CROSSBAR

DARIC – The World's First SPU

Introduction

What off-the-shelf chips are available?

Simplifying a bit for clarity, we can say there are two classes of chips that often find their way into crypto wallet: Microcontroller (MCU) and Secure Element (SE).

Microcontrollers are designed to be a jack-of-all-trades, prioritizing generality over everything else. This is extremely useful, leading to a market of about 30 billion units per year. Everything from servo motors to home appliances can be built around MCUs. They excel in applications where non-recurring costs dominate, so that design convenience is paramount. They support every type of connection and peripheral, with easy-to-use drivers. Some have basic logical security, but do not have advanced physical security. Physical security requires specialized design across the entire die, and is not suitable for typical MCU markets.

Secure Elements are in many ways opposite of an MCU. They serve price-sensitive commodity markets, and mostly focus on performing decades-old functions such as P-256 and RSA at rock-bottom price. This requires dispensing with any flexibility or extra functionality. This also serves to minimize attack surface, under a principle of "Economy of Mechanism." SE's will generally have Physical Countermeasures, which are essentially required for relevant certifications. As opposed to logical security, these defend against physical attacks such as probing, FIB, laser injection, voltage glitching, timing glitch, thermal disturbance, and so on.

What's wrong with off-the-shelf parts?

While MCU's and SE's are great for their intended markets, they are not designed for self-custody wallets. The MCU has good flexibility and processing power, but not physical security; The SE has good

security, but not processing power or flexibility. Combining the two leaves some shortcomings in a wallet application.

A typical wallet architecture is represented by the following diagram:

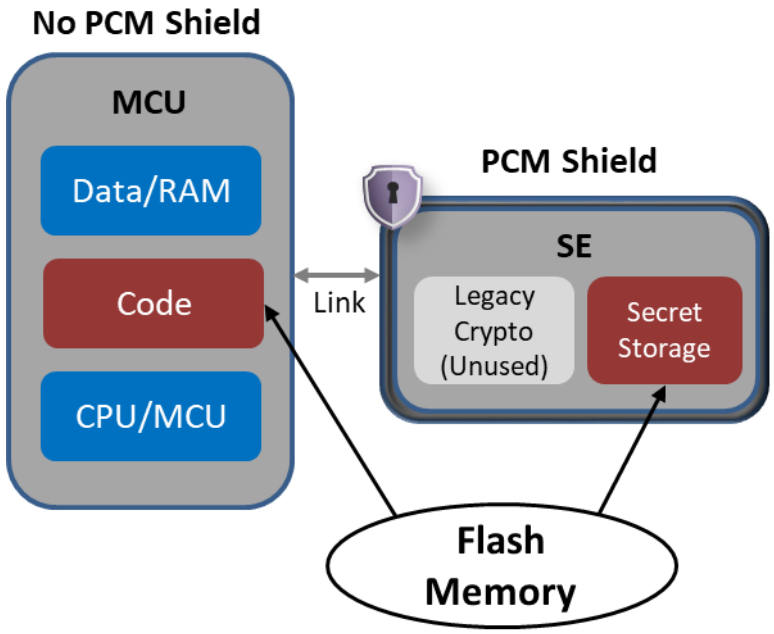

A primary problem is that Secure Elements typically do not support blockchain-specific additional functionality, such as secp256k1 curve, nor MPC-specific feature such as long RSA for Pallier keys. Therefore, the typical wallet uses the SE only as a kind of secure memory. The entire private key is reconstituted over on the MCU side, where the actual calculation is done. It seems the primary purpose of the SE in this configuration is only to be able to claim "EAL 6+." That certification seems irrelevant if no secure calculation is done there, but that is the standard industry practice[2].

The SE is also generally of low compute power, using for example an ARM M0 to M3 type core at 50-70MHz, with a small non-volatile memory and SRAM. This is sufficient for the intended purpose, but leaves no room for any kind of advanced computation.

So given these shortcomings of mass-market SEs, the MCU is forced into serving security-critical roles. Of course, first and foremost is processing the private key during signing. This is a clear attack vector, exploited for example in [4]. However there are many other avenues of attack in addition.

[2] As of this writing, the only wallet (to our knowledge) that actually performs secp256k1 signature in a Secure Element is Ledger. The SE firmware is customized slightly for them by ST Micro. This is laudable, and a huge improvement over using the SE as a memory. But it is not open source and only supports single-key signatures.

For example, key generation is also a critical and common attack vector. Blockchain security typically depends on high-entropy private key material. If the MCU is involved in key generation (and in many wallets it is solely responsible for key generation), a simple attack is to generate keys with very limited entropy, allowing the hacker to search a very limited, predetermined key space. This allows a hacker drain wallets post-deployment, without the need to interact with the user in any way [3].

Not only is the critical cryptography done in the less secure MCU, but the entire user interface runs on that side. This allows an attacker to take funds without learning the existing private key – for example, by presenting a fraudulent address to the display, or gaining control of user input, or bypassing user identity validation[3].

Broadly speaking, the lack of Physical Countermeasures (PCM) on the MCU [4],[5],[9], and the general lack of security-focused design [7], opens up many avenues of attack, of all of the aforementioned categories and more.

But in addition to the individual shortcomings of the individual MCU + SE, the simple fact that the architecture is dis-integrated presents a problem. Since the chips are separate, and the communication between them can easily be monitored (or manipulated), some strategy must be devised whereby each chip can detect if the overall environment has been tampered with.

This is more easily said than done. Reference [6] articulates some of the general challenges. In [10], Ledger points out some flaws with Trezor's method. Ledger's method has some enhancements, but reference [5] shows a demonstration of a Ledger wallet MCU being flashed with custom firmware. Again, as explained above, gaining control of either chip allows many types of attacks[4].

Note that stealing a wallet that is in use, and performing such lengthy, invasive attacks, could be logistically challenging. But a far easier avenue is to acquire wallets, modify them, then place them on the market. This can and does happen [3]. This "supply chain" method obviously removes most of that challenge.

An aside about memory

Another shortcoming of the MCU and SE chips used in typical hardware wallets is that the non-volatile storage (in both chips) is based on floating-gate flash memory. This can be manufactured in logic processes down to about 40nm[5], and so is a popular choice for chips combining non-volatile memory and logic.

[3] For example, see [6] and [8] for discussion of the importance of protecting the UI.
[4] Without delving into a maze of detail, whether extra chips contribute "defense in depth" or "multiple attack lanes" depends on the detail of the situation and the usage. Certainly if you have an MCU+SE in a disintegrated architecture, then having MCU+SE+SE is an improvement. But having at least one MCU+SE monolithic removes certain attack lanes.
[5] It is available in slightly more advanced nodes, such as 28nm, but becomes increasingly impractical such that in practice it is little-used below 40nm.

Unfortunately, this comes with two key disadvantages, especially for critical security and asset-protection roles: fragility and inspectability.

Floating gate flash is based on a trapped charge, and confining a charge in a small space is a high-energy state. The trapped electrons will seek to leak and dissipate, at any disturbance (or simply with aging). This is the same technology as your typical SD card or flash thumb drive, and most of us have experienced failures of these devices induced by aging and/or environment (temperature, radiation, magnetic field).

Also, in floating gate flash the 0's and 1's are represented by an electric charge, and therefore can be read without much difficulty by, for example, electron microscopy [11].

For further detail on these topics, see [12]

Introducing the SPU: A monolithic MCU + SE tailored for crypto

As described in the previous section, existing MCU and SE chips are not designed for HW crypto wallets, and have many shortcomings in that application. Meanwhile, advances in cryptography and system architecture create a need for even more advanced processing, even farther beyond the capability of existing embedded chips [13].

Also, the market is trending towards multi-function devices, and for good reason. Once a secure hardware device exists, it is convenient and cost-efficient to apply it to as many purposes as possible.

This situation of increasing requirements, and lack of any purpose-built silicon foundation, led Crossbar to create the world's first Secure Processing Unit (SPU): Essentially a combined MCU and SE, in advanced node process, and tailored specifically for blockchain applications. The first instantiation of this chip, codenamed "Daric", has already taped out and is being integrated into our PHSM product line. A second version "Denarius" is under development.

SPU: Secure Processing Unit

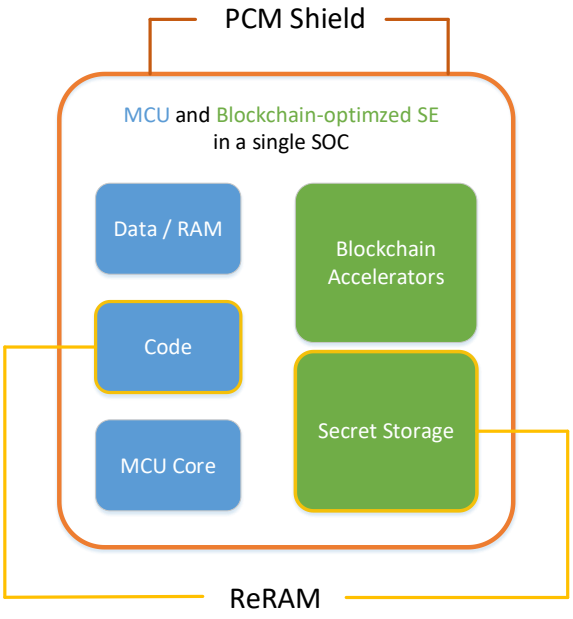

The Secure Element side of Daric is a complete crypto engine with direct hardware support for all popular blockchain curves and signature types. It also includes support for system-level functions. For example hierarchical deterministic key generation is supported, so that private keys can be generated from a mnemonic and relevant parameters, entirely within the SE environment. Also long RSA (up to 6144 bits) is supported, for Pallier key generation used in certain MPC protocols. Multiple NIST 800-90B-complaint physical TRNGs are included, with 800-90A-complaint post processing.

The MCU side contains two completely independent MCU's, to support various SW development, Open Source, and security paradigms. One is an ARM M7 with MPU, the other is a complete open source implementation of RISC-V (VexRisc) with a true MMU.

As opposed to the M0 to M3 type of MCU found in most Secure Elements, the ARM M7 is much more powerful. See [2] for full details.

	ARM M0 / M3	**ARM M7**	**Comment**
FPU	No	Yes	
DSP extensions	No	Yes	
DMIPs/MHz	0.96-1.24	2.31	
Pipeline(*)	2 or 3 stage	6 stage	
iCache	No	Yes	16kB in Daric).
dCache	No	Yes	16kB in Daric.
iTCM	No	Yes	256kB in Daric.
dTCM	No	Yes	64kB in Daric.
Bus	AHB-	AXI	

(*) In a fixed process, can achieve higher clock rate with more pipeline stages.

On the Risc-V side, the VexRisc on Daric runs at half the clock rate of the ARM M7. ARM, for whatever reason, does not support any M-series core with a true MMU (enabling memory virtualization as an isolation and security strategy). This feature is supported only on the much higher-end ARM A-series. Therefore the Risc-V on Daric allows an operating point not found in almost any current MCU. See [16] for more discussion of the benefits of this novel feature.

All non-volatile memory on our SPU is resistive (ReRAM) type. This avoids the fragility and inspectability of floating gate memory (described above). Another advantage is the main reason for existence of ReRAM in the first place: it allows use of a more advanced logic process[6]. This has all the normal benefits in cost, power, size that push all products toward more advance processes.

In our context, it has other secondary benefits. It allows us to put (for small cost) unusually large code and data memory directly on-chip. This expands boundaries in product capability, function, and security (keeping more access on-chip). It also benefits the Physical Countermeasures. The active mesh metal is at a finer pitch; gate switching current draw and EMI are lower, making them that much harder to detect.

Benefits of Integration

In addition to the aforementioned benefits of the MCU, SE, and ReRAM individually, the most powerful benefits derive from their integration on a single die, under a single supply chain validation, and protected by a common umbrella of Physical Countermeasures (PCM).

Zero Trust Architecture

In SE design, mission-critical values such as lifecycle are often routed and encoded with redundancy and protection by “glue cells” (aka canary cells). If these exist on only one chip, then complex and imperfect measures must be implemented for other chips and other functions to do the right things based on that lifecycle. But monolithic integration allows the MCU, the SE, and all MCU-SE interconnect to exist under a single umbrella of PCM protection, and avoids any such complexities.

By integrating cryptographic operations, key management, user interface, and blockchain-specific functions (e.g., secp256k1, MPC) into a single secure chip under a common PCM umbrella, the Daric SPU enforces a Zero Trust model. This eliminates reliance on external MCUs or firmware, closing attack surfaces like I^2C/SPI snooping, insecure firmware updates, and counterfeit chip substitution. Comprehensive Physical Countermeasures (PCMs) protect against physical attacks (e.g., probing, laser injection), while cryptographic provenance verification

[6] The general market for MCUs and SEs will likely move to more advanced nodes with ReRAM, and this is already starting. But currently, Crossbar is a leader in this trend, and the mainstream market is still at 40nm and above while Daric is at 22nm.

(e.g., ECDSA-based challenges) mitigates supply chain risks, such as malicious firmware. Trust boundaries are hardware-enforced, ensuring robust, verifiable security for blockchain custody.

Producing our own chip, we can put in serial number, certificates, and associated attestation keys directly on the factory floor. And these can be validated at any stage, including by the end user. And since the SPU covers not only cryptography, but also identity, user input, and user output (display), all functions are covered by supply chain certification and PCM from factory to user interface.

Atomic Security Policy Enforcement

The Daric SPU's monolithic design enables atomic enforcement of security policies by integrating lifecycle states, debug fuses, key usage policies (e.g., no-export, sign-only for private keys), and attestation roots within a single chip. This ensures security invariants, such as preventing key leakage or unauthorized debug access, are enforced indivisibly without relying on vulnerable inter-chip signaling (e.g., I^2C/SPI). Integrated tamper detection circuits monitor all subcomponents—MCU, crypto engine, and ReRAM—in real time, enabling rapid responses (e.g., wiping private keys or locking down the chip) to physical attacks like probing or laser injection. Compared to distributed MCU+SE architectures, this reduces tamper response latency and enhances reliability, providing robust, cryptographically verifiable security for blockchain custody applications.

Power Side-Channel Resistance

The Daric SPU's single-chip design enhances resistance to power side-channel attacks, such as differential power analysis (DPA), by integrating cryptographic operations, MCU, and ReRAM under unified power management. Advanced noise shaping and power delivery optimization, enabled by the 22nm process, obscure power consumption patterns during blockchain operations like ECDSA signing. In contrast, multi-chip MCU+SE systems make it trivially easy to examine the power signature of the SE in complete isolation, increasing vulnerability to leakage. This monolithic approach significantly reduces the detectability of sensitive operations, bolstering security for blockchain custody applications.

Long-Term Updateability and Lifecycle Management

The Daric SPU's monolithic design simplifies long-term firmware updateability and lifecycle management, ensuring robust security and an enhanced user experience for blockchain custody applications. By integrating cryptographic operations, MCU, and ReRAM within a single chip, the SPU enables streamlined, cryptographically verified firmware updates (e.g., using ECDSA signatures) that eliminate the complexity of coordinating updates across multiple chips in traditional MCU+SE architectures. This reduces the risk of vulnerabilities from outdated firmware, such as those exploited in multi-chip systems, and ensures consistent lifecycle state enforcement (e.g., provisioning, operational, decommissioned) across all components. For users, this translates to seamless, secure updates that maintain wallet functionality and protect

against emerging threats without the delays or errors common in distributed designs, enhancing both security and usability.

Aside on Open Source

Our SPU product line has an unmatched level of Open Source and verifiable content. That is beyond the scope of this note, but is an important distinction between our SPU and the typical "security by obscurity" business model of incumbent SE vendors. For more detail see our whitepaper [17].

Conclusion

In summary the Crossbar line of SPUs offers an unprecedented options for the design of client-side personal security devices, such as – but not limited to – our Cramium PHSM product line:

- Blockchain-specific HW acceleration;
- High performance compute for new and developing architectures and protocols;
- Large code and data memory;
- Supply chain protection from factory to user interface, covering all device functions;
- Advanced node process for enhanced cost, power, size, functionailty, and security;
- ReRAM that is permanent, robust, and private;
- MCU, SE, and interconnect under a common PCM umbrella.

The combined blockchain-tailored SE, plus powerful MCU, on a single die, gives SPU-based devices the flexibility and agility of SW, with the security of purpose-built HW.

REFERENCES:

[1] Bitkey discussion about some of the factors in deciding when to build a custom chip.
https://bitkey.build/processing-our-processor-choice/

[2] ARM comparison table published by ARM at https://documentation-service.arm.com/static/61bb37962183326f2176f8cc

[3] Example of supply chain problem
https://www.kaspersky.com/blog/fake-trezor-hardware-crypto-wallet/48155/

[4] Popular video about a voltage glitching attack
https://www.youtube.com/watch?app=desktop&v=dT9y-KQbqi4

[5] Flashing Ledger with custom firmware around 24:22 at this video; this video also has many other HW wallet attack descriptions.
https://www.youtube.com/watch?v=Y1OBIGslgGM

[6] Similar discussion of problems with MCU and SE verifying each other:
https://saleemrashid.com/2018/03/20/breaking-ledger-security-model/

[7] Example of non-SE chip that is popular and sufficient for many purposes, but not for mission-critical HW wallet private key.
https://cryptodeeptech.ru/bitcoin-bluetooth-attacks/

[8] Ledger discusses importance of protecting the display
https://www.ledger.com/academy/topics/ledgersolutions/ledger-wallets-secure-screen-security-model

[9] Example of security downgrade by light injection (physical attack)
www.usenix.org/system/files/conference/woot17/woot17-paper-obermaier.pdf

[10] Ledger note on some challenges of MCU+SE mutual validation
https://www.ledger.com/why-secure-elements-make-a-crucial-difference-to-hardware-wallet-security

[11] Reading flash memory by electon microscopy
https://www.cl.cam.ac.uk/~sps32/istfa2016_sem.pdf

[12] Crossbar whitepaper “ReRAM Advantages”

[13] !Crossbar whitepaper on distributed workflows/MPC – the custody dilemma

[14] Crossbar whitepaper “Cramium PHSM Series SDK Product Introduction”

[15] VexRisc source
https://github.com/SpinalHDL/VexRiscv

[16] Xous OS
https://xobs.io/announcing-xous-the-betrusted-operating-system/

[17] Crossbar whitepaper “Towards and Open Source Foundation for Cryptography”