic Turing Machine, x and y denote input and output, and $T \geq S$ are integers in binary notation, indicating time and space bounds, respectively
- \mathfrak{w} is a description of the nondeterministic choices of M on input x that cause it to reach output y within $\leq T$ steps, using a memory tape of size at most S (not including the read-only input tape on which \mathfrak{x} is written).

The computational integrity (CI) language L_{CI} is the projection of the binary relation R_{CI} onto its first coordinate; alternatively, $L_{\text{CI}} \triangleq \{\mathfrak{x} = (M, x, y, T, S) \mid \exists \mathfrak{w} (\mathfrak{x}, \mathfrak{w}) \in R_{\text{CI}}\}$.

The space bound S in the definition above is unneeded, because $S \leq T$. However, IOP systems for space bounded computations ($S = o(T)$) are simpler and, often, concretely more efficient (this holds for our DPM computation). Thus, we treat space-bounded computations separately and dedicate a theorem to it (Theorem 3.4) before treating the more general (and complicated) case of general computational integrity (Theorem 3.5).

3.1 Scalable Transparent IOP of Knowledge (STIK)

A STARK, defined later, is a realization of a *scalable and transparent IOP of knowledge* (STIK), discussed next. We start by recalling the IOP model as defined in [22].

Definition 3.2 (Interactive Oracle Proof (IOP)). *Let R be a binary relation induced by a nondeterministic language L and let $\epsilon \in [0, 1]$ denote soundness error. An Interactive Oracle Proof (IOP) system S for R with soundness ϵ is a pair of interactive randomized algorithms $S = (P, V)$ that satisfy the properties below; P is the prover and V is the verifier.*

- **operation:** *The input of the verifier is \mathfrak{x} , and the input of the prover is $(\mathfrak{x}, \mathfrak{w})$ for some string \mathfrak{w} . The number of interactive rounds, denoted $r(\mathfrak{x})$, is called the round complexity of the system. During a single round the prover sends a message (which may depend on \mathfrak{w} and prior messages) to which the verifier is given oracle access, and the verifier responds with a message to the prover. We denote by $\langle P(\mathfrak{x}, \mathfrak{w}) \leftrightarrow V(\mathfrak{x}) \rangle$ the output of V after interacting with P ; this output is either accept or reject.*

²²This language is called the “Computationally Sound” language in [85] and the “universal language” in [8]; we choose the name used in [13].

- **completeness** If $(\mathfrak{x}, \mathfrak{w}) \in R$ then $\Pr [\langle P(\mathfrak{x}, \mathfrak{w}) \leftrightarrow V(\mathfrak{x}) \rangle = \text{accept}] = 1$
- **soundness** If $\mathfrak{x} \notin L$ then for any P^* , $\Pr [\langle P^* \leftrightarrow V(\mathfrak{x}) \rangle = \text{accept}] \leq \epsilon$

The proof length, denoted $\ell(\mathfrak{x})$, is the sum of lengths of all messages sent by the prover. The query complexity of the protocol, denoted $q(\mathfrak{x})$, is the number of entries read by V from the various prover messages. Given witness \mathfrak{w} such that $(\mathfrak{x}, \mathfrak{w}) \in R$, prover complexity, denoted $\text{tp}(\mathfrak{x}, \mathfrak{w})$, is the complexity required to generate all prover messages, and verifier complexity, similarly defined, is denoted $\text{tv}(\mathfrak{x})$.

Next, we formally define a ZK-STIK. Most of the work described in later sections is related to constructing a new, and more efficient, ZK-STIK; similarly, Sections 2.1–2.4 describe a ZK-STIK and only Section 2.5 discusses a ZK-STARK. A (ZK-)STIK can be *proven* to be unconditionally sound, even against computationally unbounded provers; ZK-STARK systems have only computational soundness, against bounded provers, thus require additional cryptographic assumptions, discussed later.

Definition 3.3 (Scalable Transparent IOP of Knowledge (STIK)). *Let R be a binary relation induced by a nondeterministic language $L \in \text{NTIME}(T(n))$ for $T(n) \geq n$ and let $S = (P, V)$ be an IOP for L with soundness error $\epsilon(n) < 1$. We say S is*

- **transparent** if all verifier messages and queries are public random coins.
- **(fully, or doubly) scalable** if for every instance \mathfrak{x} of length n , both of the following hold:
 1. **scalable verifier:** $\text{tv}(n) = \text{poly}(n, \log T(n), \log 1/\epsilon(n))$
 2. **scalable prover:** $\text{tp}(n) = T(n) \cdot \text{poly}(n, \log T(n), \log 1/\epsilon(n))$
- **proof of knowledge** if there exists a knowledge error function $\epsilon'(n) \in [0, 1]$ and a randomized extractor E that, given oracle access to any prover P^* that causes the verifier to accept \mathfrak{x} with probability $p(n) > \epsilon'(n)$, outputs in expected time $\text{poly}\left(\frac{T(n)}{p(n) - \epsilon'(n)}\right)$ a witness \mathfrak{w} such that $(\mathfrak{x}, \mathfrak{w}) \in R$.
- **privacy preservation** if there exists a randomized simulator Sim that samples (perfectly) the distribution on transcripts of interactions between V and P , and runs in time $\text{poly}(T(n))$.

A (fully) scalable and transparent IOP of knowledge will be denoted by **STIK**. For $T(n) = \text{poly}(n)$, a privacy-preserving STIK has perfect²³ zero knowledge (ZK-STIK) but for $T(n) = n^{\omega(1)}$ it implies only the weaker notion of a witness indistinguishable proof system (wi-STIK).

A few remarks regarding the definition above:

- **transparency:** Interactive proofs in which the verifier sends only public random coins are known as *Arthur Merlin type* protocols. The term *transparent proof* was introduced in [6] and is synonymous to PCP. We choose this term because it adequately reflects the beneficial effect of public randomness on the transparency of the proof system and. Our terminology does not contradict the earlier definition of the term because transparent proofs (and PCPs) are also transparent according to the definition above.
- **scalability:** Scalable provers are called “quasi-linear” in a number of prior works and scalable verifiers are often called “succinct”. We identify both terms into a single one that reflects the beneficial effect of quasi-linear and succinct proof systems.

²³We omit the discussion of statistical zero knowledge because all known ZK-STIK systems have perfect zero knowledge.

3.2 Main Theorems

We now state the two main theorems regarding IOP systems, that our ZK-STARK system realizes; the proof of both theorems appears in Appendix B.7. Since our IOP constructions use finite fields, prover and verifier complexity are most naturally stated using arithmetic complexity over the ambient field, the size of which is derived from the size of the instance x (see Remark B.1); we use $\text{tp}^{\mathbb{F}}$ and $\text{tv}^{\mathbb{F}}$ to denote arithmetic complexity, assuming the field \mathbb{F} is understood from context.

Our first theorem is quite efficient when applied to space bounded computations, like our DPM, and indeed our implementation realizes the IOP described in the theorem (cf. Appendix B.8). Although we use asymptotic notation below, the algebraic version of this theorem (Lemma B.6), involves only explicit constants, which may be useful in the future for estimating explicit proof parameters and for asymptotic comparisons with other works.

Recall that $\text{NTimeSpace}(T(n), S(n))$ is the class of nondeterministic languages that are decidable in simultaneous time $T(n)$ and space $S(n)$.

Theorem 3.4 (ZK-STIK for space bounded computations). *Let L be a language in $\text{NTimeSpace}(T(n), S(n))$, $T(n) \geq n$ and let R be induced by L . Then R has a transparent witness indistinguishable IOP of knowledge with the following parameters, stated for soundness error function $\text{err}(n) = 2^{-\lambda(n)}$*

- perfect completeness and soundness error at most $\text{err}(n)$ for instances of size n
- knowledge error bound $\text{err}'(n) = O(\text{err}(n))$
- round complexity $r(n) = \frac{\log T(n)}{2} + O(1)$
- query complexity $q(n) = 36(\lambda + 2) \cdot (\log T(n) + S(n) + O(1))$, each query is an element of a binary field \mathbb{F} , $|\mathbb{F}| = 2^n$ for $n = \lambda + \log T(n) + O(1)$
- verifier arithmetic complexity $\text{tv}^{\mathbb{F}}(n) = 2n^2 + O(\lambda \cdot (S(n) + \log T(n)))$
- prover arithmetic complexity $\text{tp}^{\mathbb{F}}(n) = O(S(n) \cdot T(n) \cdot \log T(n))$
- proof length $O(T(n) \cdot S(n))$, measured in field elements.

In particular, for $S(n) = \text{poly log } T(n)$, this IOP is fully scalable, i.e., the system is a wi-STIK; if, additionally, $T(n) = \text{poly}(n)$, then the system has perfect ZK, i.e., it is a ZK-STIK.

For computations with super-poly-logarithmic space the theorem above is not scalable, neither for prover nor for verifier. The following theorem is fully scalable for any nondeterministic language, i.e., it can be said to be a *universal* wi-STIK.

Theorem 3.5 (wi-STIK for NEXP). *Let $L \in \text{NTIME}(T(n))$, $T(n) \geq n$ and R be induced by L . Then R has a witness-indistinguishable, fully scalable, and transparent IOP of knowledge (wi-STIK) with the following parameters, stated for soundness error function $\text{err}(n) = 2^{-\lambda(n)}$*

- perfect completeness and soundness error $\text{err}(n) \leq 2^{-\lambda(n)}$ for instances of size n
- knowledge extraction bound $\text{err}'(n) = O(\text{err}(n))$
- round complexity $r(n) = \frac{\log T(n)}{2} + O(1)$

- *query complexity* $O(\lambda(n) \cdot \log T(n))$, each query is an element of a binary field \mathbb{F} , $|\mathbb{F}| = 2^n$ for $n = \lambda(n) + \log T(n) + \log \log T(n) + O(1)$.
- *verifier arithmetic complexity* $\text{tv}^{\mathbb{F}}(n) = O(\lambda(n) \cdot \log T(n))$,
- *prover arithmetic complexity* $\text{tp}^{\mathbb{F}}(n) = O(T(n) \log^2 T(n))$,
- *proof length* $O(T(n) \log T(n))$, measured in field elements.

For $T(n) = \text{poly}(n)$ the system has perfect ZK, i.e., it is a ZK-STIK.

3.3 Scalable Transparent ARgument of Knowledge (STARK) as a realization of STIK

Definition 3.3 refers to the IOP model, in which results can be proved with no cryptographic assumptions. Indeed, most of our contributions, described in following sections (like the FRI protocol), are stated and studied in this “clean” IOP model; and a majority of our engineering work was dedicated to implementing IOP-based algorithms of a STIK system. However, we are not aware of any unconditionally secure IOP realization that is scalable, and theoretical works show that such constructions are unlikely to emerge [54]. A number of fundamental transformations have been suggested in the past to realize PCP systems using various cryptographic assumptions, and these transformations were adapted to the IOP model [22]. In all such realizations the prover must be computationally bounded, and such systems are commonly called *argument systems*, and, consequently, the realization of a STIK results in a *Scalable Transparent ARgument of Knowledge* (STARK).

The two main transformations of proof systems into realizable argument systems are:

- **Interactive STARK (iSTARK)** As shown by Kilian [74] for the PCP model, a family of collision-resistant hash functions can be used to convert a STIK into an interactive argument of knowledge system; if the STIK has perfect ZK, then the argument system has computational ZK. Any realization of a STIK using this technique will be called an *interactive STARK* (iSTARK); when one wants to emphasize that the STIK is zero knowledge, the term ZK-iSTARK will be used.
- **Non-interactive STARK (nSTARK)** As shown by Micali [85] and Valiant [105] for the PCP model, and by Ben-Sasson et al. [22] for the IOP model, any STIK can be compiled into a non-interactive argument of knowledge in the random oracle model (called a *non-interactive random-oracle proof* (NIROP) there); if the STIK had perfect zero knowledge then the resulting construction has computational zero knowledge. Any realization of a STIK using this technique will be called an *non-interactive STARK* (nSTARK); when one wants to emphasize that the STIK is zero knowledge, the term ZK-nSTARK will be used.

While non-interactive STARKs have the advantage of being comprised of a single message from the prover, they also rely on stronger assumptions. In certain settings the public may view certain blockchains (like Bitcoin’s) as a realization of both (i) an immutable public time-stamping service *and* (ii) a public beacon of randomness. Under this view, the blockchain can be used to emulate the verifier on an iSTARK system, resulting in smaller communication complexity under better (more standard) cryptographic assumptions. Thus, we leave the choice of which particular realization mode to use for a (ZK)-STIK—(ZK)-iSTARK vs. (ZK)-nSTARK—to be made by system designers based on particular use cases, and refer to both realization modes of a STIK as a STARK; to emphasize the ZK aspect of the STIK we may refer to the realization as a ZK-STARK.

3.4 Prior STIK and STARK constructions

The acronyms STIK and (ZK-)STARK may be new to this work, but IOP systems obtaining the properties that define them have been described in the past, as discussed next.

STIK PCP systems are, by definition, transparent (1-round) IOP systems. The first such system with a scalable verifier was given in the works²⁴ of Babai et al. [7, 6] and the first fully scalable PCP, i.e., the first STIK construction, appears in the works²⁵ of Ben-Sasson et al. [25, 20]. The first ZK-STIK for NP appears in the work of Ben-Sasson et al. [17], later extended to a ZK-STIK for NEXP [15].

STARK The first realization of a STIK system, i.e., the first STARK, appears in the recent work of Ben-Sasson et al. [13]; our current publication describes the *first realization of a ZK-STIK and is therefore the first ZK-STARK construction*. (See Section 1.3 for other ZK solutions).

Acknowledgements

We thank Eli Eliezer and Tal Shoenwald for providing access to the server used to measure prover performance described in Appendix A.0.1. We thank Arie Tal, Yechiel Kimchi and Gala Yadgar for help optimizing code performance. We thank the Andrea Cerulli, Venkatasubramaniam, Muthuramakrishnan, Madars Virza, and the other authors of [31, 1] for assistance in obtaining the data reported in Figure 3. We thank Shmuel (Muli) Ben-Sasson, Yuval Ishai and Uri Kolodny for commenting on an earlier draft.

²⁴The first work [7] shows this for NEXP and the second [6] scales it down to NP.

²⁵The first work [25] presents a PCP with scalable verification and quasi-linear *proof length*, the second work [20] bounds the prover running time and also proves the proof of knowledge property.

References

- [1] Scott Ames, Carmit Hazay, Yuval Ishai, and Muthuramakrishnan Venkatasubramanian. Liger: Lightweight sublinear arguments without a trusted setup. In *Proceedings of the 24th ACM Conference on Computer and Communications Security*, October 2017.
- [2] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *Journal of the ACM*, 45(3):501–555, 1998. Preliminary version in FOCS '92.
- [3] Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs: a new characterization of NP. *Journal of the ACM*, 45(1):70–122, 1998. Preliminary version in FOCS '92.
- [4] László Babai. Trading group theory for randomness. In *Proceedings of the 17th Annual ACM Symposium on Theory of Computing*, STOC '85, pages 421–429, 1985.
- [5] László Babai and Lance Fortnow. Arithmetization: A new method in structural complexity theory. *computational complexity*, 1(1):41–66, 1991.
- [6] László Babai, Lance Fortnow, Leonid A. Levin, and Mario Szegedy. Checking computations in polylogarithmic time. In *Proceedings of the 23rd Annual ACM Symposium on Theory of Computing*, STOC '91, pages 21–32, 1991.
- [7] László Babai, Lance Fortnow, and Carsten Lund. Nondeterministic exponential time has two-prover interactive protocols. In *Proceedings of the 31st Annual Symposium on Foundations of Computer Science*, FOCS '90, pages 16–25, 1990.
- [8] Boaz Barak and Oded Goldreich. Universal arguments and their applications. *SIAM Journal on Computing*, 38(5):1661–1694, 2008. Preliminary version appeared in CCC '02.
- [9] Mihir Bellare and Oded Goldreich. On defining proofs of knowledge. In *Proceedings of the 12th Annual International Cryptology Conference on Advances in Cryptology*, CRYPTO '92, pages 390–420, 1993.
- [10] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *Proceedings of the 1st ACM Conference on Computer and Communications Security*, CCS '93, pages 62–73, New York, NY, USA, 1993. ACM.
- [11] Michael Ben-Or, Oded Goldreich, Shafi Goldwasser, Johan Hr astad, Joe Kilian, Silvio Micali, and Phillip Rogaway. Everything provable is provable in zero-knowledge. In *Proceedings of the 8th Annual International Cryptology Conference*, CRYPTO '89, pages 37–56, 1988.
- [12] Eli Ben-Sasson. Universal and affordable computational integrity. <https://www.youtube.com/watch?v=Q4nWoEKUtgU&t=32s>, May 2013.
- [13] Eli Ben-Sasson, Iddo Bentov, Alessandro Chiesa, Ariel Gabizon, Daniel Genkin, Matan Hamilis, Evgenya Pergament, Michael Riabzev, Mark Silberstein, Eran Tromer, and Madars Virza. Computational integrity with a public random string from quasi-linear PCPs. *IACR Cryptology ePrint Archive*, 2016:646, 2016.
- [14] Eli Ben-Sasson, Iddo Bentov, Ynon Horesh, and Michael Riabzev. Fast Reed-Solomon interactive oracle proofs of proximity (2nd revision). *Electronic Colloquium on Computational Complexity (ECCC)*, 24:134, 2017.
- [15] Eli Ben-Sasson, Alessandro Chiesa, Michael A. Forbes, Ariel Gabizon, Michael Riabzev, and Nicholas Spooner. On probabilistic checking in perfect zero knowledge. *Electronic Colloquium on Computational Complexity (ECCC)*, 23:156, 2016.
- [16] Eli Ben-Sasson, Alessandro Chiesa, Ariel Gabizon, Michael Riabzev, and Nicholas Spooner. Short interactive oracle proofs with constant query complexity, via composition and sumcheck. *Electronic Colloquium on Computational Complexity (ECCC)*, 23:46, 2016.

- [17] Eli Ben-Sasson, Alessandro Chiesa, Ariel Gabizon, and Madars Virza. Quasilinear-size zero knowledge from linear-algebraic PCPs. In *Proceedings of the 13th Theory of Cryptography Conference*, TCC '16, pages 33–64, 2016.
- [18] Eli Ben-Sasson, Alessandro Chiesa, Christina Garman, Matthew Green, Ian Miers, Eran Tromer, and Madars Virza. Zerocash: Decentralized anonymous payments from Bitcoin. In *Proceedings of the 2014 IEEE Symposium on Security and Privacy*, SP '14, 2014.
- [19] Eli Ben-Sasson, Alessandro Chiesa, Daniel Genkin, and Eran Tromer. Fast reductions from RAMs to delegatable succinct constraint satisfaction problems. In *Proceedings of the 4th Innovations in Theoretical Computer Science Conference*, ITCS '13, pages 401–414, 2013.
- [20] Eli Ben-Sasson, Alessandro Chiesa, Daniel Genkin, and Eran Tromer. On the concrete efficiency of probabilistically-checkable proofs. In *Proceedings of the 45th ACM Symposium on the Theory of Computing*, STOC '13, pages 585–594, 2013.
- [21] Eli Ben-Sasson, Alessandro Chiesa, Daniel Genkin, Eran Tromer, and Madars Virza. SNARKs for C: Verifying program executions succinctly and in zero knowledge. In *Proceedings of the 33rd Annual International Cryptology Conference*, CRYPTO '13, pages 90–108, 2013.
- [22] Eli Ben-Sasson, Alessandro Chiesa, and Nicholas Spooner. *Interactive Oracle Proofs*, pages 31–60. Springer Berlin Heidelberg, Berlin, Heidelberg, 2016.
- [23] Eli Ben-Sasson, Alessandro Chiesa, Eran Tromer, and Madars Virza. Scalable zero knowledge via cycles of elliptic curves. In *Proceedings of the 34th Annual International Cryptology Conference*, CRYPTO '14, pages 276–294, 2014. Extended version at <http://eprint.iacr.org/2014/595>.
- [24] Eli Ben-Sasson, Alessandro Chiesa, Eran Tromer, and Madars Virza. Succinct non-interactive zero knowledge for a von Neumann architecture. In *Proceedings of the 23rd USENIX Security Symposium*, Security '14, pages 781–796, 2014. Extended version at <http://eprint.iacr.org/2013/879>.
- [25] Eli Ben-Sasson, Oded Goldreich, Prahladh Harsha, Madhu Sudan, and Salil Vadhan. Short PCPs verifiable in polylogarithmic time. In *Proceedings of the 20th Annual IEEE Conference on Computational Complexity*, CCC '05, pages 120–134, 2005.
- [26] Eli Ben-Sasson, Oded Goldreich, Prahladh Harsha, Madhu Sudan, and Salil P. Vadhan. Robust PCPs of proximity, shorter PCPs, and applications to coding. *SIAM Journal on Computing*, 36(4):889–974, 2006.
- [27] Eli Ben-Sasson and Madhu Sudan. Short PCPs with polylog query complexity. *SIAM Journal on Computing*, 38(2):551–607, 2008. Preliminary version appeared in STOC '05.
- [28] Václav E Beneš. *Mathematical theory of connecting networks and telephone traffic*, volume 17. Academic press, 1965.
- [29] Nir Bitansky, Ran Canetti, Alessandro Chiesa, and Eran Tromer. Recursive composition and bootstrapping for SNARKs and proof-carrying data. In *Proceedings of the 45th ACM Symposium on the Theory of Computing*, STOC '13, pages 111–120, 2013.
- [30] Nir Bitansky, Alessandro Chiesa, Yuval Ishai, Rafail Ostrovsky, and Omer Paneth. Succinct non-interactive arguments via linear interactive proofs. In *Proceedings of the 10th Theory of Cryptography Conference*, TCC '13, pages 315–333, 2013.
- [31] Jonathan Bootle, Andrea Cerulli, Pyrros Chaidos, Jens Groth, and Christophe Petit. Efficient zero-knowledge arguments for arithmetic circuits in the discrete log setting. In *Advances in Cryptology - EUROCRYPT 2016 - 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Vienna, Austria, May 8-12, 2016, Proceedings, Part II*, pages 327–357, 2016.

- [32] Joppe W. Bos, Onur Özen, and Martijn Stam. Efficient hashing using the aes instruction set. In Bart Preneel and Tsuyoshi Takagi, editors, *Cryptographic Hardware and Embedded Systems - CHES 2011 - 13th International Workshop, Nara, Japan, September 28 - October 1, 2011. Proceedings*, volume 6917 of *Lecture Notes in Computer Science*, pages 507–522. Springer, 2011.
- [33] Gilles Brassard, David Chaum, and Claude Crépeau. Minimum disclosure proofs of knowledge. *Journal of Computer and System Sciences*, 37(2):156–189, 1988.
- [34] Arthur Breitman. Scaling tezos. <https://hackernoon.com/scaling-tezo-8de241dd91bd>, 2017.
- [35] Benedikt Bünz, Jonathan Bootle, Dan Boneh, Andrew Poelstra, Pieter Wuille, and Greg Maxwell. Bulletproofs: Efficient range proofs for confidential transactions. Cryptology ePrint Archive, Report 2017/1066, 2017. <https://eprint.iacr.org/2017/1066>.
- [36] David G. Cantor. On arithmetical algorithms over finite fields. *J. Comb. Theory, Ser. A*, 50(2):285–300, 1989.
- [37] Lily Chen, Lily Chen, Stephen Jordan, Yi-Kai Liu, Dustin Moody, Rene Peralta, Ray Perlner, and Daniel Smith-Tone. *Report on post-quantum cryptography*. US Department of Commerce, National Institute of Standards and Technology, 2016.
- [38] Alessandro Chiesa and Zeyuan Allen Zhu. Shorter arithmetization of nondeterministic computations. *Theor. Comput. Sci.*, 600:107–131, 2015.
- [39] Alessandro Chiesa and Eran Tromer. Proof-carrying data and hearsay arguments from signature cards. In *Proceedings of the 1st Symposium on Innovations in Computer Science, ICS '10*, pages 310–331, 2010.
- [40] James W Cooley and John W Tukey. An algorithm for the machine calculation of complex fourier series. *Mathematics of computation*, 19(90):297–301, 1965.
- [41] Graham Cormode, Michael Mitzenmacher, and Justin Thaler. Practical verified computation with streaming interactive proofs. In *Proceedings of the 4th Symposium on Innovations in Theoretical Computer Science, ITCS '12*, pages 90–112, 2012.
- [42] Graham Cormode, Justin Thaler, and Ke Yi. Verifying computations with streaming interactive proofs. *Proceedings of the VLDB Endowment*, 5(1):25–36, 2011.
- [43] R. Courtland. Google aims for quantum computing supremacy [news]. *IEEE Spectrum*, 54(6):9–10, June 2017.
- [44] Joan Daemen and Vincent Rijmen. Aes proposal: Rijndael, 1999.
- [45] Ivan Bjerre Damgård. A design principle for hash functions. In G. Brassard, editor, *Advances in Cryptology—CRYPTO '89*, volume 435 of *Lecture Notes in Computer Science*, pages 416–427. Springer-Verlag, 1989.
- [46] George Danezis, Cédric Fournet, Jens Groth, and Markulf Kohlweiss. *Square Span Programs with Applications to Succinct NIZK Arguments*, pages 532–550. Springer Berlin Heidelberg, Berlin, Heidelberg, 2014.
- [47] NG de Bruijn. A combinatorial problem. *Proc. Akademie Van Westerschappen*, 49(2):758–764, 1946.
- [48] Irit Dinur and Omer Reingold. Assignment testers: Towards a combinatorial proof of the PCP theorem. In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science, FOCS '04*, pages 155–164, 2004.
- [49] Cynthia Dwork, Uriel Feige, Joe Kilian, Moni Naor, and Shmuel Safra. Low communication 2-prover zero-knowledge proofs for NP. In *Proceedings of the 11th Annual International Cryptology Conference, CRYPTO '92*, pages 215–227, 1992.
- [50] Amos Fiat and Adi Shamir. How to prove yourself: practical solutions to identification and signature problems. In *Proceedings of the 6th Annual International Cryptology Conference, CRYPTO '86*, pages 186–194, 1986.
- [51] Shuhong Gao and Todd D. Mateer. Additive fast fourier transforms over finite fields. *IEEE Trans. Information Theory*, 56(12):6265–6272, 2010.

- [52] Rosario Gennaro, Craig Gentry, and Bryan Parno. Non-interactive verifiable computing: Outsourcing computation to untrusted workers. In *Proceedings of the 30th Annual Conference on Advances in Cryptology, CRYPTO'10*, pages 465–482, Berlin, Heidelberg, 2010. Springer-Verlag.
- [53] Rosario Gennaro, Craig Gentry, Bryan Parno, and Mariana Raykova. Quadratic span programs and succinct NIZKs without PCPs. In *Proceedings of the 32nd Annual International Conference on Theory and Application of Cryptographic Techniques, EUROCRYPT '13*, pages 626–645, 2013.
- [54] Craig Gentry and Daniel Wichs. Separating succinct non-interactive arguments from all falsifiable assumptions. In *Proceedings of the Forty-third Annual ACM Symposium on Theory of Computing, STOC '11*, pages 99–108, New York, NY, USA, 2011. ACM.
- [55] Irene Giacomelli, Jesper Madsen, and Claudio Orlandi. Zkboo: Faster zero-knowledge for boolean circuits. In *25th USENIX Security Symposium (USENIX Security 16)*, pages 1069–1083, Austin, TX, 2016. USENIX Association.
- [56] K. Gödel. Über formal unentscheidbare Sätze der Principia Mathematica und verwandter Systeme I. *Monatshefte für Mathematik*, 38(1):173–198, 1931.
- [57] Oded Goldreich. Randomness and computation. In Oded Goldreich, editor, *Studies in Complexity and Cryptography. Miscellanea on the Interplay between Randomness and Computation - In Collaboration with Lidor Avigad, Mihir Bellare, Zvika Brakerski, Shafi Goldwasser, Shai Halevi, Tali Kaufman, Leonid Levin, Noam Nisan, Dana Ron, Madhu Sudan, Luca Trevisan, Salil Vadhan, Avi Wigderson, David Zuckerman*, volume 6650 of *Lecture Notes in Computer Science*, pages 507–539. Springer, 2011.
- [58] Shafi Goldwasser, Yael Tauman Kalai, and Guy N. Rothblum. Delegating computation: Interactive proofs for Muggles. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing, STOC '08*, pages 113–122, 2008.
- [59] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, 18(1):186–208, 1989. Preliminary version appeared in STOC '85.
- [60] Jens Groth. Short pairing-based non-interactive zero-knowledge arguments. In *Proceedings of the 16th International Conference on the Theory and Application of Cryptology and Information Security, ASIACRYPT '10*, pages 321–340, 2010.
- [61] Jens Groth. Efficient zero-knowledge arguments from two-tiered homomorphic commitments. In *Advances in Cryptology - ASIACRYPT 2011 - 17th International Conference on the Theory and Application of Cryptology and Information Security, Seoul, South Korea, December 4-8, 2011. Proceedings*, pages 431–448, 2011.
- [62] Jens Groth. On the size of pairing-based non-interactive arguments. In *Advances in Cryptology - EUROCRYPT 2016 - 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Vienna, Austria, May 8-12, 2016, Proceedings, Part II*, pages 305–326, 2016.
- [63] Jens Groth and Mary Maller. Snarky signatures: Minimal signatures of knowledge from simulation-extractable snarks. In *Advances in Cryptology - CRYPTO 2017 - 37th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 20-24, 2017, Proceedings, Part II*, pages 581–612, 2017.
- [64] Jens Groth and Amit Sahai. Efficient non-interactive proof systems for bilinear groups. In *Advances in Cryptology - EUROCRYPT 2008, 27th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Istanbul, Turkey, April 13-17, 2008. Proceedings*, pages 415–432, 2008.
- [65] Venkatesan Guruswami and Madhu Sudan. Improved decoding of reed-solomon and algebraic-geometry codes. *IEEE Trans. Information Theory*, 45(6):1757–1767, 1999.
- [66] Michael Hearn. Corda and SGX: a privacy update. <https://www.corda.net/2017/06/corda-sgx-privacy-update/>, 2017.
- [67] Daira Hopwood, Sean Bowe, Taylor Hornby, and Nathan Wilcox, March 2017.

- [68] Rosamond Hutt. Beyond bitcoin: 4 surprising uses for blockchain, Dec 2016.
- [69] Yuval Ishai, Eyal Kushilevitz, and Rafail Ostrovsky. Efficient arguments without short PCPs. In *Proceedings of the Twenty-Second Annual IEEE Conference on Computational Complexity*, CCC '07, pages 278–291, 2007.
- [70] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Zero-knowledge from secure multiparty computation. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 21–30. ACM, 2007.
- [71] Yuval Ishai, Mohammad Mahmoody, Amit Sahai, and David Xiao. On zero-knowledge PCPs: Limitations, simplifications, and applications, 2015. Available at <http://www.cs.virginia.edu/~mohammad/files/papers/ZKPCPs-Full.pdf>.
- [72] Yael Kalai and Ran Raz. Interactive PCP. In *Proceedings of the 35th International Colloquium on Automata, Languages and Programming*, ICALP '08, pages 536–547, 2008.
- [73] Jonathan Katz and Yehuda Lindell. *Introduction to Modern Cryptography, Second Edition*. CRC Press, 2014.
- [74] Joe Kilian. A note on efficient zero-knowledge proofs and arguments. In *Proceedings of the 24th Annual ACM Symposium on Theory of Computing*, STOC '92, pages 723–732, 1992.
- [75] Joe Kilian, Erez Petrank, and Gábor Tardos. Probabilistically checkable proofs with zero knowledge. In *Proceedings of the 29th Annual ACM Symposium on Theory of Computing*, STOC '97, pages 496–505, 1997.
- [76] Protocol Labs. Filecoin: A decentralized storage network. <https://filecoin.io/filecoin.pdf>, 2017.
- [77] F. Thomson Leighton. *Introduction to Parallel Algorithms and Architectures: Array, Trees, Hypercubes*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1992.
- [78] Rudolf Lidl and Harald Niederreiter. *Finite Fields*. Cambridge University Press, second edition edition, 1997.
- [79] Sian-Jheng Lin, Wei-Ho Chung, and Yunghsiang S. Han. Novel polynomial basis and its application to reed-solomon erasure codes. In *Proceedings of the 2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, FOCS '14, pages 316–325, Washington, DC, USA, 2014. IEEE Computer Society.
- [80] UV Linnik. On the least prime in an arithmetic progression. i. the basic theorem. *Matematicheskii Sbornik*, 15(2):139–178, 1944.
- [81] Helger Lipmaa. Progression-free sets and sublinear pairing-based non-interactive zero-knowledge arguments. In *Proceedings of the 9th Theory of Cryptography Conference on Theory of Cryptography*, TCC '12, pages 169–189, 2012.
- [82] Carsten Lund, Lance Fortnow, Howard J. Karloff, and Noam Nisan. Algebraic methods for interactive proof systems. *Journal of the ACM*, 39(4):859–868, 1992.
- [83] Ralph C. Merkle. A certified digital signature. In *Advances in Cryptology (CRYPTO '89)*, pages 218–238, Berlin - Heidelberg - New York, 1990. Springer.
- [84] S. Micali, M. Rabin, and J. Kilian. Zero-knowledge sets. In *44th Annual IEEE Symposium on Foundations of Computer Science, 2003. Proceedings.*, pages 80–91, Oct 2003.
- [85] Silvio Micali. Computationally sound proofs. *SIAM J. Comput.*, 30(4):1253–1298, 2000.
- [86] Satoshi Nakamoto. Bitcoin: a peer-to-peer electronic cash system, 2009.
- [87] Federal Bureau of Investigation. Combined dna index system (CODIS). <https://www.fbi.gov/services/laboratory/biometric-analysis/codis>. Retrieved September 2017.
- [88] Brian Parno, Craig Gentry, Jon Howell, and Mariana Raykova. Pinocchio: Nearly practical verifiable computation. In *Proceedings of the 34th IEEE Symposium on Security and Privacy*, Oakland '13, pages 238–252, 2013.

- [89] M. Peck. A blockchain currency that beats bitcoin on privacy [news]. *IEEE Spectrum*, 53(12):11–13, December 2016.
- [90] Evgenya Pergament. Algebraic RAM. Master’s thesis, Technion — Israel Institute of Technology, 2017.
- [91] Alexander Polishchuk and Daniel A. Spielman. Nearly-linear size holographic proofs. In *Proceedings of the 26th Annual ACM Symposium on Theory of Computing*, STOC ’94, pages 194–203, 1994.
- [92] FIPS PUB. Secure hash standard (shs). *FIPS PUB 180*, 4, 2012.
- [93] Alexander A Razborov. Lower bounds on the size of bounded depth circuits over a complete basis with logical addition. *Mathematical Notes of the Academy of Sciences of the USSR*, 41(4):333–338, 1987.
- [94] Omer Reingold, Guy N. Rothblum, and Ron D. Rothblum. Constant-round interactive proofs for delegating computation. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pages 49–62, 2016.
- [95] Guy N Rothblum, Salil Vadhan, and Avi Wigderson. Interactive proofs of proximity: delegating computation in sublinear time. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 793–802. ACM, 2013.
- [96] SCIPR Lab. libsnark: a C++ library for zkSNARK proofs. <https://github.com/scipr-lab/libsnark>.
- [97] Jae Hong Seo. Round-efficient sub-linear zero-knowledge arguments for linear algebra. In *Public Key Cryptography - PKC 2011 - 14th International Conference on Practice and Theory in Public Key Cryptography, Taormina, Italy, March 6-9, 2011. Proceedings*, pages 387–402, 2011.
- [98] Srinath Setty, Andrew J. Blumberg, and Michael Walfish. Toward practical and unconditional verification of remote computations. In *Proceedings of the 13th USENIX Conference on Hot Topics in Operating Systems, HotOS ’11*, pages 29–29, 2011.
- [99] Srinath Setty, Benjamin Braun, Victor Vu, Andrew J. Blumberg, Bryan Parno, and Michael Walfish. Resolving the conflict between generality and plausibility in verified computation. In *Proceedings of the 8th EuroSys Conference, EuroSys ’13*, pages 71–84, 2013.
- [100] Srinath Setty, Michael McPherson, Andrew J. Blumberg, and Michael Walfish. Making argument systems for outsourced computation practical (sometimes). In *Proceedings of the 2012 Network and Distributed System Security Symposium, NDSS ’12*, 2012.
- [101] Srinath Setty, Victor Vu, Nikhil Panpalia, Benjamin Braun, Andrew J. Blumberg, and Michael Walfish. Taking proof-based verified computation a few steps closer to practicality. In *Proceedings of the 21st USENIX Security Symposium, Security ’12*, pages 253–268, 2012.
- [102] Roman Smolensky. Algebraic methods in the theory of lower bounds for boolean circuit complexity. In *Proceedings of the nineteenth annual ACM symposium on Theory of computing*, pages 77–82. ACM, 1987.
- [103] Nick Szabo. Formalizing and securing relationships on public networks. *First Monday*, 2(9), September 1997. Retrieved October 2017.
- [104] Justin Thaler. Time-optimal interactive proofs for circuit evaluation. In *Proceedings of the 33rd Annual International Cryptology Conference, CRYPTO ’13*, pages 71–89, 2013.
- [105] Paul Valiant. Incrementally verifiable computation or proofs of knowledge imply time/space efficiency. In *Proceedings of the 5th Conference on Theory of Cryptography, TCC’08*, pages 1–18, Berlin, Heidelberg, 2008. Springer-Verlag.
- [106] Joachim von zur Gathen and Jürgen Gerhard. Fast algorithms for taylor shifts and certain difference equations. In *Proceedings of the 1997 International Symposium on Symbolic and Algebraic Computation, ISSAC ’97, Maui, Hawaii, USA, July 21-23, 1997*, pages 40–47, 1997.

- [107] Victor Vu, Srinath Setty, Andrew J. Blumberg, and Michael Walfish. A hybrid architecture for interactive verifiable computation. In *Proceedings of the 34th IEEE Symposium on Security and Privacy*, Oakland '13, pages 223–237, 2013.
- [108] Riad S. Wahby, Srinath T. V. Setty, Zuocheng Ren, Andrew J. Blumberg, and Michael Walfish. Efficient RAM and control flow in verifiable outsourced computation. In *22nd Annual Network and Distributed System Security Symposium, NDSS 2015, San Diego, California, USA, February 8-11, 2014*, 2015.
- [109] Riad S. Wahby, Ioanna Tzialla, abhi shelat, Justin Thaler, and Michael Walfish. Doubly-efficient zkSNARKs without trusted setup. Cryptology ePrint Archive, Report 2017/1132, 2017. <https://eprint.iacr.org/2017/1132>.
- [110] Michael Walfish and Andrew J. Blumberg. Verifying computations without reexecuting them. *Commun. ACM*, 58(2):74–84, 2015.
- [111] R. S. Winternitz. A secure one-way hash function built from DES. In *1984 IEEE Symposium on Security and Privacy*, pages 88–88, April 1984.
- [112] Yupeng Zhang, Daniel Genkin, Jonathan Katz, Dimitrios Papadopoulos, and Charalampos Papamanthou. A zero-knowledge version of VSQL. Cryptology ePrint Archive, Report 2017/1146, 2017. <https://eprint.iacr.org/2017/1146>.

A Measurements of the ZK-STARK for the DNA profile match

In this section we provide additional raw measurements for proving complexity and verification complexity for the DPM discussed in Section 1. We used two separate machines to measure performance: a strong server for the prover, and a “standard” laptop for the verifier.

A.0.1 Prover

Figure 6 shows the running time required to generate the ZK-STARK proofs, both in absolute terms (left) and as a multiplicative overhead in running time over naïve computation (middle). On the right we plot the size occupied by the all IOP oracles and Merkle trees used by the prover. Due to space limitations (768GB of RAM, henceforth called the “RAM threshold”), the actual space used by the prover was lower than plotted there but saving space required larger running time. Specifically, notice that at $n = 2^{14} \approx 16,000$ the total space (plotted on right) passes the RAM threshold, corresponding to (and explaining) the jump in proving time which is noticed on the middle plot at the same value of n . This jump is due to re-computing parts of the proof oracles on demand, which is required to operate within RAM limits.

The code of the prover (written in C++) has been optimized for large instances and running times, hence it uses multi-threading (MT) intensively. Our use of MT seems empirically quite efficient; in particular, disabling MT incurs a slowdown factor close to $\times 2$. The relatively large multiplicative overhead noticeable on the middle plot for small instance sizes is likely explained by the overhead that the use of MT introduces. However, since prover execution time is measured in fractions of a second for these short executions we leave further optimizations of it to future work.

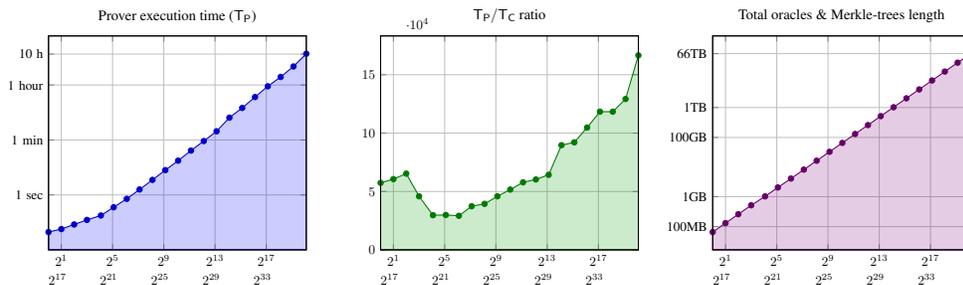


Figure 6: On the left we plot, on a double-logarithmic scale, prover running time as a function of the number of entries (n) in the DNA profile database. On the middle we plot the ratio between proving time to naïve execution time; the horizontal axis is logarithmic ($\log n$) and the vertical one measures ratio. On the right we plot, on a double-logarithmic scale, the total size of all oracles and their commitment trees (Merkle trees), generated by the prover during execution. See text inline for an explanation for the “phase transitions” seen in the middle plot.

Prover machine specifications

- CPU (2 units) : Intel(R) Xeon(R) Platinum 8168 CPU @ 2.70GHz (24 cores, 2 threads per core)
- RAM : 768GB DDR4 (32GB \times 24, Speed: 2666 MHz)
- SWAP : 3.2TB NVMe (1.6TB \times 2)
- Operating System : Red Hat Enterprise Linux 7 (3.10.0-693.5.2.el7.x86_64)

A.0.2 Verifier

Figure 7 gives the ZK-STARK verifier running time (T_V) on the left, and communication complexity on the right. The verifier is non-adaptive, which means it’s complexity can be measured even for databases sizes n that are too large to generate a proof for. The values for which an actual proof was generated are indicated by full circles in both plots.

The ZK-STARK verifier is comprised of two sub-verifiers. The first is the ZK-STIK verifier that verifies proofs in the “pure” but unrealistic IOP model. The second is the sub-verifier that checks (only) consistency of values residing in various Merkle trees with previously committed Merkle tree roots (see Section 2.5). In both plots of Figure 7 the bottom line indicates the complexity of the ZK-STIK verifier, both for time (left) and communication complexity (right). As evident from the Figure, the ZK-STIK complexity is small relative to overall complexity. Moreover, as oracle size increases, the ratio of STIK/STARK complexity grows smaller. This is because as oracles grow larger, the relevant Merkle trees grow deeper and hence there are more authentication paths, and each is of larger length.

We executed the verifier in single thread mode; the tests run by it are amenable to parallelization and faster execution time. However, since verification time is already quite small we leave these further optimizations to future work. Similarly, we point out that the additional memory consumption required by the verifier is negligible, compared to the communication-complexity. In particular, when the verifier is executed on weaker machines than the one reported here (see specificatin below), verification complexity does not increase significantly.

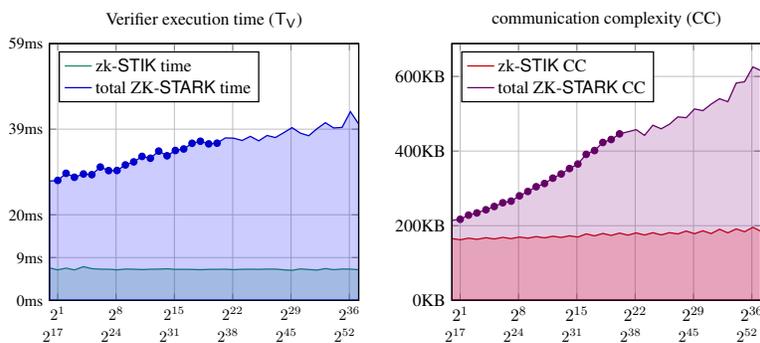


Figure 7: Verifier running time (left) and communication complexity (right) as a function of the number of entries (n) in the DNA database size; for both plots the horizontal axis is logarithmic and the vertical one is linear. In both plots, the lower line measure the complexity of the ZK-STIK verifier, and the upper line measures that of the ZK-STARK verifier. We stress that verifier measurements are performed for values of n that are (significantly) larger than those for which a proof can be generated; this is possible because our verifier is non-adaptive, thus, we tested it with randomly generated “proofs”. Full circles indicate values of n for which proofs were generated.

Verifier machine specifications

- Model : Lenovo ThinkPad W530 Laptop
- CPU : Intel(R) Core(TM) i7-3740QM CPU @ 2.70GHz (4 cores, 2 threads per core)
- RAM : 32GB DDR3 (8GB \times 4, Speed: 1600 MT/s)
- Operating System : Arch Linux (4.13.12-1-ARCH)

B From AIR to ZK-STARK

The purpose of this section is to specify in detail our IOP constructions, expanding on Section 2. The existence of a ZK-STIK system for NEXP was already established in [17, 15]; thus, our focus is on concrete efficiency and on a detailed description of the construction realized in code.

For the purposes of the current discussion, an *instance* of a computational integrity statement, denoted \mathfrak{x} , is specified by (i) a *transition relation* over a space of *machines states* and (ii) a set of *boundary constraints* (like inputs and outputs). A *witness* to the integrity of \mathfrak{x} is a *valid execution* of the computation, given by an *execution trace* — a sequence of machine states that adheres to both boundary and transition constraints of the computation. Casting CI statements like (**) in this format is a straightforward application of the Cook-Levin Theorem.

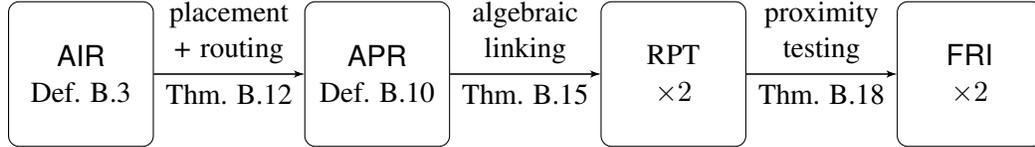


Figure 8: The reduction from an AIR instance to a pair of RPT problems, solved using the FRI protocol.

Overview Our process has 4 parts (see Figure 8):

1. The starting point is a natural *algebraic intermediate representation* (AIR) of \mathfrak{x} and w , denoted $\mathfrak{x}_{\text{AIR}}, w_{\text{AIR}}$. By “algebraic” we mean that states of the execution trace are represented by sequences of elements in a finite field \mathbb{F} , and the transition relation is specified by a set of polynomials over variables representing the “current” and “next” step of the execution trace.
2. We reduce the AIR representation into a different one, in which states of the execution trace are “placed” on nodes of an *affine graph*, so that consecutive states are connected by an edge in that graph. Informally, an affine graph is a “circuit” that has “algebraic” topology (see Appendix B.3). The process of “placing” machine states on nodes of a circuit resembles the process of *placement and routing* which is commonly used in computer and circuit design, although our design space is constrained by *algebra* rather than by physical reality. We refer to the transformation as the *algebraic placement and routing* (APR) reduction, and the resulting representation is an APR instance/witness pair $(\mathfrak{x}_{\text{APR}}, w_{\text{APR}})$. This reduction is deterministic on the verifier side, i.e., involves no verifier-side randomness and no interaction; as such, it also has perfect completeness and perfect soundness. (The prover uses randomness to obtain zero-knowledge; however, this use does not affect completeness, nor soundness.)
3. The APR representation is used to produce, via a 1-round IOP, a pair of instances of the Reed-Solomon proximity testing (RPT) problem. The instances are defined now by the parameters of the RS code. The two codes resulting from the reduction are over the same field \mathbb{F} but have different evaluation domains (L, L_{cmp}) and different code rates $(\rho, \rho_{\text{cmp}})$. The witness in this case is a pair of purported codewords $(f^{(0)}, g^{(0)})$. The verifier’s randomness in the 1-round IOP is used (among other things) to “link” the numerous constraints of the transition relation into a single (random) one. We thus refer to this step as the *algebraic linking IOP* (ALI) protocol.

4. Finally, for each of the two functions (oracles) $f^{(0)}, g^{(0)}$, we invoke the *fast RS IOP of proximity* (FRI) protocol from [14], and this completes our reduction.

Section organization After setting up notation (Appendix B.1) we formally define our starting points — the algebraic intermediate representation (AIR) of CI statements (Appendix B.2); this section also includes the main technical lemmas (Lemmas B.6 and B.9) that prove our main Theorems 3.4 and 3.5, as well as the lemma on which our ZK-STARK realization relies (Lemma B.7). In Appendix B.3 we formally define the representation of CI statements that is the output of the algebraic placement and routing (APR) reduction (Definition B.10) and the APR-reduction (leftmost arrow of Figure 8) appears in Appendix B.4. Appendix B.5 describes the next step of the reduction (middle arrow of Figure 8) starting with Definition B.14 of the RS proximity testing (RPT) problem and followed by the algebraic linking IOP (ALI) stated in Theorem B.15 and discussed later there. Appendix B.6 states the main properties of the FRI protocol that solves the RPT problem (rightmost arrow of Figure 8). Appendix B.7 uses the components defined earlier — AIR, APR, ALI, and FRI— to prove our main Theorems 3.4 and 3.5. Appendix B.8 concludes with a discussion of the particular setting of parameters that our ZK-STARK uses for the DPM computation.

B.1 Preliminaries and notation

Density The *density* of a subset S' of a finite set S is $\mu(S'/S) \triangleq |S'|/|S|$.

Sets and functions We denote by $\mathbf{0}$ the constant zero function. For S a set and f a function with domain S let $f(S) = \{f(x) \mid x \in S\}$ and for \mathcal{F} a set of functions with domain S let $\mathcal{F}(S) = \cup_{f \in \mathcal{F}} f(S)$.

Binary fields A finite field is denoted by \mathbb{F} and \mathbb{F}_q is the field of size q . A *binary field* is a finite field of characteristic 2; all fields considered in this paper are binary²⁶. When we use the term *affine space* for a subset $H \subset \mathbb{F}$ that is an additive coset of an \mathbb{F}_2 -linear space, meaning $H = \alpha + V \triangleq \{\alpha + v \mid v \in V\}$ for some fixed $\alpha \in \mathbb{F}$ and $V \subseteq \mathbb{F}$ a linear space over the two-element field \mathbb{F}_2 .

Remark B.1 (Canonical representations of binary fields). *We assume canonical representation for binary fields \mathbb{F} , given by an irreducible polynomial and a primitive element $\mathbf{g} \in \mathbb{F}$ for it (i.e., \mathbf{g} generates \mathbb{F}^*). We use the standard basis $\{1, \mathbf{g}, \mathbf{g}^2, \dots, \mathbf{g}^{n-1}\}$ to represent \mathbb{F}_{2^n} over \mathbb{F}_2 . Similarly, for a fixed \mathbb{F}_2 -subspace of dimension k , the (non-zero) polynomial of degree 2^k that vanishes on S will assumed to be known to the verifier; recall this polynomial is linearized (see [78, Section 3.4]) and has at most k non-zero coefficients.*

In particular, when discussing arithmetic complexity of prover and verifier, we assume the representations of all relevant fields, primitive elements and subspace polynomials are known to both parties. As will become evident, these depend only on the parameters of the computation (like degree and running time). This assumption can be made with no significant loss in generality, because the verifier could request the prover to present all such representations, along with a suitable “proof of primitivity” for \mathbf{g} .

Codes We view codewords in a linear error correcting code C over \mathbb{F} of blocklength n as functions with domain S , $|S| = n$ and range \mathbb{F} , i.e., $C \subseteq \mathbb{F}^S$; for $x \in S$, the x -entry of a codeword $w \in C$ is denoted by $w(x)$. For $v \in \mathbb{F}^S$ let $w_v \in C$ be a codeword that is closest to v in relative hamming distance, breaking

²⁶Many of the results here apply, with minimal modifications, to arbitrary fields of small characteristic. For the sake of simplicity we avoid dealing with this generalization.

ties arbitrarily (say, by ordering the elements of C arbitrarily and setting w_v to be the smallest $w \in C$ with minimal distance to v).

Restrictions For $S' \subset S$ we denote by $w|_{S'}$ the restriction of w to sub-domain S' and similarly define $C|_{S'} = \{w|_{S'} \mid w \in C\}$, noticing $C|_{S'} \subseteq \mathbb{F}^{S'}$.

Interpolants, evaluations and low degree extensions The *interpolant* of $w \in \mathbb{F}^S$, denoted interpolant^w , is the unique polynomial P , $\deg(P) < |S|$ satisfying $\forall x \in S, P(x) = w(x)$. The *multi-point evaluation* of $P(X)$ on domain S is the function $f : S \rightarrow \mathbb{F}$ satisfying $f(x) = P(x)$ for all $x \in S$; when $|S| > \deg(P)$, the interpolant of this evaluation is P .

Recall the definition of a low degree extension (LDE) from Definition 2.1. We use the following efficient LDE algorithm of Lin et al. [79].

Theorem B.2 (LDE arithmetic complexity[79]). *For \mathbb{F} a binary field, $S \subseteq \mathbb{F}$ an \mathbb{F}_2 -affine space and $c_1, \dots, c_n \in \mathbb{F}$ there exists an arithmetic circuit computing an advice string A that requires $(n+1)\text{polylog}|S|$ arithmetic operations in \mathbb{F} , and another arithmetic circuit that uses the advice A and computes the LDE of a function $f : S \rightarrow \mathbb{F}$ to domain $\bigcup_{c_i} S + c_i$. This latter circuit has arithmetic complexity $3(n+1)(|S| \log |S|)$ over \mathbb{F} .*

For simplicity we shall ignore the complexity of advice computation, which is negligible compared to $3(n+1)(|S| \log |S|)$.

Arithmetic complexity For an arithmetic circuit C with gates of fan-in ≤ 2 over the set of gates $\{+, \times, \div\}$, we denote by $T_{\text{arith}}(C)$ the arithmetic complexity of C , defined as the total number of gates in the arithmetic circuit. The *multiplication complexity* is the number of \times -gates and the *addition complexity* is similarly defined. For a function $f : \mathbb{F}^n \rightarrow \mathbb{F}^m$ whose ‘‘canonical circuit’’ C_f is implicitly known, we abuse notation and define $T_{\text{arith}}(f) = T_{\text{arith}}(C_f)$.

Given a set of functions $\mathcal{F} = \{f_1, f_2, \dots, f_m\}$ mapping elements of \mathbb{F}^n to \mathbb{F} , we denote by $T_{\text{arith}}(\mathcal{F})$ the (total) arithmetic complexity of \mathcal{F} and define it as the arithmetic of a circuit $C : \mathbb{F}^n \rightarrow \mathbb{F}^m$ such that for any $\vec{v} \in \mathbb{F}^n$ and any $1 \leq i \leq m$, $C(\vec{v})_i = f_i(\vec{v})$.

Functions evaluated on sets of inputs For $f : S \rightarrow \Sigma$ a function and $S' \subset S$ let $f(S') = \{f(x) \mid x \in S'\}$; when $f(S')$ is a singleton set $\{\alpha\}$ we shall simplify notation and write $\alpha = f(S')$.

B.2 Algebraic Intermediate Representation (AIR)

An *algebraic execution trace* of a program running for T steps is represented by a $w \times T$ array in which each entry is an element of a finite field. A single row describes the state of the computation at a certain point in time, and a single column tracks a ‘‘register’’ and its contents over time. It is straightforward to reduce (instance, witness) pairs for a relation that defines a language $L \in \text{NTimeSpace}(\hat{T}(n), \hat{w}(n))$ to AIR instances with algebraic execution traces of size $T(n) = O(\hat{T}(n))$ and width $w(n) = O(\hat{w}(n))$, see, e.g., [90] for examples of such reductions.

Definition B.3 (Algebraic internal representation (AIR)). *The relation R_{AIR} is the set of pairs $(\mathfrak{x}, \mathfrak{w}) = (\mathfrak{x}_{\text{AIR}}, \mathfrak{w}_{\text{AIR}})$ satisfying*

1. **Instance Format:** *the instance \mathfrak{x} is a tuple $\mathfrak{x} = (\mathbb{F}, T, w, \mathcal{P}, C, B)$ where*

- \mathbb{F} is a finite field
- T is an integer representing a bound on running time
- w is an integer representing state width
- $\mathcal{P} = \{P_1, \dots, P_s\} \subset \mathbb{F}[X_1, \dots, X_w, Y_1, \dots, Y_w]$ is a set of constraints. The degree of \mathcal{P} is $\deg(\mathcal{P}) \triangleq \max_{P \in \mathcal{P}} \deg(P)$
- C is a monotone boolean circuit over variables Z_1, \dots, Z_s with multi-input AND and OR gates. The size of the circuit is the number of gates in it, and its degree, denoted $\deg(C)$ is defined inductively thus:
 - the degree of the input variable Z_i is $\deg(Z_i) = \deg(P_i)$, where $P_i \in \mathcal{P}$ is defined above;
 - the degree of an AND gate g_j with input gates g_1, \dots, g_t is $\deg(g_j) = \max\{\deg(g_1), \dots, \deg(g_t)\}$
 - the degree of an OR gate g_j with input gates g_1, \dots, g_t is $\deg(g_j) = \deg(g_1) + \dots + \deg(g_t)$
 - finally, $\deg(C)$ is the degree of its (single) output gate.
- B is a set of boundary constraints, where each boundary constraint is a tuple (i, j, α) for $i \in [T], j \in [w], \alpha \in \mathbb{F}$

2. **Witness Format:** The witness w is a set of functions $w_1, \dots, w_w : [T] \rightarrow \mathbb{F}$; we say w satisfies the instance if and only if

- For all boundary constraints (i, j, α) we have $w_j(i) = \alpha$
- For all $t \in [T - 1]$ and we have $C(\text{IsZero}_s(\mathcal{P}(w[t], w[t + 1]))) = \text{TRUE}$ where
 - $w[t] = (w_1(t), \dots, w_w(t))$
 - $\text{IsZero} : \mathbb{F} \rightarrow \{\text{TRUE}, \text{FALSE}\}$ is the mapping that sends $0 \in \mathbb{F}$ to TRUE and all non-zero elements to FALSE, and $\text{IsZero}_k : \mathbb{F}^k \rightarrow \{\text{TRUE}, \text{FALSE}\}^k$ is the natural extension of IsZero to multiple inputs and outputs, namely, $\text{IsZero}_k(\alpha_1, \dots, \alpha_k) := (\text{IsZero}(\alpha_1), \dots, \text{IsZero}(\alpha_k))$

Finally, R_{AIR} is the set of all pairs (x, w) such that w satisfies x , and $L_{\text{AIR}} \triangleq \{x \mid \exists w, (x, w) \in R_{\text{AIR}}\}$.

Remark B.4 (Generalization to non-monotone circuits). The reason that we restrict the definition above to monotone circuits, as opposed to general circuit which are more expressive, is that adding negation gates to the boolean circuit may destroy perfect completeness, reduce efficiency, or require more rounds of interaction (depending on the way negation gates are arithmetized). Often (as is the case with our benchmarks here), it suffices to use De Morgan’s laws to effectively push negation gates to the inputs of C and incorporate them into the system of polynomials \mathcal{P} .

Our reductions will use AIRs in which \mathbb{F} is invariably a binary field, and we shall further assume the witness size and degree are “binary friendly” according to the following definition.

Definition B.5 (BAIR). A binary AIR (BAIR) instance is an AIR instance $x = (\mathbb{F}, T, w, \mathcal{P}, C, B)$ satisfying for some $n, t, d \in \mathbb{N}^+$ (i) $|\mathbb{F}| = 2^n$; (ii) $T = 2^t - 1$; and (iii) $\deg(C) \leq 2^d$. We call x an (n, t, d) -BAIR instance. The relations R_{BAIR} and language L_{BAIR} are the natural restriction of R_{AIR} and L_{AIR} to binary AIR instances.

The following lemma implies Theorem 3.4, as shown in Appendix B.7. Its statement is “information theoretic”, i.e., relies on no unproven cryptographic assumptions and holds against computationally unbounded provers. In contrast to prior ZK-IOP protocols [22, 15], it fully specifies all constants and uses no asymptotic notation.

Lemma B.6 (wi-IOP of knowledge for R_{BAIR}). *For soundness error parameter $\lambda : \mathbb{N}^+ \rightarrow \mathbb{N}^+$, rate integer $\mathcal{R} \geq 2$, and ZK parameter $k \geq 1$, let $\delta = (1 - 2^{-\mathcal{R}} (1 + 2^{-d}))$. Let relation $R_{\text{AIR}}^{(\lambda, \mathcal{R}, k)}$ denote the restriction of R_{AIR} to (n, t, d) -BAIR instances that satisfy (i) $n > \lambda + t + d + \mathcal{R} + k + 2$ and (ii) $k + t \geq 6 + \log \lambda + \log \log \frac{1}{1-\delta}$. Then $R_{\text{AIR}}^{(\lambda, \mathcal{R})}$ has a ZK-IOP of knowledge with the following parameters for (n, t, d) -BAIR instances:*

- soundness error at most $\text{err} = 2^{-\lambda}$
- knowledge error bound at most $\epsilon' = 4 \cdot \text{err}$
- prover arithmetic complexity

$$\text{tp}^{\mathbb{F}} \leq (9w(t + d + \mathcal{R} + 3) + |\mathcal{P}| + w + T_{\text{arith}}(\mathcal{P}) + 12) \cdot 2^{t+d+\mathcal{R}+3},$$

- verifier arithmetic complexity

$$\text{tv}^{\mathbb{F}} \leq 2 \left(|\mathcal{B}|^2 + \left\lceil \frac{\lambda + 2}{\log \frac{1}{1-\delta/10}} \right\rceil \cdot (8(w + |\Phi| + T_{\text{arith}}(\mathcal{P}) + |\mathcal{B}|) + 21(k + \mathcal{R} + t + d + 2)) \right),$$

- round complexity $r = \frac{t+d}{2} + 2$,
- query complexity $q \leq 8 \left\lceil \frac{\lambda+2}{\log \frac{1}{1-\delta/10}} \right\rceil \cdot (2w + k + t + d + 2)$

We conjecture that the soundness (of two separate components) in the IOP above is not tight. Consequently, we conjecture that the same soundness as above can be obtained with fewer queries, as stated next.

Lemma B.7 (wi-IOP for R_{BAIR}^λ with improved conjectured soundness). *Assuming Conjectures B.17 and B.19, the following holds. For soundness error parameter $\lambda : \mathbb{N}^+ \rightarrow \mathbb{N}^+$, rate integer $\mathcal{R} \geq 3$, and ZK parameter $k \geq 1$, let relation $R_{\text{AIR}}^{(\lambda, \mathcal{R}, k)}$ denote the restriction of R_{AIR} to (n, t, d) -BAIR instances that satisfy (i) $n > \max\{58, \lambda, t + d + \mathcal{R} + k\} + 2$ and (ii) $k + t > 5 + \log \lceil \frac{\lambda+2}{\mathcal{R}} \rceil$. Then $R_{\text{AIR}}^{(\lambda, \mathcal{R})}$ has a ZK-IOP of knowledge with the following parameters for (n, t, d) -BAIR instances:*

- soundness error and round complexity as stated in Lemma B.6,
- verifier arithmetic complexity

$$\text{tv}^{\mathbb{F}} \leq 2|\mathcal{B}|^2 + \left\lceil \frac{\lambda + 2}{\mathcal{R}} \right\rceil \cdot (16(w + |\Phi| + T_{\text{arith}}(\mathcal{P}) + 42 \cdot (k + \mathcal{R} + t + d + 2)))$$

- query complexity $q \leq 4 \cdot \left(\left\lceil \frac{\lambda+2}{\mathcal{R}+d} \right\rceil \cdot (w + 3 + k + t + d) + \left\lceil \frac{\lambda+2}{\mathcal{R}} \right\rceil \cdot (3w + 3 + k + t + d) \right)$
- prover arithmetic complexity

$$\text{tp}^{\mathbb{F}} \leq \left(9w(t + d + \mathcal{R} + 3) + |\mathcal{P}| + w + T_{\text{arith}}(\mathcal{P}) + 12 \left\lceil \frac{\lambda + 2}{n - (10 + 2 \log n)} \right\rceil \right) \cdot 2^{t+d+\mathcal{R}+3}$$

Definition B.3 assumes that the full machine state is captured by w field elements. Certain computations require as much space complexity as time complexity ($w = \Omega(T)$), leading to AIR witnesses of total size $\Omega(w \cdot T) = \Omega(T^2)$. To reach witnesses of size $O(T \log T)$, irrespective of memory consumption, we require the following definition. It follows the approach of [19, 20] and uses a *pair* of AIR constraints — one for verifying the validity of consecutive time steps and another for verifying memory consistency; the latter set of constraints is applied to a permutation of the steps of the execution trace, and part of the witness is a specification of this permutation. We call this relation the *permuted algebraic intermediate representation* (R_{PAIR}).

Definition B.8 (Permuted algebraic intermediate representation (PAIR)). *The relation R_{PAIR} is the set of pairs $(\mathfrak{x}, \mathfrak{w}) = (\mathfrak{x}_{\text{PAIR}}, \mathfrak{w}_{\text{PAIR}})$ satisfying*

1. **Instance Format:** *the instance \mathfrak{x} is a tuple $\mathfrak{x} = (\mathbb{F}, T, w, \mathcal{P}_T, \mathcal{P}_\pi, \mathcal{B})$ where*
 - \mathbb{F} is a finite field
 - T is an integer representing a bound on running time
 - w is an integer representing state width
 - $\mathcal{P}_T, \mathcal{P}_\pi \subset \mathbb{F}[X_1, \dots, X_w, Y_1, \dots, Y_w]$ are two sets of constraints.
 - \mathcal{B} is a set of boundary constraints, where each boundary constraint is a tuple (i, j, α) for $i \in [T], j \in [w], \alpha \in \mathbb{F}$
2. **Witness Format:** *The witness \mathfrak{w} is a pair (\hat{w}, π) where \hat{w} is a sequence of w functions $w_1, \dots, w_w: [T] \rightarrow \mathbb{F}$, and $\pi: [T] \rightarrow [T]$ is a permutation; we say \mathfrak{w} satisfies the instance if and only if*
 - (a) *For all boundary constraints (i, j, α) we have $w_j(i) = \alpha$*
 - (b) *For all $t \in [T - 1]$ and for all $P \in \mathcal{P}_T$ we have $P(w[t], w[t + 1]) = 0$*
 - (c) *For all $t \in [T - 1]$ and for all $P \in \mathcal{P}_\pi$ we have $P(w[t], w[\pi(t)]) = 0$*

Finally, R_{PAIR} is the set of all pairs $(\mathfrak{x}, \mathfrak{w})$ such that \mathfrak{w} satisfies \mathfrak{x} , and $\text{PAIR} \triangleq \{\mathfrak{x} \mid \exists \mathfrak{w}, (\mathfrak{x}, \mathfrak{w}) \in R_{\text{AIR}w\text{RAM}}\}$.

We define the sub-relation of binary PAIR (BPAIR) analogously to the definition of binary AIR in Definition B.5.

L_{PAIR} is NEXP-complete, as stated next. We omit the proof, which appears, e.g., in [90].

Lemma B.9 (wi-STIK for BPAIR). *For every language $L \in \text{NTIME}(T(n)), T(n) \geq n$ there exists a deterministic polynomial time reduction from L to L_{PAIR} , mapping an instance x of L to an instance $\mathfrak{x} = (\mathbb{F}, T, w, \mathcal{P}_T, \mathcal{P}_\pi, \mathcal{B})$ of L_{PAIR} that satisfies*

- $T = O(T(n))$,
- $w = O(\log T(n))$,
- $\deg(\mathcal{P}_T \cup \mathcal{P}_\pi) = O(1)$, $|\mathcal{P}_T \cup \mathcal{P}_\pi| = O(\log T(n))$, and $T_{\text{arith}}(\mathcal{P}) = O(\log T(n))$,
- $|\mathcal{B}| = O(n)$.

Furthermore, a nondeterministic witness \mathfrak{w} for \mathfrak{x} can be computed in time $O(T(n) \log T(n)) + n^{O(1)}$, given x and a nondeterministic witness w for the membership of x in L .

B.3 Algebraic placement and routing (APR)

The phase in circuit design known as *placement and routing* deals with laying out states of a computation as to optimize certain physical constraints. Our next phase does something similar, placing and routing information about the states of a computation in a way that optimizes *algebraic*, rather than physical, constraints; hence we call it the *algebraic placement and routing (APR)* binary relation.

Recall that $\text{Aff}_1(\mathbb{F})$ denotes the 1-dimensional affine group over \mathbb{F} , its set of elements is isomorphic to $\{aX + b \mid a \in \mathbb{F}^*, b \in \mathbb{F}\}$. An *affine graph* generated by a set $\mathcal{N} = \{N_1, \dots, N_s\} \subset \text{Aff}_1(\mathbb{F})$ is the directed graph with vertex set \mathbb{F} and edge set $\{(x, N(x)) \mid x \in \mathbb{F}, N \in \mathcal{N}\}$.

In the following definition, the set Φ captures both boundary, and transition, constraints of the computation. The affine graph is added to the instance description via its generating set \mathcal{N} . Due to the algebraic topology, each column of the algebraic execution trace (corresponding to a register, tracked over time) is now a Reed Solomon codeword of rate ρ and the sequence of rates (one rate per register/column) is also part of the instance description. As explained in Section 2.3, our construction differs from prior works [27, 20, 38, 13] by using a separate codeword for each algebraic register.

Definition B.10 (Algebraic placement and routing (APR) problem). *The relation R_{APR} contains all pairs $(\mathfrak{x}, \hat{\mathfrak{w}})$ satisfying the following requirements:*

1. **Instance Format:** *the instance \mathfrak{x} is a tuple $\mathfrak{x} = (\mathbb{F}, \mathcal{T}, \mathcal{N}, \Phi, L, L_{\text{cmp}}, \vec{\rho}, \rho_{\text{cmp}})$ where*

- \mathbb{F} is a finite field of characteristic 2.
- \mathcal{T} is a set of indices called the algebraic-register indices.
- $\mathcal{N} \subseteq \mathcal{T} \times \text{Aff}_1(\mathbb{F})$ is a set of pairs called neighbors, each pair contains an algebraic-register index, and a member of the affine group over \mathbb{F} (a polynomial of degree exactly 1). Given a set $S \subset \mathbb{F}$ we denote by $\mathcal{N}(S)$ the set $\mathcal{N}(S) := \{N(x) \mid x \in S, (\tau, N) \in \mathcal{N} \text{ for some } \tau \in \mathcal{T}\}$.
- $\Phi \subseteq (\mathbb{F} \times \mathbb{F}^{\mathcal{N}}) \rightarrow \mathbb{F}$ is a set of mappings over the variables $\mathcal{V} := \{X_{\text{loc}}\} \cup \{X_N\}_{N \in \mathcal{N}}$. Assignments to any $\phi \in \Phi$ expressed by mappings $\alpha: \mathcal{V} \rightarrow \mathbb{F}$. Given a set $S \subset \mathbb{F}$ and a sequence of functions indexed by elements of \mathcal{T} , $\vec{f} \in (\mathbb{F}^{\mathcal{N}(S)})^{\mathcal{T}}$, denote by $\alpha_{\vec{f}, \mathcal{N}}: S \rightarrow \mathbb{F}^{\mathcal{V}}$ the mapping defined by $[\alpha_{\vec{f}, \mathcal{N}}(x)]_{X_{\text{loc}}} := x$, and for every $(\tau, N) \in \mathcal{N}$, $[\alpha_{\vec{f}, \mathcal{N}}(x)]_{(X_N)_{(\tau, N)}} := f_{\tau}(N(x))$. Given $\phi \in \Phi$ we denote by $\phi_{\mathcal{N}}[\vec{f}] \in \mathbb{F}^S$ the vector defined by $[\phi_{\mathcal{N}}[\vec{f}]]_x := \phi(\alpha_{\vec{f}, \mathcal{N}}(x))$ for every $x \in S$.
- L and L_{cmp} are two \mathbb{F}_2 -affine subspaces of \mathbb{F} called the witness evaluation space and the composition evaluation space respectively;
- $\vec{\rho} \in (0, 1)^{\mathcal{T}}$ is a sequence of rates called the witness rates, and $\rho_{\text{cmp}} \in (0, 1)$, called the composition rate.

2. **Witness format and satisfiability:** *A witness $\hat{\mathfrak{w}}$ is a sequence of functions indexed by elements from \mathcal{T} ; formally, $\hat{\mathfrak{w}} \in (L^{\mathbb{F}})^{\mathcal{T}}$. We say $\hat{\mathfrak{w}}$ satisfies \mathfrak{x} if both of the following hold:*

- **assignment code membership:** $\forall \tau \in \mathcal{T} : \mathfrak{w}_{\tau} \in \text{RS}[\mathbb{F}, L, \rho_{\tau}]$
- **constraint code membership:** $\forall \phi \in \Phi, \phi_{\mathcal{N}}[\hat{\mathfrak{w}}] \in \text{RS}[\mathbb{F}, L_{\text{cmp}}, \rho_{\text{cmp}}]$

Let R_{APR} be the binary relation containing all pairs $(\mathfrak{x}, \hat{\mathfrak{w}})$ such that \mathfrak{x} is an instance as defined above and $\hat{\mathfrak{w}}$ satisfies \mathfrak{x} , and let APR be the nondeterministic language induced by R_{APR} ,

$$\text{APR} = \{\mathfrak{x} \mid \exists \hat{\mathfrak{w}}, (\mathfrak{x}, \hat{\mathfrak{w}}) \in R_{\text{APR}}\}. \quad (6)$$

The three properties defined next are used later to prove zero knowledge and to provide better soundness bounds.

Definition B.11 (APR properties). *We say an instance $\varkappa := (\mathbb{F}, \mathcal{T}, \mathcal{N}, \Phi, L, L_{\text{cmp}}, \vec{\rho}, \rho_{\text{cmp}}) \in \text{APR}$*

1. *is κ -independent if there is a sequence of sets $\{V_\tau\}_{\tau \in \mathcal{T}}$ such that each $V_\tau \subset \mathbb{F}^L$ is $\kappa|L|$ -independent, and for any \hat{w} such that $w_\tau \in V_\tau$, it holds that \hat{w} satisfies \varkappa . Let $T_{\text{Sampling}}(\cdot)$ denote the arithmetic complexity of the circuit, that on input $w_\tau \in V_\tau$ (and random field elements), samples a uniformly random $w \in \{V_\tau\}_{\tau \in \mathcal{T}}$.*
2. *has δ -distance if:*
 - $\delta \leq 1 - \rho_{\text{cmp}}$.
 - $\delta \leq 1 - \rho_\tau$ for any $\tau \in \mathcal{T}$.
 - *There exists some linear code $C \subset \mathbb{F}^{L_{\text{cmp}}}$ of relative distance at least δ , containing the code $\text{RS}[\mathbb{F}, L_{\text{cmp}}, \rho_{\text{cmp}}]$, such that for any sequence of mappings $\vec{f} \in (\mathbb{F}^L)^\mathcal{T}$ where $f_\tau \in \text{RS}[\mathbb{F}, L, \rho_\tau]$ for all $\tau \in \mathcal{T}$, it holds $\forall \phi \in \Phi, \phi_{\mathcal{N}}[\hat{w}] \in C$.*
3. *is Θ -overlapping if for any $y \in L$ the size of the set $S_y := \{x \in L_{\text{cmp}} \mid y \in \mathcal{N}(\{x\})\}$ is at most Θ .*

B.4 APR reduction

The following pair of theorems describe the first step of our reduction, in which a BAIR, or BPAIR, instance-witness pair, is reduced to an APR instance-witness pair. The reductions are deterministic and have no error associated with them. Since the instance-side reduction is carried out by the verifier, and the witness-side reduction by the prover, we denote the two reductions of Theorem B.12 by $\mathcal{V}^{\text{BAIR} \rightarrow \text{APR}}, \mathcal{P}^{\text{BAIR} \rightarrow \text{APR}}$, respectively, and those of Theorem B.13 are denoted $\mathcal{V}^{\text{BPAIR} \rightarrow \text{APR}}, \mathcal{P}^{\text{BPAIR} \rightarrow \text{APR}}$.

The ideas used in both reduction are based on ideas from [91, 25, 27, 19, 38] where affine graphs are used to encode and verify program traces. The current construction is most similar to the one used in [13], where a cyclic graph is used to embed known to verifier order of the trace, usually used to verify two consecutive configurations are consistent, and a back-to-back De Bruijn graph to verify a permutation not known to the verifier, usually used to verify the memory consistency. The novelty in the construction in [13] in apposed to previous construction is (i) the usage of a cyclic graph for consecutive configurations consistency testing, (ii) using back-to-back butterfly De Bruijn of to represent permutations instead of a ‘straight De Bruijn graph which is $\times 2.5$ bigger, and (iii) selectively routing on the De Bruijn graph only a small part of the configuration. All those reduce the size of the witness, effectively reducing the degrees of polynomials handled by the proof system, and the concrete efficiency of the prover. The novel technique in this work uses the techniques in [13], and additionally splits the witness to several polynomials, each of degree much smaller then what needed to represent the whole witness, effectively reducing further more the degree of polynomials handled by the proof system. As described in Appendix D.1, the proof-system protocol uses interactivity to eliminate the need to provide separate proof of RS-proximity to each witness polynomial. The reductions use the ZK-IOP construction introduced [17], preparing the APR instance and witness for the zero-knowledge protocol described in Appendix D.1.

Theorem B.12 (Algebraic placement and routing). *There are two algorithms $\mathcal{V}^{\text{BAIR} \rightarrow \text{APR}}, \mathcal{P}^{\text{BAIR} \rightarrow \text{APR}}$ such that for any $k, \mathcal{R}, t, d \in \mathbb{N}^+$ and BAIR instance $\varkappa = (\mathbb{F}, \mathcal{T}, w, \mathcal{P}, \mathcal{C}, \mathcal{B})$ with $T = 2^t - 1 > 4$, $\deg \mathcal{P} \leq 2^d$, and $|\mathbb{F}| = 2^n$ with $n > 2 + k + \mathcal{R} + t + d$:*

1. *Perfect completeness:* $(\mathfrak{x}, \mathfrak{w}) \in R_{\text{AIR}} \Rightarrow (\mathbb{V}^{\text{BAIR} \rightarrow \text{APR}}(\mathfrak{x}, k, \mathcal{R}), \mathbb{P}^{\text{BAIR} \rightarrow \text{APR}}(\mathfrak{x}, \mathfrak{w}, k, \mathcal{R})) \in R_{\text{APR}}$
2. *Perfect soundness:* $\mathfrak{x} \in \text{AIR} \Leftarrow \mathbb{V}^{\text{BAIR} \rightarrow \text{APR}}(\mathfrak{x}, k, \mathcal{R}) \in \text{APR}$
3. *Knowledge extraction:* *There is an efficient knowledge extractor $\mathbb{E}^{\text{BAIR} \rightarrow \text{APR}}$, such that for every $(\mathbb{V}^{\text{BAIR} \rightarrow \text{APR}}(\mathfrak{x}, k, \mathcal{R}), \mathfrak{w}_{\text{APR}}) \in R_{\text{APR}}$, the extractor outputs a witness $\mathfrak{w} \leftarrow \mathbb{E}^{\text{BAIR} \rightarrow \text{APR}}(\mathfrak{x}, k, \mathcal{R}, \mathfrak{w}_{\text{APR}})$ such that $(\mathfrak{x}, \mathfrak{w}) \in R_{\text{AIR}}$*
4. *Let $\mathfrak{x}_{\text{APR}} = (\mathbb{F}, \mathcal{T}, \mathcal{N}, \Phi, L, L_{\text{cmp}}, \vec{\rho}, \rho_{\text{cmp}}) = \mathbb{V}^{\text{BAIR} \rightarrow \text{APR}}(\mathfrak{x}, k, \mathcal{R})$:*
 - (a) $|\mathcal{T}| = \mathfrak{w}$
 - (b) $|\mathcal{N}| = 3\mathfrak{w}$, with same 3 neighbor polynomials for each witness index $\tau \in \mathcal{T}$
 - (c) $|\Phi| \leq 2|\mathcal{P}|$ and $\mathbb{T}_{\text{arith}}(\Phi) \leq 2\mathbb{T}_{\text{arith}}(\mathcal{P}) + 3\mathfrak{t} + O(\mathfrak{w} \cdot |\mathcal{P}| \cdot |\mathcal{B}|)$
 - (d) $\log_2 |L| = 2 + k + \mathcal{R} + \mathfrak{t} + \mathfrak{d}$
 - (e) $\log_2 |L_{\text{cmp}}| = k + \mathcal{R} + \mathfrak{t} + \mathfrak{d}$
 - (f) *The maximal rate ρ_{max} in $\vec{\rho}$ satisfies $\rho_{\text{max}} \leq 2^{-(2+\mathcal{R}+\mathfrak{d})}$*
 - (g) $\rho_{\text{cmp}} \leq 2^{-\mathcal{R}}$
 - (h) $\mathfrak{x}_{\text{APR}}$ is $(1 - 2^{-\mathcal{R}}(1 + 2^{-\mathfrak{d}}))$ -distance, 1-overlapping, and $(2^{-(2+\mathcal{R}+\mathfrak{d})}(1 - 2^{-k}))$ -independent with $\mathbb{T}_{\text{Sampling}}(\mathfrak{x}_{\text{APR}}) \leq \tilde{O}(|L|)$

5. *Arithmetic complexity over \mathbb{F} (cf. Remark B.1):*

- (a) **Verifier:** *The arithmetic complexity of $\mathbb{V}^{\text{BAIR} \rightarrow \text{APR}}(\mathfrak{x}, k, \mathcal{R})$ is that which is required to compute the polynomials $Z_{\mathcal{B},j}, \mathcal{E}_{\mathcal{B},j} \in \mathbb{F}[x]$ for all $1 \leq j \leq \mathfrak{w}$ (Appendix C.1.1), and in particular, it is at most $2|\mathcal{B}|^2$ arithmetic operations²⁷.*
- (b) **Prover:** *The arithmetic complexity of $\mathbb{P}^{\text{BAIR} \rightarrow \text{APR}}(\mathfrak{x}, \mathfrak{w}, k, \mathcal{R})$ is at most the accumulated arithmetic complexity of (i) the arithmetic complexity of $\mathbb{V}^{\text{BAIR} \rightarrow \text{APR}}(\mathfrak{x}, k, \mathcal{R})$, (ii) $2^{\mathfrak{t}}$ multiplications and additions over \mathbb{F} , (iii) $3\mathfrak{w}$ low degree extension (LDE) computations (cf. Definition 2.1), each over an affine space of dimension at most $k + \mathfrak{t}$, with evaluation over at most $2^{\mathcal{R}+\mathfrak{d}+2}$ shifts; in particular, each LDE can be computed in arithmetic complexity at most $3 \cdot (k + \mathfrak{t}) \cdot 2^{k+\mathfrak{t}+\mathcal{R}+\mathfrak{d}+2}$ (see Theorem B.2).*

The following theorem is the analog of the previous one, stated for the R_{BPAIR} rather than for R_{BAIR} .

Theorem B.13 (Algebraic placement and routing – permuted version). *There are two algorithms, denoted $\mathbb{V}^{\text{BPAIR} \rightarrow \text{APR}}, \mathbb{P}^{\text{BPAIR} \rightarrow \text{APR}}$, such that for any $k, \mathcal{R} \in \mathbb{N}$, with $\mathcal{R} > 1$, and **BPAIR** instance $\mathfrak{x} = (\mathbb{F}, \mathbb{T}, \mathfrak{w}, \mathcal{P}_{\mathbb{T}}, \mathcal{P}_{\pi}, \mathcal{B})$ with $\mathbb{T} = 2^{\mathfrak{t}} - 1 > 4$, $\max\{\deg \mathcal{P}_{\mathbb{T}}, \deg \mathcal{P}_{\pi}, 2\} \leq 2^{\mathfrak{d}}$, and $[\mathbb{F} : \mathbb{F}_2] > 2 + k + \mathcal{R} + \mathfrak{t} + \lceil \log(\mathfrak{t} + 1) \rceil + \mathfrak{d}$:*

1. *Perfect completeness:* $(\mathfrak{x}, \mathfrak{w}) \in R_{\text{BPAIR}} \Rightarrow (\mathbb{V}^{\text{BPAIR} \rightarrow \text{APR}}(\mathfrak{x}, k, \mathcal{R}), \mathbb{P}^{\text{BPAIR} \rightarrow \text{APR}}(\mathfrak{x}, \mathfrak{w}, k, \mathcal{R})) \in R_{\text{APR}}$
2. *Perfect soundness:* $\mathfrak{x} \in \text{BPAIR} \Leftarrow \mathbb{V}^{\text{BPAIR} \rightarrow \text{APR}}(\mathfrak{x}, k, \mathcal{R}) \in \text{APR}$

²⁷For large sets of boundary constraints, one may modify the construction so that $Z_{\mathcal{B},j}$ is a linearized polynomial, in which case the arithmetic complexity decreases to $O(|\mathcal{B}| \log |\mathcal{B}|)$; details omitted because typically $|\mathcal{B}|$ is small.

3. *Knowledge extraction:* There is an efficient knowledge extractor $E^{\text{BPAIR} \rightarrow \text{APR}}$, such that for every $(V^{\text{BPAIR} \rightarrow \text{APR}}(\mathfrak{x}, k, \mathcal{R}), w_{\text{APR}}) \in R_{\text{APR}}$, the extractor outputs a witness $w \leftarrow E^{\text{BPAIR} \rightarrow \text{APR}}(\mathfrak{x}, k, \mathcal{R}, w_{\text{APR}})$ such that $(\mathfrak{x}, w) \in R_{\text{BPAIR}}$
4. Let $\mathfrak{x}_{\text{APR}} = (\mathbb{F}, \mathcal{T}, \mathcal{N}, \Phi, L, L_{\text{cmp}}, \vec{\rho}, \rho_{\text{cmp}}) = V^{\text{BPAIR} \rightarrow \text{APR}}(\mathfrak{x}, k, \mathcal{R})$:
 - (a) $|\mathcal{T}| = 2(w + 1)$
 - (b) $|\mathcal{N}| \leq 11|\mathcal{T}|$, with same 11 neighbor polynomials for each witness index $\tau \in \mathcal{T}$
 - (c) $|\Phi| \leq 2|\mathcal{P}_{\mathcal{T}}| + |\mathcal{P}_{\pi}| + 8w + 11$ and $T_{\text{arith}}(\Phi) \leq 2T_{\text{arith}}(\mathcal{P}_{\mathcal{T}}) + T_{\text{arith}}(\mathcal{P}_{\pi}) + O(w(|\mathcal{P}_{\mathcal{T}}| + |\mathcal{P}_{\pi}|) + t)$
 - (d) $\log_2 |L| = 2 + k + \mathcal{R} + t + \lceil \log(t + 1) \rceil + d$
 - (e) $\log_2 |L_{\text{cmp}}| = k + \mathcal{R} + t + \lceil \log(t + 1) \rceil + d$
 - (f) The maximal rate ρ_{max} in $\vec{\rho}$ satisfies $\rho_{\text{max}} \leq 2^{-(2+\mathcal{R}+d)}$
 - (g) $\rho_{\text{cmp}} \leq 2^{-\mathcal{R}}$
 - (h) $\mathfrak{x}_{\text{APR}}$ is $(1 - 2^{-\mathcal{R}}(1 + 2^{-d}))$ -distance, 5-overlapping, and $(2^{-(2+\mathcal{R}+d)}(1 - 2^{-k}))$ -independent with $T_{\text{Sampling}}(\mathfrak{x}_{\text{APR}}) \leq \tilde{O}(|L|)$
5. *Arithmetic complexity over \mathbb{F} (cf. Remark B.1):*
 - (a) **Verifier:** The arithmetic complexity of $V^{\text{BPAIR} \rightarrow \text{APR}}(\mathfrak{x}, k, \mathcal{R})$ is that which is required to compute the polynomials $Z_{B,j}, \mathcal{E}_{B,j} \in \mathbb{F}[x]$ for all $1 \leq j \leq w$ (Appendix C.1.1), and in particular, it is at most $2|B|^2$ arithmetic operations.
 - (b) **Prover:** The arithmetic complexity of $P^{\text{BPAIR} \rightarrow \text{APR}}(\mathfrak{x}, w, k, \mathcal{R})$ is at most the accumulated arithmetic complexity of (i) the arithmetic complexity of $V^{\text{BPAIR} \rightarrow \text{APR}}(\mathfrak{x}, k, \mathcal{R})$, (ii) the complexity of routing a back-to-back De Bruijn butterfly network of dimension t (Theorem G.3) (iii) $2^t + 2^{\lceil \log(t+1) \rceil}$ multiplications and additions over \mathbb{F} , (iv) $5(w + 1)$ low degree extension (LDE) computations (cf. Definition 2.1), each over an affine space of dimension at most $k + t + \lceil \log(t + 1) \rceil$, with evaluation over at most $2^{\mathcal{R}+d+2}$ shifts; in particular, each LDE can be computed in arithmetic complexity at most $3 \cdot (k + t + \lceil \log(t + 1) \rceil) \cdot 2^{k+t+\lceil \log(t+1) \rceil + \mathcal{R}+d+2}$ (see Theorem B.2).

B.5 Algebraic linking IOP (ALI)

The output of the reductions described in the previous section are pairs in the relation R_{APR} . The next phase in our reduction uses a 1-round IOP, called the *algebraic linking IOP (ALI)*, to reduce the problem to a pair of *binary RS proximity testing (BRPT)* problems. The following definition formally defines the relation underlying the problem, adding to the informal discussion in Section 2.1).

Definition B.14 (Binary RS proximity testing (RPT)). *Instances of the RS proximity testing problem (RPT) are triples $\mathfrak{x}_{\text{RS}} = (\mathbb{F}, S, \rho)$ where $S \subseteq \mathbb{F}$ and $\rho \in [0, 1]$. A witness w_{RS} for \mathfrak{x}_{RS} is a function $w_{\text{RS}} : S \rightarrow \mathbb{F}$, and we say it satisfies \mathfrak{x}_{RS} iff and only if $w_{\text{RS}} \in \text{RS}[\mathbb{F}, S, \rho]$. The relation R_{RPT} is the set of pairs $(\mathfrak{x}_{\text{RS}}, w_{\text{RS}})$ such that w_{RS} satisfies \mathfrak{x}_{RS} .*

The binary RPT relation (R_{BRPT}) is one in which the instance satisfies:

- \mathbb{F} is a binary field

- S is an affine coset of an \mathbb{F}_2 -linear subspace of \mathbb{F} , defined by a coset shift a_0 and a basis (a_1, \dots, a_k) such that $S = \left\{ a_0 + \sum_{i=1}^k \alpha_i a_i \mid \alpha_1, \dots, \alpha_k \in \mathbb{F}_2 \right\}$
- $\rho = 2^{-\mathcal{R}}$ for $\mathcal{R} \in \mathbb{N}^+$.

The following theorem describes the main properties of the ALI protocol. This protocol, described in Appendix D.1, takes an instance-witness pair of \mathbb{R}_{APR} and uses one round of interaction to reduce it to two instance-witness pairs of the \mathbb{R}_{BRPT} relation.

Theorem B.15 (Algebraic linking IOP (ALI) properties). *Let $\mathfrak{x} := (\mathbb{F}, \mathcal{T}, \mathcal{N}, \Phi, L, L_{\text{cmp}}, \vec{\rho}, \rho_{\text{cmp}})$ be a Θ -overlapping and δ -distance APR instance with $\delta > 0$. The ALI protocol, applied to \mathfrak{x} , satisfies the following properties:*

1. Protocol Schedule

- Prover:** On input (\mathfrak{x}, w') , the output of $\mathbb{P}^{\text{ALI}}(\mathfrak{x}, w')$ is a single oracle $\mathbb{O}_{\text{assignment}}$ comprised of $|\mathcal{T}| + 1$ functions in \mathbb{F}^L and a single function in $\mathbb{F}^{L_{\text{cmp}}}$;
 - Verifier:** On input \mathfrak{x} the output of $\mathbb{V}^{\text{ALI}}(\mathfrak{x})$ is a single message R comprised of $2|\mathcal{T}| + |\Phi|$ random field elements and a pair $\mathfrak{x}_{\text{RS}} = (\mathbb{F}, L, \rho_{\text{max}}), \mathfrak{x}'_{\text{RS}} = (\mathbb{F}, L_{\text{cmp}}, \rho_{\text{cmp}})$ of instances of the BRPT relation.
 - Induced output:** The oracle $\mathbb{O}_{\text{assignment}}$ and verifier randomness define a pair of functions $f^{(0)} : L \rightarrow \mathbb{F}$ and $g^{(0)} : L_{\text{cmp}} \rightarrow \mathbb{F}$, such that each entry of $f^{(0)}$ depends on $|\mathcal{T}| + 1$ entries of $\mathbb{O}_{\text{assignment}}$ and each entry of $g^{(0)}$ depends on $|\mathcal{N}| + 1$ entries of $\mathbb{O}_{\text{assignment}}$.
2. **Completeness** If \mathfrak{x} is satisfied by w' then for any randomness R we have $f^{(0)} \in \text{RS}[\mathbb{F}, L, \rho_{\text{max}}]$ and $g^{(0)} \in \text{RS}[\mathbb{F}, L_{\text{cmp}}, \rho_{\text{cmp}}]$.

3. Soundness Let

$$\zeta \triangleq \max \left\{ 4, 1 + \Theta \frac{|L|}{|L_{\text{cmp}}|} \right\}. \quad (7)$$

If $\mathfrak{x} \notin \text{APR}$ then for every $\mathbb{O}_{\text{assignment}}$ at least one of the following holds,

- $\Pr_R \left[\Delta_{\text{H}}(f^{(0)}, \text{RS}[\mathbb{F}, L, \rho_{\text{max}}]) < \frac{\delta}{2\zeta} \right] \leq \frac{1}{|\mathbb{F}|}$
 - $\Pr_R \left[\Delta_{\text{H}}(g^{(0)}, \text{RS}[\mathbb{F}, L_{\text{cmp}}, \rho_{\text{cmp}}]) < \frac{\delta}{2\zeta} \right] \leq \frac{1}{|\mathbb{F}|}$
4. **Knowledge extraction** Assume $\rho_{\text{max}} \leq \frac{1}{4}$. Then there exists a Las Vegas (randomized) polynomial time algorithm \mathbb{E} for which the following holds. If w^* does not satisfy both conditions claimed in Items 3a and 3b above, then the output of \mathbb{E} on input (\mathfrak{x}, w^*) is a witness w' that satisfies \mathfrak{x} .
5. **Perfect Zero Knowledge** If the instance \mathfrak{x} is κ -independent, then the ALI protocol has perfect zero knowledge against any verifier that makes at most $\kappa|L|$ queries to each w_τ . This holds for any choice of RS-IOPP sub-protocols used in step 2c of the protocol (cf. Appendix D.1).

6. Arithmetic complexity

- Prover:** Assuming prover has w' satisfying \mathfrak{x} , prover arithmetic complexity is at most

$$|\mathcal{T}| \cdot |L| \cdot (3 \log |L| + 5) + 2 |L_{\text{cmp}}| \cdot (\mathbb{T}_{\text{arith}}(\Phi) + |\Phi| + 1) \quad (8)$$

where $\mathbb{T}_{\text{arith}}(\Phi)$ denotes the sum of the arithmetic complexities of the constraint set Φ .

(b) **Verifier:** V^{ALI} requires no arithmetic complexity, and involves only sampling $2|\mathcal{T}| + |\Phi|$ uniformly random field elements from \mathbb{F} .

B.5.1 On the soundness error of the ALI protocol

The probability that the “bad event” in which neither of Items 3a and 3b holds, is $\frac{1}{|\mathbb{F}|}$, translating to a lower bound on soundness error of at least $1/|\mathbb{F}|$. Repeating Step 1b for a number k of times, each with independent randomness R_1, \dots, R_k , reduces soundness error to $\frac{1}{|\mathbb{F}|^k}$. We stress that for constant k , the increase in prover arithmetic complexity is negligible. Prover complexity is dominated by the computation of Item 1a which is executed only once, and requires $O(|L| \cdot \log |L|)$ operations (first summand of Equation (8)); subsequent steps are linear in $|L|$ (second summand of Equation (8)).

We conjecture the soundness error of the ALI protocol is lower than stated Item 3 above. In particular, we believe the denominator 2ζ on the left hand side of Items 3a and 3b is not tight. Our ZK-STARK uses the FRI protocol as the RS-IOPP in the ALI protocol, thus, we state our next conjecture regarding ALI soundness in a manner that will be amenable to the soundness analysis of the FRI protocol.

Blockwise distance — notation Recall the notion of the blockwise distance measure used by FRI; cf. [14, Section 3]. In particular, for $L \subset \mathbb{F}$ an affine space and $L_0 \subset L$ a subspace of size $|L_0| = 2^\eta$, we define the L_0 -blockwise distance between $f, g : L \rightarrow \mathbb{F}$ be the fraction of cosets S of L_0 in L on which $f|_S \neq g|_S$, and let $\mathcal{S}(L_0, L)$ denote the set of cosets of L_0 that are contained in L .

With this notation in mind, we say that L_0 and L'_0 are *good* for an instance $\mathfrak{x} = (\mathbb{F}, \mathcal{T}, \mathcal{N}, \Phi, L, L_{\text{cmp}}, \vec{\rho}, \rho_{\text{cmp}})$, if (i) $L_0 \subset L$ and $L'_0 \subset L_{\text{cmp}}$, and (ii) for every neighbor $N \in \mathcal{N}$ and subset $\hat{S} \subseteq \mathcal{S}(L'_0, L_{\text{cmp}})$, it holds that

$$\left| \left\{ S \in \mathcal{S}(L_0, L) \mid S \cap N \left(\bigcup_{\hat{S} \in \hat{S}} \hat{S} \right) \right\} \right| \geq \left| \bigcup_{\hat{S} \in \hat{S}} \hat{S} \right|.$$

Informally, this condition means that attempting to change the all entries of the union of cosets in \hat{S} by changing entries of \mathfrak{w}^* in the set $N \left(\bigcup_{\hat{S} \in \hat{S}} \hat{S} \right)$ will modify at least as many cosets of $\mathcal{S}(L_0, L)$ as there are elements in $\bigcup_{\hat{S} \in \hat{S}} \hat{S}$.

Conjecture B.16. Suppose $\mathfrak{x} = (\mathbb{F}, \mathcal{T}, \mathcal{N}, \Phi, L, L_{\text{cmp}}, \vec{\rho}, \rho_{\text{cmp}})$ is an unsatisfiable instance of ALI with δ -distance, and let L_0, L'_0 be good for \mathfrak{x} and satisfy $|L_0| = |L'_0| = 2^\eta$. Let $\delta^{(0)}(f^{(0)})$ denote the distance of $f^{(0)}$ from $\text{RS}[\mathbb{F}, L, \rho_{\text{max}}]$ using the blockwise distance measure induced by L_0 . Similarly, let $\delta^{(0)}(g^{(0)})$ denote the distance of $g^{(0)}$ from $\text{RS}[\mathbb{F}, L_{\text{cmp}}, \rho_{\text{cmp}}]$ using the blockwise distance measure induced by L'_0 . Notice $\delta^{(0)}(f^{(0)})$ and $\delta^{(0)}(g^{(0)})$ are random variables that depend on the verifier randomness R . Then

$$\Pr_R \left[\delta^{(0)}(f^{(0)}) < \min \left\{ 1 - \rho_{\text{max}}, \frac{2^\eta |L_{\text{cmp}}|}{|L|} \cdot (\delta - \delta^{(0)}(g^{(0)})) \right\} \right] \leq \frac{1}{|\mathbb{F}|} \quad (9)$$

For practical security purposes, weaker conjecture like the one below may suffice. In what follows, a *pseudo-prover* P^* is a polynomially-bounded randomized machine that acts as a prover in the ALI protocol above.

Conjecture B.17. For every pseudo-prover P^* and sufficiently large instance $\mathfrak{x} = (\mathbb{F}, \mathcal{T}, \mathcal{N}, \Phi, L, L_{\text{cmp}}, \vec{\rho}, \rho_{\text{cmp}})$, we have

$$\Pr \left[\delta^{(0)}(f^{(0)}) < 1 - \rho_{\text{max}} \wedge \delta^{(0)}(g^{(0)}) < 1 - \rho_{\text{cmp}} \right] \leq \frac{1}{|\mathbb{F}|}. \quad (10)$$

The probability above is over the randomness of both P^* and V .

See Appendix D.3 for a discussion of these conjectures.

B.6 Fast Reed-Solomon (RS) IOP of Proximity (IOPP) (FRI)

We now recall the main results of the FRI interactive oracle proof of proximity (IOPP) protocol, as stated in [14]. We assume familiarity with the definition of IOPP (see, e.g., [14, Definition 1.1]).

Theorem B.18 (FRI properties). *The RS code family of rate $\rho = 2^{-\mathcal{R}}$, $\mathcal{R} \geq 2$, $\mathcal{R} \in \mathbb{N}$ has an IOPP (FRI) with the following properties, where $N = |L|$ denotes blocklength (which equals Prover's input length for a fixed $\text{RS}[\mathbb{F}, L, \rho]$ code):*

- **Schedule and rounds:** *The protocol has two phases; during the first phase (COMMIT), verifier sends public randomness and prover sends oracles; round complexity is $\lceil \frac{(\log N) - \mathcal{R}}{2} \rceil$. During the second phase (QUERY), verifier queries the oracles and reaches a decision (prover does not participate in this phase).*
- **Prover:** *prover complexity is less than $6N$ arithmetic operations in \mathbb{F} ; proof length is less than $N/3$ field elements and round complexity is at most $\frac{\log N}{2}$;*
- **Verifier:** *for query-repetition parameter $\ell \in \mathbb{N}^+$, query complexity is $2\ell \log N$; the verifier decision is computed using at most $\ell \cdot 21 \log N$ arithmetic operations over \mathbb{F}*
- **Soundness:** *There exists $\delta_0 \geq \frac{1}{4} (1 - 3\rho) - \frac{1}{\sqrt{N}}$ such that every f that is δ -far in relative Hamming distance from the code, is accepted with probability at most*

$$\text{err}_{\text{FRI}}(\delta) \triangleq \frac{3N}{|\mathbb{F}|} + (1 - \min\{\delta, \delta_0\})^\ell. \quad (11)$$

where $\ell \in \mathbb{N}^+$ is the repetition parameter defined above.

- **Parallelization:** *Each prover-message can be computed in $O(1)$ time on a Parallel Random Access Machine (PRAM) with common read and exclusive write (CREW), assuming a single \mathbb{F} arithmetic operation takes unit time.*

We conjecture that soundness is nearly equal to the blockwise distance measure $\delta^{(0)}$, for any distance value $\delta^{(0)}$, even when $\delta^{(0)} \approx 1 - \rho$. The following conjecture formalizes this; see [14] for a discussion of it.

Conjecture B.19. *Using the notation of Theorem B.18, for any $f : L \rightarrow \mathbb{F}$ of blockwise distance $\delta^{(0)}$ from $\text{RS}[\mathbb{F}, L, \rho]$, $|L| = 2^k$ and $\epsilon \in [0, \delta^{(0)}]$, with probability at least*

$$1 - \frac{(k)^2 \cdot 2^\eta}{\epsilon \cdot \eta^2 \cdot |\mathbb{F}|} \quad (12)$$

over the randomness of the verifier during the COMMIT phase, the probability of rejection during the QUERY phase is at least $1 - \delta^{(0)} (1 - \epsilon)$.

B.7 Proof of main theorems

In this section we prove Theorems 3.4 and 3.5 and Lemma B.6.

B.7.1 Proof of Main Theorem 3.4

We start with the following lemma. Its proof follows immediately from the proof of Cook–Levin Theorem, hence we omit it. Often the reduction from a general computations to the relevant BAIR instance may benefit from the ambient field structure, and lead to reductions that are concretely more efficient than a general-purpose reduction. Indeed the BAIR that corresponds to our DPM example is constructed in this way (cf. Appendix E).

Lemma B.20 (Reduction to BAIR). *Given $L \in \text{NTimeSpace}(T(n), S(n))$, there exists a reduction mapping an instance x of L of size n to an (n, t, d) -BAIR instance $x_{\text{BAIR}} = (\mathbb{F}, T, w, \mathcal{P}, C, \mathbf{B})$ in time in time $\text{poly}(n, n)$ (see Remark B.1). The resulting x_{BAIR} satisfies the conditions of Lemma B.6, and furthermore*

$$w = S(n) + O(1), \quad t = \log T(n) + O(1), \quad |\mathbf{B}| = n, \quad T_{\text{arith}}(\mathcal{P}) = O(S(n)) \text{ and } d = O(1). \quad (13)$$

The asymptotic constants above are independent of n .

Given x, x_{BAIR} as above, and a witness w such that $(x, w) \in R_L$, a BAIR witness $w_{\text{BAIR}}(x_{\text{BAIR}}, w_{\text{BAIR}}) \in R_{\text{BAIR}}$ can be computed in time $\text{poly}(T(n) \cdot S(n), n)$. Vice versa, given $x, x_{\text{BAIR}}, w_{\text{BAIR}}$ such that $(x_{\text{BAIR}}, w_{\text{BAIR}}) \in R_{\text{BAIR}}$, a witness w such that $(x, w) \in R_L$ can be computed (or extracted) in time $\text{poly}(T(n) \cdot S(n), n)$.

Proof of Theorem 3.4. Given instance x of L of size n , denote $T = T(n)$, $S = S(n)$ and $\lambda = \lambda(n)$. Set $\mathcal{R} = 3$ and let $\delta = (1 - 2^{-\mathcal{R}}(1 + 2^{-d}))$, for d given by Lemma B.20; notice δ is independent of n (which we have yet to specify). Let k be the smallest positive integer such that $t + k > 6 + \log \lambda + \log \log \frac{1}{1-\delta}$, noticing $k \leq \lambda + O(1)$. Now let $n = \lambda + t + d + \mathcal{R} + k + 3$, noticing $n \leq 2\lambda + t + O(1)$. Apply Lemma B.20 and let x_{BAIR} be the resulting (n, t, d) instance. By construction, this instance satisfies the assumptions of Lemma B.6, so we apply the ZK-IOP of knowledge specified there to x_{BAIR} . Soundness, knowledge extraction and round complexity follow directly from Lemmas B.6 and B.20. Query, verifier, and arithmetic complexity also follow directly from these two lemmas, by using the bounds stated in Equation (13) above. This completes the proof. \square

B.7.2 Proof of Main Lemma B.6

Proof of Lemma B.6. We are given an instance-witness pair $(x_{\text{BAIR}}, w_{\text{BAIR}})$ for R_{BAIR} where $x_{\text{BAIR}} = (\mathbb{F}, T, w, \mathcal{P}, C, \mathbf{B})$ satisfies the assumptions of the the lemma.

Protocol description Prover and verifier apply the following reductions:

1. the deterministic reduction of Theorem B.12: the prover executes $P^{\text{BAIR} \rightarrow \text{APR}}(x_{\text{BAIR}}, w_{\text{BAIR}})$, leading to the witness w_{APR} ; the verifier executes $V^{\text{BAIR} \rightarrow \text{APR}}(x_{\text{BAIR}})$. By Theorem B.12 the resulting instance $x_{\text{APR}} = (\mathbb{F}, \mathcal{T}, \mathcal{N}, \Phi, L, L_{\text{cmp}}, \vec{\rho}, \rho_{\text{cmp}})$ has parameters as stated in Item 4 there; additionally, x_{APR} is 1-overlapping, $2^{-(2+\mathcal{R}+d)}(1 - 2^{-k})$ -independent and has δ -distance for $\delta = (1 - 2^{-\mathcal{R}}(1 + 2^{-d}))$;
2. the 1-round ALL protocol on $(x_{\text{APR}}, w_{\text{APR}})$: Prover sends the oracle $\mathbb{O}_{\text{assignment}}$, then verifier sends public randomness that, with $\mathbb{O}_{\text{assignment}}$, induces a pair of functions $f^{(0)} : L \rightarrow \mathbb{F}, g^{(0)} : L_{\text{cmp}} \rightarrow \mathbb{F}$ that we denote by $w_{\text{RS}}, w'_{\text{RS}}$, respectively. Notice that the parameter ζ defined in Equation (7) equals 5 because x_{APR} is 1-overlapping and $\frac{|L|}{|L_{\text{cmp}}|} = 4$

3. the FRI protocol is now applied to each the two witnesses; Verifier uses instances (i) $x_{RS} = (\mathbb{F}, L, \rho_{\max})$ for w_{RS} , and (ii) $x'_{RS} = (\mathbb{F}, L_{\text{cmp}}, \rho_{\text{cmp}})$ for w'_{RS} . The repetition parameter in both cases is set to

$$\ell = \left\lceil \frac{\lambda + 2}{\log \frac{1}{1-\delta/10}} \right\rceil. \quad (14)$$

Completeness Perfect completeness follows directly from the perfect completeness of the various reductions: If $(x_{\text{BAIR}}, w_{\text{BAIR}}) \in R_{\text{BAIR}}$ then by the completeness part of Theorem B.12 the pair $(x_{\text{APR}}, w_{\text{APR}}) \in R_{\text{APR}}$; so, by the completeness of Theorem B.15 (x_{RS}, w_{RS}) and (x'_{RS}, w'_{RS}) belong to BRPT, hence, by the completeness of Theorem B.18 the verifier accepts both with probability 1.

Soundness Suppose $x_{\text{BAIR}} \notin L_{\text{BAIR}}$. Then by the soundness of Theorem B.12 we also have $x_{\text{APR}} \notin L_{\text{APR}}$. Hence, by the soundness of Theorem B.15 we have with all but probability $\epsilon_{\text{ALI}} = 1/|\mathbb{F}| \leq 2^{-(\lambda+3)}$ have that at least one of $f^{(0)}, g^{(0)}$ is $\delta^{(0)}$ -far from the relevant RS code, where $\delta^{(0)} = \frac{\delta}{2\zeta} = \frac{\delta}{10}$; assume $f^{(0)}$ is it (the analysis for $g^{(0)}$ is identical). Notice that $\delta^{(0)} < \delta_0$ for δ_0 as defined in Theorem B.18 because $\mathcal{R} \geq 2$. Therefore, by Theorem B.18, but for probability $\text{err}_{\text{COMMIT}} = \frac{3|L|}{|\mathbb{F}|} < 2^{-\lambda+3}$ over the randomness of the COMMIT phase, the probability of accepting $f^{(0)}$ during the QUERY phase, conducted with repetition parameter ℓ , is at most

$$\text{err}_{\text{QUERY}} = \left(1 - \frac{\delta}{10}\right)^\ell \leq 2^{-(\lambda+2)}. \quad (15)$$

Summing up, the total probability of error is at most $\text{err}_{\text{ALI}} + \text{err}_{\text{COMMIT}} + \text{err}_{\text{QUERY}} \leq \text{err}$, as claimed.

Knowledge extraction Let P^* be a (not necessarily honest) prover that interacts with the verifier specified by the protocol above, and which leads that verifier to accept with probability $p > 4 \cdot \text{err}$. Therefore, with probability at least $O(1/(p - \text{err}))$, the oracle $\mathbb{O}_{\text{assignment}}^*$ provided by the Prover at the end of step 1 above, leads the verifier to accept with probability $> 2\text{err}$. Fix such an $\mathbb{O}_{\text{assignment}}^*$. By the soundness of the FRI protocol as stated in Theorem B.18 and our setting of parameters, we conclude that that with probability strictly greater than $1/|\mathbb{F}|$ both $\Delta_{\text{H}}(f^{(0)}, \text{RS}[\mathbb{F}, L, \rho]) < \frac{\delta}{2\zeta}$ and $\Delta_{\text{H}}(g^{(0)}, \text{RS}[\mathbb{F}, L_{\text{cmp}}, \rho_{\text{cmp}}]) < \frac{\delta}{2\zeta}$. Noticing $\rho_{\max} = 2^{-\mathcal{R}} \leq 1/4$, the knowledge extractor of Theorem B.15 outputs with high probability a witness w_{APR} that satisfies x_{APR} . Hence, by the knowledge extraction property of Theorem B.12, Item 3, we conclude that a witness w_{AIR} can be extracted from P^* in time $\text{poly}\left(\frac{\mathcal{T}}{p-\text{err}}\right)$ as claimed.

Query complexity The query complexity of the verifier is the sum of queries to the FRI oracles; we use the protocol and set $\eta = 2$ in [14, Theorem 3.3]), leading to round complexity $r = \frac{\log |L| - \mathcal{R}}{2}$ and $2^\eta = 4$ queries to each oracle per test (there are ℓ many tests). Each query to $f^{(0)}$ is simulated by making a single query to each of the w members of w_{APR} and a single query to f_{mask} , and each query to $g^{(0)}$ is simulated by making 3 queries to each member of w_{APR} and a single query to g_{mask} . Using the value of ℓ above, and the equalities $\dim(L) = 2 + \dim(L_{\text{cmp}}) = k + \mathcal{R} + t + d + 2$ from Theorem B.12, total query complexity is thus

$$\begin{aligned} q &= 2^\eta \cdot \ell \cdot (4w + 2 + \dim(L) - \mathcal{R} + \dim(L_{\text{cmp}}) - \mathcal{R}) \\ &= 2^{\eta+1} \cdot \ell \cdot (2w + k + t + d + 2) \\ &= 8 \left\lceil \frac{\lambda + 2}{\log \frac{1}{1-\delta/10}} \right\rceil \cdot (2w + k + t + d + 2) \end{aligned} \quad (16)$$

Verifier arithmetic complexity The arithmetic complexity of the verifier is the sum of four sub-components: (i) the computation of the polynomials arising from the boundary constraints described in Item 4h, at a cost of at most $2|\mathbf{B}|^2$ arithmetic operations (see Footnote 27); (ii) the arithmetic complexity of the FRI verifier, at a cost of $21 \cdot \ell \cdot \log |L|$, (iii) the cost of computing an entry of $f^{(0)}$ from \mathbf{w} , which costs $4(\mathbf{w} + 1)$ operations per query (each test makes $2^\eta = 4$ such queries), and (vi) the computation of $g^{(0)}$, which costs $4(|\Phi| + \mathsf{T}_{\text{arith}}(\Phi) + 1)$, where $\mathsf{T}_{\text{arith}}(\Phi) = \mathsf{T}_{\text{arith}}(\mathcal{P}) + |\mathbf{B}|$; recall that here we also make 4 queries per test. Summing up, total verifier arithmetic complexity is

$$\begin{aligned} \text{tv} &\leq 2|\mathbf{B}|^2 + \ell \cdot (16(\mathbf{w} + |\Phi| + \mathsf{T}_{\text{arith}}(\Phi)) + 2 \cdot 21(\mathbf{k} + \mathcal{R} + \mathbf{t} + \mathbf{d} + 2)) \\ &= 2 \left(|\mathbf{B}|^2 + \left\lceil \frac{\lambda + 2}{\log \frac{1}{1-\delta/10}} \right\rceil \cdot (8(\mathbf{w} + |\Phi| + \mathsf{T}_{\text{arith}}(\mathcal{P}) + |\mathbf{B}|) + 21(\mathbf{k} + \mathcal{R} + \mathbf{t} + \mathbf{d} + 2)) \right) \end{aligned} \quad (17)$$

Prover arithmetic complexity Inspection of the prover arithmetic complexity shows that the components that dominate it are the $3 \cdot \mathbf{w}$ low degree extension computations, each over a space of size at most $2^{\mathbf{k}+\mathbf{t}+\mathcal{R}+\mathbf{d}+2}$. In fact, this is the only phase of the prover that requires arithmetic complexity that is super-linear in $|L|$. The computation of $f^{(0)}$ and $g^{(0)}$ require $|L| \cdot (4|\Phi| + 4\mathbf{w} + \mathsf{T}_{\text{arith}}(\mathcal{P}) + |\mathbf{B}|)$ arithmetic operations and the two invocations of the FRI prover cost $2 \cdot 6 \cdot |L|$ operations. This completes the analysis of prover complexity.

Zero knowledge By Theorem B.12, Item 4h, the instance x_{APR} is κ -independent for $\kappa = 2^{-(2+\mathcal{R}+\mathbf{d})}(1 - 2^{-\mathbf{k}})$ and $|L| = 2^{2+\mathbf{k}+\mathcal{R}+\mathbf{t}+\mathbf{d}}$. By Item 5 of Theorem B.15, if $\kappa|L|$ is greater than the number of queries made into each member w_τ of $\mathbb{O}_{\text{assignment}}$ then the protocol has perfect zero knowledge. We have

$$\kappa|L| = 2^{2+\mathbf{k}+\mathcal{R}+\mathbf{t}+\mathbf{d}-(2+\mathcal{R}+\mathbf{d})}(1 - 2^{-\mathbf{k}}) = 2^{\mathbf{k}+\mathbf{t}}(1 - 2^{-\mathbf{k}}) \geq 2^{\mathbf{k}+\mathbf{t}-1}. \quad (18)$$

The last equality holds because $\mathbf{k} \geq 1$.

A single query to $f^{(0)}$ requires a single query to $f^{(0)}_\tau$ and a single query to $g^{(0)}$ requires 3 queries to each $f^{(0)}_\tau$. Each test of the FRI verifier makes $2^\eta = 4$ queries to either $f^{(0)}$ or $g^{(0)}$ and we use the repetition parameter ℓ defined above. Therefore, the total number of queries into each w_τ is bounded $16 \cdot \ell$. Combining this number with Equation (18) shows that for $\mathbf{t} \geq 4 + \log \ell$ the protocol has perfect zero knowledge, thus, it holds for

$$\mathbf{k} + \mathbf{t} \geq 6 + \log \lambda + \log \log \frac{1}{1 - \frac{\delta}{10}}.$$

□

B.7.3 Proof of Lemma B.7

Proof of Lemma B.7. The proof is essentially the same as that of Lemma B.6, with the following changes to the protocol description. We use a smaller field, as specified in the statement of Lemma B.7. In Step 3 of the protocol above, we repeat the FRI-COMMIT phase more than once on $f^{(0)}, g^{(0)}$, leading to several sets of oracles; let s denote the number of COMMIT repetitions (to be specified below). Additionally, we use different repetition parameters for $f^{(0)}$ and $g^{(0)}$, and both are smaller than specified above. Finally, each QUERY invocation is applied to a randomly selected COMMIT oracle-set, out of s .

Denoting by $\ell_{f^{(0)}}, \ell_{g^{(0)}}$ the repetition parameters for the FRI-QUERY phase for $f^{(0)}, g^{(0)}$, respectively, we set them (and s) thus

$$s = \left\lceil \frac{\lambda + 2}{\mathbf{n} - (10 + 2 \log \mathbf{n})} \right\rceil; \quad \ell_{f^{(0)}} = \left\lceil \frac{\lambda + 2}{\mathcal{R} + \mathbf{d}} \right\rceil; \quad \ell_{g^{(0)}} = \left\lceil \frac{\lambda + 2}{\mathcal{R}} \right\rceil \quad (19)$$

Since the modified protocol follows that of Lemma B.6, completeness and round complexity do not change. The reduced query complexity allows for a smaller (i.e., better) bound on the knowledge parameter k . Recalling the analysis in the proof of Lemma B.6, to achieve zero knowledge it suffices to ensure 2^{k+t-1} is greater than the number of queries to an individual member of w , i.e., greater than

$$2^\eta \cdot (\ell_{f^{(0)}} + 3\ell_{g^{(0)}}) = 4 \cdot \left(\left\lceil \frac{\lambda+2}{\mathcal{R}+d} \right\rceil + 3 \left\lceil \frac{\lambda+2}{\mathcal{R}} \right\rceil \right) \leq 16 \cdot \left\lceil \frac{\lambda+2}{\mathcal{R}} \right\rceil.$$

Thus, this inequality holds when $k+t > 5 + \log \left\lceil \frac{\lambda+2}{\mathcal{R}} \right\rceil$.

The computation of query and verifier follow directly from the definition of $\ell_{f^{(0)}}, \ell_{g^{(0)}}$ above, and gives

$$\begin{aligned} \mathbf{q} &= 2^\eta \left(\ell_{f^{(0)}} \cdot (w+3+k+t+d) + \ell_{g^{(0)}} \cdot (3w+3+k+t+d) \right) \\ &= 4 \cdot \left(\left\lceil \frac{\lambda+2}{\mathcal{R}+d} \right\rceil \cdot (w+3+k+t+d) + \left\lceil \frac{\lambda+2}{\mathcal{R}} \right\rceil \cdot (3w+3+k+t+d) \right) \\ &\leq 8 \left\lceil \frac{\lambda+2}{\mathcal{R}} \right\rceil (2w+3+k+t+d) \end{aligned}$$

$$\mathbf{tv}^{\mathbb{F}} \leq 2|\mathbf{B}|^2 + \left\lceil \frac{\lambda+2}{\mathcal{R}} \right\rceil \cdot (16(w+|\Phi| + \mathsf{T}_{\text{arith}}(\mathcal{P}) + 42 \cdot (k+\mathcal{R}+t+d+2)))$$

Prover complexity accounts for the s repetitions of the FRI-COMMIT phase, and is

$$\mathbf{tp}^{\mathbb{F}} \leq \left(9w(t+d+\mathcal{R}+3) + |\mathcal{P}| + w + \mathsf{T}_{\text{arith}}(\mathcal{P}) + 12 \left\lceil \frac{\lambda+2}{n - (10+2\log n)} \right\rceil \right) \cdot 2^{t+d+\mathcal{R}+3}$$

Soundness error Under Conjecture B.17, with all but probability $\text{err}_{\text{ALI}} = 1/|\mathbb{F}| < 2^{-(\lambda+2)}$ we have that either $f^{(0)}$ has distance at least $\delta_{f^{(0)}} = 1 - 2^{-(\mathcal{R}+d)}$ from the code $\text{RS}[\mathbb{F}, L, \rho_{\text{max}}]$, or $g^{(0)}$ has distance at least $\delta_{g^{(0)}} = 1 - 2^{-\mathcal{R}}$ from $\text{RS}[\mathbb{F}, L_{\text{cmp}}, \rho_{\text{cmp}}]$; both $\delta_{f^{(0)}}, \delta_{g^{(0)}}$ are block-wise distances. Assume $\delta_{g^{(0)}} = 1 - 2^{-\mathcal{R}}$ (the case of $\delta_{f^{(0)}} \geq 1 - \rho_{\text{max}}$ is argued similarly). Fixing $\epsilon = 2^{-10}$ (somewhat arbitrarily) and recalling $\eta = 2$, Conjecture B.19 implies that with all but probability $\text{err}_{\text{COMMIT}} = \frac{2^{10} \cdot k^2}{|\mathbb{F}|}$ over the randomness of the verifier during the COMMIT phases, the ensuing QUERY phase, repeated $\ell_{g^{(0)}}$ times, accepts with probability at most $\delta_{f^{(0)}}^{\ell_{f^{(0)}}} \cdot \delta_{g^{(0)}}^{\ell_{g^{(0)}}}$. Assuming $\delta_{g^{(0)}} \geq 1 - 2^{-\mathcal{R}}$ and simplifying notation to $p = \text{err}_{\text{COMMIT}}, \delta = \delta_{g^{(0)}}$ and $\ell = \ell_{g^{(0)}}$, by the law of conditional probability the error probability is at most

$$\sum_{i=0}^s \binom{s}{i} p^i (1-p)^{s-i} \cdot \left(\frac{i}{s} + \frac{\delta(s-i)}{s} \right)^\ell \leq \sum_{i=0}^s (sp)^i \cdot \left(\frac{i + \delta(s-i)}{s} \right)^\ell \quad (20)$$

We claim that each of the s terms on the right hand side above is bounded by δ^ℓ/s , which is equivalent to claiming

$$s \cdot (s \cdot p)^i \leq \left(\frac{\delta s}{i + (s-i)\delta} \right)^\ell.$$

By the assumption $\delta \in [1/2, 1]$, the base of the exponent on the right hand side is at least $1/2$, so it suffices to show

$$s \cdot (s \cdot p)^i \leq 2^{-\ell}$$

taking logarithms, it suffices to show $i \log \frac{1}{p} - (i+1) \log s \geq \ell$, which we shall do even for $i = 1$. We have $\log \frac{1}{p} \geq n - (10 + 2 \log n)$ because $k \leq n$, and $\log s \leq 1 + \log \lambda \leq \log n$ thus, assuming

$$n - (12 + 4 \log n) \geq \lceil \frac{\lambda + 2}{\mathcal{R}} \rceil$$

which holds for all $n \geq 60$ because $n > \lambda$, we conclude that indeed each summand of Equation (20) is bounded by δ^ℓ/s , hence $\text{err}_{\text{FRI}} \leq \delta^\ell \leq 2^{-(\lambda+2)}$. The total error probability of the protocol is thus at most $\text{err}_{\text{ALI}} + \text{err}_{\text{FRI}} < 2^{-\lambda}$, as claimed. This completes the proof. \square

B.7.4 Proof of ZK-STIK Theorem 3.5

The proof below follows that of Theorem B.13, with minor modifications in parameters; all are due to the different starting point — a pair $(\mathbb{x}_{\text{BPAIR}}, \mathbb{w}_{\text{BPAIR}})$ — and the slightly different parameters of the end point of that particular reduction. We thus point out the parameter settings that differ from those appearing in the proof of Lemma B.6. Then we argue that asymptotic verifier complexity is strictly logarithmic, and prover complexity is strictly $O(T(n) \log^2 T(n))$.

Proof of Theorem 3.5. The proof follows that of Lemma B.6 with the following different parameter choices:

- In Item 1 of the protocol description in the proof of Lemma B.6, both parties apply the deterministic reduction of Theorem B.13: Prover executes $\text{P}^{\text{BPAIR} \rightarrow \text{APR}}(\mathbb{x}_{\text{BPAIR}}, \mathbb{w}_{\text{BPAIR}})$ leading to the witness \mathbb{w}_{APR} and verifier executes $\text{V}^{\text{BPAIR} \rightarrow \text{APR}}(\mathbb{x}_{\text{BPAIR}})$, leading to \mathbb{x}_{APR} .
- The repetition parameter is set differently than in the proof of Lemma B.6, to account for the fact that \mathbb{x}_{APR} is 5-overlapping (not 1-overlapping, as above) and thus $\zeta = 21$ (cf. Equation (7)) so $\delta^{(0)} = \frac{\delta}{42}$. This increases the repetition parameter ℓ defined in Equation (14) to

$$\ell = \lceil \frac{\lambda + 2}{\log \frac{1}{1 - \frac{\delta}{42}}} \rceil$$

so that $\text{err}_{\text{FRI}}(\delta^{(0)})$ will be bounded by $\text{err}/4$ as stated in Equation (15).

- We also have $\dim(L) = 2 + k + \mathcal{R} + t + \lceil \log(t+1) \rceil + d$, i.e., the dimension of L is greater than that used in the proof of Lemma B.6 by an additive factor of $\lceil \log(t+1) \rceil$.
- Each query to $f^{(0)}$ requires now $|\mathcal{T}| = 2(w+1) + 1$ queries to $\mathbb{O}_{\text{assignment}}$ (compared with $w+1$ in the proof above) and each query to $g^{(0)}$ requires $11|\mathcal{T}| = 22(w+1) + 11$ queries to $\mathbb{O}_{\text{assignment}}$ rather than $3|\mathcal{T}|$ as before.

Query complexity Using these parameters in the computation of query complexity requires increasing the outer most factor from 36 to 156, to account for the factor $4\frac{1}{3} = \frac{39}{9}$ increase in repetition parameter, and replacing w with $25w + 36$ in (16). This give

$$q = 156 \cdot (\lambda + 2) \cdot (25w + t + d + \lceil \log(t+1) \rceil) + 38 = O(\lambda \cdot t).$$

The last equality uses $w = O(t)$ as stated in Lemma B.9.

Verifier complexity The arithmetic complexity of the verifier is dominated by the arithmetic complexity of the FRI verifier and the computation of entries of $f^{(0)}$ and $g^{(0)}$, which should be repeated $O(\ell) = O(\lambda)$ many times. Consequently, and using the bounds on t , $d = O(1)$, $|\mathcal{P}_\tau \cup \mathcal{P}_\pi|$, $T_{\text{arith}}(\mathcal{P}) = O(\log T(n))$, and $|\mathbf{B}| = O(n)|\mathbf{B}|$, we conclude verifier arithmetic complexity is $O(\lambda \cdot t)$ as well.

Prover complexity Asymptotic prover complexity is similar to that discussed in the proof of Lemma B.6, with two important differences. First, our prover needs to compute a routing on a back-to-back De Bruijn network, costing $O(T(n) \log T(n))$ operations (see Theorem G.3); second, the dimension of L (and of L_{cmp}) is larger by an additive factor of $\lceil \log(t + 1) \rceil$. Therefore, the LDE computations are now over spaces of size $O(T(n) \log T(n))$ and total arithmetic complexity of the prover is $O(T(n) \log^2 T(n))$. \square

B.8 Realization considerations

Our code realization for the DPM described in Figure 1 (cf. Appendix A) uses the verifier-side parameter settings of Lemma B.7. The IOP is converted to a transparent argument system with computational zero knowledge, via the Kilian/Micali cryptographic compiler (cf. Section 2.5). Our code for the prover realizes *both reductions* described in Theorems 3.4 and 3.5. Our focus here is on the DPM program that requires small space, hence the former system is more efficient in this case.

For the hash function used to construct Merkle trees as commitments to oracles, we use the Davies–Meyer construction instantiated with AES128. This gives an estimated collision resistance parameter of 64 bits.

Using the notation in the proof of Lemma B.7, we set $\mathcal{R} = 3$ and $\lambda = 60$, leading to security error of at most $\text{err} \leq 2^{-60}$. The binary field we use has $|\mathbb{F}| = 2^{64}$, i.e., $n = 64$. The degree of our constraint system is 8 thus $d = 3$. We fix the ZK parameter to $k = 1$; this ensures zero knowledge for $t \geq 8$, and this is obtained for dataset sizes $n \geq 2^3$ because $t = \lceil n \cdot 62 \rceil$. For all smaller datasets, setting $k = 3$ would suffice for ZK but since our focus is on large datasets we did not implement this.

The repetition parameter for the number of FRI-COMMIT is $s = 2$, and for the FRI – QUERY we have $\ell_{f^{(0)}} = 9$ and $\ell_{g^{(0)}} = 22$, as follows from Equation (19).

C Algebraic placement and routing (APR) reduction

In this section we prove Theorems B.12 and B.13, sequentially. The proof of Theorem B.12 starts with a description of the reduction in Appendix C.1, followed by a proof that it satisfies the conditions and parameters of Theorem B.12, given in Appendix C.2. Similarly, the proof of Theorem B.13 starts with a description of the reduction in Appendix C.3, followed by the proof of Theorem B.13 in Appendix C.4.

Notation In this section we view \mathbb{F} as $\mathbb{F}_2[\mathbf{g}]/h(\mathbf{g})$ for a primitive polynomial h (see Remark B.1). For $a, b \in \mathbb{F}$ viewed as polynomials over \mathbb{F}_2 , the notation $a \% b$ represents the remainder of a divided by b in the ring of polynomials $\mathbb{F}_2[\mathbf{g}]$.

C.1 The APR reduction for space bounded computation

C.1.1 Common definitions

We first start by defining a few objects used in the construction:

- $\zeta \in \mathbb{F}_2[\mathbf{g}]$ is a primitive polynomial of degree t , notice ζ is a member of \mathbb{F} because $t < n$
- $H \subset \mathbb{F}$ is the space spanned by $\{\mathbf{g}^k \mid 0 \leq k < t\}$
- $H_0 \subset H$ is the subspace spanned by $\{\mathbf{g}^k \mid 0 \leq k < t - 1\}$
- $H_1 \subset H$ is the affine space $H_1 \triangleq H_0 + \mathbf{g}^{t-1}$
- Given an affine space $S \subset \mathbb{F}$ we define the vanishing polynomial $Z_S \in \mathbb{F}[x]$ to be the monic polynomial of degree $|S|$ such that $\forall s \in S, Z_S(s) = 0$
- For every $j \in [w]$ we define $Z_{\mathbf{B},j}, \mathcal{E}_{\mathbf{B},j} \in \mathbb{F}[x]$ by $Z_{\mathbf{B},j}(x) \triangleq \prod_{(i,j,\alpha) \in \mathbf{B}} (x - (\mathbf{g}^i \% \zeta(\mathbf{g})))$ and $\mathcal{E}_{\mathbf{B},j}$ is the polynomial of minimal degree such that for every $(i, j, \alpha) \in \mathbf{B}, \mathcal{E}_{\mathbf{B},j}(\mathbf{g}^i \% \zeta(\mathbf{g})) = \alpha$
- Given a neighbor $(\tau, N) \in \mathcal{N}$ we denote by $\widetilde{(\tau, N)}$ the expression $\widetilde{(\tau, N)} \triangleq X_{(\tau, N)} \cdot Z_{\mathbf{B},\tau}(N(X_{\text{loc}})) + \mathcal{E}_{\mathbf{B},\tau}(N(X_{\text{loc}}))$. We extend this notation further more, and denote by $(\tau_1, \tau_2, \dots, \tau_k, N)$ the sequence $(\widetilde{(\tau_1, N)}, \widetilde{(\tau_2, N)}, \dots, \widetilde{(\tau_k, N)})$.

C.1.2 Instance reduction

We now describe the (deterministic) verifier-side reduction $\mathcal{V}^{\text{BAIR} \rightarrow \text{APR}}$:

1. $\mathcal{T} \triangleq [w]$
2. $\mathcal{N} \triangleq \{(\tau, n^{\text{id}}), (\tau, n_0^{\text{cyc}}), (\tau, n_1^{\text{cyc}}) \mid \tau \in \mathcal{T}\}$ where (i) $n^{\text{id}}(x) \triangleq x$, (ii) $n_b^{\text{cyc}}(x) \triangleq gx + b\zeta$ for $b \in \{0, 1\}$.
3. For every $j \in [w]$ denote by $q_j(x, y) \triangleq xZ_{\mathbf{B},j}(y) + \mathcal{E}_{\mathbf{B},j}(y)$. Define for each $P \in \mathcal{P}$ the set Φ_P by:

$$\Phi_P \triangleq \left\{ \begin{array}{l} \frac{X_{\text{loc}}(X_{\text{loc}}-1)}{Z_{H_0}(X_{\text{loc}})} \cdot P \left(\widetilde{(1, \dots, w, n^{\text{id}})}, \widetilde{(1, \dots, w, n_0^{\text{cyc}})} \right), \\ \frac{1}{Z_{H_1}(X_{\text{loc}})} \cdot P \left(\widetilde{(1, \dots, w, n^{\text{id}})}, \widetilde{(1, \dots, w, n_1^{\text{cyc}})} \right) \end{array} \right\}$$

Finally, we define $\Phi \triangleq \bigcup_{P \in \mathcal{P}} \Phi_P$

4. Define L to be the affine space $L \triangleq \text{span}(\{g^{1+k+\mathcal{R}+t+d} \cdot (1+g)\} \cup \{g^i \mid 0 \leq i < 1+k+\mathcal{R}+t+d\}) + g^{1+k+\mathcal{R}+t+d}$
5. Define L_{cmp} to be the affine space $L_{\text{cmp}} \triangleq \text{span}\{g^i \mid 0 \leq i < k+\mathcal{R}+t+d\} + g^{1+k+\mathcal{R}+t+d}$
6. For every $j \in [w]$ define ρ_j to be $\rho_j \triangleq \frac{2^{k+t} - \deg(Z_{B,j})}{|L|}$
7. Define $\rho_{\text{cmp}} \triangleq \frac{1+2^{k+t+d}}{|L_{\text{cmp}}|}$

C.1.3 Witness reduction

Finally, we describe the prover-side (randomized) reduction $\text{P}^{\text{BAIR} \rightarrow \text{APR}}$, denoting the (input) AIR witness members by $\{w_j^{\text{AIR}}\}_j$ and the (output) APR members by $\{w_j^{\text{APR}}\}_j$. We show the construction of w_j^{APR} for every $j \in [w]$ independently. Draw uniformly random a polynomial $Q_j: \mathbb{F} \rightarrow \mathbb{F}$ of degree less than 2^{k+t} such that for every $i \in [T]$ it satisfies $Q_j(g^i \circ \zeta) = w_j^{\text{AIR}}[i]$. We define the APR witness by $w_j^{\text{APR}}(x) \triangleq \frac{Q_j(x) - \mathcal{E}_{B,j}(x)}{Z_{B,j}(x)}$, for every $x \in L$. Notice this division is well defined, as all the roots of $Z_{B,j}$ are in H by definition, and $H \cap L = \emptyset$ by construction.

C.2 Proof of Theorem B.12 for space bounded computation

C.2.1 Proof of completeness (Item 1)

Proof. We show that: (i) for all $j \in [w]$, $w_j^{\text{APR}} \in \text{RS}[\mathbb{F}, L, \rho_j]$, and (ii) $\forall \phi \in \Phi, \phi_{\mathcal{N}}[\hat{w}^{\text{APR}}] \in \text{RS}[\mathbb{F}, L_{\text{cmp}}, \rho_{\text{cmp}}]$.

Assignment code membership: For every $j \in [w]$, $\deg(Q_j - \mathcal{E}_{B,j}) < 2^{k+t}$, and by definition of satisfaction of AIR $Z_{B,j} \mid (Q_j - \mathcal{E}_{B,j})$, thus $\frac{Q_j - \mathcal{E}_{B,j}}{Z_{B,j}}$ is a polynomial of degree less than $2^{k+t} - \deg(Z_{B,j})$, showing $w_j^{\text{APR}} \in \text{RS}[\mathbb{F}, L, \rho_j]$ as required.

Constraint code membership: We notice first that for any $j \in [w]$ and any neighbor $(j, N) \in \mathcal{N}$, $\widetilde{(j, N)} = Q_j \circ N$ (denoting by $f \circ g$ the function $(f \circ g)(x) \triangleq f(g(x))$) thus it is sufficient to show that for any $P \in \mathcal{P}$ both rational functions:

$$\frac{x(x-1)}{Z_{H_0}(x)} \cdot P(Q_1(x), \dots, Q_w(x), Q_1(gx), \dots, Q_w(gx)) \quad (21)$$

$$\frac{1}{Z_{H_1}(x)} \cdot P(Q_1(x), \dots, Q_w(x), Q_1(gx - \zeta), \dots, Q_w(gx - \zeta)) \quad (22)$$

are polynomials of degree less than 2^{k+t+d} . We notice it is sufficient to show the denominator indeed divides the numerator in each of those rational functions. For that we use the observations: (i) $H = \{g^i \circ \zeta \mid 0 < i < T-1\} \cup \{0, 1\}$, and (ii) $0, 1 \in H_0$

For Equation (21) it is sufficient to show that $P(Q_1(x), \dots, Q_w(x), Q_1(gx), \dots, Q_w(gx)) = 0$ for every $g^i \in H_0 \setminus \{0, 1\}$. The claim follows by the fact that for any $g^i \in H_0$, we have $g^{i+1}\% \zeta = g^{i+1}$, thus by construction

$$P(Q_1(g^i), \dots, Q_w(g^i), Q_1(g^{i+1}), \dots, Q_w(g^{i+1})) \quad (23)$$

$$= P(w_1^{\text{AIR}}[i], \dots, w_w^{\text{AIR}}[i], w_1^{\text{AIR}}[i+1], \dots, w_w^{\text{AIR}}[i+1]) = 0 \quad (24)$$

Where the last equation follows by definition of satisfaction in AIR.

For Equation (22) the claim follows similarly, by noticing $gx - \zeta = gx\% \zeta$ for all $x \in H_1$. \square

C.2.2 Proof of soundness (Item 2)

Proof. Assume $\mathbb{V}^{\text{BAIR} \rightarrow \text{APR}}(\mathfrak{x}, k, \mathcal{R}) \in \text{APR}$ for $\mathfrak{x} = (\mathbb{F}, \mathbb{T}, w, \mathcal{P}, \mathbf{C}, \mathbf{B})$, and let \hat{w}^{APR} be a witness for it. We show $\mathfrak{x} \in \text{BAIR}$ by constructing a witness w^{BAIR} for it. We define for every $j \in [w]$ the mapping $w_j^{\text{BAIR}}: [\mathbb{T}] \rightarrow \mathbb{F}$ by

$$w_j^{\text{BAIR}}[i] \triangleq w_j^{\text{APR}}(g^i\% \zeta) \cdot Z_{\mathbf{B},j}(g^i\% \zeta) + \mathcal{E}_{\mathbf{B},j}(g^i\% \zeta) \quad (25)$$

denoting by w_j^{APR} as the low degree extension of the actual witness. We show w^{BAIR} satisfies \mathfrak{x} . Denote by $q_j(x, y) \triangleq x \cdot Z_{\mathbf{B},j}(x) + \mathcal{E}_{\mathbf{B},j}(x)$.

Boundary constraints: w^{BAIR} satisfies the boundary constraints \mathbf{B} by construction, as having $(i, j, \alpha) \in \mathbf{B}$ implies $Z_{\mathbf{B},j}(g^i\% \zeta) = 0$ and $\mathcal{E}_{\mathbf{B},j}(g^i\% \zeta) = \alpha$, thus $w_j^{\text{BAIR}}[i] = \alpha$ as required.

Consistency with \mathcal{P} : We show that for any $i \in [\mathbb{T} - 1]$ and for any $P \in \mathcal{P}$:

$$P(w_1^{\text{BAIR}}[i], \dots, w_w^{\text{BAIR}}[i], w_1^{\text{BAIR}}[i+1], \dots, w_w^{\text{BAIR}}[i+1]) = 0$$

Assume this is not the case, thus there is some i and some $P \in \mathcal{P}$ for which the above does not hold, we show in such case w^{APR} can not be a witness of $\mathbb{V}^{\text{BAIR} \rightarrow \text{APR}}(\mathfrak{x}, k, \mathcal{R})$. Denote by $(\mathbb{F}, \mathcal{T}, \mathcal{N}, \Phi, L, L_{\text{cmp}}, \vec{\rho}, \rho_{\text{cmp}}) = \mathbb{V}^{\text{BAIR} \rightarrow \text{APR}}(\mathfrak{x}, k, \mathcal{R})$. Assume $w_j^{\text{APR}} \in \text{RS}[\mathbb{F}, L, \rho_j]$, and we show $\exists \phi \in \Phi$ such that $\phi_{\mathcal{N}}[\hat{w}^{\text{APR}}] \notin \text{RS}[\mathbb{F}, L_{\text{cmp}}, \rho_{\text{cmp}}]$. Without loss of generality assume $g^i\% \zeta \in H_0 \setminus \{0, 1\}$, we show the following constraint is not a member of $\text{RS}[\mathbb{F}, L_{\text{cmp}}, \rho_{\text{cmp}}]$:

$$\phi(x) = \frac{x(x-1)}{Z_{H_0}(x)} \cdot P\left(q_1(w_1^{\text{APR}}(x), x), \dots, q_w(w_w^{\text{APR}}(x), x), q_1(w_1^{\text{APR}}(gx), gx), \dots, q_w(w_w^{\text{APR}}(gx), gx)\right)$$

Otherwise

$$\phi(x) \cdot Z_{H_0}(x) = x(x-1) \cdot P\left(q_1(w_1^{\text{APR}}(x), x), \dots, q_w(w_w^{\text{APR}}(x), x), q_1(w_1^{\text{APR}}(gx), gx), \dots, q_w(w_w^{\text{APR}}(gx), gx)\right) \quad (26)$$

for every $x \in L_{\text{cmp}}$, where both sides, as codewords over L_{cmp} , are of rate at most

$$\rho_{\text{cmp}} + \frac{|H_0|}{|L_{\text{cmp}}|} \leq \frac{2^{k+t+d} + 2^t}{2^{k+\mathcal{R}+t+d}} \leq 2^{-\mathcal{R}} (1 + 2^{-k-d}) < 1$$

But they don't agree over the entire field \mathbb{F} , as for $g^i \zeta$ the left hand of the equation vanishes as $Z_{H_0}(g^i \zeta) = 0$, while the right hand does not vanish as

$$P \left(q_1(w_1^{\text{APR}}(g^i), g^i), \dots, q_w(w_w^{\text{APR}}(g^i), g^i), \right. \\ \left. q_1(w_1^{\text{APR}}(g^{i+1}), g^{i+1}), \dots, q_w(w_w^{\text{APR}}(g^{i+1}), g^{i+1}) \right) = \quad (27)$$

$$P \left(w_1^{\text{BAIR}}[i], \dots, w_w^{\text{BAIR}}[i], w_1^{\text{BAIR}}[i+1], \dots, w_w^{\text{BAIR}}[i+1] \right) \neq 0 \quad (28)$$

The contradiction follows by showing both sides are polynomials of degree less than $\rho_{\text{cmp}}|L_{\text{cmp}}| + |H_0|$. $\deg(\phi \cdot Z_{H_0})$ is low enough by assumption. We assumed $w_j^{\text{APR}} \in \text{RS}[\mathbb{F}, L, \rho_j]$, thus

$$\deg \left(q_j \left(w_j^{\text{APR}}(x), x \right) \right) < 2^{k+t}$$

thus the right hand of Equation (26) is of degree at most

$$\deg < 2^{k+t+d} + 2 < 2^{k+t+d} + |H_0| = \rho_{\text{cmp}}|L_{\text{cmp}}| + |H_0|$$

concluding the proof. \square

C.2.3 Knowledge extraction (Item 3)

Proof. The knowledge extractor is described in Equation (25), and the soundness proof in Appendix C.2.2 shows it is indeed an extractor. \square

C.2.4 Instance properties (Item 4h)

Arithmetic complexity of Φ : While most of the summands in Item 4c are straightforward, the only non trivial dependency is the dependency in t . It is a known fact that for every linear space V , the polynomial Z_V vanishing over V is constructible in time $\text{poly}(\dim V)$, and has exactly $\dim V$ nonzero coefficients, although its degree is $|V|$. Moreover, for every constant $c \in \mathbb{F}$, the vanishing polynomial over $V + c$ is exactly $Z_{V+c} = Z_V + Z_V(c)$. Given all this, we conclude that the values of $Z_{H_0}(X_{\text{loc}})$, $Z_{H_1}(X_{\text{loc}})$ can be all computed using a single evaluation of $Z_{H_0}(X_{\text{loc}})$ and addition of a single constant value to compute $Z_{H_1}(X_{\text{loc}})$. In total, this single computation requires $3t$ multiplications and additions in the field. Additional optimization is possible when representing Z_{H_0} by a binary matrix, eliminating completely multiplications over the field from the evaluation process.

\mathbb{X}_{APR} is $(2^{-(2+\mathcal{R}+d)}(1 - 2^{-k}))$ -independent:

Proof. The witness reduction samples each Q_j from a $(2^{k+t} - 2^t)$ independent space, so by construction we notice w_j^{APR} is sampled from a $(2^{k+t} - 2^t)$ independent space as well. The claim follows by noticing $(2^{-(2+\mathcal{R}+d)}(1 - 2^{-k}))|L| = (2^{k+t} - 2^t)$. \square

\mathbb{X}_{APR} has $(1 - 2^{-\mathcal{R}}(1 + 2^{-d}))$ -distance:

Proof. $\rho_{\text{cmp}} \leq 2^{-\mathcal{R}} < 2^{-\mathcal{R}}(1 + 2^{-d})$, and $\rho_{\text{max}} \leq 2^{-(2+\mathcal{R}+d)} < 2^{-\mathcal{R}}(1 + 2^{-d})$. It is sufficient to show the existence of the code C from Item 2. We define C to be the set of mappings $w: L \rightarrow \mathbb{F}$ which are rational functions of the form $w(x) = \frac{p(x)}{q(x)}$ where $p(x)$ is a polynomial of degree less than 2^{t+d} and $q(x) \triangleq \prod_{\alpha \in H \setminus \{0,1\}} (x - \alpha)$. w is well defined over L , as none of the zeros of $q(x)$ are in L , and the distance of the

code C is the same as the distance of RS $[\mathbb{F}, L, \frac{2^{t+d}}{|L|}] = \text{RS}[\mathbb{F}, L, 2^{-(2+k+\mathcal{R})}]$, which is $1 - 2^{-(2+k+\mathcal{R})}$, and in particular grater than $1 - 2^{-\mathcal{R}} (1 + 2^{-d})$. We conclude the proof by showing $\text{RS}[\mathbb{F}, L, \rho_{\max}] \subset C$ by noticing it is exactly the sub-code where $p(x)$ is restricted to be of the form $p(x) = q(x) \cdot v(x)$, where v is a polynomial of degree less than 2^{k+t} . \square

\times_{APR} is 1-overlapping:

Proof. Define the following 3 disjoint affine subspaces of L :

- $S_x \triangleq L_{\text{cmp}} = \text{span} \{g^i \mid 0 \leq i < k + \mathcal{R} + t + d\} + g^{1+k+\mathcal{R}+t+d}$
- $S_{gx} \triangleq \text{span} \{g^{i+1} \mid 0 \leq i < k + \mathcal{R} + t + d\} + g^{2+k+\mathcal{R}+t+d}$
- $S_{gx-\zeta} \triangleq \text{span} \{g^{i+1} \mid 0 \leq i < k + \mathcal{R} + t + d\} + (1 + g^{2+k+\mathcal{R}+t+d})$

The claims follows by the observation that for any $N(x) \in \{x, gx, gx - \zeta\}$, and for any $z \in L_{\text{cmp}}$, $N(z) \in S_{N(x)}$. \square

C.3 The APR reduction for general computation

C.3.1 Common definitions

We expand the definition in Appendix C.1.1 (ζ, H). For the sake of completeness we provide those definitions again, after providing the definitions introduced first in this section. The definitions of $Z_{B,\tau}$, $\mathcal{E}_{B,\tau}$ similar to those defined in Appendix C.1.1, providing the same semantic purpose, but syntactically differ, fitting the construction described in this section. We first start by defining a few objects used in the construction:

- $\xi \in \mathbb{F}_2[g]$ is a primitive polynomial of degree $\lceil \log(t+1) \rceil$, notice ξ is a member of \mathbb{F}
- $W \subset \mathbb{F}$ is the space spanned by $\{g^{t+k} \mid 0 \leq k < \lceil \log(t+1) \rceil\}$
 - $W_0 \subset W$ is the subspace spanned by $\{g^{t+k} \mid 0 \leq k < \lceil \log(t+1) \rceil - 1\}$
 - $W_1 \subset W$ is the affine space $W_1 \triangleq W_0 + g^{t+\lceil \log(t+1) \rceil - 1}$
 - we denote by s the element $s \triangleq g^{t+1} \in W$
- Given a neighbor $(\tau, N) \in \mathcal{N}$ we denote by $\widetilde{(\tau, N)}$ the expression $\widetilde{(\tau, N)} \triangleq \widetilde{X_{(\tau, N)}} \cdot Z_{B,\tau}(N(X_{\text{loc}})) + \mathcal{E}_{B,\tau}(N(X_{\text{loc}}))$. We extend this notation further more, and denote by $(\tau_1, \tau_2, \dots, \tau_k, N)$ the sequence $\widetilde{(\tau_1, N)}, \widetilde{(\tau_2, N)}, \dots, \widetilde{(\tau_k, N)}$.

Definition identical or similar to Appendix C.1.1:

- $\zeta \in \mathbb{F}_2[g]$ is a primitive polynomial of degree t , notice ζ is a member of \mathbb{F}
- $H \subset \mathbb{F}$ is the space spanned by $\{g^k \mid 0 \leq k < t\}$
 - $H_0 \subset H$ is the subspace spanned by $\{g^k \mid 0 \leq k < t - 1\}$

- $H_1 \subset H$ is the affine space $H_1 \triangleq H_0 + \mathbf{g}^{t-1}$
- For every $\tau \in \mathcal{T}$ we define $Z_{\mathbf{B},\tau}, \mathcal{E}_{\mathbf{B},\tau} \in \mathbb{F}[x]$ by:
 - if $\tau = (j, 0)$ for $j \in [w]$ then $Z_{\mathbf{B},\tau}(x) \triangleq \prod_{(i,j,\alpha) \in \mathbf{B}} (x - (\mathbf{g}^i \% \zeta - s))$ and $\mathcal{E}_{\mathbf{B},\tau}$ is the polynomial of minimal degree such that for every $(i, j, \alpha) \in \mathbf{B}$, $\mathcal{E}_{\mathbf{B},\tau}((\mathbf{g}^i \% \zeta) + s) = \alpha$
 - otherwise $Z_{\mathbf{B},\tau} = 1$ and $\mathcal{E}_{\mathbf{B},\tau} = 0$

C.3.2 Instance reduction

We now describe $\mathbb{V}^{\text{BPAIR} \rightarrow \text{APR}}$:

1. $\mathcal{T} \triangleq ([w] \cup \{\text{ctrl}\}) \times \{0, 1\}$
2. $\mathcal{N} \triangleq \mathcal{N}_{\mathcal{P}} \cup \mathcal{N}_{\text{routing}}$, with $\mathcal{N}_{\mathcal{P}}, \mathcal{N}_{\text{routing}}$ defined as follows:

Definition of $\mathcal{N}_{\mathcal{P}}$: $\mathcal{N}_{\mathcal{P}} \triangleq \{(\tau, \mathbf{n}^{\text{id}}), (\tau, \mathbf{n}_0^{\text{cyc}}), (\tau, \mathbf{n}_1^{\text{cyc}}) \mid \tau \in \mathcal{T}\}$ where (i) $\mathbf{n}^{\text{id}}(x) \triangleq x$, (ii) $\mathbf{n}_b^{\text{cyc}}(x) \triangleq \mathbf{g}(x - \mathbf{g}^{t+1}) + b\zeta + \mathbf{g}^{t+1}$ for $b \in \{0, 1\}$.

Definition of $\mathcal{N}_{\text{routing}}$: $\mathcal{N}_{\text{routing}} \triangleq \{(\tau, \mathbf{n}_{r,c}^{\text{out},b}) \mid \tau \in \mathcal{T}; b, r, c \in \{0, 1\}\}$ where $\mathbf{n}_{r,c}^{\text{out},b}(x) \triangleq \mathbf{g} \cdot x + r \cdot (\mathbf{g}^t + 1) + c(\mathbf{g}^t \cdot \xi) + b$

3. Φ is defined as $\Phi \triangleq \Phi_{\mathcal{P}_{\mathcal{T}}} \cup \Phi_{\mathcal{P}_{\pi}} \cup \Phi_{\text{routing}}$ where $\Phi_{\mathcal{P}_{\mathcal{T}}}, \Phi_{\mathcal{P}_{\pi}}, \Phi_{\text{routing}}$ defined as follows:

- $\Phi_{\mathcal{P}_{\mathcal{T}}}$ **definition:** Define for every $P \in \mathcal{P}_{\mathcal{T}}$ define

$$\Phi_P \triangleq \left\{ \begin{array}{l} \frac{(X_{\text{loc}} - s)(X_{\text{loc}} - s - 1)}{Z_{H_0 + s}(X_{\text{loc}})} \cdot P \left(((1, 0), \dots, (\widetilde{w}, 0), \mathbf{n}^{\text{id}}), ((1, 0), \dots, (\widetilde{w}, 0), \mathbf{n}_0^{\text{cyc}}) \right), \\ \frac{1}{Z_{H_1 + s}(X_{\text{loc}})} \cdot P \left(((1, 0), \dots, (\widetilde{w}, 0), \mathbf{n}^{\text{id}}), ((1, 0), \dots, (\widetilde{w}, 0), \mathbf{n}_1^{\text{cyc}}) \right) \end{array} \right\}$$

Define $\Phi_{\mathcal{P}_{\mathcal{T}}} \triangleq \bigcup_{P \in \mathcal{P}_{\mathcal{T}}} \Phi_P$.

- $\Phi_{\mathcal{P}_{\pi}}$ **definition:** Define for every $P \in \mathcal{P}_{\pi}$ define

$$\Phi_P \triangleq \left\{ \frac{(X_{\text{loc}} - s)(X_{\text{loc}} - s - 1)}{Z_{H+s}(X_{\text{loc}})} \cdot P \left(((1, 0), \dots, (\widetilde{w}, 0), \mathbf{n}^{\text{id}}), ((1, 1), \dots, (\widetilde{w}, 1), \mathbf{n}^{\text{id}}) \right) \right\}$$

Define $\Phi_{\mathcal{P}_{\pi}} \triangleq \bigcup_{P \in \mathcal{P}_{\pi}} \Phi_P$.

4. Φ_{routing} **definition:** We define the following polynomials used by our construction:

$$q_{\text{eq1}}(x, y) \triangleq x - y \quad (29)$$

$$q_{\text{mv}}(x, y, z) \triangleq q_{\text{eq1}}(x, y)q_{\text{eq1}}(x, z) \quad (30)$$

$$q_{\text{cp}}(x, x', y, y', z, z') \triangleq q_{\text{eq1}}(x', z')q_{\text{eq1}}(x, y) + q_{\text{eq1}}(x', y')q_{\text{eq1}}(x, z) \quad (31)$$

$$S_{0,0}(x) \triangleq \frac{Z_{H_0 + \text{span}\{\mathbf{g}^t\}}(x)}{Z_{H_0 + W_0}(x)} \quad (32)$$

$$S_{1,0}(x) \triangleq \frac{Z_{H_1 + \text{span}\{\mathbf{g}^t\}}(x)}{Z_{H_1 + W_0}(x)} \quad (33)$$

$$S_{0,1}(x) \triangleq \frac{1}{Z_{H_0 + W_1}(x)} \quad (34)$$

$$S_{1,1}(x) \triangleq \frac{1}{Z_{H_1 + W_1}(x)} \quad (35)$$

and define the control constraints system by:

$$\Phi_{\text{ctrl}} \triangleq \left\{ \begin{array}{l} \frac{1}{X_{\text{loc}} - s} q_{\text{eq1}} \left(X_{((\text{ctrl},0), \mathbf{n}^{\text{id}})}, X_{((\text{ctrl},1), \mathbf{n}^{\text{id}})} \right), \\ \frac{1}{Z_{H+s}(X_{\text{loc}})} q_{\text{eq1}} \left(X_{((\text{ctrl},0), \mathbf{n}^{\text{id}})}, X_{\text{loc}} - s \right), \\ \frac{1}{Z_{H+\mathbf{g}^t}(X_{\text{loc}})} q_{\text{eq1}} \left(X_{((\text{ctrl},0), \mathbf{n}^{\text{id}})}, X_{((\text{ctrl},1), \mathbf{n}^{\text{id}})} \right) \end{array} \right\} \quad (36)$$

$$\cup \left\{ S_{r,c}(X_{\text{loc}}) \cdot q_{\text{mv}} \left(X_{((\text{ctrl},l), \mathbf{n}^{\text{id}})}, X_{((\text{ctrl},l), \mathbf{n}_{r,c}^{\text{rout},0})}, X_{((\text{ctrl},l), \mathbf{n}_{r,c}^{\text{rout},1})} \right) \mid l, r, c \in \{0, 1\} \right\} \quad (37)$$

The network flow copy constraints defined by:

$$\Phi_{\text{cp}} \triangleq \left\{ S_{r,c}(X_{\text{loc}}) q_{\text{cp}} \left(\begin{array}{l} \widetilde{((i, l), \mathbf{n}^{\text{id}})}, X_{((\text{ctrl},l), \mathbf{n}^{\text{id}})}, \\ \widetilde{((i, l), \mathbf{n}_{r,c}^{\text{rout},0})}, X_{((\text{ctrl},l), \mathbf{n}_{r,c}^{\text{rout},0})}, \\ \widetilde{((i, l), \mathbf{n}_{r,c}^{\text{rout},1})}, X_{((\text{ctrl},l), \mathbf{n}_{r,c}^{\text{rout},1})} \end{array} \right) \mid i \in [w]; l, r, c \in \{0, 1\} \right\} \quad (38)$$

We define $\Phi_{\text{routing}} \triangleq \Phi_{\text{ctrl}} \cup \Phi_{\text{cp}}$.

5. Define L to be the affine space

$$L \triangleq \text{span} \left(\left\{ \mathbf{g}^{1+k+\mathcal{R}+t+\lceil \log(t+1) \rceil + d} \cdot (1 + \mathbf{g}) \right\} \cup \left\{ \mathbf{g}^i \mid 0 \leq i < 1 + k + \mathcal{R} + t + \lceil \log(t+1) \rceil + d \right\} \right) + \mathbf{g}^{1+k+\mathcal{R}+t+\lceil \log(t+1) \rceil + d}$$

6. Define L_{cmp} to be the affine space

$$L_{\text{cmp}} \triangleq \text{span} \left\{ \mathbf{g}^i \mid 0 \leq i < k + \mathcal{R} + t + \lceil \log(t+1) \rceil + d \right\} + \mathbf{g}^{1+k+\mathcal{R}+t+\lceil \log(t+1) \rceil + d}$$

7. For every $1 \leq j \leq w$ define $\rho_{(j,0)}$ to be $\rho_{(j,0)} \triangleq \frac{2^{k+t+\lceil \log(t+1) \rceil} - \deg(Z_{\mathbf{B},j})}{|L|}$; for any other $\tau \in \mathcal{T}$ we define ρ_τ to be $\frac{2^{k+t+\lceil \log(t+1) \rceil}}{|L|} = 2^{-(2+\mathcal{R}+d)}$

8. Define $\rho_{\text{cmp}} \triangleq \frac{1+2^{k+t+\lceil \log(t+1) \rceil + d}}{|L_{\text{cmp}}|}$

C.3.3 Witness reduction

General overview Let $w = (\hat{w}, \pi)$ be the witness. We construct all functions in

$$w_{\text{APR}} = (\mathbb{W}_{(1,0)}, \mathbb{W}_{(1,1)}, \dots, \mathbb{W}_{(w,0)}, \mathbb{W}_{(w,1)}, \mathbb{W}_{(\text{ctrl},0)}, \mathbb{W}_{(\text{ctrl},1)})$$

as follows:

1. For every $\tau \in \mathcal{T}$ we define a mapping $Q_\tau: H + W \rightarrow \mathbb{F}$
2. For every $\tau \in \mathcal{T}$ we define a mapping $\tilde{Q}_\tau: \mathbb{F} \rightarrow \mathbb{F}$ as the low degree extension of Q_τ
3. For every $\tau \in \mathcal{T}$ we draw a random polynomial $R_\tau: \mathbb{F} \rightarrow \mathbb{F}$ of degree at most $2^{k+t+\lceil \log(t+1) \rceil}$ vanishing on $H + W$ and define $\bar{Q}_\tau \triangleq \tilde{Q}_\tau + R_\tau$
4. the function $w_\tau: L \rightarrow \mathbb{F}$ is the evaluation of $\frac{\bar{Q}_\tau - \mathcal{E}_{B_\tau}}{Z_{B_\tau}}$ over L

Given the above description, it is sufficient to describe the construction of Q_τ for every $\tau \in \mathcal{T}$.

Construction of $Q_{(\text{ctrl},0)}, Q_{(\text{ctrl},1)}$ Given $\pi: [\mathbb{T}] \rightarrow [\mathbb{T}]$ we define the permutation $\pi': \mathbb{F}[g]/\zeta \rightarrow \mathbb{F}[g]/\zeta$ by:

$$\pi'(0) = 0 \tag{39}$$

$$\forall 0 \leq i < \mathbb{T} : \pi'(g^i \% \zeta) = g^{\pi(i)} \% \zeta \tag{40}$$

$Q_{(\text{ctrl},0)}, Q_{(\text{ctrl},1)}$ are an affine embedding of back-to-back De Bruijn routing (Appendix G.2) of π' where $Q_{(\text{ctrl},0)}(x + s) = x$ for every $x \in H$.

Construction of $Q_{(i,b)}$ mappings Any other pair $Q_{(i,0)}, Q_{(i,1)}$ is intuitively an embedding of the same permutation π' as well, and uses exactly the same routing as $Q_{(\text{ctrl},0)}, Q_{(\text{ctrl},1)}$. The sole difference is that the $Q_{(i,0)}$ over $H + s$ is not required to contain only distinct values, but instead satisfies $Q_{(i,0)}((g^j \% \zeta) + s) = w_i[j]$.

Further optimizations The reduction above can be further optimized for concrete settings in the following manner, implemented in our ZK-STARK realization. Often, not all w algebraic registers are needed to verify memory validity; rather, a small number $l \ll w$ suffices. In this case one may save by routing on the rearrangeable network only the l needed registers (cf. [90] for more details).

C.4 Proof of Theorem B.13 for general computation

The construction of the reduction from BPAIR to APR (Appendix C.3) is very similar to the construction of reduction from BAIR to APR (Appendix C), with the main difference of using an embedding of back-to-back De Bruijn routing (Appendix G.2) to represent the permutation in the BPAIR witness. For the sake of completeness we provide below a full proof for this construction.

C.4.1 Proof of completeness (Item 1)

Proof. We show that: (i) for all $\tau \in \mathcal{T}$, $w_\tau^{\text{APR}} \in \text{RS}[\mathbb{F}, L, \rho_\tau]$, and (ii) $\forall \phi \in \Phi$, $\phi_{\mathcal{N}}[\hat{w}^{\text{APR}}] \in \text{RS}[\mathbb{F}, L_{\text{cmp}}, \rho_{\text{cmp}}]$.

Assignment code membership: For every $\tau \in \mathcal{T}$, $\deg(\bar{Q}_\tau - \mathcal{E}_{B,\tau}) < 2^{k+t+\lceil \log(t+1) \rceil}$, and by definition of satisfaction of BPAIR, we have $Z_{B,\tau} \mid (\bar{Q}_\tau - \mathcal{E}_{B,\tau})$, thus $\frac{\bar{Q}_\tau - \mathcal{E}_{B,\tau}}{Z_{B,\tau}}$ is a polynomial of degree less than $2^{k+t+\lceil \log(t+1) \rceil} - \deg(Z_{B,\tau})$, showing $w_\tau^{\text{APR}} \in \text{RS}[\mathbb{F}, L, \rho_\tau]$ as required.

Constraint code membership: We notice first that for any $\tau \in \mathcal{T}$ and any neighbor $(\tau, N) \in \mathcal{N}$, $(\tau, N) = \bar{Q}_\tau \circ N$ is of degree less than $2^{k+t+\lceil \log(t+1) \rceil}$. It is sufficient to show that every rational function $\phi \in \Phi$ is a polynomial of degree at most $2^{k+t+\lceil \log(t+1) \rceil+d}$. We notice that $\phi = \frac{p}{Z_S}$ for some multivariate polynomial p , and some set S . For simplicity we denote by $g \in \mathbb{F}[x]$ the low degree extension of $p_N[\hat{w}^{\text{APR}}]$. It is sufficient to show g is (i) of degree less than $2^{k+t+\lceil \log(t+1) \rceil+d} + |S|$, and (ii) vanishes over S , thus in particular dividable by Z_S .

The set $\Phi_{\mathcal{P}_T}$: for every $P \in \mathcal{P}_T$ denote by

$$P_0(x) \triangleq P\left(\bar{Q}_{(1,0)} \circ \mathbf{n}^{\text{id}}(x), \dots, \bar{Q}_{(w,0)} \circ \mathbf{n}^{\text{id}}(x), \bar{Q}_{(1,0)} \circ \mathbf{n}_0^{\text{cyc}}(x), \dots, \bar{Q}_{(w,0)} \circ \mathbf{n}_0^{\text{cyc}}(x)\right) \quad (41)$$

$$P_1(x) \triangleq P\left(\bar{Q}_{(1,0)} \circ \mathbf{n}^{\text{id}}(x), \dots, \bar{Q}_{(w,0)} \circ \mathbf{n}^{\text{id}}(x), \bar{Q}_{(1,0)} \circ \mathbf{n}_1^{\text{cyc}}(x), \dots, \bar{Q}_{(w,0)} \circ \mathbf{n}_1^{\text{cyc}}(x)\right) \quad (42)$$

$$P'_0(x) \triangleq \frac{(x-s)(x-s-1)}{Z_{H_0+s}(x)} \cdot P_0(x) \quad (43)$$

$$P'_1(x) \triangleq \frac{1}{Z_{H_1+s}(x)} \cdot P_1(x) \quad (44)$$

We notice $\deg P_0, \deg P_1 < 2^{k+t+\lceil \log(t+1) \rceil+d}$, thus it is sufficient to show that whenever the denominator vanishes, the numerator vanishes as well. Assume $Z_{H_0+s}(x)$ or $Z_{H_1+s}(x)$ is zero, then $x = y + s$ for some $y \in H$. If $y \in \{0, 1\}$ then $x \in H_0 + s$, thus only $Z_{H_0+s}(x) = 0$, and $(x-s)(x-s-1) = 0$ as well. Otherwise, there is some $j \in [T-1]$ such that $y = g^j \% \zeta$, and by the witness construction (Appendix C.3.3) $\bar{Q}_{(i,0)}((g^j \% \zeta) + s) = w_i[j]$. In the case $(g^j \% \zeta) + s \in H_0 + s$ we notice $\mathbf{n}_0^{\text{cyc}}((g^j \% \zeta) + s) = (g^{j+1} \% \zeta) + s$. Similarly, in the case $(g^j \% \zeta) + s \in H_1 + s$ it holds $\mathbf{n}_1^{\text{cyc}}((g^j \% \zeta) + s) = (g^{j+1} \% \zeta) + s$. Finally, the claim follows by the assumption the BPAIR witness satisfies the instance, and in particular $P(w_1[j], \dots, w_w[j], w_1[j+1], \dots, w_w[j+1]) = 0$ for all $j \in [T-1]$.

The set $\Phi_{\mathcal{P}_\pi}$: Let $P \in \mathcal{P}_\pi$ be a polynomial, denote by

$$P'(x) \triangleq P\left(\bar{Q}_{(1,0)}(x), \dots, \bar{Q}_{(w,0)}(x), \bar{Q}_{(1,1)}(x), \dots, \bar{Q}_{(w,1)}(x)\right) \quad (45)$$

$$P''(x) \triangleq \frac{(x-s)(x-s-1)}{Z_{H+s}(x)} \cdot P_0(x) \quad (46)$$

We notice $\deg P' < 2^{k+t+\lceil \log(t+1) \rceil+d}$, thus it is sufficient to show that whenever the denominator vanishes, the numerator vanishes as well. Let x be a member of $H + s$. If $x \in \{s, 1+s\}$ then $(x-s)(x-s-1) = 0$, otherwise there is some $j \in [T-1]$ such that $x = (g^j \% \zeta) + s$, thus by construction $P'(x) = P(w_1[j], \dots, w_w[j], w_1[\pi(j)], \dots, w_w[\pi(j)])$, and in particular vanishes by definition of satisfaction in BPAIR.

The set Φ_{ctrl} :

- For $\phi(x) = \frac{1}{Z_{H+s}(x)} q_{\text{eq}}(\bar{Q}_{(\text{ctrl},0)}(x), x-s)$ we have $\deg(q_{\text{eq}}(\bar{Q}_{(\text{ctrl},0)}(x), x-s)) < 2^{k+t+\lceil \log(t+1) \rceil}$ and for every $x \in H+s$, it holds by construction (Appendix C.3) $\bar{Q}_{(\text{ctrl},0)}(x) = x-s$.
- For $\phi(x) = \frac{1}{x-s} q_{\text{eq}}(\bar{Q}_{(\text{ctrl},0)}(x), \bar{Q}_{(\text{ctrl},1)}(x))$ we have $\deg(q_{\text{eq}}(\bar{Q}_{(\text{ctrl},0)}(x), \bar{Q}_{(\text{ctrl},1)}(x))) < 2^{k+t+\lceil \log(t+1) \rceil}$, and $\bar{Q}_{(\text{ctrl},0)}(s) = \bar{Q}_{(\text{ctrl},1)}(s)$ because $\pi'(0) = 0$.
- For $\phi(x) = \frac{1}{Z_{H+g^t}(x)} q_{\text{eq}}(\bar{Q}_{(\text{ctrl},0)}(x), \bar{Q}_{(\text{ctrl},1)}(x))$ we have $\deg(q_{\text{eq}}(\bar{Q}_{(\text{ctrl},0)}(x), \bar{Q}_{(\text{ctrl},1)}(x))) < 2^{k+t+\lceil \log(t+1) \rceil}$, and for every $x \in H+g^t$ it holds $\bar{Q}_{(\text{ctrl},0)}(x) = \bar{Q}_{(\text{ctrl},1)}(x)$, by the property of the back-to-back De Bruijn routing (Appendix G.2).
- For ϕ of the form $S_{r,c}(x) \cdot q_{\text{mv}}(\bar{Q}_{(\text{ctrl},l)}(x), \bar{Q}_{(\text{ctrl},l)} \circ n_{r,c}^{\text{rot},0}(x), \bar{Q}_{(\text{ctrl},l)} \circ n_{r,c}^{\text{rot},1}(x))$. Low degree follows by the properties of the back-to-back De Bruijn routing (Appendix G.2).

The set Φ_{cp} : The low degree of every $\phi \in \Phi_{\text{cp}}$ follows by construction, as all networks $\bar{Q}_{(j,0)}, \bar{Q}_{(j,1)}$ are routed exactly the same way as the embedded back-to-back De Bruijn routing in $\bar{Q}_{(\text{ctrl},1)}, \bar{Q}_{(\text{ctrl},1)}$. \square

C.4.2 Proof of soundness (Item 2)

Proof. Assume $\mathbb{V}^{\text{BPAIR} \rightarrow \text{APR}}(\mathfrak{x}, k, \mathcal{R}) \in \text{APR}$ for $\mathfrak{x} = (\mathbb{F}, \mathbb{T}, w, \mathcal{P}, \mathbf{C}, \mathbf{B})$, and let \hat{w}^{APR} be a witness for it. We show $\mathfrak{x} \in \text{BPAIR}$ by constructing a witness $w^{\text{BPAIR}} = (\hat{w}, \pi)$ for it. We define:

- for every $j \in [w]$ the mapping $w_j^{\text{BPAIR}}: [\mathbb{T}] \rightarrow \mathbb{F}$ by

$$w_j^{\text{BPAIR}}[i] \triangleq w_{(j,0)}^{\text{APR}}((g^i \% \zeta) + s) \cdot Z_{\mathbf{B},(j,0)}((g^i \% \zeta) + s) + \mathcal{E}_{\mathbf{B},j}((g^i \% \zeta) + s) \quad (47)$$

denoting by $w_{(j,0)}^{\text{APR}}$ the low degree extension of the actual witness.

- we define

$$\pi(i) = j \iff w_{(\text{ctrl},1)}^{\text{APR}}((g^i \% \zeta) + s) = (g^j \% \zeta) + s \quad (48)$$

for $i, j \in [\mathbb{T}]$. Otherwise we say $\pi(i) = \infty$.

We show w^{BPAIR} satisfies \mathfrak{x} . Let $\bar{Q}_\tau(x, y) \triangleq x \cdot Z_{\mathbf{B},\tau}(y) + \mathcal{E}_{\mathbf{B},\tau}(y)$.

Boundary constraints: w^{BPAIR} satisfies the boundary constraints \mathbf{B} by construction, as having $(i, j, \alpha) \in \mathbf{B}$ implies $Z_{\mathbf{B},(j,0)}((g^i \% \zeta) + s) = 0$ and $\mathcal{E}_{\mathbf{B},(j,0)}((g^i \% \zeta) + s) = \alpha$, thus $w_j^{\text{BPAIR}}[i] = \alpha$ as required.

General technique for rest of soundness proof: In what follows we show that in any case where π is not a permutation, or at least one of $\mathcal{P}_\mathbb{T}, \mathcal{P}_\pi$ is not satisfied by w^{BPAIR} then there exists $\phi \in \Phi$ such that $\phi_{\mathcal{N}}[\hat{w}^{\text{APR}}] \notin \text{RS}[\mathbb{F}, L_{\text{cmp}}, \rho_{\text{cmp}}]$. The general method we use is assuming by contradiction $\phi_{\mathcal{N}}[\hat{w}^{\text{APR}}] \in \text{RS}[\mathbb{F}, L_{\text{cmp}}, \rho_{\text{cmp}}]$, and denoting by $\psi: \mathbb{F} \rightarrow \mathbb{F}$ its low-degree extension. We notice the evaluation of every such ψ is over L_{cmp} representable as some rational function, implied by definition of ϕ , $\psi(x) = \frac{p(x)}{q(x)}$ where (i) $\deg p < 2^{k+t+\lceil \log(t+1) \rceil} + d$, (ii) $\deg q < 2^{t+\lceil \log(t+1) \rceil}$, and (iii) there exists some x_0 such that $q(x_0) = 0$ while $p(x_0) \neq 0$. In particular it must hold $q(x)\psi(x) = p(x)$ for all $x \in L_{\text{cmp}}$, but both equation sides are polynomial of degree less than $2^{k+t+\lceil \log(t+1) \rceil} + d + 2^{t+\lceil \log(t+1) \rceil} < |L_{\text{cmp}}|$ that do not agree on x_0 , thus they can not agree on any set of size $|L_{\text{cmp}}|$, contradicting their agreement over L_{cmp} .

Satisfaction of \mathcal{P}_T : Assume by contradiction there is $P \in \mathcal{P}_T$ and $j \in [T-1]$ such that $P(w[j], w[j+1]) \neq 0$. In case $g^j \% \zeta \in H_0$ the contradiction is achieved using $P'_0 \in \Phi_{\mathcal{P}_T}$ (Equation (43)), otherwise $g^j \% \zeta \in H_1$ and the contradiction is achieved using $P'_1 \in \Phi_{\mathcal{P}_T}$ (Equation (44)).

π is a permutation: By Theorem G.5 it is sufficient to show $\bar{Q}_{(\text{ctrl},0)}, \bar{Q}_{(\text{ctrl},1)}$ is an affine embedding of a back-to-back De Bruijn of degree t inducing a permutation π' such that (i) the domain of π' is H , (ii) $\pi'(0) = 0$. By Definition G.4 it is sufficient to show:

1. For all $x \in L$ it holds $\bar{Q}_{(\text{ctrl},0)}(x + s) = x$
2. For all $x \in L$ it holds $\bar{Q}_{(\text{ctrl},0)}(x + g^t) = \bar{Q}_{(\text{ctrl},1)}(x + g^t)$
3. For every $r, c, l \in \{0, 1\}$, and every $x \in H_r + (W_c \setminus \{0, g^t\})$, $\bar{Q}_{(\text{ctrl},l)}(x)$ equals to $\bar{Q}_{(\text{ctrl},l)} \circ n_{r,c}^{\text{rout},b}(x)$ for some $b \in \{0, 1\}$
4. $\bar{Q}_{(\text{ctrl},0)}(s) = \bar{Q}_{(\text{ctrl},1)}(s)$

We now now, using the general technique (Appendix C.4.2), that if any of the stated above constraints does not hold there is $\phi \in \Phi$ such that $\phi_{\mathcal{N}}[\hat{w}^{\text{APR}}] \notin \text{RS}[\mathbb{F}, L_{\text{cmp}}, \rho_{\text{cmp}}]$ by providing polynomials $p, q \in \mathbb{F}[x]$ as required by the technique.

1. In case Item 1 does not hold, the contradiction is achieved using $\frac{q_{\text{eq}}(X_{((\text{ctrl},0), \text{id})}, X_{\text{loc}-s})}{Z_{H+s}(X_{\text{loc}})} \in \Phi_{\text{ctrl}}$
2. In case Item 2 does not hold, the contradiction is achieved using $\frac{q_{\text{eq}}(X_{((\text{ctrl},0), \text{id})}, X_{((\text{ctrl},1), \text{id})})}{Z_{H+g^t}(X_{\text{loc}})} \in \Phi_{\text{ctrl}}$
3. In case Item 3 does not hold, the contradiction is achieved using the corresponding polynomial of the form $S_{r,c}(X_{\text{loc}}) \cdot q_{\text{mv}}(X_{((\text{ctrl},l), \text{id})}, X_{((\text{ctrl},l), n_{r,c}^{\text{rout},0})}, X_{((\text{ctrl},l), n_{r,c}^{\text{rout},1})}) \in \Phi_{\text{ctrl}}$
4. In case Item 4 does not hold, the contradiction is achieved using $\frac{q_{\text{eq}}(X_{((\text{ctrl},0), \text{id})}, X_{((\text{ctrl},1), \text{id})})}{X_{\text{loc}-s}} \in \Phi_{\text{ctrl}}$

Satisfaction of \mathcal{P}_π : For every $j \in [w]$, we notice the pair $\bar{Q}_{(j,0)}, \bar{Q}_{(j,1)}$ is routed exactly the same as $\bar{Q}_{(\text{ctrl},0)}, \bar{Q}_{(\text{ctrl},1)}$, as otherwise a contradiction is achievable using Φ_{cp} . Thus we conclude $\bar{Q}_{(\text{ctrl},0)}((g^i \% \zeta) + s) = \bar{Q}_{(\text{ctrl},1)}((g^{\pi(i)} \% \zeta) + s)$ for every $i \in [T-1]$. Assuming there is $P \in \mathcal{P}_\pi$ and $i \in [T-1]$ such that $P(w[i], w[\pi(i)]) \neq 0$, then the contradiction is achieved using the corresponding polynomial in $\Phi_{\mathcal{P}_\pi}$. \square

C.4.3 Knowledge extraction (Item 3)

Proof. The knowledge extractor is described in Equation (47) and Equation (48), and the soundness proof in Appendix C.4.2 shows it is indeed an extractor. \square

C.4.4 Instance properties (Item 4h)

The proof of Item 4h is mostly strait forward, and very similar to the proof of Item 4h in Appendix C.2.4. We provide a full proof only to it being 5-overlapping.

$\mathfrak{x}_{\text{APR}}$ is 5-overlapping:

Proof. Define the following 3 disjoint affine subspaces of L :

- $S_x \triangleq L_{\text{cmp}} = \text{span} \{ \mathbf{g}^i \mid 0 \leq i < k + \mathcal{R} + t + d \} + \mathbf{g}^{1+k+\mathcal{R}+t+d}$
- $S_{g_x} \triangleq \text{span} \{ \mathbf{g}^{i+1} \mid 0 \leq i < k + \mathcal{R} + t + d \} + \mathbf{g}^{2+k+\mathcal{R}+t+d}$
- $S_{g_{x+1}} \triangleq \text{span} \{ \mathbf{g}^{i+1} \mid 0 \leq i < k + \mathcal{R} + t + d \} + (1 + \mathbf{g}^{2+k+\mathcal{R}+t+d})$

We show there are at most 5 neighbor to each space, mapping elements of L_{cmp} to them. We notice that for every $x \in L_{\text{cmp}}$: (i) $n^{\text{id}}(x) \in S_x$, (ii) $n_0^{\text{cyc}}(x), n_{r,c}^{\text{out},b}(x) \in S_{g_x}$ whenever $b + r + c = 0$, and (iii) $n_1^{\text{cyc}}(x), n_{r,c}^{\text{out},b}(x) \in S_{g_{x+1}}$ whenever $b + r + c = 1$. Concluding our proof. \square

D Algebraic linking IOP (ALI)

In this section we describe the second part of the reduction described in Appendix B (see Figure 8), namely, the ALI protocol. This protocol receives a pair $(\mathfrak{x}_{\text{APR}}, \mathfrak{w}_{\text{APR}})$ as described in Appendix B.3 and Definition B.10, uses a single round of interaction in which the verifier sends public randomness, and ends with a pair of instances of the RS proximity testing problem from Definition B.14.

The ALI protocol is described in Appendix D.1. The main properties achieved by it were presented earlier in Theorem B.15, and we prove this Theorem in Appendix D.2. The section ends in Appendix D.3 with a justification for Conjectures B.16 and B.17 that suggest better soundness (and smaller query and verifier complexity) for the ALI protocol.

D.1 The Algebraic Linking IOP (ALI) protocol

The protocol below is a generalization of the “duplex PCP” 2-round IOP protocol from [17, Section 6.2] to the case of κ -independent APR instances. There, the verifier sent randomness to reach a random linear combination of the assignment and a “mask” polynomial. Since we have numerous assignments and constraints, the verifier adds randomness to check a *random* linear combination of *all assignments* and another *random* linear combination to check *all constraints*. The end result of this is that even though the APR witness has many algebraic registers, and a single codeword per register, totaling $|\mathcal{T}| + 1$ many RS-codewords, at the end of the ALI protocol we need to test proximity only for a *pair* of purported RS codewords.

ALI protocol

Input:

- Verifier has a κ -independent instance $\mathfrak{x} = \mathfrak{x}_{\text{APR}} = (\mathbb{F}, \mathcal{T}, \mathcal{N}, \Phi, L, L_{\text{cmp}}, \vec{\rho}, \rho_{\text{cmp}})$
- Prover has \mathfrak{x} and a witness $\mathfrak{w} = \mathfrak{w}_{\text{APR}} = \{ \mathfrak{w}_\tau \in V_\tau, \tau \in \mathcal{T} \}$ satisfying \mathfrak{x} , where each \mathfrak{w}_τ was sampled uniformly from the $\kappa|L|$ -independent space V_τ described in Definition B.11.

Protocol:

1. Prover samples uniformly and independently random
 - $f_{\text{mask}} \in \text{RS}[\mathbb{F}, L, \rho_{\text{max}}]$;
 - $g_{\text{mask}} \in \text{RS}[\mathbb{F}, L_{\text{cmp}}, \rho_{\text{cmp}}]$;

- and sends to verifier the oracle $\mathbb{O}_{\text{assignment}} \triangleq (\mathbb{w}, f_{\text{mask}}, g_{\text{mask}})$
2. Verifier performs the following:
 - (a) sample and send to prover a sequence R of $2|\mathcal{T}| + |\Phi|$ uniformly random \mathbb{F} -elements containing:
 - $r_{\tau,0}, r_{\tau,1}$ for every $\tau \in \mathcal{T}$
 - r_ϕ for every $\phi \in \Phi$
 - (b) define the *random constraint* $\phi_R: \mathbb{F}^\mathcal{V} \rightarrow \mathbb{F}$ by $\phi_R(\alpha) \triangleq \sum_{\phi \in \Phi} r_\phi \cdot \phi(\alpha)$
 - (c) invoke the two following RS-IOPP sub-protocols, accepting if and only if both sub-protocols accept:
 - i. verify proximity of the function $f^{(0)}: L \rightarrow \mathbb{F}$ to the code $\text{RS}[\mathbb{F}, L, \rho_{\text{max}}]$ where

$$\forall x \in L, \quad f^{(0)}(x) \triangleq f_{\text{mask}}(x) + \sum_{\tau \in \mathcal{T}} \left(r_{\tau,0}^\gamma + r_{\tau,1}^\gamma \cdot x^{|L| \cdot (\rho_{\text{max}} - \rho_\tau)} \right) \cdot \mathbb{w}_\tau(x) \quad (49)$$
 - ii. verify proximity of the function $g^{(0)}: L_{\text{cmp}} \rightarrow \mathbb{F}$ to the code $\text{RS}[\mathbb{F}, L_{\text{cmp}}, \rho_{\text{cmp}}]$ where

$$\forall x \in L_{\text{cmp}}, \quad g^{(0)}(x) \triangleq g_{\text{mask}}(x) + \phi_R(\alpha_{\mathbb{w}, \mathcal{N}}(x)) \quad (50)$$

D.2 Proof of the ALI reduction Theorem B.15

Our proof follows the order of items stated in Theorem B.15. As typical for IP and PCP statements, the most intricate parts of our proof are those dealing with soundness, and, to a lesser extent, those discussing zero knowledge and knowledge extraction.

D.2.1 Completeness — Part 2

Suppose $\hat{\mathbb{w}}$ satisfies \times according to Definition B.10. Then by definition of a κ -independent APR, the (random) witness \mathbb{w} sampled by the prover in step 1 of the ALI satisfies \times as well. This means that the following holds:

- $\mathbb{w}_\tau \in \text{RS}[\mathbb{F}, L, \rho_\tau]$ for each $\tau \in \mathcal{T}$ and $f^{(0)}$ is in the linear span of \mathbb{w} ; thus $f^{(0)} \in \text{RS}[\mathbb{F}, L, \rho_{\text{max}}]$.
- $\phi_{\mathcal{N}}[\mathbb{w}] \in \text{RS}[\mathbb{F}, L_{\text{cmp}}, \rho_{\text{cmp}}]$ for each $\phi \in \Phi$ and ϕ_R is in the linear span of $\{\phi_{\mathcal{N}}[\mathbb{w}]\}$, thus $\phi_R \in \text{RS}[\mathbb{F}, L_{\text{cmp}}, \rho_{\text{cmp}}]$ as well, concluding $g^{(0)} \in \text{RS}[\mathbb{F}, L_{\text{cmp}}, \rho_{\text{cmp}}]$.

Therefore, the completeness of the RS-IOPP used in step 2c of the ALI, implies that our verifier accepts the proof above with probability 1, i.e., the ALI protocol has perfect completeness.

D.2.2 Soundness — Part 3

Our proof of soundness requires a few preliminary statements, stated next. The proof of soundness follows in the next sub-section.

Preliminaries The following claim discusses interpolants, using the definition and notation from Appendix B.1.

Claim D.1. *Let $S \subseteq \mathbb{F}$ and $d, k \in \mathbb{N}$ satisfy $d + 2k < |S|$. Given $f: S \rightarrow \mathbb{F}$ define $\hat{f}(x) \triangleq x^k \cdot f(x)$. Suppose that both $\deg(\text{interpolant}^f) < d + k$ and $\deg(\text{interpolant}^{\hat{f}}) < d + k$. Then we also have $\deg(\text{interpolant}^f) < d$.*

Proof. Let $P(X) = \text{interpolant}^{\hat{f}}$ and let $Q(X) \triangleq X^k \cdot P(X)$; both P, Q are viewed as members of $\mathbb{F}[X]$. By assumption $\deg(P) < d + k$, so $\deg(Q) < d + 2k < |S|$. The multi-point evaluation of Q on domain S is precisely the function \hat{f} . The uniqueness of the interpolant, along with the observation $\deg(Q) < |S|$, imply that $\text{interpolant}^{\hat{f}} = Q$. Therefore, the assumption $\deg(\text{interpolant}^{\hat{f}}) < d + k$ gives $\deg(Q) < d + k$. By construction $\deg(P) = \deg(Q) - k$ and this completes the proof. \square

The next lemma says that linear spaces whose members are “close on average” to a linear error correcting code, have small support.

Lemma D.2 (Proximity to codes implies small support). *Let $C \subset \mathbb{F}^S$ be an \mathbb{F} -linear code of blocklength $\leq |\mathbb{F}|$ and relative distance δ . Fix $c > 6/\delta$. Suppose $V \subset \mathbb{F}^S$ satisfies*

$$\Pr_{v \in \text{span}(V)} \left[\Delta_{\text{H}}(v, C) \leq \frac{1}{c} \right] > 1/|\mathbb{F}|.$$

Then there exists $S' \subset S$ of density $\mu(S'/S) \geq 1 - \frac{2}{c} > 1 - \frac{\delta}{3}$ such that $V|_{S'} \subseteq C|_{S'}$.

The proof of the lemma above requires a result from [95] (stated as Lemma 1.6 there). We recall and prove that lemma next, then prove Lemma D.2.

Lemma D.3 (Average distance amplification). *Let $C \subset \mathbb{F}^S$ be a linear space. If $f_1, \dots, f_k \in \mathbb{F}^S$ are such that there exists f_i that is ϵ -far from C in relative Hamming distance, then*

$$\Pr_{r_1, \dots, r_k \in \mathbb{F}} \left[\Delta_{\text{H}} \left(\sum_{i=1}^k r_i f_i, C \right) \leq \epsilon/2 \right] \leq 1/|\mathbb{F}|$$

Proof of Lemma D.2. By Lemma D.3 we have

$$\forall v \in \text{span}(V), \quad \Delta_{\text{H}}(v, C) \leq \frac{2}{c}. \quad (51)$$

Since $2/c < \delta/2$ by assumption, we conclude that the codeword of C that is closest to $v \in \text{span}(V)$ is unique, denote it by \bar{v} . Define $S_v \triangleq \{x \in S \mid v(x) \neq \bar{v}(x)\}$ and for $V' \subseteq \text{span}(V)$ let $S_{V'} = \cup_{v \in V'} S_v$. To prove Lemma D.2 it suffices to show

$$\mu(S_{\text{span}(V)}/S) \leq \frac{2}{c} \quad (52)$$

by setting $S' = S \setminus S_{\text{span}(V)}$.

We prove (52) by way of contradiction, namely, we show that $\mu(S_{\text{span}(V)}/S) > \frac{2}{c}$ and (51) together imply that (51) is false, so (52) holds.

Write $v = \bar{v} + v'$ where v' has relative Hamming weight $\mu(S_{v'}/S)$. Abusing notation, we identify v with v' and henceforth assume the codeword closest to v is $\mathbf{0}$ and that S_v denotes the support of v , i.e., the set of its nonzero entries.

Since (51) implies $\mu(S_v/S) < \delta/3$, if (52) is false then there exists some $V' \subset V$ such that

$$\mu(S_{V'}/S) \in \left(\frac{2}{c}, \frac{4}{c} \right) \subseteq \left(\frac{2}{c}, \delta - \frac{2}{c} \right)$$

The containment follows because $4/c < 2\delta/3 < \delta - 2/c$. By linearity of expectation, the expected support size of a random word in $\text{span}(V')$ is precisely $(1 - 1/|\mathbb{F}|)|S_{V'}|$ which is strictly greater than $(1 - 1/|S_{V'}|)|S_{V'}|$ because $|S_{V'}| < |S| \leq |\mathbb{F}|$. Thus, it must be the case that some $v \in \text{span}(V')$ is fully supported on $S_{V'}$ which means that the relative support size of v is in $(\frac{2}{c}, \delta - \frac{2}{c})$; we conclude v has relative distance $\mu(S_{V'}/S) > 2/c$ from C , contradicting (51) and completing the proof. \square

Soundness analysis

Proof of Item 3 of Theorem B.15. We prove the contrapositive: If neither item 3a nor item 3b of Theorem B.15 hold, which means both of the following items hold:

1. $\Pr \left[\Delta_{\text{H}} \left(f^{(0)}, \text{RS}[\mathbb{F}, L, \rho_{\max}] \right) \leq \frac{\delta}{2\zeta} \right] > 1/|\mathbb{F}|$
2. $\Pr \left[\Delta_{\text{H}} \left(g^{(0)}, \text{RS}[\mathbb{F}, L_{\text{cmp}}, \rho_{\text{cmp}}] \right) \leq \frac{\delta}{2\zeta} \right] > 1/|\mathbb{F}|$

Then $\mathfrak{x} \in \text{APR}$. Details follow.

We apply Lemma D.2 to Item 1 above, while setting the constant c in that lemma to

$$c \triangleq \frac{2\zeta}{\delta} \quad (53)$$

The assumptions of Lemma D.2 hold because $\zeta > 3$ (cf. Equation (7)) and $\delta \leq 1 - \rho_{\max}$, hence $c > 6/(1 - \rho_{\max})$ as required by Lemma D.2. By that lemma we deduce the existence of a set $S \subset L$, $\mu(S/L) \leq \frac{2}{c}$ such that for all $\tau \in \mathcal{T}$ we have both

$$\mathfrak{w}_{\tau}|_{L \setminus S} \in \text{RS}[\mathbb{F}, L, \rho_{\max}]|_{L \setminus S} \text{ and } \left(x^{|L| \cdot (\rho_{\max} - \rho_{\tau})} \cdot \mathfrak{w}_{\tau} \right)|_{L \setminus S} \in \text{RS}[\mathbb{F}, L, \rho_{\max}]|_{L \setminus S}. \quad (54)$$

Let $d = \deg(\mathfrak{w}_{\tau}|_{L \setminus S}) < \rho_{\max} \cdot |L|$ and $k = |L| \cdot (\rho_{\max} - \rho_{\tau})$. We have

$$d + 2k < |L|(\rho_{\max} + 2(\rho_{\max} - \rho_{\tau})) < |L| \cdot 3\rho_{\max} < |L \setminus S|.$$

The last inequality follows from $\rho_{\max} \leq 1/4$ and $|S| \leq 2/c \leq 1/4$ (see Equations (7) and (53)). So by Claim D.1 we conclude

$$\forall \tau \in \mathcal{T} \quad \mathfrak{w}_{\tau}|_{L \setminus S} \in \text{RS}[\mathbb{F}, L, \rho_{\tau}]|_{L \setminus S}. \quad (55)$$

Let $H_0 = \{x \in L_{\text{cmp}} \mid \mathcal{N}(x) \cap S \neq \emptyset\}$; a union bound gives $|H_0| \leq \Theta|S|$. Let w'_{τ} be the low degree extension of $w_{\tau}|_{L \setminus S}$ to domain L , noticing $w'_{\tau} \in \text{RS}[\mathbb{F}, L, \rho_{\tau}]$. Let $\mathfrak{w}' = \{w_{\tau} \mid \tau \in \mathcal{T}\}$. Recall the assumption that \mathfrak{x} has δ -distance, and let C be the linear code of minimal distance δ that contains $\text{RS}[\mathbb{F}, L_{\text{cmp}}, \rho_{\text{cmp}}]$ as required by Item 2). By Equation (55) and Item 2 we conclude $\phi_{\mathcal{N}}[\mathfrak{w}'] \in C$ for each $\phi \in \Phi$. Thus, to complete our soundness analysis, we need only show that $\phi_{\mathcal{N}}[\mathfrak{w}'] \in \text{RS}[\mathbb{F}, L_{\text{cmp}}, \rho_{\text{cmp}}]$ for each $\phi \in \Phi$.

Consider the set of functions

$$\{\phi_{\mathcal{N}}[\mathfrak{w}'] \mid \phi \in \Phi\} \cup \{g_{\text{mask}}\}$$

and let V denote the linear span of this set. Since, by assumption, $\phi_{\mathcal{N}}[\mathfrak{w}']$ agrees with $\phi_{\mathcal{N}}[\mathfrak{w}]$ on $L_{\text{cmp}} \setminus H_0$, we conclude that $g^{(0)}|_{L_{\text{cmp}} \setminus H_0} \in V_{L_{\text{cmp}} \setminus H_0}$. Therefore, if there exists even one member of $\{\phi_{\mathcal{N}}[\mathfrak{w}'] \mid \phi \in \Phi\}$ that does not belong to $\text{RS}[\mathbb{F}, L_{\text{cmp}}, \rho_{\text{cmp}}]$, then Lemma D.3 implies

$$\Pr \left[\Delta_{\text{H}} \left(g^{(0)}, \text{RS}[\mathbb{F}, L_{\text{cmp}}, \rho_{\text{cmp}}] \right) \leq \frac{1}{2} \cdot (\delta - \mu(H_0/L_{\text{cmp}})) \right] \leq 1/|\mathbb{F}|$$

which contradicts Item 2 stated at the beginning of this proof, because $\frac{\delta}{2\zeta} \leq \frac{1}{2} (\delta - \mu(H_0/L_{\text{cmp}}))$ by our choice of ζ in Equation (7). Therefore we conclude that

$$\forall \phi \in \Phi, \quad \phi_{\mathcal{N}}[\mathfrak{w}'] \in \text{RS}[\mathbb{F}, L_{\text{cmp}}, \rho_{\text{cmp}}] \quad (56)$$

and hence \mathfrak{w}' satisfies \mathfrak{x} , completing the soundness analysis. \square

Remark D.4 (Potential for improvement of soundness analysis). *The $\times 2\zeta$ loss in soundness, compared to distance δ , is due to two factors. Lemma D.3 “costs” a $\times 2$ factor, and the union bound used in the proof above incurs another $\times \left(1 + \Theta \frac{|L|}{|L_{\text{cmp}}|}\right)$ loss. It remains an interesting open problem to decide if either factor is actually required (cf. Conjectures B.16 and B.17).*

D.2.3 Knowledge Extraction — Part 4

The proof of Item 4 of Theorem B.15 relies on the following lemma. After proving it, we complete the proof of knowledge extraction. Recall we assume \mathfrak{x} has δ -distance.

Lemma D.5. *Suppose there exists $L'_{\text{cmp}} \subseteq L_{\text{cmp}}$ such that all of the following hold:*

1.
 - $\mu(L'_{\text{cmp}}/L_{\text{cmp}}) > 1 - \delta$
 - $\phi_{\mathcal{N}}[\mathfrak{w}]|_{L'_{\text{cmp}}} \in \text{RS}[\mathbb{F}, L_{\text{cmp}}, \rho_{\text{cmp}}]|_{L'_{\text{cmp}}}$ for every $\phi \in \Phi$
2. For $L' := \mathcal{N}(L'_{\text{cmp}})$ and each $\tau \in \mathcal{T}$:
 - $\mu(L'/L) \geq \sqrt{\rho_{\tau}}$
 - $\mathfrak{w}_{\tau}|_{L'} \in \text{RS}[\mathbb{F}, L, \rho_{\tau}]|_{L'}$

Then $\mathfrak{x} \in \text{APR}$ and the assignment $\mathfrak{w}' := \{\mathfrak{w}'_{\tau} \mid \tau \in \mathcal{T}\}$ where \mathfrak{w}'_{τ} is the low-degree extension of $\mathfrak{w}_{\tau}|_{L'}$ to L witnesses $\mathfrak{x} \in \text{APR}$.

Proof. By definition of \mathfrak{w}' and by assumption 2 we conclude $\mathfrak{w}'_{\tau} \in \text{RS}[\mathbb{F}, L, \rho_{\tau}]$ for all $\tau \in \mathcal{T}$. By assumption 1 we know there is some codeword $w \in \text{RS}[\mathbb{F}, L_{\text{cmp}}, \rho_{\text{cmp}}]$ that is within distance δ of $\phi_{\mathcal{N}}[\mathfrak{w}']$, while $\phi_{\mathcal{N}}[\mathfrak{w}']$ is a codeword of C having relative distance at least δ , as \mathfrak{x} is a δ -distance instance, thus $\phi_{\mathcal{N}}[\mathfrak{w}'] \in \text{RS}[\mathbb{F}, L, \rho_{\tau}]$. Finally we conclude $(\mathfrak{x}, \mathfrak{w}') \in \text{R}_{\text{APR}}$ by definition of R_{APR} . \square

Lemma D.6. *There exists a Las Vegas (randomized) polynomial time algorithm E that satisfies the following condition. Given as input a δ -distance instance \mathfrak{x} with $\rho_{\max} \leq 1/4$ and $\mathfrak{w} \subset (\mathbb{F}^L)^{\mathcal{T}}$ for which there exists L'_{cmp} that satisfies the assumptions of Lemma D.5 with respect to \mathfrak{x} and \mathfrak{w} , the output of E on input $(\mathfrak{x}, \mathfrak{w})$ is a witness $\mathfrak{w}' \subset (\mathbb{F}^L)^{\mathcal{T}}$ that satisfies \mathfrak{x} .*

Proof of Theorem B.15, Part 4. As argued in the proof of soundness above (and using the notation there), if neither item 3a nor item 3b of Theorem B.15 hold, then Equation (55) holds, and moreover, the low degree extension \mathfrak{w}' of $\mathfrak{w}|_{L \setminus S}$ (for S defined there) is a satisfying assignment. Thus, our extractor will find this \mathfrak{w}' (with high probability), using the polynomial time Guruswami-Sudan (GS) list-decoding algorithm [65]. Recall that for $\text{RS}[\mathbb{F}, S, \rho]$, $|S| = n$ and $f : S \rightarrow \mathbb{F}$, the GS algorithm runs in time $\text{poly}(n)$ and outputs a list $L_f \subset \text{RS}[\mathbb{F}, S, \rho]$, $|L_f| = \text{poly}(n)$ that contains every codeword that agrees with f on more than a $\sqrt{\rho}$ -fraction of entries. In other words, when given $f : L \rightarrow \mathbb{F}$ that agrees with some codeword $w \in \text{RS}[\mathbb{F}, L, \rho]$ on more than a $\sqrt{\rho}$ -fraction of entries, the GS algorithm will return w as part of L_f .

Notice that in our case, the assumption $\rho_{\max} \leq 1/4$ along with the upper bound $|S| \leq |L|/4$, which follows from Equations (7) and (53), implies that the set of inputs on which $w \in \mathfrak{w}$ and $w' \in \mathfrak{w}'$ is of size at least $\sqrt{\rho_{\tau}}$. We apply the following extractor.

$E(\mathfrak{x}, \mathfrak{w})$

1. Sample uniformly random $\vec{a} \in \mathbb{F}^{\mathcal{T}}$ and compute $f : L \rightarrow \mathbb{F}$ by $f(x) \triangleq \sum_{\tau \in \mathcal{T}} a_{\tau} w_{\tau}(x)$
2. Let L_f be the output of the GS list-decoding algorithm on f .
3. For each $g \in L_f$,
 - (a) let $S_g = \{x \in L_{\text{cmp}} \mid \forall (\tau, N) \in \mathcal{N}, f(N(x)) = g(N(x))\}$;
 - (b) let $\mathfrak{w}' = \{\mathfrak{w}'_{\tau}\}_{\tau \in \mathcal{T}}$ where \mathfrak{w}'_{τ} is the low-degree extension of $w_{\tau}|_{\mathcal{N}(S_g)}$;

(c) if w' satisfies \times then output w' and return “success”

4. return “fail”

Clearly E runs in polynomial time in its input because $|L_f| = \text{poly}(|L|)$. Additionally, it is a one-sided error algorithm because if it returns “success” then w' indeed satisfies \times . Thus it only remains to analyze the probability of failure, given by the following claim.

Claim D.7. *Suppose L'_{cmp} satisfies all properties of Lemma D.5 and is of maximal size with respect to property 2 of that Lemma. Then*

$$\Pr_{\bar{a}} [\exists g \in L_f, S_g = L'_{\text{cmp}}] \geq 1 - \frac{|L_{\text{cmp}}|\delta}{|\mathbb{F}|}. \quad (57)$$

Assuming the claim, when examining $g \in L_f$ with $S_g = L'_{\text{cmp}}$ we have $w_\tau|_{\mathcal{N}(S_g)} \in \text{RS}[\mathbb{F}, L, \rho_\tau]|_{\mathcal{N}(S_g)}$ for any $\tau \in \mathcal{T}$ by assumption and the low degree extension w'_τ of $w_\tau|_{\mathcal{N}(S_g)}$ witnesses $\times \in \text{APR}$. This completes the proof of Lemma D.6 but for the proof of Claim D.7, which appears next. \square

Proof of Claim D.7. The assumption that L'_{cmp} is maximal with respect to property 2 of Lemma D.5 means that for each $\alpha \in L_{\text{cmp}} \setminus L'_{\text{cmp}}$ there exists $\beta_\alpha \in \mathcal{N}(\alpha)$ and some $\tau_\alpha \in \mathcal{T}$ such that $w_{\tau_\alpha}(\beta_\alpha)$ does not agree with the low-degree extension of $w_{\tau_\alpha}|_{\mathcal{N}(L'_{\text{cmp}})}$ that we shall denote by w'_{τ_α} . Let $e_{\tau_\alpha} : L \rightarrow \mathbb{F}$ be the “error function” related to w'_{τ_α} , defined as $e_{\tau_\alpha}(x) = w'_{\tau_\alpha}(x) - w_{\tau_\alpha}(x)$. We have

$$\Pr_{\bar{a}} \left[f(\beta_\alpha) = \sum_{\tau \in \mathcal{T}} a_\tau w'_\tau(\beta_\alpha) \right] = \Pr_{\bar{a}} \left[\sum_{\tau \in \mathcal{T}} a_\tau e_\tau(\beta_\alpha) = 0 \right] = \frac{1}{|\mathbb{F}|}$$

the last equality holds because by assumption $e_{\tau_\alpha}(\beta_\alpha) \neq 0$. Applying a union bound to $\alpha \in L_{\text{cmp}} \setminus L'_{\text{cmp}}$ we conclude that (57) holds with probability at least $1 - \frac{|L_{\text{cmp}}| - |L'_{\text{cmp}}|}{|\mathbb{F}|} \geq 1 - \frac{|L_{\text{cmp}}|\delta}{|\mathbb{F}|}$, and this completes the proof. \square

D.2.4 Perfect Zero-Knowledge — Part 5

Our proof follows [17, Section 6] but we make the simplifying assumption that the verifier’s first message is the sequence of randomness R mentioned in Step 2a (see Remark D.8). The ALI protocol assumes two RS-IOPP systems, used in steps 2(c)i and 2(c)ii there, each with its own prover and verifier, let P_f, P_g denote the two provers, respectively. Our STIK (and STARK) instantiates both P_f, P_g to be the FRI prover (for different RS code parameters) but our proof of zero knowledge works for *any* choice of RS-IOPP because our simulator uses the relevant RS-IOPP prover(s) in a black-box manner. We assume messages from the verifier have a canonical format that indicates which RS-IOPP prover is being addressed, and which oracle is being queried among $\mathbb{O}_{\text{assignment}}$ and the various oracles produced by the pair of RS-IOPP protocols.

The simulator Our straight-line PZK simulator is denoted Sim . Given a verifier V^* and instance \times , the simulator starts by sampling uniformly random functions $f^{(0)} \in \text{RS}[\mathbb{F}, L, \rho_{\text{max}}]$ and $g^{(0)} \in \text{RS}[\mathbb{F}, L_{\text{cmp}}, \rho_{\text{cmp}}]$ and recording them. The simulator also instantiates the two RS-IOPP provers — P_f and P_g corresponding to steps Item 2(c)i, Item 2(c)ii of the ALI protocol — with $f^{(0)}$ and $g^{(0)}$, respectively. It also invokes V^* and records the first message, which is the randomness R provided by V^* . The simulator now continues to run V^* . All messages and queries directed by V^* to one of the two RS-IOPP protocols (dealing with $f^{(0)}$ and $g^{(0)}$) are managed by Sim by invoking the corresponding IOPP prover(s) P_f, P_g .

To complete the description of Sim we need only explain how it answers queries to $\mathbb{O}_{\text{assignment}} = (\mathbb{w}, f_{\text{mask}}, g_{\text{mask}})$. Recall that \mathbb{w} is a collection of functions, each with domain L , and this is also the domain of f_{mask} ; the domain of g_{mask} is L_{cmp} . As in [17, Section 6], Sim maintains a set of partial functions $\mathbb{w}^*, f_{\text{mask}}^*, g_{\text{mask}}^*$ (with the same domains as $\mathbb{w}, f_{\text{mask}}, g_{\text{mask}}$); all these functions are initialized with $*$ values that indicate “undetermined”. When a function in $\mathbb{O}_{\text{assignment}}$ is queried, Sim answers with the value recorded in $\mathbb{w}^*, f_{\text{mask}}^*, g_{\text{mask}}^*$, if determined, and will otherwise determine it, i.e., change its value from $*$ to some element of \mathbb{F} . The process by which undetermined values get determined is described next:

1. A query $x_0 \in L$ sent to a function $f \in \mathbb{w} \cup \{f_{\text{mask}}\}$, is determined jointly for all functions $f' \in \mathbb{w} \cup \{f_{\text{mask}}\}$. Consider Equation (49). The term on the left hand side, $f^{(0)}(x_0)$, is already fixed by Sim. On the right hand side, the terms

$$\left\{ r_{\tau,0}^\gamma + r_{\tau,1}^\gamma \cdot x_0^{|L| \cdot (\rho_{\text{max}} - \rho_\tau)} \mid \tau \in \mathcal{T} \right\} \quad (58)$$

are all fixed. The only undetermined values are those of $f_{\text{mask}}^*(x_0)$ and $\{\mathbb{w}_\tau^*(x_0) \mid \tau \in \mathcal{T}\}$. Thus, our simulator determines these undetermined values by sampling a uniformly random solution to the linear constraint imposed by $f^{(0)}(x_0)$ and Equation (58).

2. A query $y_0 \in L_{\text{cmp}}$, sent to the function g_{mask} , is determined thus. The left hand side of Equation (50) is already determined by Sim. For all $x_0 \in \mathcal{N}(y_0)$, our simulator determines the value of all $\{\mathbb{w}_\tau^*(x_0) \mid \tau \in \mathcal{T}\}$ and $f_{\text{mask}}^*(x_0)$ using the process described in Item 1. After all such values $\mathbb{w}_\tau^*(x_0)$ are determined for $x_0 \in \mathcal{N}(y_0)$, notice that the rightmost summand of Equation (50) is also determined. Thus, Sim determines $g_{\text{mask}}^*(y_0)$ to be the unique field element that causes the linear constraint of Equation (50) to be satisfied.

Perfect zero knowledge First, notice Sim is straight-line, i.e., it never restarts \mathbb{V}^* . To prove perfect zero knowledge, we shall show that the distribution sampled by Sim interacting with \mathbb{V}^* on a satisfiable instance \mathfrak{x} , is the same as the distribution on transcripts of the interaction between \mathbb{V}^* and an honest prover holding a witness for \mathfrak{x} and operating as described in ALL. Notice the following facts about the distribution supplied by the honest prover:

1. Each \mathbb{w}_τ is sampled uniformly and independently from a $\kappa|L|$ -wise independent space;
2. $\phi_R(\alpha_{\mathbb{w}, \mathcal{N}(y_0)})$ is determined by R and $\{\mathbb{w}_\tau(x_0) \mid x_0 \in \mathcal{N}(y_0), \tau \in \mathcal{T}\}$;
3. the pair $(f_{\text{mask}}, g_{\text{mask}})$ is sampled uniformly from $\text{RS}[\mathbb{F}, L, \rho_{\text{max}}] \times \text{RS}[\mathbb{F}, L_{\text{cmp}}, \rho_{\text{cmp}}]$;
4. consequently, independently of R and \mathbb{w} , the pair $(f^{(0)}, g^{(0)})$ is sampled uniformly from $\text{RS}[\mathbb{F}, L, \rho_{\text{max}}] \times \text{RS}[\mathbb{F}, L_{\text{cmp}}, \rho_{\text{cmp}}]$

Item 4 above relies on the completeness property, which says that if \mathbb{w} satisfies \mathfrak{x} then the rightmost summand of Equation (50) is a codeword of $\text{RS}[\mathbb{F}, L_{\text{cmp}}, \rho_{\text{cmp}}]$.

Consequently, for every fixing of the first verifier message R , and for every subset $S \subset L, |S| < \kappa|L|$, the distribution on $\mathbb{w}|_S, f_{\text{mask}}|_S, f^{(0)}|_S$ generated by the honest ALL prover is the *uniform distribution* on field elements satisfying the linear constraint of Equation (49) for each $x \in S$. By construction, the distribution supplied by Sim invoking \mathbb{V}^* which makes these queries S , is precisely the same distribution.

Next, assume the aforementioned S includes all $x_0 \in \mathcal{N}(S')$, where $S' \subset L_{\text{cmp}}$ is the set of queries made by \mathbb{V}^* to g_{mask} . By Item 2, the distribution on the rightmost term of Equation (50) generated by the

honest ALL prover is the exact same distribution as that supplied by Sim invoking V^* . By Items 3 and 4 above, the distribution on $g^{(0)}|_{S'}, g_{\text{mask}}|_{S'}$ is thus the uniform distribution on field elements satisfying the linear constraint of Equation (50) for every $y_0 \in S'$. By construction, Sim produces the same distribution on $g^{(0)}|_{S'}, g_{\text{mask}}|_{S'}$.

Finally, the distribution of messages between V^* and the sub-provers P_1, P_2 used as part of the RS-IOPP protocols are, by construction, the same distribution as provided by Sim invoking V^* because both the honest ALL prover and the simulator invoke the same sub-provers P_1, P_2 and supply them with the exact same uniformly random inputs $f^{(0)}$ and $g^{(0)}$.

We have shown that the distribution output by the straight-line simulator Sim invoking V^* is equal to the distribution output by V^* interacting with an honest prover on a satisfiable instance. This completes the proof of Item 5 of Theorem B.15.

Remark D.8. *Inspection reveals that the proof of perfect zero knowledge appearing in [17, Section 6] can be adapted to our case (details omitted). That proof is more complicated, as it is designed to address the case where the verifier may query the first oracle ($\mathbb{O}_{\text{assignment}}$) even before sending the randomness R . We point out that in all concrete STARK realizations — both interactive (iSTARK) and non-interactive (nSTARK) — this assumption is unrealistic.*

D.2.5 Arithmetic complexity — Part 6

In this section we prove Item 6 of Theorem B.15. Verifier complexity is as stated because the verifier's only action during the ALL protocol is to sample the randomness R . Regarding prover complexity, we follow the steps of the protocol:

1. In step 1 the prover samples random functions $w, f_{\text{mask}}, g_{\text{mask}}$. For each member of w this requires $3 \cdot |L| \log |L|$ arithmetic operations, as stated in Item 4h of Theorem B.12, and Theorem B.2.
2. In the beginning of Item 2(c)i the prover computes $f^{(0)}$, by performing point wise computation using Equation (49), at a cost of $2|L||\mathcal{T}|$ multiplications, and $2|L|(|\mathcal{T}| + 1)$ additions over \mathbb{F} . The total arithmetic complexity of this part is less than $5|\mathcal{T}| \cdot |L|$.
3. In step 2(c)ii the prover computes $g^{(0)}$ and (its proximity proof is discussed later), by performing a complete point wise computation using $|L_{\text{cmp}}| \cdot (\mathsf{T}_{\text{arith}}(\Phi) + |\Phi|)$ multiplications, and $|L_{\text{cmp}}| \cdot (\mathsf{T}_{\text{arith}}(\Phi) + |\Phi| + 1)$ additions over \mathbb{F} . This accounts for the rightmost summand of Equation (8).

Summing up, the total cost is as stated in Equation (8), and this completes our complexity analysis for the ALL protocol.

D.3 Conjectured soundness

In this section we discuss the rationale behind Conjectures B.16 and B.17, starting with the latter one.

Pure pseudo-prover Informally, Conjecture B.16 captures the intuition that *pure* pseudo-provers achieve the largest soundness-error against general instances. A pure pseudo-prover is one that generates *pure* pseudo-assignments $\mathbb{O}_{\text{assignment}}^* = (w^*, f_{\text{mask}}^*, g_{\text{mask}}^*)$; a pure pseudo-assignment satisfies

$$\text{span}(w^*, f_{\text{mask}}^*) \subseteq \text{RS}[\mathbb{F}, L, \rho_{\text{max}}] \bigvee \text{span}(\{\phi(\alpha_{w^*, \mathcal{N}}) \mid \phi \in \Phi\}, g_{\text{mask}}^*) \subseteq \text{RS}[\mathbb{F}, L_{\text{cmp}}, \rho_{\text{cmp}}].$$

In words, a pure pseudo-assignment selects $\mathbb{O}_{\text{assignment}}^*$ such that either each w_τ^* is a codeword of the relevant code $\text{RS}[\mathbb{F}, L, \rho_\tau]$, in which case the resulting constraint polynomials will be, with very high probability, maximally far from $\text{RS}[\mathbb{F}, L_{\text{cmp}}, \rho_{\text{cmp}}]$, or else P^* fixes the intended values of each $\phi \in \Phi$ to correspond to a member $g_\phi \in \text{RS}[\mathbb{F}, L_{\text{cmp}}, \rho_{\text{cmp}}]$ and then find, for each $y \in L_{\text{cmp}}$, a setting for the values of $w(\mathcal{N}(y))$ so that $\phi(\alpha_{w^*, \mathcal{N}})(y) = g_\phi(y)$; in this case, the resulting w will likely be maximally far from $\text{RS}[\mathbb{F}, L, \rho_{\text{max}}]$.

Mixed pseudo-prover The “attack” imagined to support Conjecture B.17 is the following: on input \mathfrak{x} , the pseudo-prover starts with some $w^* \subseteq \text{RS}[\mathbb{F}, L, \rho_{\text{max}}]$, which leads to $\{\phi(\alpha_{w^*, \mathcal{N}}) \mid \phi \in \Phi\}$ being maximally far from $\text{RS}[\mathbb{F}, L_{\text{cmp}}, \rho_{\text{cmp}}]$. Next, for some ϵ fraction of the entries of L_{cmp} , the attacker “reverse-engineers” a change to an ϵ -fraction of the entries of w^* as to make each member of $\{\phi(\alpha_{w^*, \mathcal{N}}) \mid \phi \in \Phi\}$ ϵ -closer to low-degree. Assuming the spaces L_0, L'_0 are good for \mathfrak{x} , as we do in Conjecture B.17, this modification “ruins” an $\epsilon \cdot 2^\eta \frac{|L_{\text{cmp}}|}{L}$ -fraction of the entries of w^* , as expressed by Equation (9). We stress that we do not know how to efficiently instantiate this attack for general APR instances because the “reverse-engineering” step above may be hard to solve on general instances.

E An algebraic intermediate representation of the DNA profile match

Let us describe the algebraic intermediate representation (abbrev. AIR, see Section 2.2) of the DNA profile match (DPM) program, that uses the Rijndael-160 based hash function. This exposition will show how we achieve the quantities that are specified in Tables 4 and 5 (i.e., the width w , the cycles c , etc.), by providing a “bottom up” description of our algebraic construction.

We first recount the Rijndael block cipher, but from an algebraic perspective (Appendix E.1). We then explain our AIR constraints for the Rijndael cipher (Appendices E.2 and E.3). Following that, we describe the AIR for the transformation from a block cipher to a cryptographic hash function (Appendix E.4). Finally, we show our to implement the AIR of the logic of the DPM program, that performs an exhaustive search to compare the loci pairs that are stored in the elements of a hashchain (Appendix E.5).

The Advanced Encryption Standard (AES) instantiates Rijndael with 128-bit block size and 128-bit, 192-bit, or 256-bit key sizes. We denote by Rijndael-160 the cipher with 160-bit block size and 160-bit key size, and hence output (cipher-text) size of 160 bits. Assuming that Rijndael-160 is an ideal cipher (cf. Appendix E.4), it can be used to build a collision-resistant hash function (CRHF) with an 80-bit security parameter.

It should be noted that Rijndael with 192-bit (256-bit) block and key sizes can be used to build CRHF with a security parameter of 96 bits (128 bits), and that these stronger parameters entail a rather mild overhead in our algebraic construction (see Table 4). However, the STARK construction that we benchmark has 60 bits of security, and therefore the stronger hash functions will not provide better security with our benchmarked system.

Our reference code will be available at <https://github.com/elibensasson/STARK>.

Notation. Let g denote a primitive element of $\mathbb{F}_{2^{64}}$, i.e., $\langle g \rangle = \mathbb{F}_{2^{64}}^*$. We assume throughout that field elements are represented according to a standard basis $(1, g, g^2, g^3, \dots, g^{63})$, rather than a normal basis. We denote by $R_{(t)} \in \mathbb{F}_{2^{64}}$ the content of the algebraic register R at cycle t of the execution (e.g., $K00_{(0)}$ is the field element that resides in the register $K00$ during the first cycle).

E.1 Algebraic description of the Rijndael cipher

The input to the Rijndael cipher can be regarded as a plain-text array of $4n$ elements and a key array of $4n$ elements, such that each element resides in \mathbb{F}_{2^8} . Rijndael executes in $n + 6$ rounds, where each round consists of the following four steps (except for the last round that skips Step 3):

1. **SubBytes** - Given byte x , compute $y = Mx^{-1} + b$, where $M \in \mathbb{F}_{2^8}^{8 \times 8}$, $b \in \mathbb{F}_{2^8}^{8 \times 1}$ are constants.
2. **ShiftRows** - For $i \in \{1, 2, 3, 4\}$, perform i cyclic shifts (rightwards) of the i th row of the plain-text matrix.
3. **MixColumns** - For $j \in 1, 2, \dots, n$, multiply the j th column of the plain-text matrix by the following constant circulant MDS matrix:

$$\begin{bmatrix} P0[j]_{(t+1)} \\ P1[j]_{(t+1)} \\ P2[j]_{(t+1)} \\ P3[j]_{(t+1)} \end{bmatrix} = \begin{bmatrix} \mathbf{g}_0 & \mathbf{g}_1 & 1 & 1 \\ 1 & \mathbf{g}_0 & \mathbf{g}_1 & 1 \\ 1 & 1 & \mathbf{g}_0 & \mathbf{g}_1 \\ \mathbf{g}_1 & 1 & 1 & \mathbf{g}_0 \end{bmatrix} \begin{bmatrix} P0[j]_{(t)} \\ P1[j]_{(t)} \\ P2[j]_{(t)} \\ P3[j]_{(t)} \end{bmatrix}.$$

Here, \mathbf{g}_0 is a specific field element of \mathbb{F}_{2^8} with $|\langle \mathbf{g}_0 \rangle| = 51$, and $\mathbf{g}_1 \triangleq \mathbf{g}_0 + 1$ generates $\mathbb{F}_{2^8}^*$. We note that MixColumns also can be defined as computing the following linear combinations:

$$\begin{aligned} P0[j]_{(t+1)} &= \mathbf{g}_0 \cdot P0[j]_{(t)} + \mathbf{g}_1 \cdot P1[j]_{(t)} + P2[j]_{(t)} + P3[j]_{(t)} \\ P1[j]_{(t+1)} &= P0[j]_{(t)} + \mathbf{g}_0 \cdot P1[j]_{(t)} + \mathbf{g}_1 \cdot P2[j]_{(t)} + P3[j]_{(t)} \\ P2[j]_{(t+1)} &= P0[j]_{(t)} + P1[j]_{(t)} + \mathbf{g}_0 \cdot P2[j]_{(t)} + \mathbf{g}_1 \cdot P3[j]_{(t)} \\ P3[j]_{(t+1)} &= \mathbf{g}_1 \cdot P0[j]_{(t)} + P1[j]_{(t)} + P2[j]_{(t)} + \mathbf{g}_0 \cdot P3[j]_{(t)} \end{aligned}$$

4. AddRoundKey:

- Key-Scheduler

- Using $Rcon(t) \triangleq \mathbf{g}_0^{t-1}$, the first column in the new key-matrix is computed according to:

$$\begin{bmatrix} K00_{(t+1)} \\ K10_{(t+1)} \\ K20_{(t+1)} \\ K30_{(t+1)} \end{bmatrix} = \begin{bmatrix} \text{SubBytes}(K14_{(t)}) + K00_{(t)} + Rcon(t) \\ \text{SubBytes}(K24_{(t)}) + K10_{(t)} \\ \text{SubBytes}(K34_{(t)}) + K20_{(t)} \\ \text{SubBytes}(K04_{(t)}) + K30_{(t)} \end{bmatrix}.$$

- The other key elements are computed as: $K[i, j]_{(t+1)} = K[i, j-1]_{(t+1)} + K[i, j]_{(t)}$.
- The new key is added by combining each byte of the current plain-text with the corresponding byte of the key, using bitwise exclusive-OR.

E.2 Implementation technique of the Rijndael cipher

Per our complexity measures (cf. Section 2.2), we wish to construct an efficient representation of a hash function by using algebraic constraints. However, the Rijndael cipher is computed over \mathbb{F}_{2^8} with field operations modulo the irreducible polynomial $x^8 + x^4 + x^3 + x + 1$, while the operations in our IOP system are over $\mathbb{F}_{2^{64}}$, defined using a different primitive polynomial. The properties of finite fields entail that for

any field \mathbb{F}_{p^m} and $k|m$, there exists a subfield \mathbb{F}_{p^k} . Therefore, there is an isomorphism between \mathbb{F}_{2^8} and a subfield of $\mathbb{F}_{2^{64}}$.

We obtain such an isomorphism by mapping a primitive element of $\mathbb{F}_{2^{64}}$ to a primitive element of the subfield with $F' \cong \mathbb{F}_{2^8}$, so that the mapping is implied by their powers. This is done by finding an element of $\mathbb{F}_{2^{64}}$ with order $2^8 - 1$.

We then transform all constants needed for Rijndael to their representation in F' , and perform all the field operations in F' . Importantly, this enables an efficient constraint for the SubBytes step of Rijndael, since we can represent the field inverse via a single multiplication in F' . Specifically, by using an auxiliary element $z \in F'$, the constraint $y = x^{-1}$ can be represented via $y \cdot z = x$. By contrast, a naïve implementation of the inverse operation would require auxiliary elements $\{b_i\}_{i=0}^7$, booleanity constraints $\cup\{b_i(b_i + 1)\}_{i=0}^7$, and a polynomial of degree 8 with 256 summands.

The full SubBytes S-box is defined according to $x \rightarrow M \cdot x^{-1} + b$, where $M \in \mathbb{F}_2^{8 \times 8}$ and $b \in \mathbb{F}_2^{8 \times 1}$ are constants. Adding the constant b is a simple field addition in F' , whereas the multiplication by the constant matrix M can be represented using a linear transformation $T : \mathbb{F}_{2^n} \rightarrow \mathbb{F}_{2^n}$. Using algebraic properties, we have that any linear transformation can be represented by a linearized polynomial [78, Chaper 3.4]. We obtain the linearized polynomial $C(x) = \sum_{i=0}^7 c_i x^{2^i}$ by finding coefficients $\{c_i\}_{i=0}^7$ that satisfy $C(a_i) = b_i$, where (a_0, a_1, \dots, a_7) is a basis for the domain of T and (b_0, b_1, \dots, b_7) is a basis for the range of T .

While the degree of $C(x)$ is 128, the degree of our constraint polynomial for the entire Rijndael-160 computation is in fact only 8. At the high-level, the degree reduction is achieved via a decomposition $C(x) = C_1(C_2(C_3(x)))$ with 3 auxiliary field elements, using an AIR such as $\{z' + C_3(x), z'' + C_2(z'), z = C_1(z'')\}$. Per Section 2.3, this AIR is translated into a single constraint $(z' = C_3(z)) \wedge (z'' = C_2(z')) \wedge (z = C_1(z''))$, where the logical-AND is accomplished using the ALI protocol, i.e., random coefficients that are picked by the verifier and sent to the prover in the next round of interaction (this round is used simultaneously for zero-knowledge masking, cf. Appendix D.1). For better efficiency, the exact implementation uses repeated squaring/quadrupling of the coefficients $\{c_i\}_{i=0}^7$, rather than the polynomial composition $C_1(C_2(C_3(x)))$.

The ShiftRows operation is implemented together with SubBytes, by simply placing the results of the SubBytes S-box in the appropriate registers for the next cycle (cf. Appendix E.3). The MixColumns operations is implemented in a single cycle, using the linear combination that we described above to perform field additions and multiplications by the constant g_0 . The AddRoundKey operation is done at the same cycle that we compute MixColumns, using the aforementioned efficient SubBytes S-box implementation.

E.3 State machine of the Rijndael cipher

The algebraic execution trace (cf. Section 2) of the Rijndael-160 hash function is shown in the following table. We enforce boundary constraints on the first and last rows (i.e., cycle 0 and cycle T).

P00 \dots P34 ₍₀₎	K00 \dots K34 ₍₀₎	INV1 \dots INV5 ₍₀₎	W11 \dots W53 ₍₀₎	F1 ₍₀₎	F2 ₍₀₎	RC ₍₀₎	INVRC ₍₀₎	STATE ₍₀₎
⋮							⋮	
P00 \dots P34 _(t)	K00 \dots K34 _(t)	INV1 \dots INV5 _(t)	W11 \dots W53 _(t)	F1 _(t)	F2 _(t)	RC _(t)	INVRC _(t)	STATE _(t)
P00 \dots P34 _(t+1)	K00 \dots K34 _(t+1)	INV1 \dots INV5 _(t+1)	W11 \dots W53 _(t+1)	F1 _(t+1)	F2 _(t+1)	RC _(t+1)	INVRC _(t+1)	STATE _(t+1)
⋮							⋮	
P00 \dots P34 _(T)	K00 \dots K34 _(T)	INV1 \dots INV5 _(T)	W11 \dots W53 _(T)	F1 _(T)	F2 _(T)	RC _(T)	INVRC _(T)	STATE _(T)

$$\begin{aligned}
& (\text{INV1}_{(t)}\text{P00}_{(t)} + 1)(\text{P00}_{(t)} \wedge \text{INV1}_{(t)}) \wedge \\
& (\text{W11}_{(t)} + \text{INV1}_{(t)}^4) \wedge (\text{W12}_{(t)} + \text{W11}_{(t)}^4) \wedge (\text{W13}_{(t)} + \text{W12}_{(t)}^4) \wedge \\
& (\text{P00}_{(t+1)} + c_0 \cdot \text{INV1}_{(t)} + c_1 \cdot \text{INV1}_{(t)}^2 + c_2 \cdot \text{W11}_{(t)} + c_3 \cdot \text{W11}_{(t)}^2 \\
& \quad + c_4 \cdot \text{W12}_{(t)} + c_5 \cdot \text{W12}_{(t)}^2 + c_6 \cdot \text{W13}_{(t)} + c_7 \cdot \text{W13}_{(t)}^2 + b)
\end{aligned}$$

Figure 9: Constraints polynomial for the Rijndael-160 SubBytes S-box.

The 20 registers	$ \begin{bmatrix} \text{P00} & \text{P01} & \text{P02} & \text{P03} & \text{P04} \\ \text{P10} & \text{P11} & \text{P12} & \text{P13} & \text{P14} \\ \text{P20} & \text{P21} & \text{P22} & \text{P23} & \text{P24} \\ \text{P30} & \text{P31} & \text{P32} & \text{P33} & \text{P34} \end{bmatrix} $	contain the 160 bits of the plain-text.
The 20 registers	$ \begin{bmatrix} \text{K00} & \text{K01} & \text{K02} & \text{K03} & \text{K04} \\ \text{K10} & \text{K11} & \text{K12} & \text{K13} & \text{K14} \\ \text{K20} & \text{K21} & \text{K22} & \text{K23} & \text{K24} \\ \text{K30} & \text{K31} & \text{K32} & \text{K33} & \text{K34} \end{bmatrix} $	contain the 160 bits of the key.

Each of these 40 registers is a field element of $\mathbb{F}_{2^{64}}$ that resides in the subfield $F' \cong \mathbb{F}_{2^8}$, and hence contains only 8 bits of information. Per Appendix E.2, this is done to support a native inversion operation for the SubBytes S-box. In each cycle, the registers $\text{INV1}, \text{INV2}, \text{INV3}, \text{INV4}, \text{INV5}$ are used primarily as the auxiliary field elements that compute the inverses for SubBytes.

For $i \in \{1, \dots, 5\}$, the registers $\text{Wi1}, \text{Wi2}, \text{Wi3}$ store the repeated quadrupling that are used to compute the powers of $\text{INV}i$: $\text{Wi1} = \text{INV}i^4, \text{Wi2} = \text{Wi1}^4 = \text{INV}i^{16}, \text{Wi3} = \text{Wi2}^4 = \text{Wi1}^{16} = \text{INV}i^{64}$. Our constraints will then also square these registers, for example $\text{Wi3} \cdot \text{Wi3} = \text{INV}i^{128}$.

The registers $\text{F1}, \text{F2}$ are inner flags that specify the current step in the Rijndael loop. Every round of Rijndael takes 4 steps, and our algebraic constraints use the values of $\text{F1}, \text{F2}$ to enforce the requirements of the current step.

The register RC is used to compute $\text{Rcon}(i)$ in round i of the Rijndael loop. The register INVRC is used for the inverse of RC , in order to tell when to stop the Rijndael iterations. The register STATE is an external flag that specifies whether we compute the Rijndael cipher or some additional logic (i.e., STATE would be unnecessary for a single invocation of Rijndael-160).

We provide an excerpt of the algebraic constraints of a single SubBytes S-box in Figure 9.

Overall, the width of the computation is 65, per the above description. The Rijndael-160 cipher requires 11 rounds where each round consists of SubBytes, ShiftRows, MixColumns, and AddRoundKey (except for the last round that lacks MixColumns). Each round takes 5 cycles in our implementation, hence an entire invocation of Rijndael-160 takes 55 cycles. The prover needs to compute a total of $55 \cdot 65 = 3575$ field elements for a single invocation of Rijndael-160.

E.4 From encryption to hash function: Davies-Meyer

The Rijndael-160 block cipher can be converted into a hash function by using the Davies-Meyer transformation: $\text{hash}(B, K) = E_K(B) \oplus B$, where E is the Rijndael-160 cipher in our case. The resulting $\text{hash}(B, K)$ is collision-resistant under the assumption that for any key K the function $E_K(\cdot)$ is an independent random

permutation, see [73, Theorem 6.5]. Let us also note that constructions with additional overhead can give a CRHF under milder assumptions (e.g., that $E_K(\cdot)$ is a pseudo-random permutation), see for example [32, Table 3].

To implement the Davies-Meyer construction, the execution trace (cf. Section 2.2) requires saving the 160 bits of B while computing the output of the Rijndael-160 cipher. As discussed, our implementation expands B into 20 registers that contain elements of $\mathbb{F}_{2^{64}}$. Since each of these 20 registers holds only 8 bits of information, we can compress and save B in 3 registers that hold 64 bits of information.

The compression method is quite simple: by treating $\mathbb{F}_{2^{64}}$ as an extension field of $F' \cong \mathbb{F}_{2^8}$ of degree 8, we set a basis of the extension field and encode 8 registers of B into an element of $\mathbb{F}_{2^{64}}$ by using them as the coefficients of the basis elements. It can in fact be proved that $(1, g_0, g_0^2, \dots, g_0^7)$ is such a basis. Thus, the encoding is done as $B_0 = \sum_{k=0}^7 p_k \cdot g_0^k$, where the values p_k are taken from the registers P_{ij} . We encode B_2, B_3 in the same manner, except that for B_3 the last 4 coefficients are set to 0.

In order to feed the output digest into the next Rijndael-160 invocation (see Appendix E.5), we require the output to be in the “uncompressed” form that spans 20 registers. Thus, for $i \in \{0, 1, 2\}$, we decompress B_i after the Rijndael-160 cipher computation, and then add the uncompressed values to the cipher’s output. This is done by letting the prover supply 20 field element $\{p_k\}_{k=0}^{19}$ nondeterministically (with $p_k = 0$ for $k \in \{20, 21, 22, 23\}$), and enforcing the algebraic constraint $B_i + \sum_{k=0}^7 p_{k+8i} \cdot g_0^k = 0$ for $i \in \{0, 1, 2\}$. However, this linear combination is unique only if the coefficients p_k reside in F' , and therefore the verifier must also validate $p_k \in F'$ for $k \in \{0, \dots, 19\}$. This is achieved using Fermat’s little theorem, which states that for any finite field F , $\forall x \in F : x^{|F|} = x$. In our case the constraints are $\cup \{p_k^{256} + p_k = 0\}_{k=0}^{19}$, and we again reduce their total degree to 8 by using repeated quadrupling.

The state machine thus requires 3 additional registers for every cycle:

$B_{0(0)}$	$B_{1(0)}$	$B_{2(0)}$
\vdots	\vdots	\vdots
$B_{0(t)}$	$B_{1(t)}$	$B_{2(t)}$
$B_{0(t+1)}$	$B_{1(t+1)}$	$B_{2(t+1)}$
\vdots	\vdots	\vdots
$B_{0(T)}$	$B_{1(T)}$	$B_{2(T)}$

We show the crux of the algebraic constraints for compression and decompression in Figure 10.

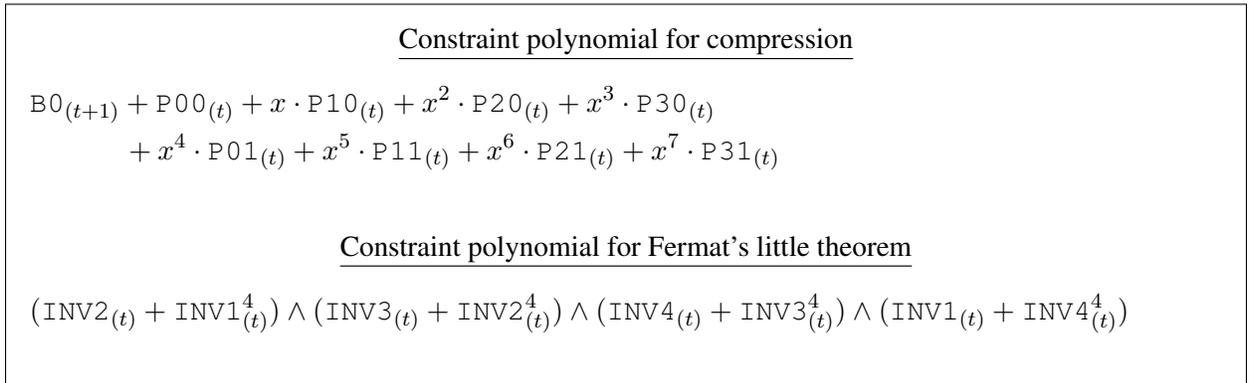


Figure 10: Constraint polynomials for compression and decompression

The compression requires one cycle before each computation of the Rijndael-160 cipher, and the decompression requires two cycles afterwards. Thus, the number of cycles for the Rijndael-160 hash function is 58, the width is 68, and the total number of field elements that the prover needs to compute is $58 \cdot 68 = 3944$.

E.5 DNA profile match (DPM)

The high-level pseudo-code of the DNA profile match is given as Program 1. The database records are assumed to reside in a hashchain of N elements, as illustrated in Figure 11. More precisely, the hashchain is computed using the Merkle-Damgard construction with Rijndael-160 as the compression function. The hash of the last element commits to the entire database, and is verified as a boundary constraint (line XV). Notice that we harness the power of nondeterminism to supply the values of the chain elements during the exhaustive search, which implies that for an arbitrary N the program can operate with only a small constant number of auxiliary variables. Due to the collision resistance of the underlying compression function and the Merkle-Damgard construction [83, 45], this use of nondeterminism is secure.

An element of the chain is a DNA profile according to the Combined DNA Index System (CODIS) format; it is comprised of Short Tandem Repeat (STR) counts for 20 ‘‘core loci’’; we use an 8 bit integer to record a single STR value, and we encode the integer by a single element of F' (cf. Appendix E.2). Since a single DNA profile requires 20 *pairs* of STR values (2 per loci), each record (profile) is stored in two consecutive elements of the hashchain. Thus, a database $D(n)$ of n profiles requires $N = 2n$ chain elements, and Program 1 consists of logic that alternates between odd and even elements.

Program 1 DNA profile match

Explicit inputs: n, cm_t, cm_p

Nondeterministic inputs: $\{\text{VAL}_{i,j}\}_{i \in \{1,2\}, j \in \{1,2,\dots,20\}}, \{W_i\}_{i \in \{0,1,2,\dots,N\}}$

```

I: if  $cm_p \neq \text{hash160}(\{\text{VAL}_{1,j}\}_{j \in \{1,2,\dots,20\}}, \{\text{VAL}_{2,j}\}_{j \in \{1,2,\dots,20\}})$  then return false end if
II:  $k \leftarrow 1, \text{flag} \leftarrow 0, h \leftarrow 0, T \leftarrow 0, N \leftarrow 2n$ 
III: while  $k \neq g^N$  do
IV:   Parse  $(L_1, R_1, L_2, R_2, \dots, L_{10}, R_{10}) = W_j$   $\triangleright j = \log_g k$ 
V:   if  $\text{flag} = 0$  then
VI:      $T_1 \leftarrow \text{CheckPairs}(\text{VAL}_{1,1}, \text{VAL}_{1,2}, L_1, R_1, \text{VAL}_{2,1}, \text{VAL}_{2,2}, L_2, R_2, \dots, \text{VAL}_{10,1}, \text{VAL}_{10,2}, L_{10}, R_{10})$ 
VII:   else
VIII:     $T_2 \leftarrow \text{CheckPairs}(\text{VAL}_{11,1}, \text{VAL}_{11,2}, L_1, R_1, \text{VAL}_{12,1}, \text{VAL}_{12,2}, L_2, R_2, \dots, \text{VAL}_{20,1}, \text{VAL}_{20,2}, L_{10}, R_{10})$ 
IX:     $T \leftarrow \text{MatchingResult}(T_1, T_2, T)$ 
X:   end if
XI:    $h \leftarrow \text{hash160}(h, W_j)$ 
XII:   $k \leftarrow g \cdot k$ 
XIII:  $\text{flag} \leftarrow 1 - \text{flag}$ 
XIV: end while
XV: if  $cm_t \neq h$  then return false else return T end if

```

To validate that the prover does not skip over some prefix of the chain in the exhaustive search, the total number of hash invocations N is also checked by the verifier as a boundary constraint. We also note that Program 1 increments its counter via a field multiplication with the generator g of \mathbb{F}_{264}^* , thereby avoiding integer arithmetics.

The register T stores the output, and is verified in the last cycle via a boundary constraint. The output

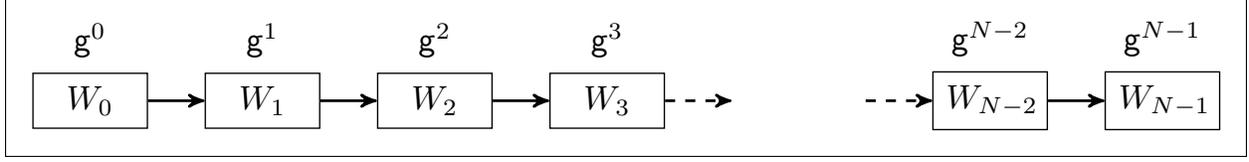


Figure 11: Illustration of the data structure used in the DNA profile match program.

can be either 1) “perfect match”, meaning that an exact match between the input (i.e., the commitment cm_p that the prover decommits in ZK into 20 STR pairs) and a profile in the hashchain was found, or 2) “partial match”, meaning that the exhaustive search found a profile in the hashchain such that $\forall j \in \{1, \dots, 20\}$ at least one STR value of its j th pair matches a value of the j th STR pair of the input, or 3) “no match”.

We provide high-level pseudocode of the perfect/partial match logic in Figures 12 and 13, and the corresponding algebraic constraints in Figures 14 and 15. The $\text{InnerMatch}_{(t)}$ and $\text{Match}_{(t)}$ registers of Figure 15 correspond to T_1 and T_2 in Figure 13, and $\text{Match}_{(t+1)}$ corresponds to T . In Figure 14, the registers $\text{PAIR1}_{(t)}$ and $\text{PAIR2}_{(t)}$ are compared to the i th loci pair of the input. Each such comparison relies on 5 auxiliary registers: $\text{INVTMP1}_{(t)}, \dots, \text{INVTMP4}_{(t)}$ to compute field inverses, and $\text{NewFlag}_{(t)}$ that represents the current inner matching (the register $\text{LastFlag}_{(t)}$ is derived from the previous iteration).

Since the prover (but not the verifier) knows the decommitment of cm_p , the first invocation of Rijndael-160 in Program 1 is executed with completely nondeterministic inputs (that occupy 40 registers), and the output is constrained to be equal to cm_p . Due to the ZK guarantee of our proof system, this implies semantic security (by contrast, comparing the database records to 40 explicit constants that commit to STR strings is not semantically secure because the same value may appear more than once, especially if Program 1 is executed multiple times with the same committed database cm_t). In each of the next invocations, those nondeterministic values should be compared against the current database record (i.e., 20 registers that are supplied nondeterministically by the prover). However, keeping the initial 40 values throughout the entire execution of Program 1 is inefficient. Since each of those 40 registers contains only 8 bits of information, we use the same technique as in Appendix E.4 to compress this data into 5 auxiliary registers.

$L0_{(0)}$	$L1_{(0)}$	$L2_{(0)}$	$L3_{(0)}$	$L4_{(0)}$
\vdots	\vdots	\vdots	\vdots	\vdots
$L0_{(t)}$	$L1_{(t)}$	$L2_{(t)}$	$L3_{(t)}$	$L4_{(t)}$
$L0_{(t+1)}$	$L1_{(t+1)}$	$L2_{(t+1)}$	$L3_{(t+1)}$	$L4_{(t+1)}$
\vdots	\vdots	\vdots	\vdots	\vdots
$L0_{(T)}$	$L1_{(T)}$	$L2_{(T)}$	$L3_{(T)}$	$L4_{(T)}$

After each invocation of Rijndael-160, we decompress these 5 registers in order to execute the comparison logic (cf. Figures 12 to 15) for STR pairs with the newly supplied database record. Here too we must verify the decompressed registers by using Fermat’s little theorem, and this requires two cycles as there are not enough temporary registers to compute the repeated quadrupling in a single cycle.

Program 1 also requires an extra register for the counter $k = g^j$, two registers for handling the perfect/partial match constraints, three more registers that handle the possible states, and two additional auxiliary registers. Overall, the width of the witness is 81, the number of cycles is 62, and the total number of field elements that the prover computes is $81 \cdot 62 \cdot N = 5022 \cdot N$.

```

CheckPairs( $VAL_{1,1}, VAL_{1,2}, L_1, R_1, VAL_{2,1}, VAL_{2,2}, L_2, R_2, \dots, VAL_{10,1}, VAL_{10,2}, L_{10}, R_{10}$ )

I:  $t \leftarrow 2$ 
II: for  $j \in \{1, \dots, 10\}$  do
III:   if ( $VAL_{j,1} = L_j$ ) AND ( $VAL_{j,2} = R_j$ ) then Continue end if
IV:   if ( $VAL_{j,1} = R_j$ ) AND ( $VAL_{j,2} = L_j$ ) then Continue end if
V:   if ( $VAL_{j,1} = L_j$ ) OR ( $VAL_{j,1} = R_j$ ) OR ( $VAL_{j,2} = L_j$ ) OR ( $VAL_{j,2} = R_j$ ) then
VI:      $t \leftarrow 1$ 
VII:   else
VIII:     return 0
IX:   end if
X: end for
XI: return  $t$ 

```

Figure 12: CheckPairs subroutine of Program 1.

```

MatchingResult( $T_1, T_2, T$ )

I: if  $T_1 = T_2 = 2$  then return 2 end if
II: if ( $T_1 = 0$ ) OR ( $T_2 = 0$ ) then return  $T$  end if
III: if  $T = 0$  then return 1 end if

```

Figure 13: MatchingResult subroutine of Program 1.

$$\begin{aligned}
& (\text{PAIR1}_{(t)} + \text{values}_{i,0}) \cdot (\text{PAIR1}_{(t)} + \text{values}_{i,1}) \cdot [\\
& \quad (\text{PAIR2}_{(t)} + \text{values}_{i,0}) \cdot (\text{PAIR2}_{(t)} + \text{values}_{i,1}) \cdot \text{NewFlag}_{(t)} \wedge \\
& \quad \text{LastFlag}_{(t)} \cdot (\text{NewFlag}_{(t)} + 1) \cdot [\\
& \quad \quad ((\text{PAIR2}_{(t)} + \text{values}_{i,0}) \cdot \text{INVTMP1}_{(t)} + 1) \wedge \\
& \quad \quad ((\text{PAIR2}_{(t)} + \text{values}_{i,1}) \cdot \text{INVTMP2}_{(t)} + 1)] \wedge \\
& (\text{PAIR2}_{(t)} + \text{values}_{i,0}) \cdot (\text{PAIR2}_{(t)} + \text{values}_{i,1}) \cdot [\\
& \quad \text{LastFlag}_{(t)} \cdot (\text{NewFlag}_{(t)} + 1) \cdot [\\
& \quad \quad ((\text{PAIR1}_{(t)} + \text{values}_{i,0}) \cdot \text{INVTMP3}_{(t)} + 1) \wedge \\
& \quad \quad ((\text{PAIR1}_{(t)} + \text{values}_{i,1}) \cdot \text{INVTMP4}_{(t)} + 1)] \wedge \\
& (\text{NewFlag}_{(t)} + \text{LastFlag}_{(t)}) \cdot [\\
& \quad ((\text{PAIR1}_{(t)} + \text{values}_{i,1}) \cdot \text{INVTMP4}_{(t)} + 1) \cdot ((\text{PAIR2}_{(t)} + \text{values}_{i,0}) \cdot \text{INVTMP1}_{(t)} + 1) \wedge \\
& \quad ((\text{PAIR1}_{(t)} + \text{values}_{i,0}) \cdot \text{INVTMP3}_{(t)} + 1) \cdot ((\text{PAIR2}_{(t)} + \text{values}_{i,1}) \cdot \text{INVTMP2}_{(t)} + 1)] \wedge \\
& (\text{LastFlag}_{(t)} + 1) \cdot (\text{LastFlag}_{(t)} + X) \cdot \text{NewFlag}_{(t)}
\end{aligned}$$

Figure 14: Algebraic constraints for one pair in the CheckPairs subroutine.

$$\begin{aligned}
& \text{InnerMatch}_{(t)} \cdot (\text{InnerMatch}_{(t)} + 1) \cdot (\text{Match}_{(t+1)} + X) \wedge \\
& \text{Match}_{(t)} \cdot (\text{Match}_{(t)} + 1) \cdot (\text{Match}_{(t+1)} + X) \wedge \\
& \text{InnerMatch}_{(t)} \cdot (\text{InnerMatch}_{(t)} + 1) \cdot (\text{Match}_{(t)} + X) \cdot (\text{Match}_{(t+1)} + 1) \wedge \\
& (\text{InnerMatch}_{(t)} + 1) \cdot (\text{InnerMatch}_{(t)} + X) \cdot (\text{Match}_{(t)} + 1) \cdot (\text{Match}_{(t)} + X) \cdot \text{Match}_{(t+1)}
\end{aligned}$$

Figure 15: Algebraic constraints for the MatchingResult subroutine.

F The AIR of the SHA2 hash function

The popular Secure Hash Algorithm 2 (SHA2) family [92] requires modular addition and cyclic shifts which are not particularly “binary field friendly”. Nevertheless, we constructed a rather efficient AIR for it (first row of Table 4), using field-specific constraints. A notable example is a constraint system that “extracts” the i th bit from $\alpha \in \mathbb{F}_{2^t}$ for any $i < t$; this system uses only a pair of constraints of degree 2 (notice the number of constraints and their degree is independent of t); we believe this bit-extraction constraint set will be useful for other computations.

We provide here notation and several basic lemmas, which facilitate the bit extraction technique that our efficient SHA-256 hash function implementation is based on.

The trace of an element $y \in \mathbb{F}_{2^m}$ is defined as $\text{Tr}_{m|2}(y) \triangleq \sum_{i=0}^{m-1} y^{2^i}$.

Proposition 1. For any \mathbb{F}_2 -linear function $f : \mathbb{F}_{2^m} \rightarrow \mathbb{F}_2$, there exists a field element $\alpha_f \in \mathbb{F}_{2^m}$ such that $\forall y \in \mathbb{F}_{2^m} : f(y) = \text{Tr}_{m|2}(\alpha_f \cdot y)$.

Proposition 2. For every $c \in \mathbb{F}_{2^m}$, the equation $y^2 + y = c$ has solutions in \mathbb{F}_{2^m} if and only if $\text{Tr}_{m|2}(c) = 0$. Notice that if y_0 is a solution of $y(y + 1) = c$ then $y_0 + 1$ is the other solution, since the field characteristic is 2.

Definition 1. Let $\text{isZero}(\alpha, w, v)$ be the polynomial $\text{isZero}(\alpha, w, v) \triangleq w^2 + w + \alpha \cdot v$.

Lemma 1. The i th coefficient in the standard basis representation of $y \in \mathbb{F}_{2^m}$ is 0 if and only if there exists $w \in \mathbb{F}_{2^m}$ such that $0 = \text{isZero}(\alpha_i, w, y)$, where $\alpha_i \in \mathbb{F}_{2^m}$ is some field element that depends only on i .

Proof. The function $f_i : \mathbb{F}_{2^m} \rightarrow \mathbb{F}_2$, $f_i(y) = \begin{cases} 0, & i\text{th coefficient of } y \text{ is } 0 \\ 1, & i\text{th coefficient of } y \text{ is } 1 \end{cases}$ is \mathbb{F}_2 -linear. It follows from Proposition 1 and Proposition 2 that the required α_i exists. \square

Furthermore, it is straightforward to pre-compute the constant α_i by solving linear equations for the trace of basis elements.

G Affine rearrangeable networks

G.1 Combinatorial representation of back-to-back De Bruijn routing

The following classical results regarding De Bruijn networks [47] can be found, e.g., in [77]. We use Theorem G.3 below in the statement and proof of Theorem B.13. Therefore, we provide a self-contained account of the needed results for the sake of completeness.

Definition G.1 (De Bruijn butterfly network). *The De Bruijn butterfly network of dimension n is a directed graph $G = (V, E)$ over the vertices $V \triangleq \{0, \dots, n\} \times \{0, 1\}^n$ and edges*

$$E \triangleq \{[(i, w), (i + 1, \text{csr}(w))], [(i, w), (i + 1, \text{csr}(w) \oplus 1)] \mid 0 \leq i < n, w \in \{0, 1\}^n\}$$

where $\text{csr} : \{0, 1\}^n \rightarrow \{0, 1\}^n$ is the cyclic shift-right operation, and $[x \oplus 1] : \{0, 1\}^n \rightarrow \{0, 1\}^n$ flips the leftmost bit in x , and does not change any other bit. Let \mathcal{C} be a set of colors. We say $\chi : V \rightarrow \mathcal{C}$ is a

coloring of G if for every vertex $v = (i, w)$ with $i < n$, there is a successor u of v (i.e. $vu \in E$) such that $\chi(v) = \chi(u)$.

Definition G.2 (back-to-back De Bruijn routing). A back-to-back De Bruijn routing of dimension n is a pair of coloring functions $\bar{\chi} = (\chi_0, \chi_1)$ of the De Bruijn butterfly network of dimension n , such that for all $w \in \{0, 1\}^n$ (i) $\chi_0(n, w) = \chi_1(n, w)$, and (ii) for any $w' \neq w$, $\chi_0(0, w) \neq \chi_0(0, w')$. In case $|\mathcal{C}| = 2^n$ we define the function induced by $\bar{\chi}$, $\pi_{\bar{\chi}}: \mathcal{C} \rightarrow \mathcal{C}$ to be $\pi_{\bar{\chi}}(\chi_0(0, w)) = \chi_1(0, w)$.

Theorem G.3 (back-to-back De Bruijn is rearrangeable). For every set of colors \mathcal{C} of size 2^n , the set of functions induced by back-to-back De Bruijn routings of dimension n and colors from \mathcal{C} is exactly the set of permutations over \mathcal{C} . Moreover, given a permutation $\pi: \mathcal{C} \rightarrow \mathcal{C}$, there is an algorithm constructing a back-to-back De Bruijn routing $\bar{\chi}$ inducing π with time complexity $O(n \cdot 2^n)$.

Proof. By Definition G.2, the function induced by a back-to-back routing must be a permutation, thus it is sufficient to show an algorithm producing a back-to-back De Bruijn routing given a permutation π . The algorithm is recursive and based on the Beneš network routing algorithm [28], and provided in Algorithm 3. The algorithm as provided in Algorithm 3 does not reach the stated complexity, because it contains many move operations, which can be easily eliminated using simple index translation. \square

Procedure 2 Auxiliary method for routing

Inputs: two vectors u, v of length 2^n indexed by $\{0, 1\}^n$ with distinct entries, where u is a permutation of v

Output: vectors v^0, v^1, u^0, u^1 each of length 2^{n-1} indexed by $\{0, 1\}^{n-1}$ with (i) each element of v is in v^0 or v^1 , (ii) for every $w \in \{0, 1\}^{n-1}$ and $i, b \in \{0, 1\}$, $v_{bw} = v_x^i$ or $u_{bw} = u_x^i$ only if $x = w$, and (iii) u^i is a permutation of v^i .

```

I:  function SPLIT( $n, u, v$ )
II:  initialize  $v^0, v^1, u^0, u^1$  as partial mappings
III:  while  $v^0, v^1, u^0, u^1$  not yet well defined do
IV:    if  $\exists w \in \{0, 1\}^{n-1}, b, i \in \{0, 1\}$  such that  $v_{b \cdot w} \in u^i$  but  $v_{b \cdot w} \notin v^i$  then
V:       $v_w^i \leftarrow v_{b \cdot w}$ 
VI:      $v_w^i \leftarrow v_{\bar{b} \cdot w}$ 
VII:    continue
VIII:  else if  $\exists w \in \{0, 1\}^{n-1}, b, i \in \{0, 1\}$  such that  $u_{b \cdot w} \in v^i$  but  $u_{b \cdot w} \notin u^i$  then
IX:     $u_w^i \leftarrow u_{b \cdot w}$ 
X:      $u_w^i \leftarrow u_{\bar{b} \cdot w}$ 
XI:    else
XII:     pick arbitrary  $w \in \{0, 1\}^{n-1}$  such that  $v_{0 \cdot w} \notin \{v_w^0, v_w^1\}$ 
XIII:     $v_w^i \leftarrow v_{i \cdot w}$  for  $i \in \{0, 1\}$ 
XIV:   end if
XV:   end while
XVI:  return  $v^0, v^1, u^0, u^1$ 
XVII: end function

```

G.2 Affine embedding of back-to-back De Bruijn routing

Definition G.4 (Algebraic De Bruijn). Let \mathbb{F} be a field of characteristic 2 with primitive element \mathbf{g} , and fix some integer t . Let $[\log(t+1)]$ be a primitive binary polynomial of degree $[\log(t+1)]$, Define $H \triangleq$

Program 3 Routing permutation over back-to-back De Bruijn

Inputs: two vectors u, v of length 2^n indexed by $\{0, 1\}^n$ and distinct entries, where u is a permutation of v

Output: back-to-back De Bruijn routing $\bar{\chi}$, represented as a pair of $2^n \times (n + 1)$ tables, with induced function satisfying $\pi_{\bar{\chi}}(v_w) = u_w$ for all w

```
I: function ROUTE_DE_BRUIJN( $n, u, v$ )
II:                                     ▷ Recursion stopping condition
III:   if  $n = 1$  then
IV:     if  $v_0 = u_0$  then
V:       return  $\chi_0 = \chi_1 = \begin{bmatrix} v_0 & v_0 \\ v_1 & v_1 \end{bmatrix}$ 
VI:     else
VII:       return  $\chi_0 = \begin{bmatrix} v_0 & v_0 \\ v_1 & v_1 \end{bmatrix}, \chi_1 = \begin{bmatrix} v_1 & v_0 \\ v_0 & v_1 \end{bmatrix}$ 
VIII:    end if
IX:   end if
X:                                     ▷ General case
XI:    $(v^0, v^1, u^0, u^1) \leftarrow \text{SPLIT}(n, u, v)$ 
XII:   $\bar{\chi}^0 \leftarrow \text{ROUTE\_DE\_BRUIJN}(n - 1, u^0, v^0)$ 
XIII:  $\bar{\chi}^1 \leftarrow \text{ROUTE\_DE\_BRUIJN}(n - 1, u^1, v^1)$ 
XIV:  for all  $w \in \{0, 1\}^n$  do
XV:     $\chi_0(0, w) \leftarrow v_w$ 
XVI:     $\chi_1(0, w) \leftarrow u_w$ 
XVII: end for
XVIII: for all  $(i, w) \in [n] \times \{0, 1\}^{n-1}$  do
XIX:   represent  $w$  by  $w = x \cdot y$  where  $|y| = i - 1$ 
XX:     $\chi_0(i, x \cdot 0 \cdot y) \leftarrow \chi_0^0(i - 1, w)$ 
XXI:    $\chi_0(i, x \cdot 1 \cdot y) \leftarrow \chi_0^1(i - 1, w)$ 
XXII:   $\chi_1(i, x \cdot 0 \cdot y) \leftarrow \chi_1^0(i - 1, w)$ 
XXIII:  $\chi_1(i, x \cdot 1 \cdot y) \leftarrow \chi_1^1(i - 1, w)$ 
XXIV: end for
XXV: end function
```

$\text{span}\{\mathbf{g}^i \mid 0 \leq i < t\}$, and $W \triangleq \text{span}\{\mathbf{g}^{t+i} \mid 0 \leq i < \lceil \log(t+1) \rceil\}$. A pair of mappings $\chi_0, \chi_1: H \oplus W \rightarrow \mathbb{F}$ is an affine embedding of back-to-back De Bruijn of degree t if:

- for every $x \neq x' \in H$, $\chi_0(x + \mathbf{g}^{t+1}) \neq \chi_0(x' + \mathbf{g}^{t+1})$
- for every $x \in H$, $\chi_0(x + \mathbf{g}^t) = \chi_1(x + \mathbf{g}^t)$
- for every $b \in \{0, 1\}$, $x = \sum x_i \mathbf{g}^i \in \text{span}\{\mathbf{g}^i \mid 0 \leq i < t\}$ and every non zero $y = \sum y_i \mathbf{g}^i \in \text{span}\{\mathbf{g}^i \mid 0 \leq i < \lceil \log(t+1) \rceil\}$, denoting by $z = x + \mathbf{g}^t y$, if $(x_{t-1}, y_{\lceil \log(t+1) \rceil - 1}) = (r, c)$ for $r, c \in \{0, 1\}$ then $\chi_b(z)$ equals to one of:
 - $\chi_b(\mathbf{g} \cdot z + r \cdot (\mathbf{g}^t + 1) + c \cdot (\mathbf{g}^t \cdot \xi))$
 - $\chi_b(\mathbf{g} \cdot z + r \cdot (\mathbf{g}^t + 1) + c \cdot (\mathbf{g}^t \cdot \xi) + 1)$

Let \mathcal{C} be the image of χ over H . The function induced by χ_0, χ_1 is a mapping $\pi: \mathcal{C} \rightarrow \mathbb{F}$ mapping defined by $\chi_0(x + \mathbf{g}^{t+1}) \mapsto \chi_1(x + \mathbf{g}^{t+1})$ for every $x \in H$.

Theorem G.5 (Affine permutation). *The set of all mappings induced by affine embeddings of back-to-back De Bruijn of degree t is exactly the set of all permutations over 2^t elements of \mathbb{F} .*