Extending FRIDA Beyond Unique Decoding for Free

Nicolas Mohnblatt* Benedikt Wagner†

May 26, 2026

Abstract

Hall-Andersen, Simkin and Wagner (CiC 1:4) show how to construct data availability sampling schemes from code commitments. Later, in FRIDA (CRYPTO'24), the same authors describe a compiler that takes an interactive oracle proof of proximity (IOPP) for a code and produces a secure code commitment. Chaining both results allows to construct efficient data availability sampling schemes from IOPPs.

In this short note, we give a novel security analysis that extends the results of FRIDA beyond the unique decoding radius of the code being used. This strict improvement leads to data availability sampling schemes with smaller commitments.

Towards our novel analysis, we define a variant of the opening-consistency property introduced in FRIDA, which we name *opening-consistency with assign*. Crucially, our new property does not depend on the unique decoding radius of the code. We then show that the FRIDA compiler can be applied to IOPPs that have opening-consistency with assign to produce secure code commitments. Finally, we show that under mutual correlated agreement, the batched FRI protocol (FOCS'20) satisfies opening-consistency with assign. This latter result is enabled by a recent analysis of FRI by Garreta, Mohnblatt and Wagner (ePrint 2025/1993).

Contents

1 Introduction **2**

2 Preliminaries **5**

3 From IOPP to code commitment, revisited **6**
3.1 Opening-consistency with assign 6
3.2 New security analysis 7

4 Instantiation from batched FRI **8**
4.1 Non-adaptive 8
4.2 Query-selection 8
4.3 Opening-consistency with assign 8

5 Evaluation **11**

A Additional comments on DEEP-FRIDA **16**

*ZKSecurity, nico@zksecurity.xyz
†Ethereum Foundation, benedikt.wagner@ethereum.org

Disclaimer: this document functions as an extension to FRIDA [HASW24b] incorporating the findings of [GMW25]. It is not intended to be self-contained or be read as a standalone document.

1 Introduction

Data availability sampling (DAS) [ASBK21, HASW24a] has recently attracted significant attention [HASW24a, HASW24b, WZ24, EMA25, EA25, WZ25, BNNP25, CSK25, YZD26]. Intuitively, it enables parties to verify that a large dataset associated with a succinct commitment is *available* in a peer-to-peer network, without requiring them to download the entire dataset.

In FRIDA [HASW24b], the authors showed that a certain class of interactive oracle proofs of proximity (IOPPs) can be transformed into data availability sampling schemes. More precisely, if an IOPP satisfies the notions of *non-adaptivity*, *query-selection* and *opening-consistency*, then it can be converted into a *code commitment scheme (CC)*, which can in turn be compiled into a DAS scheme using the compiler from [HASW24a]. The authors further demonstrated that the celebrated FRI protocol [BBHR18] and its batched variant [BCI+20] satisfy non-adaptivity, query-selection and opening-consistency. However, the definition and analysis of opening-consistency in their work required operating in the *unique decoding regime*, which leads to inefficient concrete parameters, in particular when it comes to the commitment size.

In this short note, we extend the analysis beyond the unique decoding regime, resulting in concretely more efficient DAS constructions. Our new analysis is enabled by a recent security proof for FRI [GMW25] that establishes additional, previously unproven structural properties of the protocol. In particular, we use the fact that for every proximity parameter for which mutual correlated agreement holds, the commit phase of the FRI protocol acts as a commitment scheme binding the prover to a unique (and extractable) codeword. Notably, our analysis does not require *out-of-domain sampling*, a technique often used in the IOPP literature to bind the prover to a unique codeword, even in the list decoding regime [BGKS20, ACFY24, ACFY25].

Roadmap. Our results follow this roadmap:

1. We remove direct references to unique decoding from the definition of opening-consistency. Our modified definition, named *opening-consistency with assign*, is presented in Definition 3.1.

2. We show in Theorem 3.4 that the IOPP-to-CC compiler of FRIDA [HASW24c] yields a secure code commitment when replacing the opening-consistency property with our opening-consistency with assign.

3. Finally, we show that the IOPP-to-CC compiler can be applied to batched FRI [BCI+20]. To do so, we show that batched FRI is non-adaptive, has query-selection and has opening-consistency with assign, within *and beyond* the unique decoding regime. In particular, we show in Theorem 4.9 that the latter holds for all distance parameters for which Reed-Solomon codes have satisfactory mutual correlated agreement error.

The above definitions and theorems, and their relations to previous work are summarized in Figure 1. Our results provide an alternative IOPP-to-CC path that yields a strict improvement over the path given in FRIDA.

Evaluation. Although we do not implement the resulting DAS scheme, we give estimates for the scheme's performance using the same evaluation script used in FRIDA (with appropriate

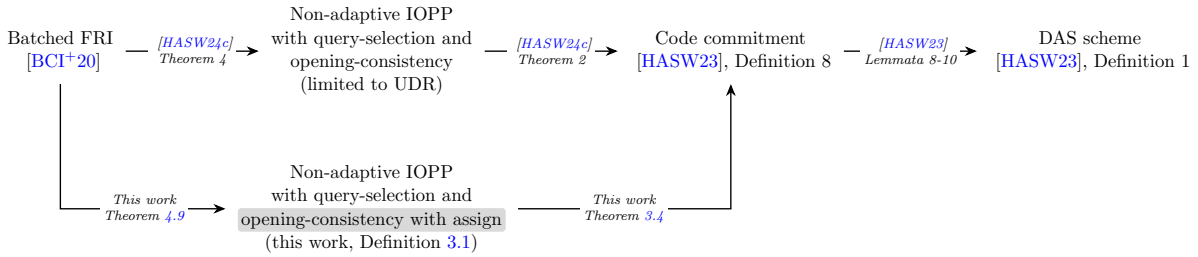

Figure 1: Overview of our modifications of the IOPP-to-DAS pipeline from *Foundations of Data Availability Sampling* [HASW24a] and *FRIDA* [HASW24b]. Arrows are labelled with the theorems and lemmas that prove these relationships (in the respective ePrint versions [HASW23, HASW24c]), omitting proofs of non-adaptivity and query-selection.

modifications). As expected, our calculations show a strict improvement on commitment size, reducing it by a modest 8.8% for rate 1/2 and by more than 40% at rates 1/8 and below. We discuss our evaluation in Section 5.

Concurrent work. In concurrent work DEEP-FRIDA, Yang, Zhang and Deng [YZD26] also propose to extend FRIDA beyond unique decoding. Similar to us, they extend the opening-consistency definition to work beyond the unique decoding regime. They then show that the DEEP-FRI protocol [BGKS20] satisfies their notion of opening-consistency. There lies the essential difference with our work: DEEP-FRI is a variant of FRI that relies on out-of-domain samples, specifically using the *quotienting* technique. As we show in this work, out-of-domain samples are not necessary to extend FRIDA beyond unique decoding. By avoiding out-of-domain samples, we can use IOPPs that run in fewer interaction rounds, and require less prover and verifier work. This improvement comes at no expense since our results hold in the same parameter range as DEEP-FRIDA and with smaller soundness errors. We additionally comment on a minor technical issue with DEEP-FRIDA in Appendix A.

DAS from other IOPPs. A natural question is to ask what other IOPPs can be used to construct DAS schemes. In *Foundations of Data Availability Sampling* [HASW24a], Hall-Andersen, Simkin and Wagner explore a Ligero [AHIV17]/Brakedown [GLS+23]-like approach. FRIDA [HASW24b] then gives a generic compiler and applies it to FRI [BBHR18].

Since then, many new IOPPs have been proposed, improving on the above either in proof size, verifier efficiency or prover efficiency [ABN23, ZCF24, ACFY24, ACFY25, NA25, BMMS25a, BCF+25, BFRW25, BMMS25b, GNR26, BFMM26]. With only a few exceptions, these new protocols rely on *code switching* [RR24], either implicitly (*e.g.*, STIR [ACFY24], WHIR [ACFY25], Ligerito [NA25]) or explicitly (*e.g.*, Blaze [BCF+25], FICS [BMMS25a]). While we believe it is possible to prove opening-consistency for many of these protocols, we do not know of any protocol that uses code switching and satisfies query-selection. Indeed, a code switching IOPP will typically sample queries to the initial codeword at an early round in the interactive proof rather than at the final verifier message, as required by query-selection. As a consequence, we cannot apply FRIDA's IOPP-to-CC compiler to these code switching IOPPs[1].

Open problems. We conclude by stating some open problems related to the above discussion:

[1] A recent work, ISTIRDA [WSZ+26], proves opening-consistency for STIR but fails to discuss query-selection. As a result, the work incorrectly claims to construct a DAS scheme from STIR. We contacted the authors expressing our concerns but were left with no response at the time of publication.

1. Can we design or modify an IOPP to benefit from code switching while still retaining query-selection?

2. Can we construct an alternative IOPP-to-CC compiler that does not require query-selection?

3. Are there benefits to constructing DAS schemes from IOPPs that do not use code switching, such as those that generalize FRI to any foldable code [ABN23] and allow to additionally enforce a sumcheck constraint on the codeword [ZCF24]?

2 Preliminaries

We use the same notation and preliminary definitions as in FRIDA [HASW24b], specifically the extended version available on ePrint [HASW24c]. In addition, we use definitions given in the body of the FRIDA paper such as

- *non-adaptive IOPP* [HASW24c, Definition 7],
- *query-selection of IOPPs* [HASW24c, Definition 8],
- *suitable transcripts* [HASW24c, Definition 9] and
- *opening-consistency of IOPPs* [HASW24c, Definition 10].

As this note intends to update FRIDA with recent results on the round-by-round soundness of FRI [GMW25], we introduce the necessary concepts and notation below.

Mutual correlated agreement. *Mutual correlated agreement* formalizes the intuition that if a random linear combination of functions agrees with a codeword on a "large enough" set, then each of the functions also agrees with a codeword on this same set, except with negligible probability. We refer readers to Definition 2.2 of [GMW25] for a full definition. The definition assumes that the random batching factors are computed by feeding a random seed to a *proximity generator* Gen. The error with which the property might not hold for the generator Gen is a function of the code $\mathcal{C}$, the number of functions being batched $m \in \mathbb{N}$ and $\gamma \in [0, 1]$, which parametrizes the notion of "large enough". It is denoted $\varepsilon^{\mathsf{mca}}_{\mathsf{Gen}}(\mathcal{C}, m, \gamma)$.

As in [GMW25, Definition 2.3], we define $\mathsf{Powers}(m; r)$ to be the generator that returns the first m univariate powers of the seed r:

$$\mathsf{Powers}(m; r) := (1, r, \ldots, r^{m-1}) \in \mathbb{F}^m.$$

Round-by-round soundness. *Round-by-round (RBR) soundness* is a property of interactive oracle proofs (IOPs) that strengthens the regular notion of soundness. This strengthening is used to prove that non-interactive proofs compiled from IOPs [BCS16] are also sound.

Intuitively, RBR soundness states that it is hard to break any *single round* of the IOP. This intuition is captured by defining a state function State that takes as input an instance x and interaction transcript tr and outputs a bit, 0 or 1. The former indicates that the verifier will reject, while the latter indicates that the verifier will accept. To prove RBR soundness, we consider each round in isolation and bound the probability that the state function changes from 0 to 1 when a transcript is extended with this round's verifier message.

Notation. Formally, we use the definitions of state function and RBR soundness respectively given in Definition 3.4 and Definition 3.5 of [GMW25]. Note, however, that the syntax for an IOP transcript is different in FRIDA [HASW24b] and [GMW25]. Where we write (x, tr) in [GMW25] to denote an instance and interaction transcript, FRIDA [HASW24b] writes T to denote the whole instance-transcript tuple. Since these notations are equivalent and interchangeable, we keep the notation from FRIDA [HASW24b] in order to make our new definition of opening-consistency easily comparable to the original paper.

3 From IOPP to code commitment, revisited

We first briefly revisit the IOPP to code commitment compiler from the FRIDA paper [HASW24b], by generalizing the opening-consistency notion and showing how to adapt the proof for this more general notion.

3.1 Opening-consistency with assign

The following definition generalizes the opening-consistency definition from [HASW24b]. In particular, the original opening-consistency definition stated that for a suitable transcript, the alleged codeword $\boldsymbol{c}$ is within the unique decoding radius of the code. Here, we only require that we can uniquely assign some codeword to it.

Definition 3.1. Let $\mathcal{C} \subseteq \Sigma^n$ be a code, let (P, V) be an IOPP for $\mathcal{C}$ with r rounds, and let $\mathcal{T}$ denote the set of final round verifier-turn transcripts. Then, (P, V) is said to have *opening-consistency with assign* with errors $\mathsf{err}_1, \mathsf{err}_2$ if there are sets Bad and Lucky, and a function $\mathsf{Assign} : \mathcal{T} \to \mathcal{C} \cup \{\perp\}$ such that the following properties hold:

- **No Luck.** For every $i \in [r]$, and any partial verifier-turn transcript $T := (\boldsymbol{c}, \rho_1, \pi_1, \ldots, \rho_{i-1}, \pi_{i-1})$ that contains $2i - 1$ elements, we have
$$\Pr_{\rho_i}[T \circ \rho_i \in \mathsf{Lucky}] \leq \mathsf{err}_1.$$
- **Bad is Rejected.** Let $T := (\boldsymbol{c}, \rho_1, \pi_1, \ldots, \rho_r, \pi_r)$ be a transcript with $T \in \mathsf{Bad}$. Then, we have
$$\Pr_{\rho_{r+1}}\left[\ \mathsf{V}^{\boldsymbol{c}, \pi_1, \ldots, \pi_r}(\rho_1, \ldots, \rho_r, \rho_{r+1}) = 1\ \right] \leq \mathsf{err}_2.$$
- **Suitable Assignment.** Let T be a suitable transcript with respect to Bad and Lucky. Then $\mathsf{Assign}(T) \neq \perp$.
- **Inconsistent is Rejected.** Let T be a suitable transcript with respect to Bad and Lucky. Let $\boldsymbol{c}^* := \mathsf{Assign}(T)$ be the assignment of T[2]. Let ρ_{r+1} such that there is a position[3] $j \in \mathcal{Q}_0(T \circ \rho_{r+1})$ with $\boldsymbol{c}_j \neq \boldsymbol{c}^*_j$. For any such transcript, it holds that:
$$\mathsf{V}^{\boldsymbol{c}, \pi_1, \ldots, \pi_r}(\rho_1, \ldots, \rho_r, \rho_{r+1}) = 0.$$

Remark 3.2. Let δ^* denote the unique decoding radius of $\mathcal{C}$. We can recover the original definition of opening-consistency [HASW24b] from the above by setting Assign to be the function that returns the unique closest codeword if $\Delta(\boldsymbol{c}, \mathcal{C}) \leq \delta^*$ and returns $\perp$ otherwise. Given this assignment rule, the *suitable assignment* property is equivalent to *suitable is close*, and *inconsistent is rejected* is defined with respect to the unique closest codeword.

Remark 3.3 (Round-by-round soundness)**.** As explained in FRIDA, opening-consistency is a stronger requirement than round-by-round soundness. We draw a closer link between the two notions in our analysis of batched FRI in Section 4. As we will see, the Bad set and Lucky set are defined with respect to the IOPP's state function. The *no luck* and *bad is rejected* properties follow directly from the definition of round-by-round soundness, regardless of the specific state function. On the other hand, the *suitable assignment* and *inconsistent is rejected* properties require the IOPP to have special properties beyond round-by-round soundness.

[2] This assignment exists and is not $\perp$ because of the *suitable assignment* property.

[3] Recall that $\mathcal{Q}_0$ is shorthand notation for the verifier's query set, as per Definition 6 of [HASW24c].

3.2 New security analysis

The security analysis in [HASW24b] proceeded in the following steps:

1. Show that any opening-consistent IOPP (P, V) for $\mathcal{C}$ (non-adaptive and query-selectable) yields a secure erasure-code commitment scheme for $\mathcal{C}$. Secure here means (a) *position-binding* and (b) *code-binding*.

2. Use the result from [HASW24a] to turn this into a data availability sampling scheme.

We need to revisit the first step, as we changed the opening-consistency notion. It should be noted that (a) position-binding directly follows from position-binding of Merkle trees (this is shown in [HASW24b]) and therefore this proof still applies without any change. For (b) code-binding we show how to adapt the proof. We refer the reader to [HASW24b] for the construction of the erasure-code commitment scheme CC from (P, V).

Theorem 3.4. *Let $\mathcal{C}$ be a code. Assume that (P, V) is a non-adaptive query-selectable IOPP for $\mathcal{C}$, and assume that it is opening-consistent with assign with errors $\mathsf{err}_1, \mathsf{err}_2$. Let $\mathsf{H}\colon \{0,1\}^* \to \{0,1\}^\lambda$, $\mathsf{H}' : \{0,1\}^* \to \{0,1\}^\lambda$, and $\hat{\mathsf{H}} : \{0,1\}^* \to \{0,1\}^z$ be random oracles.*

Then, for every algorithm $\mathcal{A}$ that makes at most Q_{H}, $Q_{\mathsf{H}'}$, $Q_{\hat{\mathsf{H}}}$ queries[4] to the random oracle H, H', $\hat{\mathsf{H}}$ respectively, we have

$$\mathsf{Adv}^{\mathsf{code-bind}}_{\mathcal{A},\mathsf{CC}}(\lambda) \leq \frac{2Q_{\mathsf{H}'}^2 + Q_{\mathsf{H}'}Q_{\hat{\mathsf{H}}} + 2Q_{\mathsf{H}}^2 + Q_{\mathsf{H}}Q_{\mathsf{H}'}}{2^\lambda} + Q_{\mathsf{H}'} \max\{\mathsf{err}_1, \mathsf{err}_2^L\}.$$

Proof. The proof is almost identical to the equivalent theorem in FRIDA [HASW24b] (ePrint version: [HASW24c]). We follow the same game-hopping strategy using the same games $\mathbf{G}_0$-$\mathbf{G}_4$. Our new proof only differs in game $\mathbf{G}_4$, at the claim labelled *Second Claim*. We replace the claim in FRIDA by the following:

> *Second Claim.* Assume that game $\mathbf{G}_4$ outputs 1 and let $\boldsymbol{c}^* \in \mathcal{C}$ denote the output of $\mathsf{Assign}(T_r^*)$, which is not $\perp$ due to the claim labelled *First Claim*[5] and the *suitable assignment* property. Then, there is a $j^* \in J$ such that $\boldsymbol{c}_{j^*} \neq \boldsymbol{c}^*_{j^*}$ and $j^* \in \mathcal{Q}_0(T_r^* \circ \rho_{j^*,r+1})$.

The proof of our modified claim is identical to the proof of the original *Second Claim*. Finally, the proof of our theorem above is completed identically to the proof for the original FRIDA theorem, by applying the *inconsistent is rejected* property. □

[4] These queries include the queries made by the experiment to check the winning condition.
[5] recall that *First Claim* establishes that T_r^* is suitable with respect to Bad and Lucky.

4 Instantiation from batched FRI

We now show that Theorem 3.4 can be applied to batched FRI [BCI+20]. To do so, we need to prove that batched FRI is a non-adaptive IOPP (Section 4.1), that it has query-selection (Section 4.2) and that it satisfies our notion of opening-consistency with assign (Section 4.3). In contrast to FRIDA [HASW24b], we go beyond the unique decoding radius, which yields better parameters. The reader should consult Section 4 and Appendix A of [GMW25] respectively for a description of FRI and batched FRI as well as their security properties.

We use the same notation and parameters as given in [GMW25]. In particular, for r folding rounds, we define the RS codes $\mathcal{C}_0, \ldots, \mathcal{C}_r$ on the domains $\mathcal{L}_0, \ldots, \mathcal{L}_r$ with degree bounds $d_0, \ldots, d_r$, and let $m_1, \ldots, m_r$ denote the folding factors. Throughout this section, we consider the batched FRI IOPP with batch size b and proximity generator $\mathsf{Gen} : \mathsf{S} \to \mathbb{F}^b$ for the batching round with mutual correlated agreement error $\varepsilon^{\mathsf{mca}}_{\mathsf{Gen}}$.

Remark 4.1 (Repetition parameters)**.** It is typical for IOPPs to run multiple repetitions of a query phase in order to obtain satisfactory soundness guarantees. In our reference FRI analysis [GMW25], the number of repetitions is a parameter of the security analysis and is denoted t. Similarly, the IOPP-to-CC compiler of FRIDA [HASW24b] has a repetition parameter L that controls the number of executions of the final round of the IOPP. When applying the IOPP-to-CC compiler to batched FRI, both parameters t and L control the same variable: the number of queries that the FRI verifier makes to the initial codeword. To avoid redundant parameters, we set $t = 1$ and adjust the number of queries using L as per Theorem 3.4.

4.1 Non-adaptive

Informally, we say that an IOPP is non-adaptive (see [HASW24c, Definition 7]) if the locations of all of the queries made by the verifier are identical for transcripts $T := (\boldsymbol{c}, \rho_1, \pi_1, \ldots, \rho_r, \pi_r, \rho_{r+1})$ and $T' := (\boldsymbol{c}', \rho_1, \pi'_1, \ldots, \rho_r, \pi'_r, \rho_{r+1})$. In other words, the definition states that the initial function and the prover's messages do not affect the locations of the verifier's queries.

By construction, the queries made by the batched FRI verifier are determined by the final verifier message, which is sampled uniformly at random from a fixed set. Therefore, batched FRI satisfies non-adaptivity.

4.2 Query-selection

Informally, we say that an IOPP has query-selection (see [HASW24c, Definition 8]) if there exists a deterministic algorithm $\mathsf{QSelect}$ that, given the verifier randomness $\rho_1, \ldots, \rho_r$ and a target query location $s \in \mathcal{L}_0$, outputs a final verifier message ρ_{r+1} such that the IOPP verifier will query location s.

Recall from Remark 4.1 that we are in the setting where FRI is run with a single repetition of the query phase. Thus, the final verifier message is a single value $\rho_{r+1} \in \mathcal{L}_0$ that the verifier will query. Therefore, for batched FRI, $\mathsf{QSelect}$ simply returns the target location s.

4.3 Opening-consistency with assign

We now prove the key technical contribution of this note, namely that batched FRI has opening-consistency with assign. Note that our analysis is given with respect to a parameter θ. This parameter reflects an inherent tradeoff in the security analysis of (batched) FRI: one can trade soundness in the folding phase for better soundness in the query phase, and vice-versa. This tradeoff carries over to the code-commitment construction where err_1 is an increasing function of θ while err_2 is a decreasing function of θ. Given a target security level, practitioners can

minimize the number of repetitions L by setting θ to be as large as possible while still keeping the quantity err_1 below the target level.

Definition 4.2 (Lucky set for batched FRI)**.** Consider the batched FRI IOPP, let $\theta \in [0,1]$ be a tradeoff parameter and State_θ denote the state function for batched FRI [GMW25, Appendix A]. For every $i \in [r]$, verifier-turn transcript $T := (\boldsymbol{c}, \rho_1, \pi_1, \ldots, \rho_{i-1}, \pi_{i-1})$ and verifier message ρ_i, we say that $T \circ \rho_i$ is in Lucky_θ if and only if:

- $\mathsf{State}_\theta(T) = 0$, and;
- $\mathsf{State}_\theta(T \circ \rho_i) = 1$.

Definition 4.3 (Bad set for batched FRI)**.** Consider the batched FRI IOPP, let $\theta \in [0,1]$ be a tradeoff parameter and State_θ denote the state function for batched FRI [GMW25, Appendix A]. We say that $T := (\boldsymbol{c}, \rho_1, \pi_1, \ldots, \rho_r, \pi_r)$ is in Bad_θ if and only if

$$\Pr_{\rho_{r+1}}[\mathsf{State}_\theta(T \circ \rho_{r+1}) = 1] < 1 - \theta.$$

Definition 4.4 (Assign function for batched FRI)**.** Consider the batched FRI IOPP, let $\theta \in [0, 1-\rho]$ be a tradeoff parameter and State_θ denote the state function for batched FRI [GMW25, Appendix A]. We define the Assign_θ function using the following deterministic procedure on input $T := (\boldsymbol{c}, \rho_1, \pi_1, \ldots, \rho_r, \pi_r)$:

1. if $\boldsymbol{c} \in {\mathcal{C}_0}^{\equiv b}$, return $\boldsymbol{c}$;
2. else,
   (a) compute the set $\mathsf{AccQueries} \subseteq \mathcal{L}_0$ of accepting final round verifier queries for the transcript T;
   (b) if $|\mathsf{AccQueries}| < (1-\theta) \cdot |\mathcal{L}_0|$, return $\perp$;
   (c) else, run the extractor $\mathbf{E}$ (as defined in the proof of Corollary A.4 of [GMW25]). That is, output $\boldsymbol{c}^* := \mathbf{E}(T)$.

Lemma 4.5 (No Luck)**.** *For every tradeoff parameter $\theta \in [0,1]$, the batched FRI IOPP satisfies the* no luck *property of opening-consistency with assign (Definition 3.1) with Lucky_θ as in Definition 4.2 and error err_1:*

$$\mathsf{err}_1 \leq \max\{\varepsilon^{\mathsf{mca}}_{\mathsf{Gen}}(\mathcal{C}_0, b, \theta), \varepsilon^{\mathsf{mca}}_{\mathsf{Powers}}(\mathcal{C}_0, m_0, \theta), \ldots, \varepsilon^{\mathsf{mca}}_{\mathsf{Powers}}(\mathcal{C}_k, m_k, \theta)\}.$$

Proof. Consider $i \in [r]$, and any partial verifier-turn transcript $T := (\boldsymbol{c}, \rho_1, \pi_1, \ldots, \rho_{i-1}, \pi_{i-1})$ that contains $2i-1$ elements. It follows by the definition of round-by-round soundness that

$$\Pr_{\rho_i}[\mathsf{State}_\theta(T) = 0 \wedge \mathsf{State}_\theta(T \circ \rho_i) = 1] \leq \varepsilon_i(\delta).$$

Taking the maximum over all rounds and substituting $\varepsilon_i(\delta)$ with the upper bound from [GMW25] proves the lemma. □

Lemma 4.6 (Bad is Rejected)**.** *For every tradeoff parameter $\theta \in [0,1]$, the batched FRI IOPP satisfies the* bad is rejected *property of opening-consistency with assign (Definition 3.1) with Bad_θ as in Definition 4.3 and error $\mathsf{err}_2 < 1 - \theta$.*

Proof. Consider $T \in \mathsf{Bad}_\theta$. Since $T \circ \rho_{r+1}$ is a complete transcript, it holds by the definition of the state function State_θ that if $\mathsf{State}_\theta(T \circ \rho_{r+1}) = 0$, then $\mathsf{V}^{\boldsymbol{c},\pi_1,\ldots,\pi_r}(\rho_1,\ldots,\rho_r,\rho_{r+1}) = 0$. Therefore

$$\Pr_{\rho_{r+1}}[\mathsf{V}^{\boldsymbol{c},\pi_1,\ldots,\pi_r}(\rho_1,\ldots,\rho_r,\rho_{r+1}) = 1] \leq \Pr_{\rho_{r+1}}[\mathsf{State}_\theta(T \circ \rho_{r+1}) = 1].$$

By definition of Bad_θ, it holds that $\Pr_{\rho_{r+1}}[\mathsf{State}_\theta(T \circ \rho_{r+1}) = 1] < 1 - \theta$. Chaining both inequalities proves the lemma. □

Lemma 4.7 (Suitable Assignment)**.** *For every tradeoff parameter $\theta \in [0, 1-\rho]$, the batched FRI IOPP satisfies the* suitable assignment *property of opening-consistency with assign (Definition 3.1) with Bad_θ as in Definition 4.3, Lucky_θ as in Definition 4.2 and Assign_θ as in Definition 4.4.*

Proof. Consider a suitable transcript T. First, notice that if $\boldsymbol{c} \in {\mathcal{C}_0}^{\equiv b}$, then *suitable assignment* holds trivially by Item 1 of the Assign function. Therefore, we focus on the case $\boldsymbol{c} \notin {\mathcal{C}_0}^{\equiv b}$.

We now argue that

$$\mathsf{State}_\theta(T) = 0. \tag{1}$$

By definition of round-by-round soundness, it holds that $\mathsf{State}_\theta(\boldsymbol{c}) = 0$. Assume towards contradiction that $\mathsf{State}_\theta(T) = 1$. Then, there exists a round i and prefix T' such that $\mathsf{State}_\theta(T') = 0$ and $\mathsf{State}_\theta(T' \circ \rho_i) = 1$. Therefore, the prefix $T' \circ \rho_i$ is in Lucky_θ, contradicting the premise that T is a suitable transcript.

By the definition of a suitable transcript, $T \notin \mathsf{Bad}_\theta$. Therefore,

$$\Pr_{\rho_{r+1}}[\mathsf{State}_\theta(T \circ \rho_{r+1}) = 1] \geq 1 - \theta. \tag{2}$$

Given Equations (1) and (2), we can apply the additional properties of batched FRI from [GMW25] with proximity parameter $\delta^* = \theta$. Using the same argument as for the *proximity* property, it holds that $|\mathsf{AccQueries}| \geq (1-\theta) \cdot |\mathcal{L}_0|$. Therefore, Assign_θ does not output $\perp$ on Item 2b. Finally, we complete our proof of the *suitable assignment* property by noting that $\mathbf{E}$ does not abort, by definition. □

Lemma 4.8 (Inconsistent is Rejected)**.** *For every tradeoff parameter $\theta \in [0, 1-\rho]$, the batched FRI IOPP satisfies the* inconsistent is rejected *property of opening-consistency with assign (Definition 3.1) with Bad_θ as in Definition 4.3, Lucky_θ as in Definition 4.2 and Assign_θ as in Definition 4.4.*

Proof. Consider a suitable transcript T. First, notice that if $\boldsymbol{c} \in {\mathcal{C}_0}^{\equiv b}$, then $\boldsymbol{c}^* = \boldsymbol{c}$ by definition of the Assign_θ function. Therefore there are no inconsistent queries.

In the case $\boldsymbol{c} \notin {\mathcal{C}_0}^{\equiv b}$, then Equations (1) and (2) hold. We can therefore apply the *erasure detection* property of [GMW25] with distance parameter $\delta^* = \theta$ and conclude that if there exists a position $j \in \mathcal{Q}_0(T \circ \rho_{r+1})$ such that $\boldsymbol{c}_j \neq \boldsymbol{c}^*_j$, then $j \notin \mathsf{AccQueries}$ and $\mathsf{V}^{\boldsymbol{c},\pi_1,\ldots,\pi_r}(\rho_1,\ldots,\rho_r,\rho_{r+1}) = 0$. □

Theorem 4.9 (Batched FRI has opening-consistency with assign)**.** *For every tradeoff parameter $\theta \in [0, 1-\rho]$, FRI has opening-consistency with assign with errors $\mathsf{err}_1, \mathsf{err}_2$ with respect to the sets Bad_θ as in Definition 4.3, Lucky_θ as in Definition 4.2 and the function Assign_θ as in Definition 4.4, where:*

$$\begin{aligned}
&\mathsf{err}_1 \leq \max\{\varepsilon^{\mathsf{mca}}_{\mathsf{Gen}}(\mathcal{C}_0, b, \theta), \varepsilon^{\mathsf{mca}}_{\mathsf{Powers}}(\mathcal{C}_0, m_0, \theta), \ldots, \varepsilon^{\mathsf{mca}}_{\mathsf{Powers}}(\mathcal{C}_k, m_k, \theta)\},\\
&\mathsf{err}_2 \leq 1 - \theta.
\end{aligned}$$

5 Evaluation

To evaluate the effect of our new analysis, we give a comparison of different instantiations of FRIDA in Table 1. We vary the rate ρ of the Reed-Solomon code, and the tradeoff parameter θ used for the analysis. The latter is set either in the unique decoding regime (UDR), where the original FRIDA analysis also applies, or in the Johnson bound regime (JBR) where our novel analysis holds. Other than that, we consider the same setting and metrics as in the original FRIDA paper [HASW24c], Section 5. To generate the table and plots we have incorporated our new analysis into the scripts from the FRIDA paper. Our modified code can be found at

https://github.com/nmohnblatt/da-efficiency-scripts

Table 1: Comparison of different instantiations of FRIDA, for different rates and analyzed in the unique decoding (UDR) or Johnson bound (JBR) regime.

				Communication	
Rate	Regime	Commitment [KB]	Encoding [MB]	Per Query [KB]	Total [MB]
$\rho = 1/2$	UDR	735.33	499.12	3.81	610.08
	JBR	670.59	499.12	3.81	610.08
$\rho = 1/4$	UDR	464.83	1031.80	3.94	444.34
	JBR	338.88	1031.80	3.94	444.34
$\rho = 1/8$	UDR	394.59	2130.71	4.07	394.70
	JBR	232.03	2130.71	4.07	394.70
$\rho = 1/16$	UDR	369.28	4395.63	4.19	374.07
	JBR	180.64	4395.63	4.19	374.07
$\rho = 1/32$	UDR	363.26	9059.70	4.32	365.06
	JBR	151.58	9059.70	4.32	365.06

Discussion. We can interpret the results as follows:

- *Erasure code contribution.* We first comment on the standard behavior of erasure codes. As the rate decreases, the code introduces more redundancy. Therefore, as shown in the table, the encoding size increases. At the same time, the added redundancy means that a smaller total number of samples is required to achieve the DAS security guarantees, which explains the reduction in total communication size.
- *Moving to the JBR.* As expected, using JBR instead of UDR for the analysis reduces the commitment size (for fixed rate), while keeping all other metrics the same.
- *UDR vs JBR as rate decreases.* The reduction mentioned above is small for large rates and increases as the rate gets smaller. This effect is illustrated in Figure 2 where we plot the commitment size against the inverse rate $1/\rho$. It is explained by noting that the gap between the UDR and JBR grows as ρ approaches 0.

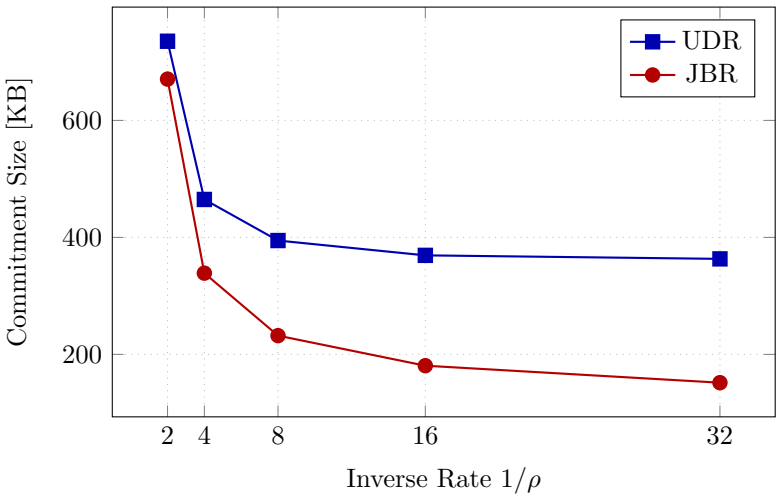

Figure 2: Commitment size of FRIDA for UDR and JBR across different rates.

Acknowledgements

We thank Albert Garreta, Mark Simkin and Ignacio Manzur for early discussions on the topic.

References

[ABN23] Daniel Augot, Sarah Bordage, and Jade Nardi. Efficient multivariate low-degree tests via interactive oracle proofs of proximity for polynomial codes. *DCC*, 91(3):1111–1151, 2023. doi:10.1007/s10623-022-01134-z.

[ACFY24] Gal Arnon, Alessandro Chiesa, Giacomo Fenzi, and Eylon Yogev. STIR: Reed-solomon proximity testing with fewer queries. In Reyzin and Stebila [RS24], pages 380–413. doi:10.1007/978-3-031-68403-6_12.

[ACFY25] Gal Arnon, Alessandro Chiesa, Giacomo Fenzi, and Eylon Yogev. WHIR: Reed-solomon proximity testing with super-fast verification. In Fehr and Fouque [FF25], pages 214–243. doi:10.1007/978-3-031-91134-7_8.

[AHIV17] Scott Ames, Carmit Hazay, Yuval Ishai, and Muthuramakrishnan Venkitasubramaniam. Ligero: Lightweight sublinear arguments without a trusted setup. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 2017*, pages 2087–2104. ACM Press, October / November 2017. doi:10.1145/3133956.3134104.

[ASBK21] Mustafa Al-Bassam, Alberto Sonnino, Vitalik Buterin, and Ismail Khoffi. Fraud and data availability proofs: Detecting invalid blocks in light clients. In Nikita Borisov and Claudia Díaz, editors, *FC 2021, Part II*, volume 12675 of *LNCS*, pages 279–298. Springer, Berlin, Heidelberg, March 2021. doi:10.1007/978-3-662-64331-0_15.

[BBHR18] Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. Fast reed-solomon interactive oracle proofs of proximity. In Ioannis Chatzigiannakis, Christos Kaklamanis, Dániel Marx, and Donald Sannella, editors, *ICALP 2018*, volume 107 of *LIPIcs*, pages 14:1–14:17. Schloss Dagstuhl, July 2018. doi:10.4230/LIPIcs.ICALP.2018.14.

[BCF+25] Martijn Brehm, Binyi Chen, Ben Fisch, Nicolas Resch, Ron D. Rothblum, and Hadas Zeilberger. Blaze: Fast SNARKs from interleaved RAA codes. In Fehr and Fouque [FF25], pages 123–152. doi:10.1007/978-3-031-91134-7_5.

[BCI+20] Eli Ben-Sasson, Dan Carmon, Yuval Ishai, Swastik Kopparty, and Shubhangi Saraf. Proximity gaps for reed-solomon codes. In *61st FOCS*, pages 900–909. IEEE Computer Society Press, November 2020. doi:10.1109/FOCS46700.2020.00088.

[BCS16] Eli Ben-Sasson, Alessandro Chiesa, and Nicholas Spooner. Interactive oracle proofs. In Martin Hirt and Adam D. Smith, editors, *TCC 2016-B, Part II*, volume 9986 of *LNCS*, pages 31–60. Springer, Berlin, Heidelberg, October / November 2016. doi:10.1007/978-3-662-53644-5_2.

[BFMM26] Anubhav Baweja, Giacomo Fenzi, Pratyush Mishra, and Tushar Mopuri. Field-agnostic SNARKs with small proofs via encode-repeat-accumulate (ERA) codes. Cryptology ePrint Archive, Paper 2026/864, 2026. URL: https://eprint.iacr.org/2026/864.

[BFRW25] Benedikt Bünz, Giacomo Fenzi, Ron D. Rothblum, and William Wang. Tensor-Switch: Nearly optimal polynomial commitments from tensor codes. Cryptology ePrint Archive, Report 2025/2065, 2025. URL: https://eprint.iacr.org/2025/2065.

[BGKS20] Eli Ben-Sasson, Lior Goldberg, Swastik Kopparty, and Shubhangi Saraf. DEEP-FRI: Sampling outside the box improves soundness. In Thomas Vidick, editor, *ITCS 2020*, volume 151, pages 5:1–5:32. LIPIcs, January 2020. doi:10.4230/LIPIcs.ITCS.2020.5.

[BMMS25a] Anubhav Baweja, Pratyush Mishra, Tushar Mopuri, and Matan Shtepel. FICS and FACS: Fast IOPPs and accumulation via code-switching. Cryptology ePrint Archive, Report 2025/737, 2025. URL: https://eprint.iacr.org/2025/737.

[BMMS25b] Anubhav Baweja, Pratyush Mishra, Tushar Mopuri, and Matan Shtepel. Query-optimal IOPPs for linear-time encodable codes. Cryptology ePrint Archive, Report 2025/1588, 2025. URL: https://eprint.iacr.org/2025/1588.

[BNNP25] Dan Boneh, Joachim Neu, Valeria Nikolaenko, and Aditi Partap. Data availability sampling with repair. Cryptology ePrint Archive, Report 2025/1414, 2025. URL: https://eprint.iacr.org/2025/1414.

[CSK25] Seunghyun Cho, Eunyoung Seo, and Young-Sik Kim. Locally recoverable data availability sampling. Cryptology ePrint Archive, Report 2025/1851, 2025. URL: https://eprint.iacr.org/2025/1851.

[EA25] Alex Evans and Guillermo Angeris. The accidental computer: Polynomial commitments from data availability. Cryptology ePrint Archive, Report 2025/918, 2025. URL: https://eprint.iacr.org/2025/918.

[EMA25] Alex Evans, Nicolas Mohnblatt, and Guillermo Angeris. ZODA: Zero-overhead data availability. Cryptology ePrint Archive, Report 2025/034, 2025. URL: https://eprint.iacr.org/2025/034.

[FF25] Serge Fehr and Pierre-Alain Fouque, editors. *EUROCRYPT 2025, Part IV*, volume 15604 of *LNCS*. Springer, Cham, May 2025.

[GLS+23] Alexander Golovnev, Jonathan Lee, Srinath T. V. Setty, Justin Thaler, and Riad S. Wahby. Brakedown: Linear-time and field-agnostic SNARKs for R1CS. In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part II*, volume 14082 of *LNCS*, pages 193–226. Springer, Cham, August 2023. doi:10.1007/978-3-031-38545-2_7.

[GMW25] Albert Garreta, Nicolas Mohnblatt, and Benedikt Wagner. A simplified round-by-round soundness proof of FRI. Cryptology ePrint Archive, Report 2025/1993, 2025. URL: https://eprint.iacr.org/2025/1993.

[GNR26] Kobi Gurkan, Andrija Novakovic, and Ron D. Rothblum. Bolt: Faster SNARKs from sketched codes. Cryptology ePrint Archive, Paper 2026/310, 2026. URL: https://eprint.iacr.org/2026/310.

[HASW23] Mathias Hall-Andersen, Mark Simkin, and Benedikt Wagner. Foundations of data availability sampling. Cryptology ePrint Archive, Report 2023/1079, 2023. URL: https://eprint.iacr.org/2023/1079.

[HASW24a] Mathias Hall-Andersen, Mark Simkin, and Benedikt Wagner. Foundations of data availability sampling. *CiC*, 1(4):34, 2024. doi:10.62056/a09qudhdj.

[HASW24b] Mathias Hall-Andersen, Mark Simkin, and Benedikt Wagner. FRIDA: Data availability sampling from FRI. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part VI*, volume 14925 of *LNCS*, pages 289–324. Springer, Cham, August 2024. doi:10.1007/978-3-031-68391-6_9.

[HASW24c] Mathias Hall-Andersen, Mark Simkin, and Benedikt Wagner. FRIDA: Data availability sampling from FRI. Cryptology ePrint Archive, Report 2024/248, 2024. URL: https://eprint.iacr.org/2024/248.

[NA25] Andrija Novakovic and Guillermo Angeris. Ligerito: A small and concretely fast polynomial commitment scheme. Cryptology ePrint Archive, Report 2025/1187, 2025. URL: https://eprint.iacr.org/2025/1187.

[RR24] Noga Ron-Zewi and Ron Rothblum. Local proofs approaching the witness length. *J. ACM*, 71(3):18, 2024. doi:10.1145/3661483.

[RS24] Leonid Reyzin and Douglas Stebila, editors. *CRYPTO 2024, Part X*, volume 14929 of *LNCS*. Springer, Cham, August 2024.

[WSZ+26] Jiaxi Wang, Wenbo Sun, Ziyuan Zhou, Shihua Wu, Jiang Xu, and Shan Ji. Istirda: An efficient data availability sampling scheme for lightweight nodes in blockchain. *Computers, Materials and Continua*, 87(1), 2026. URL: https://www.sciencedirect.com/science/article/pii/S1546221826001979, doi:10.32604/cmc.2025.073237.

[WZ24] Benedikt Wagner and Arantxa Zapico. A documentation of ethereum's PeerDAS. Cryptology ePrint Archive, Report 2024/1362, 2024. URL: https://eprint.iacr.org/2024/1362.

[WZ25] Benedikt Wagner and Arantxa Zapico. Proving the security of PeerDAS without the AGM. Cryptology ePrint Archive, Report 2025/1683, 2025. URL: https://eprint.iacr.org/2025/1683.

[YZD26] Ming Yang, Xinxuan Zhang, and Yi Deng. Deep-frida: transparent data availability sampling with smaller size. *Cybersecur.*, 9(1):145, 2026. URL: https://doi.org/10.1186/s42400-026-00564-9, doi:10.1186/S42400-026-00564-9.

[ZCF24] Hadas Zeilberger, Binyi Chen, and Ben Fisch. BaseFold: Efficient field-agnostic polynomial commitment schemes from foldable codes. In Reyzin and Stebila [RS24], pages 138–169. doi:10.1007/978-3-031-68403-6_5.

A Additional comments on DEEP-FRIDA

In addition to our earlier comparison between DEEP-FRIDA [YZD26] and our work, we note a minor difference and technical issue that is easily resolved. First, their modification of opening-consistency allows a probabilistic algorithm to assign a codeword to suitable transcripts (as defined in [HASW24c, Definition 9]). This is in contrast to our definition (Definition 3.1) where the Assign function is deterministic. Second, their work does not show that the IOPP-to-CC compiler of FRIDA is secure under their modified definition of opening-consistency. Our security proof of the IOPP-to-CC compiler (see Theorem 3.4) relies on the fact that the Assign function is deterministic and always assigns a unique, fixed codeword per suitable transcript. It is therefore unclear whether a similar result holds for a probabilistic Assign function (or "extractor" in the terminology of DEEP-FRIDA).

This issue is only a minor technicality, which is resolved by adopting our notion of opening-consistency with assign. Indeed, Yang, Zhang and Deng [YZD26] prove that DEEP-FRI has opening-consistency (according to their definition) using a deterministic codeword-assigning algorithm, thus also satisfying our definition.