Foundations and Frontiers of Probabilistic Proofs (Summer 2021)
Worksheet B.4: Intro to FRI Protocol
Date: 2021.07.30

Problem 1. (Radix-3 FFT) We analyze a slight variation of the well-known Fast Fourier Transform algorithm (which in turn inspires the FRI protocol). Recall that in the Discrete Fourier Transform problem we receive $x_0, \ldots, x_{n-1} \in \mathbb{F}$ and must compute $\widehat{x}_0, \ldots, \widehat{x}_{n-1} \in \mathbb{F}$, where $\widehat{x}_j := \sum_{\ell=0}^{n-1} \omega^{j\ell} \cdot x_\ell$ and ω generates a size-n multiplicative subgroup of $\mathbb{F}^*$. (It is not necessary that $\mathbb{F}$ be finite; indeed, taking $\mathbb{F} = \mathbb{C}$ and $\omega = e^{-2\pi i/n}$ is not only valid but also useful in practice.)

We focus on the case where $n = 3^k$.

1. Show that we rearrange each $\widehat{x}_j$ as $\alpha_j + \omega \cdot \beta_j + \omega^2 \cdot \gamma_j$ where $\alpha_j = \alpha_j(\omega^3), \beta_j = \beta_j(\omega^3), \gamma_j = \gamma_j(\omega^3)$ are (evaluations of) polynomials on ω^3.
2. Show that $\{\alpha_j, \beta_j, \gamma_j\}_{j \in [n]}$ can be computed by Discrete Fourier Transforms of vectors of size $n/3$.
3. Conclude that this yields a recursive algorithm for the DFT problem; give its time complexity (in terms of field operations) and its recursion depth. Why is it useful to have radix-r FFTs for r other than 2?

Problem 2. (Attack on FRI) In this problem we will consider generic attacks on FRI. Let $\mathbb{F}$ be an arbitrary field with some power-of-two multiplicative subgroup L. Let $d \in \mathbb{N}$ a power of 2, $\delta < \frac{1}{2}(1 - d/|L|)$.

1. Give a function $f \colon L \to \mathbb{F}$ that is δ-far from $\mathrm{RS}[\mathbb{F}, L, d]$ and a strategy which convinces the verifier to accept f with probability at least $\max\{1/|\mathbb{F}|, (1-\delta)^t\}$ (recall that t is the number of consistency checks performed by the verifier). *Hint: consider a function that is zero on a large fraction of its domain and linear on the rest.*
2. As above, but convince the verifier with probability at least $\max\{1 - (1 - 1/|\mathbb{F}|)^{\log d}, (1-\delta)^t\}$. *Hint: consider the polynomial* $p_d(X) = \sum_{i=0}^{d-1} \beta^i X^i$*, for some* $\beta \in \mathbb{F}$.
3. **Mini open problem:** Can you do better?

Problem 3. (Simplification of FRI) Consider a modification to the FRI protocol where, instead of performing consistency checks at the end, the verifier performs consistency checks in every round. More precisely, when the verifier sends the i-th field element α_i and the prover replies with a function $f_i \colon L^{2^i} \to \mathbb{F}$, the verifier samples $O(\log d)$ uniformly random points $\mu \in L^{2^{i-1}}$ for each such point and checks that $f_i(\mu^2) = \frac{f_{i-1}(\mu) + f_{i-1}(-\mu)}{2} + \alpha_i \cdot \frac{f_{i-1}(\mu) - f_{i-1}(-\mu)}{2\mu}$, rejecting immediately if any check fails.

1. What is the query complexity of this protocol? (How does it compare to the FRI protocol?)
2. Argue that this protocol has perfect completeness: if $f_0 \colon L \to \mathbb{F}$ is in the Reed–Solomon code $\mathrm{RS}[\mathbb{F}, L, d]$ and the prover is honest, then the verifier always accepts.
3. Argue that this protocol is sound: for every constant $\delta > 0$ there exists a constant $\varepsilon > 0$ such that, if $f_0 \colon L \to \mathbb{F}$ is δ-far from the Reed–Solomon code $\mathrm{RS}[\mathbb{F}, L, d]$, then the verifier accepts with probability at most ε. In the proof of the soundness, you can use the following fact (that you will see in the next lecture): if f is at a sufficiently small distance δ from $\mathrm{RS}[\mathbb{F}, L, d]$, then, with probability $(1 - \frac{\mathsf{L}}{\mathbb{F}})$ over α, $\mathsf{Fold}(f, \alpha)$ is at distance δ from $\mathrm{RS}[\mathbb{F}, L^2, d/2]$.