

SoK: Message Authentication Codes in Constrained Machine-to-Machine Communication

Eric Wagner*, Shalini Banerjee^{†‡}, David Heye[§], Jan Bauer[¶], Martin Serror[¶],
Lennart Bader[¶], Klaus Wehrle[§], Andy Rupp^{*‡}, Jan Pennekamp[§], and Vincent Lenders*

*University of Luxembourg [†]Karlsruhe Institute of Technology [‡]KASTEL Security Research Labs

[§]RWTH Aachen University [¶]Fraunhofer FKIE

{firstname.lastname}@uni.lu {firstname.lastname}@kit.edu

{lastname}@comsys.rwth-aachen.de {firstname.lastname}@fkie.fraunhofer.de

Abstract—Message Authentication Codes (MACs) are the de facto standard to ensure authenticity and integrity in constrained machine-to-machine (M2M) communication. Yet, research on optimizing MACs for these environments has received little systematic treatment, and deployed protocols still rely on generic constructions such as HMAC-SHA256 and CMAC. This SoK paper provides the first comprehensive analysis of MACs for constrained M2M communication. We survey authentication mechanisms across 10 deployment domains, identify 20 lightweight MAC schemes through a systematic literature review, and evaluate 13 of them on embedded hardware, showing that software-implemented lightweight MACs can outperform hardware-accelerated conventional schemes and authenticated encryption for the short payloads prevalent in practice. We further consolidate research on MAC aggregation and MAC-based multicast authentication, the latter growing more relevant as post-quantum requirements push digital signatures further out of reach. From this, we derive open challenges and provide actionable guidance for researchers and system designers.

Index Terms—Message Authentication Code, Cyber-Physical Systems, Source Authentication

1. Introduction

Message Authentication Codes (MACs) are fundamental to securing the authenticity and integrity of machine-to-machine (M2M) communication, yet their deployment in constrained environments remains challenging [1]. Devices such as Programmable Logic Controllers (PLCs), sensors, and actuators are often limited in terms of computational power [2], energy supply [3], latency [4], and communication bandwidth [5]. These constraints directly impact the feasibility and efficiency of cryptographic operations, whose design has traditionally prioritized strong security guarantees over resource efficiency. With the rapid proliferation of the (industrial) Internet of Things (IoT) and Cyber-Physical Systems (CPSs), the need for MACs that are both secure and lightweight has become increasingly pressing [1].

The research community and standardization bodies have devoted significant effort to designing lightweight cryptographic primitives. However, initiatives such as the NIST Lightweight Cryptography standardization process [6] or the CAESAR competition [7] have focused primarily on authenticated encryption (AE) schemes, which combine confidentiality and integrity protection into a single primitive. While AE schemes offer advantages in reducing implementation complexity and amortizing cryptographic costs, they are not universally applicable to constrained M2M communication. In many M2M scenarios, confidentiality is either unnecessary, *e.g.*, in power grid communication to save processing overhead [8], or even undesirable, *e.g.*, in broadcasted Global Navigation Satellite System (GNSS) signals for civil applications [9].

In contrast, integrity protection is critical, particularly for CPSs, where adversaries can manipulate control messages or sensor data to induce unsafe physical behavior [10]. Real-world incidents, such as attacks on the Ukrainian power grid, demonstrate how unauthorized command injection can cause large-scale disruptions to critical infrastructure [11]. In such settings, preventing forgery and tampering of messages is crucial, even if messages are publicly observable. Fittingly, standalone MACs provide message integrity and authenticity without the overhead of full AE schemes [12].

M2M communication, however, imposes additional requirements on MAC schemes. Constrained devices frequently transmit messages as short as 1 byte over bandwidth-constrained links [13], making the size of authentication tags a critical factor, especially as limited energy is primarily consumed for data transmission [14]. Consequently, techniques such as tag truncation [15], aggregation of authentication across multiple messages [16], and protocol-aware optimizations [17] reduce communication overhead while maintaining acceptable security levels. Processing and latency constraints simultaneously drive the need for low-complexity schemes [4].

Despite this active research landscape, the field remains fragmented. Prior surveys treat standalone MACs only as a footnote to broader discussions of lightweight cryptography, leaving researchers and practitioners without a clear

Domain	Typical Protocols	Msg. Size [byte]	Latency Req. [ms]	Data Rate [bit/s]	Devices [#]
 Body Area Networks	802.15.6, Zigbee, BLE	1–256	100	10–50 M	≤ 64
 Intravehicular	CAN, CAN-FD, LIN	8–64	< 20	20 k–8 M	≤ 100
 Smart Home	ZigBee, Thread, BLE	8–1024	> 100	> 250 k	≤ 100
 Industrial Control Systems	5G, 802.15.4, MIOTY	1–100	< 1	M–G	diverse
 Smart City	LoRaWAN, Sigfox, MIOTY	< 222	diverse	1–M	> 1000
 Underwater Networks	JANUS, GUWAL	8–16	> 1000	100–10 k	< 100
 Power Grid	GOOSE, SMV	> 100	0.2–3	≥ 100 M	≤ 1000
 Aviation	ADS-B, MODE S	7–14	–	≈ 1 M	$\gg 1000$
 Satellite Navigation	GPS, Galileo	64–122 ¹	–	25–125	∞

¹ Minimal number of message bits required for position calculation.

TABLE 1: Constrained environments differ significantly in their requirements. For example, Body Area Networks (BANs) use short messages with relaxed latency constraints, whereas power grids require low latency for larger messages supported by high data rates. Further information on the respective domains can be found in Appendix A.

view of the MAC design space, its trade-offs, and which schemes suit which applications. As such, comparatively less attention has been given to systematically comparing security guarantees across design paradigms or relating the underlying assumptions to deployment constraints. Consequently, advances in cryptography rarely translate into practical deployments, even where they could substantially alleviate resource pressure. This disconnect extends to the role of symmetric authentication in multicast scenarios. Although traditionally addressed with digital signatures, their cost remains prohibitive in constrained environments and increases further under post-quantum requirements [18]; MAC-derived multicast authentication can often fill this role in practice, despite weaker security properties.

This Systemization of Knowledge (SoK) paper comprehensively analyzes MAC constructions tailored to constrained environments. We first map the requirements of M2M deployments (Section 2) and position our work against prior work on lightweight cryptography (Section 3). Afterward, we introduce MACs, their security definitions, and deployment considerations (Section 4), and then survey how authentication is implemented in practice across domains (Section 5). We then consolidate and critically evaluate work on resource-efficient MACs (Section 6), bandwidth optimizations (Section 7), and group communication (Section 8). Security is analyzed under the standard EUF-CMA paradigm, clarifying how deployment constraints stress different components of the (EUF-CMA) bound in practice. Finally, we derive open challenges for future research (Section 9). This paper thus serves both as a foundation for future research and a practical guide for system designers.

Open Science. Artifacts are available at <https://doi.org/10.5281/zenodo.22839391>.

2. Constrained M2M Communication

M2M communication generally refers to the periodic or event-driven exchange of information (*i.e.*, messages) between a large number of electronic devices to enable sens-

ing, processing, and actuation activities [19], [20]. Depending on the realized service or functionality, such message exchanges are subject to particular performance requirements, *e.g.*, regarding latency. Especially stringent demands in these dimensions are commonly referred to as *mission-critical* M2M communication [21]. Meeting these requirements typically necessitates communication protocols adapted to the application environment and the available communication and computational resources. To understand these requirements, we look at the constraints related to cryptographic processing, *i.e.*, message sizes, latency requirements, and data rates for different network scales, from BANs to Global Area Networks (GANs) for satellite navigation.

Table 1 summarizes our analysis, which is discussed further in Appendix A. It shows that constrained environments cannot easily be grouped. Power grids, on the one hand, have tight latency requirements to prevent power outages, but they can take advantage of high-bandwidth channels. On the other hand, BAN and underwater communication can only transfer a few small messages, but with more relaxed latency requirements. Yet other environments, such as Industrial Control Systems (ICSs), are constrained by message sizes and latency requirements alike. Consequently, optimizations have to occur in different dimensions to address the limitations of current technologies.

Insight

Constrained environments differ in both type and scale and therefore require different optimization.

Regarding the security of M2M communication, message authentication is of utmost importance [22]. It fulfills two roles: ensuring that messages stem from the alleged source and that they have not been tampered with during transmission. For example, for a power grid operator, reliable and timely control of assets is critical. To achieve this, confidentiality is gladly forgone [8]. Even more, broadcast communication, as seen in *e.g.*, intra-vehicular communication or satellite navigation, explicitly forgoes confidentiality to not exclude devices without decryption capabil-

ities (legacy devices or those in the process of retrieving keys). Hence, mission-critical M2M communication needs optimized security to provide (standalone) message authentication, for which the de facto standard is MACs.

3. Related Work

The need for lightweight cryptography in constrained environments is well-established. However, consolidating work has largely focused on encryption and AE. Meanwhile, standalone MACs have received barely any attention.

Surveys and evaluations of lightweight cryptography have primarily focused on block ciphers. Where MAC evaluation is a focus, it has been limited to CMAC instantiated with different block ciphers [23]. The state of lightweight cryptography has been surveyed more broadly [24], yet only SipHash and Chaskey are discussed as standalone lightweight MAC algorithms. Work centered on lightweight MAC constructions and their security properties exists [25], but does not address deployment or performance in constrained settings. Likewise, other surveys on lightweight cryptography restrict their scope to block and stream ciphers, leaving MAC schemes largely unexamined [26], [27].

Current benchmarking evaluates candidates from the NIST Lightweight Cryptography standardization process, further reinforcing the community's prevailing focus on AE. These evaluations examine FPGA and ASIC implementations [28], the hardware efficiency and energy consumption of candidate schemes [29], and software implementations targeting RISC-V microarchitectures [30].

Insight

The current body of work does not address the distinct requirements and trade-offs that arise when message authentication is needed without confidentiality, a common scenario in constrained M2M communication where overhead budgets are tight.

4. Message Authentication Codes (MACs)

We begin by giving a general introduction to MACs, their general security definitions, deployments, and recommendations. This sets the stage to establish their limitations in constrained environments and the resulting impact.

4.1. Notations

We let $\{0, 1\}^n$ denote the set of n -bit strings. By $|x|$, we denote the bit-length of string x . We use 0^ℓ to denote the all-zero string of length ℓ . For a finite set S , we use $x \leftarrow S$ to denote the uniformly random sampling of an element x from S . For a probabilistic algorithm $\mathcal{A}$, we let $y \leftarrow \mathcal{A}(x)$ denote the process of running $\mathcal{A}(x)$ with fresh randomness and assigning the result to y . We let $\|_{i=1}^n(a_i)$ denote the concatenation $a_1 \| \dots \| a_n$ of the sequence $(a_1, \dots, a_n)$. We use $\text{Trunc}_s(x)$ for the truncation of x to s bits. We let $\lambda \in \mathbb{N}$ denote the security parameter. We use PPT to denote a randomized algorithm running in time $a \cdot \lambda^c$ for constants a

and c . A function $\text{negl} : \mathbb{N} \rightarrow \mathbb{R}^+$ is called negligible in λ , if it grows slower than λ^{-c} for every constant c .

4.2. Security Definitions

MACs are cryptographic primitives that authenticate the source of a message, *i.e.*, verify that a message originates from the claimed sender. In addition, they provide message integrity, ensuring that messages are not modified in transit. At a high-level, a MAC scheme consists of a triplet of efficient algorithms $\Pi_{\text{MAC}} = (\text{Gen}, \text{Tag}, \text{Vrfy})$. The key generation algorithm Gen samples a key k that is pre-shared between a sender and a receiver. The tag generation algorithm Tag is run by the sender, which takes as input the key k and a message m and outputs a tag tag . The verification algorithm Vrfy is run by the receiver, which takes as input a key k , a message m and a tag tag and outputs 0 (invalid) or 1 (valid). Correctness means that a MAC scheme can reliably validate a tag when given the same key and message used to generate the tag as input.

The intuitive idea behind Existential UnForgeability under Chosen Message Attack (EUF-CMA) security is that no polynomial time adversary $\mathcal{A}$ should be able to generate a valid tag on any “new” message (*i.e.*, a message not seen before), even if $\mathcal{A}$ learns valid tags for messages of their choosing. To capture this, we allow the adversary to adaptively query the tag generation oracle $\text{OTag}(k, \cdot)$ with polynomially many messages and record the responses. We formalize this via the security game $\text{GEUF-CMA}^{\mathcal{A}, \Pi_{\text{MAC}}}(\lambda)$ in Definition 1.

Definition 1. A MAC scheme $\Pi_{\text{MAC}} = (\text{Gen}, \text{Tag}, \text{Vrfy})$ is a triplet of efficient algorithms with a key space $\mathcal{K}$, a message space $\mathcal{M}$, and a tag space $\mathcal{T}$, if the following holds:

$k \leftarrow \text{Gen}(1^\lambda)$: The key generation algorithm Gen takes as input a security parameter $\lambda \in \mathbb{N}$, and outputs a secret key $k \in \mathcal{K}$.

$\text{tag} \leftarrow \text{Tag}(k, m)$: The tag generation algorithm takes as input a secret key $k \in \mathcal{K}$, a message $m \in \mathcal{M}$, and outputs a tag $\text{tag} \in \mathcal{T}$.

$\{0, 1\} \leftarrow \text{Vrfy}(k, m, \text{tag})$: The verification algorithm Vrfy takes as input a secret key $k \in \mathcal{K}$, a message $m \in \mathcal{M}$, and a tag $\text{tag} \in \mathcal{T}$, and outputs 0 or 1.

Furthermore, the algorithms must satisfy the following correctness and security properties.

Correctness. For every message m , the following holds:

$$\Pr \left[\text{Vrfy}(k, m, \text{Tag}(k, m)) = 1 : \begin{array}{l} k \leftarrow \text{Gen}(1^\lambda) \\ \text{tag} \leftarrow \text{Tag}(k, m) \end{array} \right] = 1$$

where probability is taken over choice of k and randomness of Tag algorithm.

EUF-CMA Security. A MAC scheme Π_{MAC} is EUF-CMA secure, if for every PPT adversary $\mathcal{A}$, there exists a negligible function $\text{negl}(\lambda)$, such that for all $\lambda \in \mathbb{N}$, the following holds:

$$\Pr [\text{GEUF-CMA}^{\mathcal{A}, \Pi_{\text{MAC}}}(\lambda) = 1] \leq \text{negl}(\lambda)$$

where the security game is defined as follows:

$\text{GEUF-CMA}^{\mathcal{A}, \Pi_{\text{MAC}}}(\lambda) :$

- 1) $k \leftarrow \text{Gen}(1^\lambda)$
- 2) $(m^*, \text{tag}^*) \leftarrow \mathcal{A}^{\text{OTag}(k, \cdot)}(1^\lambda)$
where $\text{OTag}(k, m) = \text{Tag}(k, m)$
- 3) if $\text{Vrfy}(k, m^*, \text{tag}^*) = 1$ and $m^* \notin \mathcal{Q}$
where $\mathcal{Q}$ is the set of oracle queries made by $\mathcal{A}$
return 1
- 4) else return 0

The strictly stronger notion of Strong EUF-CMA (SUF-CMA) security ensures that $\mathcal{A}$ cannot produce a new valid tag for any message, including messages previously submitted to $\text{OTag}(k, \cdot)$. However, SUF-CMA is generally not a meaningful notion for nonce-based MACs, as a nonce is a part of the authenticated input—a fresh nonce on a previously queried message always yields a tag never returned by the oracle, trivially satisfying the SUF-CMA winning condition. Hence, we focus on the EUF-CMA security when analyzing nonce-based MACs (Appendix B).

The EUF-CMA notion provides asymptotic security guarantees. Nevertheless, it does not capture the actual probability of forgery in real systems. Emphasizing this, the $\text{GEUF-CMA}^{\mathcal{A}, \Pi_{\text{MAC}}}(\lambda)$ game allows q tagging queries but only *one* forgery attempt by $\mathcal{A}$. But in real systems, an adversary can repeatedly query the verification interface t (without rate limits), thereby changing its success probability. Hence, security analysis is important for choosing the optimal parameters for real-world instantiations.

Theorem 1. Let $\Pi_{\text{MAC}} = (\text{Gen}, \text{Tag}, \text{Vrfy})$ be a MAC scheme with a tag space $\mathcal{T} = \{0, 1\}^\ell$. Let $\mathcal{A}$ be an EUF-CMA adversary who makes at most $q = q(\lambda)$ oracle calls to OTag and at most $q_v = q_v(\lambda)$ verification attempts. Let $\text{PRF} : \mathcal{K} \times \mathcal{M} \mapsto \mathcal{T}$ be an underlying pseudorandom function used to construct the MAC. Then there exists a PPT adversary $\mathcal{B}$ against PRF, such that the success probability of $\mathcal{A}$ is bounded by

$$\text{Adv}_{\text{MAC}}^{\text{EUF-CMA}}(\mathcal{A}) \leq \text{Adv}_{\text{F}}^{\text{PRF}}(\mathcal{B}) + \frac{q_v}{2^\ell}$$

Proof Sketch. The proof proceeds via a standard hybrid argument. We first replace the PRF with a random function R . Then the distinguishability between the two experiments is bounded by $\text{Adv}_{\text{PRF}}(\mathcal{B})$ for some adversary $\mathcal{B}$. In the modified experiment, tags on fresh messages are uniform and random; hence, the probability that $\mathcal{A}$'s guess tag^* matches $R(m^*)$ is exactly the probability of guessing an ℓ -bit tag, i.e., at most $2^{-\ell}$. Applying a Union Bound over at most q_v messages, the success probability is at most $\frac{q_v}{2^\ell}$. $\square$

In short, Theorem 1 says that the advantage of any EUF-CMA adversary is bounded by the sum of the security of the underlying PRF and the probabilities of guessing a correct ℓ -bit tag across q_v verification attempts. The q -tagging queries contribute no additional terms since PRF outputs are pseudorandom and leak no useful information. Hence, it makes sense to rate-limit verification attempts (minimize q_v) and use large tag sizes ($\ell \geq 128$). This bound holds for PRF-based deterministic MACs where tagging

queries leak no useful information. The standard EUF-CMA game allows for a single forgery attempt; we extend Theorem 1 to q_v verification attempts to capture a practical scenario where $\mathcal{A}$ repeatedly queries the verification interface.

Key secrecy is assumed in the EUF-CMA definition, and attacks such as brute-forcing the key are accounted for within the security of the underlying primitive (typically bounded by a term negligible in λ), and not considered separately in concrete security. However, at an implementation level, using the same MAC key for a large number of messages may increase the overall success probability of $\mathcal{A}$ by allowing more forgery attempts due to accumulation of q_v over the lifetime of the key. Rekeying limits the number of interactions per key, ensuring $\frac{q_v}{2^\ell}$ remains negligible.

Remark. Theorem 1 establishes an EUF-CMA bound for PRF-based MACs, with an advantage against the underlying primitive as $\text{Adv}_{\text{F}}^{\text{PRF}}(\mathcal{B})$ and an irreducible tag-guessing probability of $\frac{1}{2^\ell}$, where ℓ is the tag length in bits. The remaining classifications incur additional paradigm-specific bounds, which we detail in Section 6.2 and Appendix B.

4.3. Replay Protection and Nonces

MACs do not provide security against *replay attack*, where an adversary intercepts and records (m, tag) pairs and later reuses them. Such replayed (m, tag) pairs would be valid given they were produced by a valid sender. More specifically, the EUF-CMA adversary in the unforgeability game cannot output a pair it has already queried before (cf. Definition 1) whereas a replay adversary reuses what the oracle already gave, and hence does not violate the EUF-CMA winning condition.

The general approach to protect against replay attack is to integrate a nonce (number used once) into the MAC scheme as follows: the tag generation algorithm Tag takes a nonce n as input, and computes $\text{tag} \leftarrow \text{Tag}(k, n, m)$, where n is a public value, and the verification algorithm simply computes $\text{Vrfy}(k, n, m, \text{tag})$ to check its validity. Correctness states that $\text{Vrfy}(k, n, m, \text{tag}(n, m))$ always outputs 1. EUF-CMA security requires that no PPT adversary can produce a new tuple (n^*, m^*, tag^*) that correctly verifies, where (n^*, m^*) was never queried before. Such an adversary is called *nonce-respecting* and requires using a fresh nonce in every query as nonce-reuse may leak structure and break the EUF-CMA security (e.g., Carter-Wegman (CW)-style MAC constructions [31]). Another class called *nonce-misuse* adversary provides strictly stronger security by allowing nonce reuse (e.g., EWCDM MAC [32]). Uniqueness of nonces is usually ensured at the protocol level using counters, timestamps, pseudorandom values, and hybrid approaches. We note that our MAC analyses use EUF-CMA for both classes of adversaries—nonce-respecting ones cannot repeat (m, tag) , while a nonce-misuse can repeat only the nonce in queries. The winning condition (m^*, tag^*) is still the same in both cases, i.e., EUF-CMA never allows querying a previously queried pair.

Protocols designed for constrained environments typically favor counters over pseudorandom values or times-

tamps as nonces. This preference is motivated by several factors: the limited entropy available to constrained devices performing repetitive tasks [33]; the larger nonce fields required for random nonces to achieve a comparable collision probability [34]; the relatively high computational cost of invoking Pseudorandom Number Generators (PRNGs) on constrained devices [33]; and the lower complexity of tracking nonce reuse with counters [35]. Counter values that fall below the expected range indicate replayed messages, while only a short history of previously received counter values needs to be retained to accommodate message reordering.

Over reliable channels, the counter can be maintained *implicitly* by both parties and omitted from transmissions. Over unreliable channels, Luk *et al.* [35] proposed transmitting only the n least-significant bits of the counter, while the upper bits of the full counter are incremented upon overflow. This approach guarantees synchronization as long as fewer than 2^n consecutive messages are lost. Similar approaches are used in, *e.g.*, DTLS 1.3 [17] and LoRaWAN [15].

The nonce size plays a critical role in the concrete security analysis of nonce-based MACs. For a randomly sampled r -bit nonce and q generated messages, collisions (repeating nonce) occur with a probability of approximately $\frac{q^2}{2^r}$, by the birthday bound. Counter-based nonces wrap around after 2^r messages, violating the freshness assumption. Furthermore, reusing the same key-nonce pair across two different messages may compromise EUF-CMA security. Rekeying limits the number of messages authenticated per key, thereby bounding both nonce usage and forgery attempts.

4.4. Truncation & Tag Length Recommendations

The length of MAC tags bounds the maximum achievable security levels (*cf.* Section 4.2), as short tags are more likely to be guessed at random. Meanwhile, many constrained protocols, such as LoRaWAN or AUTOSAR SecOC, employ truncation to reduce MACs overhead. Therefore, a tag is first fully computed and then partially discarded, a possibility for deterministic MAC schemes that rely on recomputing tags for verification. Truncated tags thus incur a proportional loss of security [36]. LoRaWAN, for example, truncates tags to only 32 bit to avoid overburdening the limited bandwidth [15]. Thus, a valid tag can be guessed with probability of 1 in $2^{32} = 4.2e9$, compared to 1 in $2^{128} = 3.4e38$ for strong 128-bit tags. In general, a randomly guessed tag has a probability of 1 in 2^ℓ of being validated by the verification algorithm for an ℓ -bit long tag [37]. Consequently, tags can be truncated to a desired length with a controlled drop in the achieved security level.

To determine adequate security levels, key cybersecurity agencies provide recommendations for tag lengths in practice that we summarize in Table 2. Modern protocols such as TLS use 128 bit tags and thus fall well within the recommendations [43]. Even the newly standardized lightweight Authenticated Encryption with Associated Data (AEAD) cipher ASCON uses a default authentication tag length of 128 bit but supports truncation to a minimum of 32 bit [40]. Such support for short tags is not uncommon in constrained

Agency	Year	Recommended	Acceptable*	Source
ANSSI	2026	≥ 128 bit	—	[38]
BSI	2025	≥ 96 bit	< 96 bit	[39]
NIST	2025	≥ 64 bit	≥ 32 bit	[37], [40], [41]
ENISA	2014	128 – 256 bit	≥ 80 bit	[42]

* for constrained/legacy systems, after carefully evaluating the impact

TABLE 2: Tag lengths recommendations by key cybersecurity agencies.

environments. Still, they fall outside the recommended tag lengths. They should only be used after careful consideration of the resulting security impact in a given scenario, as done, for example, for Galileo’s Open Service Navigation Message Authentication (OSNMA) [36].

5. MACs in Constrained Environments

MACs require cryptographic processing and yield relatively long authentication tags (often 128 bit). These two properties make deploying message authentication in constrained environments challenging, where processing resources are scarce, and bandwidth must be conserved. We thus survey the landscape of protocols that *provide* message authentication to understand their practical implications.

5.1. Protocol Survey

We derive the set of protocols considered in this survey from the constrained environments introduced in Section 2. For each environment, we include either the previously given example protocols, if they support message authentication, or the corresponding security extension, if it exists. Table 3 summarizes the resulting protocol landscape. For each protocol, we analyze the MAC mechanisms employed, whether authentication is combined with encryption, the amount of authenticated data per Protocol Data Unit (PDU), tag and nonce lengths, and whether the protocol supports group communication. Our analysis enables a cross-domain comparison of how message authentication is performed in practice across constrained environments.

5.2. Research Challenges

Our analysis reveals several recurring design patterns and limitations across domains. Despite substantially different deployment settings, many protocols converge on similar cryptographic constructions and parameter choices. From these observations, we highlight key gaps between current practice and the requirements of constrained environments.

Across domains, most protocols rely on a few well-established MAC schemes such as CMAC and HMAC. These constructions provide mature security guarantees, but they were not specifically designed for constrained environments. Nevertheless, they remain prevalent even in environments where cryptographic processing overhead is a primary concern. For example, 5G makes security processing optional

Protocols	Domain	MAC Schemes	MAC Input [byte]	Tag [bit]	Nonce [bit] (trans.)	Encryption	Group Key
802.15.6 [44]		CBC-MAC-AES	16–272	32	104 (8)	✓	✓
Zigbee [45]		CBC-MAC-AES	5–123	32–128	104 (40)	✓	✓
BLE [46]		CBC-MAC-AES	2–252	32	104 (0)	(X)	✓
AUTOSAR SecOC (CAN) [47]		CMAC-AES	10–17	24	0–64 (0–8)	X	✓
AUTOSAR SecOC (CAN-FD) [47]		CMAC-AES	10–73	24	0–64 (0–8)	X	✓
OMS (excl. Mode 13) [48]		CMAC-AES CBC-MAC-AES	6–2 ¹⁴	64	0–104 (0–96)	(X)	X
Matter [49]		CBC-MAC-AES	9–2 ¹⁶	128	104 (40–104)	✓	✓
802.15.4 [50]		CBC-MAC-AES	3–121	32–128	104 (0–104)	(X)	✓
5G [51]		SNOW 3G, CMAC-AES, ZUC	> 4	32	32 (12–18)	✓	X
DTLS 1.3 [17]		GMAC-AES, CBC-MAC-AES, Poly1305-ChaCha20	≥ 32	64–128	48 (8–16)	✓	X
MIOTY [52]		CMAC-AES	10–150	8–48	32 (24)	✓	✓
LoRaWAN [15]		CMAC-AES	8–251	32	32 (16)	✓	✓
Sigfox [53]		CBC-MAC-AES	6–18	16–40	12	(X)	X
IEC 62351 (GOOSE) [8]		HMAC-SHA256, GMAC-AES	≥ 68	128	–	X	✓
IEC 62351 (SMV) [8]		HMAC-SHA256, GMAC-AES	≥ 56	128	–	X	✓
OSNMA [54], [55]		HMAC-SHA256, CMAC-AES	18–69	20–40	*	X	X

(X): optional encryption *no traditional nonce, but short key validity plus precise timestamps

TABLE 3: Protocols in constrained environments rarely adopt optimized MAC schemes and often employ shorter tags than recommended without properly considering the impact of these deployments.

due to potential overhead [51] and IEC 62351 explicitly discourages the use of encryption in certain power-grid deployments solely due to processing and latency concerns [8]. In Section 6.5, we show that hardware acceleration alone is often insufficient to offset the computational costs of these constructions compared with more lightweight alternatives. Meanwhile, protocol standards rarely adopt MAC schemes tailored to the characteristics of constrained environments, *e.g.*, small payloads or low-frequency rekeying, likely due to a lack of standards, insufficient understanding of the potential benefits, and limited awareness of the associated security implications.

Insight

Protocols mostly rely on generic MACs rather than schemes tailored to constrained deployments.

Many surveyed protocols employ authentication tags that are significantly shorter than the commonly recommended security margins (*cf.* Table 2). Although application-specific risk assessments are recommended when selecting shortened tags, we find little evidence that such analyses are systematically performed in practice. A notable exception is the work of Fernandez *et al.* on GNSS authentication [36],

which derives suitable tag lengths of 28–40 bit based on concrete threat assumptions. Consequently, deployments may (unknowingly) operate with insufficient security margins, leading to layered and thus inefficient protocol stacks, such as (D)TLS over LoRaWAN, to satisfy security requirements [56], [57].

Insight

Many protocols employ authentication tags below commonly recommended security margins without domain-specific analysis of resulting risks.

Moreover, many protocols (recommended to) silently discard packets with invalid tags, such as DTLS 1.3 [17], AUTOSAR SecOC [47], or 802.15.4 [50]. While these events are usually logged (in DTLS 1.3, they “MAY” be logged [17]), the logs of embedded devices are not necessarily closely monitored. Thus, an attacker may make many forgery attempts during a single session without detection.

Insight

Protocols often silently discard invalid MAC tags to avoid expensive rekeying procedures.

Many protocols operate in multicast or broadcast settings

in which multiple devices share a common authentication key. Such deployments enable efficient group communication but inherently weaken source authentication guarantees: once a single group member is compromised, that member can impersonate any other participant in the group. Some systems explicitly address this problem. For example, Galileo’s OSNMA employs TESLA-based delayed-disclosure mechanisms to mitigate the risk of forgery in broadcast authentication [54]. However, comparable protections are absent in most other protocols. Meanwhile, replacing symmetric authentication with digital signatures that offer source authentication is often impractical due to bandwidth, latency, and computational overhead constraints.

Insight

Group key-based authentication remains widely used in constrained multicast deployments despite its inherent impersonation risks. The security implications of such designs and alternative efficient solutions are not yet sufficiently studied.

Beyond these challenges, we make several broader observations regarding the operational characteristics of message authentication in constrained environments. First, small payload sizes reduce the number of primitive invocations required by constructions operating on fixed block sizes. This could enable specialized implementation optimizations or constructions for constrained systems.

Insight

Authenticated inputs in constrained protocols are typically small, often well below 100 byte, even when protocol headers and nonces are included.

Second, several protocols intentionally omit encryption to reduce latency and processing overhead or to maintain compatibility with legacy devices that do not support cryptographic processing. Consequently, research on standalone MAC constructions remains relevant despite the strong recent focus on (lightweight) AEAD schemes.

Insight

Message authentication is often deployed independently of encryption in constrained environments.

Third, many protocols transmit only partial nonces to protect against replay. While these mechanisms are generally sufficient for expected communication patterns, their security critically depends on appropriate rekeying intervals and careful management of counter exhaustion.

Insight

Constrained protocols commonly employ compact nonce and replay-protection mechanisms to minimize communication overhead.

Overall, our protocol survey identifies a recurring mismatch between the constraints faced by real-world deployments and the authentication mechanisms standardized for them. In particular, the continued reliance on generic MAC constructions, aggressive tag truncation, and symmet-

ric group authentication motivates the research directions explored in the remainder of this paper.

6. Optimizing Processing Times

As traditional MAC schemes such as HMAC-SHA256 often incur significant processing overhead, research has proposed a range of more efficient schemes. In the following, we present an overview of these optimized schemes identified through a systematic literature survey.

6.1. Systematic MAC Scheme Identification

To fairly represent the landscape of MAC schemes for constrained environments, we conducted an iterative systematic literature review (SLR) on January 6th, 2026. We start from an initial publication and iteratively explore the references and citations of all included publications. We include publications that explicitly design MAC schemes for constrained environments, identifying them by examining their titles and, where necessary, their abstracts and introductions. We include publications that focus on MAC scheme design and explicitly mention constrained environments. The identification of citing and cited papers is automated through the Semantic Scholar API. We choose this method because it provides broad coverage of the relevant literature while avoiding the ambiguity of the term MAC (which can refer to either Message Authentication Code or Medium Access Control) in keyword-based searches.

As a starting point, we select the popular Poly1305 [58] scheme, arguably one of the most prominent lightweight MAC schemes, owing to its integration into, *e.g.*, DTLS 1.3 [17]. The SLR converges after five iterations, in the sense that no additional qualifying publications are identified through further exploration of the references and citations of the included publications. Moreover, starting the review with any publication in the resulting set would yield the same set of publications, although the number of iterations required to reach convergence may differ. This convergence, however, is limited to the citation-connected literature reachable from the selected seed. We cannot rule out the existence of disconnected clusters of relevant publications that are not directly linked to the identified literature; however, we are not aware of any such publications. In total, we identify 20 MAC schemes across 22 publications, after examining 1 402 publications.

6.2. MAC Schemes

The 20 MAC schemes identified by our survey are summarized in Table 4, classified by the underlying construction: Universal Hash-based MACs (U), Block Cipher-based MAC (C), PRF-based MACs (P), and Permutation/Sponge-based MACs (S). A short description and EUF-CMA security analysis of all schemes are provided in Appendix B. Several observations relevant to MAC security in constrained platforms emerge from the analysis.

MAC Schemes	Year	Code	Standard	Const.
Badger [59]	2005	✗	✗	
Poly1305 [58]	2005	[60]	RFC 8439 [61]	
UMAC [62]	2006	[63]	RFC 4418 [64]	
VMAC [65]	2006	[66]	✗	
ShMAC [67], [68]	2009	✗	✗	
TuLP [69], [70]	2009	✗	✗	
SipHash ¹ [71]	2012	[72]	✗	
Auth256 [73]	2014	✗	✗	
Chaskey [74]	2014	[75]	ISO/IEC 29192-6 [76]	
LightMAC [77]	2016	[78]	ISO/IEC 29192-6 [76]	
MergeMAC ² [79]	2018	✗	✗	
Mac611 [25]	2019	[80]	✗	
LedMAC [81]	2021	✗	✗	
LDMAC [82]	2022	[83]	✗	
BP-MAC [84]	2022	[85]	✗	
ElIMAC [86]	2023	[87]	✗	
PetitMAC [88]	2024	[89]	✗	
LeMAC [88]	2024	[89]	✗	
ASCON-MAC [90]	2024	[91]	✗	
SMAC [92]	2025	[93]	✗	

¹ Only computes 64 bit long tags ² broken MAC scheme [94]

TABLE 4: Various lightweight MAC schemes, optimized for constrained environments, have been proposed.

The surveyed schemes can be classified as deterministic or randomized (nonce-based) in tag generation. In deterministic MACs, the tag is a fixed function of a secret key and message. In nonce-based MACs, tag generation additionally requires a nonce, and EUF-CMA security is defined with respect to a nonce-respecting or nonce-misuse adversary. Both paradigms present complementary security challenges in constrained environments. For deterministic MACs, security bounds typically degrade with the total number of authenticated messages under a fixed key. Constrained platforms often involve high message volumes—and every message authenticated with a fixed key is a potential oracle query from an EUF-CMA adversary. Thus, as q grows, the birthday term $\frac{q^2}{2^b}$ grows, where b is the internal state or block size of the underlying primitive. Independently, as q_v grows, repeated verification attempts without rate-limiting increase the tag-guessing probability of $\frac{q_v}{2^t}$. Rekeying limits the number of messages that can be authenticated under a single key, thereby directly bounding q .

Nevertheless, secure key distribution may be costly in constrained settings. Rate limiting verification attempts is additionally necessary to bound q_v . In contrast, nonce-based MACs shift part of the security burden to nonce management. Constrained devices often have limited storage and

weak entropy sources, increasing the risk of nonce reuse, which in turn may collapse/degrade the security of CW-style MACs. Intermittent connectivity further complicates nonce synchronization between a sender and a receiver.

Beyond the above distinction, the EUF-CMA analyses of the surveyed schemes exhibit a similar reduction: the adversary’s forgery advantage is expressed as the underlying primitive-specific failure event, together with the unavoidable tag-guessing probability. Universal Hash-based MACs reduce to the PRP advantage of the masking block cipher with a PRP-PRF switching term, a hash collision term from the AU/AXU property of the hash family. Block Cipher-based MACs reduce to the PRP security of the underlying block cipher. They may incur two additive birthday terms $\frac{q^2}{2^b}$ due to the PRP-PRF switching lemma and internal-state collisions, where b is the block size of the underlying cipher. PRF-based MACs reduce to PRF hardness with no additional terms unless the PRF is instantiated with a block cipher, in which case a PRP-PRF switching term is introduced. Permutation/Sponge-based constructions reduce to the inability of a PPT adversary to find collisions in the hidden internal state. The security bounds, therefore, include a birthday term proportional to the size of the state. Furthermore, round-reduced constructions such as TuLP and SMAC require resistance against differential cryptanalysis. Their security analysis, therefore, includes the probability of differential trails leading to collisions in the internal state.

From a security perspective, different environments stress different components of the EUF-CMA bound. For example, in Universal Hash-based MACs, the hash collision term may become significant in high-throughput environments. Hence, practical security may depend substantially on the universality parameter of the hash family, rather than primarily on the underlying block cipher. Similarly, if a tag is heavily truncated, then the tag-guessing probability can dominate over all terms in the EUF-CMA bound.

Insight

The security of a MAC scheme is limited by the deployment assumption that is most likely to fail in practice. A tight EUF-CMA bound offers little assurance if the dominant failure mechanism—primitive distinguishing, hash collisions, hidden-state collisions, nonce reuse, or tag guessing—is not adequately addressed in the deployment.

6.3. MAC Security in Constrained Deployments

The constrained platforms surveyed in Section 2 exhibit several deployment constraints that directly affect MAC security. Across these environments, we identify *six* recurring constraints: (i) *computational constraints and resulting latency constraints*—limited processing power restricts the use of computationally expensive algorithms (ii) *memory constraints*—limited memory capacity (e.g., few kilobyte) affects key storage, state size, lookup tables, (iii) *communication constraints*—low bandwidth channels impose strict limits on tag length, nonce transmission, (iv) *state*

management constraints—increases risk of nonce-reuse due to power cycles and resets, (v) *energy constraints*—nodes operating on battery/harvested energy have strictly limited budget which allows few primitive evaluations per authenticated byte, and (vi) *key-management constraints*—nodes accumulate large query counts q and verification attempts q_v under a fixed key affecting the EUF-CMA bound.

While the EUF-CMA bound characterizes forgery resistance under idealized assumptions, different deployment constraints amplify different components of the bound, favoring distinct MAC designs. Computational and memory constraints primarily affect the choice of primitives underlying the security term $\text{Adv}_{\text{MAC}}^{\text{PRF/PRP}}(\mathcal{B})$. With less processing power, round-reduced constructions such as TuLP and SMAC are employed, replacing formal PRP/PRF reductions with heuristic differential cryptanalysis. With limited memory, primitives with large key scheduling are replaced with compact-state primitives, which may offer weaker or heuristic security guarantees. Energy and key-management constraints primarily stress the birthday and tag-guessing components of the EUF-CMA bound. As the number of authenticated messages q grows, the birthday term increases. Similarly, q_v grows without rate-limiting, increasing the tag-guessing probability. Communication constraints directly affect the tag-guessing term: bandwidth limitations necessitate tag truncation, reducing ℓ and increasing $\frac{q_v}{2^\ell}$. In extreme scenarios, such as underwater acoustic and satellite navigation, tag guessing may dominantly contribute to the security bound. State-management constraints affect the security of nonce-based MACs, including all CW-style MACs except for EWCDM variants such as Mac611.

Having identified the constraint categories and their security implications, we now discuss the MAC paradigms best suited to each. *Universal hash-based MACs* excel in high-throughput environments, as the underlying universal hash layer is highly parallelizable and efficiently handles large message volumes. However, the hash collision term grows with query count, requiring careful bound management in long-lived deployments. They are thus attractive for ICSSs and smart city infrastructure. *Block cipher-based MACs* are well suited for environments with short, fixed-length messages and existing block-cipher hardware, making them natural choices for intra-vehicular networks, power grids, and aviation. In long-lived deployments with infrequent rekeying, the birthday term becomes dominant. *PRF-based MACs* are well-suited to computationally and memory-constrained environments processing short messages, owing to their minimal additional structure. *Permutation/sponge-based MACs* are well-suited to hardware-constrained platforms, benefiting from the absence of key scheduling and from permutation reuse, as in BANs.

Insight

Deployment constraints determine which term of the EUF-CMA bound dominates in practice. Consequently, MAC selection depends upon the dominant constraints of the target constrained platform.

AEAD Schemes	Year	Code	Standard
ASCON [97], [98]	2014	[91]	NIST SP 800-232 [40]
ChaCha20-Poly1305	2018		RFC 8439 [61]

TABLE 5: Finalist of NIST Lightweight Cryptography Challenge for AE schemes.

6.4. Performance Evaluation Methodology

To evaluate the practical efficiency of MAC schemes, we here focus on processing time, a key metric alongside other factors such as memory usage and code size. Unlike existing evaluation frameworks such as FELICS [95], which emphasize block and stream ciphers, we assess standalone MAC schemes for constrained environments. Our benchmark uses the Zolertia RE-Mote (ARM Cortex-M3 @ 32 MHz, 16 kB RAM) and LP-MSP430FR2476 (MSP430 @ 16 MHz, 8 kB RAM) platforms to measure the average tag generation times for the 11 MAC schemes with publicly available code¹. We additionally evaluate two AEAD schemes (*cf.* Table 5). ASCON is the winner of the NIST AEAD Cipher competition and ChaCha20-Poly1305 is deployed in (D)TLS 1.3. For these AEAD ciphers, we only use the associated data and do not encrypt any data. Finally, we include hardware implementations of HMAC-SHA256 and CMAC-AES on the Zolertia RE-Mote, the most commonly used MAC schemes in constrained environments (*cf.* Table 3). While GMAC-AES is also commonly used in constrained environments, GHASH, its underlying hash function, can be inefficient on constrained processors because they often lack hardware support for carry-less multiplication [96]. Where available, we use implementations optimized for the given architecture.

6.5. Performance Evaluation Results

Fig. 1 presents the results of our evaluation for both platforms. The x-axis shows the input length passed to the $\text{Tag}(k, m)$ algorithm of each MAC scheme; note that $\text{Vrfy}(k, m, \text{tag})$ requires only a single additional comparison. For each input length, we compute 1 000 tags and report the average runtime as the number of clock cycles required to process one byte.

We first examine hardware-accelerated MAC implementations, as highlighted in Fig. 1c. Recent embedded chips provide hardware support for CMAC-AES and HMAC-SHA256, which are allegedly capable of outperforming software implementations. We confirm that hardware acceleration yields substantial speedups over corresponding software implementations; nevertheless, neither CMAC-AES nor HMAC-SHA256 match the fastest software-only schemes evaluated here. Notably, CMAC-AES performs well on short payloads, whereas HMAC-SHA256 becomes competitive only for longer inputs, lengths that frequently

1. No code to run PetitMAC and LeMAC on embedded hardware was made available, even upon request.




Figure 1: Evaluating the processing time of various MAC schemes on incremental small input sizes on two constrained processors reveals that (i) software implementation of lightweight schemes can outperform hardware-accelerated HMAC and CMAC, (ii) schemes scale differently, and (iii) AEAD ciphers add significant overhead compared to standalone schemes.

exceed those observed in constrained environments (*cf.* Section 2). Hardware acceleration nonetheless offers practical benefits, reducing both the program memory footprint and the compute load, even when latency is higher.

Insight Software implementations of lightweight MAC schemes can outperform hardware-accelerated traditional schemes, particularly for short inputs.

A second observation is highlighted in Fig. 1d. Not only do MAC schemes scale differently when looking at hardware implementations. The most striking example is BP-MAC, which is explicitly designed for short inputs: it leads the field for inputs of a few byte but approaches its performance ceiling even at 50 byte inputs. Similarly, ASCON-MAC is outperformed by multiple schemes as inputs grow.

Insight MAC schemes do not scale uniformly with input length; the optimal choice depends strongly on the expected message sizes in the target deployment.

Unlike established MAC evaluation practices, which re-

port performance at a handful of discrete input lengths [86], [90], [92], we evaluate every input length in our range. Fig. 1e shows why this should be standard practice: selecting representative lengths that coincide with full processing blocks omits the padding overhead incurred at most other lengths. For example, on the ARM Cortex M3, SipHash is outperformed by Chaskey by only 4.3 % for inputs between 17 and 22 byte, yet by an average of 21.4 % and 19.4 % for inputs five byte shorter and longer, respectively. Chaskey, meanwhile, outperforms SipHash by 50.6 % at exactly 16 byte. These results illustrate that point measurements can substantially misrepresent the practical performance of MAC schemes.

Insight The high variance in cycles-per-byte, even within a narrow range of input lengths, shows that this metric should be evaluated over a broad range of input lengths when comparing MAC schemes.

Including two widely used AEAD ciphers in our evaluation yields additional insights. Although their combined confidentiality and integrity guarantees are desirable in many settings, they introduce measurable overhead when

only integrity protection is required. As shown in Fig. 1f, ASCON-MAC achieves a 23.9% average reduction in processing time compared to full ASCON, consistent with the results reported by the original authors [90]. This reduction can be explained by standalone MACs requiring less branching than AEAD schemes and by specific optimizations, such as reducing robustness properties [90].

Insight

AEAD ciphers impose significant overhead relative to standalone MAC schemes, even when confidentiality is not employed.

Schemes such as BP-MAC [84] and Poly1305 [58] support *preprocessing*: idle periods between transmissions are exploited to perform preliminary computations, reducing tag generation latency when data must be sent. Such idle periods are common in constrained networks, where devices transmit sensor readings at periodic intervals; yet in latency-critical scenarios, these readings must be forwarded with minimal delay [99]. Preprocessing applies to a broader class of MAC schemes. CW-style constructions, for instance, restrict cryptographic operations to the nonce (typically a predictable counter), which can be computed in advance. Similarly, predictable prefixes in MAC inputs can be exploited for partial precomputation [99].

Insight

Latency-critical deployments frequently exhibit idle periods between transmissions that many MAC schemes can exploit through preprocessing.

Despite being generally ill-advised (*cf.* Section 4.4), many protocols use heavily truncated tags. Truncated tags are typically produced by computing the full-length tag and discarding the excess. Some schemes, *e.g.*, SipHash, BP-MAC, UMAC, can natively generate shorter tags, reducing computational cost compared to post hoc truncation.

Insight

When short MAC tags are used, schemes that natively produce truncated outputs can reduce computational overhead relative to truncating a full tag.

The processing platform, potential hardware accelerators, expected input sizes, desired tag lengths, and the availability of idle time for preprocessing all influence the optimal MAC scheme in terms of processing time, without even considering other potential constraints.

6.6. Beyond Processing Time

While Section 6.5 focuses on authentication tag computation time, bootstrapping time, memory footprint, code size, and energy consumption are also decisive constraints in practice. A fast MAC scheme that exceeds a device's memory budget or drains a battery-powered node faster than alternative schemes may ultimately be less practical.

Code size and memory footprint directly determine whether a scheme can be accommodated within the often-limited resources of constrained devices, independent of

its processing speed. Constructions such as BP-MAC, for example, trade memory for speed by precomputing a $2^{|m|}$ -entry lookup table, thereby reducing tag generation to $|m|$ XOR operations. Consequently, BP-MAC is only practical for devices that can accommodate the required memory overhead, which restricts deployments to short messages. Such trade-offs are rarely apparent from processing-time measurements alone, yet they can have a decisive influence on MAC scheme selection in practice. Moreover, execution speed, code size, and memory footprint are often subject to competing implementation trade-offs.

Energy consumption is likewise an important consideration for MAC schemes. Although energy spent on wireless transmissions and receptions can constitute a major part of the energy budget in many wireless deployments [14], cryptographic processing is not negligible. It can also be relevant in wired communication. Quantifying the energy consumed by cryptographic processing is therefore important for assessing the deployability of a MAC scheme. Processing cycles provide only a proxy for energy consumption, as the relationship depends, among other factors, on the target platform's sleep behavior, memory footprint, and use of hardware acceleration.

Despite the importance of these dimensions, hardly any existing work analyzes them in depth (*cf.* Section 3). A comprehensive evaluation across these dimensions is therefore essential for informed MAC scheme selection and future deployment planning. They enable designers to make optimal use of limited resources in constrained environments while balancing security guarantees, processing speed, device capabilities, deployment costs, and device lifetime.

7. Optimizing Bandwidth Usage

Beyond processing overhead, tag sizes pose a significant burden in constrained environments: tags typically require 16 byte to achieve adequate security, often exceeding the payload itself (*cf.* Table 2). Simple truncation reduces security in proportion (*cf.* Section 4.4), motivating two decades of research into MAC aggregation [100], [101], [102].

7.1. The Idea of MAC Aggregation

MAC aggregation reduces authentication overhead without sacrificing security by jointly authenticating multiple messages [100], [101], [102]. Given two tags $\text{tag}_1 \leftarrow \text{Tag}(k_1, m_1)$ and $\text{tag}_2 \leftarrow \text{Tag}(k_2, m_2)$ of deterministic MAC schemes, an aggregated, *e.g.*, via XOR, tag ($\text{tag} \leftarrow \text{tag}_1 \oplus \text{tag}_2$) allows a verifier knowing k_1 and k_2 (a single key for both messages may be chosen) to jointly verify both messages, as any modification to either causes verification failure. For ideal MAC schemes with independent tags, this incurs no security loss over individual tags [102]. The limitation is that if any message is lost or corrupted, verification fails. Aggregation was therefore initially studied only for reliable channels, where it offers a clean bandwidth-latency trade-off [103], [104], [105], [106]. However, as constrained communication increasingly migrates to lossy

(wireless) stacks [107], message loss becomes a primary concern largely neglected in current evaluations. Today, MAC aggregation is thus arguably more promising for lossy communication channels than for the reliable channels it was originally designed for.

7.2. Challenges in Lossy Environments

Deploying MAC aggregation in lossy environments introduces several challenges beyond the processing and memory overhead of tracking and aggregating messages, as the resulting improvements depend heavily on the underlying packet loss. The key metric is *authenticated goodput*, i.e., fully authenticated payload byte relative to total byte sent, which depends not only on the overall Packet Error Rate (PER) but also on the packet-loss distribution and the scheme parameterization, e.g., the number of jointly authenticated messages. Higher loss rates and longer payloads reduce effectiveness; for short payloads, meaningful gains begin to emerge at PERs below 10%, well in line with practical deployments [108]. Aggregation scheme selection and parameterization are critical, and dynamic adaptation is preferable when channel conditions shift due to mobility or weather. Additional deployment considerations include verification delays and jitter, as well as increased susceptibility to jamming attacks [108].

Insight

MAC aggregation applies to lossy channels, but scheme selection and parameterization must account for channel quality and security goals.

7.3. Defining MAC Aggregation Schemes

While MAC aggregation is conceptually simple, its design space is extensive. We introduce a formal definition that captures and modularizes the different constructions in this area. A MAC aggregation scheme extends the traditional MAC definition by incorporating multiple messages $M \subset \mathcal{M}$ into the $\text{Tag}(k, M)$ and $\text{Vrfy}(k, M, \text{tag})$ algorithms. Two core properties characterize any such scheme: the *dependencies* that define M for a given message m , and the *MAC computation* based on multiple messages rather than a single one. We detail each in turn.

7.3.1. Dependencies. Tag generation in MAC aggregation schemes depends on a set of messages M . We can describe M_i for a message m_i relative to the current message m as a *dependency set* $\mathcal{D}$ [109]:

$$\mathcal{D} = \{k \mid m_{i-k} \in M_i\}.$$

With this notation, non-aggregating MAC schemes have $\mathcal{D} = \{0\}$, as only the current message is an input to the Tag and Vrfy functions. Overall, a dependency set $\mathcal{D}$ means that to verify a tag all messages $m_{i-k}, k \in \mathcal{D}$ must be received. Likewise, the verification of all tags $\text{tag}_{i+k}, k \in \mathcal{D}$ is necessary to fully validate the message m_i .

Different proposed MAC aggregation schemes use different dependency sets. Most commonly, a set of consecutive messages is used [110], [111], [112], [113], [114]. While these sets reduce the memory needed to track the history of prior relevant messages, they also become vulnerable to message loss and selective jamming attacks that amplify Denial of Service (DoS) attacks [109]. As a remedy, the R2D2 scheme identifies optimal dependency sets to distribute the effect of message loss and also protect against selective jammers [109]. However, this results in noticeably higher verification delays [108]. As an alternative strategy to thwart selective jamming, 2D MAC computes multiple tags over different dependency sets and utilizes superposition coding for efficient encoding [115]. Thus, the selection of the dependency set is crucial to strike a balance between resilience to packet loss and authentication delays.

7.3.2. MAC Aggregation Algorithm. Beyond identifying which messages to aggregate, a scheme must define how to combine them into a single tag. Two approaches are common: bitwise XOR and hash-based accumulators [116].

Bitwise XOR. Most proposed schemes use an XOR-based approach [109], [111], [112], [113], [114], [115]. This approach first computes an ordinary authentication tag for each message. The tags for the messages in the dependency set are then either partially or fully aggregated by XORing them. The inclusion of a nonce with each tag is crucial to prevent not only replay attacks but also mix-and-match attacks within a set of jointly authenticated messages [103]. If individual tags are pseudorandom, e.g., HMAC [117], an attacker would need to know every single ordinary tag to compute a valid aggregated tag [118]. However, this aggregation approach may not be secure, e.g., for Universal Hash-based MACs, which are often more performant in constrained environments.

Hash-based Accumulators. To circumvent this problem, an alternative is to compute the aggregated authentication tag directly from the messages [106], [110], [113]. For such an aggregation, the communicating entities need to keep a history of all messages still to be used for computing aggregated tags, which can require significantly more memory. As a benefit, the resulting aggregation fulfills stronger security properties [110], [113].

7.4. MAC Aggregation in Practice

Despite nearly two decades since its introduction, MAC aggregation attracted limited interest as wired network bandwidth scaled faster than constraints, making authentication overhead negligible. Efforts instead focused on optimizing MAC computation (cf. Section 6). The rise of long-range and high-frequency wireless communication has changed this: a 128-bit tag now occupies a significant fraction of a frame, driving the heavy truncation seen in deployed protocols (cf. Section 5) at the cost of security.

Recent work has renewed interest: Graphene [116] and Diamond [119] offer flexible AEAD schemes with aggregation support. Simulations and dataset-driven evalua-

tions also highlight significant potential savings in wireless environments if parameterized correctly [108]. Backward-compatible integration into DTLS 1.3 yields goodput improvements of up to 50 % and energy savings of up to 17 % [16], with 802.15.4 and Wi-Fi integrations showing energy savings of 21–37 % for short payloads [110]. Early prototypical deployments thus demonstrate clear potential, though large-scale validation is still needed.

Insight

Early deployments demonstrate real-world gains by MAC aggregation in goodput and energy efficiency, but large-scale real-world validation is still needed.

7.5. Progressive Authentication

MAC aggregation achieves strong security with short tags but introduces verification delays, conflicting with the low-latency requirements of environments such as ICSs (*cf.* Section 2). Progressive MACs (ProMACs) [110], [111], [114] address this via *optimistic security* [120]: security processing is deferred under the optimistic assumption that attacks are absent, with retrospective detection triggering appropriate action (*e.g.*, emergency shutdown). Unlike Intrusion Detection Systems (IDSs), ProMACs offer faster, more reliable attack detection with no false alarms.

Insight

ProMACs bring the bandwidth savings of MAC aggregation to low-latency scenarios.

ProMACs have been proposed for drone control [110], [121], [122], haptic feedback [110], [123], tactile networks [110], [124], IEEE 802.11a [110], Intel SGX memory encryption [110], [125], and network-on-chip communication [110], [126]. Related work on optimistic security contextualizes acceptable authentication delays [127], [128], [129], [130], [131]. As such, delays of up to 100 messages are potentially acceptable for intra-vehicular throttle control, given high sampling rates and physical inertia [127], [128]; and delays of up to 16 messages can easily be recovered from [129]. Likewise, reaction times of 60 ms (corresponding to 100 subsequent packets) are proposed for some ICS settings [130]. Furthermore, most recent MAC aggregation schemes can serve as ProMAC instantiations [109], [110], [111], [112], [113], [114], [115], [122]. Yet despite this breadth, we found no evidence of (prototypical) deployments of ProMACs for productive systems.

Insight

ProMACs have been proposed for several settings, and optimistic security has been discussed broadly. However, concrete scenarios that strictly require ProMAC deployments remain elusive.

8. MACs in Group Communication Systems

Group communication in multicast and broadcast scenarios poses unique authentication challenges: a sender must simultaneously authenticate a single message to multiple

recipients. Digital signatures are the standard solution, but incur significant computational and bandwidth overhead. Protocols in constrained environments, therefore, often rely on shared group keys to reduce cost (*cf.* Section 5), sacrificing source authentication as any group-key holder can forge indistinguishable messages. Thus, we study MAC-based alternatives, but begin with the challenges of deploying digital signatures in constrained environments.

8.1. The Drawbacks of Digital Signatures

Digital signatures are often impractical for M2M communication due to their high computational and communication overhead compared to MACs. MACs require only efficient symmetric-key operations and produce shorter tags. In contrast, signatures incur higher signing and verification costs and yield larger authentication metadata, often by an order of magnitude in both processing cost and transmitted byte [18]. Consequently, many security protocols for constrained environments rely on MACs despite group keys preventing source authentication (*cf.* Section 5), as public verifiability and non-repudiation are often unnecessary.

This gap widens under post-quantum requirements: post-quantum signature schemes introduce substantially larger signatures, public keys, and/or computational costs [132], whereas symmetric primitives often require only moderate parameter increases to maintain security [133]. MAC-based authentication thus remains the practical choice for scalable or resource-constrained group communication.

Insight

Post-quantum digital signatures impose even greater resource demands relative to MACs, making MAC-based authentication increasingly favorable for constrained environments.

8.2. MAC-based Source Authentication Schemes

Various strategies to adapt MACs for source authentication have been proposed. To achieve this additional security property, they require some form of asymmetry between the sender and receivers of messages [134].

8.2.1. Time Asymmetry. The most prominent approach employs *time asymmetry* via delayed key disclosure. In TESLA [135], [136] and its variants [137], [138], the sender authenticates messages over a fixed interval using a key disclosed only after that interval; receivers reject any subsequent messages authenticated with that key. The main challenges are that (i) disclosed keys must be bound to the sender, (ii) receivers must buffer messages until disclosure, and (iii) verification is inherently delayed. Nevertheless, TESLA underpins OSNMA, the message authentication scheme of Galileo [54], with further integrations proposed for Unmanned Aerial Vehicles (UAVs) [139], smart grids [140], and Underwater Acoustic Networks (UANs) [141].

MAC-based source authentication is relevant for many constrained environments that rely on group keys, as digital signatures remain too expensive.

8.2.2. Information Asymmetry. A second class exploits *information asymmetry*, where each receiver can verify only a portion of the multiplexed authentication data [142]. The baseline construction computes one tag per receiver, which scales poorly. Tags can be combined into a single shorter tag while preserving partial per-receiver verifiability [143], [144], or replaced by a single tag with a per-receiver witness [145]. A recent scheme [146] reduces this to two MACs with truncation support, but incurs significant hash-chain overhead and is restricted to very short messages. In general, these schemes quickly exceed the overhead of classical digital signatures as the receiver count grows.

8.2.3. Verifier Asymmetry. CAIBA [147] splits verification between two entities: the sender produces an aggregated tag combining a source-authenticating component and an integrity-protecting component. A dedicated authenticator, sharing a secret key pair with each sender, strips the source-authenticating component, leaving only the group-key-based tag for receivers. CAIBA has been tested on a slow CAN bus, but is limited in its applications as it requires an authenticator that can both observe all traffic and modify messages in transit.

While TESLA's delayed key disclosure has been adopted in Galileo's OSNMA, MAC-based source authentication schemes usually introduce significant limitations compared to group-key MACs in order to boost security. Their applicability thus depends heavily on environment-specific constraints and conditions (*cf.* Section 2).

9. Challenges and Future Directions

This SoK reveals a persistent gap between research and practice in MAC-based authentication for constrained M2M communication. We round off this paper by identifying the most pressing open challenges.

Lightweight Standalone MAC Schemes. Despite the wealth of lightweight MAC proposals identified in our survey, their adoption in deployed protocols remains minimal. Filling this gap requires not only comprehensive performance evaluations across all relevant dimensions (*e.g.*, processing time, latency, memory, code size, and security foundation) but also standardization efforts and accompanying guidance for system designers on how to exploit the characteristics of constrained environments, such as small and predictable payloads.

Understanding Suitable Tag Lengths. Many constrained protocols operate with shorter tags than recommended. There is a clear need for principled, domain-specific analyses, analogous to Fernandez et al.'s work for GNSS [36], that characterize the practical security implications of shortened tags and keys.

MAC Aggregation in Lossy Environments. While early results are promising, the potential of MAC aggregation in practical wireless deployments has not been fully characterized. Open questions include the dynamic adaptation of aggregation parameters to changing channel conditions, the interaction between aggregation schemes and existing protocol stacks, and the evaluation of ProMACs in concrete deployment scenarios that demonstrably benefit from these schemes.

MAC-based Source Authentication. As post-quantum security requirements render digital signatures increasingly impractical in constrained environments, MAC-based source authentication deserves renewed attention. Developing efficient, standards-ready constructions that provide meaningful source authentication guarantees without prohibitive overhead remains an open and practically important problem.

10. Conclusion

This SoK paper provides the first comprehensive analysis of MACs for constrained M2M communication. Through a survey of 9 deployment domains, a systematic literature review identifying 20 lightweight MAC schemes, and an embedded hardware evaluation, we show that lightweight software-implemented MACs can outperform hardware-accelerated conventional schemes for the short payloads prevalent in practice, and that AEAD ciphers introduce unnecessary overhead where confidentiality is not required. We further consolidate research on MAC aggregation and MAC-based multicast authentication, the latter growing increasingly relevant as post-quantum requirements render digital signatures impractical, and derive open challenges to guide future research and system designers.

Acknowledgments

This work was supported by the Luxembourg National Research Fund (FNR) under the PEARL project P24/IS/18423066/SCICT. This work was funded by the Topic Engineering Secure Systems of the Helmholtz Association (HGF), and supported by KASTEL Security Research Labs, Karlsruhe, Germany. This work was funded by the Federal Ministry of Research, Technology, and Space (BMFTR) in Germany under the grant number 16KIS2243 of the SUSTAINET-guardian project. The responsibility for the content of this publication lies with the authors.

References

- [1] M. Serror, S. Hack, M. Henze, M. Schuba, and K. Wehrle, "Challenges and Opportunities in Securing the Industrial Internet of Things," *IEEE Transactions on Industrial Informatics*, vol. 17, no. 5, 2021.
- [2] B. Arazi, "Message authentication in computationally constrained environments," *IEEE Transactions on Mobile Computing*, vol. 8, no. 7, pp. 968–974, 2009.

- [3] M. T. Lazarescu, "Design of a WSN Platform for Long-term Environmental Monitoring for IoT Applications," *IEEE Journal on emerging and selected topics in circuits and systems*, vol. 3, no. 1, 2013.
- [4] M. Luvisotto, Z. Pang, and D. Dzung, "Ultra High Performance Wireless Control for Critical Applications: Challenges and Directions," *IEEE Transactions on Industrial Informatics*, vol. 13, no. 3, 2016.
- [5] S. Vitturi, C. Zunino, and T. Sauter, "Industrial Communication Systems and Their Future Challenges: Next-Generation Ethernet, IIoT, and 5G," *Proceedings of the IEEE*, vol. 107, no. 6, 2019.
- [6] National Institute of Standards and Technology (NIST), "Submission Requirements and Evaluation Criteria for the Lightweight Cryptography Standardization Process," Tech. Rep., 2018. [Online]. Available: <https://csrc.nist.gov/csrc/media/Projects/lightweight-cryptography/documents/final-lwc-submission-requirements-august2018.pdf>
- [7] "CAESAR: Competition for Authenticated Encryption: Security, Applicability, and Robustness," <https://competitions.cr.yp.to/caesar.html>, accessed: 2026-06-03.
- [8] International Electrotechnical Commission (IEC), "Power systems management and associated information exchange – Data and communications security," Standard IEC 62351:2025, 2025.
- [9] D. Schmidt, K. Radke, S. Camtepe, E. Foo, and M. Ren, "A survey and analysis of the GNSS spoofing threat and countermeasures," *ACM Computing Surveys (CSUR)*, vol. 48, no. 4, pp. 1–31, 2016.
- [10] W. Duo, M. Zhou, and A. Abusorrah, "A survey of cyber attacks on cyber physical systems: Recent advances and challenges," *IEEE/CAA Journal of Automatica Sinica*, vol. 9, no. 5, pp. 784–800, 2022.
- [11] The MITRE Corporation, "Detection Engineering in Industrial Control Systems—Ukraine 2016 Attack: Sandworm Team and Industroyer Case Study," MITRE Corporation, Tech. Rep., 2022. [Online]. Available: <https://www.mitre.org/sites/default/files/2022-04/pr-22-0094-detection-engineering-in-industrial-control-systems-ukraine-2016-attack-case-study.pdf>
- [12] J. P. Mattsson, G. Selander, S. Paavolainen, F. Karakoç, M. Tiloca, and R. Moskowit, "Proposals for Standardization of the Ascon Family," in *Proceedings of the 6th Lightweight Cryptography Workshop*, 2023.
- [13] B. Galloway and G. P. Hancke, "Introduction to Industrial Control Networks," *IEEE Communications Surveys & Tutorials*, vol. 15, no. 2, 2012.
- [14] F. K. Shaikh and S. Zeadally, "Energy harvesting in wireless sensor networks: A comprehensive review," *Renewable and Sustainable Energy Reviews*, vol. 55, 2016.
- [15] LoRa Alliance, "LoRaWAN™ 1.1 Specification," Tech. Rep., 2017.
- [16] E. Wagner, D. Heye, M. Serror, and K. Wehrle, "Integrating MAC Aggregation over Lossy Channels in DTLS 1.3," in *Proceedings of the 33rd IEEE International Conference on Network Protocols (ICNP'25)*, 2025.
- [17] E. Rescorla, H. Tschofenig, and N. Modadugu, "The Datagram Transport Layer Security (DTLS) Protocol Version 1.3," IETF, RFC 9147, 2022.
- [18] A. Braeken, "Public key versus symmetric key cryptography in client–server authentication protocols," *International Journal of Information Security*, vol. 21, no. 1, 2022.
- [19] A. Barki, A. Bouabdallah, S. Gharout, and J. Traoré, "M2M Security: Challenges and Solutions," *IEEE Communications Surveys & Tutorials*, vol. 18, no. 2, pp. 1241–1254, 2016.
- [20] O. A. Amodu and M. Othman, "Machine-to-Machine Communication: An Overview of Opportunities," *Computer Networks*, vol. 145, pp. 255–276, 2018.
- [21] P. Suriyachai, U. Roedig, and A. Scott, "A Survey of MAC Protocols for Mission-Critical Applications in Wireless Sensor Networks," *IEEE Communications Surveys & Tutorials*, vol. 14, no. 2, pp. 240–264, 2011.
- [22] National Institute of Standards and Technology (NIST), "Guide to Operational Technology (OT) Security," Tech. Rep., 2026.
- [23] S. Shin, M. Kim, and T. Kwon, "Experimental performance analysis of lightweight block ciphers and message authentication codes for wireless sensor networks," *International Journal of Distributed Sensor Networks*, vol. 13, no. 11, 2017.
- [24] A. Biryukov and L. Perrin, "State of the Art in Lightweight Symmetric Cryptography," *IACR Cryptol. ePrint Arch.*, 2017. [Online]. Available: <https://eprint.iacr.org/2017/511.pdf>
- [25] S. Duval and G. Leurent, "Lightweight MACs from Universal Hash Functions," in *Proceedings of the International Conference on Smart Card Research and Advanced Applications (CARDIS'19)*. Springer, 2019, pp. 195–215.
- [26] S. Khan, P. A. Ferreira Lopes martins, B. Sousa, and V. Pereira, "A Comprehensive Review on Lightweight Cryptographic Mechanisms for Industrial Internet of Things Systems," *ACM Computing Surveys*, vol. 58, no. 1, 2025.
- [27] I. Bhardwaj, A. Kumar, and M. Bansal, "A review on lightweight cryptography algorithms for data security and authentication in IoTs," in *Proceedings of the 4th International Conference on Signal Processing, Computing and Control (ISPCC'17)*. IEEE, 2017.
- [28] E. Konstantopoulou, G. Athanasiou, and N. Sklavos, "Review and Analysis of FPGA and ASIC Implementations of NIST Lightweight Cryptography Finalists," *ACM Computing Surveys*, vol. 57, no. 10, 2025.
- [29] I. Elsadek, S. Aftabjahani, D. Gardner, E. MacLean, J. R. Wallrabenstein, and E. Y. Tawfik, "Hardware and Energy Efficiency Evaluation of NIST Lightweight Cryptography Standardization Finalists," in *Proceedings of the IEEE International Symposium on Circuits and Systems (ISCAS'22)*. IEEE, 2022, pp. 133–137.
- [30] F. Campos, L. Jellema, M. Lemmen, L. Müller, A. Sprengels, and B. Viguier, "Assembly or optimized C for lightweight cryptography on RISC-V?" in *Proceedings of the International Conference on Cryptology and Network Security (CANS'20)*. Springer, 2020, pp. 526–545.
- [31] M. N. Wegman and J. L. Carter, "New Hash Functions and Their Use in Authentication and Set Equality," *Journal of Computer and System Sciences*, vol. 22, no. 3, 1981.
- [32] B. Cogliati and Y. Seurin, "EWCDM: an efficient, beyond-birthday secure, nonce-misuse resistant MAC," in *Proceedings of the Annual International Cryptology Conference (CRYPTO'16)*. Springer, 2016, pp. 121–149.
- [33] P. Kietzmann, T. C. Schmidt, and M. Wählisch, "A Guideline on Pseudorandom Number Generation (PRNG) in the IoT," *ACM Computing Surveys (CSUR)*, vol. 54, no. 6, 2021.
- [34] E. Zenner, "Nonce Generators and the Nonce Reset Problem," in *International Conference on Information Security*. Springer, 2009.
- [35] M. Luk, G. Mezzour, A. Perrig, and V. Gligor, "MiniSec: a secure sensor network communication architecture," in *Proceedings of the 6th International Symposium on Information Processing in Sensor Networks (IPSN'07)*, 2007.
- [36] I. Fernández-Hernández, T. Ashur, and V. Rijmen, "Analysis and recommendations for MAC and key lengths in delayed disclosure GNSS authentication protocols," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 57, no. 3, pp. 1827–1839, 2021.
- [37] National Institute of Standards and Technology (NIST), "Recommendation for Block Cipher Modes of Operation: The CMAC Mode for Authentication," Tech. Rep., 2005.

- [38] French Cybersecurity Agency (ANSSI), “RÈGLES ET RECOMMANDATIONS CONCERNANT LE CHOIX ET LE DIMENSIONNEMENT DES MÉCANISMES CRYPTOGRAPHIQUES,” Tech. Rep. ANSSI-PG-083, 2026.
- [39] Federal Office for Information Security Germany (BSI), “Technical Guideline – Cryptographic Mechanisms: Recommendations and Key Lengths,” Tech. Rep. BSI TR-02102-1, 2025.
- [40] National Institute of Standards and Technology (NIST), “Ascon-Based Lightweight Cryptography Standards for Constrained Devices: Authenticated Encryption, Hash, and Extendable Output Functions,” Tech. Rep., 2025.
- [41] —, “Keyed-Hash Message Authentication Code (HMAC) Specification of HMAC and Recommendations for Message Authentication,” Tech. Rep., 2024.
- [42] European Union Agency for Cybersecurity (enisa), “Algorithms, key size and parameters report,” Tech. Rep., 2014.
- [43] E. Rescorla, “The Transport Layer Security (TLS) Protocol Version 1.3,” IETF, RFC 8446, 2018.
- [44] International Organization for Standardization (ISO), “ISO/IEC/IEEE International Standard - Information technology – Telecommunications and information exchange between systems – Local and metropolitan area networks – Specific requirements – Part 15-6: Wireless body area network,” Tech. Rep., 2018.
- [45] Connectivity Standards Alliance, “Zigbee Specification,” Connectivity Standards Alliance, Technical Report Zigbee Document 05-3474-23, 2023.
- [46] Core Specification Working Group, “Bluetooth Core Specification,” Bluetooth Special Interest Group (Bluetooth SIG), Tech. Rep. v6.3, May 2026, version date: 2026-05-05; Bluetooth Specification. Prepared by the Core Specification Working Group.
- [47] AUTOSAR, “Specification of Secure Onboard Communication Protocol, AUTOSAR FO R20-11,” Standard, 2020, https://www.autosar.org/fileadmin/standards/R20-11/FO/AUTOSAR_PRS_SecOcProtocol.pdf.
- [48] Open Metering System Group, “Open Metering System Specification Vol. 2 — Primary Communication,” Open Metering System Group e.V., Specification Issue 5.0.1, Dec. 2023. [Online]. Available: <https://oms-group.org/en/specification/>
- [49] Connectivity Standards Alliance, “Matter Specification: Core Specification, Version 1.2,” Connectivity Standards Alliance, Davis, CA, USA, Technical Report 23-27349-003_Matter-1.2-Core-Specification, Oct. 2023.
- [50] Institute of Electrical and Electronics Engineers (IEEE), “IEEE Standard for Low-Rate Wireless Networks,” *IEEE Std 802.15.4-2024 (Revision of IEEE Std 802.15.4-2020)*, pp. 1–967, Dec 2024.
- [51] 3GPP, “5G; Security architecture and procedures for 5G System (3GPP TS 33.501 version 17.5.0 Release 17),” ETSI, Technical Specification ETSI TS 133 501 V17.5.0, May 2022. [Online]. Available: https://www.etsi.org/deliver/etsi_ts/133500_133599/133501/17.05.00_60/ts_133501v170500p.pdf
- [52] ETSI, “ETSI TS 103 357-1 V2.1.1 (2024-06): Technical specification — Short Range Devices; Low Throughput Networks (LTN); Protocols for radio interface A; Part 1: Lfour and DD-UNB protocol families,” ETSI, 650 Route des Lucioles, F-06921 Sophia Antipolis Cedex, France, Tech. Rep. TS 103 357-1 V2.1.1, Jun. 2024.
- [53] SIGFOX SA, “Sigfox connected objects: Radio specifications,” SIGFOX SA, Technical report (Ref.: EP-SPECS; Rev. 1.7), Apr. 2023.
- [54] European Union Agency for the Space Programme, “GALILEO OPEN SERVICE NAVIGATION MESSAGE AUTHENTICATION (OSNMA) - INFO NOTE,” Tech. Rep., 2021.
- [55] European Union Agency for the Space Programme (EUSPA), “European GNSS (Galileo) Open Service Signal-In-Space Interface Control Document (OS SIS ICD),” EUSPA, Tech. Rep. Issue 2.2, 2025.
- [56] C. Zenger, C. Paar, and S. Wans, “Method for establishing a secure end-to-end connection via LoRa (WAN) TM,” German Patent DE102017 117 170B3, 2017.
- [57] M. Rademacher, H. Linka, J. Konrad, T. Horstmann, and K. Jonas, “Bounds for the Scalability of TLS over LoRaWAN,” in *Proceedings of the 26th ITG-Symposium on Mobile Communication-Technologies and Applications*. VDE, 2022, pp. 1–6.
- [58] D. J. Bernstein, “The Poly1305-AES Message-Authentication Code,” in *Proceedings of the International Workshop on Fast Software Encryption (FSE’05)*, 2005.
- [59] M. Boesgaard, T. Christensen, and E. Zenner, “Badger – A Fast and Provably Secure MAC,” in *Proceedings of the International Conference on Applied Cryptography and Network Security (ACNS’05)*. Springer, 2005.
- [60] “poly1305-donna,” <https://github.com/floodyberry/poly1305-donna>, accessed: 2025-11-06.
- [61] Y. Nir and A. Langley, “ChaCha20 and Poly1305 for IETF Protocols,” IETF, RFC 8439, 2018.
- [62] J. Black, S. Halevi, H. Krawczyk, T. Krovetz, and P. Rogaway, “UMAC: Fast and Secure Message Authentication,” in *Proceedings of the Annual International Cryptology Conference (CRYPTO’99)*. Springer, 1999, pp. 216–233.
- [63] “UMAC: Fast and Provably Secure Message Authentication,” <https://fastcrypto.com/umac/>, accessed: 2026-10-14.
- [64] T. Krovetz, “UMAC: Message authentication code using universal hashing,” IETF, RFC 4418, 2006.
- [65] —, “Message Authentication on 64-bit Architectures,” in *Proceedings of the International Workshop on Selected Areas in Cryptography (SAC’06)*. Springer, 2006, pp. 327–341.
- [66] “An unoptimized, straightforward reference implementation of VMAC,” <https://www.fastcrypto.org/vmac/>, accessed: 2026-10-14.
- [67] J. Garay, V. Kolesnikov, and R. McLellan, “MAC Precomputation with Applications to Secure Memory,” in *Proceedings of the International Conference on Information Security (ISC’09)*. Springer, 2009, pp. 427–442.
- [68] J. A. Garay, V. Kolesnikov, and R. McLellan, “MAC Precomputation with Applications to Secure Memory,” *ACM Transactions on Privacy and Security (TOPS)*, vol. 19, no. 2, pp. 1–21, 2016.
- [69] Z. Gong, P. Hartel, S. Nikova, and B. Zhu, “Towards Secure and Practical MACs for Body Sensor Networks,” in *Proceedings of the International Conference on Cryptology in India (INDOCRYPT’09)*. Springer, 2009, pp. 182–198.
- [70] Z. Gong, P. Hartel, S. Nikova, S.-H. Tang, and B. Zhu, “TuLP: A Family of Lightweight Message Authentication Codes for Body Sensor Networks,” *Journal of Computer Science and Technology*, vol. 29, no. 1, 2014.
- [71] J.-P. Aumasson and D. J. Bernstein, “SipHash: A Fast Short-Input PRF,” in *Proceedings of the International Conference on Cryptology in India (INDOCRYPT’12)*, 2012.
- [72] “SipHash in C,” <https://github.com/majek/csiphash/>, accessed: 2026-10-14.
- [73] D. J. Bernstein and T. Chou, “Faster Binary-Field Multiplication and Faster Binary-Field MACs,” in *Proceedings of the International Conference on Selected Areas in Cryptography (SAC’14)*. Springer, 2014, pp. 92–111.
- [74] N. Mouha, B. Mennink, A. Van Herrewege, D. Watanabe, B. Preneel, and I. Verbauwhede, “Chaskey: an efficient MAC algorithm for 32-bit microcontrollers,” in *Proceedings of the International Conference on Selected Areas in Cryptography (SAC’14)*, 2014.
- [75] “Chaskey-12 reference C implementation,” <https://mouha.be/wp-content/uploads/chaskey12.c>, accessed: 2025-11-04.

- [76] International Organization for Standardization (ISO) and International Electrotechnical Commission (IEC), "Information security — Lightweight cryptography — Part 6: Message authentication codes (MACs)," ISO/IEC ISO/IEC 29192-6:2019, 2019.
- [77] A. Luykx, B. Preneel, E. Tischhauser, and K. Yasuda, "A MAC Mode for Lightweight Block Ciphers," in *Proceedings of the International Conference on Fast Software Encryption (FSE'16)*. Springer, 2016, pp. 43–59.
- [78] "Reference implementation of the LightMAC mode of operation," <https://github.com/KULeuven-COSIC/lightmac>, accessed: 2026-10-14.
- [79] R. Ankele, F. Böhl, and S. Friedberger, "MergeMAC: a MAC for authentication with strict time constraints and limited bandwidth," in *Proceedings of the International Conference on Applied Cryptography and Network Security (ACNS'18)*, 2018.
- [80] "MAC611 implementation," <https://github.com/Cryptosaurus/MAC611>, accessed: 2025-11-06.
- [81] Y. Shen, L. Wang, and D. Gu, "LedMAC: More efficient variants of lightMAC," *Cryptology ePrint Archive*, 2021.
- [82] M. Eichlseder, A. C. Mert, C. Rechberger, and M. Schofnegger, "Small MACs from Small Permutations," *Cryptology ePrint Archive*, 2022.
- [83] "ldmac," <https://github.com/mschof/ldmac>, accessed: 2025-11-04.
- [84] E. Wagner, M. Serror, K. Wehrle, and M. Henze, "BP-MAC: Fast Authentication for Short Messages," in *Proceedings of the 15th ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec'22)*, 2022.
- [85] "BP-MAC implementation and benchmarking," <https://github.com/fkie-cad/bpmac>, accessed: 2025-11-04.
- [86] C. Dobraunig, B. Mennink, and S. Neves, "Elimac: speeding up lightmac by around 20%," *IACR Transactions on Symmetric Cryptology*, vol. 2023, no. 2, 2023.
- [87] "EliMAC," <https://github.com/jedisct1/elimac>, accessed: 2025-11-04.
- [88] A. Bariant, J. Baudrin, G. Leurent, C. Pernot, L. Perrin, and T. Peyrin, "Fast AES-Based Universal Hash Functions and MACs: Featuring LeMac and PetitMac," *IACR Transactions on Symmetric Cryptology*, vol. 2024, no. 2, 2024.
- [89] "Implementation of LeMac and PetitMac," https://github.com/AugustinBariant/Implementations_LeMac_PetitMac, accessed: 2025-11-04.
- [90] C. Dobraunig, M. Eichlseder, F. Mendel, and M. Schläffer, "Ascon MAC, PRF, and short-input PRF: lightweight, fast, and efficient pseudorandom functions," in *Cryptographers' Track at the RSA Conference*, 2024.
- [91] "Reference, highly optimized, masked C and ASM implementations of Ascon (NIST SP 800-232)," <https://github.com/ascon/ascon-c>, accessed: 2025-11-04.
- [92] D. Wang, A. Maximov, P. Ekdahl, and T. Johansson, "A New Stand-Alone MAC Construct Called SMAC," *IACR Transactions on Symmetric Cryptology*, vol. 2025, no. 1, 2025.
- [93] "Tools, demo, and test vectors for SMAC," <https://github.com/ONG/smac-tools>, accessed: 2025-11-06.
- [94] T. Iwata, V. Lallemand, G. Leander, and Y. Sasaki, "Universal Forgery and Multiple Forgeries of MergeMAC and Generalized Constructions," in *Proceedings of the Cryptographers' Track at the RSA Conference*. Springer, 2019, pp. 293–312.
- [95] D. Dinu, A. Biryukov, J. Großschädl, D. Khovratovich, Y. Le Corre, and L. Perrin, "FELICS-Fair Evaluation of Lightweight Cryptographic Systems," in *Proceedings of the NIST Workshop on Lightweight Cryptography*, vol. 128, 2015.
- [96] H. Kim and H. Seo, "Optimizing AES-GCM on 32-Bit ARM Cortex-M4 Microcontrollers: Fixslicing and FACE-based Approach," *ACM Transactions on Embedded Computing Systems*, vol. 24, no. 6, 2025.
- [97] C. Dobraunig, M. Eichlseder, F. Mendel, and M. Schläffer, "Ascon v1 Submission to the CAESAR Competition," Submission to CAESAR: Competition for Authenticated Encryption: Security, Applicability, and Robustness, 2014. [Online]. Available: <https://competitions.cr.yp.to/round1/asconv1.pdf>
- [98] —, "Ascon v1.2 Submission to NIST," Submission to NIST Lightweight Cryptography Standardization Process, 2021. [Online]. Available: <https://csrc.nist.gov/CSRC/media/Projects/lightweight-cryptography/documents/finalist-round/updated-spec-doc/ascon-spec-final.pdf>
- [99] J. Hiller, M. Henze, M. Serror, E. Wagner, J. N. Richter, and K. Wehrle, "Secure Low Latency Communication for Constrained Industrial IoT Scenarios," in *Proceedings of the 42nd IEEE Conference on Local Computer Networks (LCN'18)*, 2018.
- [100] D. Ma and G. Tsudik, "Forward-secure sequential aggregate authentication," in *Proceedings of the IEEE Symposium on Security and Privacy (S&P'07)*, 2007, pp. 86–91.
- [101] A. C. Chan and C. Castelluccia, "On the (im) possibility of aggregate message authentication codes," in *Proceedings of the IEEE International Symposium on Information Theory (ISIT'08)*. IEEE, 2008.
- [102] J. Katz and A. Y. Lindell, "Aggregate Message Authentication Codes," in *Proceedings of the Cryptographers' Track at the RSA Conference*, 2008.
- [103] O. Eikemeier, M. Fischlin, J.-F. Götzmann, A. Lehmann, D. Schröder, P. Schröder, and D. Wagner, "History-Free Aggregate Message Authentication Codes," in *Proceedings of the International Conference on Security and Cryptography for Networks (SCN'10)*, 2010.
- [104] V. Kolesnikov, "MAC Aggregation with Message Multiplicity," in *Proceedings of the International Conference on Security and Cryptography for Networks (SCN'12)*, 2012.
- [105] S. Hirose and J. Shikata, "Non-adaptive Group-Testing Aggregate MAC Scheme," in *Proceedings of the International Conference on Information Security Practice and Experience (ISPEC'18)*, 2018.
- [106] A. A. Yavuz and P. Ning, "Self-sustaining, efficient and forward-secure cryptographic constructions for unattended wireless sensor networks," *Ad Hoc Networks*, vol. 10, no. 7, pp. 1204–1220, 2012.
- [107] Y. Chen and T. Kunz, "Performance Evaluation of IoT Protocols under a Constrained Wireless Access Network," in *Proceedings of International Conference on Selected Topics in Mobile & Wireless Networking (MoWNeT'16)*, 2016.
- [108] E. Wagner, M. Serror, K. Wehrle, and M. Henze, "When and How to Aggregate Message Authentication Codes on Lossy Channels?" in *Proceedings of the International Conference on Applied Cryptography and Network Security (ACNS'24)*, 2024.
- [109] E. Wagner, J. Bauer, and M. Henze, "Take a Bite of the Reality Sandwich: Revisiting the Security of Progressive Message Authentication Codes," in *Proceedings of the 15th ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec'22)*, 2022.
- [110] F. Armknecht, P. Walther, G. Tsudik, M. Beck, and T. Strufe, "ProMACs: Progressive and Resynchronizing MACs for Continuous Efficient Authentication of Message Streams," in *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security (CCS'20)*, 2020.
- [111] H. Li, V. Kumar, J.-M. J. Park, and Y. Yang, "Cumulative Message Authentication Codes for Resource-Constrained Networks," in *Proceedings of the IEEE Conference on Communications and Network Security (CNS'20)*, 2020.

- [112] H. Li, V. Kumar, J.-M. Park, and Y. Yang, "Cumulative Message Authentication Codes for Resource-Constrained IoT Networks," *IEEE Internet of Things Journal*, vol. 8, no. 15, 2021.
- [113] P. Ginzboorg, V. Niemi, and J. Ott, "Authentication of fragments with short tags," *Theoretical Computer Science*, vol. 1009, 2024.
- [114] J. Schmandt, A. T. Sherman, and N. Banerjee, "Mini-MAC: Raising the bar for vehicular security with a lightweight message authentication protocol," *Vehicular Communications*, vol. 9, 2017.
- [115] S. Kashani, S. W. Kim, and A. Khokhar, "Two-Dimensional Compound Message Authentication Code in Lossy Channels," in *Proceedings of the IEEE International Conference on Communications (ICC'25)*, 2025.
- [116] S. E. Nouma and A. A. Yavuz, "Lightweight and Breach-Resilient Authenticated Encryption Framework for Internet of Things," in *Proceedings of the IEEE Military Communications Conference (MILCOM'25)*, 2025, pp. 1–6.
- [117] P. Gaži, K. Pietrzak, and M. Rybár, "The exact PRF-security of NMAC and HMAC," in *Proceedings of the 34th Annual Cryptology Conference (CRYPTO'14)*, 2014.
- [118] M. Bellare, R. Guérin, and P. Rogaway, "XOR MACs: New Methods for Message Authentication Using Finite Pseudorandom Functions," in *Proceedings of the Annual International Cryptology Conference (CRYPTO'95)*, 1995.
- [119] S. E. Nouma, G. Mumcu, and A. A. Yavuz, "Diamond: Design and Implementation of Breach-Resilient Authenticated Encryption Framework For Internet of Things," *arXiv preprint arXiv:2601.00353*, 2026.
- [120] D. Povey, "Optimistic Security: A New Access Control Paradigm," in *Proceedings of the 1999 workshop on New Security Paradigms (NSPW'99)*, 1999.
- [121] C. Bachhuber, E. Steinbach, M. Freundl, and M. Reisslein, "On the Minimization of Glass-to-Glass and Glass-to-Algorithm Delay in Video Communication," *IEEE Transactions on Multimedia*, vol. 20, no. 1, 2017.
- [122] J. Yun, S. Yoon, B. Kim, D. Kim, and Y. Kang, "Design of a Progressive MAC for Drone RC Data Protection," in *Proceedings of the 16th International Conference on Information and Communication Technology Convergence (ICTC'25)*. IEEE, 2025, pp. 100–102.
- [123] E. Steinbach, S. Hirche, M. Ernst, F. Brandi, R. Chaudhari, J. Kammerl, and I. Vittorias, "Haptic Communications," *Proceedings of the IEEE*, vol. 100, no. 4, 2012.
- [124] K. Antonakoglou, X. Xu, E. Steinbach, T. Mahmoodi, and M. Dohler, "Toward Haptic Communications Over the 5G Tactile Internet," *IEEE Communications Surveys & Tutorials*, vol. 20, no. 4, 2018.
- [125] S. Gueron, "Memory Encryption for General-Purpose Processors," *IEEE Security & Privacy*, vol. 14, no. 6, 2016.
- [126] S. Moriam, E. Franz, P. Walther, A. Kumar, T. Strufe, and G. Fettweis, "Protecting Communication in Many-Core Systems against Active Attackers," in *Proceedings of the Great Lakes Symposium on VLSI (GLSVLSI'18)*, 2018.
- [127] C. Szilagyi and P. Koopman, "Low Cost Multicast Authentication via Validity Voting in Time-Triggered Embedded Control Networks," in *Proceedings of the 5th Workshop on Embedded Systems Security (WESS'10)*, 2010.
- [128] C. Szilagyi, "Low cost multicast network authentication for embedded control systems," Ph.D. dissertation, Carnegie Mellon University, 2012.
- [129] D. K. Nilsson, U. E. Larson, and E. Jonsson, "Efficient In-Vehicle Delayed Data Authentication Based on Compound Message Authentication Codes," in *Proceedings of the 68th Vehicular Technology Conference (VTC-Fall'08)*, 2008.
- [130] J. H. Castellanos, D. Antonioli, N. O. Tippenhauer, and M. Ochoa, "Legacy-Compliant Data Authentication for Industrial Control System Traffic," in *Proceedings of the International Conference on Applied Cryptography and Network Security (ACNS'17)*, 2017.
- [131] C. Szilagyi and P. Koopman, "Flexible Multicast Authentication for Time-Triggered Embedded Control Network Applications," in *Proceedings of the IEEE/IFIP International Conference on Dependable Systems & Networks (DSN'09)*, 2009.
- [132] D. Sikeridis, P. Kampanakis, and M. Devetsikiotis, "Assessing the overhead of post-quantum cryptography in TLS 1.3 and SSH," in *Proceedings of the 16th International Conference on emerging Networking EXperiments and Technologies (CoNEXT'20)*, 2020, pp. 149–156.
- [133] S. Bootsma and M. De Vries, "A survey on the quantum security of block cipher-based cryptography," *IEEE Access*, vol. 12, pp. 194 711–194 727, 2024.
- [134] Y. Challal, H. Bettahar, and A. Bouabdallah, "A taxonomy of multicast data origin authentication: Issues and solutions," *IEEE Communications Surveys & Tutorials*, vol. 6, no. 3, 2004.
- [135] A. Perrig, R. Canetti, J. D. Tygar, and D. Song, "Efficient Authentication and Signing of Multicast Streams over Lossy Channels," in *Proceedings of the IEEE Symposium on Security and Privacy (S&P'00)*, 2000.
- [136] A. Perrig, R. Canetti, D. Song, and J. D. Tygar, "Efficient and Secure Source Authentication for Multicast," in *Proceedings of the Network and Distributed Systems Security Symposium (NDSS'01)*, 2001.
- [137] D. Liu and P. Ning, "Multilevel μ TESLA: Broadcast authentication for distributed sensor networks," *ACM Transactions on Embedded Computing Systems*, vol. 3, no. 4, 2004.
- [138] A. Studer, F. Bai, B. Bellur, and A. Perrig, "Flexible, extensible, and efficient VANET authentication," *Journal of Communications and Networks*, vol. 11, no. 6, pp. 574–588, 2009.
- [139] J. C. P. García, A. Benslimane, A. Braeken, and Z. Su, " μ tesla-based authentication for reliable and secure broadcast communications in iot using blockchain," *IEEE Internet of Things Journal*, vol. 10, no. 20, pp. 18 400–18 413, 2023.
- [140] U. Tefek, E. Esiner, D. Mashima, B. Chen, and Y.-C. Hu, "Caching-based multicast message authentication in time-critical industrial control systems," in *Proceedings of the IEEE Conference on Computer Communications (INFOCOM'22)*. IEEE, 2022, pp. 1039–1048.
- [141] M. Kutylowski, A. Cinal, P. Kubiak, and D. Korniienko, "Authenticated multicast in tiny networks via an extremely low-bandwidth medium," *Applied Sciences*, vol. 14, no. 17, p. 7962, 2024.
- [142] R. Canetti, J. Garay, G. Itkis, D. Micciancio, M. Naor, and B. Pinkas, "Multicast Security: A Taxonomy and some Efficient Constructions," in *Proceedings of the 18th Annual Joint Conference on Computer Communications (INFOCOM'99)*, 1999.
- [143] F. Ye, H. Luo, S. Lu, and L. Zhang, "Statistical en-route filtering of injected false data in sensor networks," *IEEE Journal on selected areas in communications*, vol. 23, no. 4, pp. 839–850, 2005.
- [144] T. Wu, Y. Cui, B. Kusi, A. Ledeczi, J. Sallai, N. Skirvin, J. Werner, and Y. Xue, "A Fast and Efficient Source Authentication Solution for Broadcasting in Wireless Sensor Networks," in *New Technologies, Mobility and Security*, 2007.
- [145] X. Yao, X. Han, X. Du, and X. Zhou, "A lightweight multicast authentication mechanism for small scale IoT applications," *IEEE Sensors Journal*, vol. 13, no. 10, pp. 3693–3701, 2013.
- [146] S. Woo, M. Seo, and C.-S. Park, "mMIC: Multicast Message Integrity Code for Securing Low-Entropy Control Command," *IEEE Access*, 2025.
- [147] E. Wagner, F. Basels, J. Bauer, T. Zimmermann, K. Wehrle, and M. Henze, "CAIBA: Multicast Source Authentication for CAN Through Reactive Bit Flipping," in *Proceedings of the 2025 IEEE 10th European Symposium on Security and Privacy (EuroS&P'25)*, 2025.

- [148] S. Movassaghi, M. Abolhasan, J. Lipman, D. Smith, and A. Jamalipour, "Wireless body area networks: A survey," *IEEE Communications surveys & tutorials*, vol. 16, no. 3, pp. 1658–1686, 2014.
- [149] R. Punj and R. Kumar, "Technological aspects of WBANs for health monitoring: a comprehensive review," *Wireless Networks*, vol. 25, no. 3, pp. 1125–1157, 2019.
- [150] I. Straw, A. Polamarasetty, and M. Jaafar, "Abusing the Internet of Medical Things: Evaluating Threat Models and Forensic Readiness for Multi-Vector Attacks on Connected Healthcare Devices," in *Proceedings of the European Symposium on Security and Privacy (EuroS&P'26)*, 2026.
- [151] H. Malik, M. M. Alam, Y. Le Moullec, and A. Kuusik, "NarrowBand-IoT performance analysis for healthcare applications," *Procedia computer science*, vol. 130, pp. 1077–1083, 2018.
- [152] B. Latré, B. Braem, I. Moerman, C. Blondia, and P. Demeester, "A survey on wireless body area networks," *Wireless networks*, vol. 17, no. 1, pp. 1–18, 2011.
- [153] M. Hernandez, R. Kohno, T. Kobayashi, and M. Kim, "Technical Requirements Document," Tech. Rep. IEEE 802.15 Working Group for Wireless Specialty Networks, 2021.
- [154] R. N. Charette, "This car runs on code," *IEEE spectrum*, vol. 46, no. 3, p. 3, 2009.
- [155] International Organization for Standardization (ISO), "Road vehicles - Controller area network (CAN) - Part 1: Data link layer and physical signalling," ISO 11898-1:2015, 2015.
- [156] K. Tindell and A. Burns, "Guaranteeing Message Latencies on Controller Area Network (CAN)," in *Proceedings of the 1st International CAN Conference*. Citeseer, 1994, p. 187.
- [157] International Organization for Standardization (ISO), "Road vehicles - Local Interconnect Network (LIN) - Part 1: General information and use case definition," ISO 17987-1:2016, 2016.
- [158] D. Mocrii, Y. Chen, and P. Musilek, "IoT-based smart homes: A review of system architecture, software, communications, privacy and security," *Internet of Things*, vol. 1-2, pp. 81–98, 2018.
- [159] A. A. Zaidan, B. B. Zaidan, M. Y. Qahtan, O. S. Albahri, A. S. Albahri, A. Mussab, F. M. Jumaah, M. Tala, K. L. Tan, W. L. Shir, and L. C. K., "A survey on communication components for IoT-based technologies in smart homes," *Telecommunication Systems*, vol. 69, no. 1, pp. 1–25, 2018.
- [160] M. H. Mazhar and Z. Shafiq, "Characterizing Smart Home IoT Traffic in the Wild," in *Proceedings of the IEEE/ACM Fifth International Conference on Internet-of-Things Design and Implementation (IoTDI'20)*, 2020, pp. 203–215.
- [161] M. B. Attia, K. K. Nguyen, and M. Cheriet, "Dynamic QoS-Aware Scheduling for Concurrent Traffic in Smart Home," *IEEE Internet of Things Journal*, vol. 7, no. 6, pp. 5412–5425, 2020.
- [162] A. G. Ramonet and T. Noguchi, "IEEE 802.15.4 Now and Then: Evolution of the LR-WPAN Standard," in *Proceedings of the International Conference on Advanced Communication Technology (ICACT'20)*, 2020, pp. 1198–1210.
- [163] M. Wollschlaeger, T. Sauter, and J. Jasperneite, "The Future of Industrial Communication: Automation Networks in the Era of the Internet of Things and Industry 4.0," *IEEE Industrial Electronics Magazine*, vol. 11, no. 1, 2017.
- [164] T. Zhang, G. Wang, C. Xue, J. Wang, M. Nixon, and S. Han, "Time-Sensitive Networking (TSN) for Industrial Automation: Current Advances and Future Directions," *ACM Comput. Surv.*, vol. 57, no. 2, 2024. [Online]. Available: <https://doi.org/10.1145/3695248>
- [165] T. Sauter, "The three generations of field-level networks—Evolution and compatibility issues," *IEEE Transactions on Industrial Electronics*, vol. 57, no. 11, pp. 3585–3595, 2010.
- [166] R. Glebke, M. Henze, K. Wehrle, P. Niemietz, D. Trauth, P. Mattfeld, and T. Bergs, "A Case for Integrated Data Processing in Large-Scale Cyber-Physical Systems," in *Proceedings of the 52nd Hawaii International Conference on System Sciences (HICSS'19)*, 2019.
- [167] D. Eckhoff and I. Wagner, "Privacy in the Smart City—Applications, Technologies, Challenges, and Solutions," *IEEE Communications Surveys & Tutorials*, vol. 20, no. 1, pp. 489–516, 2018.
- [168] I. Yaqoob, I. A. T. Hashem, Y. Mehmood, A. Gani, S. Mokhtar, and S. Guizani, "Enabling Communication Technologies for Smart Cities," *IEEE Communications Magazine*, vol. 55, no. 1, pp. 112–120, 2017.
- [169] A. R. Javed, F. Shahzad, S. ur Rehman, Y. B. Zikria, I. Razzak, Z. Jalil, and G. Xu, "Future smart cities: requirements, emerging technologies, applications, challenges, and future aspects," *Cities*, vol. 129, p. 103794, 2022.
- [170] A. Pal, F. Campagnaro, K. Ashraf, M. R. Rahman, A. Ashok, and H. Guo, "Communication for Underwater Sensor Networks: A Comprehensive Summary," *ACM Transactions on Sensor Networks*, vol. 19, no. 1, 2022, DOI: 10.1145/3546827.
- [171] J. Potter, J. Alves, D. Green, G. Zappa, I. Nissen, and K. McCoy, "The JANUS underwater communications standard," in *Proceedings of Underwater Communications and Networking (UComms'14)*, 2014, pp. 1–4.
- [172] M. Goetz, I. Nissen, R. Otnes, and P. van Walree, "Performance Analysis of Underwater Network Protocols within International Sea Trial," in *Proceedings of OCEANS – Genova*, 2015.
- [173] M. Goetz and I. Nissen, "GUWMANET – Multicast routing in Underwater Acoustic Networks," in *Proceedings of the Military Communications and Information Systems Conference (MCC'12)*, 2012.
- [174] L. Lux, E. Padilla, and J. Bauer, "Trust in the Deep: a Survey on Authenticating Acoustic Underwater Communication," in *Proceedings of the IEEE Virtual Conference on Communications (VCC'25)*, 2025.
- [175] P. H. Nardelli, N. Rubido, C. Wang, M. S. Baptista, C. Pomalaza-Raez, P. Cardieri, and M. Latva-aho, "Models for the modern power grid," *The European Physical Journal Special Topics*, vol. 223, no. 12, pp. 2423–2437, 2014.
- [176] X. Fang, S. Misra, G. Xue, and D. Yang, "Smart Grid — The New and Improved Power Grid: A Survey," *IEEE Communications Surveys & Tutorials*, vol. 14, no. 4, pp. 944–980, 2012.
- [177] H. León, C. Montez, M. Stemmer, and F. Vasques, "Simulation Models for IEC 61850 Communication in Electrical Substations Using GOOSE and SMV Time-critical Messages," in *Proceedings of the IEEE World Conference on Factory Communication Systems (WFCS'16)*, 2016.
- [178] J. Gaspar, T. Cruz, C.-T. Lam, and P. Simões, "Smart Substation Communications and Cybersecurity: A Comprehensive Survey," *IEEE Communications Surveys & Tutorials*, vol. 25, no. 4, pp. 2456–2493, 2023.
- [179] A. d. Santos, B. Soares, C. Fan, M. Kuipers, S. Sabino, A. M. R. C. Grilo, P. R. B. A. Pereira, M. S. Nunes, and A. Casaca, "Characterization of Substation Process Bus Network Delays," *IEEE Transactions on Industrial Informatics*, vol. 14, no. 5, pp. 2085–2094, 2018.
- [180] International Civil Aviation Organization (ICAO), "Aeronautical Communication – Volume IV: Surveillance and Collision Avoidance Systems, Annex 10, Edition 5," International Standards and Recommended Practices, 2014.
- [181] Z. Wu, T. Shang, and A. Guo, "Security Issues in Automatic Dependent Surveillance – Broadcast (ADS-B): A Survey," *IEEE Access*, vol. 8, pp. 122 147–122 167, 2020.
- [182] M. Strohmeier, V. Lenders, and I. Martinovic, "On the Security of the Automatic Dependent Surveillance-Broadcast Protocol," *IEEE Communications Surveys & Tutorials*, vol. 17, no. 2, pp. 1066–1087, 2015.
- [183] E. D. Kaplan and C. Hegarty, *Understanding GPS/GNSS: principles and applications*. Artech house, 2017.

- [184] U.S. Space Force, “IS-GPS-200: NAVSTAR GPS Space Segment/Navigation User Segment Interfaces,” U.S. Space Force, Tech. Rep., 2022.
- [185] M. Bellare, J. Kilian, and P. Rogaway, “The security of cipher block chaining,” in *Proceedings of the Annual International Cryptology Conference (CRYPTO’94)*. Springer, 1994, pp. 341–358.
- [186] J. Black and P. Rogaway, “A Block-Cipher Mode of Operation for Parallelizable Message Authentication,” in *Proceedings of the Annual International Conference on the Theory and Applications of Cryptographic Techniques (Eurocrypt’02)*. Springer, 2002, pp. 384–397.
- [187] G. Leurent, “Improved differential-linear cryptanalysis of 7-round chaskey with partitioning,” in *Proceedings of the Annual International Conference on the Theory and Applications of Cryptographic Techniques (Eurocrypt’16)*, 2016.

A Considered Constrained Environments

Body Area Networks are poised to transform health-care, the military, and sports through continuous monitoring of vital functions [148], [149], though their sensitivity to attacks on device functionality and personal health data poses significant security challenges [150]. Consisting of implanted and surface sensors connected to gateways like smartphones, BANs operate with small payloads (often 1–4 byte [151], [152]) and polling frequencies from a few times per minute to thousands per second [152], with latencies under 100 ms and data rates up to 50 Mbit/s across up to 64 battery-constrained devices [152], [153]. Communication relies on IEEE 802.15.6, proprietary protocols, and general-purpose IoT standards such as ZigBee and BLE [44], [149].

Intra-vehicular networks coordinate up to 100 Electronic Control Units (ECUs) handling everything from powertrain to infotainment and driver assistance [154]. Safety-critical messages are usually short Controller Area Network (CAN) bus control frames (just a few byte) due to the protocol’s simplicity and fault tolerance [155], as ECUs must also operate reliably under harsh conditions. Some hard real-time safety messages need deterministic responses under 20 ms [156]. For less critical low-bandwidth controls, *e.g.*, window operation, Local Interconnect Network (LIN) is used to minimize costs and weight [157].

Smart Homes interconnect heterogeneous IoT devices to automate functions like lighting, climate, and security [158]. These may include dozens to hundreds of sensors, actuators, appliances, and gateways communicating over local wireless networks. Traffic ranges from small, periodic sensor updates to continuous multimedia streams at several Mbit/s [159], [160]. Latency needs are diverse, with interactive uses requiring responses below 100 ms, but routine automation tolerates higher delays [161]. Low-power protocols like IEEE 802.15.4 limit data rates to 250 kbit/s and support over 100 devices [162]. In addition to Wi-Fi for high throughput, smart homes use efficient protocols such as ZigBee, Thread, and Bluetooth Low Energy (BLE).

Industrial Control Systems monitor and control processes in sectors, including manufacturing, energy, and water treatment [163], comprising up to hundreds of components [22], such as PLCs, Remote Terminal Units (RTUs),

sensors, and actuators. Communication consists mostly of small periodic messages (up to a few dozen byte), with critical applications requiring latencies below 1 ms and supervisory traffic tolerating tens to hundreds of milliseconds [4], [164]. Many ICSs rely on resource-constrained legacy devices [165], with aggregate bandwidth up to Gbit/s [166], using industrial Ethernet, fieldbuses, and wireless technologies such as 5G and IEEE 802.15.4.

Smart Cities integrate heterogeneous CPSs across urban domains, including transportation and environmental monitoring, to improve efficiency, sustainability, and quality of life [167]. They interconnect a vast number of distributed IoT devices, *e.g.*, traffic sensors, surveillance cameras, smart meters, and public infrastructure components. Communication patterns range from low-rate periodic measurements (*e.g.*, sensing or metering) to high-bandwidth real-time video streams, resulting in data rates from a few bit/s up to several Mbit/s per device [168]. Latency requirements span from near real-time, *e.g.*, for safety-critical services, to highly delay-tolerant, *e.g.*, for environmental monitoring [169]. Given their geographic scale, smart city networks combine technologies such as 5G, LoRaWAN, and Sigfox.

Underwater Acoustic Networks interconnect submerged sensors, autonomous underwater vehicles, and fixed subsea infrastructure for environmental monitoring, offshore engineering, and defense applications [170]. Communication relies predominantly on acoustics, employing standards such as JANUS [171] and vendor-specific stacks, *e.g.*, GUWAL [172], with multi-hop routing [173]. The extremely limited bandwidth yields only a few bit/s to tens of kbit/s, with traffic dominated by periodic sensing, event-driven alarms, and command-and-control [170]. Propagation delays of several seconds over tens of kilometers favor sparse, delay-tolerant operation, and authentication has yet to be established in current standards [174].

Power Grids transport and distribute electrical energy from generation units to end users [175]. Traditionally centralized with unidirectional flow, these networks are shifting toward decentralized architectures with bidirectional power flows driven by distributed renewable energy sources [176], introducing higher coordination complexity and increased demand for reliable, secure communication. Mission-critical M2M communication typically takes place within substations, *e.g.*, to implement protection functions [177], [178], with stringent latency requirements as low as 3 ms for GOOSE trip commands and 208.3 μ s for SMV measurements, with typical payloads of 100–300 byte [179].

Aviation Aviation surveillance infrastructures wirelessly interconnect aircraft, ground stations, and traffic control centers to support safety-critical situational awareness. Weight and certification constraints limit protocol and hardware complexity, while protocols such as ADS-B [180] and Mode S involve tens of thousands of transponders broadcasting several short, fixed-format messages (*e.g.*, 56–112 bits) per second [181], [182], achieving per-aircraft data rates of a few kbit/s, though aggregate load can be high in dense airspaces. The protocols rely on periodic broadcasts, tolerat-

ing occasional packet loss via redundancy and multi-sensor fusion, yet offer only limited cryptographic protection despite high integrity and authenticity requirements [182].

✿ **Satellite Navigation** systems such as GPS or Galileo provide global positioning, navigation, and timing services by broadcasting signals from medium-Earth orbit to passive receivers [183]. Receivers are strictly listen-only, decoding low-power spread-spectrum signals at very low data rates (25–125 bit/s), with message payloads of a few hundred bits carrying ephemeris, clock corrections, and integrity data [55], [184]. Latency is inherently bounded by signal propagation and acquisition time rather than protocol design.

B Details on MAC Constructions

We now discuss the MAC constructions where we consider a single forgery attempt, and the q_v term from Theorem 1 applies universally as an additive overhead when multiple verification attempts are considered.

🌀 **Universal Hash-based MACs.** UHFs are commonly used primitives for constructing efficient and secure MACs, particularly the CW-style constructions. UHFs are a family of hash functions $\mathcal{H} = \{h_k\}$ from a universe $\mathcal{U}$ to a range $\{0, \dots, N - 1\}$, such that for any fixed distinct pair of inputs $(x, y) \in \mathcal{U}$, a uniformly random function $h_k \in \mathcal{H}$ collides with a probability approximately $\frac{1}{N}$. While this is true for a “perfectly” universal hash family, a less expensive choice with a negligible loss in security is provided by ϵ -Almost Universal (ϵ -AU) hash family, where the collision probability ϵ is a small value derived from the hash’s algebraic structure and may scale linearly with the length of the input messages. A strictly stronger notion of ϵ -Almost XOR Universal (ϵ -AXU) requires that for any two distinct inputs $x, y \in \mathcal{U}$, and for any fixed output difference Δ , the probability that $h_k(x) \oplus h_k(y) = \Delta$ over uniform random choices of h is at most ϵ . Since UHF security is purely statistical (based on uniform random choices over h_k from $\mathcal{H}$), and does not rely on any cryptographic hardness assumptions, the functions can be implemented using fast, parallelizable, and hardware-friendly algebraic operations. CW-style MACs are nonce-respecting and combine a UHF $\mathcal{H}$ with a pseudorandom function F to compute a tag: $\text{CW}_{k_1, k_2}(n, m) = \mathcal{H}_{k_1}(m) \oplus F_{k_2}(n)$, where k_1, k_2 are secret keys, and n is a nonce. While UHF ensures low collision probability, F masks the hash value. Hence, the success probability of an EUF-CMA adversary for forging a tag on a new message is bound by the UHF’s collision probability and the security of F . Nevertheless, security requires that nonces be unique to ensure that mask values are never repeated.

Poly1305 [58]: A highly optimized instantiation of the CW paradigm defined as: $\text{Poly1305}_{k_1, k_2}(n, m) = \mathcal{H}_{k_1}(m) + \text{AES}_{k_2}(n) \bmod 2^{128}$. Here, the UHF is instantiated via a polynomial $f(\cdot)$ generated by parsing $m = m_1 \parallel \dots \parallel m_n$ and evaluating $f(k_1)$ as $\sum_{i=1}^n m_i \cdot k^i$, while the masking is instantiated with a block cipher AES. Security relies on the computational indistinguishability of AES output from

truly random bits of the same length. However, since AES is a PRP, its outputs are bijective and never collide, whereas a truly random function may. The PRP-PRF Switching Lemma bounds the distinguishing advantage between the two by $\frac{q^2}{2^{128}}$, which propagates into the EUF-CMA bound as an additive term. Poly1305 is very fast in software implementations as it uses integer arithmetic modulo a pseudo Mersenne Prime. This choice enables efficient modular reduction using just addition and multiplication. Furthermore, the regular structure of polynomial evaluation allows for vectorization and parallelization on modern processors.

UMAC [62]: A CW-style MAC that uses a keyed ϵ -AU hash function UHASH combined with a pseudorandom pad, achieving high speed and parallelization along with provable security guarantees. It is primarily designed for 32-bit architectures with SIMD support, enabling execution of many hash functions in parallel using their fastest instructions, resulting in high throughput (approx. 1 CPU cycle/byte). UHASH constitutes a 3-layered construction that forms an ϵ -almost universal hashing family, where the internal layer uses a software-optimized hash family NH defined as: $\text{NH}_k(m) = \sum_{i=1}^{u/2} (k_{2i-1} + m_{2i-1}) \cdot (k_{2i} + m_{2i}) \bmod 2^{2w}$, where $m = (m_1, \dots, m_u)$ and $k = (k_1, \dots, k_v)$, $m_i, k_i \in \{0, \dots, 2^w - 1\}$. In particular, additions and multiplications are standard word-level integer operations that modern CPUs can perform efficiently, while achieving very high parallelism through the simultaneous processing of the message blocks m_i . Nevertheless, the NH family alone provides limited universality guarantees. Layer 2 uses a randomized polynomial hash over a finite field that combines outputs of layer 1, and gives a tight universal bound on the collision probability via polynomial evaluation at a secret point. Layer 3 uses an inner-product hash for final compression, ensuring that the output fits the desired tag size while preserving the construction’s universality properties. Finally, the hash output is masked using a PRF, producing a provably secure MAC. EUF-CMA security follows from the ϵ -AU hashing collision bounds (concretely, 2^{-60} for a 64-bit tag), PRF security and uniqueness of nonces.

VMAC [65]: A CW-style MAC optimized for 64-bit processors with advanced SIMD support, producing extremely high throughput (0.5 CPU cycles per byte). VMAC instantiates UHF with a keyed hash function family VHASH, while the PRF is usually replaced by an AES cipher. VHASH is essentially a UHASH, redesigned for 64-bit vector hardware. Since VHASH is an ϵ -AU hash family, the collision probability for a 64-bit VMAC is 2^{-60} (approx.). Hence, as long as the underlying PRF has indistinguishable outputs from truly random bits of the same length and nonce uniqueness is guaranteed, the success of an EUF-CMA adversary stays below 2^{-60} (alternatively 2^{-120} for 128-bit tags).

BP-MAC [84]: Another MAC in the CW setting, precisely suited for low-latency transmission of *short* messages. The intuitive idea is to offload the computationally expensive algorithms at the setup phase, such that tag generation requires only a small number of lightweight operations. Concretely, BP-MAC instantiates the UHF with AES, which precomputes tags for every $v \in \{0, 1\}$ as $\text{AES}_{k_1}(j \parallel v)$, where

$j \in \{0, \dots, |m| - 1\}$, and stores them in a table of $2|m|$ entries. Then, for a message m , the tag is computed as $\bigoplus_{i=0}^{|m|-1} \text{AES}_{k_1}(j \| m[j]) \oplus \text{AES}_{k_2}(n)$. The key connection is that $\text{AES}_{k_1}(j \| m[j])$ is not a fresh AES evaluation—it is simply a lookup into the precomputed table, reducing the overall Tag algorithm to ℓ XOR operations. The EUF-CMA security bound is approximately the summation of collision probability of $\frac{q^2}{2^{128}}$, the PRP advantage of AES, and the tag-guessing probability of $\frac{1}{2^\ell}$.

Mac611 [25]: An extension of the CW framework (called the EWCDM paradigm [32]) that addresses two of its limitations. In standard CW, the output of UHF is combined with an XOR, which leads to i) a birthday term from the PRP-PRF Switching Lemma, and ii) EUF-CMA security breaks down assuming a nonce-misuse adversary. Mac611 computes $\text{tag} = F_{k_2}(\mathcal{H}_{k_1}(m) \| n)$, feeding the concatenation of hash output and nonce directly into F . The UHF is instantiated with XPoly, a polynomial hash over $\mathbb{F}_{2^{61}-1}$ that produces a 61 bit digest, concatenated with a 67 bit nonce, and fed into AES-128. Since AES operates on 128-bit message-dependent input rather than the nonce alone, collisions are governed by the birthday bound over 2^{128} , reducing the switching loss to a negligible $\frac{q^2}{2^{128}}$. Additionally, even if a nonce repeats, the AES inputs remain distinct unless there is a collision in the UHF itself, *i.e.*, $\mathcal{H}_{k_1}(m) = \mathcal{H}_{k_1}(m')$ —which is bounded by the ϵ -AU of XPoly. In a nutshell, the EUF-CMA security reduces to breaking AES as a PRP, finding a hash collision on a repeated nonce, or guessing a valid ℓ -bit tag.

Auth256 [73]: A CW-style MAC designed to provide 256-bit security via a universal hash HASH256. A message $m = m_1 \| \dots \| m_n$ is compressed using an ϵ -AXU hash over $\mathbb{F}^{256}$, and the resulting digest is masked using a pseudorandom key stream. The core hashing component is based on a pseudo dot product $(m_1 + r_1)(m_2 + r_2) \dots (m_n + r_n)$, with only 0.5 multiplications per block - significantly fewer than polynomial hashing (1 multiplication per block). To achieve high performance for *long* messages, field multiplications are implemented using FFT techniques, and the construction exploits the additive structure of the pseudo dot products to merge the inverse FFT computations across multiple products. The EUF-CMA security reduces to the probability of finding a collision in the underlying ϵ -AXU hash function, distinguishing a pseudorandom masking stream from a random one, along with the tag-guessing probability.

LeMAC [88]: Built on top of an ϵ -AU hash function, LeMAC uses the EWCDM paradigm and achieves high performance on modern Intel processors with AES-NI hardware support. LeMAC maintains 9×128 -bit internal state registers, along with auxiliary registers, and absorbs message blocks via repeated applications of the AES round function $A(\cdot)$. Following this step, several finalization rounds are applied, yielding an intermediate compressed value h . Finally, the tag is computed as $\text{tag} = \text{AES}_{k_1}(h + \text{AES}_{k_2}(n) + n)$, where n is the nonce. The EUF-CMA security is based on the failure of two events: inducing a collision in the AES round-based compression state, which is bounded by the active S-box

analysis, and the generic probability of guessing a valid tag. **PetitMAC** [88]: A lightweight MAC derived from the same design methodology as LeMAC, precisely tailored for applications where a high degree of parallelism offers limited benefit or for *short* messages. Whereas LeMAC uses nine parallel registers, PetitMAC uses a single AES round-based state together with five auxiliary 128-bit registers, processing one message block per iteration—making it a suitable choice for memory-constrained devices. Like LeMAC, PetitMAC performs message absorption, keyed finalization, state compression, and a masked nonce-dependent AES evaluation to generate the tag—their difference being mainly architectural. Consequently, the EUF-CMA security is analyzed using state collision attacks and the tag-guessing probability. **Badger** [59]: A CW-style MAC that is fast for long and short messages. At its core, a fast NH hash family hashes the message, and the result is masked with a PRF-derived pad (keyed and nonce-dependent). For long messages, Badger uses a tree structure. Multiple levels of NH hashing iteratively reduce the output until a short, fixed-size value remains, making the collision term the dominant contributor to the forgery bound in high-throughput environments. The EUF-CMA security is reduced to the almost-universal bound of the NH family, the PRF-distinguishing advantage, and the tag-guessing probability.

EliMAC [86]: A UHF-based deterministic MAC suitable for modern software platforms with AES-NI support, achieving high throughput in memory-constrained devices. At its core, EliMAC uses a parallel universal hash function EliHASH built over two UHFs, $\mathcal{H} : \{0, 1\}^{k_1} \times [1, \dots, n] \rightarrow \{0, 1\}^{k_2}$ and $\mathcal{I} : \{0, 1\}^{k_2} \times \{0, 1\}^b \rightarrow \{0, 1\}^b$, where $\mathcal{I}$ is an ϵ -AXU hash family, and $\mathcal{H}$ is an n -independent hash family which informally says that no n outputs of $\mathcal{H}$ are simultaneously predictable. For a message $m = m_1 \| \dots \| m_n$, EliHASH is a nested UHF evaluation defined as $\bigoplus_{i=1}^n \mathcal{I}(\mathcal{H}_{k_1}(i), m_i)$. To compute a tag, EliMAC feeds the EliHASH digest into a block cipher $E_{k_2} : \{0, 1\}^b \rightarrow \{0, 1\}^b$ and outputs a truncated ℓ -bit tag as $\text{Trunc}_\ell(E_{k_2}(\text{EliHASH}_{k_1}(m)))$. The EUF-CMA proof proceeds by replacing the final block cipher call with a random permutation (incurring a PRP advantage), and then shows that forgery is possible only if two distinct inputs collide; the n -independence of $\mathcal{H}$ and ϵ -AXU security of $\mathcal{I}$ bounds the latter.

Block Cipher-based MACs. Block ciphers are symmetric key primitives that operate on fixed-size blocks, where for each key $k \in \{0, 1\}^k$, the function $E_k : \{0, 1\}^b \rightarrow \{0, 1\}^b$ defines a permutation. A block cipher-based MAC runs a block cipher to compute a tag. Security relies on an ideal model where a block cipher behaves as a good PRP—*i.e.*, on an adversary’s inability to distinguish between a black-box for $E_k(\cdot)$ on a random k and a black-box on a random permutation $\pi(\cdot)$. Security is argued in the standard model under the PRP assumption, unlike permutation-based MACs, which use the ideal permutation model.

Block cipher-based MACs can be further categorized into CBC-MAC [185] and PMAC [186] types of constructions. As a canonical example of CBC-MAC, let $E_k : \{0, 1\}^b \rightarrow \{0, 1\}^b$ be a block cipher under a key k . Then,

for a message $m = m_1 \parallel \dots \parallel m_n$, with $|m_i| = b$, define $y_0 = 0^b$ and $y_i = E_k(y_{i-1} \oplus m_i)$ for $i \in [n]$. Then CBC-MAC is defined as: $\text{CBC-MAC} = E_k(m)[\ell] = \text{Trunc}_\ell(y_n)$, where $\ell \leq b$ is the tag length. Quite intuitively, since each block depends on the processing of the previous block, CBC-MACs are strictly sequential. Security is limited by the birthday bound due to collisions among hidden states. In particular, under the PRP assumption, the chaining states behave like random b -bit values, and the probability of collision is bounded by $\frac{q^2}{2^b} + 2^{-\ell}$, where the latter is the tag-guessing probability.

PMACs use fully parallelizable block ciphers for MACs. Let $E_k : \{0, 1\}^b \rightarrow \{0, 1\}^b$ be a block cipher. For a message $m = m_1 \parallel \dots \parallel m_n$, each block is first masked using a block-dependent offset $\Delta_i = \lambda_i \cdot L$, where $L = E_k(0^b)$. Each masked block is then independently encrypted as $Y_i = E_k(M_i \oplus \Delta_i)$. The intermediate value is computed as $S = \bigoplus Y_i$, and the final tag is obtained as $\text{tag} = E_k(S)$ (possibly truncated to ℓ bits). Finally, the EUF-CMA security of PMAC can be reduced to the PRP security of the underlying block cipher up to a loss of order $\frac{q^2}{2^b}$, where q is the total number of message blocks queried.

LightMAC [77]: A PMAC-type MAC used for short/medium-length messages, which is standardized by the ISO/IEC [76] for lightweight cryptography. Let $E : \{0, 1\}^b \rightarrow \{0, 1\}^b$ be a block cipher instantiated with independent keys k_1 and k_2 . Let ctr be a global counter of size s . Then, a message $m = m_1 \parallel \dots \parallel m_n$, where each block is $(b - s)$ bits, is encrypted as $Y_i = E_{k_1}(\langle \text{ctr}_i \rangle_s \parallel m_i)$, where $\langle \text{ctr}_i \rangle_s$ is the s -bit encoding of the counter ctr_i . The intermediate value is $S = ((\bigoplus_{i=1}^{n-1} Y_i) \oplus m_n 10^*)$, where the last message block is padded with 1 followed by zeroes so that the length of the bitstring becomes exactly b . The final step is to compute the tag as $E_{k_2}(S)$. The ctr guarantees that inputs to E_{k_1} are always distinct, and since E_{k_1} is a permutation, there are no collisions among Y_i for a fixed k_1 . Security is based on the linear dependence of the outputs, which occur with a probability of $\frac{q^2}{2^b}$.

LedMAC [81]: A more efficient variant of LightMAC that eliminates the inefficiencies in handling the last message block, achieving the same security level as LightMAC. Recall that LightMAC parses a message in blocks of $b - s$ bits, each concatenated with s -bit counter value, before being fed to E_{k_1} . However, the last message is padded to $m_n 10^*$ satisfying $|m_n| \leq b - s - 1$. Hence, the usable payload capacity of the final block is reduced by 1 bit. This redundancy in the final representation ensures that the encoding is prefix-free and that the final block is uniquely decodable. As a consequence, even if m is a multiple of $b - s$, an entire extra block of padding is appended. LedMAC eliminates this by using a single-bit domain separator—it appends a 0 if the last block is full and 1, otherwise. This allows the maximal length of the message to be $\leq 2^s(b - s) + s + 1$ bits as compared to $\leq 2^s(b - s)$ bits for LightMAC. The EUF-CMA security reduction replaces E_{k_1} and E_{k_2} with two independent random permutations. The remaining advantage is bounded by the probability of collisions.

LDMAC [82]: A parallel CBC-MAC suitable for low-data scenarios—the number of message blocks to be authenticated is small. The core idea is to split a state into sub-states and to apply a permutation to the sub-states. Let $E : \{0, 1\}^b \times \{0, 1\}^b \rightarrow \{0, 1\}^b$ be a block cipher and let $P : \{0, 1\}^b \rightarrow \{0, 1\}^b$ be a permutation. A secret N -bit state st is computed by running E_k on some fixed input. Split $\text{st} = \text{st}_1 \parallel \dots \parallel \text{st}_t$, where $|\text{st}_i| = b$. Then, for a message $m = m_1 \parallel \dots \parallel m_n$ where $|m_i| = b$, each b -bit part of tag is computed in parallel as: $\text{tag}_i = E_{k^{(i)}}(P(\dots P(P(\text{st}_i \oplus m_1) \oplus m_2) \dots) \oplus m_n)$, where $k^{(i)}$ is a b -bit key used for the i th block cipher call. The LDMAC security analysis does not rely on PRF security because its chaining structure, which uses an unkeyed permutation, leaks structural information. It allows birthday collisions after roughly $2^{b/2}$ samples per branch. Hence, if the total number of processed blocks is less, i.e., $q \cdot n \leq 2^{b/2}$, the collision probability stays negligible. Finally, the absence of internal collisions reduces an EUF-CMA forgery to guessing a $b \cdot t$ -bit tag. This reduction implicitly assumes the independence of t parallel branches, which is plausible given distinct st_i and $k^{(i)}$ per branch.

TuLP [69], [70]: A family of tunable lightweight MACs built on the PRESENT block cipher, designed primarily for wireless body area networks requiring strong authenticity under stringent computational, memory, and energy constraints. TuLP derives from a modified version of the ALFRED construction - a serial CBC-MAC style framework using only r rounds of PRESENT round function per message block, with a tunable parameter r offering a flexible security-performance trade-off. For a PRESENT encryption E , the state is initialized as $\text{st}_0 = E_k(IV)$, where IV is an initialization vector. Then, for a message $m = m_1 \parallel \dots \parallel m_n$, m_i is XORed with state st_{i-1} as a part of round key input, and fed into an r round E_k . Finally, the tag is extracted as $\text{Trunc}_\ell(\text{st}_n)$. The EUF-CMA security reduces to distinguishing PRESENT from a truly random permutation, the birthday collision term $\frac{q^2}{2^{64}}$ —collisions become likely after 2^{32} blocks—addressed by the TuLP-128 variant, which reduces this to $\frac{q^2}{2^{128}}$. The concrete security of TuLP depends on the strength of the r -round PRESENT instance against differential attacks—to be verified for the choice of r .

SMAC [92]: A standalone MAC construction based on the Finite State MACHINE (FSM) component of the SNOW stream cipher family. The key observation is that the non-linear AES-based FSM of SNOW-V can be isolated and repurposed as a MAC compression function Π , enabling the reuse of existing SNOW hardware implementations that achieve high throughput in both software and hardware environments. The construction uses three 128-bit registers (A_1, A_2, A_3) as its internal state updated as $\Pi(A_1^i, A_2^i, A_3^i, m_i) := A_1^{i+1} = \sigma(A_2^i \oplus A_3^i \oplus m_i)$, $A_2^{i+1} = \text{AESR}(A_1^i, m_i)$, $A_3^{i+1} = \text{AESR}(A_2^i, m_i)$, where σ is a 16-byte permutation, AESR is the AES round function, and $m = m_1 \parallel \dots \parallel m_n$, $|m_i| = 128$. In the initialization phase, registers are loaded with the key k and nonce n . In the compression phase, the message blocks are injected into the register update by applying Π to the full state per block. Finally,

tag is computed as $\text{tag} = \text{Trunc}_\ell(A_1^{n+d}, A_2^{n+d}, A_3^{n+d})$, where d dummy calls are additionally executed. The EUF-CMA security of SMAC is evaluated by its resistance to differential forgery attacks with a probability of $2^{-6\alpha}$ for α AES S-boxes, in addition to the tag-guessing probability of $2^{-\ell}$. Unlike formal reduction to standard hardness assumptions, the security argument is heuristic—based on the absence of known differential trials with probability exceeding $2^{-6\alpha}$.

◆ **PRF-based MACs.** PRFs are a family of functions $F_k : \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Y}$, such that for a uniformly random $k \in \mathcal{K}$, the efficiently computable function $F_k(\cdot)$ is computationally indistinguishable from a truly random function $\mathcal{R}(\cdot)$. PRF-based MACs instantiate tag generation with a direct evaluation of a pseudorandom function on a message, within minimal additional structure. Its security proof reduces to the underlying PRF security rather than relying on statistical collision bounds of a UHF or security properties of chaining modes. Formally, $\text{Adv}_{\text{MAC}}^{\text{EUF-CMA}}(\mathcal{A}) \leq \text{Adv}_F^{\text{prf}}(\mathcal{B}) + 2^{-\ell}$, where ℓ is the length of the tag. In practice, F_k is often instantiated using block cipher or Addition-Rotation-XOR (ARX)-based primitives. The overall security proceeds as: MAC to PRF security, and PRF to security of the underlying primitive (e.g., PRP for block ciphers). ARX-based constructions avoid PRP-PRF switching as they are designed directly as PRFs rather than being adapted from block ciphers.

MergeMAC [79]: A PRF-based MAC for time-invariant low-bandwidth platforms. The core intuition is to partition a message into low- and high-entropy components $m = m_1 \| m_2$, and process them independently using two keyed PRFs, F and $\tilde{F}$. The resulting intermediate output is then combined through a public Merge function: $\rho_1 = F_{k_1}(m_1)$, $\rho_2 = \tilde{F}_{k_2}(m_2)$, $\tau = \text{Merge}(\rho_1, \rho_2)$. This split-and-merge paradigm enables efficient caching of PRF evaluations for message components that remain unchanged across transitions, thereby reducing computation overhead during tag generation. Additionally, the low-entropy component of m is not communicated to save bandwidth. Verification is an exhaustive search until a valid tag is obtained, constraining the space of low-entropy content and shifting computational burden to the receiver. The EUF-CMA security follows from: (i) F and $\tilde{F}$ are secure PRFs, and (ii) Merge satisfies a property—on uniformly random inputs, the output distribution is computationally indistinguishable from random. Concretely, F and $\tilde{F}$ are instantiated with PRESENT or PRINCE in the CMAC mode, and Merge with a Chaskey variant—the construction was later broken by two attacks: (a) a generic universal forgery adapted from CBC-like MAC attacks, and (b) a low data universal forgery that exploits the weak diffusion of a 3-round Chaskey [94].

SipHash [71]: A family of PRFs optimized for short inputs of a few byte—hash table keys, packet headers, etc. The construction uses a 256-bit state by XORing a 128-bit key with public constants. The message blocks are processed sequentially using ARX mixing rounds, accumulating the message bits into the state. After all blocks are processed, a domain separator initiates the final phase, which ensures additional mixing rounds before the state is collapsed into a

64-bit tag. The EUF-CMA security reduces to breaking the PRF security and tag guessing.

ShMAC [67], [68]: A PRF-based MAC designed to authenticate memory that may be controlled by an adversary. A precomputation phase is performed by evaluating F on nonce n known a priori, where F is instantiated with a full AES evaluation (10 rounds). Online authentication reduces to a lightweight computation $\text{tag} = G(m \oplus F(n))$, where G is a permutation instantiated with a round-reduced AES, implementable as a small-depth circuit. Security is based on SDU—a property of G ensuring that for any non-zero input difference (known to the adversary), the output difference distribution is uniformly bounded with overwhelming probability. The EUF-CMA security reduces to breaking the PRF. If this succeeds, then the adversary knows the entire input $(m + F(n))$ to G —violating the SDU assumption, where an adversary only knows the input difference.

◆ **Permutation/Sponge-based MACs.** Permutation or Sponge-based MACs are based on a public cryptographic permutation, which is an unkeyed bijection $\pi : \{0, 1\}^b \rightarrow \{0, 1\}^b$. Concretely, the permutation is designed to achieve diffusion, confusion, and resistance to structural cryptanalysis. In sharp contrast to keyed block cipher-based MACs, where the secret key forms an intrinsic input to the primitive and requires dedicated key-scheduling algorithms, often increasing implementation complexity, hardware cost, and latency. Permutation-based MACs are commonly realized through keyed permutation constructions, such as the Even-Mansour (EM) paradigm. It considers a permutation $\pi(m \oplus k_1) \oplus k_2$ —where keys are XORed before and after applying π . The EUF-CMA security is argued in the ideal permutation model— π is modeled as a truly random permutation accessible to all parties, including the adversary. Any successful forgery must find a key-dependent internal state collision, i.e., produce a message that, when XORed with the key and processed through π , yields a valid tag. The forgery probability is bounded by the birthday term over the internal state.

In contrast, sponge-based constructions maintain a large internal state, divided into rate r and capacity c components. The key is absorbed into the initial state; the message blocks are XORed into the rate portion iteratively, and the full state is updated through repeated applications π after each block. Finally, the tag is extracted by truncating the final r , while c is hidden—it accumulates a hidden mixture of the key and message blocks through π . An adversary must produce two different messages that produce the same full internal state, including the hidden c -bit component. Since c is never visible, finding a collision requires $2^{c/2}$ queries.

Chaskey12 [74]: Chaskey is a permutation-based MAC optimized for 32 bit processors, and intended for applications requiring 128-bit security. The construction transforms an unkeyed ARX permutation π into a keyed block cipher via the EM paradigm: $E_{k||k}(m) = \pi(m \oplus k) \oplus k$ used in a CBC-MAC mode. Message $m = m_1 \| m_2 \| \dots \| m_n$ is processed as $E_{k||k}(h_i \oplus m_i)$, while the last block uses modified keys derived from subkeys. The security proof is an EM argument extended to a CBC-MAC setting: the $E_{k||k}$ evaluations are

indistinguishable from a random permutation in the ideal world, after which the construction becomes a CBC chaining of truly random permutations. Then the primary attack is to find two messages that collide in the internal chaining values. Differential-linear attacks were shown on 6- and 7-round variants of Chaskey, indicating that the original 8 round construction offers a limited security [187]. In response, Chaskey12 was proposed, which restores an adequate security margin.

ASCON-MAC [90]: A sponge-based MAC built on top of a 320-bit ASCON permutation—the primitive underlying the ASCON family selected by NIST as the lightweight cryptographic standard. Unlike the classical sponge-based constructions, where the key is absorbed into the r portion of the state, and the security is derived from the hidden c , the ASCON-MAC is instantiated as an ASCON PRF with a fixed 128-bit output that allows the key to span across both r and c —reducing the amount of publicly known state information to the adversary. The EUF-CMA forgery is bounded by the probability that an adversary correlates MAC and permutation queries to match the hidden c component, together with the probability of guessing a valid tag.

C Ethics Considerations

This paper surveys publicly available cryptographic constructions and protocol standards. While we discuss security shortcomings of deployed MAC schemes, this information is well-known within the cryptographic community and does not provide meaningful uplift to adversaries. Our goal is instead to raise awareness among system operators and designers, who may be unaware of the security implications of current deployment practices.

D Meta-Review

The following meta-review was prepared by the program committee for the 2027 IEEE Symposium on Security and Privacy (S&P) as part of the review process as detailed in the call for papers.

D.1. Summary

This SoK provides a comprehensive analysis of MACs for constrained M2M communication, by surveying authentication mechanisms across various deployment domains. In particular, it explores ten deployment domains and identifies 20 lightweight MAC schemes via a systematic literature review, and evaluate 13 on embedded hardware. The evaluation shows that, for short messages, sw-implemented lightweight MACs can outperform hardware-accelerated schemes and authenticated encryption. It concludes with open challenges and actionable guidance.

D.2. Scientific Contributions

- Independent Confirmation of Important Results with Limited Prior Research
- Provides a Valuable Step Forward in an Established Field
- Creates a New Tool to Enable Future Science
- Addresses a Long-Known Issue

D.3. Reasons for Acceptance

- 1) The paper provides a valuable comprehensive analysis on M2M environment requirements and the lightweight MAC schemes suitable for constrained deployments. In particular, the authors provide actionable insight, mapping M2M operational requirements to MAC scheme design properties.
- 2) The research community has spent a decade optimizing AEAD while the protocols that actually ship in constrained networks are still running CMAC-AES and HMAC-SHA256, and the paper provides evidence for this gap. Several of the findings in the paper are counterintuitive (*e.g.*, that software lightweight MACs beat hardware-accelerated CMAC and HMAC at short inputs), and directly actionable.