

Practical relativistic bit commitment

T. Lunghi,¹ J. Kaniewski,^{2,3} F. Bussi eres,¹ R. Houlmann,¹ M. Tomamichel,^{2,4} S. Wehner,^{2,3} and H. Zbinden¹

¹*Group of Applied Physics, University of Geneva,
Chemin de Pinchat 22, CH-1211 Gen eve 4, Switzerland*

²*Centre for Quantum Technologies, National University of Singapore, 3 Science Drive 2, Singapore 117543*

³*QuTech, Delft University of Technology, Lorentzweg 1, 2628 CJ Delft, Netherlands*

⁴*School of Physics, The University of Sydney, Sydney 2006, Australia*

(Dated: June 26, 2015)

Bit commitment is a fundamental cryptographic primitive in which Alice wishes to commit a secret bit to Bob. Perfectly secure bit commitment between two mistrustful parties is impossible through asynchronous exchange of quantum information. Perfect security is however possible when Alice and Bob each split into several agents exchanging classical information at times and locations suitably chosen to satisfy specific relativistic constraints. In this Letter we first revisit a previously proposed scheme [1] that realizes bit commitment using only classical communication. We prove that the protocol is secure against quantum adversaries for a duration limited by the light-speed communication time between the locations of the agents. We then propose a novel multi-round scheme based on finite-field arithmetic that extends the commitment time beyond this limit, and we prove its security against classical attacks. Finally, we present an implementation of these protocols using dedicated hardware and we demonstrate a 2 ms-long bit commitment over a distance of 131 km. By positioning the agents on antipodal points on the surface of the Earth, the commitment time could possibly be extended to 212 ms.

Bit commitment is a fundamental primitive with several applications such as coin tossing [2], secure voting [3], contract signing or honesty-preserving auctions [4]. In a bit commitment protocol, Alice commits a secret bit to Bob which she can choose to reveal some time later. Security here means that if Alice is honest, then her bit is perfectly concealed from Bob until she decides to open the commitment and reveal her bit. Furthermore, if Bob is honest, then it should be impossible for Alice to change her mind once the commitment is made. That is, the only bit she can unveil is the one she originally committed herself to. Information-theoretically secure bit commitment in a setting where the two mistrustful parties exchange classical messages in an asynchronous fashion is impossible. An extensive amount of work was devoted to study asynchronous quantum bit commitment, for which perfect security was ultimately shown to be impossible [5–8]. Note however that arbitrarily long commitments are possible if one makes the assumption that the quantum memory of the dishonest party is bounded [9, 10] or noisy [11, 12].

Alternatively, bit commitment with split agents exchanging classical information was proposed as early as 1988 [13]. Security against classical attacks was proved under the condition that no communication was possible between some of the agents. This protocol was later simplified [1], and the new scheme called simplified-BGKW, *s*BGKW [13] was proven secure against classical and a restricted class of quantum attacks. The possibility of enforcing the no-communication condition using relativistic constraints on the timing of the classical communication was formulated in [14]. This later led to the proposal of relativistic protocols based on the exchange of quantum and classical information [15, 16], which were proved to be secure against quantum ad-

versaries [17, 18]. Such protocols were experimentally demonstrated recently [19, 20]. However, the commitment time achievable using these protocols is fundamentally bounded by half the time required to send light signals between the remote agents, i.e. at most ~ 21 ms if they are constrained to be on the surface of the Earth.

The possibility of extending the commitment to an arbitrary duration was proposed in 1999 [14]. It relies on positioning one agent of Alice \mathcal{A}_1 near an agent of Bob \mathcal{B}_1 at an agreed upon location, and similarly agents \mathcal{A}_2 and \mathcal{B}_2 at another location. Carefully timed classical communication between \mathcal{A}_i and \mathcal{B}_i allows Alice to commit to a bit that she later reveals at a time of her choosing. This requires several rounds of communication, and the amount of communication increases exponentially with the number of rounds making it impractical. This limitation was later mitigated, at least in principle, using a compression scheme that requires only a constant communication rate [21]. Security argument against classical adversaries presented in Ref. [21] is of asymptotic nature and, therefore, not sufficient for implementation purposes.

In this Letter, we first revisit the *s*BGKW bit commitment protocol [1] that uses classical communication only. We show that successful cheating is equivalent to winning a non-local game analyzed in Ref. [22], thereby proving the security of this protocol against quantum adversaries. To the best of our knowledge, this is the first entirely classical protocol to be proven secure against arbitrary quantum adversaries. To extend the duration of the commitment beyond the communication time between the locations of the agents (which constitutes the relativistic constraint in the *s*BGKW scheme), we introduce a novel multi-round scheme based on finite-field arithmetic and we prove its security against classical adversaries.

Our scheme is simple and efficient and the security argument leads to a natural, algebraic problem for which we prove explicit and quantitative bounds (see Proposition B.2 in the Supplemental Material (SM)). Finally, we present practical implementations of both the sBGKW scheme and the multi-round variant, and show how this could be used to realize commitments of duration reaching up to ~ 212 milliseconds.

Security definition We take $n \in \mathbb{N}$ to be the security parameter and we interpret n -bit strings as elements of the finite field \mathbb{F}_{2^n} (for compactness we write 0 to denote $0^n = 00\dots 0$). We denote addition by “ \oplus ” (in this case it is just the bitwise XOR) and multiplication by “ \cdot ”. Moreover, if d is a bit and b is an n -bit string then we define

$$d \cdot b = \begin{cases} 0 & \text{if } d = 0, \\ b & \text{if } d = 1. \end{cases}$$

All secret strings used in the protocol are chosen uniformly at random from $\{0, 1\}^n$.

Let Alice (who makes the commitment) and Bob (who receives the commitment) have agents at two distinct locations (\mathcal{A}_1 and \mathcal{B}_1 at Location 1; \mathcal{A}_2 and \mathcal{B}_2 at Location 2) and let $d \in \{0, 1\}$ be the bit that honest Alice wants to commit to. The protocol consists of multiple rounds which alternate between the two locations and the timing is chosen such that every two consecutive rounds are space-like separated. Hence, no message sent during a certain round from one location can reach the other location in time for the next round.

Security for honest Alice is quantified by Bob’s ability to guess her commitment *immediately before* the open phase (assuming he might deviate arbitrarily from the honest protocol). All the protocols considered in this paper are *perfectly hiding*, which means that Bob remains completely ignorant about Alice’s commitment (his guessing probability equals $\frac{1}{2}$).

Security for honest Bob is quantified through a scenario in which Alice performs an arbitrary action in the commit phase and is *immediately after* challenged to open one of the bits. Given a particular strategy adopted by Alice in the commit phase we define p_d to be the optimal probability of successfully unveiling d . The protocol is ε -binding if

$$p_0 + p_1 \leq 1 + \varepsilon$$

for all strategies of dishonest Alice in the commit phase. Note that this is a weak, non-composable definition of security. In Appendix C we discuss how to formalize these definitions in the relativistic setting. (For a general overview see Ref. [17].)

Security of the sBGKW scheme We now present the scheme proposed in Ref. [1] and we prove its security against quantum adversaries. Before the protocol begins \mathcal{A}_1 and \mathcal{A}_2 must share a secret n -bit string a . Note that \mathcal{B}_1 also needs a secret string b but it can be generated before or during the protocol. The protocol consists of two rounds:

1. (commit) \mathcal{B}_1 sends b to \mathcal{A}_1 . \mathcal{A}_1 returns $(d \cdot b) \oplus a$ to \mathcal{B}_1 .
2. (open) \mathcal{A}_1 unveils the committed bit d to \mathcal{B}_1 while \mathcal{A}_2 sends a to \mathcal{B}_2 .

To check whether the commitment should be accepted \mathcal{B}_1 and \mathcal{B}_2 need to communicate (e.g. through an authenticated channel) and verify that the string returned by \mathcal{A}_1 in the commit phase equals $(d \cdot b) \oplus a$.

Security for honest Alice comes from the fact that the only message that Bob receives in the commit phase is a uniformly random string.

Security for honest Bob in the classical case is fairly intuitive: in order for \mathcal{A}_2 to be able to unveil both commitments she would need to know both a and $a \oplus b$, hence, she would know b . However, since b is chosen uniformly at random by Bob this must be difficult. This argument can be made rigorous [1] to show that the protocol is ε -binding for $\varepsilon = 2^{-n}$ (and this is actually tight: the trivial strategy of always outputting 0 gives $p_0 = 1$ and $p_1 = 2^{-n}$). Unfortunately, this reasoning does not work against quantum adversaries since \mathcal{A}_2 could have two distinct measurements that reveal a and $a \oplus b$, respectively, but since they could be incompatible this would not have direct implications on her ability to guess b .

To find an explicit bound on $p_0 + p_1$ we formulate cheating as a non-local game in which \mathcal{A}_1 receives b , \mathcal{A}_2 receives d (the bit she is required to unveil) and the XOR of their outputs is supposed to equal $d \cdot b$. Winning such a game with probability p_{win} corresponds to a cheating strategy that achieves $p_0 + p_1 = 2p_{\text{win}}$. More concisely, the rules of the non-local game are [1]:

1. \mathcal{A}_1 receives $b \in \{0, 1\}^n$, \mathcal{A}_2 receives $d \in \{0, 1\}$ (both chosen uniformly at random).
2. \mathcal{A}_1 outputs $a_1 \in \{0, 1\}^n$, \mathcal{A}_2 outputs $a_2 \in \{0, 1\}^n$ and they win iff $a_1 \oplus a_2 = d \cdot b$.

This game has been considered in Ref. [22] under the name CHSH_n and it has been shown that

$$p_{\text{win}}(n) \leq \frac{1}{2} + \frac{1}{\sqrt{2^{n+1}}},$$

which is sufficient for our purposes as it implies that

$$p_0 + p_1 \leq 1 + \sqrt{2} \cdot 2^{-n/2}$$

for all strategies of dishonest Alice. Therefore, the protocol is ε -binding with $\varepsilon = 2^{(1-n)/2}$ decaying exponentially in n (but note that the decay rate is half of the decay rate against classical adversaries).

The two-round protocol is mapped onto a non-local game precisely because of the assumption of no communication. More specifically, we require that \mathcal{A}_1 outputs the answer outside of the future of \mathcal{A}_2 receiving the input and vice versa.

A new multi-round protocol To extend the commitment time we propose a multi-round protocol and prove its security against classical adversaries. In principle, the commitment time can be made arbitrarily long. However, security depends on the number of rounds of the protocol, which is proportional to the length of the commitment. Therefore, the longer the commitment, the more resources (randomness and communication bandwidth) are required to achieve a given level of security.

Suppose that Alice and Bob want to execute the protocol with $m + 1$ rounds and we use k as a label for the round under consideration. Then \mathcal{A}_1 and \mathcal{A}_2 must share m secret strings denoted by $\{a_k\}_{k=1}^m$. Similarly, Bob's agents need one secret string for every round denoted by $\{b_k\}_{k=1}^m$ but, again, these can be generated locally during the protocol. All the rounds before the open phase ($1 \leq k \leq m$) have the same communication pattern: first \mathcal{B}_i sends an n -bit string to \mathcal{A}_i and then she replies with another n -bit string. In the last round \mathcal{A}_i sends \mathcal{B}_i a bit (her commitment) and an n -bit string (proof of her commitment). We will denote the n -bit string announced by Bob (Alice) in the k round by x_k (y_k) regardless of whether he/she is honest or not. The protocol is:

1. (commit, $k = 1$) \mathcal{B}_1 sends $x_1 = b_1$ to \mathcal{A}_1 . \mathcal{A}_1 returns $y_1 = d \cdot x_1 \oplus a_1$.
2. (sustain, $2 \leq k \leq m$) \mathcal{B}_i sends $x_k = b_k$ to \mathcal{A}_i . \mathcal{A}_i returns $y_k = (x_k * a_{k-1}) \oplus a_k$.
3. (open, $k = m + 1$) \mathcal{A}_i sends d and $y_{m+1} = a_m$ to \mathcal{B}_i .

To check whether the commitment should be accepted \mathcal{B}_1 and \mathcal{B}_2 communicate and verify the following relation:

$$y_{m+1} = y_m \oplus b_m * y_{m-1} \oplus b_m * b_{m-1} * y_{m-2} \oplus \dots \oplus b_m * b_{m-1} * \dots * b_2 * y_1 \oplus d \cdot b_m * b_{m-1} * \dots * b_1.$$

Security for honest Alice is a direct consequence of the fact that every message she announces is masked by a fresh secret n -bit string, which implies that the transcripts corresponding to $d = 0$ and $d = 1$ are statistically indistinguishable (see Proposition C.1 in the SM).

Proving security for honest Bob is a more challenging task, because we require security immediately after round $k = 1$. We first state the main result and then outline the idea behind the proof (for details refer to Sections B.2 and C.2 in the SM). The multi-round protocol with $m + 1$ rounds is ε -binding for $\varepsilon = c_m$ defined as

$$c_m = \begin{cases} 2^{-n} & \text{for } m = 1, \\ \frac{1}{2^{n+1}} + \sqrt{c_{m-1}} & \text{for } m \geq 2. \end{cases} \quad (1)$$

The security argument is conceptually simple: in the classical scenario the *sequential* cheating game in the multi-round protocol is *equivalent* to a game in which multiple players act *in parallel* which allows us to disregard the causal structure of the protocol. We show that cheating in a protocol with $m + 1$ rounds is at least

as difficult as winning the following m -player game. Let X_1, X_2, \dots, X_m be independent random variables drawn uniformly from the set of n -bit strings $\{0, 1\}^n$ and the k player receives all the variables except for X_k and outputs an n -bit string. The game is won if the XOR of the outputs equals $X_1 * X_2 * \dots * X_m$. The bounds we obtain decay exponentially in $n/2^m$. This means that they become significantly weaker as the number of players increases, which ultimately limits the maximum number of rounds that can be implemented in practice. The tightness of these bounds is an interesting open problem and it is briefly discussed in Appendix B. Note that no explicit cheating strategy is known, whose winning probability would approach our security bounds.

Implementation We implemented the two-round and the multi-round protocols described above. Each party has agents at two distinct locations: one at the Group of Applied Physics of the University of Geneva and one at the Institute of Applied Physics of the University of Berne. The straight-line distance between the two locations is $s = 131$ km, corresponding to a time separation of $437 \mu\text{s}$. The hardware installed in Geneva is conceptually represented in FIG. 1(a) and identical to the one in Berne. Each of the classical agents is a standalone computer equipped with a field-programmable gate array (FPGA) programmed to execute the necessary steps of the protocol. Each FPGA is synchronized to the Coordinated Universal Time (UTC) via a Global Positioning System clock (GPS clock), which consists of a GPS receiver and a Oven-Controlled Quartz-Crystal Oscillator (OCXO) generating a 10 MHz sinusoidal waveform. Through its GPS connection, the receiver outputs one electronic pulse per second (PPS), which is used to discipline the OCXO. The receiver is locked to the GPS signal with a time accuracy better than 150 ns. The 10 MHz signal generated by the OCXO is fed into the FPGA board and it is used to generate a 125 MHz signal using a phase-locked loop. This 125 MHz signal then serves as the time basis for the computations performed on the FPGA. The FPGA also receives the PPS signal to monitor the synchronization with the GPS clock. In particular, the number of cycles between two successive PPS signals is confirmed to be 125×10^6 plus or minus one, where each cycle corresponds to 8 ns. Therefore, the FPGA tolerates fluctuations up to 24 ns on the arrival time of the PPS synchronization signal. The GPS clock also provides the FPGA with a universal time stamp of every PPS signal, allowing Alice and Bob to locate their actions in time.

Before either the two-round or the multi-round protocol starts, \mathcal{A}_1 and \mathcal{A}_2 (and similarly \mathcal{B}_1 and \mathcal{B}_2) share an appropriate number of random n -bit strings. At time t_1 , which was agreed upon by both parties, \mathcal{B}_1 sends the random string x_1 through the optical link. For a string of 512 bits communicated through the 2.5 Gbps optical link, this requires 205 ns. \mathcal{A}_1 's FPGA then computes the string y_1 and sends it to \mathcal{B}_1 ; see FIG. 1(b). The relativistic constraint requires space-like separation between

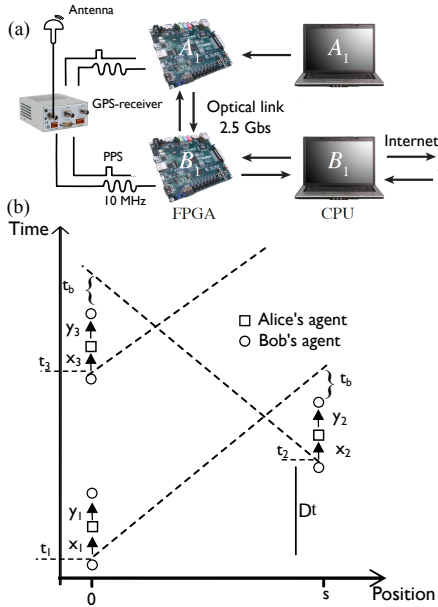


FIG. 1. (a) Experimental setup. (b) Space-time diagram of the experimental setup.

every two consecutive rounds, which means that the entire second round must be outside of the future light cone of the first bit of x_1 leaving the FPGA of \mathcal{B}_1 . The commitment begins when the last bit of y_1 is recorded by the FPGA of \mathcal{B}_1 . With $n = 512$ bits, the security parameter of the two-round protocol is $\varepsilon \approx 10^{-77}$.

In the two-round protocol, \mathcal{A}_2 unveils the commitment in the second round, at time $t_2 = t_1 + \Delta t$. She does so by sending the string a_1 to \mathcal{B}_2 , along with the committed bit d . \mathcal{B}_2 checks that the last bit of a_1 is received outside the future light cone of the beginning of the protocol. If this is the case, \mathcal{B}_2 communicates a_1 and d to \mathcal{B}_1 through an authenticated channel. Finally, \mathcal{B}_1 verifies that $y_1 \oplus a_1 = d \cdot x_1$ and accepts the commitment. If the relativistic constraint is not respected, or if \mathcal{B}_1 's verification fails, the protocol aborts.

In the multi-round protocol, \mathcal{A}_1 and \mathcal{A}_2 successively sustain the commitment until the last round. All rounds (except the first and last rounds) proceed as follows. Let us consider the k^{th} round, with k even (odd rounds are similar). Between rounds k and $k - 2$, the string x_k is loaded in the memory of \mathcal{B}_2 's FPGA, and strings a_{k-1} and a_k are loaded in \mathcal{A}_2 's FPGA. At time $t_k = t_1 + (k - 1)\Delta t$, \mathcal{B}_2 communicates x_k through the optical link. Then \mathcal{A}_2 sustains the commitment by computing y_k with the FPGA and sending it to \mathcal{B}_2 . The time between the communication of x_k and the reception of y_k is $6.1 \mu\text{s}$. \mathcal{B}_2 checks that the reception of y_k is outside the future light cone of the beginning of the communication between \mathcal{B}_1 to \mathcal{A}_1 that happened in round $k - 1$. We used $\Delta t = 400 \mu\text{s}$ (see Fig. 1), which is $37 \mu\text{s}$ shorter than the light-speed separation between the Berne and Geneva locations. Considering the $6.1 \mu\text{s}$, the absolute inaccuracies of the GPS clock ($\leq 150 \text{ ns}$), and the tolerance in the

fluctuations of the synchronization signals ($\leq 24 \text{ ns}$) the round is completed $\approx 30.7 \mu\text{s}$ before the relativistic constraint expires. In the final $(m + 1)^{\text{th}}$ round, \mathcal{A}_1 (or \mathcal{A}_2) opens the commitment at time t_{m+1} by sending the string a_{m+1} along with the committed bit d . To verify the commitment, \mathcal{B}_2 sends to \mathcal{B}_1 all the strings communicated by \mathcal{A}_2 through an authenticated channel. \mathcal{B}_1 then checks if the commitment should be accepted as outlined above. Authentication is based on an information-theoretic secure message-authenticator code which consists of a combination of polynomial hashing, and a strongly-universal family of hash functions [23].

In the multi-round scheme, we aimed to maximize the number of rounds with a reasonable value for the security parameter ε . The limit of $n = 512$ bits and $m + 1 = 6$ rounds was ultimately set by the performance that we could achieve with the FPGA at our disposal. This yields a security parameter of $\varepsilon \approx 2.3 \times 10^{-10}$. Using these parameters, we realized a commitment of 2 ms duration, which extends beyond the $437 \mu\text{s}$ limit of the two-round protocol. Because synchronizing rounds over longer durations is a simple task for our hardware, it is straightforward to achieve significantly longer commitment times using more distant agents. For example, 150 ms could be easily achieved using Geneva and Singapore as the locations (these locations were used in our previous demonstration of quantum-relativistic bit commitment [19]), while 212 ms could be achieved using antipodal locations on the Earth.

Summary We have shown that classical relativistic protocols allow us to implement information-theoretically secure commitment schemes in a straightforward fashion.

The commitment scheme we implemented belongs to the class of timed commitments, i.e. commitments that expire after a certain period of time. Even though they cannot be used to implement primitives whose security is required to hold forever (e.g. oblivious transfer), they are known to have other important applications, e.g. contract signing, honesty-preserving auctions or secure voting [3, 4] (see also Appendix A).

For the sBGKW scheme we obtain an explicit, quantitative security bound by making a connection to a non-local game analyzed previously. We also propose a multi-round scheme which is secure against classical adversaries. We note that the number of rounds that we implemented here could have been higher using better optimized hardware. However, the scaling of the security bound with the number of rounds (1) prohibits a much larger number of rounds. An important problem is therefore to find a multi-round protocol whose security exhibits better scaling with the number of rounds, or, ideally, no dependence at all. This would allow us to obtain longer (or maybe even arbitrarily long) commitments while only using simple, commercially available digital devices.

Acknowledgments: We thank Mohammad Bavarian, Gilles Brassard, Iordanis Kerenidis, Raghav Kulkarni, Troy Lee, Laura Manćinska, Miklos Santha and Sarvagya Upadhyay for useful discussions. JK especially thanks

Igor Shparlinski for sharing his ideas about Proposition B.2 of the SM and subsequent discussions. We thank André Stefanov and Daniel Weber for helping to install the setup in Berne. JK, MT and SW are funded by the Ministry of Education (MOE) and National Research Foun-

dation Singapore, as well as MOE Tier 3 Grant "Random numbers from quantum processes" (MOE2012-T3-1-009). Financial support is provided by the Swiss NCCR QSIT.

Author contributions: TL is the first experimental author and JK is the first theoretical author.

-
- [1] C. Crépeau, L. Salvail, J.-R. Simard, and A. Tapp, Proc. 17th ASIACRYPT, Lecture Notes on Computer Science (2011).
- [2] M. Blum, in *Advances in Cryptology: A Report on CRYPTO'81, Santa Barbara, California, USA* (1981) pp. 11–15.
- [3] A. Broadbent and A. Tapp, in *Proceedings of the IAVoSS Workshop On Trustworthy Elections (WOTE 2008)* (<http://eprint.iacr.org/2008/266>, 2008).
- [4] D. Boneh and M. Naor, in *Proc. 20th CRYPTO*, Vol. 1880 (2000) p. 236254.
- [5] D. Mayers, Phys. Rev. Lett. **78**, 3414 (1997).
- [6] H.-K. Lo and H. F. Chau, Phys. Rev. Lett. **78**, 3410 (1997).
- [7] G. M. D'Ariano, D. Kretschmann, D. Schlingemann, and R. F. Werner, Phys. Rev. A **76**, 032328 (2007).
- [8] S. Winkler, M. Tomamichel, S. Hengli, and R. Renner, Phys. Rev. Lett. **107**, 090502 (2011).
- [9] I. B. Damgård, S. Fehr, L. Salvail, and C. Schaffner, Proc. 46th IEEE FOCS, 449 (2005).
- [10] I. B. Damgård, S. Fehr, R. Renner, L. Salvail, and C. Schaffner, Proc. 27th CRYPTO, 360 (2007).
- [11] S. Wehner, C. Schaffner, and B. M. Terhal, Phys. Rev. Lett. **100**, 220502 (2008).
- [12] R. König, S. Wehner, and J. Wullschleger, IEEE Trans. Inf. Theory **58**, 1962 (2012).
- [13] M. Ben-Or, S. Goldwasser, J. Kilian, and A. Wigderson, in *Proc. ACM STOC* (ACM Press, New York, New York, USA, 1988) pp. 113–131.
- [14] A. Kent, Phys. Rev. Lett. **83**, 1447 (1999).
- [15] A. Kent, New Journal of Physics **13**, 113015 (2011).
- [16] A. Kent, Phys. Rev. Lett. **109**, 130501 (2012).
- [17] J. Kaniewski, M. Tomamichel, E. Hänggi, and S. Wehner, Information Theory, IEEE Transactions on **59**, 4687 (2013).
- [18] S. Croke and A. Kent, Phys. Rev. A **86**, 052309 (2012).
- [19] T. Lunghi, J. Kaniewski, F. Bussi eres, R. Houlmann, M. Tomamichel, A. Kent, N. Gisin, S. Wehner, and H. Zbinden, Phys. Rev. Lett. **111**, 180504 (2013).
- [20] Y. Liu, Y. Cao, M. Curty, S.-K. Liao, J. Wang, K. Cui, Y.-H. Li, Z.-H. Lin, Q.-C. Sun, D.-D. Li, H.-F. Zhang, Y. Zhao, T.-Y. Chen, C.-Z. Peng, Q. Zhang, A. Cabello, and J.-W. Pan, Phys. Rev. Lett. **112**, 010504 (2014).
- [21] A. Kent, Journal of Cryptology **18**, 313 (2005).
- [22] J. Sikora, A. Chailloux, and I. Kerenidis, Phys. Rev. A **89**, 022334 (2014).
- [23] J. Carter and M. N. Wegman, Journal of Computer and System Sciences **18**, 143 (1979).
- [24] H. Buhrman and S. Massar, Phys. Rev. A **72**, 052103 (2005).
- [25] M. Bavarian and P. W. Shor, in *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science, ITCS '15* (ACM, New York, NY, USA, 2015) pp. 123–132.
- [26] J. Ford and A. Gal, Comput. Complex. **22**, 595 (2013).

Appendix A: Preliminaries

1. How useful is a relativistic bit commitment protocol?

The commitment scheme we implement belongs to the class of timed commitments, i.e. commitments that are only valid for a period of time but then ultimately *expire*. Such commitments cannot be used in reductions implementing primitives whose security is required to hold forever (e.g. oblivious transfer or secure function evaluation) but they are known to have other applications. For example, Boneh and Naor [4] study commitments which after some fixed Δt automatically open, i.e. the committed value is revealed. They show that such commitments can be used for contract signing or honesty-preserving auctions. In our case after Δt the commitment simply vanishes, i.e. Bob should not accept any opening (because Alice could unveil either bit with unit probability) and the originally committed value (if there was one) remains secret. Therefore, it gives more power to the committer by giving her the freedom not to open the commitment and, hence, protect her privacy. Generally speaking, such temporary secrecy is sufficient if the goal is not to preserve secrecy forever but to force parties to act simultaneously (in the sense that their respective actions should not depend on each other) even if the communication model is sequential. Our commitment might be particularly useful for multi-party protocols which are robust against a certain fraction of dishonest parties (then we would simply call dishonest any party that refuses to open the commitment). A prime application of this type would be the task of secure voting as presented in Ref. [3].

2. Notation

Let $[n] = \{1, 2, \dots, n\}$. Generally, we use uppercase letters for random variables and lowercase letters to denote particular values. For $j, k \in \mathcal{S}$ we use $\sum_{j \neq k}$ as shorthand notation for $\sum_{j \in \mathcal{S}} \sum_{k \in \mathcal{S} \setminus \{j\}}$.

3. The Cauchy-Schwarz inequality for probabilities

Let X be a random variable distributed uniformly over $[n]$ and let $\{E_j\}_{j \in [m]}$ be a family of events defined on X .

Lemma A.1. *Let p be the average probability of the family of events*

$$p := \frac{1}{m} \sum_{j \in [m]} \Pr[E_j]$$

and c be the cumulative size of the pairwise intersections

$$c := \sum_{j \neq k} \Pr[E_j \wedge E_k].$$

Then the following inequality holds

$$p \leq \frac{1 + \sqrt{1 + 4c}}{2m}.$$

Proof. Each event can be represented by an n -dimensional, real vector whose entries are labelled by the possible values that X can take. If a particular value of X belongs to the event, we set the corresponding component to $1/\sqrt{n}$ and if it does not we set it to 0

$$[s_j]_x = \begin{cases} \frac{1}{\sqrt{n}} & \text{if } x \in E_j, \\ 0 & \text{otherwise.} \end{cases}$$

Moreover, let n be the normalised, uniform vector: $[n]_x = 1/\sqrt{n}$ for all $x \in [n]$. It is straightforward to check that with these definitions we have

$$\Pr[E_j] = \langle s_j, n \rangle = \langle s_j, s_j \rangle \quad \text{and} \quad \Pr[E_j \wedge E_k] = \langle s_j, s_k \rangle.$$

Clearly, we have $\langle s_j, s_k \rangle \geq 0$. Due to linearity of the inner product we have

$$p = \frac{1}{m} \sum_{j \in [m]} \Pr[E_j] = \frac{1}{m} \sum_j \langle s_j, n \rangle = \frac{1}{m} \langle \sum_j s_j, n \rangle,$$

which can be upper bounded using the Cauchy-Schwarz inequality. Since $\langle n, n \rangle = 1$ we have

$$\langle \sum_j s_j, n \rangle^2 \leq \sum_{jk} \langle s_j, s_k \rangle = \sum_j \langle s_j, s_j \rangle + \sum_{j \neq k} \langle s_j, s_k \rangle = mp + c,$$

which gives the following quadratic constraint

$$p^2 \leq \frac{p}{m} + \frac{c}{m^2}.$$

Solving for p gives the desired bound. □

Appendix B: Finite-field multiplication in the ‘‘Number on the Forehead’’ model

We introduce a family of multiplayer games which are a natural generalisation of the two-player family introduced in Ref. [24] and generalised in Ref. [25]. Since these games rely on finite-field arithmetic we first state some basic properties of finite fields, then we define the game and show that finding the optimal winning probability corresponds to a natural algebraic problem concerning multivariate polynomials over finite fields. Finally, we prove upper bounds on the optimal winning probability and discuss their tightness.




FIG. 2. In the “Number on the Forehead” model there are m inputs X_1, X_2, \dots, X_m and \mathcal{P}_k (the k player out of m) receives all the inputs except for X_k . We denote the output of \mathcal{P}_k by Y_k . Vertical lines remind us that no communication between the players is allowed.

1. Finite-field arithmetic

Let \mathbb{F}_q denote the finite field of order $q = p^k$ (p is a prime and k is an integer) and let 0 denote the zero element of \mathbb{F}_q . Operations over finite-field satisfy the following properties

1. Multiplication by zero gives zero
 $x \cdot 0 = 0 \quad \forall x \in \mathbb{F}_q$
2. Multiplication is distributive over addition
 $x(y + z) = (xy) + (xz) \quad \forall x, y, z \in \mathbb{F}_q$

2. Definition of the game

Consider a one-round game with m -players denoted by $\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_m$. With every player we associate an *input* and an *output*, e.g. for \mathcal{P}_k these are denoted by X_k and Y_k , respectively. Let each of X_1, X_2, \dots, X_m be drawn independently, uniformly at random from \mathbb{F}_q . In the “Number on the Forehead” model \mathcal{P}_k receives *all the inputs except for the k one* (denoted by $X_{[m] \setminus \{k\}}$) as shown in FIG. 2. Each player is required to output an element of \mathbb{F}_q (denoted by Y_k) and the game is won if

$$\prod_{k=1}^m X_k = \sum_{k=1}^m Y_k. \quad (\text{B1})$$

In the classical setting the optimal winning probability can be achieved when each player adopts a deterministic strategy, i.e. a function $f : \mathbb{F}_q^{(m-1)} \rightarrow \mathbb{F}_q$. If \mathcal{P}_k employs a strategy represented by f_k , i.e. he outputs $Y_k = f_k(X_{[m] \setminus \{k\}})$, then the winning probability equals

$$\omega_m(f_1, f_2, \dots, f_m) := \Pr \left[\prod_{k=1}^m X_k = \sum_{k=1}^m f_k(X_{[m] \setminus \{k\}}) \right]$$

and we define ω_m to be the optimal winning probability

$$\omega_m := \max_{f_1, f_2, \dots, f_m} \omega_m(f_1, f_2, \dots, f_m), \quad (\text{B2})$$

where the maximisation is taken over all functions from $\mathbb{F}_q^{(m-1)}$ to \mathbb{F}_q .

3. Characterisation through multivariate polynomials over a finite field

First, note that since the probability distribution of inputs is uniform then the winning probability is proportional to the number of inputs (x_1, x_2, \dots, x_m) on which the condition (B1) is satisfied

$$\prod_{k=1}^m x_k = \sum_{k=1}^m f_k(x_{[m] \setminus \{k\}}).$$

Alternatively, the winning probability can be deduced by counting the number of zeroes of the following function

$$P(x_1, x_2, \dots, x_m) = \prod_{k=1}^m x_k - \sum_{k=1}^m f_k(x_{[m] \setminus \{k\}}).$$

By the Lagrange interpolation method every function from $\mathbb{F}_q^n \rightarrow \mathbb{F}_q$ (for arbitrary $n \in \mathbb{N}$) can be written as a polynomial. Therefore, the question concerns the number of zeroes of the polynomial P . Different strategies employed by the players give rise to different polynomials and we need to characterise what polynomials are “reachable” in this model. The output of \mathcal{P}_k is an arbitrary polynomial of $x_{[m] \setminus \{k\}}$, hence, it only contains terms that depend on *at most* $m - 1$ variables. This means that the part of P that depends on *all* m variables comes solely from the first term and equals $\prod_{k=1}^m x_k$. Therefore, finding the optimal winning probability of the game is equivalent to finding the polynomial with the largest number of zeroes, whose only term that depends on all m variables equals $\prod_{k=1}^m x_k$. This reduces the problem of finding the optimal strategy to a purely algebraic problem about properties of polynomials over finite fields.

4. A recursive upper bound on the optimal winning probability

Here, we show how to find explicit upper bounds on ω_m through an induction argument. First, note that for $m = 1$ there is only one term on the right-hand side of Eq. (B1) and since this term takes no arguments it is actually a constant. Since X_1 is uniform we have

$$\omega_1 := \max_{c \in \mathbb{F}_q} \Pr[X_1 = c] = \frac{1}{q}.$$

Now, we show how to prove an upper bound on ω_m in terms of ω_{m-1} . For a fixed strategy $\{f_1, f_2, \dots, f_m\}$ the winning probability can be written as

$$\begin{aligned} \omega_m(f_1, f_2, \dots, f_m) &= \Pr \left[X_1 X_2 \dots X_m = \sum_{k=1}^m f_k(X_{[m] \setminus \{k\}}) \right] \\ &= \sum_{y \in \mathbb{F}_q} \Pr[X_m = y] \cdot \Pr \left[X_1 X_2 \dots X_m = \sum_{k=1}^m f_k(X_{[m] \setminus \{k\}}) \mid X_m = y \right] \\ &= q^{-1} \sum_y \Pr \left[X_1 X_2 \dots X_m = \sum_{k=1}^m f_k(X_{[m] \setminus \{k\}}) \mid X_m = y \right]. \end{aligned}$$

Conditioning on a particular value of X_m leads to events that only depend on X_1, X_2, \dots, X_{m-1} . In particular, we can define for $X_m = y$ the event F_y

$$F_y \iff X_1 X_2 \dots X_{m-1} y = \sum_{k=1}^{m-1} f_k(X_{[m-1] \setminus \{k\}}, y) + f_m(X_{[m-1]}),$$

which satisfies

$$\Pr[F_y] = \Pr \left[X_1 X_2 \dots X_m = \sum_{k=1}^m f_k(X_{[m] \setminus \{k\}}) \mid X_m = y \right]. \quad (\text{B3})$$

We can use Lemma A.1 to find a bound on $\omega_m(f_1, f_2, \dots, f_m) = q^{-1} \sum_y \Pr[F_y]$ as long as we are given bounds on $\Pr[F_y \wedge F_z]$ for $y \neq z$.

Proposition B.1. *For $y \neq z$ we have $\Pr[F_y \wedge F_z] \leq \omega_{m-1}$.*

Proof. Eq. (B3) defines F_y through a certain equation in the finite field. If the equations corresponding to F_y and F_z are satisfied simultaneously then clearly any linear combination of these equations is also satisfied. More specifically, we define a new event

$$G_{yz} \iff X_1 X_2 \dots X_{m-1} (y - z) = \sum_{k=1}^{m-1} f_k(X_{[m-1] \setminus \{k\}}, y) - f_k(X_{[m-1] \setminus \{k\}}, z) \quad (\text{B4})$$

and since $F_y \wedge F_z \implies G_{yz}$ we are guaranteed that $\Pr[F_y \wedge F_z] \leq \Pr[G_{yz}]$.

To find an upper bound on $\Pr[G_{yz}]$ we give the players more power by allowing a more general expression on the right-hand side. In Eq. (B4) the k term is a function of $X_{[m-1]\setminus\{k\}}$, y and z , so let us replace it by an arbitrary function of these variables

$$f_k(X_{[m-1]\setminus\{k\}}, y) - f_k(X_{[m-1]\setminus\{k\}}, z) \rightarrow g_k(X_{[m-1]\setminus\{k\}}, y, z).$$

Under this relaxation, we arrive at the following equality

$$X_1 X_2 \dots X_{m-1} (y - z) = \sum_{k=1}^{m-1} g_k(X_{[m-1]\setminus\{k\}}, y, z).$$

Clearly, $(y - z)$ is a constant (non-zero) multiplicative factor known to each player. Dividing the equation through by $(y - z)$ leads to the same game as considered before but the number of players has decreased by one: there are only $m - 1$ players now. Therefore,

$$\Pr[F_y \wedge F_z] \leq \Pr[G_{yz}] \leq \omega_{m-1}.$$

□

Now, we can state and prove our main technical result.

Proposition B.2. *The optimal winning probability of the game defined in (B2) satisfies the following recursive relation*

$$\omega_m \leq \frac{1 + \sqrt{1 + 4q(q-1)\omega_{m-1}}}{2q}. \quad (\text{B5})$$

Proof. The statement follows directly from combining Lemma A.1 with Proposition B.1. □

Since we know that $\omega_1 = q^{-1}$, we can obtain a bound on ω_m by recursive evaluation of Eq. (B5). More precisely, we get $\omega_m \leq c_m$ for

$$c_m = \begin{cases} q^{-1} & \text{for } m = 1, \\ \frac{1 + \sqrt{1 + 4q(q-1)c_{m-1}}}{2q} & \text{for } m \geq 2. \end{cases} \quad (\text{B6})$$

Note that this bound is always non trivial, i.e. $c_m < 1$ for all values of q and m . To obtain a slightly weaker but simpler form presented as Eq. (1) in main text we note that $1 - 4qc_{m-1} \leq 0$ and set $q = 2^n$.

5. Discussion

Having proved an explicit upper bound on ω_m we would like to investigate its tightness. It can be shown that in the regime interesting from the cryptographic point of view ($q \gg 1$) the leading behaviour of c_m is

$$c_m \propto q^{-2^{-m}}. \quad (\text{B7})$$

In other words, c_m decays exponentially in q but the value of the exponent depends on the number of players: every time we add a player we lose half of the decay exponent. This might seem unexpectedly weak but it has recently been shown (Theorem 6.5 in Ref. [25]) that

$$\omega_2 = \Omega(q^{-\frac{2}{3}}).$$

In fact, for $q = p^k$ where k is even it can be improved (Theorem 1.3 in Ref. [25]) to give

$$\omega_2 = \Omega(q^{-\frac{1}{2}}).$$

This shows that for $m = 2$ the asymptotic decay of $q^{-1/2}$ is the best we can hope for (at least for an upper bound that holds for both odd and even values of q). Moreover, as far as we know, the best explicit upper bound on ω_2 is the quantum upper bound (Theorem 1.2 in Ref. [25])

$$c_2^{\text{qm}} = \frac{1}{q} + \frac{q-1}{q} \frac{1}{\sqrt{q}}.$$

On the other hand, evaluating c_2 according to Eq. (B6) gives

$$c_2 = \frac{\sqrt{1 - \frac{3}{4q}}}{\sqrt{q}} + \frac{1}{2q}.$$

By noting that

$$\begin{aligned} \sqrt{1 - \frac{3}{4q}} &\leq 1 - \frac{3}{8q} \quad \text{for } q \geq \frac{3}{4} \quad \text{and} \\ \frac{1}{\sqrt{q}} - \frac{3}{8q\sqrt{q}} + \frac{1}{2q} &\leq \frac{1}{\sqrt{q}} + \frac{1}{q} - \frac{1}{q\sqrt{q}} \quad \text{for } q \geq \frac{25}{16} \end{aligned}$$

we conclude that our bound is strictly tighter, $c_2 < c_2^{\text{qm}}$, for all $q \geq 2$. (Note that since c_2 only applies to classical strategies, this comparison has no implications on the tightness of c_2^{qm} in the quantum setting.)

To close the discussion, let us mention that the ‘‘Number on the Forehead’’ model has been extensively studied in the communication complexity setting. In fact, for certain Boolean functions related to finite-field multiplication (in finite fields of characteristic 2, i.e. $q = 2^n$) a lower bound of the form $\Omega(n/2^m)$ was recently shown [26]. As it appears strikingly similar to (B7) it would be interesting to investigate whether the two scenarios can be related in a rigorous way.

Appendix C: Relativistic bit commitment protocols

All the n -bit strings that appear in the protocols below should be interpreted as elements of the finite field \mathbb{F}_{2^n} . The addition and multiplication are denoted by \oplus and $*$, respectively. Note that the addition is exactly the bitwise XOR but the multiplication *does not* correspond to taking bitwise AND. Moreover, if d is a bit and b is an n -bit string we define

$$d \cdot b = \begin{cases} 0 & \text{if } d = 0, \\ b & \text{if } d = 1. \end{cases}$$

1. Causal structure of the protocol

The relativistic protocols we consider require Alice (who makes the commitment) and Bob (who receives the commitment) to delegate agents to exchange information at two distant locations labelled by 1 and 2. We refer to Alice’s (Bob’s) agent at the i location by \mathcal{A}_i (\mathcal{B}_i). Odd (even) rounds take place at Location 1 (2) and the timing is chosen such that every pair of consecutive rounds is space-like separated (see FIG. 3), which means that Alice’s message in the $(k + 1)$ round must not depend on Bob’s message in the k round and vice versa. These are the *causal constraints* of the protocol and any actions that do not violate them are allowed. Note that these constraints are *strictly more restrictive* than any form of no-signalling between the two locations. We will see in Section C.4 how these constraints restrict the power of the dishonest party.

2. Security model and definitions

Here, we describe the class of cheating strategies allowed for the dishonest parties and how to quantify security for the honest ones.

a. Honest Alice and dishonest Bob

If Alice is honest, she will make an honest commitment to d and Bob should remain completely ignorant about d until the open phase *regardless of his behaviour during the protocol* (for an explanation of the phase structure of relativistic bit commitment schemes see Ref. [17]). More specifically, we define the knowledge of Bob as the knowledge of \mathcal{B}_1 and \mathcal{B}_2 pooled together. (To see why it is not sufficient to require that *each* agent remains ignorant note that if Alice aborts the protocol immediately before the open phase we want Bob to remain ignorant about d for an indefinite




FIG. 3. The protocol requires each pair of consecutive rounds to be space-like separated. Here, we show a valid space-time arrangement for rounds $(k-1)$, k and $(k+1)$. Dots represent rounds and dashed lines represent future light cones. Clearly, both pairs $(k-1, k)$ and $(k, k+1)$ are space-like separated.

period of time. Clearly, in this setting there is enough time for \mathcal{B}_1 and \mathcal{B}_2 to combine their knowledge.) Dishonest Bob is limited only by the causal constraints (explained in Section C.1): the messages announced by his classical agents may be arbitrary functions of all messages exchanged in the past (including any randomness generated before the protocol begins available to both \mathcal{B}_1 and \mathcal{B}_2). Quantum agents are, in addition, allowed to preshare a quantum state (of arbitrary dimension) and then throughout the protocol perform arbitrary measurements on it.

The security definition for honest Alice is based on the *transcript*, which contains complete information about all the messages exchanged in the protocol at both locations and, hence, represents the combined knowledge of \mathcal{B}_1 and \mathcal{B}_2 . Since in our case all the messages are classical the transcript is just a classical random variable denoted by T . Perfect security for Alice means that Bob should not be able to extract any information about her commitment. More precisely, we require that the distributions of transcripts are identical (statistically indistinguishable). We say that a protocol is *perfectly hiding* if

$$\Pr[T = t|d = 0] = \Pr[T = t|d = 1] \quad \forall t$$

for all strategies of dishonest Bob. Note that d is not a random variable, it is Alice's input to the protocol. Therefore, $\Pr[T = t|d = 0]$ should be understood as the probability of seeing the transcript t *given* that Alice has decided to commit to 0.

b. Honest Bob and dishonest Alice

If Bob is honest, no strategy of dishonest Alice should allow her to successfully unveil both bits with high probability. Dishonest Alice is, again, limited only by the causal constraints: the messages she announces may be arbitrary functions of all messages exchanged in the past (including any randomness generated before the protocol begins available to both \mathcal{A}_1 and \mathcal{A}_2). Since we want our protocol to force Alice to become committed in the commit phase we must show that even if Alice decides on the value of the commitment *immediately* after the commit phase she will still fail. Therefore, we consider a model in which her behaviour in the commit phase must be independent of the bit she will attempt to unveil later, denoted by d , but the messages exchanged later might depend on it. Given a particular strategy adopted by Alice in the commit phase we define p_d to be the optimal probability of successfully unveiling d . The protocol is ε -*binding* if

$$p_0 + p_1 \leq 1 + \varepsilon$$

for all strategies of dishonest Alice in the commit phase. Note that this is a weak, non-composable definition of security and that in some other scenarios stronger security definitions can be used [17].

3. Security of sBGKW scheme against quantum adversaries

The protocol requires \mathcal{A}_1 and \mathcal{A}_2 to share a secret n -bit string, $a \in \{0, 1\}^n$, chosen uniformly at random, which is consumed in the protocol.

1. (commit) \mathcal{B}_1 sends \mathcal{A}_1 an n -bit string, $b \in \{0, 1\}^n$, chosen uniformly at random. \mathcal{A}_1 returns $d \cdot b \oplus a$ to \mathcal{B}_1 .
2. (open) \mathcal{A}_1 unveils to \mathcal{B}_1 the committed bit d while \mathcal{A}_2 unveils to \mathcal{B}_2 the secret string a .
3. (verify) Bob collects data from \mathcal{B}_1 and \mathcal{B}_2 and accepts the commitment iff the string returned by \mathcal{A}_1 in the commit phase equals $d \cdot b \oplus a$.

Security for honest Alice follows from the fact that the knowledge of Bob (more precisely, the knowledge of his two agents pooled together) before the open phase is restricted to one n -bit string: the string that \mathcal{B}_1 receives in the commit phase, which in the honest scenario equals $d \cdot b \oplus a$. It follows that as long as Alice is honest and chooses a uniformly this string is distributed uniformly regardless of the value of her commitment.

Security for honest Bob against dishonest Alice who is restricted to classical cheating strategies is fairly intuitive: in order for \mathcal{A}_2 to be able to unveil both commitments she would need to know both a and $a \oplus b$, which implies that she would know b . However, since b is chosen uniformly at random by Bob this must be difficult. This argument can be made rigorous [1] to show that the protocol is ε -binding for $\varepsilon = 2^{-n}$ (and this is actually tight: the trivial strategy of always outputting the string of all zeroes gives $p_0 = 1$ and $p_1 = 2^{-n}$). Unfortunately, this reasoning does not work against quantum adversaries since it could be the case that \mathcal{A}_2 has two distinct measurements (one that reveals a and another one that reveals $a \oplus b$) but since they are incompatible this does not imply anything about her ability to guess b .

To find an explicit bound on $p_0 + p_1$ in the quantum setting we formulate cheating as a non-local game in which \mathcal{A}_1 receives b , \mathcal{A}_2 receives d (the bit she is required to unveil, chosen uniformly at random) and the XOR of their outputs is supposed to equal $d \cdot b$. Winning such a game with probability p_{win} corresponds to a cheating strategy that satisfies $p_0 + p_1 = 2p_{\text{win}}$. More concisely, the rules of the non-local game are

1. \mathcal{A}_1 receives $b \in \{0, 1\}^n$, \mathcal{A}_2 receives $d \in \{0, 1\}$ (both chosen uniformly at random).
2. \mathcal{A}_1 outputs $a_1 \in \{0, 1\}^n$, \mathcal{A}_2 outputs $a_2 \in \{0, 1\}^n$ and they win iff $a_1 \oplus a_2 = d \cdot b$.

This is exactly the game considered in Ref. [22] under the name CHSH_n . They show that

$$p_{\text{win}}(n) \leq \frac{1}{2} + \frac{1}{\sqrt{2^{n+1}}},$$

which is sufficient for our purposes as it implies that

$$p_0 + p_1 \leq 1 + \sqrt{2} \cdot 2^{-n/2}$$

for all strategies of dishonest Alice. Therefore, the protocol is ε -binding with $\varepsilon = 2^{(1-n)/2}$ decaying exponentially in n (but note that the decay rate is half of the decay rate against classical adversaries).

The two-round protocol gets mapped onto a non-local game precisely because of the assumption of no communication. More specifically, we require that \mathcal{A}_1 outputs the answer outside of the future of \mathcal{A}_2 receiving the input and vice versa.

4. A new multi-round protocol based on finite-field arithmetic

The protocol presented in Section C.3 implements a bit commitment scheme that is provably secure against quantum adversaries. Unfortunately, the commitment time is limited by s/c , where s is the spatial separation between Locations 1 and 2 and c is the speed of light. If the two sites are constrained to be on the surface of the Earth then the commitment can only be valid for approximately 42 milliseconds. Here, we present a new, multi-round scheme which, by adding extra intermediate rounds, allows for an arbitrarily long commitment and we prove its security against classical adversaries. Note, however, that the security guarantee depends on the number of rounds of the protocol (which is proportional to the length of the commitment): the longer the commitment, the more resources (randomness and communication bandwidth) are required to achieve the same level of security.

The protocol consists of $m + 1$ rounds labelled by $k \in [m + 1]$, which obey the causal structure described in Section C.1. The commitment is initiated in the first round ($k = 1$), it is sustained for $k = 2, 3, \dots, m$ and is eventually

opened in the last round ($k = m + 1$). Let us emphasise that we want Alice to become committed as soon as the first round is over and if she were able to decide on the value of her commitment in the second round we would consider that cheating. This is necessary to argue that the commitment really begins in the first round. All the rounds before the open phase ($1 \leq k \leq m$) have the same communication pattern: first \mathcal{B}_i sends an n -bit string to \mathcal{A}_i and then she replies with another n -bit string. In the last round \mathcal{A}_i sends \mathcal{B}_i a bit (her commitment) and an n -bit string (proof of her commitment). We will denote the n -bit string announced by Bob (Alice) in the k round by x_k (y_k) regardless of whether he/she is honest or not.

Before the protocol begins Alice generates m strings of n -bits (private from Bob), denoted by $\{a_k\}_{k=1}^m$, drawn independently, uniformly at random from $\{0, 1\}^n$, and distributes them to \mathcal{A}_1 and \mathcal{A}_2 . Similarly, Bob generates m strings of n -bits (private from Alice), denoted by $\{b_k\}_{k=1}^m$, drawn independently, uniformly at random from $\{0, 1\}^n$, and distributes them to \mathcal{B}_1 and \mathcal{B}_2 .

The protocol goes as follows

1. (commit) In the first round \mathcal{B}_1 sends $x_1 = b_1$ to \mathcal{A}_1 and she replies with $y_1 = d \cdot x_1 \oplus a_1$. This initiates the commitment.
2. (sustain) In the k round (for $2 \leq k \leq m$) \mathcal{B}_i sends \mathcal{A}_i the string $x_k = b_k$ and she replies with $y_k = (x_k * a_{k-1}) \oplus a_k$.
3. (open) In the $(m + 1)$ round \mathcal{A}_i sends d and $y_{m+1} = a_m$ to \mathcal{B}_i .
4. (verify) Bob collects data from \mathcal{B}_1 and \mathcal{B}_2 and accepts the commitment iff the strings announced by \mathcal{A}_1 and \mathcal{A}_2 satisfy

$$\begin{aligned} y_{m+1} &= y_m \oplus b_m * y_{m-1} \oplus b_m * b_{m-1} * y_{m-2} \oplus \dots \\ \dots \oplus b_m * b_{m-1} * \dots * b_2 * y_1 \oplus d \cdot b_m * b_{m-1} * \dots * b_1. \end{aligned} \quad (\text{C1})$$

Note that the private strings of Alice and Bob as well as the messages they exchange are actually random variables and that is how they must be treated in the analysis.

a. Correctness

It is easy to verify (by induction) that if Alice and Bob follow the protocol then the condition (C1) is satisfied for any choice of strings $\{a_k\}_{k=1}^m$ and $\{b_k\}_{k=1}^m$.

b. Security for honest Alice

We start with a lemma which formalises the intuition that if we take an arbitrary random variable taking values in a finite field and perform (finite field) addition with a uniform and uncorrelated random variable then there will be no correlations between the input and the output (or any function thereof). More specifically, in the following lemma Y is a random variable from which the input is generated using function g , X is the fresh (finite field) randomness and h is a function allowing us to condition on a certain subset of values of Y .

Lemma C.1. *Let $\mathcal{X} = \mathbb{F}_q$ and \mathcal{Y}, \mathcal{Z} be arbitrary finite sets. Let X and Y be two random variables taking values in \mathcal{X} and \mathcal{Y} , respectively, such that X is uniform and independent from Y*

$$\Pr[X = x, Y = y] = q^{-1} \cdot \Pr[Y = y], \quad (\text{C2})$$

for all $x \in \mathcal{X}$ and $y \in \mathcal{Y}$. Then for arbitrary functions $g : \mathcal{Y} \rightarrow \mathcal{X}$, $h : \mathcal{Y} \rightarrow \mathcal{Z}$ and arbitrary fixed $x \in \mathcal{X}$, $z \in \mathcal{Z}$ it holds that

$$\Pr[X + g(Y) = x \mid h(Y) = z] = q^{-1}.$$

Proof. Note that

$$\begin{aligned} \Pr[X + g(Y) = x, h(Y) = z] &= \sum_{y \in \mathcal{Y}} \Pr[X = x - g(y), h(y) = z, Y = y] \\ &= \sum_{\substack{y \in \mathcal{Y} \\ h(y) = z}} \Pr[X = x - g(y), Y = y] = \sum_{\substack{y \in \mathcal{Y} \\ h(y) = z}} q^{-1} \cdot \Pr[Y = y] = q^{-1} \cdot \Pr[h(Y) = z], \end{aligned}$$

where the second last equality follows from applying the assumption (C2) to every term of the sum. \square

Proposition C.1. *If Alice is honest then the protocol is perfectly hiding.*

Proof. As explained in Section C2a we need to show that the transcripts for $d = 0$ and $d = 1$ after m rounds (immediately before the open phase) are indistinguishable

$$\Pr[Y_1 = y_1, Y_2 = y_2, \dots, Y_m = y_m | d = 0] = \Pr[Y_1 = y_1, Y_2 = y_2, \dots, Y_m = y_m | d = 1]$$

for all y_1, y_2, \dots, y_m . In fact, we will show by induction that

$$\Pr[Y_1 = y_1, Y_2 = y_2, \dots, Y_t = y_t | d = b] = 2^{-nt}, \quad (\text{C3})$$

for all $t \in [m]$ regardless of the value of $b \in \{0, 1\}$, which clearly satisfies the indistinguishability condition.

Honest Alice will follow the protocol, which means that $\{A_k\}_{k=1}^m$ are drawn independently, uniformly at random from $\{0, 1\}^n$, the value of her commitment is d and then Alice's message in k round is

$$Y_k = \begin{cases} d \cdot X_1 \oplus A_1 & \text{for } k = 1, \\ Y_k = (X_k * A_{k-1}) \oplus A_k & \text{for } 2 \leq k \leq m. \end{cases} \quad (\text{C4})$$

Bob, on the other hand, is only limited by the causal constraints, which means that his message in the k round might depend on some randomness preshared between \mathcal{B}_1 and \mathcal{B}_2 , denoted by R_B , and all the responses of Alice which belong to the past of the k round. Therefore, without loss of generality his message in the k round is

$$X_k = f_k(R_B, Y_1, Y_2, \dots, Y_{k-2}) \quad (\text{C5})$$

for some arbitrary function f_k (we include all randomness used by Bob in R_B so f_k is deterministic).

In this scenario the full transcript is a deterministic function of Alice's commitment d , her private randomness $\{A_k\}_{k=1}^m$ and Bob's preshared randomness R_B . For every string announced by Alice and Bob we can explicitly find the subset of random variables it may depend on as listed in the table below

message	random variables it might depend on
X_1	R_B
X_2	R_B
X_3	d, R_B, A_1
\vdots	\vdots
X_k	$d, R_B, A_1, A_2, \dots, A_{k-2}$
\vdots	\vdots
X_m	$d, R_B, A_1, A_2, \dots, A_{m-2}$
Y_1	d, R_B, A_1
Y_2	d, R_B, A_1, A_2
Y_3	d, R_B, A_1, A_2, A_3
\vdots	\vdots
Y_k	$d, R_B, A_1, A_2, \dots, A_k$
\vdots	\vdots
Y_m	$d, R_B, A_1, A_2, \dots, A_m$

First, we verify that Eq. (C3) holds for $t = 1$

$$\Pr[Y_1 = y_1 | d = b] = \Pr[b \cdot X_1 \oplus A_1 = y_1] = \Pr[b \cdot f_1(R_B) \oplus A_1 = y_1] = 2^{-n},$$

where the first two equalities follow from Eqs. (C4) and (C5), respectively. The last equality is a direct consequence of Lemma C.1 (in a simplified form: no conditioning) applied to $X = A_1$, $Y = (b, R_B)$, $g(Y) = b \cdot f_1(R_B)$. Now, suppose that Eq. (C3) holds for $t = k$. Then

$$\begin{aligned} & \Pr[Y_1 = y_1, \dots, Y_{k+1} = y_{k+1} | d = b] \\ &= \Pr[Y_{k+1} = y_{k+1} | d = b, Y_1 = y_1, \dots, Y_k = y_k] \cdot \Pr[Y_1 = y_1, \dots, Y_k = y_k | d = b] \\ &= \Pr[(X_{k+1} * A_k) \oplus A_{k+1} = y_{k+1} | d = b, Y_1 = y_1, \dots, Y_k = y_k] \cdot 2^{-nk} \\ &= 2^{-n} \cdot 2^{-nk} = 2^{-(n+1)k}, \end{aligned}$$

where the second last inequality follows from applying Lemma C.1 to

$$\begin{aligned} X &= A_{k+1}, \\ Y &= (b, R_B, A_1, \dots, A_k), \\ g(Y) &= X_{k+1} * A_k, \\ h(Y) &= (Y_1, Y_2, \dots, Y_k). \end{aligned} \tag{C6}$$

Note that it is not immediately obvious and the reader should verify (using the table presented above) that the quantities on the right-hand side of Eq. (C6) are functions of Y alone, and therefore satisfy the assumptions of the lemma. This shows that Eq. (C3) holds for $t = k + 1$ and so by induction it must hold for all $t \in [m]$. Therefore, even just before the open phase the transcript contains no information about Alice's commitment and the protocol is perfectly hiding. \square

c. Security for honest Bob

Proposition C.2. *If Bob is honest then the protocol is ε -binding for $\varepsilon = \omega_m$ defined in Eq. (B2).*

Proof. Honest Bob will follow the protocol so $X_k = B_k$, where $\{B_k\}_{k=1}^m$ are drawn independently, uniformly at random from $\{0, 1\}^n$. Let R_A be any randomness preshared by \mathcal{A}_1 and \mathcal{A}_2 before the protocol begins. The most general cheating strategy for Alice allowed by the security model described in Section C 2 b is a collection of (deterministic) functions, $\{f_k\}_{k=1}^{m+1}$, all of which output an n -bit string while their inputs are as described below.

- Alice's message in the commit phase might depend on the preshared randomness and the first message of Bob

$$Y_1 = f_1(R_A, B_1).$$

- Alice's messages during the sustain phase ($k \in \{2, 3, \dots, m\}$) might depend on the preshared randomness, Bob's messages from the past and the bit she is trying to unveil d

$$Y_k = f_k(R_A, B_1, B_2, \dots, B_{k-2}, B_k, d).$$

Note that Y_k must not depend on B_{k-1} because, by assumption, it does not belong to the past of the k round.

- Alice's message in the open phase might depend on the preshared randomness, Bob's messages from the past and the bit she is trying to unveil d

$$Y_{m+1} = f_{m+1}(R_A, B_1, B_2, \dots, B_{m-1}, d).$$

Again, Y_{m+1} must not depend on B_m .

The commitment to d will be accepted iff (C1) is satisfied for that value of d and let us denote this event by H_d . By definition $p_d = \Pr[H_d]$ and since both events are defined over $(R_A, B_1, B_2, \dots, B_m)$ it is meaningful to talk about $H_0 \vee H_1$ and $H_0 \wedge H_1$. (Note that this reasoning does not work in the quantum setting since H_0 and H_1 are *not defined* simultaneously.) To bound $p_0 + p_1$ we use $\Pr[H_0] + \Pr[H_1] = \Pr[H_0 \vee H_1] + \Pr[H_0 \wedge H_1] \leq 1 + \Pr[H_0 \wedge H_1]$. The event $H_0 \wedge H_1$ happens if (C1) is satisfied for both values of d . Define K to be the event that the XOR of the two conditions (i.e. Eq. (C1) for $d = 0$ and $d = 1$) is satisfied

$$\begin{aligned} K \iff & B_1 * B_2 * \dots * B_m = g_{m+1}(R_A, B_1, B_2, \dots, B_{m-1}) \oplus g_m(R_A, B_1, B_2, \dots, B_{m-2}, B_m) \\ & \bigoplus_{k=2}^{m-1} B_m * B_{m-1} * \dots * B_{k+1} * g_k(R_A, B_1, B_2, \dots, B_{k-2}, B_k), \end{aligned}$$

where $g_k(R_A, B_1, B_2, \dots, B_{k-2}, B_k) = f_k(R_A, B_1, B_2, \dots, B_{k-2}, B_k, d = 0) \oplus f_k(R_A, B_1, B_2, \dots, B_{k-2}, B_k, d = 1)$. Note that since $H_0 \wedge H_1 \implies K$ we have $\Pr[H_0 \wedge H_1] \leq \Pr[K]$.

To bound $\Pr[K]$ note that the right-hand side contains exactly m terms, but each of them depends on $(m - 1)$ B 's; none of the terms depends on *all* B 's *simultaneously*. The terms corresponding to $2 \leq k \leq m - 1$ have some internal

structure (e.g. the dependence on B_m is *not* arbitrary) but we can relax the problem to the case where the k term is an arbitrary function of all the B 's except for B_k denoted by h_k . The winning condition for the relaxed game is

$$B_1 * B_2 * \dots * B_m = \bigoplus_{k=1}^m h_k(B_{[m] \setminus \{k\}})$$

and clearly the winning probability is an upper bound on $\Pr[K]$. In Section B.2 we define the optimal winning probability for this game to be ω_m . This concludes the proof since

$$p_0 + p_1 \leq 1 + \Pr[K] \leq 1 + \omega_m.$$

□

Note that a non-trivial upper bound on ω_m (for an arbitrary m) can be obtained using a recursive argument presented in Section B.4.