Comparison of FRI vs. Ligero proof sizes

Justin Thaler

Amongst polynomial commitment schemes that use no cryptographic primitives besides hashing (Merkle-hashing and the Fiat-Shamir transformation), FRI is thought to be the one with the shortest proofs. The reality is that this depends on the degree of the polynomial being committed. For polynomials of degree up to about 2^{18}, FRI proofs are actually bigger than an alternative called Ligero. This holds without basing the security of Ligero on any unproven conjectures about statistical security, and while basing the security of FRI on those conjectures (see Item 7 of the blog post associated with this note).

Degree-2^{18} is admittedly lower than what most projects use today. But some projects plan to use polynomials of this degree in the future to control the substantial memory costs of the FRI prover (which are many GBs).

If one does not base the security of FRI on the conjectures mentioned above, then Ligero proofs remain smaller than FRI proofs until the degree is larger than about 2^{20}.

Amortization. Comparisons are further complicated because most current FRI-based projects apply it to many polynomials simultaneously, amortizing some costs. Although Ligero also has amortization capabilities, they aren't as effective as FRI's. However, Ligero may still be a viable option in these scenarios, as addressed later in this note.

Background on Ligero and FRI. Ligero works over exactly the same fields as FRI and is based on similar techniques (namely, Reed-Solomon encodings, Merkle-hashing, and the Fiat-Shamir transformation). Its proofs consist of $O(\sqrt{n\lambda})$ field elements and a handful of hash evaluations. A clean description of the Ligero polynomial commitment can be found in Section 4.2 of the Brakedown paper, with security analysis given in Appendix B; see also recent work of Diamond and Posen for an excellent exposition of additional optimizations not described in the Brakedown paper.

My guess is that Ligero is not yet popular because people see the $\sqrt{n}$ term and assume its proofs are too big to be useful (the initial implementation was also rather unoptimized, which may have given an inaccurate impression of performance). However, while the dependence on the degree bound n of Ligero is asymptotically worse than FRI ($O(\sqrt{n})$ vs. $O(\log(n)^2)$), Ligero's dependence on the security parameter λ is better. Moreover, $\sqrt{n}$ looks a lot like $\log(n)^2$ for small values of n. In fact, $\sqrt{n}$ is smaller than $log(n)^2$ when n is less than 2^{16}.

Ligero has additional simplicity and performance benefits. For example, its prover naturally does about $O(\sqrt{n/\lambda})$ FFTs of size $O(\sqrt{n\lambda})$, which results in simpler and better parallelization than the single FFT of length $O(n)$ in FRI. Ligero's prover also does less hashing than FRI's.

Quantitative comparison. FRI proofs at λ bits of security (under the aforementioned conjectures that known attacks are exactly optimal) consist of roughly $\lambda \log(n/\rho)^2/(2\log(1/\rho))$ hash values, where $1/\rho$ is the so-called "FRI blowup factor" that controls the tradeoff between prover time and proof size. Setting the blowup factor to 1/4 and λ to 128, this translates to a proof size of about 400KBs for degree $n = 2^{18}$.

This is an overestimate: Some savings are possible due to techniques such as Merkle capping, which shortens the length of Merkle authentication paths at the cost of increasing commitment size. I estimate that these techniques can shave about 33% off the proof size, yielding an estimate of about 270 KiBs for degree $n = 2^{18}$. Without the conjectures mentioned in Item 7 of , the proofs would be more than twice as large.

Ligero proofs (without analogous soundness conjectures) consist of about $2\sqrt{n\lambda/\log(2/(1+\rho))}$ field elements and a small number of hash values. Many FRI-based projects today work over a 128-bit field, although this is slightly too small to provide 128 bits of security without conjectures. Working over a 128-bit field, and setting the blowup factor to 1/4 and λ to 128 as above, translates to proofs of 225 KiBs for the same degree, $n = 2^{18}$.

The qualitative comparison I am making here has been corroborated in recent work. For example, the top of page 10 of this survey on FRI reports that for 128 bits of security (note that this is without invoking unproven conjectures about statistical security beyond the Johnson bound), and a polynomial of degree only 2^{12}, FRI proofs using a blowup factor of 8 are well over 300 KiBs. At the same security level, Diamond and Posen report a Ligero proof size of about 270 KiBs for a polynomial of degree 2^{16} (i.e., sixteen times bigger than 2^{12}), and using a blowup factor of 4, which implies at least a 2× faster prover than using a blowup factor of 8. Qualitatively similar findings are also mentioned in the related work section of the recent extended version of Ligero itself.

Deployment in amortized settings. In most deployments today, FRI is used to commit to some number k of polynomials, where k is between about 50 and several hundred (see for example section 5 of this writeup on RISC Zero, where the number of "columns" corresponds to the number of committed polynomials). Both FRI and Ligero have non-trivial amortization, whereby their proof sizes are much less than k times bigger than the non-amortized case (i.e., k=1).

I estimate FRI proofs in these settings (under aggressive security conjectures) to consist of roughly

$$3\lambda \log(n/\rho)/\log(1/\rho) + \lambda \log(n/\rho)^2/(2\log(1/\rho))$$

hash values plus $k\lambda/\log(1/\rho)$ field elements, with some savings possible due to techniques such as Merkle-capping. Ligero proofs are dominated by $2\sqrt{kn\lambda/\log(2/(1+\rho))}$ field elements.

The comparison in this amortized setting shifts in the favor of FRI, but Ligero remains competitive. For $\lambda = 128$, $k = 60$, a blowup factor of 4, a field size of 128 bits, and degree equal to $n = 2^{18}$, FRI proofs are about 415 KiBs under aggressive security conjectures, and 830 KiBs without the conjectures. For Ligero, it's about 1.75 MiBs with no conjectures.

Most deployments of FRI today use SNARK recursion. Hence, it's not essential that the proof size be as small as possible. What is important is that the FRI proofs are small enough, and verification fast enough, that recursion is not the prover time bottleneck. This may render Ligero preferable to FRI even in amortized settings, as it has a faster prover and its proof size is within a factor 2-4 of FRI's for degree $n = 2^{18}$.

Grinding. The proof lengths of the Ligero and FRI commitment schemes can both benefit modestly from a technique called "grinding," which makes known attacks on the Fiat-Shamir transformation more expensive without increasing proof size (see my previous blog post for details). Confusingly, "grinding technique" in this context refers to a defense against the "grinding attack" described in Item 13. FRI benefits more from this grinding technique than Ligero, owing to FRI's worse dependence on the security parameter λ. Some but not all FRI deployments today use grinding.

Summary. Projects applying FRI to polynomials of degree 2^{18} or less should consider switching to Ligero for improved performance and avoidance of strong security conjectures.