Proximity Gaps and Correlated Agreement: The Core of FRI Security Proof

- Jade Xie jade@secbit.io
- Yu Guo yu.guo@secbit.io

This article is mainly inspired by the video Proximity Gaps & Applications to Succinct Proofs, combined with the paper [BCIKS20], introducing the concept of Proximity Gaps and the closely related Correlated Agreement theorem, which play a very important role in the security proof of FRI.

In the FRI protocol, for a polynomial $f : \mathcal{D} \to \mathbb{F}_q$, let $f(x) = a_0 + a_1 x + a_2 x^2 + \ldots + a_{k-1} x^{k-1}$, which is a polynomial of degree less than k, evaluated on the domain $\mathcal{D}$, where $|\mathcal{D}| = n$, then $f \in \mathrm{RS}[\mathbb{F}_q, \mathcal{D}, k]$. The Prover wants to prove to the Verifier that the degree of $f(x)$ is indeed less than k. If $f \in \mathrm{RS}[\mathbb{F}_q, \mathcal{D}, k]$, the Verifier outputs `accept`, if f is δ far from the corresponding code space $\mathrm{RS}[\mathbb{F}_q, \mathcal{D}, k]$, it outputs `reject`. What the Verifier can obtain is the oracle about a series of functions, and what the FRI protocol wants to achieve is that the Verifier queries the oracle as little as possible and can distinguish which of the above situations f belongs to.

Let's assume $k-1$ is even, then

$$
\begin{aligned}
f(x) &= a_0 + a_1 x + a_2 x^2 + \ldots + a_{k-1} x^{k-1} \\
&= (a_0 + a_2 x^2 + \cdots + a_{k-1} x^{k-1}) + x(a_1 + a_3 x^2 + \cdots + a_{k-2} x^{k-3}) \\
&:= g(x^2) + x h(x^2)
\end{aligned} \tag{1}
$$

We can find that the functions

$$
\begin{aligned}
g(x) &= a_0 + a_2 x + \cdots + a_{k-1} x^{\frac{k-1}{2}} \\
h(x) &= a_1 + a_3 x + \cdots + a_{k-2} x^{\frac{k-3}{2}}
\end{aligned} \tag{2}
$$

Initially, the Prover wants to prove to the Verifier that the degree of $f(x)$ is less than k, now it can be decomposed into three sub-problems:

1. Prove that the degree of function $g(x)$ is less than $k/2$, i.e., $g(x) \in \mathrm{RS}[\mathbb{F}_q, \mathcal{D}^{(1)}, k/2]$
2. Prove that the degree of function $h(x)$ is less than $k/2$, i.e., $h(x) \in \mathrm{RS}[\mathbb{F}_q, \mathcal{D}^{(1)}, k/2]$
3. Prove that $f(x) = g(x^2) + x \cdot h(x^2)$

where $|D^{(1)}| = n/2$. The third item is to prove that the odd-even splitting is correct. Similarly, $g(x)$ and $h(x)$ can be decomposed into odd and even terms like $f(x)$, decomposing them into two polynomials of degree less than $k/4$, so we need to prove that 4 polynomials are of degree less than $k/4$, until finally decomposing to prove constant polynomials. This process is shown in the figure below, and we can see that the polynomials to be proved are growing in the form of powers of 2. In this process, in order to prove that the odd-even splitting is not problematic, we need to send oracles about all these polynomials to the Verifier, and we can imagine that there are too many polynomials being sent, which grow explosively as k increases.

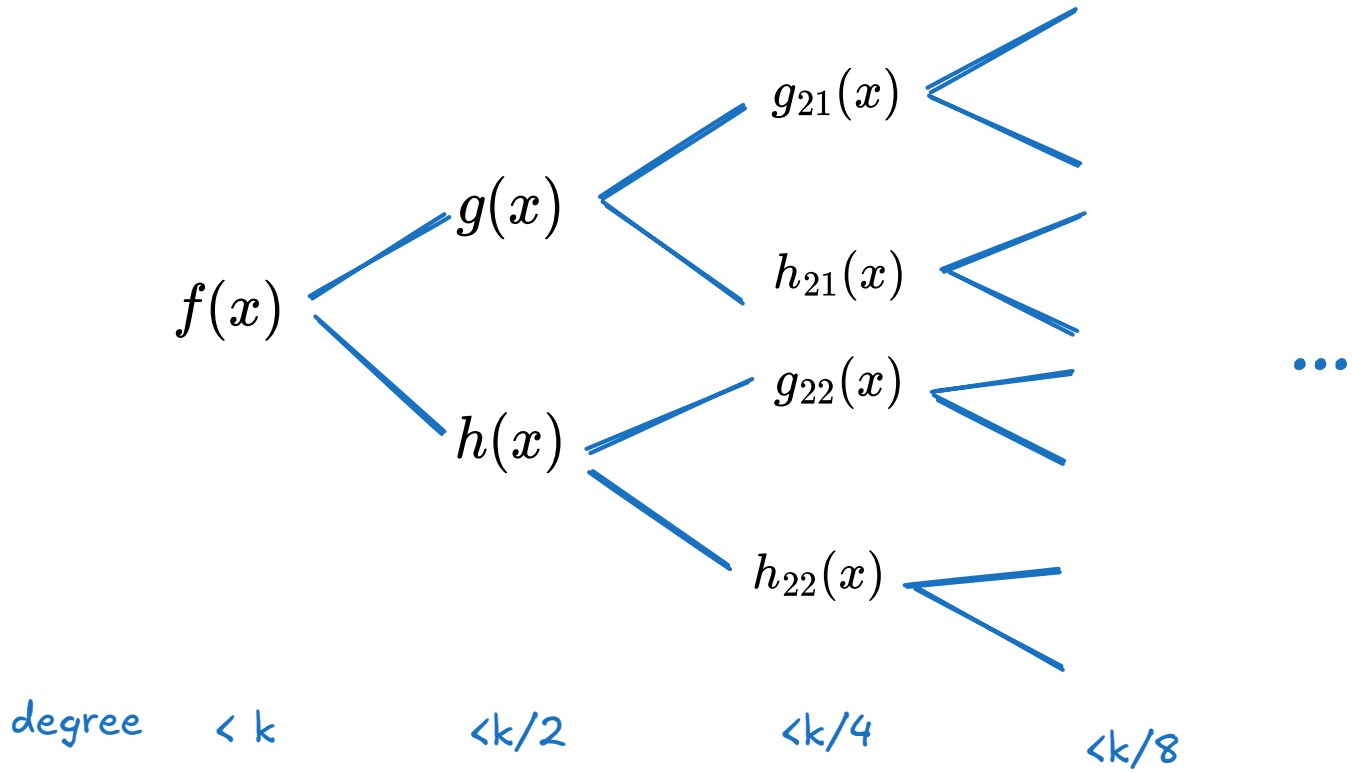

Since our purpose is to prove that the degree of the polynomial is less than a certain number, our idea is that we don't want to split the problem of $f(x)$ like above, splitting it into two polynomials, we want to prove in the next step that a polynomial is of degree less than $k/2$, which can greatly reduce the number of polynomials sent. How to do this? We can ask the Verifier for a random number $r \in \mathbb{F}$, make a linear combination of $g(x)$ and $h(x)$, get $g(x) + r \cdot h(x)$, and decompose the problem of $f(x)$ being of degree less than k into:

1. The degree of $f^{(1)}(x) = g(x) + r \cdot h(x)$ is less than $k/2$, i.e., $f^{(1)}(x) \in \mathrm{RS}[\mathbb{F}_q, \mathcal{D}^{(1)}, k/2]$

At this time, the graph of the polynomials to be sent becomes like the figure below, and you can see that the oracle of the polynomials to be sent is greatly reduced.

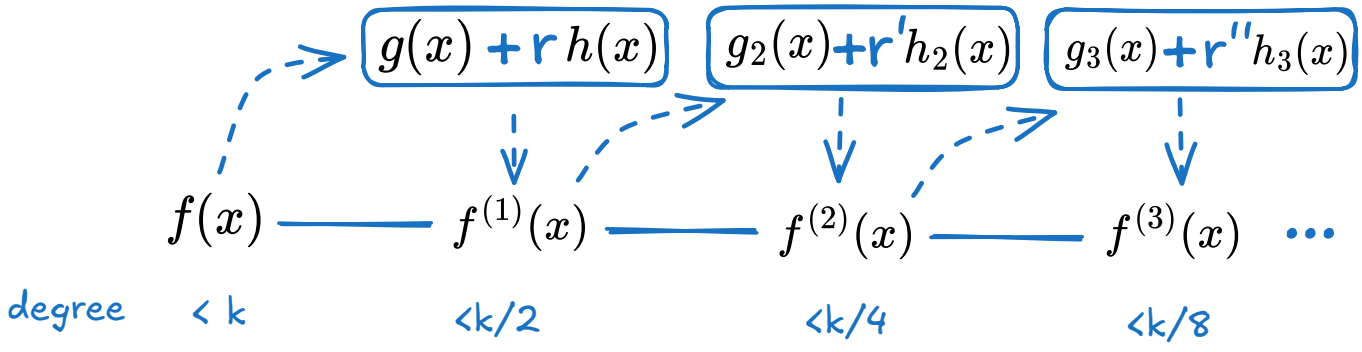

Now the remaining question is, is this equivalent to the original method? Of course, if the Prover is honest, according to the linearity of RS encoding, $g(x), h(x) \in \mathrm{RS}[\mathbb{F}_q, \mathcal{D}^{(1)}, k/2]$, then their linear combination is still in $\mathrm{RS}[\mathbb{F}_q, \mathcal{D}^{(1)}, k/2]$. But what if the Prover cheats? For example, if $g(x)$ is δ far from the code space $\mathrm{RS}[\mathbb{F}_q, \mathcal{D}^{(1)}, k/2]$, we hope that after the linear combination with the random number r, $g(x) + r \cdot h(x)$ is still δ far, so that the Verifier can discover the Prover's cheating. What we don't want is that the folded $g(x) + r \cdot h(x)$ becomes closer to the corresponding code space. Proximity Gaps tells us that the probability of this happening is very small, like winning the lottery, so we can boldly use random numbers

for folding.

Proximity Gaps

Above we considered the case of folding two polynomials, in practice we will use random numbers to fold multiple times or batch multiple polynomials. Here let's consider the general case, assuming there are m vectors $(u_0, \ldots, u_{m-1})$, for each $u_i \in \mathbb{F}_q^{\mathcal{D}}$, it can be seen as a polynomial on $\mathcal{D} \to \mathbb{F}$, or as a vector of dimension $|\mathcal{D}| = n$. Make a linear combination of these m vectors, denoted as $A = \text{span}\{u_0, \ldots, u_{m-1}\}$, where A is an affine space in $\mathbb{F}^{\mathcal{D}}$, and let the code space be $V := \text{RS}[\mathbb{F}, \mathcal{D}, k]$.

We are concerned about the distance relationship between elements in A and the code space V. As shown in the figure below, represent all codes in the code space V as points, draw a sphere with these points as the center and δ as the radius. The space formed by A is represented by a two-dimensional plane. If the elements in A have a relative Hamming distance less than or equal to δ from some codes in V, it means that they intersect with some Hamming balls in the figure, and all these intersections combined form the green shaded area in the figure. In other words, for each element a in the shaded area $S \subset A$, there must exist a $v \in V$, such that $\Delta(a, v) \leq \delta$.

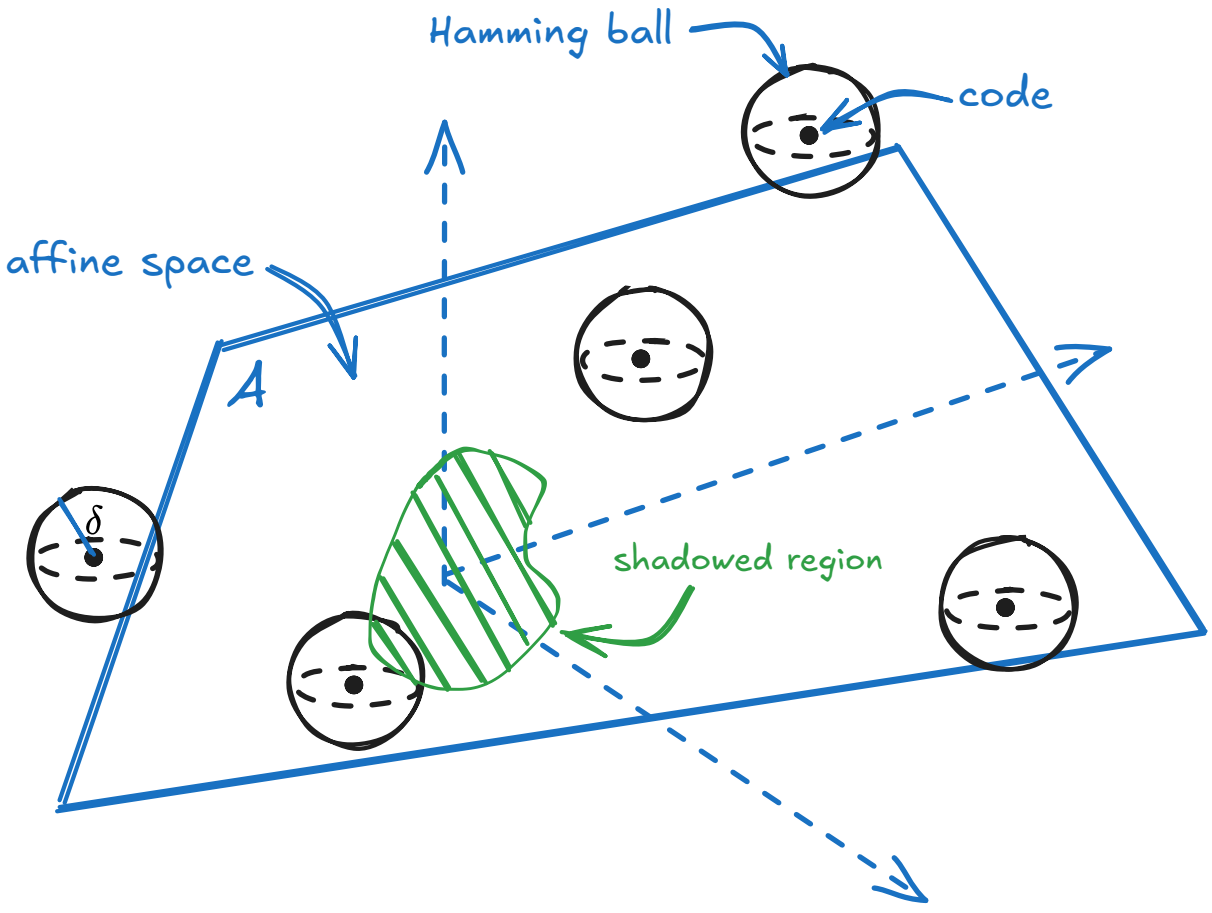

Let's form a set C_{Affine} consisting of all affine spaces in $\mathbb{F}^{\mathcal{D}}$. The Proximity Gaps conclusion [BCIKS20, Theorem 1.2] tells us that for any $A \in \text{C}_{\text{Affine}}$ (such as $A = \text{span}\{u_0, \ldots, u_{m-1}\}$), either all elements in A are in the shaded area, or only a very small part of the elements in A are in the shaded area. It's impossible to say that half of the elements in A are in the shaded area while the other half are not. Expressed in formula, it can only conform to one of the following two situations:

1. $\Pr_{a \in A}[\Delta(a, V) \leq \delta] \leq \epsilon$

2. $\Pr_{a\in A}[\Delta(a,V)\le\delta]=1$

We call δ the proximity parameter, and ϵ the error parameter, which is a very small number. Of course, there is a specific expression for ϵ, which is related to q,n,ρ,δ, that is, $\epsilon=\epsilon(q,n,\rho,\delta)$, where ρ represents the code rate, $\rho=\frac{k}{n}$.

So what does the shaded area here represent? What is the relationship between this conclusion and the security analysis of FRI? Let's analyze the application of the Proximity Gaps conclusion for the cases of honest Prover and cheating Prover.

Honest Prover

If it's an honest Prover, then for each vector in $(u_0,\dots,u_{m-1})$, we have $u_i\in V$.

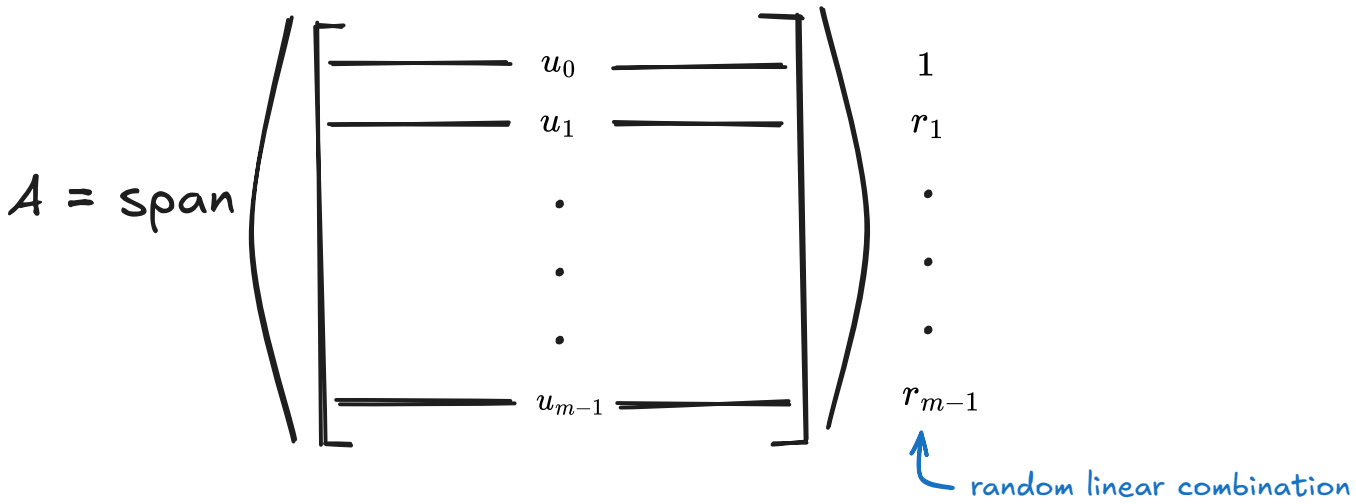

Due to the linearity of RS encoding, we know that after linear combination, it must still be in the code space V, so $A\subset V$. At this time, all elements in A are in V, so when the Verifier makes a random linear combination and arbitrarily selects a point $a\in A$, they will always get $a\in V$, and the Verifier will definitely accept. This situation corresponds to the second case in Proximity Gaps, taking $\delta=0$, at this time

$$\Pr_{a\in A}[\Delta(a,V)=0]=1 \tag{3}$$

Malicious Prover

If the Prover cheats, suppose one vector in the m vectors $\vec{u}=(u_0,\dots,u_{m-1})$ sent by the Prover to the Verifier is δ far from V, that is

$$\exists u_i^*\in\vec{u},\quad \Delta(u_i^*,V)>\delta \tag{4}$$

Then in $A=\text{span}\{u_0,u_1,\dots,u_{m-1}\}$, take $a^*=u_i^*\in A$, we must have

$$\exists a^*\in A,\quad \Delta(a^*,V)>\delta \tag{5}$$

At this time, according to the Proximity Gaps conclusion, there is already an element in A that is not in the shaded area, so the case $\Pr_{a\in A}[\Delta(a,V)\le\delta]=1$ is excluded, and it can only be $\Pr_{a\in A}[\Delta(a,V)\le\delta]\le\epsilon$. This also means that even if only one of the m vectors is δ far from the corresponding code space, most elements in A are δ far from V. In other words, a point a randomly selected from A can represent the farthest distance from V among the m vectors.

Now the Verifier randomly selects a point $a \in A$ to check whether $\Delta(a, V)$ is greater than δ. Two situations will occur. One is that it falls into the shaded area in the figure, and the other is that it falls outside the shaded area.

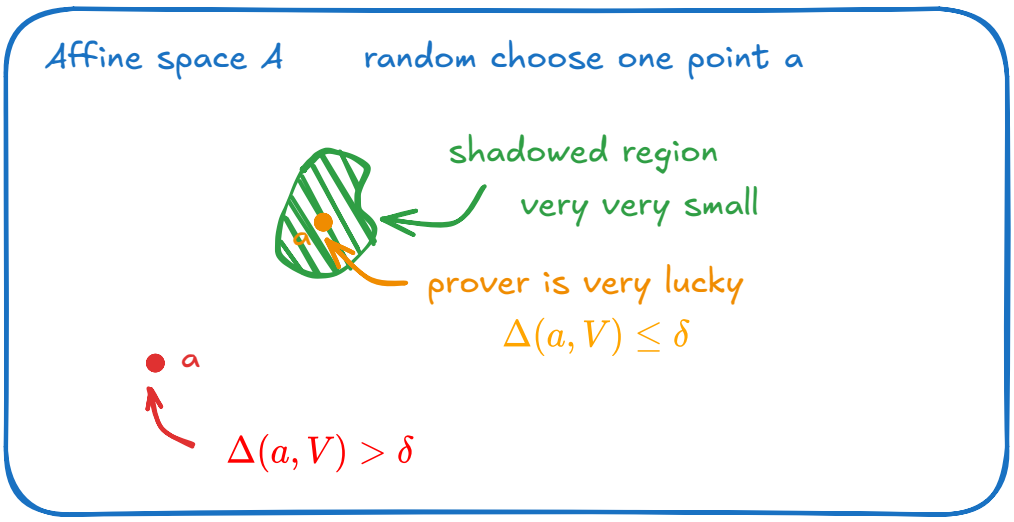

Case 1: $\Delta(a, V) \leq \delta$**.** At this time, the point a selected by the Verifier is in the shaded area. We say that the Prover is very lucky at this time. Although the Prover provided an incorrect witness, which is δ far from the code space, after random linear combination, it becomes δ close to the code space, and at this time, the Prover can successfully deceive the Verifier. The occurrence of this situation is not good for the Verifier, but fortunately, the Proximity Gaps conclusion tells us that $\mathrm{Pr}_{a \in A}[\Delta(a, V) \leq \delta] \leq \epsilon$, which means that the probability of randomly choosing a point that can enter the shaded area is very, very small. The Prover needs to be as lucky as winning the lottery, that is, at this time, the probability that the Prover can successfully deceive the Verifier will not exceed ϵ.

Case 2: $\Delta(a, V) > \delta$**.** At this time, the point a selected by the Verifier is outside the shaded area. Can the Prover still succeed in cheating? There is still a chance, because the Verifier received the oracle about a, but will not check all the values in a, only wants to query some values to see if it is in V. If the Verifier only queries once, since $\Delta(a, V) > \delta$, more than δ proportion of the components in a are not equal to the corresponding components of v, at this time the Verifier has a probability greater than δ to catch the Prover cheating, which means that at this time the probability that the Prover can cheat successfully does not exceed $1 - \delta$.

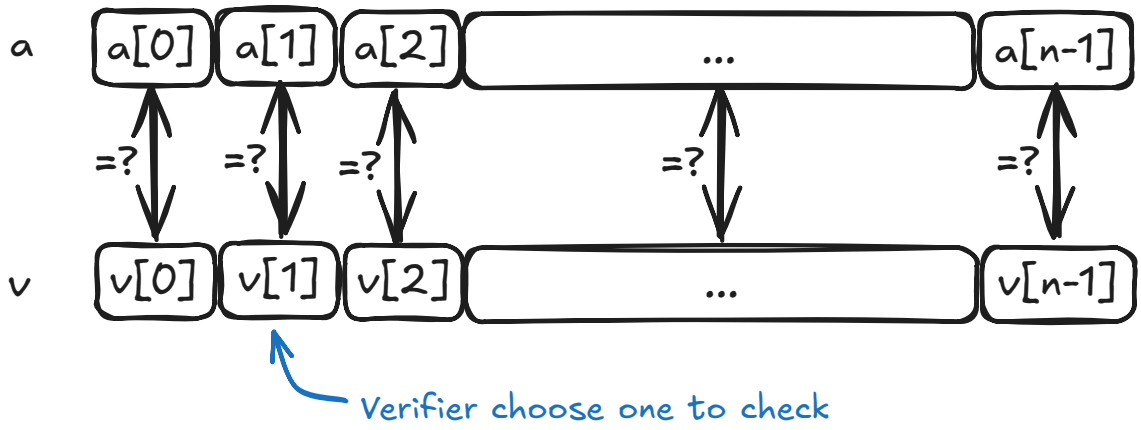

If the Verifier repeats the query κ times, the probability that the Prover can cheat successfully will not exceed $(1 - \delta)^{\kappa}$.

So, the probability that a cheating Prover can succeed is the joint probability of the above two cases, that is, it will not exceed

$$\epsilon + (1-\delta)^{\kappa} \tag{6}$$

The above analysis is actually the general idea of the soundness analysis of the FRI protocol. In the paper, the occurrence of case 1 is called the occurrence of some "bad" events, and then assuming that the "bad" events did not occur, estimate the probability of case 2, and finally combine the two for analysis.

We know that the FRI protocol is divided into two stages, one is the Commit stage and the other is the Query stage. We can correspond the above two cases to these two stages:

1. The above case 1 occurs in the Commit stage, where the Verifier will select random numbers to let the Prover fold the polynomials.
2. The above case 2 corresponds to the Query stage, where the Verifier will randomly select some points for query checks.

If it's a batched version of the FRI protocol, to prove multiple polynomials $f_0^{(0)}, f_1^{(0)}, \ldots, f_t^{(0)}$ are all polynomials of degree less than k, we can first use random numbers $\{x_1, \ldots, x_t\}$ for aggregation, obtaining

$$f^{(0)}(x) = f_0^{(0)} + \sum_{i=1}^{t} x_i \cdot f_i^{(0)} \tag{7}$$

Then apply the general FRI protocol to $f^{(0)}(x)$ to prove that it is a polynomial of degree less than k. The soundness analysis here also corresponds to the above case 1, that is, there may exist a situation where due to the selection of random numbers, $f^{(0)}(x)$ is no longer δ far from the corresponding RS code space.

Impact of Increasing δ

Let's analyze what impact the increase of the proximity parameter δ will bring. We have already analyzed that the probability of a cheating Prover successfully deceiving the Verifier does not exceed

$$\epsilon + (1-\delta)^{\kappa} \tag{8}$$

This probability consists of two parts, the increase of δ will lead to:

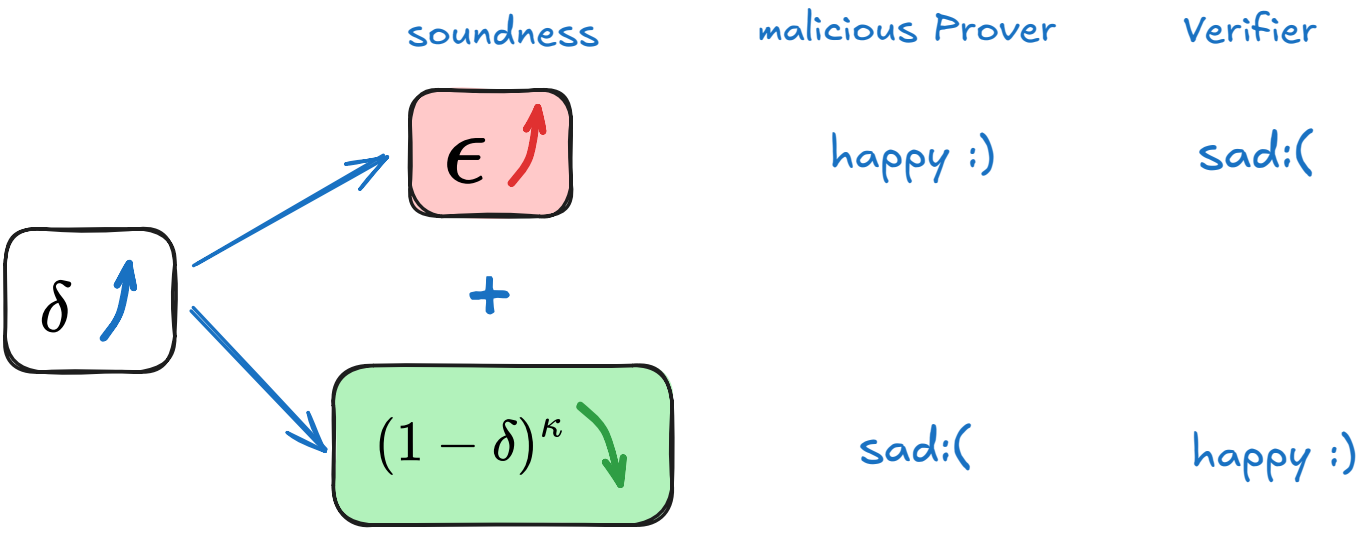

1. $\epsilon \uparrow$. From a graphical understanding, δ controls the radius of each Hamming ball. If δ increases, then the Hamming balls become larger, and their intersection with the affine space A should be larger, which means the shaded area increases, which implies that ϵ will increase.
    - This is good news for the cheating Prover :). Because at this time, the Prover becomes luckier than before, with a greater probability of entering the green shaded area, and can successfully deceive the Verifier.
    - Naturally, this is bad news for the Verifier :(.
2. $(1-\delta)^\kappa \downarrow$. This expression is directly related to δ, if δ increases, then $(1-\delta)^\kappa$ will decrease.
    - This is bad news for the cheating Prover :(. Because at this time, the probability of the Prover's successful cheating will decrease.
    - This is good news for the Verifier :). At this time, there is a greater probability of catching the Prover cheating. Under the same security requirements, the Verifier only needs fewer rounds of polling to meet the requirements.

It can be seen that the increase of δ causes ϵ to increase and $(1-\delta)^\kappa$ to decrease. In practice, ϵ is very small, and $(1-\delta)^\kappa$ accounts for a larger proportion in the whole sum, so the overall will still decrease, which means that for the entire FRI protocol, the soundness decreases, indicating that it will be more secure.

The above analysis is from the perspective of soundness. The video Proximity Gaps & Applications to Succinct Proofs also mentions a point that the increase of δ will make the corresponding Correlated Agreement related conclusions weaker. Correlated Agreement is a stronger conclusion than Proximity Gaps (so far, their equivalence has not been proven). Let's introduce the Correlated Agreement conclusion below.

Correlated Agreement

For the affine space $A = \text{span}\{u_0, u_1, \ldots, u_{m-1}\}$ mentioned earlier, to maintain consistency with [BCIKS20, Theorem 1.6], we don't use a random number before the first vector u_0, let $A = u_0 + \text{span}\{u_1, \ldots, u_{m-1}\}$.

The Correlated Agreement theorem ([BCIKS20, Theorem 1.6]) says that if $\delta \in (0, 1-\sqrt{\rho})$ and

$$\Pr_{a \in A}[\Delta(a, V) \leq \delta] > \epsilon, \tag{9}$$

where ϵ is the ϵ given in the Proximity Gaps conclusion, then there exist $\mathcal{D}' \subset \mathcal{D}$ and $v_0, \ldots, v_{m-1} \in V$ such that

1. **Density**: $\frac{|\mathcal{D}'|}{|\mathcal{D}|} \geq 1-\delta$,
2. **Agreement**: For any $i \in \{0, \ldots, m-1\}$, we have $u_i|_{\mathcal{D}'} = v_i|_{\mathcal{D}'}$.

This means that if there are many elements falling into the shaded area, with a proportion larger than ϵ in the Proximity Gaps conclusion, then there exist codewords $v_0, \ldots, v_{m-1}$ in V, and there will be a subset $\mathcal{D}'$ in the domain $\mathcal{D}$ with a very large proportion (more than $1-\delta$), where each u_i is consistent with the corresponding v_i on $\mathcal{D}'$.

According to the conclusion of Proximity Gaps, the elements in A fall into two categories:

1. $\Pr_{a \in A}[\Delta(a, V) \leq \delta] \leq \epsilon$

2. $\Pr_{a\in A}[\Delta(a,V)\le\delta]=1$

Now, if the proportion of elements falling into the shaded area is greater than ϵ , we can naturally exclude the first case. This leads to the conclusion that all elements in A fall within the shaded area, i.e.,

$$\Pr_{a\in A}[\Delta(a,V)\le\delta]=1.$$

The correlated agreement theorem provides a more specific conclusion. It describes the relationship between the elements u_i before folding and the codewords v_i found in the encoding space V .

For example, if the Prover wants to prove that a polynomial $f\in \mathrm{RS}[\mathbb{F}_q,\mathcal{D}^{(0)},k]$, let $\mathcal{D}^{(0)}=\{x_1,\dots,x_n\}$, calculate $\{f(x_1),\dots,f(x_n)\}$, the Prover will send the oracle of these values to the Verifier. In practice, Merkle trees are used to implement the oracle.

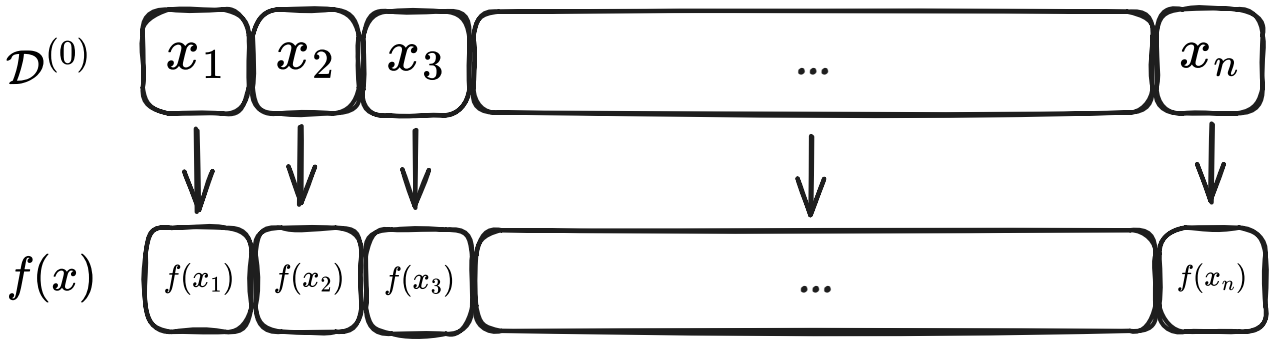

Split f to obtain two polynomials $g(x)$ and $h(x)$. In the honest case, $g,h\in \mathrm{RS}[\mathbb{F}_q,\mathcal{D}^{(1)},k/2]$, where $|\mathcal{D}^{(1)}|=|\mathcal{D}^{(0)}|/2=n/2$.

The Correlated Agreement conclusion tells us that for the affine space $A=\{g+z\cdot h : z\in\mathbb{F}\}$ formed by $g(x)$ and $h(x)$, if more than ϵ proportion of elements in A fall into the "shaded area", i.e., satisfying $\Delta(a,V)\ge\delta$, then there exist $\mathcal{D}'$ as shown in the figure below, and $\bar{g},\bar{h}\in \mathrm{RS}[\mathbb{F}_q,\mathcal{D}^{(1)},k/2]$. Let's assume $\mathcal{D}'=\{\alpha_1,\alpha_2,\dots,\alpha_i\}$, then according to the conclusion $|\mathcal{D}'|/|\mathcal{D}^{(1)}|\ge 1-\delta$, we have index $i\ge(1-\delta)n/2$. On all $\mathcal{D}'$, g is consistent with $\bar{g}$, and h is consistent with $\bar{h}$, which is represented in green in the figure, meaning that when evaluating at these points in the $\mathcal{D}'$ set, their values are the same.

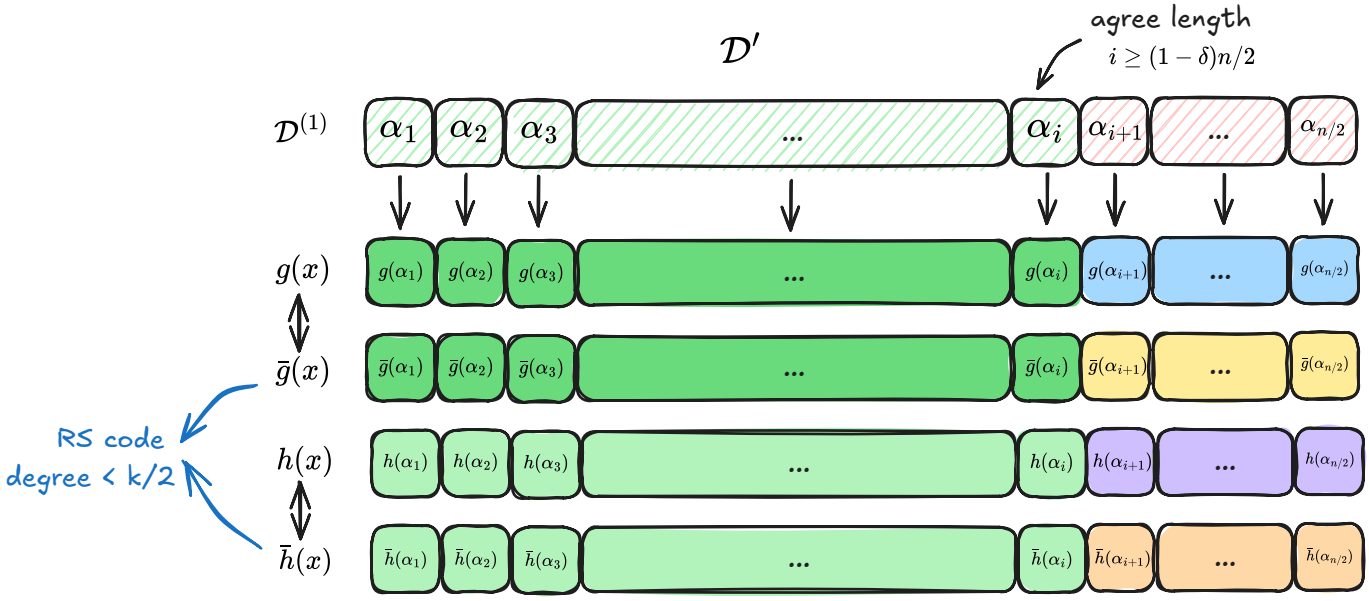

Back to the analysis of δ increase, we can see that as δ increases, the $1-\delta$ in the first condition **Density** of the Correlated Agreement conclusion will become smaller, which makes the subset $\mathcal{D}'$ that can be ensured to exist in the conclusion smaller, making the obtained conclusion weaker.

In the [BCIKS20] paper, it is said that the Proximity Gap theorem ([BCIKS20, Theorem 1.2]) is derived from the Correlated Agreement theorem ([BCIKS20, Theorem 1.6]), but it is not known yet whether the Proximity Gap theorem can derive the Correlated Agreement theorem. If the Proximity Gap cannot derive the Correlated Agreement theorem, it means that the Correlated Agreement theorem is a stronger conclusion than the Proximity Gap theorem. If it can be derived, it means that these two theorems are equivalent.

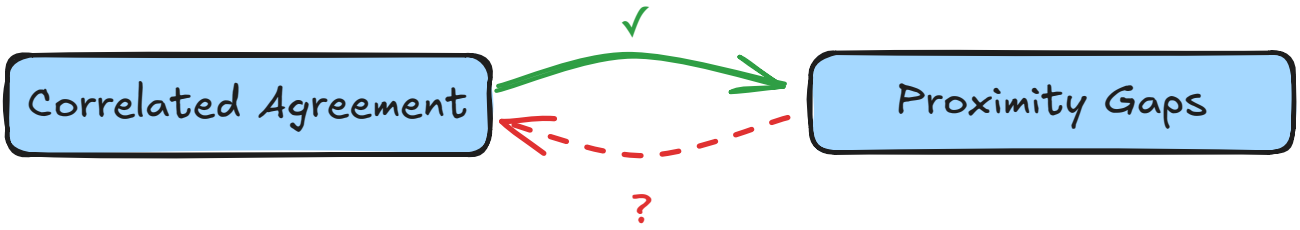

In fact, there are many versions of the Correlated Agreement theorem, taking different A can lead to different theorems, A can be:

1. Lines: $A = \{u_0 + zu_1 : z \in \mathbb{F}\}$
2. Low-degree parameterized curves: $\mathbf{curve(u)} = \left\{u_z := \sum_{i=0}^{m-1} z^i \cdot u_i | z \in \mathbb{F}_q\right\}$
3. Affine space: $u_0 + \mathbf{span}\{u_1, \cdots, u_{m-1}\}$

At the same time, regarding the condition of the Correlated Agreement theorem

$$\Pr_{a \in A}[\Delta(a, V) \leq \delta] > \epsilon, \tag{10}$$

Here we measure the relative Hamming distance between a and V, we can also make this measure more general by adding weights. Give a weight function $\mu : \mathcal{D} \to [0, 1]$, define the relative μ *-agreement* between two vectors u and v as

$$\mathrm{agree}_\mu(u, v) := \frac{1}{|\mathcal{D}|} \sum_{x : u(x) = v(x)} \mu(x) \tag{11}$$

When taking $\mu \equiv 1$,

$$\mathrm{agree}_\mu(u, v) = \frac{1}{|\mathcal{D}|} \sum_{x : u(x) = v(x)} \mu(x) = \frac{1}{|\mathcal{D}|} \sum_{x : u(x) = v(x)} 1 = 1 - \Delta(u, v) \tag{12}$$

The value of this measure is exactly equal to 1 minus the relative Hamming distance. Similarly, define the maximum agreement between a vector u and the code space V as

$$\mathrm{agree}_\mu(u, V) := \max_{v \in V} \mathrm{agree}_\mu(u, v) \tag{13}$$

Changing the condition in the theorem to:

$$\Pr_{a \in A}[\mathrm{agree}_\mu \leq \alpha] > \epsilon, \tag{14}$$

We will get the corresponding Weighted correlated agreement theorem (see [BCIKS20, Section 7]). It can be seen that the Correlated agreement theorem is very flexible. In the paper [BCIKS20, Theorem 8.3], for the soundness proof of the batched FRI protocol, it first defines the required weight function μ, uses the Weighted Correlated Agreement theorem to prove, rather than using the Proximity Gap theorem to prove. And this theorem generally appears in proof by contradiction, it can powerfully help us find the codewords v_i of the code space, and satisfy the properties mentioned in the theorem conclusion, which can help us find contradictions through derivation.

Application of Correlated Agreement Theorem in Soundness

Here's a brief description of the application of the Correlated Agreement theorem in the soundness proof, which is not so rigorous, and the actual security analysis will be more complex.

As mentioned before, the soundness analysis of the FRI protocol is divided into two parts:

1. In the batch stage or Commit stage, due to the improper selection of random numbers, polynomials that were originally far from the code space become closer to the corresponding code space after folding, that is, entering the "shaded area".
2. In the Query stage, due to random checks, the Prover's cheating was not caught.

The Correlated Agreement theorem is mainly applied in the probability analysis of the first part. It will first define the "bad" event $E^{(i)}$: before folding $\Delta^*(f^{(i)}, \mathrm{RS}^{(i)}) > \delta$, split $f^{(i)}$ into $g^{(i+1)}$ and $h^{(i+1)}$, then use a random number $r \in \mathbb{F}$ for folding to get $\mathrm{fold}_r(f^{(i)})$, and the following occurs

$$\Delta(\mathrm{fold}_r(f^{(i)}), \mathrm{RS}^{(i+1)}) \leq \delta \tag{15}$$

Here Δ^* is used, its definition is different from the Hamming distance, it is related to the random query of the Query stage of FRI, which will not be explained in detail here. Assume that the probability of a "bad" event $E^{(i)}$ occurring does not exceed ϵ, that is

$$\Pr[E^{(i)}] = \Pr_{r \in \mathbb{F}}[\Delta(\mathrm{fold}_r(f^{(i)}), \mathrm{RS}^{(i+1)}) \leq \delta] \leq \epsilon \tag{1}$$

If the FRI protocol folds d times, then the probability of some "bad" events occurring does not exceed $d \cdot \epsilon$, that is

$$\bigcup_{i=0}^{d} \Pr[E^{(i)}] \leq d \cdot \epsilon \tag{16}$$

This way, the probability analysis of the first part is done, then assume that these "bad" events do not occur, analyze the probability of the second part, and finally combine the two parts of probability to get the conclusion of soundness.

Now the remaining key problem is how to prove equation (1), that is, to prove that if $\Delta^*(f^{(i)}, \mathrm{RS}^{(i)}) > \delta$, we have

$$\Pr_{r \in \mathbb{F}}[\Delta(\mathrm{fold}_r(f^{(i)}), \mathrm{RS}^{(i+1)}) \leq \delta] \leq \epsilon \tag{2}$$

The idea is to use proof by contradiction, assuming that equation (2) does not hold, that is

$$\Pr_{r\in\mathbb{F}}[\Delta(\text{fold}_r(f^{(i)}), \text{RS}^{(i+1)}) \leq \delta] > \epsilon \tag{17}$$

This satisfies the condition of the Correlated Agreement theorem, which means that there exist $\mathcal{D}' \subset \mathcal{D}^{(i+1)}$, and $\bar{g}^{(i+1)}, \bar{h}^{(i+1)} \in \text{RS}^{(i+1)}$ satisfying

$$\bar{g}^{(i+1)}|_{\mathcal{D}'} = g^{(i+1)}|_{\mathcal{D}'}, \quad \bar{h}^{(i+1)}|_{\mathcal{D}'} = h^{(i+1)}|_{\mathcal{D}'} \tag{18}$$

and $|\mathcal{D}'| \geq (1-\delta)|\mathcal{D}^{(i+1)}|$. With these codewords $\bar{g}^{(i+1)}$ and $\bar{h}^{(i+1)}$ in the code space, we can get a polynomial $\bar{f}^{(i)}$,

$$\bar{f}^{(i)}(x) = \bar{g}^{(i+1)}(x^2) + x \cdot \bar{h}^{(i+1)}(x^2) \tag{19}$$

Due to the linearity of encoding, $\bar{f}^{(i)}$ must also be a codeword, and $\bar{f}^{(i)} \in \text{RS}^{(i)}$, and at the same time we have

$$\bar{f}^{(i)}|_{\mathcal{D}'} = f^{(i)}|_{\mathcal{D}'} \tag{20}$$

Since $|\mathcal{D}'| \geq (1-\delta)|\mathcal{D}^{(i+1)}|$, we can get $\Delta^*(f^{(i)}, \text{RS}^{(i)}) \leq \Delta^*(f^{(i)}, \bar{f}^{(i)}) \leq \delta$, which contradicts the assumption, so equation (2) holds.

Summary

Proximity gap plays a crucial role in the FRI protocol, allowing us to confidently use random numbers to fold polynomials, which greatly reduces the number of oracles sent by the Prover and also reduces the number of queries by the Verifier. In addition, Proximity gap is closely related to the Correlated Agreement theorem and plays a key role in the soundness analysis of FRI.

References

- [BCIKS20] Eli Ben-Sasson, Dan Carmon, Yuval Ishai, Swastik Kopparty, and Shubhangi Saraf. Proximity Gaps for Reed–Solomon Codes. In *Proceedings of the 61st Annual IEEE Symposium on Foundations of Computer Science*, pages 900–909, 2020.
- Video: Proximity Gaps & Applications to Succinct Proofs