



MAX PLANCK INSTITUTE
FOR SECURITY AND PRIVACY

Formosa Crypto – high-assurance crypto software in practice

Peter Schwabe

February 20, 2024

Cryptographic software

- Primitives, no protocols
- “Secure-channel” primitives

Cryptographic software

- Primitives, no protocols
- “Secure-channel” primitives
- Only software-visible side channels

Cryptographic software

- Primitives, no protocols
- “Secure-channel” primitives
- Only software-visible side channels
- Big CPUs

Back in the days. . .

- Use X25519, Ed25519
- Use SHA2, ChaCha20, Poly1305

Back in the days. . .

- Use X25519, Ed25519 (or NISTP256-ECDH, ECDSA)
- Use SHA2, ChaCha20, Poly1305 (or AES, HMAC)

- Use X25519, Ed25519 (or NISTP256-ECDH, ECDSA)
- Use SHA2, ChaCha20, Poly1305 (or AES, HMAC)
- Follow “constant-time” paradigm
 - No secret-dependent branches
 - No memory access at secret-dependent location
 - No variable-time arithmetic (e.g., DIV)

- Use X25519, Ed25519 (or NISTP256-ECDH, ECDSA)
- Use SHA2, ChaCha20, Poly1305 (or AES, HMAC)
- Follow “constant-time” paradigm
 - No secret-dependent branches
 - No memory access at secret-dependent location
 - No variable-time arithmetic (e.g., DIV)
- Fairly little code, doesn't even need function calls!

- More assumptions, more schemes, more parameters, **more code**
- More complexity in implementations, protocols, and proofs

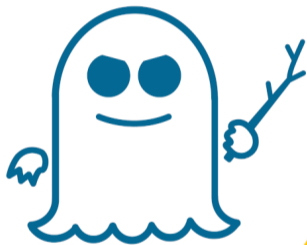
- More assumptions, more schemes, more parameters, **more code**
- More complexity in implementations, protocols, and proofs
- Initially many bugs that were not caught by functional testing
- Early personal intuition:
 - no big-integer arithmetic \rightarrow no “rare” bugs
 - Confidence in functional correctness through test vectors ...?

- More assumptions, more schemes, more parameters, **more code**
- More complexity in implementations, protocols, and proofs
- Initially many bugs that were not caught by functional testing
- Early personal intuition:
 - no big-integer arithmetic → no “rare” bugs
 - Confidence in functional correctness through test vectors ... ?
- Shattered by Hwang, Liu, Seiler, Shi, Tsai, Wang, and Yang (CHES 2022): *Verified NTT Multiplications for NISTPQC KEM Lattice Finalists: Kyber, SABER, and NTRU.*

Advanced microarchitectural side channels



MELTDOWN



Hertzbleed



CACHE OUT

Tools that aren't built for crypto

“... implementations shall consist of source code written in ANSI C.”

—NIST PQC Call for Proposals, 2017

- No memory safety
- Finicky semantics
 - Undefined behavior
 - Implementation-specific behavior
 - Context-specific behavior
- No mandatory initialization
- No (optional) runtime checks

Tools that aren't built for crypto

“... implementations shall consist of source code written in ANSI C.”

—NIST PQC Call for Proposals, 2017

- No memory safety
- Finicky semantics
 - Undefined behavior
 - Implementation-specific behavior
 - Context-specific behavior
- No mandatory initialization
- No (optional) runtime checks

but... Rust!

- Memory safe
- More clear semantics (?)
- Mandatory variable initialization
- (Optional) runtime checks for, e.g., overflows

Lack of security features

- **No concept of secret vs. public data**
- No preservation of “constant-time”
- Limited protection against microarchitectural attacks
- Limited support for erasure of sensitive data

“We argue that we must stop fighting the compiler, and instead make it our ally.”

—Simon, Chisnall, Anderson, 2018

Secure erasure in LLVM

- Simon, Chisnall, Anderson implement secure erasure in LLVM
- Code available at <https://github.com/lmrs2/zerostack>
- **Not adopted in mainline LLVM**

Secret types in Rust + LLVM

- Initiative at HACS 2020: secret integer types in Rust, C++, **and LLVM**
- Rust draft RFC online at <https://github.com/rust-lang/rfcs/pull/2859>
- Implementation in LLVM is massive effort, **no real progress, yet**

Spectre protections in LLVM

- Carruth, 2019: Spectre v1 countermeasure in LLVM¹ (see later in the talk)
- *“does not defend against secret data already loaded from memory and residing in registers”*

¹<https://llvm.org/docs/SpeculativeLoadHardening.html>

²*Ultimate SLH: Taking Speculative Load Hardening to the Next Level.* USENIX Security, 2023

Spectre protections in LLVM

- Carruth, 2019: Spectre v1 countermeasure in LLVM¹ (see later in the talk)
- *“does not defend against secret data already loaded from memory and residing in registers”*
- Zhang, Barthe, Chuengsatiansup, Schwabe, Yarom, 2023: More principled approach²
- Report and proposed patches to LLVM in March 2022
- September 2022: **Status: WontFix (was: New)**

¹<https://llvm.org/docs/SpeculativeLoadHardening.html>

²*Ultimate SLH: Taking Speculative Load Hardening to the Next Level.* USENIX Security, 2023



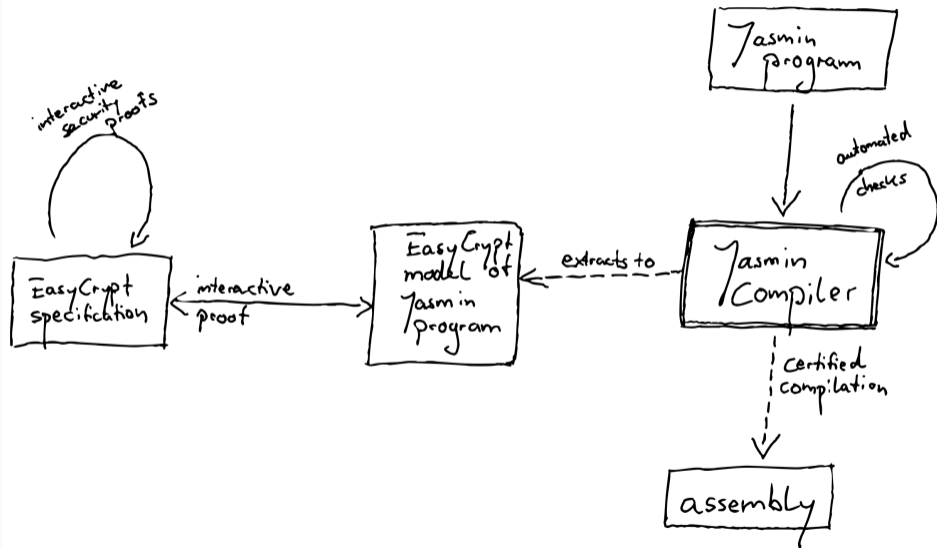
FORMOSA CRYPTO

- Effort to formally verify crypto
- Goal: **verified PQC ready for deployment**
- Three main projects:
 - EasyCrypt proof assistant
 - Jasmin programming language
 - Libjade (PQ-)crypto library
- Core community of \approx 30–40 people
- Discussion forum with $>$ 200 people

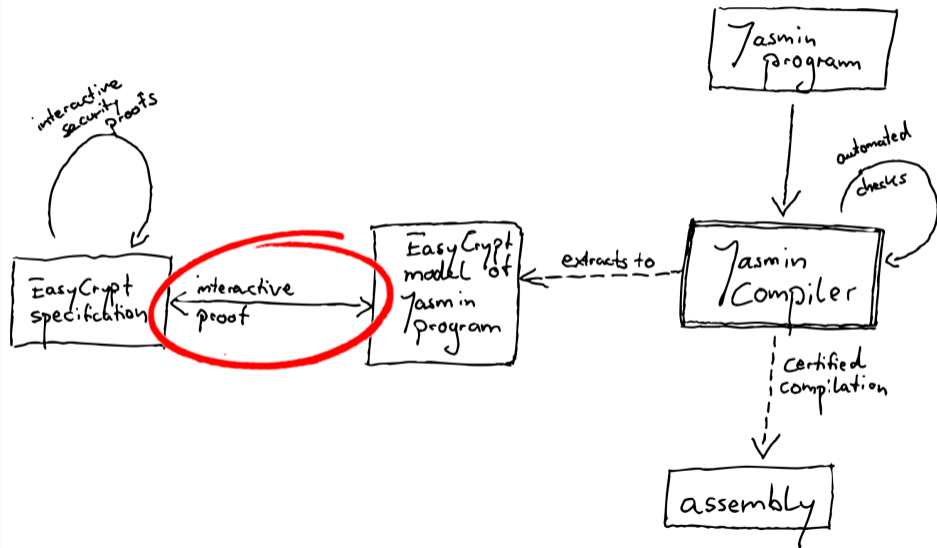


Aaron Kaiser, Adrien Koutsos, Alley Stoughton, Amber Sprenkels, Andreas Hülsing, Antoine Séré, Basavesh Ammanaghatta Shivakumar, **Benjamin Grégoire**, Benjamin Lipp, Bo-Yin Yang, Bow-Yaw Wang, Chitchanok Chuengsatiansup, Christian Doczkal, Daniel Genkin, Denis Firsov, Fabio Campos, François Dupressoir, Gilles Barthe, Hugo Pacheco, Jack Barnes, **Jean-Christophe Léchenet**, José Bacelar Almeida, Kai-Chun Ning, Lionel Blatter, Lucas Tabary-Maujean, Manuel Barbosa, Matthias Meijers, Miguel Quaresma, Ming-Hsien Tsai, Peter Schwabe, Pierre Boutry, Pierre-Yves Strub, Ruben Gonzalez, Rui Qi Sim, Sabrina Manickam, **Santiago Arranz Olmos**, Sioli O'Connell, Sunjay Cauligi, Swarn Priya, Tiago Oliveira, Vincent Hwang, **Vincent Laporte**, William Wang, Yi Lee, Yuval Yarom, Zhiyuan Zhang

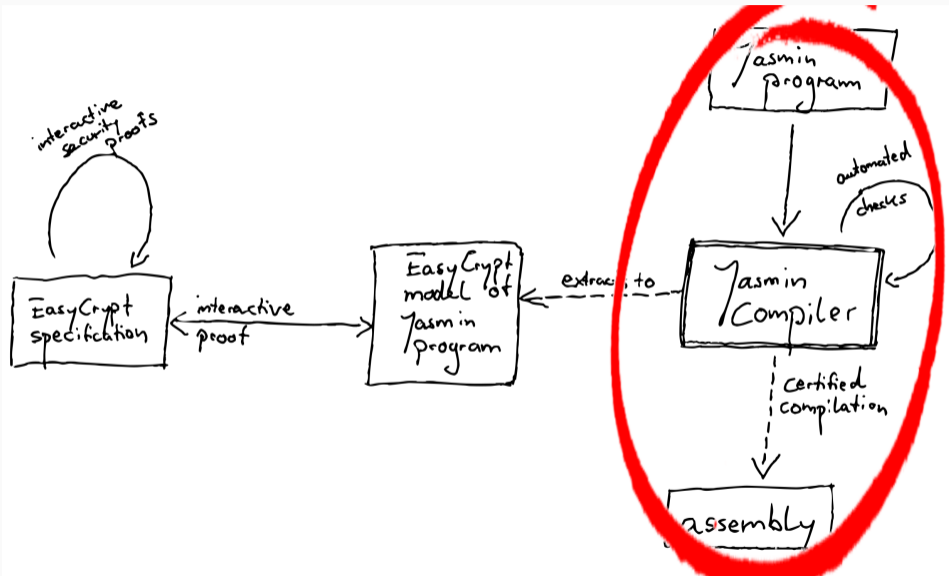
The toolchain and workflow



The toolchain and workflow



The toolchain and workflow

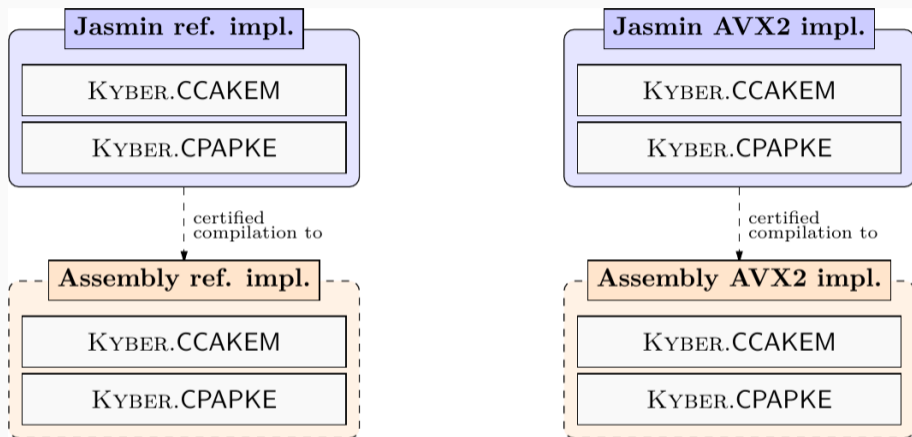


“The public-key encryption and key-establishment algorithm that will be standardized is CRYSTALS-KYBER.”

—NIST IR 8413-upd1

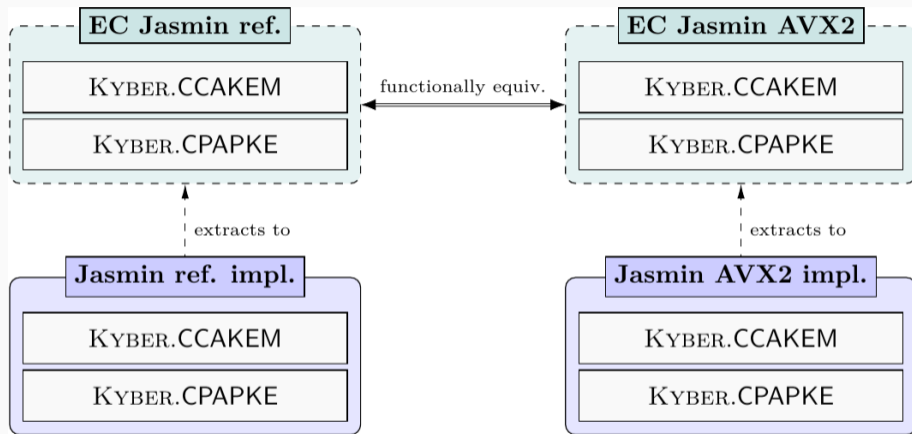
- Lattice-based KEM, joint work with Avanzi, Bos, Ding, Ducas, Kiltz, Lepoint, Lyubashevsky, Schanck, Schwabe, Seiler, and Stehlé.
- Three parameter sets; “recommended” is **Kyber768**
- FIPS draft standard public for comments:
<https://csrc.nist.gov/pubs/fips/203/ipd>
- Already deployed in TLS by Google and Cloudflare

Functional correctness of Kyber implementations



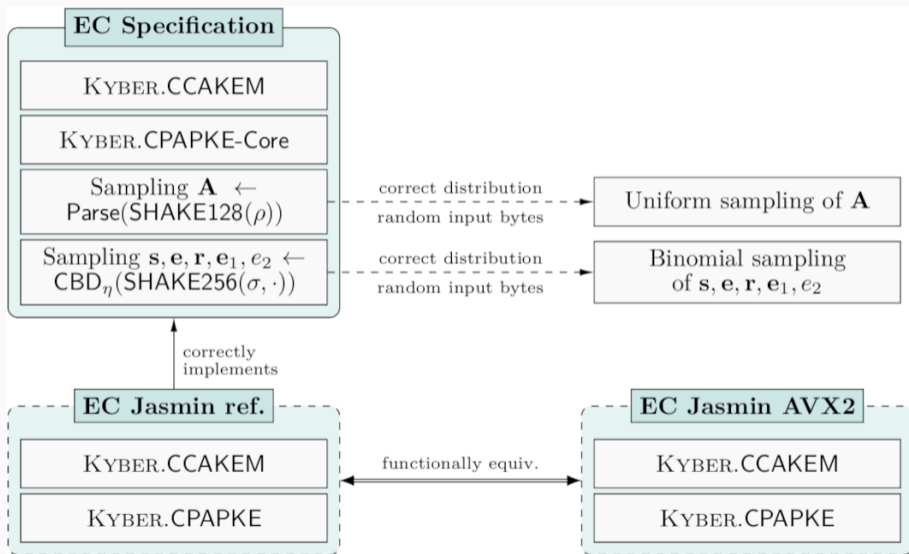
Almeida, Barbosa, Barthe, Grégoire, Laporte, Léchenet, Oliveira, Pacheco, Quesma, Schwabe, Séré, and Strub. *Formally verifying Kyber – Episode IV: Implementation Correctness*. TCHES 2023-3.

Functional correctness of Kyber implementations



Almeida, Barbosa, Barthe, Grégoire, Laporte, Léchenet, Oliveira, Pacheco, Quesma, Schwabe, Séré, and Strub. *Formally verifying Kyber – Episode IV: Implementation Correctness*. TCHES 2023-3.

Functional correctness of Kyber implementations



Almeida, Barbosa, Barthe, Blot, Grégoire, Laporte, Oliveira, Pacheco, Schmidt, Strub. *Jasmin: High-Assurance and High-Speed Cryptography*. ACM CCS 2017

- Language with “C-like” syntax
- Programming in Jasmin is much closer to assembly:
 - Generally: 1 line in Jasmin → 1 line in assembly
 - A few exceptions, but highly predictable
 - Compiler does not schedule code
 - Compiler does not spill registers

³Barthe, Grégoire, Laporte, and Priya. *Structured Leakage and Applications to Cryptographic Constant-Time and Cost*. ACM CCS 2022

Almeida, Barbosa, Barthe, Blot, Grégoire, Laporte, Oliveira, Pacheco, Schmidt, Strub. *Jasmin: High-Assurance and High-Speed Cryptography*. ACM CCS 2017

- Language with “C-like” syntax
- Programming in Jasmin is much closer to assembly:
 - Generally: 1 line in Jasmin → 1 line in assembly
 - A few exceptions, but highly predictable
 - Compiler does not schedule code
 - Compiler does not spill registers
- Compiler is formally proven to preserve semantics
- Static (trusted) safety checker
- Compiler is formally proven to preserve constant-time property³

³Barthe, Grégoire, Laporte, and Priya. *Structured Leakage and Applications to Cryptographic Constant-Time and Cost*. ACM CCS 2022

- Can do (almost) everything you can do in assembly
- Architecture-specific implementations
- Small limitations to enable static safety checking (no raw pointers)

- Can do (almost) everything you can do in assembly
- Architecture-specific implementations
- Small limitations to enable static safety checking (no raw pointers)
- Easier to write and maintain than assembly
 - Loops, conditionals
 - Function calls (optional: inline)
 - Function-local variables
 - Register and stack arrays
 - Register and stack allocation

Performance of Kyber code

Implementation	operation	Skylake	Haswell	Comet Lake
C/asm AVX2	keygen	49572	47280	41682
	encaps	60018	62900	55956
	decaps	45854	47784	43906
Jasmin AVX2 (fully verified)	keygen	106578	96296	93244
	encaps	119308	111536	107474
	decaps	105336	98328	96564
Jasmin AVX2 (fully optimized)	keygen	50004	48800	45046
	encaps	65132	63988	59496
	decaps	50340	51444	48172

Security – “constant time”

- Enforce constant-time on Jasmin source level
- Every piece of data is either `secret` or `public`
- Flow of secret information is traced by type system

“Any operation with a secret input produces a secret output”

Security – “constant time”

- Enforce constant-time on Jasmin source level
- Every piece of data is either `secret` or `public`
- Flow of secret information is traced by type system
 - *“Any operation with a secret input produces a secret output”*
- Branch conditions and memory indices need to be `public`

Security – “constant time”

- Enforce constant-time on Jasmin source level
- Every piece of data is either `secret` or `public`
- Flow of secret information is traced by type system

“Any operation with a secret input produces a secret output”

- Branch conditions and memory indices need to be `public`
- In principle can do this also in, e.g., Rust (`secret_integers` crate)
- **Remember: Jasmin compiler is verified to preserve constant-time!**

Security – “constant time”

- Enforce constant-time on Jasmin source level
- Every piece of data is either `secret` or `public`
- Flow of secret information is traced by type system

“Any operation with a secret input produces a secret output”

- Branch conditions and memory indices need to be `public`
- In principle can do this also in, e.g., Rust (`secret_integers` crate)
- **Remember: Jasmin compiler is verified to preserve constant-time!**
- Explicit `#declassify` primitive to move from `secret` to `public`

Security – Spectre v1 (“Speculative bounds-check bypass”)

```
stack u8[10] public;
stack u8[32] secret;
reg u8 t;
reg u64 r, i;

i = 0;
while(i < 10) {
    t = public[(int) i] ;
    r = leak(t);
    ...
}
```

Extending the type system

- Type system gets three security levels:
 - `secret`: secret
 - `public`: public, also during misspeculation
 - `transient`: public, but possibly secret during misspeculation

Extending the type system

- Type system gets three security levels:
 - `secret`: `secret`
 - `public`: `public`, also during misspeculation
 - `transient`: `public`, but possibly `secret` during misspeculation
- Don't branch or index memory based on `secret` **or** `transient` data

Extending the type system

- Type system gets three security levels:
 - `secret`: secret
 - `public`: public, also during misspeculation
 - `transient`: public, but possibly secret during misspeculation
- Don't branch or index memory based on `secret` **or** `transient` data
- Guide programmer to protect code
- Selective speculative load hardening (selSLH):
 - Misspeculation flag in register
 - Mask "transient" values with flag before leaking them

Extending the type system

- Type system gets three security levels:
 - `secret`: secret
 - `public`: public, also during misspeculation
 - `transient`: public, but possibly secret during misspeculation
- Don't branch or index memory based on secret **or transient** data
- Guide programmer to protect code
- Selective speculative load hardening (selSLH):
 - Misspeculation flag in register
 - Mask "transient" values with flag before leaking them
- Overhead for Kyber768 (on Intel Comet Lake):
 - 0.28% for Keypair
 - 0.55% for Encaps
 - 0.75% for Decaps
- Exploits synergies with protections against "traditional" timing attacks

Ammanaghatta Shivakumar, Barthe, Grégoire, Laporte, Oliveira, Priya, Schwabe, and Tabary-Maujean. *Typing High-Speed Cryptography against Spectre v1*. IEEE S&P 2023.

“... A cryptographic module shall provide methods to zeroize all plaintext secret and private cryptographic keys”

—FIPS 140-3, Section 9.7.A

“... A cryptographic module shall provide methods to zeroize all plaintext secret and private cryptographic keys”

—FIPS 140-3, Section 9.7.A

Goal of zeroization

Scrub all (sensitive) data from memory (stack) and registers when crypto routine returns.

“... A cryptographic module shall provide methods to zeroize all plaintext secret and private cryptographic keys”

—FIPS 140-3, Section 9.7.A

Goal of zeroization

Scrub all (sensitive) data from memory (stack) and registers when crypto routine returns.

Failure modes

0. Don't perform any zeroization

“... A cryptographic module shall provide methods to zeroize all plaintext secret and private cryptographic keys”

—FIPS 140-3, Section 9.7.A

Goal of zeroization

Scrub all (sensitive) data from memory (stack) and registers when crypto routine returns.

Failure modes

0. Don't perform any zeroization
1. Dead-store elimination

“... A cryptographic module shall provide methods to zeroize all plaintext secret and private cryptographic keys”

—FIPS 140-3, Section 9.7.A

Goal of zeroization

Scrub all (sensitive) data from memory (stack) and registers when crypto routine returns.

Failure modes

0. Don't perform any zeroization
1. Dead-store elimination
2. Only API-level stack zeroization

“... A cryptographic module shall provide methods to zeroize all plaintext secret and private cryptographic keys”

—FIPS 140-3, Section 9.7.A

Goal of zeroization

Scrub all (sensitive) data from memory (stack) and registers when crypto routine returns.

Failure modes

0. Don't perform any zeroization
1. Dead-store elimination
2. Only API-level stack zeroization
3. Don't scrub source-level invisible data

“... A cryptographic module shall provide methods to zeroize all plaintext secret and private cryptographic keys”

—FIPS 140-3, Section 9.7.A

Goal of zeroization

Scrub all (sensitive) data from memory (stack) and registers when crypto routine returns.

Failure modes

0. Don't perform any zeroization
1. Dead-store elimination
2. Only API-level stack zeroization
3. Don't scrub source-level invisible data
4. Mis-estimate stack space when scrubbing from caller

Solution in Jasmin compiler

Zeroize used stack space and registers when returning from export function

Arranz Olmos, Barthe, Gonzalez, Grégoire, Laporte, Léchenet, Oliveira, Schwabe: *High-assurance zeroization*.
TCHES 2024-1.

Solution in Jasmin compiler

Zeroize used stack space and registers when returning from export function

- Make use of multiple features of Jasmin:
 - Compiler has global view
 - All stack usage is known at compile time
 - Entry/return point is clearly defined

Arranz Olmos, Barthe, Gonzalez, Grégoire, Laporte, Léchenet, Oliveira, Schwabe: *High-assurance zeroization*.
TCHES 2024-1.

Solution in Jasmin compiler

Zeroize used stack space and registers when returning from export function

- Make use of multiple features of Jasmin:
 - Compiler has global view
 - All stack usage is known at compile time
 - Entry/return point is clearly defined
- Performance overhead for Kyber768:
 - 0.59% for Keypair
 - 0.24% for Encaps
 - 1.04% for Decaps

Arranz Olmos, Barthe, Gonzalez, Grégoire, Laporte, Léchenet, Oliveira, Schwabe: *High-assurance zeroization*.
TCHES 2024-1.

<https://github.com/formosa-crypto/libjade>

- Collection of primitive implementations rather than library
- “A library to be used by libraries”

<https://github.com/formosa-crypto/libjade>

- Collection of primitive implementations rather than library
- “A library to be used by libraries”
- Example:

```
cd src/crypto_kem/kyber/kyber768/amd64/ref/ && make
```

will build

```
src/crypto_kem/kyber/kyber768/amd64/ref/kem.s
```

with API described in

```
src/crypto_kem/kyber/kyber768/amd64/ref/include/api.h
```

- Releases contain
 - compiled assembly files + headers
 - jasmin files
 - usage examples written in C
- Latest release: 2023.05-1

- Releases contain
 - compiled assembly files + headers
 - jasmin files
 - usage examples written in C
- Latest release: 2023.05-1
- Plans for next release:
 - Integrate EasyCrypt proofs (covered by CI)
 - Integrate/consolidate various features
 - Special focus on Kyber-768

More proof automation!

- Integrate with CryptoLine (<https://github.com/fmlab-iis/cryptoline>)⁴
 - (semi-)automated proof of branch-free arithmetic
 - “Prove without understanding code”
- Automated equivalence proving. . .

⁴Fu, Liu, Shi, Tsai, Wang, and Yang. Signed Cryptographic Program Verification with Typed CryptoLine. ACM CCS 2019

More proof automation!

- Integrate with CryptoLine (<https://github.com/fmlab-iis/cryptoline>)⁴
 - (semi-)automated proof of branch-free arithmetic
 - “Prove without understanding code”
- Automated equivalence proving. . .

Beyond Spectre v1

- Spectre v2: Avoid by not using indirect branches
- Spectre v4: Use SSBD: <https://github.com/tyhicks/ssbd-tools>
- **Spectre protection requires separation of crypto code!**

⁴Fu, Liu, Shi, Tsai, Wang, and Yang. Signed Cryptographic Program Verification with Typed CryptoLine. ACM CCS 2019

Support more architectures

- 32-bit Arm (ARMv7ME): works, still “experimental”
- Opentitan’s OTBN: work in progress
- 64-bit ARM and RISC-V: very early WIP

Support more architectures

- 32-bit Arm (ARMv7ME): works, still “experimental”
- Opentitan’s OTBN: work in progress
- 64-bit ARM and RISC-V: very early WIP

Secure interfacing

- Currently use C function-call ABI (caller/callee contract through documentation)
- Check/Enforce caller requirements?
- Stronger safety notions (e.g., interfacing with Rust)

Make high-assurance tools mainstream/default!

<https://formosa-crypto.org>