

SWIFFT: A Modest Proposal For FFT Hashing

Vadim Lyubashevsky

Daniele Micciancio

Chris Peikert

Alon Rosen

UCSD

UCSD

SRI International

Herzliya IDC


The image shows the title page of Jonathan Swift's satirical work, 'A Modest Proposal'. The text is centered on a light-colored, aged paper background. At the top, there is a decorative flourish. The title 'A MODEST PROPOSAL' is written in a large, bold, serif font. Below it, the author's name 'Jonathan Swift' is written in a slightly smaller, elegant serif font. The subtitle, 'FOR PREVENTING THE CHILDREN OF POOR PEOPLE IN IRELAND FROM BEING A BURDEN TO THEIR PARENTS OR COUNTRY, AND FOR MAKING THEM BENEFICIAL TO THE PUBLICK', is in a smaller, all-caps serif font. At the bottom, the location 'DUBLIN, IRELAND' is written in an italicized serif font. There are two more decorative flourishes, one above and one below the subtitle.

A MODEST PROPOSAL

by
Jonathan Swift

FOR PREVENTING THE CHILDREN OF POOR PEOPLE IN IRELAND
FROM BEING A BURDEN TO THEIR PARENTS OR COUNTRY, AND
FOR MAKING THEM BENEFICIAL TO THE PUBLICK

DUBLIN, IRELAND



SWIFFT

A collection of compression functions

- Efficient
 - Highly parallelizable
 - Supporting proof of security
- } FFT

Not an “all-purpose” function

- In particular, it is linear: $f(x+y) = f(x) + f(y)$
- However, has many desirable properties
 - a) Cryptographic
 - b) Statistical

Our Starting Point

At a (very) high-level:

- ❑ Key: m random $\deg < n$ polynomials in α
- ❑ Input: m polynomials w/ binary 0-1 coefficients
- ❑ Function: compute sum of products

All arithmetic modulo p and $(\alpha^n + 1)$

$$R = \mathbb{Z}_p[\alpha] / (\alpha^n + 1)$$

- ❑ Key: $\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_m \in R$
- ❑ Input: $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_m \in \{0, 1\}^n \subset R$
- ❑ Function: $f_A(\mathbf{X}) = \sum_{i=1}^m \mathbf{a}_i \cdot \mathbf{x}_i \in R$

Supporting Proof of Security

For random key \mathbf{A} , the function is collision resistant, assuming worst-case hardness in cyclic/ideal lattices [PR06, LM06].

- ❑ Continues a long line of works [Ajtai96,...,Mic02]
- ❑ proof is asymptotic
- ❑ meaningful only for large parameters

In this work:

- ❑ Concrete parameters ($m=16, n=64, p=257$)
- ❑ Function maps 1024 bits to 528 bits
- ❑ Very efficient implementation.

- ❑ Security proof suggests that design is sound
- ❑ Heuristic analysis suggests that parameters are sound

Towards Efficient Implementation

Central Observations:

1. Polynomial multiplication \Leftrightarrow FFT

$$\mathbf{a}_i \cdot \mathbf{x}_i = \text{FFT}^{-1}(\text{FFT}(\mathbf{a}_i) \odot \text{FFT}(\mathbf{x}_i))$$

2. Can pre-compute $\text{FFT}(\mathbf{a}_1), \text{FFT}(\mathbf{a}_2), \dots, \text{FFT}(\mathbf{a}_m)$
3. No need to compute FFT^{-1}
4. Specifics of $\mathbb{Z}_p[\alpha]/(\alpha^n + 1)$ allow FFT optimization
 - a) Can perform *modular* FFT (NTT)
 - b) FFT of dimension n is sufficient

Resulting function is completely equivalent (security-wise).

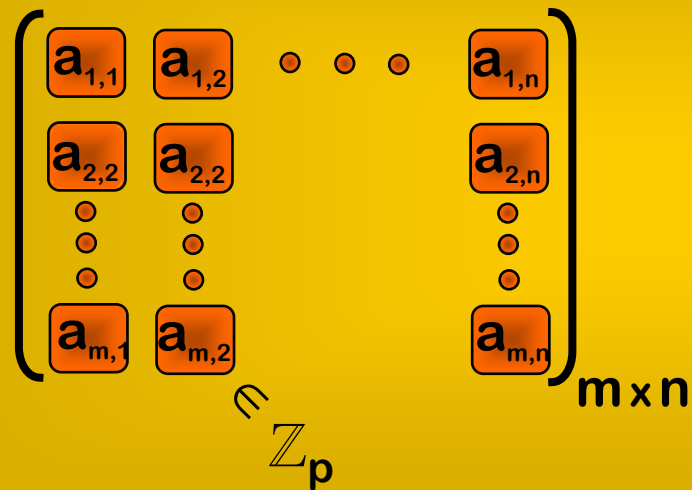
SWIFFT

Parameters: n, m, p

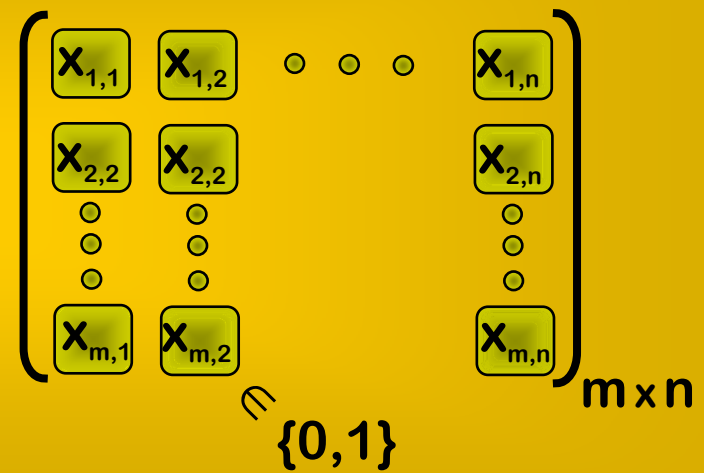
Key: $m \times n$ matrix $(a_{i,j}) \in \mathbb{Z}_p^{m \times n}$

Input: $m \times n$ binary matrix $(x_{i,j}) \in \{0,1\}^{m \times n}$

Key



Input

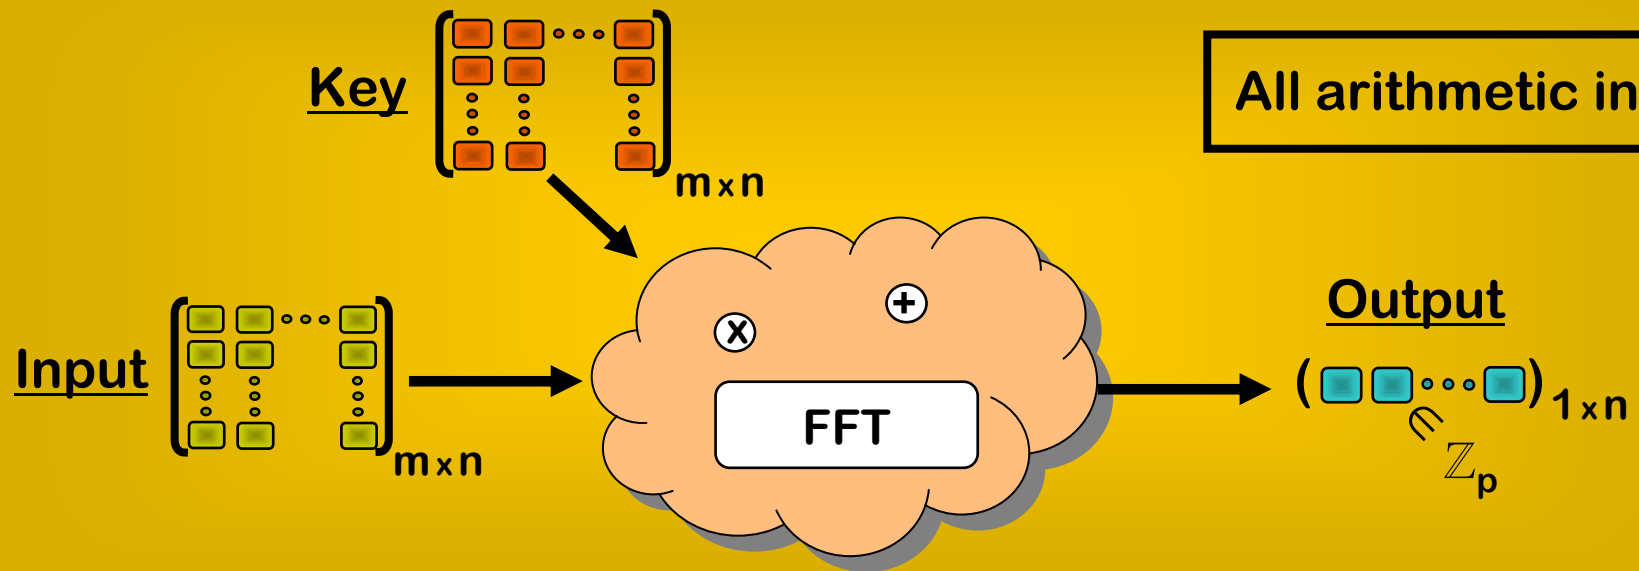


SWIFFT


Parameters: n, m, p

Key: $m \times n$ matrix $(a_{i,j}) \in \mathbb{Z}_p^{m \times n}$

Input: $m \times n$ binary matrix $(x_{i,j}) \in \{0,1\}^{m \times n}$



SWIFFT (m=3)

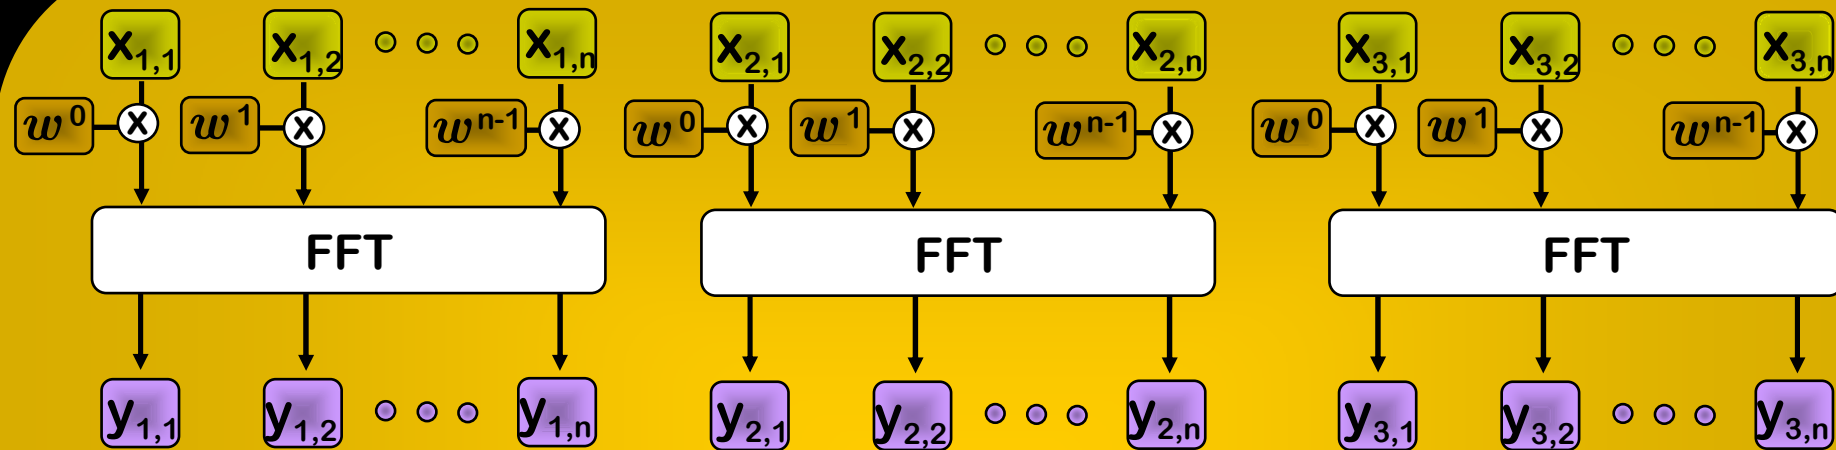


Step 1: For each row $i = 1, \dots, m$ compute:

$$(y_{i,1}, \dots, y_{i,n}) = \text{FFT}(\omega^0 \cdot x_{i,1}, \dots, \omega^{n-1} \cdot x_{i,n})$$

where $\omega \in \mathbb{Z}_p$ is a $2n^{\text{th}}$ root of unity in \mathbb{Z}_p

SWIFFT (m=3)

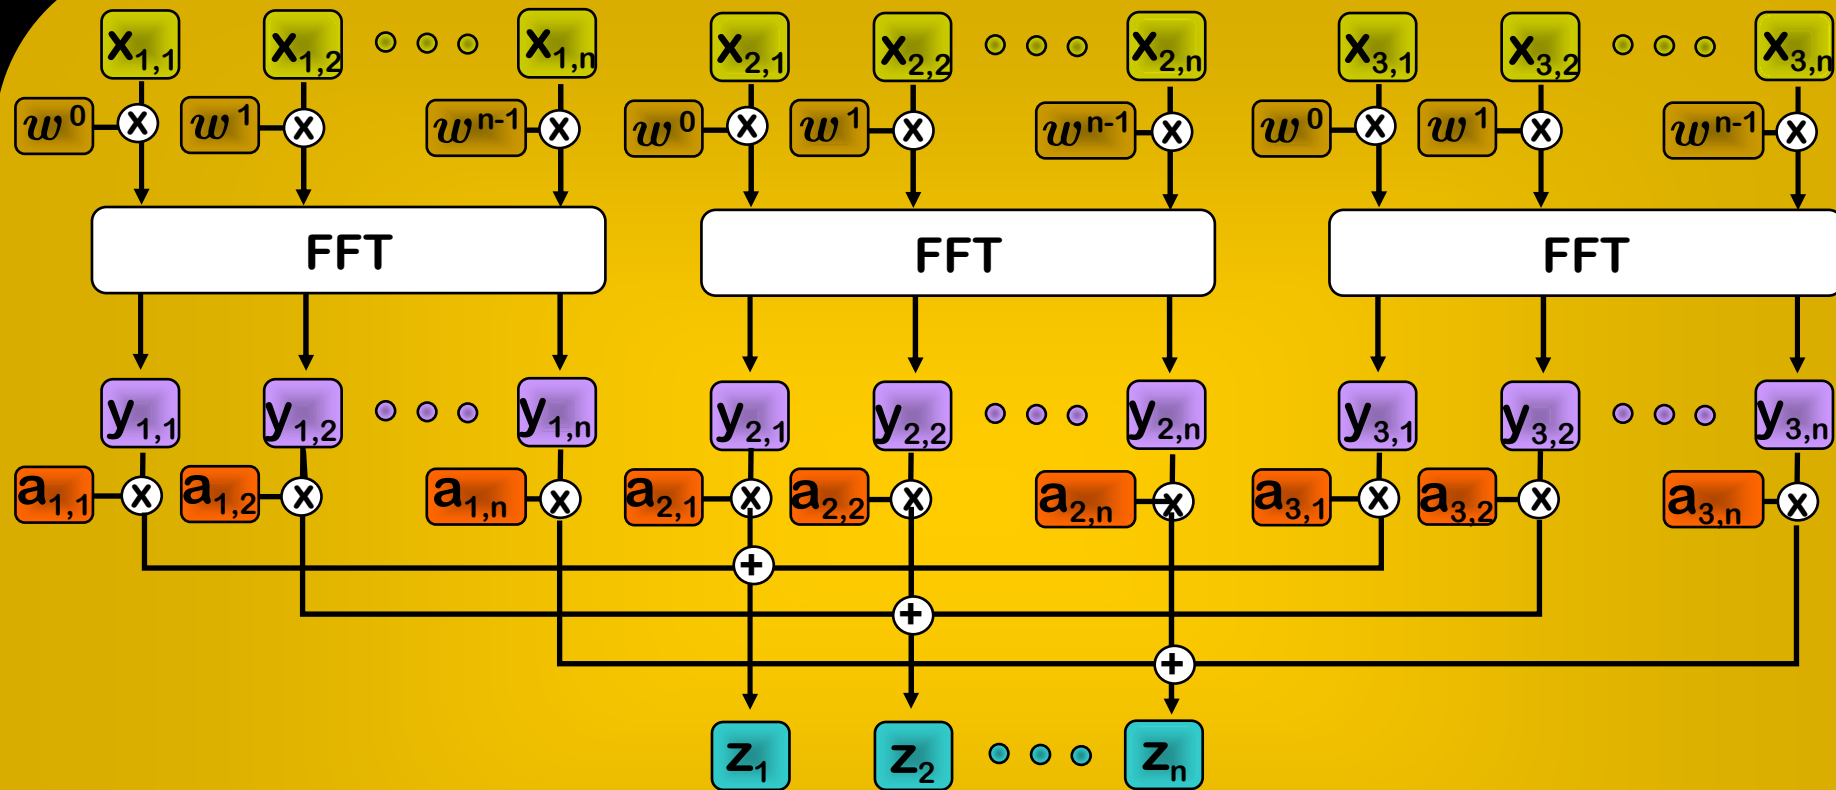


Step 1: For each row $i = 1, \dots, m$ compute:

$$(y_{i,1}, \dots, y_{i,n}) = \text{FFT}(\omega^0 \cdot x_{i,1}, \dots, \omega^{n-1} \cdot x_{i,n})$$

where $\omega \in \mathbb{Z}_p$ is a $2n^{\text{th}}$ root of unity in \mathbb{Z}_p

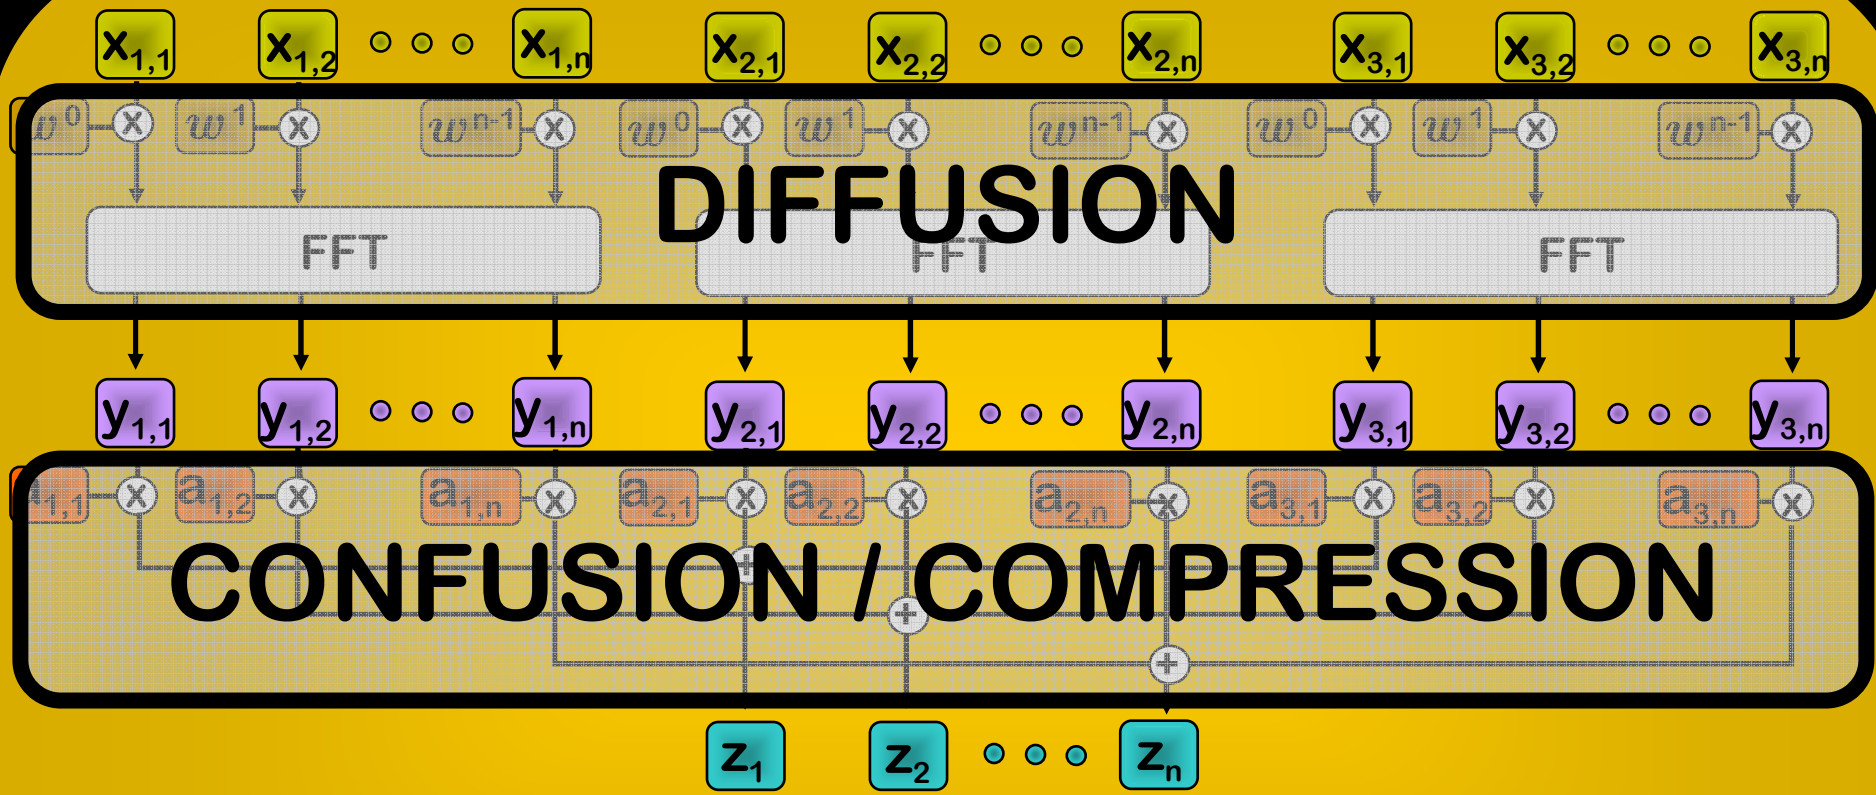
SWIFFT (m=3)



Step 2: For each column $j = 1, \dots, n$ compute:


$$z_j = a_{1,j} \cdot y_{1,j} + \dots + a_{m,j} \cdot y_{m,j}$$

SWIFFT (m=3)



SWIFFT ($m = 16, n = 64, p = 257$)

$nm = 1024$ bits
Hard to solve z_1, \dots, z_n simultaneously



$n \log_2 p \sim 528$ bits

Easy to find solution $(y_{1,j}, \dots, y_{m,j})$ to each $z_j = \sum_{i=1}^m a_{i,j} \cdot y_{i,j}$ individually

□ The reason: $(y_{i,1}, \dots, y_{i,n})$ are highly constrained

1. Dependency through FFT
2. Need to find binary $(x_{i,1}, \dots, x_{i,n})$

□ This way of “breaking linearity” is different from previous proposals for FFT hashing [S91, S92, SV93, V92].

Choice of Parameters (m=16, n=64, p=257)


Security considerations:

- ❑ Subset-sum instance from 1024 to 528 bits.
- ❑ $n = 2^k \Leftrightarrow$ modulus polynomial $(\alpha^n + 1)$ is irreducible (over \mathbb{Q})
 1. crucial for security proof
 2. otherwise can find collisions (LASH, [Mic02] OWF).
 3. Enables to avoid straightforward weaknesses

Performance considerations:

- ❑ p=257 is a prime of the form $p = 4n+1$
- ❑ enables efficient 64-dimensional modular FFT.
 1. \mathbb{Z}_{257} is a field
 2. $\omega \in \mathbb{Z}_{257}$ is a 128th root of unity
 3. Odd powers of w are the roots of $(\alpha^n + 1)$

Fast Implementation



To improve performance:

1. Use lookup tables inside FFT.
2. Parallelize atomic operations (+,x).
3. Avoid modular reductions (whenever possible).
4. Multiply by w powers using left shifts.


Speeding up the FFT

Input to FFT is a binary vector:

- ❑ Few possible intermediate values.
- ❑ Can pre-compute and store in lookup table.

FFT is highly parallelizable:


- ❑ Reduce FFT_{64} to 8 parallel FFT_8
- ❑ Use SIMD (single-instruction multiple-data) instructions to perform operations in parallel.
- ❑ Point-wise vector addition/multiplication on 8 dim. registers (w/ 16 bit signed integer entries).



Further Optimizations

Use of \mathbb{Z}_{257} :

- ❑ $\omega = 42$ is a 128th root of unity mod 257.
- ❑ FFT_8 uses $\omega^{16} = 2^2 \pmod{257}$.
- ❑ Multiplications by powers of ω^{16} using left shifts.
- ❑ Can avoid most modular reductions w/out overflow.
- ❑ Use SIMD for parallel modular reduction.



Multi-core processors:

- ❑ FFTs are completely independent.
- ❑ We did not exploit multi-core capabilities yet



Performance

Implemented and tested:

- ❑ On 3.2 GHz Intel Pentium 4.
- ❑ Written in C (using INTEL intrinsics for SSE2).
- ❑ Compiled using gcc 4.1.2 on Linux kernel 2.6.19.

Compared to SHA256:

- ❑ Same system
- ❑ `openssl` version 0.9.8 speed benchmark

Results:

- ❑ **SWIFFT** - Throughput ~40 MB/s
- ❑ **SHA256** - Throughput ~47MB/s

Statistical/Cryptographic Properties

Statistical properties (no computational assumptions):

1. Universal hashing
2. Regularity
3. Randomness extraction

Cryptographic properties:

1. One-wayness
2. Second-preimage resistance
3. Target collision resistance
4. Collision resistance



We aim for 2^{100} security. We do NOT claim:

- 2^{528} security against inversion attacks
- $2^{528/2}$ security against collision attacks

Concrete Security Analysis

A convenient way to view SWIFFT is as a subset-sum instance:

$$\mathbf{a} \cdot \mathbf{x} \in R \iff \begin{bmatrix} a_0 & -a_{n-1} & \cdots & -a_1 \\ a_1 & a_0 & & -a_2 \\ \vdots & & \ddots & \\ a_{n-1} & & \cdots & a_0 \end{bmatrix} \begin{bmatrix} x_0 \\ x_1 \\ \vdots \\ x_{n-1} \end{bmatrix} \pmod p$$

Our function can be viewed as multiplying a vector $\mathbf{x} \in \{0,1\}^{mn}$ with a matrix $\mathbf{A} = [\mathbf{A}_1 \mid \cdots \mid \mathbf{A}_m]_{n \times mn}$ where \mathbf{A}_i is the skew-circulant matrix that corresponds to \mathbf{a}_i

- ❑ Best known attack:
 - Wagner's generalized birthday attack.
 - Has complexity 2^{106}
- ❑ Lattice reduction algorithms do not do as well.

Conclusions

SWIFFT: FFT-based hashing

- Provably secure design (under worst-case assumption)
- Concrete instantiation w/ heuristic security analysis
- Highly efficient implementation

Future directions

- Further cryptanalysis (possibly algebraic)
- Faster implementation
- Shorter output/smaller description
- Exploiting linearity for applications