

Practical Non-interactive Multi-signatures, and a Multi-to-Aggregate Signatures Compiler

Matthieu Rambaud
Télécom Paris

Christophe Levrat
Inria Saclay

Abstract—In a fully non-interactive multi-signature, resp. aggregate-signature scheme (fNIM, resp. fNIA), signatures issued by many signers on the same message, resp. on different messages, can be succinctly “combined”, resp. “aggregated”. fNIMs are used in the Ethereum consensus protocol, to produce the certificates of validity of blocks which are to be verified by billions of clients. fNIAs are used in some PBFT-like consensus protocols, such as the production version of Diem by Aptos, to replace the forwarding of many signatures by a new leader. In this work we address three complexity bottlenecks. (i) fNIAs are costlier than fNIMs, e.g., we observe that verification time of a 3000-wise aggregate signature of BGLS (Eurocrypt’03), takes 300x longer verification time than verification of a 3000-wise pairing-based multisignature. (ii) fNIMs impose that each verifier processes the setup published by the group of potential signers. This processing consists either in verifying proofs of possession (PoPs), such as in Pixel (Usenix’20) and in the IETF’22 draft inherited from Ristenpart-Yilek (Eurocrypt’07), which costs a product of pairings over all published keys. Or, it consists in re-randomizing the keys, such as in SMSKR (FC’24). (iii) Existing proven security bounds on efficient fNIMs do not give any guarantee in practical curves with 256bits-large groups, such as BLS12-381 (used in Ethereum) or BLS12-377 (used in Zexe). Thus, computing in much larger curves is required to have provable guarantees.

Our first contribution is a new fNIM called *dms*, it addresses both (ii) and (iii). It is as simple as adding Schnorr PoPs to the schoolbook pairing-based fNIM of Boldyreva (PKC’03). (ii) For a group of 1000 signers, verification of these PoPs is: 5+ times faster than for the previous pairing-based PoPs; and 3+ times faster than the Verifier’s processing of the setup in SMSKR (and contrary to the latter, needs not be re-started when a new member joins the group). (iii) We prove a tight reduction to the discrete logarithm (DL), in the algebraic group model (AGM). Given the current estimation of roughly 128 bits of security for the DL in both the curves BLS12-381 and BLS12-377, we deduce a probability of forgery of *dms* no higher than about 2^{-93} for a time 2^{80} adversary. This reduction is our main technical contribution. The only related proof before was for an interactive Schnorr-based multi-signature scheme, using Schnorr PoPs. Our approach easily fills a gap in this proof, since we take into account that the adversary has access to a signing oracle even before publishing its PoPs. But in our context of pairing-based multi-signatures, extraction of the keys of the adversary is significantly more complicated, since the signing oracle produces a correlated random string. We finally provide another application of *dms*, which is that it can be plugged in recent threshold signatures without setup (presented by Das et al at CCS’23, and Garg et al at SP’24), since these schemes implicitly build on any arbitrary BLS-based fNIM.

Our second contribution addresses (i), it is a very simple compiler: *MtoA* (multi-to-aggregate). It turns any fNIM into an fNIA, suitable for aggregation of signatures on messages with a prefix in common, with the restriction that a signer must not sign

twice using the same prefix. The resulting fNIA is post-quantum secure as soon as the fNIM is, such as Chipmunk (CCS’23). We demonstrate the relevance for Diem by applying *MtoA* to *dms*: the resulting fNIA enables to verify 39x faster an aggregate of 129 signatures, over messages with 7 bits-long variable parts, than BGLS.

I	Introduction	1
II	Preliminaries	4
III	New definitions	5
IV	<i>MtoA</i> : Multi to Aggregate Compiler	7
V	<i>dms</i>	7
VI	Proof of Theorem 4	9
VII	Evaluation and Comparison	11
VIII	Acknowledgements	13
	Appendix A: Further formalization and Optimization of <i>MtoA</i>	16
	Appendix B: Details for Applications of <i>dms</i>	16
	Appendix C: Details for application of <i>MtoA</i> to consensus	17
	Appendix D: Further details on related works	19

I. INTRODUCTION

In an aggregate signature scheme [23, 15, 77, 37, 73, 59, 25, 63, 30, 45, 68, 80, 89], a single short string takes the place of n individual signatures by n signers on n messages. Our focus is what we call *fully non-interactive* aggregate signature schemes (fNIAs) [23, 15, 37, 73, 63, 45, 68, 80, 89]. They offer an aggregation algorithm Ag which, roughly, takes as input any multiset of triples of public key - message - valid individual signature: $(pk_i, m_i, \Sigma_i)_{i \in [n]}$, and outputs a single

aggregate signature Σ . Finally, there is a public verification algorithm $\forall f$ which takes as input a purported aggregate signature Σ with respect to a multiset of pairs of public key - message: $(pk_i, m_i)_{i \in [n]}$, and outputs a bit denoting acceptance or rejection. Such schemes have the unforgeability property, that acceptance implies that for any pair (pk, m) in the multiset such that the key \hat{X} was generated by some honest process, then it must have signed the corresponding message \hat{m} . *Fully non-interactive multisignature schemes* (fNIMs) enable aggregation only over *identical* messages. We then denote N the number of messages, instead of n , and dub aggregation the combination algorithm: $Cb : m, (pk_i, \Sigma_i)_{i \in [N]} \rightarrow \Sigma$. Pairing-based fNIMs are used in the consensus protocol of Ethereum [48, 50], they enable clients to verify that a block was voted by enough validators [3].

Goal (i): reducing the efficiency gap between the verification times of fNIAs and of fNIMs. In all pairing-based fNIAs [15, 37, 73, 68, 80], following BGLS [23], the complexity of the Verifier is at least a product of $n+1$ pairings. By contrast, in nearly all pairing-based fNIMs [20, 77, 85, 21][22, §6][47, 56, 9], the online verification complexity of the Verifier is mostly two pairings. Our terminology “online” is to differentiate from the other task of the Verifier, which we will call “processing of the group setup” and discuss separately. Concretely, in Table 8 we observe that the online verification time of the verifier in any existing pairing-based fNIA, for 3000 signatures over different message contents, is at least $300\times$ higher than the online verification of any of the previous pairing-based fNIMs over 3000 signatures on identical message contents. Turning to lattice-based fNIMs, the state of the art called Chipmunk [52, 51] enjoys a verification $6.5\times$ faster for $N = 8192$ than the naive verification of individual Falcon signatures. Whereas, there exists no public evaluation, to our knowledge, of the Verifier runtimes of the state of the art lattice-based fNIAs [89, 1]. They consist of SNARKs of signatures, using the system called Labrador [19], which has linear Verifier complexity. Apart from them, the fNIA [45] was recently broken [29].

This efficiency gap is best illustrated by blockchain consensus algorithms, say, among n_C processes. The fastest ones are known as “leader-based”: [61, 71, 58, 44] (their liveness requires partial synchrony). The most recent implementation used in production is the one of Facebook’s Diem21 ([44]) by Aptos. It proceeds by iterations called “rounds”. Under good conditions, the leader of a new round in Diem21 only has to combine $(2/3)n_C$ identical votes with a fNIM, into a multi-signature dubbed a “quorum certificate” (QC), which it multicasts. But if the leader of the previous round was corrupt or the network not synchronous, then the current leader must aggregate $(2/3)n_C$ signatures over *different* so-called “timeout messages”. It multicasts the aggregate signature, called a “timeout certificate”. Aggregation is done by Aptos ([5]) with BGLS. Hence, already for $(2/3)n_C = 129$, we observe in Table 8 that verification of a timeout certificate is $49\times$ slower than verification of a QC. For convenience we recall Diem21 in Figure 9, for simplicity with aggregation instead instantiated as a naive concatenation of signatures.

Goal (ii): reducing the complexity of processing of the group setup. All pairing-based fNIMs [85, 24, 47, 21, 56, 9] require costly additional tasks from the Verifier. Namely, the group of potential signers must initially publish the outcome

of their group setup, which we call generically the “keys of the group” and denote \mathcal{KG} . The Verifier must then process \mathcal{KG} , we call this task *processing of the group setup*.

In a first category of fNIMs ($MSP-pop$ [85, 22, 24], Pixel and $ASMP-pop$ [85, 47, 22, 24]), each signer incorporates a so-called *proof of possession* (PoP): π into its public key: $pk = (X, \pi)$. The purpose of the PoP is to enforce (possibly with a loose reduction) that its issuer “knows” a secret key corresponding to X . PoPs thus somehow emulate the model called “knowledge-of-secret-key” (kosk). The kosk assumes that the adversary gives to the reduction a secret key for every public key appearing in its forgery (excepted the target one). Interest of the kosk model is that the security of the two fastest known fNIMs: [20, 77] is proven only in the kosk. The fNIM of Boldyreva [20] has verification complexity of only two pairings, but without the kosk it is vulnerable to so-called “rogue key attacks” [21]. There, the adversary creates a forgery involving public keys, other than the target one, for which it does not know corresponding secret keys. In all previous works, the PoP π is equal to a pairing-based signature (BLS [26]) on the public key itself. The Verifier then has to verify the PoPs of all the keys: \mathcal{KG} of the group of signers. In Table 6 we estimate that their (batched & optimized) verification for an 2702-sized group takes 1947ms on a laptop. As clear from Table 8, this time is orders of magnitude longer than the online verification of a multi-signature.

In a second category of fNIMs ($ASMP/ASMP-pop$ [22], SMSKR [9] and SIG_1 [27]), the processing of the group setup consists in (re-)computing a so-called verification key for the group, out of the list of their published keys: \mathcal{KG} . In $ASMP$ and SIG_1 this verification key is of constant size (at the cost of an interactive setup). While in SMSKR, each signer in the group re-randomizes its secret signing key based on all other published keys: \mathcal{KG} . The Verifier then has to compute the re-randomized public keys accordingly: these will be the ones used for verifying signatures (both individual and combined). Again, as evidenced in Tables 6 and 8, this task is the bottleneck of the Verifier since it takes three orders of magnitude longer than the online verification of signatures. Worse: unlike verification of PoPs, the group key must be re-computed each time there is a new group member, since re-randomization depends on the list of published keys: \mathcal{KG} .

A last category of fNIMs (the blog version $MSP-blog$ [21, 56]) does not require processing of the group setup tasks, beyond verification of membership of the keys in the subgroup \mathbb{G}_2 . But its online verification requires to compute a combined verification key, equal to the sum of the re-randomizations of the public keys of the signers. We note that this computation is comparatively faster than computing re-randomized keys separately, since it can be done in one single N -sized multi-exponentiation (plus N times $N\kappa$ -sized hashes).

Goal (iii): achieving provable security for use with the curves used in practice. In Table 7 we observe that no existing fNIM is proven safe to use with the curves used in practice: BLS12-381, adopted by Ethereum, and BLS12-377, proposed by Zeux [31]. Despite these curves having a discrete-log-in-subgroup (DL) problem of estimated security close to 128 bits [70, 6], when instantiating with them the fNIM $MSP-pop$ [85, 22, 24] (IETF draft standard 2022), we find that the proven

bound on the probability of a forgery after 2^{80} clock cycles is higher than 2^{-13} . This estimate is optimistic, since under the assumption that co-cdh would be as hard as DL. Other pairing-based fNIMs [47, 21, 56, 9] are not either proven secure with such curves (the proven formulas would give an upper-bound higher than 1). Finally, the lattice-based fNIMs [52, 51] have public keys of logarithmic size in the number of signatures allowed, since they are equal to commitments to vectors of one-time public keys. A first exception is the fNIA called $\mathcal{AS}\text{-}4$ (long version of Bellare-Namprempre-Neven [15], following a trick of Katz-Wang [72]). But its verification complexity is prohibitive, since it is a product of $N + 1$ pairings, as in all related fNIA schemes. The second and last exception is the recent scheme SIG_1 of [27], but which has a quadratic Verifier’s processing of keys (see Sec. VII and below). So this raises the natural question of finding a fNIM with tight reduction to the discrete logarithm (DL), and thus usable with practical curves. This has been achieved so far with multisignatures [42, 13] requiring two rounds of interaction, and thus which are not usable in consensus protocols such as Diem21 [44].

A. An Efficient Multi-to-Aggregate Compiler

We address Goal (i) by introducing \mathcal{MtoA} : it is a compiler which transforms any fNIM into a special-purpose fNIA. The resulting fNIA applies to messages divided into two parts: a common prefix τ , called the *tag*, and the remaining message v_i , of bitlength denoted $|v|$. The resulting fNIA is particularly efficient when the length $|v|$ of the variable contents v_i is only of a few bits. \mathcal{MtoA} operates as follows. Each signer prepares and publishes $2|v|$ public verification keys. To sign a message $m_i = (\tau, v_i)$, it signs τ $|v|$ times: each time using the public key indicated by the j -th bit of v_i , for $j \in [|v|]$. As all the signatures of all signers are on the same τ , the verification cost is equal to verifying a combined multisignature. We see that a signer should not sign two different messages: v_i and v'_i that share the same tag τ (otherwise the adversary could cherry-pick signed bits from both v_i and v'_i to forge another signature). For this reason, the resulting fNIA is called “one-time-tagged”.

1) Performance and Applications of \mathcal{MtoA} ; and related works: We resume the example of the production implementation [5] of Diem21 [44] by Aptos. A timeout message of a process \mathcal{P}_j is formatted as the signed message: $\{r, r_j\}_j$. The tag r is the current round number, while $r_j < r - 1$ is roughly the highest round number in which \mathcal{P}_j saw a QC: $qc_{high,j}$. Hence, r_j can be advantageously encoded as the difference $v_j := r - r_j - 1$ (this improves by -1 a nice idea of [60]). Thus, after the network becomes synchronous and assuming that each leader is honest with probability $2/3$, then the expected value of v_j is only 0.5. Hence, the timeout messages fall in the regime where \mathcal{MtoA} is efficient. In Table 8 we report on the Verifier’s runtime for an \mathcal{MtoA} aggregate over 129 signers, calibrated with a variable parts of bitlength $|v| = 7$ (which is overkill for Diem21, by the above considerations). Thus, verification consists of verifying a $N = 7n$ -wise multisignature. We used as input any of the three BLS-based fNIMs: $\mathcal{MSP}\text{-}pop$, SMSKR and our dms (below), which have the same online verification complexity for a given curve. We achieve an online verification runtime of 3.4ms. It

is close to the batch verification of a naive concatenation of 129 Schnorr signatures. This is $39\times$ faster than the verification time: 116.6ms of an 129-wise BGLS aggregate signature, as used so far in the implementation of Diem21 [5].

One-time-tagged fNIA schemes were considered in [4, 69, 60], but all based on non-post-quantum assumptions (the latter inspired \mathcal{MtoA}). By contrast, \mathcal{MtoA} has post-quantum security whenever applied to any post-quantum fNIM, e.g. to [51]. Another use-case where \mathcal{MtoA} is advantageous is suggested by [4]: they consider connected devices signing short measurements (such as the temperature, weight, speed), one-time-tagged with the time of the measure.

B. A Faster and Tightly Secure fNIM

We introduce a fully non-interactive multisignature scheme which achieves both goals (ii) and (iii), called Dynamic Multisignature with Schnorr proofs of possession (dms). It is as simple as augmenting the pairing-based (BLS) multisignature scheme of Boldyreva [20] with proofs of possession (PoP) consisting of Schnorr signatures of processes on their public keys. Since batch verifying N Schnorr signatures [14] is much faster than computing $N + 1$ pairings, it is not surprising that the verification of PoPs in dms achieves a $> 5\times$ speedup over the most optimized verification of the ones of $\mathcal{MSP}\text{-}pop$ [85, 22, 24]: see Table 6. As any PoP-based fNIM, dms comes with the bonus of being *dynamic*, i.e., the signing algorithm needs not taking as input a *group* of potential signers \mathcal{KG} . In turn, individual signatures can be combined without the restriction that the signers agreed together on some common group \mathcal{KG} . Finally, as in any PoP-based fNIM, when new potential signers, say, 14, publish their keys, the marginal cost of the Verifier is only to verify 14 PoPs. In Table 6 we evaluated this task to be $> 500\times$ faster than the Verifier’s processing in SMSKR when new members join the group, since it must re-randomize all published keys.

a) Proving tight reduction of dms to DL (in the AGM):

The proof does not follow from previous works, it is one of our main contributions. The literature of the last 25 years suggests that mere Schnorr ZK PoKs of secret keys might not simply provably thwart rogue key attacks ([79] Micali-Ohta-Reyzin “That is, for the simulator to be polynomial time, there can be at most logarithmically many signers” and Bellare-Neven [16] “one would require ZK PoKs extractable under such concurrent conditions. This eliminates many standard protocols, including standard POKs of discrete logarithms.”). Worse, the attempt [8] had its proof invalidated by [46], and finally a proof attempt, on a related issue (CCA security from Schnorr PoKs of randomness) [18], required a sophisticated revisiting [54], despite being in the algebraic adversary model. We consider the *algebraic group model* (AGM), in line with [83, 2, 18, 54, 53, 11, 12, 13, 9, 7, 41]. Namely, our security bounds hold against so-called *algebraic adversaries*. Heuristics partly supporting the AGM are that, according to [70, §4][6], no better attacks are known against DL in BLS12-377/381, than the generic square-root algebraic one (see Sec. VII-2).

The main difficulty of our proof consists, upon being given a forgery w.r.t. a set of public keys $(X_i, \pi_i)_{i \in [N]}$, in extracting from the forger all the secret keys, i.e., ξ s.t. $X_i = \xi \cdot G_2$ where G_2 is a public generator of the subgroup \mathbb{G}_2 (except the one

of the challenge public key \hat{X}). The last step from there is that the reduction builds a valid individual signature for \hat{X} on the challenge message, thus breaking unforgeability of BLS. This last step is as in [20], with the minor twist that the set $(X_i)_{i \in [N]}$ could well contain multiple copies of \hat{X} , without the forger giving \hat{x} to the reduction. The *main difficulty* is singled-out in Sec. VI as an abstract extractability game, called SSC (Schnorr straight-line extraction despite correlations). It considers an adversary having access to a signing oracle for the BLS signature of a given key X (playing the role of \hat{X}). Then it outputs one or several keys appended with valid proofs of possession: $(X^*, R^*, z^*), \dots$ (other than the one of the challenge). It wins the game if the challenger fails to extract in straight-line one of the corresponding secret keys x^* . The game is not comparable to [53, §6], where the adversary had instead to forge a BLS signature. In Theorem 4 we show that the advantage of the adversary is bounded by the advantage in the DL game. The proof is strictly more difficult than in [42, §A], since their adversary does not have access to a signing oracle for the target key \hat{X} . Absence of this oracle in [42, §A] invalidates the security proof of their multisignature (page 10, step “key registration”), because in the unforgeability definition ([42, §5.1] and our Sec. III) the adversary *has* the power to query such an oracle *potentially before* it chooses the set of keys of its forgery. This power models that the adversary could, in practice, engage in signing sessions with concurrent groups containing the same target key \hat{X} , before it registers the set of keys $(X_i)_{i \in [N]}$ of its forgery. On the one hand, we observe that [42, p10] can easily be fixed (we notified this on 01-23 2024 to the authors). Indeed, the proof of [42, §A] would go unchanged after adding the necessary Schnorr-signing oracle, since it would return uniformly random group elements. Whereas our signing oracle produces a *correlated* sequence of group elements: $\{(H(m_j), \hat{x} \cdot H(m_j))_j, \hat{X} = \hat{x} \cdot G_2, G_2\}$, j running over the signing queries. These correlations make necessary for our reduction to DL to follow two alternative behaviors (the second, called D, is designed to cope with the event which we call “very bad”, defined in Eq. (16)), instead of one single behavior in [42, §A].

b) Applications of dms: It enables to divide by N the online storage size of certificates of validity for blocks [48, 50, 3], compared to a batch of Schnorr signatures. Although fNIMs are already used in Ethereum for this purpose, Table 7 shows that previous fNIMs were proven secure only in curves much larger than BLS12-381 or -377. This would imply non-standard curves, larger storage size, and longer verification time. Moreover the $5 \times$ processing of the group setup speedup of dms (for equal curves) directly impacts billions of verifiers. Contrary to a common belief, individual verification of signatures of a fNIM can be made faster than the one of Schnorr signatures. *First*, as observed in [38] and confirmed by Table 8, *batch verification* of BLS individual signatures (as used in *MSP-pop*, *SMSKR*, *dms*) is $3 \times$ faster than batch verification of Schnorr signatures, for $n = 3073$ signatures. The Verifier simply combines the signatures (the cost of these n additions is negligible) then verifies the obtained multisignature. Batch verification virtually always succeeds in use-cases such as signatures published in blocks (otherwise, the validators of the blocks would be severely punished). *Second*, for use-cases where invalid signatures often occur, then it was observed by [56, 34] that signers can add a Chaum-Pedersen proof of

equality of discrete logs to their BLS signatures, enabling an individual verification time comparable to the one of a Schnorr signature. We give further details in Sec. B. In Sec. B we also explain how dms can be advantageously plugged in the threshold signature schemes [43, 57], in place of *MSP-pop*.

II. PRELIMINARIES

We use the formalism of games [17], with the simplification that we merge the finalization in the main body, as in [53], and that we use the more mainstream meaning of the *advantage of an adversary \mathcal{A} in a game \mathbf{g}* [12], denoted $\mathbb{P}(\mathbf{g}^{\mathcal{A}} = 1)$, to designate the probability that \mathcal{A} wins the game, i.e., that \mathbf{g} sets the flag win $\leftarrow 1$.

a) Bilinear groups: A bilinear group description ([55]) is a tuple $\mathbb{G} = (\mathbb{G}_1, \mathbb{G}_2, G_T, e, \phi, \psi, p)$ such that \mathbb{G}_i is a cyclic group of prime order p for $i \in \{1, 2, T\}$, in additive notation; e is a non-degenerate bilinear map $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$, i.e., for all $a, b \in \mathbb{Z}_p$ and all generators G_1 of \mathbb{G}_1 and G_2 of \mathbb{G}_2 we have that $G_T := e(G_1, G_2)$ generates \mathbb{G}_T and $e(a \cdot G_1, b \cdot G_2) = ab \cdot e(G_1, G_2) = ab \cdot G_T$; $\phi : \mathbb{G}_1 \rightarrow \mathbb{G}_2$ is an isomorphism, and $\psi : \mathbb{G}_2 \rightarrow \mathbb{G}_1$ is an isomorphism. All group operations and the bilinear map e must be efficiently computable. \mathbb{G} is of Type 1 if the maps ϕ and ψ are efficiently computable, in which case we will consider without loss of generality that $\mathbb{G}_1 = \mathbb{G}_2$; \mathbb{G} is of Type 2 if there is no efficiently computable map ϕ ; and \mathbb{G} is of Type 3 if there are no efficiently computable maps ϕ and ψ .

In line with [37], all our assumptions and statements are with respect to a choice of fixed public generators G_1, G_2 . Note that since we are in the random oracle model, we have a fortiori a public uniform random string (URS). So G_1, G_2 could be fixed by seeding any public uniform sampling algorithm with the URS.

b) The algebraic group model (AGM): In line with [83, 2, 18, 54, 53, 11, 12, 13, 9, 7, 41], we consider provable security against adversaries known as *algebraic algorithms*. We recall the most recent model in the setting of bilinear groups, from [11, Def 2]. An algorithm \mathcal{A} executed in a security game ([17]) is called algebraic if for all group elements $Z \in \mathbb{G}_i$, $i \in \{1, 2, T\}$ that \mathcal{A} outputs to any oracle of the game, it additionally provides a representation in terms of received group elements in \mathbb{G}_i and those from groups from which there is an efficient mapping to \mathbb{G}_i . In particular for $Z \in \{\mathbb{G}_1, \mathbb{G}_2\}$: if $U_0, \dots, U_\ell \in \mathbb{G}_1$ and $V_0, \dots, V_m \in \mathbb{G}_2$ denote the group elements in \mathbb{G}_1 and \mathbb{G}_2 received so far by \mathcal{A} , then it provides a list of coefficients in \mathbb{Z}_p $(\mu_i)_{i \in [\ell]}$, and possibly $(\zeta_j)_{j \in [m]}$, such that, depending on the case:

- $Z \in \mathbb{G}_1$ (Type 1 and 2): $Z = \sum_i \mu_i U_i + \sum \zeta_j \psi(V_j)$
- (Type 3): $Z = \sum_i \mu_i U_i$
- $Z \in \mathbb{G}_2$ (Type 1): $Z = \sum_i \mu_i U_i + \sum \zeta_j V_j$
- (Type 2 and 3): $Z = \sum_j \zeta_j V_j$

c) Signatures, example of BLS: We recall the standard notion of a digital signature scheme with existential unforgeability under chosen message attacks (EUF-CMA, [62, 20]). It is the data of algorithms for key generation Kg , signature Sign and verification Vf , the latter returning a bit denoting acceptance or rejection. They furthermore have the following properties. Correctness requires that $\forall (\text{sk}, \text{pk}) \leftarrow \text{Kg}(), \forall m,$

$\text{Vf}(\text{pk}, m, \text{Sign}(\text{sk}, m)) = 1$. Unforgeability requires that a polynomial adversary \mathcal{B} , dubbed forger, is unable to forge valid a signature on a message m for which it did not query a signature. More formally it is defined by a game, which for concreteness we exemplify in Figure 1 on the example of the well-known BLS signature scheme [26, 37]. BLS is parametrized by a bilinear group with public generators: $G_1 \in \mathbb{G}_1$, $G_2 \in \mathbb{G}_2$, and by a hash-to-curve map $m \rightarrow \text{H}(m) \in \mathbb{G}_2$ which we model as a random oracle.

- $\text{BLS.Kg}()$: sample $x \xleftarrow{\$} \mathbb{Z}_p$, output $(\text{sk}, \text{pk}) = (x, x.G_2)$;
- $\text{BLS.Sign} : m \rightarrow \text{sk.H}(m)$;
- $\text{BLS.Vf} : \text{pk}, m, \Sigma \rightarrow e(\Sigma, G_2) == e(\text{H}(m), X)$.

BLS-uf	
$x \xleftarrow{\$} \mathbb{Z}_p, X = x.G_2$	
$(m^*, \Sigma^*) \leftarrow \mathcal{B}^{\text{SIGN}, \text{H}}(X)$	
return $(m \notin Q_{\text{sig}} \wedge e(\Sigma^*, G_2) == e(\text{H}(m^*), X))$	
oracle $\text{H}(m)$	oracle $\text{SIGN}(m)$
return $M \leftarrow \text{H}(m)$	$Q_{\text{sig}} \leftarrow Q_{\text{sig}} \cup \{m\}$
	return $x.H(m)$

Figure 1: EUF-CMA game for BLS

In the random oracle (RO) model, the security of BLS has a loose reduction to the computational Diffie Hellman (CDH) problem. More precisely, over bilinear groups of type I and II it has a loose reduction [26] to the problem called co-cdh: given $(Q \xleftarrow{\$} \mathbb{G}_1, a.G_2)$ for $a \xleftarrow{\$} \mathbb{Z}_p$, compute $a.Q$. Whereas over bilinear groups of type III, it has a loose reduction [37] to the problem called co-cdh*: given $(Q \xleftarrow{\$} \mathbb{G}_1, a.G_1, a.G_2)$ for $a \xleftarrow{\$} \mathbb{Z}_p$, compute $a.Q$. Note that in BDN18 [22], co-cdh and co-cdh* are respectively renamed ψ -co-cdh and co-cdh. In the RO+AGM model, the security of BLS has a tight reduction [53, §6] to the discrete logarithm (DL) in \mathbb{G}_1 . We observe that over type II groups, a minor adaptation of the proof shows a tight reduction to the DL in \mathbb{G}_2 (use ψ to port the DL challenge in \mathbb{G}_1). We observe that over type III groups, the same proof shows a reduction to the type III variant of DL known as (1,1)-DL [11]: given $(a.G_1, a.G_2)$, compute a .

III. NEW DEFINITIONS

We first define fully non-interactive multisignature schemes (fNIM), then fully non-interactive one-time-tagged aggregate signature schemes (fNIA), which are output by our compiler \mathcal{M} to \mathcal{A} . Our definitional choices aim at capturing most existing non-interactive schemes, capturing the remaining ones would require only straightforward adaptations. Our syntax is meant to be safely used as such by the Combiner/Aggregator or the Verifier, no hidden additional check is required from them. In the last subsection we discuss our choices w.r.t. related specifications. All algorithms take as additional input a multiset of public keys denoted \mathcal{KG} , called the keys of the group (of potential signers). A scheme is called dynamic when this input plays no role. In turn, all \mathcal{KG} -dependent restrictions in the specifications and security definitions are then removed. For simplicity we formalize our definitions and results with schemes in which the setup is non-interactive (provided implicit coordination within the group \mathcal{KG} for non-dynamic

schemes). Our compiler \mathcal{M} to \mathcal{A} obviously extends to schemes with interactive setup, of which the three existing ones are [77], \mathcal{ASMP} [22, §4] and $\mathcal{ASMP-pop}$ [22, §6]. Both n and N denote arbitrary positive integers. The definitions involve only n - (or N -)sized multisets of tuples, such as $(\text{pk}_i, \Sigma_i)_{i \in [N]}$ or $(\text{pk}_i, v_i, \Sigma_i)_{i \in [n]}$, i.e., without ordering between each other. So our indices 1, i etc. are just here to name elements. The numbers N and n of messages to be combined/aggregated are variable inputs to the algorithms.

Fully Non-Interactive Multi-signatures

Definition 1 (fNIM). A fully non-interactive multisignature scheme consists of the following local algorithms.

- $\text{Kg}() \rightarrow (\text{sk}, \text{pk})$ (key generation)
- $\text{Sign}(\mathcal{KG}, \text{sk}_i, m) \rightarrow \Sigma_i$ (signing)
- $\text{iVf}(\mathcal{KG}, \text{pk}_i, m, \Sigma_i) \rightarrow 0/1$ (individual signature verification)
- $\text{Cb}(\mathcal{KG}, m, (\text{pk}_i, \Sigma_i)_{i \in [N]}) \rightarrow \Sigma$ (combination)
- $\text{Vf}(\mathcal{KG}, (\text{pk}_i)_{i \in [N]}, m, \Sigma) \rightarrow 0/1$ (verification)

Moreover they should satisfy the following three properties.

Individual completeness states that for any message m , then a correctly generated signature on m by any correctly generated key with respect to any group \mathcal{KG} to which it belongs, must pass individual verification. Formally:

$$(1) \quad \mathbb{P} \left(0 \leftarrow \text{iVf}(\mathcal{KG}, \text{pk}, m, \Sigma) \mid \begin{array}{l} \text{sk}, \text{pk} \leftarrow \text{Kg}() \\ \wedge \text{pk} \in \mathcal{KG} \\ \wedge \Sigma \leftarrow \text{Sign}(\mathcal{KG}, \text{sk}, m) \end{array} \right) = \text{negl}$$

Combination-robustness states that combination must return a valid multisignature whenever applied on any number n of individual signatures on any message m , with respect to a common group \mathcal{KG} containing all signers, as soon as they all pass individual verification. Formally, for any polynomial adversary \mathcal{A} :

$$(2) \quad \mathbb{P} \left(0 \leftarrow \text{Vf}(\mathcal{KG}, (\text{pk}_i)_{i \in [n]}, m, \Sigma) \mid \begin{array}{l} \mathcal{KG}, (\text{pk}_i, \Sigma_i)_{i \in [n]} \leftarrow \mathcal{A} \\ \wedge (\text{pk}_i)_{i \in [N]} \subset \mathcal{KG} \\ \wedge \text{iVf}(\mathcal{KG}, \text{pk}_i, m, \Sigma_i) = 1 \forall i \in [N] \\ \wedge \Sigma \leftarrow \text{Cb}(\mathcal{KG}, m, (\text{pk}_i, \Sigma_i)_{i \in [N]}) \end{array} \right) = \text{negl}$$

Unforgeability states that any polynomial forger \mathcal{F} has negligible advantage in the following game denoted **m-uf**. A challenger exhibits to \mathcal{F} an honestly generated key pk , then \mathcal{F} commits to a group of public keys \mathcal{KG} , then it is granted access to a pk-signing oracle with respect to \mathcal{KG} . The forger wins if it can create a valid multi-signature on behalf of some subgroup of signers containing pk , on a message never queried to the oracle.

Note that the definition in the non-dynamic case is weak, because the forger can make signing queries only *after* it has committed to its challenge group \mathcal{KG} of keys. So this non-dynamic restriction rules-out a forger which would concurrently interact with multiple groups. But this restriction appears

<p>m-uf</p> <hr/> $(\text{sk}, \text{pk}) \leftarrow \text{Kg}()$ $\mathcal{KG} \leftarrow \mathcal{F}$ $((\text{pk}_i)_{i \in [N]}, m^*, \Sigma^*) \leftarrow \mathcal{F}^{\text{SIGN}}(\text{pk})$ return $(m \notin Q_{\text{sig}} \wedge \text{pk} \in (\text{pk}_i)_{i \in [N]} \subset \mathcal{KG}$ $\wedge \text{Vf}(\mathcal{KG}, (\text{pk}_i)_{i \in [N]}, m, \Sigma^*) = 1)$
<p>oracle $\text{SIGN}(m)$</p> <hr/> if $\text{pk} \in \mathcal{KG}$ $Q_{\text{sig}} \leftarrow Q_{\text{sig}} \cup \{m\}$ return $\dot{\Sigma} \leftarrow \text{Sign}(\mathcal{KG}, \text{sk}, m)$

Figure 2: Multi-signatures unforgeability game. Instructions removed in dynamic schemes are shaded-out

only in [9] to our knowledge, it is absent from most security definitions, even those of non-dynamic fNIMs (\mathcal{ASMP} [22, §4.1],[13]). As observed in the Introduction, the definition of [42, §5.1] does not either make this restriction, although their security proof implicitly makes it.

Fully Non-Interactive (One-time-tagged) Aggregate-signatures

The following definition applies to messages to be signed which come as $m_i = (\tau_i | v_i)$ where the τ_i are called their tags and v_i their variable parts. Again, the number n of signatures aggregated is a variable input. Compared to classical aggregate signatures, aggregation is enabled only on messages with the same tag τ , and unforgeability is guaranteed only if honest signers do not sign two different messages with the same tag.

Definition 2 (fNIA). A fully non-interactive one-time-tagged aggregate signature scheme consists of the following local algorithms.

- $\text{Kg}() \rightarrow (\text{sk}, \text{pk})$ (key generation)
- $\text{Sign}(\mathcal{KG}, \text{sk}_i, \tau_i, v_i) \rightarrow \Sigma_i$ (signing)
- $\text{iVf}(\mathcal{KG}, \text{pk}_i, \tau_i, v_i, \Sigma_i) \rightarrow 0/1$ (individual signature verification)
- $\text{Ag}(\mathcal{KG}, \tau, (\text{pk}_i, \Sigma_i, v_i)_{i \in [n]}) \rightarrow \Sigma$ (aggregation)
- $\text{Vf}(\mathcal{KG}, \tau, (\text{pk}_i, v_i)_{i \in [n]}, \Sigma) \rightarrow 0/1$ (verification)

Moreover they should satisfy the following three properties.

Individual completeness requires that for any tagged message $(\tau | v)$, then a correctly generated signature on $(\tau | v)$ by any correctly generated key with respect to any group \mathcal{KG} to which it belongs, must pass individual verification. Formally:

$$(3) \quad \mathbb{P} \left(0 \leftarrow \text{iVf}(\mathcal{KG}, \text{pk}, \tau, v, \dot{\Sigma}) \mid \begin{array}{l} \text{sk}, \text{pk} \leftarrow \text{Kg}() \\ \wedge \text{pk} \in \mathcal{KG} \\ \dot{\Sigma} \leftarrow \text{Sign}(\mathcal{KG}, \text{sk}, \tau, v) \end{array} \right) = \text{negl}$$

Aggregation-robustness requires that aggregation must return a valid aggregate signature whenever applied on any number n of individual signatures on any identically tagged

messages $m_i = (\tau, v_i)$, with respect to a common group \mathcal{KG} containing all signers, as soon as they all pass individual verification. Formally, for any polynomial adversary \mathcal{A} :

$$(4) \quad \mathbb{P} \left(0 \leftarrow \text{Vf} \left(\mathcal{KG}, \tau, (\text{pk}_i, v_i)_{i \in [n]}, m, \Sigma \right) \mid \begin{array}{l} \mathcal{KG}, \tau, (\text{pk}_i, v_i, \Sigma_i)_{i \in [n]} \leftarrow \mathcal{A} \\ \wedge (\text{pk}_i)_{i \in [n]} \subset \mathcal{KG} \\ \wedge \text{iVf}(\mathcal{KG}, \text{pk}_i, \tau, v_i, \Sigma_i) = 1 \forall i \in [n] \\ \wedge \Sigma \leftarrow \text{Ag}(\mathcal{KG}, \tau, (\text{pk}_i, v_i, \Sigma_i)_{i \in [n]}) \end{array} \right) = \text{negl}$$

Unforgeability requires that any polynomial forger \mathcal{F} has negligible advantage in the following game denoted **a-uf**. A challenger exhibits to \mathcal{F} an honestly generated key pk , then \mathcal{F} commits on a group of public keys \mathcal{KG} , then it is granted access to a pk -signing oracle with respect to \mathcal{KG} , which however refuses to sign twice with the same tag. The forger wins if it can create a valid aggregate signature on behalf of some subgroup of signers on a list of key - messages pairs containing some (pk, v^*) , for some tag τ^* , such that the oracle never delivered a signature on the tagged message $(\tau^* | v^*)$. Formally:

<p>a-uf</p> <hr/> $(\text{sk}, \text{pk}) \leftarrow \text{Kg}()$ $\mathcal{KG} \leftarrow \mathcal{F}$ $(\mathcal{KG}, \tau^*, (\text{pk}_i, v_i)_{i \in [n]}, \Sigma^*) \leftarrow \mathcal{F}^{\text{SIGN}}(\text{pk})$ return $(\exists v^* \text{ s.t. } (\text{pk}, v^*) \in (\text{pk}_i, v_i)_{i \in [n]} \wedge (\tau^*, v^*) \notin Q_{\text{sig}}$ $\wedge (\text{pk}_i)_{i \in [n]} \subset \mathcal{KG} \wedge \text{Vf}(\mathcal{KG}, \tau^*, (\text{pk}_i, v_i)_{i \in [n]}, \Sigma^*) = 1)$
<p>oracle $\text{SIGN}(\tau, v)$</p> <hr/> 1: if $\text{pk} \in \mathcal{KG} \wedge \nexists (\tau, v' \neq v) \in Q_{\text{sig}}$ 2: $Q_{\text{sig}} \leftarrow Q_{\text{sig}} \cup \{(\tau, v)\}$ 3: return $\dot{\Sigma} \leftarrow \text{Sign}(\mathcal{KG}, \text{sk}, \tau, v)$

Figure 3: Aggregate signatures unforgeability game. Instructions removed in dynamic schemes are shaded-out

Comments on Definitions and Related Works

Blob of signatures and messages. Virtually all existing fNIM schemes allow the combiner to take as separate inputs a multiset of keys, and a multiset of messages: $\text{Cb}(\mathcal{KG}, m, (\text{pk}_i)_{i \in [N]}, (\Sigma_i)_{i \in [N]})$. The reason why our syntax requires to regroup them in pairs: $(\text{pk}_i, \Sigma_i)_{i \in [n]}$ is because it is necessary in the fNIA called $\mathcal{AS-4}$ [16], which we observe is a fortiori a fNIM.

Key-aggregation. In all existing fNIM schemes, the verification can be factored in a first step called key-aggregation, which shrinks the multiset of the keys of the subgroup of signers into a short aggregate key: $\text{KAg} : (\mathcal{KG}, (\text{pk}_i)_{i \in [n]}) \rightarrow \text{apk}$. Then, verification proceeds using only apk : $\text{Vf-lazy}(\text{apk}, m, \Sigma) \rightarrow 0/1$. We purposely did not include this decomposition in the specifications, in line with [56, 51], because we believe it to be bug-prone. Indeed, some existing specifications only specify the last part Vf-lazy , but it would be insecure to run such a Verifier algorithm without any

additional safe means of checking, in a way or another, that apk was correctly output by KAg from the keys of the purported signers. A more subtle example is the scheme RSMS-pop in [56, §4.3]: its Verifier algorithm, VerifyMul, does not check validity of the PoPs on the public keys, although it does aggregate them with KAg. The problem is that its KAg does not either check validity of the PoPs. In conclusion, it seems to us that rogue-key attacks [21] invalidate the unforgeability of RSMS-pop, in the sense of [56, def. 3.3]. Of course solving this apparent issue would just require to specify that the RSMS-pop.VerifyMul verifies PoPs, in a way or another.

Interactive setups. The issue of enabling verification of apk is even clearer in the cases, not captured by our definitions, where KAg would be an interactive protocol between the signers. Fortunately the only such fNIM schemes to our knowledge are MOR01 [79] and *ASMP* & *ASMP-pop* [22, §4 & §6], and in all of them the protocol is *publicly verifiable* as long as it is executed over a broadcast channel.

Concurrent groups. Although our unforgeability definitions in Figures 2 and 3 follow the ones of [9], in which the adversary must commit on a single group \mathcal{KG} and has then access to a signing oracle w.r.t. this group only, this limitation is relaxed in the specifications of [22, §4]. More precisely, provided interactive group setups, they show that the security of their schemes *ASMP/ASMP-pop* decreases linearly only in the total number of signers over all groups in which the target key pk is involved. Of course all these issues disappear in dynamic schemes.

Not fully non-interactive schemes. In the scheme *MSP* [22], the Sign algorithm takes as input the exact list of the keys which are meant to sign the message: $(\text{pk}_i)_{i \in [N]}$. Upon receiving N signatures on some message m generated by the keys $(\text{pk}_i)_{i \in [N]}$, the combination Cb can combine them only if they were all generated with input $(\text{pk}_i)_{i \in [N]}$. Said in the other way: Cb cannot combine any signature on m generated with input $(\text{pk}_i)_{i \in [N]}$, unless receiving such signatures from *all* the intended signers $(\text{pk}_i)_{i \in [N]}$. Said otherwise: aggregation fails as soon as one of the intended signers aborts (this issue is lifted in the blog version *MSP-blog* [21]). This limitation also shows-up in the two-round schemes ([46, 42, 13]): although they require interaction only in a first round which is message-independent, the signatures produced in the second round can be combined only if they are produced by all participants of the first round.

Aggregation robustness. This is a property which we credit to [56, 51]. Prior specifications guaranteed a successful aggregation only over correctly generated signatures and keys.

IV. *MtoA*: MULTI TO AGGREGATE COMPILER

We convey all ideas of *MtoA*, further formalism and the (obvious) proof can be found in Sec. A. We consider any fNIM: M , and describe the resulting one-time-tagged fNIA, called A . Each potential signer i generates then outputs a list of $2|v|$ M-public keys: $\{\text{pk}_i^{j,b}, j \in |v|, b \in \{0, 1\}\}$. To sign some $(\tau|v_i)$, the signer i : parses $v_i = (v_i^j)_{j \in [|v|]}$, the bit decomposition of the variable part, then outputs a signature *on* τ under each public key pk_i^{j,v_i^j} for $j \in [|v|]$. That is, the data of the variable part v_i of the message intended to be

signed is *not* encoded by the actual content signed, which is just equal to the fixed part τ , but instead by the list of the *keys* which signed τ . The Aggregator: upon receiving signatures over n messages $(v_i)_n$ from n signers, where we recall each signature consists of a $|v|$ -sized list of M-individual signatures, applies the M-combination algorithm: M.Cb on all $N = |v|n$ signatures received. Finally, the verifier checks the multisignature Σ received against the multiset of keys \mathcal{PK} which it reads from the messages $(v_i)_{i \in [n]}$. That is, for each message $v_i = (v_i^j)_{j \in [|v|]}$ it appends $\mathcal{PK} \leftarrow \mathcal{PK} \cup \{\text{pk}_i^{j,v_i^j}\}_{j \in [|v|]}$, then outputs M.Vf(\mathcal{PK}, m, Σ).

V. dms

dms is specified over any bilinear group $(\mathbb{G}_1, G_1, \mathbb{G}_2, G_2, \mathbb{G}_T, e)$ and operates with any hash-to-curve random oracle $H : \{0, 1\}^* \rightarrow \mathbb{G}_1$ and random oracle $H_{\text{pop}} : \{0, 1\}^* \rightarrow \mathbb{Z}_p$. We will show that its security tightly reduces to the one of BLS signatures (Figure 1). Themselves tightly reduce to the hardness of the DL, in the AGM, by [53, §6]. The syntax of dms does not contain any group \mathcal{KG} , hence, it is a *dynamic* fNIM. Note that in the description below of dms, if one removes the PoP: π from the Kg, and its verification: kVf from both the Cb and Vf algorithms, then we are brought back to the schoolbook BLS fNIM [20] (also recalled in Sec. D-2).

Kg(): sample $x \xleftarrow{\$} \mathbb{Z}_p$, set $X \leftarrow x.G_2$; sample $r \xleftarrow{\$} \mathbb{Z}_p$, $R \leftarrow r.G_2$, set $c \leftarrow H_{\text{pop}}(X, X, R)$ and $z := r + c.x$, let $\pi \leftarrow (R, z)$ (the PoP), output $\text{sk} := x$ and $\text{pk} := (X, \pi)$.

Since PoPs are to be verified by both the individual verification algorithm (iVf) and the combination algorithm (Cb), we factor-out their verification with the following helper function:

kVf(X): parse $(X, \pi) \leftarrow \text{pk}$ and $(R, z) \leftarrow \pi$; $c \leftarrow H_{\text{pop}}(X, X, R)$, output $(X \in \mathbb{G}_2 \wedge z.G_2 == R + c.X)$.

[leftmargin=0pt,topsep=0pt, label=] Noticeably, and unlike pairing-based PoPs, no subgroup membership in \mathbb{G}_2 is to be performed on the PoP π , since the verified relation $R = z.G_2 - c.X$ automatically implies membership of R in \mathbb{G}_2 .

Sign(sk, m) = $\text{sk}.H(m)$;

iVf(pk_i, m, Σ_i): parse $\text{pk}_i \leftarrow (X_i, \pi_i)$, return $\Sigma_i \in \mathbb{G}_1 \wedge \text{kVf}(\text{pk}_i) == 1 \wedge e(\Sigma_i, G_2) == e(H(m), X_i)$

Cb(m, (pk_i, Σ_i)_{i ∈ [N]}): if $\{\text{iVf}(\text{pk}_i, m, \Sigma_i) == 1 \forall i \in [N]\}$, then return $\sum_{i \in [N]} \Sigma_i$.

Vf((pk_i)_{i ∈ [N]}, m, Σ): parse $(X_i, \pi_i) \leftarrow \text{pk}_i \forall i \in [N]$; return $\Sigma \in \mathbb{G}_1 \wedge \{\text{kVf}(\text{pk}_i) == 1 \forall i \in [N]\} \wedge e(\Sigma, G_2) == e(H(m), \sum_{i \in [N]} X_i)$.

Theorem 3. *dms is a fNIM in the AGM. For all three types of bilinear groups, its unforgeability tightly reduces to hardness of the discrete logarithm (DL) problem.*

The proofs of Individual completeness and aggregation-robustness are identical to the ones of standalone BLS multisignatures [20] (plus correctness of Schnorr signatures used as PoPs), so we skip them. In Sec. V-A we state the main ingredient of unforgeability of dms. It states that whenever an algebraic adversary \mathcal{A} outputs a public key X^* along with

a Schnorr signature on itself: (R^*, z^*) , then the discrete logarithm of X^* , i.e., a secret key, can be efficiently computed from the decomposition of X^* given by the adversary. Then in Sec. V-B we conclude the proof of unforgeability from this ingredient.

A. Extracting Schnorr despite correlations

We formalize the game called SSC in Figure 4, which stands for ‘‘Schnorr Straight-line extraction in presence of BLS Correlations’’. The game samples a public key $X = x.G_2$, which we dub the ‘‘honest key’’, generates a Schnorr signature on it: $\pi = (R, z)$, and shows (X, π) to the adversary \mathcal{A} . Then the adversary is given access to hash-into- \mathbb{G}_1 and hash-into- \mathbb{Z}_p random oracles: $H : m_i \rightarrow M_i$ and H_{pop} ; and to a BLS signing oracle for the honest secret key $\text{SIGN} : m_i \rightarrow \Sigma_i := x.H(m_i)$. All in all, up to delivering to \mathcal{A} both replies from SIGN and H for every queried m , we have that \mathcal{A} is delivered a random string with a hidden structure: $(M_i, x.M_i)_{i \in [q_H]}$, in addition to the Schnorr proof of knowledge π for the same exponent: x of X . In the game SSC in Figure 4, the challenger tries to extract in straight-line a discrete logarithm x^* upon being submitted some X^* and some Schnorr signature (R^*, z^*) on it, valid for X^* . The goal of the adversary is to defeat this extraction: it wins as soon as one extraction attempt fails over all its submissions. The extractor is defined as follows. Upon submitting some (X^*, R^*, z^*) , the adversary gives the decompositions of X^* , R^* in terms of all group elements received so far, of which all hashes M_i and signatures Σ_i returned by the oracles so far:

$$(5) \quad X^* = \alpha.G + \beta.X + \sum_{i \in [q_H]} \gamma_i \Sigma_i + \sum_{i \in [q_H]} \delta_i M_i$$

$$(6) \quad R^* = \alpha'.G + \beta'.X + \sum_{i \in [q_H]} \gamma'_i \Sigma_i + \sum_{i \in [q_H]} \delta'_i M_i$$

Note that, without loss of generality, we assumed that R does not appear in the decompositions, since $R = z.G - c.X$. As will be precised in the proof, these decompositions are to be understood as those which \mathcal{A} gave when outputting X^* and R^* for the first time. The extractor is then simply defined as the function which returns α . So this extractor outputs a correct discrete logarithm if $X^* = \alpha.G_2$, else, this means that the adversary wins. The following theorem states that escaping this extractor is as hard as solving DL. The proof is done in Sec. VI and is one of our main technical contributions.

Theorem 4. *From any SSC-adversary \mathcal{A} with advantage ϵ and making q_H RO queries, one can build an adversary \mathcal{E} against the discrete logarithm (DL) in \mathbb{G}_2 with advantage $\epsilon' \geq (1 - \frac{q_H+1}{p})(\epsilon - q_H/p)$ and with at most twice the running time.*

B. Tight reduction of dms to DL.

The proof follows from the following chain of tight reductions:

$$(7) \quad \mathbf{m-uf}(\text{dms}) \xrightarrow{\text{Lem. 5}} \text{SSC} \wedge \mathbf{BLS-uf} \xrightarrow{\text{Thm. 4}} \mathbf{DL} \wedge \mathbf{BLS-uf} \xrightarrow{[53]} \mathbf{DL}$$

We now outline each of the reductions. The game on the left is the unforgeability of dms. The first reduction is stated more precisely as follows, and will be proven below:

<p>SSC</p> <hr/> <p>$x \xleftarrow{\\$} \mathbb{Z}_p, X = \delta.G_2$ $r \xleftarrow{\\$} \mathbb{Z}_p, R \leftarrow r.G_2, c \leftarrow H_{\text{pop}}(X, X, R)$ $z \leftarrow r + cx, \pi \leftarrow (R, z)$ foreach $(X^*, R^*, z^*) \leftarrow \mathcal{A}^{\text{SIGN}, H, H_{\text{pop}}}(X, \pi) \triangleright$ up to q_H attempts $c^* \leftarrow H(X^*, X^*, R^*)$ if $z^*.G_2 = R^* + c^*.X^* \wedge (X^*, R^*, z^*) \neq (X, R, z)$ receive the decomposition (5): $X^* = \alpha.G_2 + \dots$ (see above) if $X^* \neq \alpha.G_2$ then win $\leftarrow 1$ return win</p> <hr/> <p>oracle $\text{SIGN}(m)$ return $\Sigma \leftarrow x.H(m)$</p>

Figure 4: Schnorr Straight-line extraction in presence of Correlations oracles. H and H_{pop} are random oracles: into \mathbb{G}_1 and into \mathbb{Z}_p ; SIGN is a BLS signing oracle.

Lemma 5. *For any algebraic forger \mathcal{F} in the unforgeability game $\mathbf{m-uf}$ of dms, with advantage ϵ , running time t and making at most q_H RO queries; then there exists a forger \mathcal{B} in the unforgeability game $\mathbf{BLS-uf}$ of BLS signatures, which has advantage $\epsilon' \geq \epsilon - q_H/p - \text{UB}^{\text{SSC}}(t)$ and at most twice the running time t , where $\text{UB}^{\text{SSC}}(t)$ denotes the upper-bound on the advantage of a time- t adversary in game SSC.*

The third reduction, labelled by Thm. 4 is because Theorem 4 shows a tight reduction from SSC to DL.

The last reduction, from BLS unforgeability to the hardness of DL, is proven in [53, §6]. Note the their proof holds for type I bilinear groups, and could be easily adapted to type II. On the other hand, for their proof to carry over type III groups, the DL problem to be considered is when the adversary is given challenges in both groups: $(\ell.G_1, \ell.G_2)$. This type III DL problem is formalized in [11] as (1, 1)-DL. Note that in the plain model (non-AGM), then BLS-uf in type III groups reduces (non-tightly) to an analogous variant of computational Diffie-Hellman, which is called co-cdh* in [37], and simply ‘‘co-DH’’ in [22, Def. 2].

Now, all what remains to be proven is Lemma 5.

Proof of Lemma 5. The following proof holds for all three types of bilinear groups. \mathcal{B} receives the challenge honest key X and simulates the required Schnorr PoP π on X using the standard strategy. Namely: it samples $(r, z) \leftarrow \mathbb{Z}_p^2$, programs the random oracle as $c \leftarrow H_{\text{pop}}(X, X, R)$, up to the q_H/p -probability event where \mathcal{A} would already have queried $H_{\text{pop}}(X, X, R)$, then outputs $\pi \leftarrow (R := r.G_2, z)$. Then it gives $\text{pk} \leftarrow (X, \pi)$ as the challenge honest key to \mathcal{F} . The rest of the proof closely follows [20]. The only difference, which is all the point of our work, is that instead of being given for free all valid secret keys from \mathcal{F} , \mathcal{B} extracts them (except those identical to X) thanks to an obvious reduction to SSC. Whenever \mathcal{F} makes an $\mathbf{m-uf-SIGN}$ signature query on some message m , \mathcal{B} simply makes the query on the same m to its own $\mathbf{BLS-uf-SIGN}$ oracle, then forwards the result to \mathcal{F} . At some point, \mathcal{F} outputs a dms forgery: $((\text{pk}_i)_{i \in [N]}, m, \Sigma)$, in particular such that m was not queried before to SIGN .

For simplicity, let us re-index the keys such that all those *different* from X come first: $(pk_i)_{i \in [N_1]}$, then the $N - N_1$ ones identical to X come last. Note that by definition of a forgery, $N - N_1 \geq 1$.

Claim: Parse $(X_i, \pi_i) \leftarrow pk_i \forall i \in [N]$. Then, \mathcal{F} is able to extract the discrete logarithms $(x_i)_{i \in [N_1]}$ of the $(X_i)_{i \in [N_1]}$ from the queries of \mathcal{F} to \mathbb{H} and \mathbb{H}_{pop} , except with probability $\text{UB}^{\text{SSC}}(t)$, in at most twice the running time t of \mathcal{F} .

proof of the Claim. Consider formally the adversary \mathcal{F}^\perp , which is equal to \mathcal{F} except that the last two outputs of its forgery are removed: $((pk_i)_{i \in [N]}, \perp, \perp)$. Then \mathcal{F}^\perp is a game-SSC adversary, which concludes the Claim.

End of the proof. Informally, the reduction \mathcal{B} removes, from the forgery Σ , the contributions of the individual signatures from the non- X keys. Then, what remains is valid signature on m for the key $(N - N_1) \cdot X$, so it scales it down to a valid signature for X . Formally, \mathcal{B} computes the individual signatures $\Sigma_i \leftarrow x_i \cdot \mathbb{H}(m)$, then outputs the forgery:

$$(8) \quad \Sigma_X \leftarrow \frac{1}{N - N_1} \left(\Sigma - \sum_{i \in [N_1]} \Sigma_i \right).$$

Let us formally verify that Σ_X is indeed a valid BLS signature on m for the challenge key X , which will conclude the proof. For readability we multiply everywhere by $(N - N_1)$. By construction:

$$(9) \quad (N - N_1)e(\Sigma_X, G_2) = e\left(\Sigma - \sum_{i \in [N_1]} x_i \cdot \mathbb{H}(m), G_2\right).$$

On the other hand, Σ being a valid dms multi-signature, we have

$$(10) \quad e(\Sigma, G_2) = e\left(\mathbb{H}(m), \sum_{i \in [N_1]} x_i \cdot G_2 + (N - N_1)X \cdot G_2\right).$$

Developping by linearity the RHS of Eq. (9), then replacing $e(\Sigma, G_2)$ by Eq. (10), we obtain that this RHS is equal to

$$(11) \quad e\left(\mathbb{H}(m), \sum_{i \in [N_1]} x_i \cdot G_2 + (N - N_1)X \cdot G_2\right) - e\left(\sum_{i \in [N_1]} x_i \cdot \mathbb{H}(m), G_2\right)$$

Finally, cancelling-out all equalities $e(\mathbb{H}(m), x_i \cdot G_2) = e(x_i \cdot \mathbb{H}(m), G_2)$, we obtain:

$$(12) \quad (N - N_1)e(\Sigma_X, G_2) = e\left(\mathbb{H}(m), (N - N_1)X\right),$$

which, after dividing by $(N - N_1)$, proves the validity of Σ_X .

VI. PROOF OF THEOREM 4

As in most works on BLS [53, 9, 7], we describe only the proof in the case of type I bilinear groups, so in what follows we identify $\mathbb{G} := \mathbb{G}_1 = \mathbb{G}_2$ and $G := G_1 = G_2$, excepted in the formal descriptions of Figures 4 and 5. We will explain in the end why the proof is much simpler in the cases of type II and III groups.

Simplifications w.l.o.g. Note that in the definition of the game SSC, we could assume without loss of generality that \mathcal{A} submits at most one Schnorr signature to the game SSC, instead of $q_{\mathbb{H}}$ submissions. Indeed it can make submissions to itself and run the extractor on itself to check if it won or not. For this reason, without loss of generality (w.l.o.g.), we now consider that \mathcal{A} makes at most one submission to the game. Also, w.l.o.g., consider that when the adversary makes a query

m to \mathbb{H} then it immediately makes the same query m to SIGN , and conversely.

The reduction \mathcal{E} against **DL** receives a challenge L and its goal is to output the exponent ℓ , i.e., s.t. $L = \ell \cdot G$. To this end, it runs the SSC adversary \mathcal{A} and simulates SSC to it, as follows. We call \mathcal{E} the *master reduction* because it tosses a coin and, depending on its value 1 or 0, behaves towards \mathcal{A} as reduction C or D, both described in Figure 5. We now convey their intuition. As will be clear from their description, both are (almost) perfect simulations of game SSC, and furthermore C and D are information-theoretically indistinguishable from each other. The difference between them, hidden to \mathcal{A} , is how the challenge L is embedded. Both reductions use procedures, denoted $\tilde{\mathbb{H}}$, $\tilde{\text{SIGN}}$ and $\tilde{\mathbb{H}}_{\text{pop}}$, to simulate to \mathcal{A} the responses to its queries to oracles \mathbb{H} , Sign and \mathbb{H}_{pop} respectively.

Reduction C embeds the **DL** challenge L as the honest public key of **SSC**, i.e., sets $X := L$. It simulates the required Schnorr PoP: π on X following the standard technique, namely: samples $(r, z) \leftarrow \mathbb{Z}_p^2$, programs the random oracle as $c \leftarrow \tilde{\mathbb{H}}_{\text{pop}}(X, X, R)$, then outputs $\pi \leftarrow (R := r \cdot G_2, z)$. So its simulation of SSC is perfect, up to the $q_{\mathbb{H}}/p$ -probability event where \mathcal{F} would have queried $\mathbb{H}_{\text{pop}}(X, X, R)$ before it was programmed on c .

Reduction D honestly generates the public key X as $x \leftarrow \mathbb{Z}_p$, $X \leftarrow x \cdot G_2$. It embeds the DL challenge in the simulated hash-to-curve oracle $\tilde{\mathbb{H}}$, using the following trick of [53, §6]. Upon queried a new message m_i , $\tilde{\mathbb{H}}$ samples $(\hat{h}_i, b_i) \leftarrow \mathbb{Z}_p^2$ then returns $M_i \leftarrow b_i \cdot L + \hat{h}_i \cdot G$. In particular, the output M_i varies uniformly in \mathbb{G} , so $\tilde{\mathbb{H}}$ perfectly simulates a random oracle. In conclusion, it can perfectly simulate the signing oracle, as: $m \leftarrow x \cdot \tilde{\mathbb{H}}(m)$. Note that since D knows the secret key x , it could also honestly generate a Schnorr PoP as $r \leftarrow \mathbb{Z}_p$, $R \leftarrow r \cdot G_2$, $c \leftarrow \mathbb{H}_{\text{pop}}(X, X, R)$ then $z \leftarrow r + c \cdot x$. So it would not have to program \mathbb{H}_{pop} , so its simulation of SSC would be perfect. But we make the choice to specify that D does instead generate a simulated Schnorr proof and programs $\tilde{\mathbb{H}}$, like C does. The reason for this choice is that this makes the view of the adversary identically distributed against C and D, so this will simplify the proof.

Strategy of C to win against DL. In what follows we consider the event $C^{\mathcal{A}} = 1$ where \mathcal{A} outputs a winning triple (X^*, R^*, z^*) to C. Namely, it passes verification of Schnorr proofs, i.e., s.t. for $c^* \leftarrow \tilde{\mathbb{H}}_{\text{pop}}(X^*, X^*, R^*)$ then $z^* \cdot G = R^* + c^* \cdot X^*$. Being algebraic, \mathcal{A} also submits a linear decomposition of X^* and R^* on all group elements which it was delivered so far. We are now more precise, and specify that the decompositions given in equations (5) (6) are those that \mathcal{A} submitted when it outputted X^* and R^* for the first time, i.e., either to oracle $\tilde{\mathbb{H}}_{\text{pop}}$ or directly to the main C procedure. Since \mathcal{A} won, then α cannot be the only nonzero coefficient in X^* . With these notations, recall that the goal of C is to find the exponent $x := \ell$ of $X = L$, i.e., s.t. $x \cdot G = X$. To explain how C tries to find it efficiently, start from the relation $z^* \cdot G = R^* + c^* \cdot X^*$, substitute R^* and X^* by their decompositions in Eq. (5), then we obtain:

$$(13) \quad z^* \cdot G = c^* \left(\alpha \cdot G + \beta \cdot X + \sum_i \gamma_i \cdot \Sigma_i + \sum_i \delta_i \cdot M_i \right)$$

$$+ \alpha'.G + \beta'.X + \sum_i \gamma'_i.\Sigma_i + \sum_i \delta'_i.M_i .$$

Replacing the oracle responses by their values: $\Sigma_i = h_i.X$ and $M_i = h_i.G$, and substituting $X = x.G$, we obtain:

$$(14) \quad z^*G = x \left(c^*\beta + c^* \sum_i \gamma_i h_i + \beta' + \sum_i \gamma'_i h_i \right) + c^*\alpha + c^* \sum_i \delta_i h_i + \alpha' + \sum_i \delta'_i h_i .$$

Thus C can efficiently recover $\ell = x$ by division by the scalar:

$$(15) \quad f^* := c^* \left(\beta + \sum_i \gamma_i h_i \right) + \beta' + \sum_i \gamma'_i h_i$$

... unless this scalar is zero.

To analyze this bad ($f^* = 0$) event, the important observation is that in all games considered (both SSC and its reductions C and D), the decompositions (5)(6) of X^* and R^* , were handed-out by the adversary \mathcal{A} strictly before c^* was sampled uniformly at random. Let us prove it on the example of C. Either \mathcal{A} queried (X^*, X^*, R^*) to \tilde{H}_{pop} before outputting (X^*, R^*, z^*) to the main C procedure, then c^* was sampled by \tilde{H}_{pop} just after. Or, \mathcal{A} gave (X^*, R^*, z^*) to the main C procedure without having queried $\tilde{H}_{\text{pop}}(X^*, X^*, R^*)$ before, then C makes the query to its internal procedure $\tilde{H}_{\text{pop}}(X^*, X^*, R^*)$ just after, which then samples c^* .

Lemma 6. *Consider, as before, the event (up to probability q_H/p) where [no query $\tilde{H}_{\text{pop}}(X, X, R)$ was made before \tilde{H}_{pop} was programmed as $\tilde{H}_{\text{pop}}(X, X, R) \rightarrow c$]. Consider any query $\tilde{H}_{\text{pop}}(\dot{X}, \dot{X}, \dot{R})$ made for the first time. Denote the decompositions of \dot{X} and \dot{R} as in (5) (6) (so we omit adding a dot above the coefficients $\alpha, \beta, \gamma, \dots$), and the response \dot{c} . Consider the indices $i = 1, \dots, i_0$ of all queries m_i to \tilde{H} and $\tilde{\text{SIGN}}$ (responding \tilde{M}_i and $\tilde{\Sigma}_i$, respectively) which were made before the query $\tilde{H}_{\text{pop}}(\dot{X}, \dot{X}, \dot{R})$, i.e., before \dot{c} was sampled. Then:*

\dot{c} is sampled independently from $(\alpha, \beta, (\delta_i, \delta'_i, \gamma_i, \gamma'_i, h_i)_{i \leq i_0})$

and, for all $i > i_0$: $\delta_i = \delta'_i = \gamma_i = \gamma'_i = 0$.

The rest of the proof strategy is as follows. Let us consider the event where the adversary wins against the master reduction: $(\mathcal{E}^{\mathcal{A}} = 1) := (C^{\mathcal{A}} = 1 \wedge \mathcal{E} = C) \vee (D^{\mathcal{A}} = 1 \wedge \mathcal{E} = D)$, where $\mathcal{E} = C$ and $\mathcal{E} = D$ denote the events where the coin was 1 or 0, i.e., where \mathcal{E} behaves as C or D. We are going to consider the sub-event, denoted $\neg V \subset (\mathcal{E}^{\mathcal{A}} = 1)$, defined as the non-vanishing of at least one of the coefficient in f^* (Eq. (15)), i.e., $\neg V := (\beta + \sum_i \gamma_i h_i \neq 0) \vee (\beta' + \sum_i \gamma'_i h_i \neq 0)$. An immediate consequence of Lemma 6 is that, under C, in the sub-event $\neg V \wedge (C^{\mathcal{A}} = 1)$, then the bad event ($f^* = 0$) (almost) never happens, and thus C is (almost) always able to find the DL challenge $x = \ell$. The ‘‘almost’’ will be quantified later. So what remains to conclude the proof is to show that, in the complementary event $V \subset (\mathcal{E}^{\mathcal{A}} = 1)$ which we call ‘‘very bad’’, then the reduction D will (almost always) be able to find ℓ efficiently. Our first task is to formalize a predicate equivalent to V and which is well-defined under D. To do so, we multiply the relations defining $\neg V$ by X then

take the negation, which yields:

$$(16) \quad V := \left\{ \beta.X + \sum_i \gamma_i M_i = 0 \wedge \beta'.X + \sum_i \gamma'_i M_i = 0 \right\} .$$

This equivalent predicate being purely in terms of the view of the adversary, it is also meaningful under D.

In order to conclude, we recall the fact that C and D are perfectly indistinguishable from the adversary \mathcal{A} . A consequence is that the coin tossed by the master-reduction \mathcal{E} , i.e., its choice of behavior C or D, is independent of which event V or $\neg V$ happens (otherwise the adversary could distinguish between C and D). In conclusion, each time the adversary wins, i.e., $(\mathcal{E}^{\mathcal{A}} = 1)$, whatever V or $\neg V$ is the most likely to happen, the master-reduction \mathcal{E} will have almost probability 1/2 to extract the DL challenge ℓ . We now formalize the above claims as Lemma 7, then formalize the above conclusion of Theorem 4 from it, then prove Lemma 7.

Lemma 7. *There exists reductions C and D from SSC to the DL game, such that for any SSC-adversary \mathcal{A} :*

- both the views of \mathcal{A} against C and D, are identically distributed as in SSC, except with q_H/p probability;
- the views of \mathcal{A} against C and D are identically distributed;
- C and D enjoy the following probabilities of success, i.e., of $(C^{\mathcal{A}} = 1)$ and $(D^{\mathcal{A}} = 1)$:

$$(17) \quad \mathbb{P}(\mathbf{DL}^C = 1) = \left(1 - \frac{q_H}{p}\right) \cdot \mathbb{P}(C^{\mathcal{A}} = 1 \wedge \neg V)$$

$$(18) \quad \mathbb{P}(\mathbf{DL}^D = 1) = \left(1 - \frac{q_H + 1}{p}\right) \cdot \mathbb{P}(D^{\mathcal{A}} = 1 \wedge V)$$

Assuming Lemma 6, let us conclude Theorem 4. Since the master reduction \mathcal{E} behaves as C or D with probability 1/2 each, we have $\mathbb{P}(\mathbf{DL}^{\mathcal{E}} = 1) = 1/2 \mathbb{P}(\mathbf{DL}^C = 1) + 1/2 \mathbb{P}(\mathbf{DL}^D = 1)$. By Equations (17) and (18), it is in turn $\geq \left(1 - \frac{q_H + 1}{p}\right) [\mathbb{P}(C^{\mathcal{A}} = 1 \wedge \neg V) + \mathbb{P}(D^{\mathcal{A}} = 1 \wedge V)]$. Now, note that for any fixed \mathcal{A} , we have $\mathbb{P}(C^{\mathcal{A}} = 1 \wedge \neg V) = \mathbb{P}(D^{\mathcal{A}} = 1 \wedge \neg V)$. Indeed if not, then an unlimited adversary \mathcal{A} could distinguish between C and D, a contraction. Substituting, we obtain: $\mathbb{P}((\mathbf{DL})^{\mathcal{E}} = 1) \geq 1/2 \left(1 - \frac{q_H + 1}{p}\right) [\mathbb{P}(C^{\mathcal{A}} = 1)]$. By the first claim of Lemma 7, since the view of \mathcal{A} against C is q_H/p -close to its view in SSC, we have that $\mathbb{P}(C^{\mathcal{A}} = 1) \geq \epsilon - q_H/p$. Replacing in the previous formula of $\mathbb{P}((\mathbf{DL})^{\mathcal{E}} = 1)$ yields the theorem.

a) *Proof of Lemma 7:* The first claim was already argued along with the definitions of C and D, namely, the only difference between the view against SSC is in the event where the adversary had already queried $\tilde{H}_{\text{pop}}(X, X, R)$ before it was programmed.

Proof of Eq. (17). Since the bound is to be proven under reduction C only, we consider only the reduction C, i.e., we condition on the event $\mathcal{E} = C$. By Lemma 6, for each query $\tilde{H}_{\text{pop}}(\dot{X}, \dot{X}, \dot{R})$ in the execution, and denoting \dot{f} as defined in the formula (15) (here w.r.t. $(\dot{X}, \dot{X}, \dot{R})$), we have $\mathbb{P}(\dot{f} = 0 | \neg V) = 1/p$ where the probability is taken over the sampling of the answer c^* . Taking the union bound over all q_H queries in the execution, we thus have probability q_H/p that none of

their \hat{f} is equal to 0. In particular, for the specific f^* of the winning triple, we thus have

$$(19) \quad \mathbb{P}(f^* = 0 | \neg V) \leq q_H/p.$$

Since C is able to extract $x = \ell$ when $f^* \neq 0$, this concludes the proof.

Proof of Eq. (18). Since the bound is to be proven under reduction D only, we consider only the reduction D, i.e., we condition on the event $\mathcal{E} = D$. Let us start from Eq. (13) and, since we assumed V , simplify by Eq. (16). Replacing $M_i = b_i.L + \hat{h}_i.G$ and $L = \ell.G$ we obtain

$$(20) \quad z^*.G = c^* \left(\alpha.G + \sum_i \delta_i(b_i \ell.G + \hat{h}_i.G) \right) + \alpha'.G + \sum_i \delta'_i(b_i \ell.G + \hat{h}_i.G)$$

Thus D can efficiently recover ℓ by division by the scalar:

$$(21) \quad \lambda := c^* \sum_i \delta_i b_i + \sum_i \delta'_i b_i$$

... unless this scalar λ is zero.

Let us assume that it is the case, then we cannot be in the event $W := \{(\delta_i = 0 \wedge \delta'_i = 0) \forall i \wedge V\}$. Indeed, substituting in Eq. (5) those vanishings and those of Eq. (16), would yield $X^* = \alpha.G$, contradicting that \mathcal{A} wins. Hence, we must be in the event $\neg W \subset V$ where at least one of the coefficients δ_i or δ'_i is nonzero. To conclude, we apply the same kind of reasoning as in the proof of Eq. (17). Let us consider one query $\tilde{H}_{\text{pop}}(\dot{X}, \dot{X}, \dot{R})$ in the execution. By Lemma 6, $\mathbb{P}(\dot{c} = 0 | \neg W) = 1/p$ where the probability is taken over the sampling of c^* and on the coins of the adversary. Taking the union bound over all q_H queries in the execution, we thus have probability q_H/p that none of their \dot{c} is equal to 0. In particular, for the specific c^* of the winning triple, we thus have

$$(22) \quad \mathbb{P}(c^* = 0 | \neg W) \leq q_H/p.$$

In the event $c^* \neq 0$, we Claim that:

$$(23) \quad \mathbb{P}(\lambda = 0 | c^* \neq 0 \wedge \neg W) \leq 1/p,$$

which concludes the proof. The Claim follows from the fact that, by construction, all $(b_i)_i$ are information-theoretically hidden from the adversary, hence they are independent of $(\delta_i, \delta'_i)_i$, of which at least one is nonzero by definition of $\neg W$.

b) Comments on SSC and on the proof: Compared to [54], which consider a forger against Schnorr signatures, the goal of our adversary in game SSC is easier, and thus our proof apparently harder. Indeed, their forger has to forge a Schnorr signature for a *given target key*. Whereas, our forger succeeds as long as it outputs *any* Schnorr proof, such that the discrete logarithm cannot be extracted.

We credit to [54] the crucial observation that c^* is sampled after the adversary first returns the decomposition of the Schnorr proof that it submits. Notice that it would be fallacious to conclude that, for a given winning triple output by the adversary: (X^*, R^*, z^*) , then the $c^* \leftarrow \tilde{H}(X^*, X^*, R^*)$ would be independent from the decompositions of X^* and R^* . Indeed, the adversary could well make a unique output to the game: (X^*, R^*, z^*) , chosen among possibly many winning triples, as the one which maximizes the number of digits of c^* in common with, e.g., the coefficient δ_1 in the decomposition of X^* . Such correlations are captured by the overhead q_H in

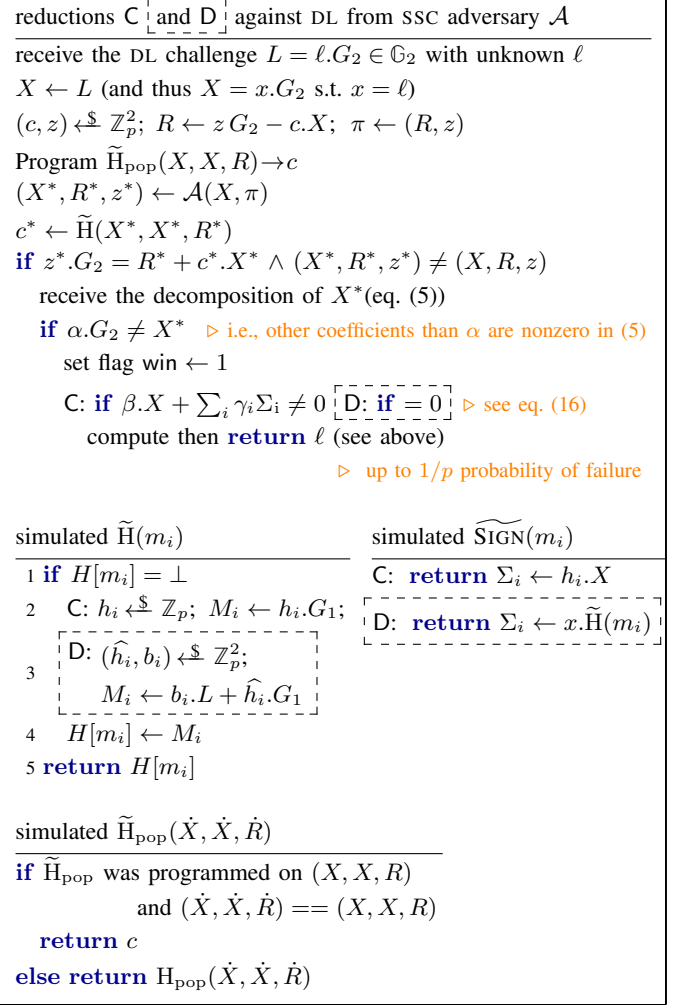


Figure 5

the probabilities of our bad events, i.e., they have probability q_H/p instead of just $1/p$. This overhead captures all possible queries (X^*, X^*, R^*) that could have been made to \tilde{H}_{pop} in order to find winning triples.

Although the proof immediately reduces to the case where \mathcal{A} submits at most one Schnorr signature to the game SSC, we defined SSC with multi-submissions to make it easier to use in the analysis of dms.

The proof can be much simplified in the case of Type II or III bilinear groups. The DL challenge L is in \mathbb{G}_2 , but in type II and III groups the algebraic adversary is further restricted to decompose \mathbb{G}_2 elements in \mathbb{G}_2 only. So all the complicated terms in \mathbb{G}_1 in the decompositions Equations (5) and (6) disappear (so the reduction needs not anymore program the hash-to-curve).

VII. EVALUATION AND COMPARISON

Our implementations¹ were run on a laptop with Core i5-8265U (8 cores at 1.6GHz), 16GB of RAM, with the library gnark-crypto on Go [28]. The curve used was BLS12-377,

¹The code is available at <https://anonymous.4open.science/r/MtoA-830A/>.

offering a pairing of type III, for which the uncompressed size of a point in \mathbb{G}_1 is 768 bits and of a point in \mathbb{G}_2 is 1536 bits. Compressed points, i.e., their x -coordinate plus one bit, are twice smaller. Each number is the mean over 10 executions.

1) *Comparing processing of the group setup runtimes:* We dub “Verifier” the verification function $\text{Vf}(\mathcal{KG}, (\text{pk}_i)_{i \in [N]}, m, \Sigma) \rightarrow 0/1$ of a fNIM. We call processing of the group setup the tasks of the Verifier which can be done straight upon learning the group of potential signers: \mathcal{KG} ; and online verification the remaining tasks performed upon learning the actual subset of signers $(\text{pk}_i)_{i \in [N]} \subset \mathcal{KG}$ and the signed message (m, Σ) . Of course, in dynamic fNIMs such as *MSP-pop* [85][22, §6][24] and *dms*, the Verifier does not take any group of public keys \mathcal{KG} as input. What we call processing of the group setup in dynamic fNIMs is the task of verifying the proofs of possession (PoP) of the published keys \mathcal{KG} . Recall that in *dms* we formalized this task as the *key verification* function $\text{kVf}(\text{pk}_i), \forall i \in [N]$. In Table 6 we consider the three fNIMs which have the fastest online verification: *SMSKR* [9], *MSP-pop* and *dms*. The online verification is identical in all of them, i.e., returns $\Sigma \in \mathbb{G}_1 \wedge e(\Sigma, G_2) == e\left(\text{H}(m), \sum_{i \in [N]} X_i\right)$. On the other hand, most of the runtime of the Verifiers in both *MSP-pop* and *SMSKR* (recalled in Sections D-3 and D-4) is explained by their processing of the group setups. As evidenced in Table 6, *dms* removes this bottleneck. We now detail the figures.

On the first line we measure the time of the processing of the group setup of a group of $|\mathcal{KG}| = 2702$ keys all-at-once. In what follows we denote $N = |\mathcal{KG}| = 2702$ for simplicity. This number was chosen as $2702 = 2|v|.193$, for $|v| = 7$. These numbers illustrate the use-case of the compiler *MtoA* applied to a group of $n_C = 193$ potential signers and to messages of $|v| = 7$ -bits-long variable parts. We batched the verifications of the N pairing-based PoPs as follows. First, we used the trick of [35] for reducing batch verification of N BLS signatures into a product of N pairings (recalled in Sec. D-4). Second, we computed this product using the optimized algorithm of gnark-crypto, inherited from [64]. Analogously, we batched the verification of the N PoPs in *dms*, by using the method of [14] for batch verification of Schnorr signatures (recalled in Sec. D-5).

On the second line we consider the incremental processing of the group setup in the scenario where: there is a group \mathcal{KG} of $N = 2702$ keys, over which processing of the group setup was already performed, and then there are 14 new keys $(\text{pk}'_i)_{i \in [14]}$ which join the group, in place of the old keys $(\text{pk}_i)_{i \in [14]}$. Note that this corresponds to the same use-case of *MtoA*, for messages of variable parts $|v| = 7$ bits, when one of the $n = 193$ real group members leaves and is replaced by a new member. This results in the new group $\mathcal{KG}' = (\mathcal{KG} \setminus (\text{pk}_i)_{i \in [14]}) \cup (\text{pk}'_i)_{i \in [14]}$. Since the resulting group \mathcal{KG}' is different from \mathcal{KG} , in *SMSKR* the Verifier needs to compute again all the 2702 rerandomized keys relatively to the new group \mathcal{KG}' , in addition to checking \mathbb{G}_1 membership of the 14 new keys. Hence, we see that the incremental processing of the group setup of *SMSKR* is nearly as costly as the processing of the group setup of a whole new group of keys, as evidenced by the first column of Table 6. Whereas in both *MSP-pop* and *dms*, the incremental processing of the group

setup only consists in verifying the PoPs of the 14 new keys (and the $\mathbb{G}_1/\mathbb{G}_2$ -memberships).

in ms	SMSKR [9]	<i>MSP-pop</i> [85][22, §6][24]	<i>dms</i>
Batch $ \mathcal{KG} $ keys	1134.9	1947.4	366.6
14 new keys	828.7	12.5	3.3

Table 6: processing of the group setup runtimes (in ms) for each new group of $|\mathcal{KG}| = 2702$ public keys, in three fNIMs. First line: for a group \mathcal{KG} of completely new keys. Second line: incremental processing of the group setup when 14 new keys join the group \mathcal{KG} .

2) *Comparison with previous provable securities:* In Table 7 we state proven upper-bounds: $\text{UB}^{\text{m-uf}}$ on the probability to forge a multisignature, i.e., the advantage in the m-uf game.

- *The first column* states the formulas of $\text{UB}^{\text{m-uf}}$, not directly in terms of the running time t of the adversary, but instead in terms of: q_H the number of its random oracle requests (hash-to-curve); q_s the number of its signing requests; and of the upper-bounds: UB^{DL} , $\text{UB}^{\text{co-cdh}}$, $\text{UB}^{\text{co-bdh}}$ and $\text{UB}^{\text{BLS-uf}}$ on the advantage in the games of DL, co-cdh, co-bdh ([27]) and of forgery against standalone BLS signatures. The upper-bounds are for any adversary with roughly the same running time t as the forger, neglecting additive time overheads. In the proofs, such additive overheads typically amount to roughly $+q_H \cdot \tau_{\text{exp}}$, where τ_{exp} is a scalar multiplication, also known as exponentiation. In line with [22, Thm 5], we corrected the bound for *MSP-pop* displayed in [85, Thm 4.1], in which q_H was replaced by the much smaller number q_s of signature queries. We also corrected a bug in the bounds of [9, Thms 1–4] (as confirmed by the authors on 26/1/2024).

- *The second column* states the models in which the formulas are proven: RO stands for random oracle, and “RMSS” is, for simplicity, the assumption that the “random modular subset sum” [9, Def. 3] is at least as hard as DL ([9, §C.3]).

- *The third column* are numerical applications when assuming furthermore the AGM. Concretely, under the AGM then $\text{UB}^{\text{co-bdh}} = \text{UB}^{\text{co-cdh}} = \text{UB}^{\text{dl}}$ [11]. We took $q_H = 2^{80}$; the number of clock cycles $t = 2^{80} \tau_{\text{exp}}$; $q_s = 2^{30}$ (in line with [13]); groups of size $p \sim 2^{253}$, and the hardness of DL estimated as $\text{UB}^{\text{dl}} \leq t^2/p$. The latter formula is shown for generic groups in [86] with t the number of group operations. Whereas, we apply it more conservatively to t the number of clock cycles. This estimate seems recently validated [70, 6] for both the popular curves BLS12-377 ($p \sim 2^{253}$), used in Zexe [31], and BLS12-381 ($p \sim 2^{255}$), used in Ethereum, since they are both estimated to have close to 126 bits of security. It is however estimated by Duquesne-Barbulescu and NCC group [10, 65] that both those BLS12 curves, in order to match this security, should be instantiated with a base prime q of size at least 460bits.

3) *Comparing verification times of MtoA + multi-BLS; vs BGLS:* In the first three lines of Table 8 we compare the online verification times of three aggregate signature schemes, for an aggregate signature over n messages. The messages are of the form $m_i = (\tau, v_i)$ with have any arbitrary common prefix τ and variable suffixes v_i , all of size $|v| = 7$ bits. On the last column with display the size of signatures, including the data of the public keys of the signers. Given a known

	Proven security: UB^{m-uf}			Verifier's efficiency		Dynamic	PoP needed
	theoretical	Model	$\log_2(p) \approx 253$	Group setup	Online		
[20, 77] if kosk model	UB^{BLS-uf}	RO	2^{-93} (if AGM)	●	●	●	
\mathcal{AS} -4 [72, 16]	UB^{co-cdh}	RO	2^{-93} (if AGM)	●	○	●	
\mathcal{MSP} -pop [85, Th 4.1][22, §6][24] & \mathcal{ASM} -pop [22, §6] & Pixel [47]	$q_H \cdot UB^{co-cdh}$	RO	2^{-13} (if AGM)	○	○	●	yes
\mathcal{MSP} -blog [21, 56]	$q_H^{3/2} \cdot \sqrt{UB^{co-cdh}}$	RO	1	●	○	●	
\mathcal{ASM} [22, §4.2]	$q_H^{3/2} \cdot \sqrt{UB^{co-cdh}}$	RO	1	○	●	○	
SMSKR [9]	$q_H^{3/2} \cdot \sqrt{UB^{dl}}$	RO+AGM+RMSS	1	○	●	○	
SIG_1 [27] (Sec. D-6)	$q_s \cdot UB^{co-bdh}$	RO	2^{-63} (if AGM)	○	○	○	
Squirrel & Chipmunk [52, 51]		(Lattice-based)				●	
dms (this work)	UB^{dl}	RO+AGM	2^{-93}	●	●	●	yes

Table 7: fNIMs

group of public keys \mathcal{KG} , the public keys of the subgroup of n signers can be encoded as a $|\mathcal{KG}|$ -sized array of bits. We approximated $|\mathcal{KG}| \cong n$, since $|\mathcal{KG}| = n_C = (3/2)n$ in Diem21-like consensus algorithms. On the last two lines, in grey: we display the times for BLS multisignatures and threshold signatures, which is of course not an apples-to-apples comparison. *In more detail:*

First line: BGLS [23, 15, 37, 73, 22], of which the verification takes $n+1$ pairings (recalled in Sec. D-1). More precisely, we evaluated verification of the fastest variant of BGLS: $e(\Sigma, \mathbb{G}_2) == \sum_{i \in [n]} e(H(m_i), X_i)$. This variant, dubbed \mathcal{AS} -1 in [15], is restricted to pairwise different messages m_i . Hence, for a fair comparison with \mathcal{MtoA} + dms, which is unrestricted and has tight security, we should have instead evaluated the costlier variant of BGLS called \mathcal{AS} -4 in [15]. The verification would then have taken even more time. The signature is a \mathbb{G}_1 element, which has uncompressed size equal to 92bytes = 768bits.

Second line: \mathcal{MtoA} instantiated with any BLS-based fNIM with optimal online verification, i.e., either \mathcal{MSP} -pop, SMSKR or our dms. Namely, a \mathcal{MtoA} signature comes as a multisignature over $N = |v|n = 7n$ keys, and its verification is as in dms without verification of PoPs, i.e., as in [20], recalled in Sec. D-2. As in consensus protocols, we considered that the Verifier knows the tag τ in advance, and thus could pre-compute $H(\tau)$.

Third line: naive concatenation of n Schnorr signatures. Since no pairing is necessary, we used the faster curve secp256k1. Each signature is of the form (R, z) , where the \mathbb{G}_1 element R is now only of size 64bytes, and the \mathbb{Z}_p -element z is of size 32bytes.

Fourth line: multisignature over n signers with any BLS-based fNIM with optimal online verification, i.e., either \mathcal{MSP} -pop, SMSKR or our dms.

Fifth line: BLS threshold signature [20, 12]. The signature output is a (standard, single-key) BLS signature.

VIII. ACKNOWLEDGEMENTS

We thank Zhuolun Xiang for informing of the use of aggregate signatures in the production implementation of Diem21 by Aptos [5].

times in ms	$n = 129$	$n = 3073$	size (bits)
BLS aggregate sig.	116.6	2661.6	$768 + n$
\mathcal{MtoA} with BLS multisig.	3.4	35.9	$768 + n$
Batch Schnorr	1.9	22.4	$768n + n$
BLS multisig.	2.4	7.1	$768 + n$
BLS threshold sig.	1.9	1.9	768

Table 8: Online verification times over n messages (see above)

REFERENCES

- [1] M. A. Aardal, D. F. Aranha, K. Boudgoust, S. Kolby, and A. Takahashi. *Aggregating Falcon Signatures with LaBRADOR*. ePrint 2024/311. 2024.
- [2] M. Abdalla, F. Benhamouda, and P. MacKenzie. “Security of the J-PAKE Password-Authenticated Key Exchange Protocol”. In: *IEEE SP*. 2015.
- [3] S. Agrawal, J. Neu, E. N. Tas, and D. Zindros. “Proofs of Proof-Of-Stake with Sublinear Complexity”. In: *AFT*. 2023.
- [4] J. H. Ahn, M. Green, and S. Hohenberger. “Synchronized aggregate signatures: new definitions, constructions and applications”. In: *CCS*. 2010.
- [5] Aptos. *Implementation of Aptos consensus, following Diem*. https://github.com/aptos-labs/aptos-core/blob/main/consensus/consensus-types/src/timeout_2chain.rs Retrieved on June 23, 2024. 2024.
- [6] D. F. Aranha, Y. E. Housni, and A. Guillevic. *A survey of elliptic curves for proof systems*. Des. Codes, Cryptogr. 2022.
- [7] R. Bacho and J. Loss. “On the Adaptive Security of the Threshold BLS Signature Scheme”. In: *CCS*. 2022.
- [8] A. Bagherzandi, J.-H. Cheon, and S. Jarecki. “Multisignatures Secure under the Discrete Logarithm Assumption and a Generalized Forking Lemma”. In: *CCS*. 2008.
- [9] F. Baldimtsi, K. K. Chalkias, F. Garillot, J. Lindstrom, B. Riva, A. Roy, A. Sonnino, P. Waiwitlikhit, and J. Wang. “Subset-optimized BLS Multi-signature with Key Aggregation”. In: *FC*. 2024.
- [10] R. Barulescu and S. Duquesne. *Updating key size estimations for pairings*. JOC. 2018.

- [11] B. Bauer, G. Fuchsbauer, and J. Loss. “A Classification of Computational Assumptions in the Algebraic Group Model”. In: *CRYPTO*. 2020.
- [12] M. Bellare, E. Crites, C. Komlo, M. Maller, S. Tessaro, and C. Zhu. “Better than Advertised Security for Non-interactive Threshold Signatures”. In: *CRYPTO*. 2022.
- [13] M. Bellare and W. Dai. “Chain Reductions for Multi-Signatures and the HBMS Scheme”. In: *ASIACRYPT*. 2021.
- [14] M. Bellare, J. A. Garay, and T. Rabin. “Fast batch verification for modular exponentiation and digital signatures”. In: *EUROCRYPT*. 1998.
- [15] M. Bellare, C. Namprempre, and G. Neven. “Unrestricted Aggregate Signatures”. In: *ICALP*. long version. 2007.
- [16] M. Bellare and G. Neven. “Multi-signatures in the plain public-key model and a general forking lemma”. In: *CCS*. 2006.
- [17] M. Bellare and P. Rogaway. “Code-Based Game-Playing Proofs and the Security of Triple Encryption”. In: *EUROCRYPT*. 2006.
- [18] D. Bernhard, M. Fischlin, and B. Warinschi. “On the Hardness of Proving CCA-security of Signed ElGamal”. In: *PKC* (2016).
- [19] W. Beullens and G. Seiler. “LaBRADOR: Compact Proofs for R1CS from Module-SIS”. In: *CRYPTO*. 2023.
- [20] A. Boldyreva. “Threshold Signatures, Multisignatures and Blind Signatures Based on the Gap-Diffie-Hellman-Group Signature Scheme”. In: *PKC*. Latest long version at <https://faculty.cc.gatech.edu/~aboldyre/papers/b.pdf>. 2003.
- [21] D. Boneh, M. Drijvers, and G. Neven. *Bls multi-signatures with public-key aggregation*. <https://crypto.stanford.edu/~dabo/pubs/papers/BLSmultisig.html>. 2018.
- [22] D. Boneh, M. Drijvers, and G. Neven. “Compact Multi-signatures for Smaller Blockchains”. In: *ASIACRYPT*. 2018.
- [23] D. Boneh, C. Gentry, B. Lynn, and H. Shacham. “Aggregate and Verifiably Encrypted Signatures from Bilinear Maps”. In: *EUROCRYPT*. 2003.
- [24] D. Boneh, S. Gorbunov, R. S. Wahby, H. Wee, C. A. Wood, and Z. Zhang. *BLS Signatures draft-irtf-cfrg-bls-signature-05*. <https://datatracker.ietf.org/doc/draft-irtf-cfrg-bls-signature/>. 2022.
- [25] D. Boneh and S. Kim. *One-Time and Interactive Aggregate Signatures from Lattices*. https://crypto.stanford.edu/~skim13/agg_ots.pdf. 2020.
- [26] D. Boneh, B. Lynn, and H. Shacham. “Short Signatures from the Weil Pairing”. In: *J. Cryptol.* (2004).
- [27] D. Boneh, A. Partap, and B. Waters. *Accountable Multi-Signatures with Constant Size Public Keys*. ePrint 2023/1793. 2023.
- [28] G. Botrel, T. Piellard, Y. E. Housni, A. Tabaie, and I. Kubjas. *ConsenSys/gnark-crypto: v0.6.1*. 2022.
- [29] K. Boudgoust, E. Gachon, and A. Pellet-Mary. “Some Easy Instances of Ideal-SVP and Implications on the Partial Vandermonde Knapsack Problem”. In: *CRYPTO*. 2022.
- [30] K. Boudgoust and A. Takahashi. *Sequential Half-Aggregation of Lattice-Based Signatures*. ePrint 2023/159. 2023.
- [31] S. Bowe, A. Chiesa, M. Green, I. Miers, P. Mishra, and H. Wu. “ZEXE: Enabling Decentralized Private Computation”. In: *IEEE SP*. 2020.
- [32] M. Bravo, G. V. Chockler, and A. Gotsman. “Liveness and Latency of Byzantine State-Machine Replication”. In: *DISC*. 2022.
- [33] M. Bravo, G. V. Chockler, and A. Gotsman. “Making Byzantine consensus live”. In: *Distributed Comput.* (2022).
- [34] J. Burdges, O. Ciobotaru, S. Lavasani, and A. Stewart. *Efficient Aggregatable BLS Signatures with Chaum-Pedersen Proofs*. ePrint 2022/1611. 2022.
- [35] J. Camenisch, S. Hohenberger, and M. Ø. Pedersen. “Batch Verification of Short Signatures”. In: *EUROCRYPT*. 2007.
- [36] B. Y. Chan and R. Pass. “Simplex Consensus: A Simple and Fast Consensus Protocol”. In: *TCC*. 2023.
- [37] S. Chatterjee, D. Hankerson, E. Knapp, and A. Menezes. “Comparing Two Pairing-Based Aggregate Signature Schemes”. In: *Des. Codes, Cryptogr.* (2010).
- [38] H. Cheng, Y. Lu, Z. Lu, Q. Tang, Y. Zhang, and Z. Zhang. *JUMBO: Fully Asynchronous BFT Consensus Made Truly Scalable*. 2024.
- [39] G. Chockler. *Modular Construction of Live Byzantine Consensus Protocols*. an abstract appearing on <https://lp.jetbrains.com/sptdc-2023/>. 2023.
- [40] S. Cohen, R. Gelashvili, E. Kokoris-Kogias, Z. Li, D. Malkhi, A. Sonnino, and A. Spiegelman. “Be Aware of Your Leaders”. In: *FC*. 2022.
- [41] E. Crites, C. Komlo, and M. Maller. “Fully Adaptive Schnorr Threshold Signatures”. In: *CRYPTO*. 2023.
- [42] E. Crites, C. Komlo, and M. Maller. *How to Prove Schnorr Assuming Schnorr: Security of Multi- and Threshold Signatures*. ePrint 2021/1375. Merged into CRYPTO’22 “Better than advertised security”. 2021.
- [43] S. Das, P. Camacho, Z. Xiang, J. Nieto, B. Bunz, and L. Ren. “Threshold Signatures from Inner Product Argument: Succinct, Weighted, and Multi-threshold”. In: *CCS*. 2023.
- [44] Diem. *DiemBFT v4: State Machine Replication in the Diem Blockchain*. <https://developers.diem.com/papers/diem-consensus-state-machine-replication-in-the-diem-blockchain/2021-08-17.pdf>. 2021.
- [45] Y. Doröz, J. Hoffstein, J. H. Silverman, and B. Sunar. *MMSAT: A Scheme for Multimessage Multiuser Signature Aggregation*. ePrint 2020/520. 2020.
- [46] M. Drijvers, K. Edalatnejad, B. Ford, E. Kiltz, J. Loss, G. Neven, and I. Stepanovs. “On the Security of Two-Round Multi-Signatures”. In: *IEEE Security and Privacy*. 2019.
- [47] M. Drijvers, S. Gorbunov, G. Neven, and H. Wee. “Pixel: Multi-signatures for Consensus”. In: *USENIX security*. 2020.
- [48] B. Edgington. *Upgrading Ethereum*. <https://eth2book.info/latest/book.pdf>. 2023.
- [49] M. F. Esgin, O. Ersoy, V. Kuchta, J. Loss, A. Sakzad, R. Steinfeld, X. Yang, and R. K. Zhao. “A New Look at Blockchain Leader Election: Simple, Efficient, Sustainable and Post-Quantum”. In: *AsiaCCS*. 2023.

- [50] Ethereum. *Ethereum Altair upgrade*. <https://github.com/ethereum/consensus-specs/blob/dev/specs/altair/bls.md>. 2023.
- [51] N. Fleischhacker, G. Herold, M. Simkin, and Z. Zhang. “Chipmunk: Better Synchronized Multi-Signatures from Lattices”. In: *CCS*. 2023.
- [52] N. Fleischhacker, M. Simkin, and Z. Zhang. “Squirrel: Efficient Synchronized Multi-Signatures from Lattices”. In: *CCS*. 2022.
- [53] G. Fuchsbauer, E. Kiltz, and J. Loss. “The Algebraic Group Model and its Applications”. In: *CRYPTO*. 2018.
- [54] G. Fuchsbauer, A. Plouviez, and Y. Seurin. “Blind Schnorr Signatures and Signed ElGamal Encryption in the Algebraic Group Model”. In: *EUROCRYPT*. 2020.
- [55] S. D. Galbraith, K. G. Paterson, and N. P. Smart. “Pairings for cryptographers”. In: *Discret. Appl. Math.* (2008).
- [56] D. Galindo and J. Liu. “Robust Subgroup Multi-signatures for Consensus”. In: *CT-RSA*. 2022.
- [57] S. Garg, A. Jain, P. Mukherjee, R. Sinha, M. Wang, and Y. Zhang. “hinTS: Threshold Signatures with Silent Setup”. In: *IEEE SP*. 2024.
- [58] R. Gelashvili, L. Kokoris-Kogias, A. Sonnino, A. Spiegelman, and Z. Xiang. “Jolteon and Ditto: Network-Adaptive Efficient Consensus with Asynchronous Fallback”. In: version 2024-04-30, fixing the FC’22 version and the 2023-12 version. 2024.
- [59] C. Gentry, A. O’Neill, and L. Reyzin. *A Unified Framework for Trapdoor-Permutation-Based Sequential Aggregate Signatures*. ePrint 2018/070. 2018.
- [60] N. Giridharan, H. Howard, I. Abraham, N. Crooks, and A. Tomescu. *No-commit proofs: Defeating livelock in bft*. eprint 2021/1308. 2021.
- [61] G. Golan-Gueta, I. Abraham, S. Grossman, D. Malkhi, B. Pinkas, M. K. Reiter, D. Seredinschi, O. Tamir, and A. Tomescu. “SBFT: A Scalable and Decentralized Trust Infrastructure”. In: *DSN*. 2019.
- [62] S. Goldwasser, S. Micali, and R. L. Rivest. “A Digital Signature Scheme Secure Against Adaptive Chosen-Message Attacks”. In: *SIAM J. Comput.* (1988).
- [63] B. Goodell and A. Feickert. *Fusion One-Time Non-Interactively-Aggregatable Digital Signatures From Lattices*. ePrint 2023/303. 2023.
- [64] R. Granger and N. P. Smart. “On Computing Products of Pairings”. In: *eprint 2006/172* (2006).
- [65] N. Group. *Zcash Overwinter Consensus and Sapling Cryptography Review*. https://research.nccgroup.com/wp-content/uploads/2020/07/NCC_Group_Zcash2018_Public_Report_2019-01-30_v1.3.pdf. 2019.
- [66] B. Guo, Y. Lu, Z. Lu, Q. Tang, J. Xu, and Z. Zhang. “Speeding Dumbo: Pushing Asynchronous BFT Closer to Practice”. In: *NDSS*. 2022.
- [67] K. Guo, K. Hu, and Z. Zhang. “Liveness Attacks on HotStuff: The Vulnerability of Timer Doubling Mechanism”. In: *The Computer Journal* (2024).
- [68] C. Héban and D. Pointcheval. “Traceable Constant-Size Multi-Authority Credentials”. In: *SCN*. 2022.
- [69] S. Hohenberger and B. Waters. “Synchronized Aggregate Signatures from the RSA Assumption”. In: *EUROCRYPT*. 2018.
- [70] Y. E. Housni and A. Guillevic. *Optimized and secure pairing-friendly elliptic curves suitable for one layer proof composition*. CANS. 2020.
- [71] M. M. Jalalzai, J. Niu, C. Feng, and F. Gai. “Fast-HotStuff: A Fast and Resilient HotStuff Protocol”. In: *IEEE Transactions on Dependable and Secure Computing* (2023).
- [72] J. Katz and N. Wang. “Efficiency improvements for signature schemes with tight security reductions”. In: *CCS*. 2003.
- [73] M. Lacharité. “Security of BLS and BGLS signatures in a multi-user setting”. In: *Cryptogr. Commun.* (2018).
- [74] A. Lewis-Pye and I. Abraham. “Fever: OptiFmal Responsive View Synchronisation”. In: *Opodis*. 2023.
- [75] A. Lewis-Pye, D. Malkhi, O. Naor, and K. Nayak. “Lumiere: Making Optimal BFT for Partial Synchrony Practical”. In: *Podc*. 2024.
- [76] Z. Li, A. Sonnino, and P. Jovanovic. “Performance of EdDSA and BLS Signatures in Committee-Based Consensus”. In: *ApPLIED at PODC*. 2023.
- [77] S. Lu, R. Ostrovsky, A. Sahai, H. Shacham, and B. Waters. “Sequential Aggregate Signatures and Multisignatures Without Random Oracles”. In: *EUROCRYPT*. 2006.
- [78] D. Malkhi, C. Stathakopoulou, and M. Yin. “BBCA-CHAIN: One-Message, Low Latency BFT Consensus on a DAG”. In: *FC*. 2024.
- [79] S. Micali, K. Ohta, and L. Reyzin. “Accountable-Subgroup Multisignatures: Extended Abstract”. In: *CCS*. 2001.
- [80] O. Mir, B. Bauer, S. Griffy, A. Lysyanskaya, and D. Slamanig. “Aggregate Signatures with Versatile Randomization and Issuer-Hiding Multi-Authority Anonymous Credentials”. In: *CCS*. 2023.
- [81] O. Naor, M. Baudet, D. Malkhi, and A. Spiegelman. “Cogsworth: Byzantine View Synchronization”. In: *arxiv 1909.05204* (2019).
- [82] O. Naor and I. Keidar. “Expected Linear Round Synchronization: The Missing Link for Linear Byzantine SMR”. In: *DISC*. 2020.
- [83] P. Paillier and D. Vergnaud. “Discrete-Log-Based Signatures May Not Be Equivalent to Discrete Log”. In: *ASIACRYPT*. 2005.
- [84] M. Rambaud. “(Section 6 of version 2020-11-29) Malicious Security Comes for Free in Consensus with Leaders”. In: *eprint 2020/1480* (2020).
- [85] T. Ristenpart and S. Yilek. “The Power of Proofs-of-Possession: Securing Multiparty Signatures against Rogue-Key Attacks”. In: *EUROCRYPT*. 2007.
- [86] V. Shoup. “Lower Bounds for Discrete Logarithms and Related Problems”. In: *EUROCRYPT*. 1997.
- [87] A. Spiegelman, N. Giridharan, A. Sonnino, and L. Kokoris-Kogias. “Bullshark: DAG BFT Protocols Made Practical”. In: *CCS*. 2022.
- [88] X. Sui, S. Duan, and H. Zhang. “Marlin: Two-Phase BFT with Linearity”. In: *DSN*. 2022.
- [89] T. Tomita and J. Shikata. *Compact Aggregate Signature from Module-Lattices*. ePrint 2023/471. 2023.
- [90] M. Yin, D. Malkhi, M. K. Reiter, G. Golan-Gueta, and I. Abraham. “HotStuff: BFT Consensus with Linearity and Responsiveness”. In: *PODC*. we refer to the arxiv v6 long version. 2019.

APPENDIX A

FURTHER FORMALIZATION AND OPTIMIZATION OF \mathcal{M} to \mathcal{A}

All ideas were conveyed in Sec. IV, we now put them in the formalism of Sec. III.

Theorem 8. *Let $M = (\text{Kg}, \text{Sign}, \text{iVf}, \text{Cb}, \text{Vf})$ be any fNIM and $|v| > 0$ an integer; then the following scheme $A := (A.\text{Kg}, A.\text{Sign}, A.\text{iVf}, A.\text{Ag}, A.\text{Vf})$ is a one-time-tagged fNIA supporting messages with variable parts of bitlength $|v|$.*

- $A.\text{Kg}()$: $(\text{sk}_i^{j,b}, \text{pk}_i^{j,b}) \leftarrow \text{Kg}() \quad \forall (j \in [|v|], b \in \{0, 1\})$;
 $\text{sk}_i \leftarrow (\text{sk}_i^{j,b})_{j \in [|v|], b \in \{0,1\}}$; $\text{pk}_i \leftarrow (\text{pk}_i^{j,b})_{j \in [|v|], b \in \{0,1\}}$
output $(\text{sk}_i, \text{pk}_i)$
- $A.\text{Sign}(\mathcal{KG}, \text{sk}_i, \tau_i, v_i)$: *decompose in bits* $v_i = (v_i^j)_{j \in [|v|]}$;
output $\Sigma_i \leftarrow \left(\text{Sign}(\mathcal{KG}, \text{sk}_i^{j,v_i^j}, \tau_i) \right)_{j \in [|v|]}$
- $A.\text{iVf}(\mathcal{KG}, \text{pk}_i, \tau_i, v_i, \Sigma_i)$: *decompose in bits*
 $v_i = (v_i^j)_{j \in [|v|]}$;
output $\bigwedge_{j \in [|v|]} \text{iVf}(\mathcal{KG}, \text{pk}_i^{j,v_i^j}, \tau_i, \Sigma_i^j)$
- $A.\text{Ag}(\mathcal{KG}, \tau, (\text{pk}_i, \Sigma_i, v_i)_{i \in [n]})$:
decompose in bits $v_i = (v_i^j)_{j \in [|v|]} \quad \forall i \in [n]$;
output $\text{Cb}(\mathcal{KG}, \tau, (\text{pk}_i^{j,v_i^j}, \Sigma_i^j)_{j \in [|v|], i \in [n]})$
- $A.\text{Vf}(\mathcal{KG}, \tau, (\text{pk}_i, v_i)_{i \in [n]}, \Sigma)$:
decompose in bits $v_i = (v_i^j)_{j \in [|v|]} \quad \forall i \in [n]$;
output $\text{Vf}(\mathcal{KG}, (\text{pk}_i^{j,v_i^j})_{i \in [n], j \in [|v|]}, \tau, \Sigma)$

Proof: Both individual completeness and robustness follow straightforwardly from the ones of M , let us prove unforgeability. Consider a forger \mathcal{F} in the game **a-uf** of Figure 3, and the event where it wins the game. Namely, it outputs $(\mathcal{KG}, \tau^*, (\text{pk}_i, v_i)_{i \in [n]}, \Sigma^*)$, such that there exists $(\text{pk}, v^*) \in (\text{pk}_i, v_i)_{i \in [n]}$ such that no query on (τ^*, v^*) was ever responded by the oracle SIGN. By assumption, SIGN responded to at most one query prefixed by τ^* . Without loss of generality, we can assume that exactly one such query was made and responded to: $\Sigma \leftarrow (\tau^*, v')$. By the above $v' \neq v^*$, thus there is a bit index $j_0 \in [|v|]$ at which $b_0 := (v^*)^{j_0} \neq (v')^{j_0}$, thus no signature of pk^{j_0, b_0} on τ was ever delivered by SIGN. Since the other keys of the honest signer: $(\text{pk}^{j,b})_{(j,b) \neq (j_0, b_0)}$ are generated independently from pk^{j_0, b_0} , the intuitive conclusion is that Σ constitutes a M-forgery on the message τ with respect to the signers $(\text{pk}_i^{j,v_i^j})_{i \in [n], j \in [|v|]}$ and target key pk^{j_0, b_0} .

Of course the indices (j_0, b_0) are not known in advance. So to make this argument rigorous and build from \mathcal{F} a forger against M , the reduction must choose at random an index $(j_0, b_0) \in [|v|] \times \{0, 1\}$, and embed its target key in pk^{j_0, b_0} . Thus we have at most $2|v|$ loss compared to the $(2|v|n)$ -unforgeability of M . ■

a) Optimization for variable parts of variable lengths:

In the case of messages with shorter variable bitlengths than the maximum: $|v'| < |v|$, is very easy to enable multisignatures with respect to the shorter set of keys $(\text{pk}_i^{j,b})_{j \in [|v'|], b \in \{0,1\}}$ for each signer. A global adjustment consists in encoding $|v'|$

in the tag for all messages. It is possible to achieve further flexibility and let each signer i adjust the length of its signature depending on the length of its variable part v_i . To this end, add to each public key a list of $|v|$ keys, where the j -th key is used for signing the end-of-string at position $j \in [|v|]$.

APPENDIX B

DETAILS FOR APPLICATIONS OF dms

a) To Blockchain Consensus: The main elementary operation in all such protocols, e.g., [58, 87, 66, 78, 36], is: one (or several) designated Combiner(s) wait(s) to receive a sufficiently large number of signatures, say N , on the same message content: m , then combine the signatures into Σ . Then it multicasts Σ to all the participants to the consensus, dubbed the *processes*. Moreover, Σ is often meant and verified by billions of external clients, since in most cases Σ attests validity of a block. Since in addition Σ is stored on-chain, it is therefore a first-class requirement that Σ be both small and fast to verify. As shown in Table 8, pairing-based multisignatures schemes are the most advantageous instantiation of Σ with this respect, since they take close to $N \times$ less storage space than a naive concatenation of Schnorr signatures. Furthermore, at least for $N \geq 3073$, they have $3 \times$ smaller Verifier runtime. The last hurdle to their adoption, as stressed in [88], was the runtime of $O(n)$ pairings required to verify pairing-based PoPs. This hurdle is now removed by dms. Still, surprisingly, a number of implementations of consensus protocols instantiate fNIMs as mere concatenations of signatures [58, 87, 66, 38], instead of BLS-based multisignatures. The reason invoked [76, 38] is the verification time of an individual BLS signature [26, 37](algorithm iVf), which takes 2 pairings. We observe that there exists known ways around this potential runtime gap. First, in case the Combiner would receive individual BLS signatures faster than it can verify individually, then it can simply combine them and check them as a multisignature: this was empirically confirmed by [38]. In the rare events where one ill-formed signature would make this batch verification fail, the cheater which issued it would be publicly identified so this is a strong deterrence. Second, there is a more or less known method enabling a faster iVf, which is proposed in [56, 34]. The signer, in addition to its signature: $\sigma_i = \text{sk}_i.H(m)$, appends to it a ‘‘Chaum-Pedersen’’ proof: π , of knowledge of a common discrete logarithm: sk_i between σ_i and the public key $X_i := \text{sk}_i.G_2$. Then the verification algorithm: iVf^{DLEQ} verifies only π against σ_i and X_i , not anymore σ_i against X_i . Our implementation of iVf^{DLEQ} , with the same configuration as in Sec. VII (same machine, gnark-crypto, BLS-377 curve), shows a runtime of 0.785ms, down from 1.9ms for the iVf of a standard BLS signature (Table 8). So this reduces the gap w.r.t. our verification time of a Schnorr signature on secp256k1, which is of 0.220ms.

b) To Threshold Signatures: The recent weighted threshold signature schemes [43, 57] are constructed on the top of BLS multisignatures, they both operate as follows. The list of published keys is denoted $(X_i)_{i \in [N]}$ and their weights $(w_i)_{i \in [N]}$. The Combiner collects valid individual BLS signatures: $(\Sigma_i)_{i \in I}$ on some message m , issued by a subset $I \subset [N]$, totalizing some desired weight: $w := \sum_{i \in I} w_i$. It outputs the public multisignature $\Sigma := \sum_{i \in I} \Sigma_i$ and the public weight w . It also outputs a proof of knowledge π of

$I \subset [N]$, encoded as a N -sized binary vector $(b_i)_{i \in [N]}$, verifying the following (bi)linear relations: (i) $w = \sum_{i \in [N]} b_i w_i$ and (ii) $e(\Sigma, G_2) = e(H(m), \sum_{i \in [N]} b_i X_i)$. Note that (i) and (ii) together ensure that Σ passes the BLS multisignature verification (m-blsVf*) against the public keys $(X_i)_{i \in [N]}$, and that they totalize weight w . However, the sole passing of (m-blsVf*) does not guarantee unforgeability: as the reader knows well, some processing of the group setup must be done on the group of keys. In [43] the authors suggest using *MSP-pop* [85, 22], where the Verifier verifies pairing-based PoPs appended to the published keys $(X_i)_{i \in [N]}$. Using instead dms, i.e., Schnorr-based PoPs, divides by $> 5 \times$ this latter runtime, as demonstrated in Table 6.

APPENDIX C

DETAILS FOR APPLICATION OF \mathcal{M}_{toA} TO CONSENSUS

In Figure 9 we further recall the consensus Diem21 [44] among $n_C = 3f + 1$ processes, of which f are corrupt. Diem21 was used in production by Meta, today by Aptos, and should not be confused with previous versions of Diem, as presented in [58, Fig. 1], which instead followed Hotstuff [90]. Then in Figure 10 we formalize how \mathcal{M}_{toA} can be straightforwardly plugged: either in place of naive concatenation of $(2/3)n_C$ signatures (in [44]) or in place of the BGLS [23] aggregate signature (in the production version [5]).

Diem21 proceeds by iterations called rounds, each with a designated process called the *leader*. We borrow freely from the terminology of [58, 36]. Although our presentation follows Jolteon (of which the timeout certificates were very recently fixed in [58]), we stick to the unusual specification of a new-round appearing in Diem21 [44]. We highlight it (in red) in Figure 9. The reason for not choosing the mainstream specification ([61, 58, 78]) of a new-round message, is that the latter mainly consists of a signature on a quorum certificate (QC). Since a QC is typically a multisignature, these objects are not efficiently aggregatable.

We however depart from Diem21 in that our model abstracts-out the view-synchronizer. Let us recall in more detail that a *view-synchronizer* [74] is a protocol enabling players to advance their local round numbers: $r_{\text{new}} \leftarrow r$ in two ways. Either (i) upon receiving a round- $(r_{\text{new}}-1)$ -QC from the consensus protocol, which we left explicit in the protocol; or, (ii) upon outputting, from the view-synchronizer protocol, a signal of the form: $(\text{NEWROUND}, r_{\text{new}})$ for $r_{\text{new}} > r$. A view-synchronizer protocol should guarantee that, eventually, honest players are in the same round for a sufficiently long time, and that this happens infinitely often. Liveness of a consensus protocol is then conditioned to this guarantee. Hence, we consider the hybrid model, where a process goes to a round r_{new} either upon receiving a r_{new} -QC, or a signal $(\text{NEWROUND}, r_{\text{new}})$ from a black box view-synchronizer. In this hybrid model, Diem21 enjoys a linear number of messages per round since communications are star-shaped around leaders.

a) *Why abstracting-out view-synchronization ?*: Since the main claim of the Hotstuff consensus [90], i.e., linear communication complexity and responsiveness, is stated in the hybrid model of an abstract view-synchronizer, we choose the same model in order to make an apples-to-apples comparison. Even though an unproven implementation of view-synchronizer was suggested in [90], under

the name ‘‘Pacemaker’’, an attack breaking its liveness was recently shown in [67]. Our choice is also motivated by readability, since our contributions are orthogonal from view-synchronization. Last, even though the protocol Diem21 [44], which we use as baseline, innovated with a nice ‘‘Bracha’’ view-synchronizer, it is now advocated by specialists ([39]) to instead abstract-out view synchronizers, and delegate their implementation to recent dedicated papers with thorough proofs and tight performances [81, 75, 33, 82, 32, 74, 75]. The one of [75] has communication complexity in nf' , where f' is the actual number of faults in the execution. As a side-remark, it is actually not hard to imagine how to divide the complexity of the view-synchronizer of Diem21 [44] as follows. Instead of appending their highest QC to the new-view message which they multicast, processes need only appending it to one which they send to the next leader. Thus, provided an implementation of QCs with mere concatenation of signatures, the total communication complexity and verification complexity would both drop from $O(n^3)$ down to $O(n^2)$.

1) *Terminology*: Multicast is the instruction to send a message to all, so nothing prevents processes from receiving different messages if the sender is corrupt. We denote $\langle m \rangle_i$ a standalone signature of \mathcal{P}_i on the message m , we then say that \mathcal{P}_i is the signer of the ‘‘signed message $\langle m \rangle_i$ ’’. We denote $\{m\}_i$ an individual signature of \mathcal{P}_i on the message m , and $\{m\}$ a $(2f + 1)$ -multisignature on m . It is a triple consisting of: m , a $(2f + 1)$ -sized subset $J \subset [n]$, and a multisignature Σ on m which is valid w.r.t. the public keys of J .

- **Round Number**. The protocol runs in sequential iterations called rounds $r = 1, 2, 3, \dots$ where each player starts in round $r = 1$. Note that each player may advance through rounds at a different speed, and at any given time, two players may be in two different rounds due to network delay (since we are in the partially synchronous setting). As local state, each player $\mathcal{P} \in (\mathcal{P}_i)_{i \in [n]}$ keeps track of which round r it is currently in (formerly denoted r_{curr} in [58]). It also stores all of the certified blocks that it has seen thus far, to be defined below. Additionally, we assume that each round r has a pre-determined block proposer called leader: $\text{lead}(r) \in (\mathcal{P}_i)_{i \in [n]}$. It may be randomly or deterministically chosen ahead of time, e.g., [40, 49]; this is referred to as a leader election oracle.
- **Block format**. A block is formatted as $b = [id, qc, rc, r, txn]$ where:
 - $id = H(qc, r, v, txn)$ is the unique hash digest of (qc, r, v, txn) ;
 - qc is a quorum certificate (QC: defined below) of the parent block of b ;
 - r is the round number of b ;
 - rc is either (1) a round- $(r-1)$ new-round (see below), or (2) $rc = \perp$ if qc is a round- $(r-1)$ -QC, i.e., $qc.r = r-1$;
 - txn is a batch of new transactions;

Note that when describing the protocol, it suffices to specify qc and r for a new block, since txn and id follow the definitions. We will use $b.x$ to denote the element x of b .

- **Quorum certificate (QC)**. A QC: qc for a block b is a multi (or threshold) signature for the message $(b.id, b.r)$, produced by combining the individual signatures $\{b.id, b.r\}$ from any set of $2f + 1$ players. The round number of a QC: qc for a block b is denoted by $qc.r$ which is equal to $b.r$. QCs are

ranked by their round numbers, hence, we abuse notation and shorten as $qc \geq qc'$ the relation $qc.r \geq qc'.r$. Since the QC contained in a block determines its unique parent block, and since the genesis block is the common ancestor of all blocks, the total data structure forms a tree of blocks. A branch is called a “blockchain”. We use $b \leftarrow b'$ to denote a 2-chain, i.e., a block b' of which the parent is the block b , i.e., such that $b'.qc$ is a QC of b . Each player stores the highest QC: qc_{high} , which is the QC with the highest round number, which it ever received or formed. For convenience we denote $r_{high} := qc_{high}.r$ its round number.

- **New-round message and New-round certificate (RC).** Jolteon & Diem21, [58, 44] make use of data structures called a Timeout Message (tmo) and a Timeout Certificate (TC). We present their definitions as appearing in Diem21 [44]. We rename them a *new-round message* and a *new-round certificate* RC), for at least two reasons.
 - First, for compatibility with the syntax of view-synchronizers, we specify that players send their round- r new-round message just after entering the new round r , instead of, in [58, 44], sending their round- $(r-1)$ timeout message upon timing-out in the old round $r-1$. This difference is merely syntactical, processes take the same actions. Namely: both in [58, 44] and in our presentation, players do not either cast anymore round- $(r-1)$ votes after sending their timeout message.
 - Second, because these data structures do not play anymore a role in the implementation of the view-synchronizer, as they did in [58, 44], since in our model the view-synchronizer is abstracted-out.

A round- r *new-round message* by a player i consists of the players’s qc_{high} , and of its individual signature on the pair (r, r_{high}) : $\langle r, r_{high} \rangle_i$, where we recall $r_{high} := qc_{high}.r$. A round- r RC is meant to be formed out of $2f+1$ new-rounds. In the written specifications [44] (recalled in Figure 9), a round- r RC consists of the naive concatenation of $2f+1$ signed pairs from distinct issuers: $rc \leftarrow \left[\langle r, r_j \rangle_j : j \in J \right]$, where $J \subset [n]$ is a $(2f+1)$ -sized subset: we follow this. Whereas in the production version [5], the RC is the BGLS aggregate signature $rc \leftarrow \text{BGLS.Ag} \left[\langle r, r_j \rangle_j : j \in J \right]$. Note that this presentation simplifies the slightly more compact encoding in [44], where the common prefix r is factored out of the $2f+1$ signed messages. Following [84], note that rc plays the role of a *proof of non-supermajority* because it guarantees that no set of $f+1$ honest players, upon entering round r , could have previously voted for a block containing a QC of round strictly higher than $r_{max} := \max(r_j, j \in J)$.

Diem21 with black-box view-synchronizer

Instructions for each player $\mathcal{P}_i, i \in [n]$ in (local) round r (denoted r_{curr} in Jolteon). It keeps the highest voted round r_{vote} , the highest QC: qc_{high} and the highest locked round $r_{high} := qc_{high}.r$. Players initialize $r_{vote} = 0, r_{high} = 0, qc_{high}$ as the QC of the genesis block of round 0, and enter round $r = 1$. The leader of round r is denoted $lead(r)$.

Propose Upon entering round r , if \mathcal{P}_i is the leader $lead(r)$, then it waits until the first of the following two events happens:

- receiving or forming a round- $(r-1)$ QC, i.e., $qc_{high}.r = r-1$
 - ▷ e.g., if it entered round r upon receiving qc_{high}
 - Then it sets $rc \leftarrow \perp$; or
- receiving $2f+1$ valid round- r new-round messages: $\left[\langle (r, r_j)_j, qc_{high,j} \rangle : j \in J \right]$, where $J \subset [n]$ is of size $2f+1$. In this case it sets:
$$rc \leftarrow \left[\langle r, r_j \rangle_j : j \in J \right]$$
 - ▷ proof of non-supermajority
 - ▷ In the production version [5], rc is instead the BGLS aggregate.

Then it multicasts a block $b = [id, qc_{high}, rc, r, txn]$.

Vote Upon receiving the first proposal $b = [id, qc, rc, r, txn]$ from $lead(r)$ while in round r execute **Lock**, and **Advance Round**, and then **Commit**, as instructed below. If $r > r_{vote}$ and $\left\{ \text{either (1): } r = qc.r + 1; \text{ or (2) } rc = \left[\langle r, r_j \rangle_j : j \in J \right] \text{ with } |J| = 2f+1 \text{ and } qc.r \geq \max\{r_j \mid j \in J\} \right\}$ then it votes for b by sending the individual signature $\langle id, r \rangle$ to $lead(r+1)$, and updates $r_{vote} \leftarrow r$.

Lock Upon receiving or forming a QC: qc , update $qc_{high} \leftarrow \max(qc_{high}, qc)$ (and thus $r_{high} \leftarrow \max(r_{high}, qc.r)$).

Commit (2-chain commit rule) Whenever there exists two adjacent certified blocks $b \leftarrow b'$ in the chain with consecutive round numbers, i.e., $b'.r = b.r + 1$, commit b and all its ancestors.

Advance Round ▷ Dotted box = the view-synchronizer model

- ▷ $\sqrt{r_{new} > r}$, update the current round number $r \leftarrow r_{new}$.
- ▷ implying that it stops voting for round- $\leq (r_{new}-1)$ proposals
- either upon receiving or forming a round- $(r_{new}-1)$ QC: qc
- or, upon a $(\text{NEWROUND}, r_{new})$ signal
- in this latter case, i.e., $(\text{NEWROUND}, r_{new})$, it then sends to $lead(r)$ a round- r new-round message:
$$\left(\langle r, r_{high} \rangle_i, qc_{high} \right)$$
 - ▷ where $\langle r, r_{high} \rangle_i$ is a (standalone) signature of \mathcal{P}_i on (r, r_{high}) ,
 - ▷ and where we recall $r_{high} := qc_{high}.r$.

Timer and Timeout (Implemented (NEWROUND) signals)

Figure 9: Differences with Jolteon: **new-round messages and RCs (highlighted)**. Differences with both Jolteon & Diem21: the black-box view-synchronizer (dotted-boxed), which replaces the explicit implementation of (NEWROUND) from timeout certificates in Diem21/Jolteon.

Diem21 + \mathcal{MtoA}

Instructions for each process \mathcal{P}_i . Same initialization as in Diem. We consider $\mathbf{A} := (\text{Kg}, \text{Sign}, \text{iVf}, \text{Ag}, \text{Vf})$ a one-time-tagged aggregate signature scheme, e.g., obtained from \mathcal{MtoA} .

Propose Upon entering round r , if \mathcal{P}_i is the leader $\text{lead}(r)$, then it waits until the first of the following two events happens:

- receiving or forming a round- $(r-1)$ QC: qc_{high} . Then $rc \leftarrow \perp$;
- or, receiving $2f+1$ round- r new-round messages: $\left[\left((r, r_j)_{j \in J}, qc_{high,J} \right) : j \in J \right]$. Then it aggregates the signatures w.r.t. the fixed tag r :

$$rc \leftarrow \left[r, (r_j)_{j \in J}, J, \mathbf{A}.\text{Ag}(r, \langle r_j \rangle_{j \in J}) \right]$$

\triangleright proof of non-supermajority

Then it multicasts a block $b = [id, qc_{high}, rc, r, txn]$.

Vote Upon receiving the first proposal $b = [id, qc, rc, r, txn]$ from $\text{lead}(r)$ while in round r , **Lock**, and **Advance Round**, and then **Commit**. If $r > r_{vote}$ and $\left\{ \begin{array}{l} \text{either (1): } r = qc.r + 1; \text{ or} \\ \text{(2) } r = rc.r + 1 \text{ and } rc = [r, (r_j)_{j \in J}, J, \Sigma] \text{ with } |J| = 2f + 1 \\ \text{and } \Sigma \text{ a valid } \mathbf{A}\text{-signature on the tagged messages } (r, r_j)_{j \in J} \\ \text{w.r.t. the public keys of } J \text{ and } qc.r \geq \max\{r_j \mid j \in J\} \end{array} \right\}$

then it votes for b by sending the individual signature $\langle id, r \rangle$ to $\text{lead}(r+1)$, and updates $r_{vote} \leftarrow r$.

Lock, Advance round, Commit Same as in Diem21.

New-round Upon receiving a signal ($\text{NEWROUND}, r_{new}$) from the view-synchronization mechanism, advance the current round number $r \leftarrow r_{new}$. Send to $\text{lead}(r)$ a new-round message:

$$\left(\{r, r_{high}\}_i, qc_{high} \right)$$

where $\{r, r_{high}\}_i \leftarrow \mathbf{A}.\text{Sign}(\text{sk}_i, r, r_{high})$ is an individual signature of player \mathcal{P}_i , and where we recall $r_{high} := qc_{high}.r$.

Figure 10: Differences with Figure 9 are highlighted.

APPENDIX D

FURTHER DETAILS ON RELATED WORKS

1) $\mathcal{AS}\text{-}3$ ([15]): We recall below the verification algorithm, tagged ($\mathcal{AS}\text{-}3.\text{Vf}$), of the fNIA called $\mathcal{AS}\text{-}3$ in [15]. $\mathcal{AS}\text{-}3$ is a verifier-unrestricted variant of the seminal BGLS [23]. $\mathcal{AS}\text{-}3$ is defined over a bilinear group $(e, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T)$ with public generators G_1, G_2 , and for a given hash-to-curve H . The X_i are the public keys.

$$(\mathcal{AS}\text{-}3.\text{Vf}) \quad \left((X_i, m_i)_{i \in [n]}, \Sigma \right) \longrightarrow \\ e(\Sigma, G_2) == \sum_{i \in [n]} e\left(H(X_i | m_i), X_i \right) \wedge (\Sigma \in \mathbb{G}_1)$$

2) *Pairing-based fNIM of Boldyreva* [20]: We recall below its verification formula, tagged (m-blsVf^*). Recall that this fNIM is secure only in the kosk model. The formula is the same as the one of dms (Sec. V), without the checks of PoPs on public keys (the function which we called kVf).

$$(\text{m-blsVf}^*) \quad \left((X_i)_{i \in [N]}, m, \Sigma \right) \longrightarrow \\ e(\Sigma, G_2) == e\left(H(m), \sum_{i \in [N]} X_i \right) \wedge \Sigma \in \mathbb{G}_1$$

3) *SMSKR* ([9]): The key published by each member of the group of potential signers, consists of a raw \mathbb{G}_2 element: $\text{pk} =$

$X \leftarrow x.G_2$. Once all keys of the group have been published: \mathcal{KG} , each signer re-randomizes its secret key:

$$(24) \quad x \leftarrow H(\mathcal{KG}|X).x$$

Hence, the individual signatures which it will generate with its re-randomized secret key x , will be valid w.r.t. the re-randomized public key:

$$(25) \quad X \leftarrow H(\mathcal{KG}|X).X \quad \forall X \in \mathcal{KG}$$

Likewise, combined signatures will be verified against the sum of the re-randomized keys of the subgroup of signers. Hence, the Verifier must compute the rerandomized public keys. Namely, the processing of the group setup consists of computing Eq. (25) for each $X \in \mathcal{KG}$.

4) *MSP-pop* [85, 24][22, §6], and *batch verification of group setup*: Each key $\text{pk}_i = (X_i, \pi_i)$ comes appended with a PoP equal to a BLS signature on X_i : $\Pi_i \leftarrow x_i.X_i$, where $X_i \leftarrow x_i.G_2$. Their verification cost is dominated by two pairings for each key:

$$(26) \quad \text{kVf}(\text{pk}_i) \rightarrow (e(\Pi_i, G_2) == e(H(X_i), X_i)) \wedge X_i \in \mathbb{G}_2 \wedge \Pi_i \in \mathbb{G}_1)$$

In our benchmarks (Table 6), we first used the $2\times$ speedup of [35] for batch verification of BLS signatures: the Verifier samples random numbers $(e_i)_{i \in [N]} \leftarrow_{\$} \mathbb{Z}_p^{[N]}$, then checks

$$(27) \quad e\left(\sum_{i \in [N]} e_i.\Pi_i, G_2 \right) == \sum_{i \in [N]} e\left(e_i.H(X_i), X_i \right)$$

Second, we sped-up the right-hand sum with the optimized implementation of products of pairings in gnark-crypto, inherited from [64].

5) *dms: Batch verification of group setup*: In our benchmarks (Table 6) we sped-up the verification of the PoPs of dms, i.e., $\text{kVf}(\text{pk}_i) \forall i \in [N]$, using the method of [14] for batch verification of Schnorr signatures. Namely: parse $(X_i, \pi_i) \leftarrow \text{pk}_i$ and $(R_i, z_i) \leftarrow \pi_i \forall i \in [N]$; $c_i \leftarrow H_{\text{pop}}(X_i, X_i, R_i)$; sample $(e_i)_{i \in [N]} \leftarrow_{\$} \mathbb{Z}_p^{[N]}$; output

$$(28) \quad X_i \in \mathbb{G}_2 \quad \forall i \in [N] \wedge \\ \left(\sum_{i \in [N]} e_i z_i \right) G_2 == \sum_{i \in [N]} e_i R_i + \sum_{i \in [N]} (e_i c_i).X_i$$

6) SIG_1 [27]: The recent (non-dynamic) fNIM called SIG_1 [27] has a processing of the group setup runtime which is one order of magnitude higher than the three previous fNIMs, i.e., SMSKR, *MSP-pop* and dms. Verification of N keys (each N -sized) requires $O(N^2)$ pairings and $O(N^2)$ group membership tests, instead of $N+1$ pairings in *MSP-pop* and $O(N)$ group membership tests in all three previous fNIMs. Then, computing the verification key of a group \mathcal{KG} of N keys takes N multi-additions, each with N terms.