

Average-Case Sparse Fourier Transform Algorithms

Piotr Indyk

MIT


Ghazi, Hassanieh, Indyk, Katabi, Price, Lixin, “Sample-Optimal Average-Case Sparse Fourier Transform in 2D”

Fourier Transform

- Discrete Fourier Transform:
 - Given: a signal $x[1\dots n]$
 - Goal: compute the frequency vector x^{\wedge} where

$$x^{\wedge}_f = \sum_t x_t e^{-2\pi i t f/n}$$

- Sparse Fourier Transform
 - Only “few” “large” coefficients
 - GL89, KM90, Mansour’92, GGIMS’02, AGS’03, GMS’05, Iwen’10, Akavia’10, HIKP’12, BCGLS’12, LWC’12...



— Sampled Audio Data (Time)
— DFT of Audio Samples (Frequency)

SFFT

[Hassanieh, Indyk, Katabi, Price'12]

- All algorithms randomized, with constant probability of success, n is a power of 2
- Exactly k -sparse case : $O(k \log n)$
(SFFT 3.0)
- Approximately k -sparse case
 - Let $\text{Err}_2^k(x^\wedge) = \min_{k\text{-sparse } z^\wedge} \|x^\wedge - z^\wedge\|_2$
 - l_2/l_2 guarantee $\|x^\wedge - y^\wedge\|_2 \leq C \text{Err}_2^k(x^\wedge)$: $O(k \log(n) \log(n/k))$
time
(SFFT 4.0)
 - Weaker results for l_∞/l_2 guarantee
(SFFT 1.0 and SFFT 2.0)

Sample complexity?

Algorithm	Time	Samples	Lower bound
SFFT 3.0 (exact)	$O(k \log n)$	$O(k \log n)$	k
SFFT 4.0 (robust)	$O(k \log(n) \log(n/k))$	$O(k \log(n) \log(n/k))$	$k \log(n/k)$


How Does Sparse FFT Work?

1- Bucketize



Divide spectrum into a few buckets

2- Estimate



Estimate the large coefficient (position and value) in each non-empty bucket



Bucketization




- For $F = \text{Sinc} \times \text{Gaussian}$, $|\text{supp}^{>1/n}(F)| * |\text{supp}^{>1/n}(F^\wedge)| = n \log n$
 - This leads to $O(k \log n)$ sample complexity



- Spike train is better (optimal): $|\text{supp}(F)| * |\text{supp}(F^\wedge)| = n$ (see Mark Iwen's talk)
- However, can't hash multiple times
 - Frequencies i and j fall into the same bucket iff $i=j \pmod B$
 - Invariant under affine transform
- But what if the coefficients are distributed at **random**? E.g., k -non-zeros in random places
 - Lawlor-Wang-Christlieb'12: $O(k)$ samples, but sampling at arbitrary points

Bucketization using spike train \approx aliasing

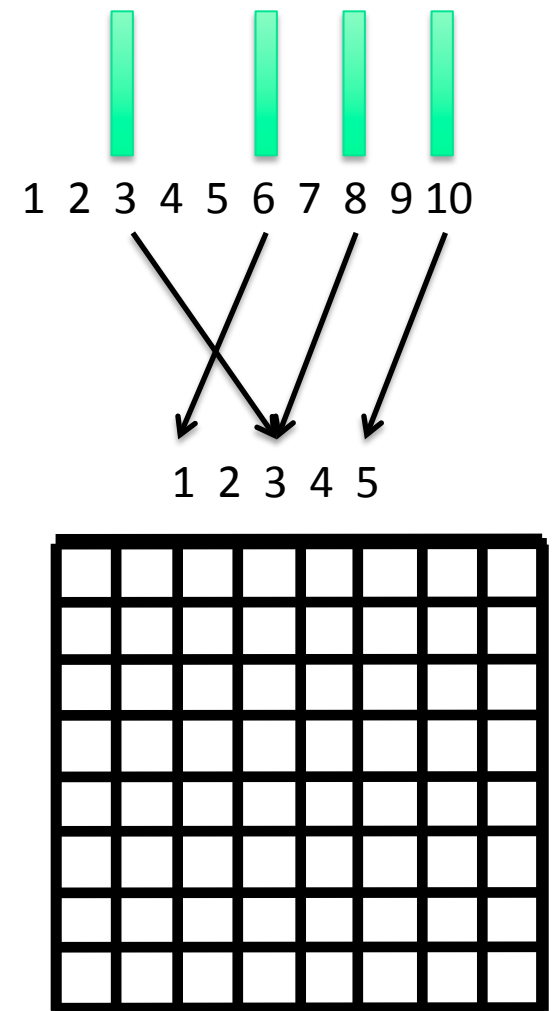
- Frequencies i and j fall into the same bucket iff $i=j \pmod B$




- What happens if we alias random frequencies ?

Spectral balls and bins


- k balls (coefficients), B buckets/bins
 - Need $B > k^2$ to achieve isolation
 - Sample complexity $> k^2$, i.e., bad
- Issue: we need more than one way of aliasing frequencies
- This is possible if, e.g., we have
 - **2D** transform over an $n \times n$ grid
 - Suppose that $k \approx n$
- Then we can define buckets to be either single columns or single rows



Alternating rows/columns

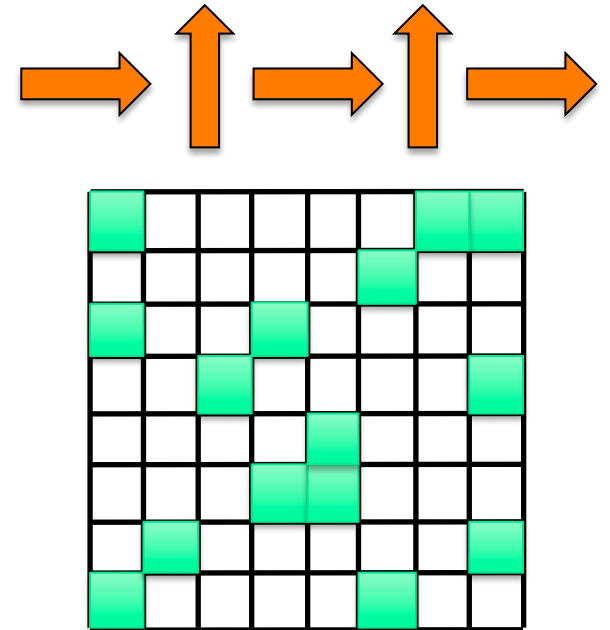


n x n grid



Analysis

- Analysis
 - With good probability (over the input) the process converges after $O(\log n)$ steps for $k \approx n$
 - Each step requires $O(n) = O(k)$ samples
 - But we can use **the same samples** in all steps (no rehashing)
 - So $O(k)$ samples total
 - Time: $O(k \log n)$
- Generalizes to
 - $k < n$: $O(k)$ samples
 - l_2/l_2 recovery (Gaussian noise): $O(k \log n)$ samples



Conclusions

- Sparse FFT with running times $O(k \log n)$ or $O(k \log(n) \log(n/k))$
 - Improves over FFT for $k \ll n$
- Can improve sample complexity to $O(k)$ or $O(k \log n)$ in **2D average case**
 - Bonus: the algorithm is really simple
- Questions:
 - $k \log n$ time for **approximately** sparse signals?
 - Not clear: $k \log(n/k)$ samples needed, extra $\log n$ for FT
 - Better sample complexity in the **worst case**?
 - Deterministic ? (best known runtime $> k^2$)
 - Model-based (coefficients cluster in blocks) ?
 - ...

Empirical evaluation: SFFT 3.0 (exact sparsity)

