

Practical Relativistic Zero-Knowledge for NP

Claude Crépeau^{1*}, Arnaud Massenet^{2**},
Louis Salvail^{3*}, Lucas Stinchcombe^{4**}, and Nan Yang^{5***}

¹ McGill University, Montréal, Québec, Canada. crepeau@cs.mcgill.ca

² University of Oxford, Oxford, Oxfordshire, UK. arnaud.massenet@mail.mcgill.ca

³ Université de Montréal, Montréal, Québec, Canada. salvail@iro.umontreal.ca

⁴ Bloomberg L.P, Tokyo, Japan. lucas.stinchcombe@mail.mcgill.ca

⁵ Concordia University, Montreal, Quebec, Canada. na_yan@encs.concordia.ca

Abstract. In this work we consider the following problem: in a Multi-Prover environment, how close can we get to prove the validity of an **NP** statement in Zero-Knowledge? We exhibit a set of two novel Zero-Knowledge protocols for the 3-COLORability problem that use two (*local*) provers or three (*entangled*) provers and only require them to reply two trits each. This greatly improves the ability to prove Zero-Knowledge statements on very short distances with very minimal equipment.

1 Introduction

The idea of using distance and special relativity (a theory of motion justifying that the speed of light is a sort of asymptote for displacement) to prevent communication between participants to multi-prover proof systems can be traced back to Kilian[1]. Probably, the original authors (Ben Or, Goldwasser, Kilian and Wigderson) of [2] had that in mind already, but it is not explicitly written anywhere. Kent was the first author to venture into *sustainable* relativistic commitments [3] and introduced the idea of arbitrarily prolonging their life span by playing some ping-pong protocol between the provers (near the speed of light). This idea was made considerably more practical by Lunghi *et al.* in [4] who made commitment sustainability much more efficient. This culminated into an actual implementation by Verbanis *et al.* in [5] where commitments were sustained for more than a day!

As nice as this may sound, such *long-lasting* commitments have found so far very little practical use. Consider for instance the zero-knowledge proof for Hamiltonian Cycle as introduced by Chailloux and Leverrier[6]. Proving in Zero-Knowledge that a 500-node graph contains a Hamiltonian cycle would require transmitting 250 000 bit commitments (each of a couple hundreds of bits in length) and eventually sustaining them before the verifier can announce his choice of unveiling the whole adjacency matrix or just the Hamiltonian cycle. For a graph of $|V|$ vertices, this would require an estimated $200|V|^2$ bits of communication before the verifier can announce his choice *chall* (see Fig. 1). This makes the application prohibitively expensive. If you use a larger graph, you will need more time to commit, leading to more distance to implement the protocol of [6]. Either a huge separation is necessary between the provers (so that one of them can unveil according to the verifier's choice *chall* before he finds out the committal information B used by the other prover while the former must commit all the necessary information before he can find out the verifier's choice *chall*) or we must achieve extreme communication speeds between prover-verifier pairs. This would only be possible by vastly parallelizing communications between them at high cost. Modern (expensive) top-of-line communication equipment may reach throughputs of roughly 1Tbits/sec. A back of the envelope calculation estimates that the distance between the verifiers must be at least 100 km to transmit 250 000 commitments at such a rate.

* Supported in part by FRQNT (INTRIQ) and NSERC.

** This research was performed while a student at McGill University under C.C.'s supervision.

*** Supported in part by Professors Jeremy Clark, and Claude Crépeau.




Fig. 1. Space-Time diagrams of [6]’s ZK-MIP* for NP. (45° diagonals are the speed of light.)

In the above two diagrams, V_1 at a first location sends a random matrix B to P_1 who uses each entry to commit an entry of the adjacency matrix Y of G . At another location, V_2 sends a random challenge $chall$ to P_2 who unveils all or some commitments as A . At all times, V_1 and V_2 must make sure that the answers they get from P_1 and P_2 come early enough that the direct communication line between V_1 and V_2 (even at the speed of light) is not crossed. The transition from left to right shows that increasing the number of nodes (and thus increasing the total commit time) pushes the verifiers further away from each other. In [6] the distance must increase quadratically with the number of nodes in the graph.

In this work we consider the following problem: in a Multi-Prover environment, how close can we get the provers in a Zero-Knowledge IP showing the validity of an NP statement? We exhibit a set of (3) novel Zero-Knowledge protocols for the 3-COLORability problem that use two (*local*) provers or three (*entangled*) provers and only require them to communicate two trits each after having each received an edge and two trits each from the verifier. This greatly improves the ability to prove Zero-Knowledge statements on very short distances with very little equipment. In comparison, the protocol of [6] would require transmitting millions of bits between a prover and his verifier before the latter may disclose what to unveil or not. This implies the provers would have to be very far from each other because all of these must reach the verifier *before* the former can communicate with its partner prover.

Although certain algebraic zero-knowledge multi-prover interactive proofs for NP and NEXP using explicitly no commitments at all have been presented before in [7], [8] (sound against local provers) and [9],[10] (sound against entangled provers), in the local cases making these protocols entanglement sound is absolutely non-trivial, whereas in the entangled case the multi-round structure and the amount of communication in each round makes implementing the protocol completely impractical as well. (To their defense, the protocols were not designed to be *practical*).

The main technical tool we use in this work is a general Lemma of Kempe, Kobayashi, Matsumoto, Toner, and Vidick[11] to prove soundness of a three-prover protocol when the provers are *entangled* based on the fact that a two-prover protocol version is sound when the provers are

only *local*. More precisely, they proved this when the three-prover version is the same as the two-prover version but augmented with an extra prover who is asked exactly the same questions as one of the other two at random and is expected to give the same exact answers.

Our protocols build on top of the earlier protocol due to Cleve, Høyer, Toner and Watrous[12] who presented an extremely simple and efficient solution to the 3-COL problem that uses only two provers, each of which is queried with either a node from a common edge, or twice the same node. In the former case, the verifier checks that the two ends of the selected edge are of distinct colours, while in the latter case, he check only that the provers answer the same colour given the same node. On the bright side, their protocol did not use commitments at all but unfortunately it did not provide Zero-Knowledge either. Moreover, it is a well established fact that this protocol cannot possibly be sound against entangled provers, because certain graph families have the property that they are not 3-colourable while having entangled-prover pairs capable of winning the game above with probability one. This was already known at the time when they introduced their protocol. The reason this protocol is not zero-knowledge follows from the undesirable fact that dishonest verifiers can discover the (random) colouring of non-edge pairs of nodes in the graph, revealing if they are of the same colour or not in the provers' colouring.




Fig. 2. Space-Time diagram of our ZK-MIP* for NP. (45° diagonals are the speed of light.)

We are able to remedy to the zero-knowledge difficulty by allowing the provers to use commitments for the colour of their nodes. However they use these commitments in an innovative way that we call the *unveil-via-commit principle* (of independent interest) explained below. For this purpose we use commitments similar to those of Lunghi *et al.*[4] but in their simplest form possible, over the field \mathbb{F}_3 (or \mathbb{F}_4 if you insist working in binary), and thus with extremely weak binding property but also minimal in communication cost: a complete execution of the basic protocol transmits exactly two node numbers (using only $\log |V|$ bits each) and two trits from verifiers to provers and two trits back from the provers to verifiers (see Fig. 2). This implies that for a fixed communication speed, the minimal distance of the provers in our protocol increases logarithmically with the number of nodes whereas the same parameter grows quadratically in [6]. Nevertheless, this is good enough to obtain a zero-knowledge version of the protocol that remains sound against *local* pairs of provers. The main idea being that the provers will each commit to the colours of two requested nodes only if they form an edge of the graph. To unveil the colour of any node, the verifiers must request commitment of the *same* node by *both* provers but using different randomizations. This way the verifiers may compute the colour of a node from the *linear*

system established by the two commitments and not by explicitly requesting anyone to unveil. This is the unveil-via-commit principle (very similar to the double-spending mechanism of the untraceable electronic cash of Chaum, Fiat and Naor[13]). We then use the Lemma of [11] to prove soundness of the three-prover version of this protocol even when the provers are *entangled*. A positive side of the protocol of [6], however, is the fact that only two provers are necessary while we use three. Zero-Knowledge follows from the fact that only two edge nodes can be unveiled by requesting the same edge to both provers. Otherwise only a single node may be unveiled. Finally, we show that even the three-prover version of this protocol retains the zero-knowledge property: requesting any three edges from the provers may allow the verifiers to unveil the colours of a triangle in the graph but never two end-points that do not form an edge (going to four provers would however defeat the zero-knowledge aspect).

An actual physical implementation of this protocol is currently being developed in collaboration with Pouriya Alikhani (McGill), Nicolas Brunner, Sébastien Designolle, Weixu Shi, and Hugo Zbinden (Université de Genève).

1.1 Implementations Issues

Traditionally in the setup of Multi-Prover Interactive Proofs, there is a single verifier interacting with the many provers. However, when implementing no-communication via spatial separation (the so called relativistic setting) it is standard to break the verifier in a number of verifiers equal to the number of provers, each of them interacting at very short distance from their own prover. The verifiers can use the timing of the replies of their respective provers to judge their relative distance. In practice, this means that we can implement MIPs under relativistic assumptions if the verifier are “split” into multiple verifiers, each locally interacting with its corresponding prover. The verifiers use the distance between *themselves* to enforce the impossibility of the provers to communicate: no message from a verifier can be used to reply to another verifier faster than the speed of light *wherever the provers are located*.

Moreover, multi-prover interactive proof systems may have several rounds in addition to several provers. In general, protocols with several rounds may cause a treat to the inherent assumption that the provers are not allowed to communicate during the protocol’s execution. Nevertheless, most of the existing literature resolves this issue by providing an honest verifier that is *non-adaptive*. To simplify this task, most of the protocols are actually single-round. We stick to these guidelines in this work. Moreover, in order to prove soundness of our protocols against entangled provers, we use a theorem that is currently only proven for single-round protocols. The protocols we describe are indeed single-round and non-adaptive.

2 Preliminaries

2.1 Notations

Random variables $A, B \in \Gamma$ are said to be equivalent, denoted $A = B$, if for all $x \in \Gamma$, $\Pr(A = x) = \Pr(B = x)$. The class of probabilistic polynomial-time Turing machines will be denoted PPT in the following. A PPT Turing machine is one having access to a fresh infinite read-only tape of random values (uniform values from the set of input symbols) at the outset of the computation. In the following, adversaries will also be allowed (in some cases) to be quantum machines. The precise ways quantum and classical machines are defined is not important in the following.

For M a Turing machine, we denote by $M(x)$ its execution with x on its input tape (x being a string of the tape alphabet symbols). A Turing machine (quantum or classical) augmented with read-only auxiliary-input tapes and write-only auxiliary-output tapes is called an *interactive Turing machine* (ITM). Read-only input tapes provide incoming messages while the

write-only output tapes allow to send messages. Interactive Turing machine M_1 and M_2 are said to *interact* when for each of them, one of its write-only auxiliary-output tape corresponds to one read-only auxiliary-input tape of the other Turing machine. An execution of interactive Turing machines M_1, \dots, M_k on common input x is denoted $[M_1 \dots M_k](x)$. For $1 \leq i \leq k$, machine M_i *accepts* the interactive computation on input x if it stops in state **accept** after the execution $[M_1 \dots M_k](x)$. When the ITM M_i that accepts a computation is clear from the context, we say that $[M_1 \dots M_k](x)$ accepts when M_i 's final state is **accept**. In this scenario, $\Pr([M_1 \dots M_k](x) = \text{accept})$ denotes the probability that M_i terminates in state **accept** upon common input x . Quantum machines are also interacting through communication tapes the same way than for classical machines. When a quantum machine M_1 interacts with a classical machine M_2 , we suppose that the write-only auxiliary tape and the read-only auxiliary tape of M_1 used to communicate with M_2 are classical. This is the situation we will be addressing almost all the time in the following. A quantum machine M is also allowed to have a quantum auxiliary read-only input tape that may contain a part of a quantum state shared with other machines. This allows to model machines sharing entanglement at the outset of an interactive computation. Henceforth, we suppose that the (main) input tape of all machines (quantum or classical) is classical.

In the following, $G = (V, E)$ denotes an undirected graph with vertices V and edges E . If $n = |V|$ then we denote the set of vertices in G by $V = \{1, 2, \dots, n\}$. We suppose that $(i, i) \notin E$ for all $1 \leq i \leq n$ (i.e. G has no loop). We denote uniquely each edge in E as (i, j) with $j > i$. For $i \in V$, let $\text{Edges}(i) := \{(j, i) \in E\}_{j < i} \cup \{(i, j) \in E\}_{j > i}$ be the set of edges connecting vertex i in G . For $e, e' \in E$, we define $e \cap e' = i \in V$ if e and e' have only one vertex $i \in V$ in common. When e and e' have four distinct vertices in V , we set $e \cap e' = 0$. Finally, when $e = e'$, we set $e \cap e' := \infty$. For readability, we use the following special notations: $(a, b) \neq (c, d)$ means $a \neq c$ **and** $b \neq d$, while as always, $(a, b) \neq (c, d)$ simply means $a \neq c$ **or** $b \neq d$.

2.2 Non-local Games, Multi-Prover Interactive Proofs, and Relativistic Proofs

Multi-provers interactive protocols are protocols involving a set of *provers* modelled by interactive Turing machines, each of them interacting with an interactive PPT Turing machine called the verifier V . Although all provers may share an infinite read-only auxiliary input tape at the outset of their computation, they do not interact with each other. When the provers are quantum, an extra auxiliary read-only quantum input tape is given and can be entangled with other provers at the beginning.

Definition 1. Let P_1, \dots, P_k be computationally unbounded interactive Turing machines and let V be an interactive PPT Turing machine. The P_i 's share a joint, infinitely long, read-only random tape (and an auxiliary reads-only quantum input tape if the provers are quantum). Each P_i interacts with V but cannot interact with P_j for any $1 \leq j \neq i \leq k$. We call $[P_1, \dots, P_k, V]$ a k -prover interactive protocol (k -prover IP).

A $[P_1, \dots, P_k, V]$ k -prover interactive protocol is a *multi-prover interactive proof system* for L if it can be used to show V that a public input x is such that $x \in L$. At the end of its computation, V concludes $x \in L$ if and only if it ends up in state **accept**. We restrict our attention to interactive proof systems with perfect completeness since all our protocols have this property.

Definition 2. The k -prover interactive protocol $\Pi = (P_1, \dots, P_k, V)$ is said to be a k -prover interactive proof system with perfect completeness for L if there exists $q(n) < 1 - \frac{1}{\text{poly}(n)}$ such that following holds:

perfect completeness: $(\forall x \in L) [\Pr([P_1, \dots, P_k, V](x) = \text{accept}) = 1]$,
soundness: $(\forall x \notin L)(\forall \tilde{P}_1, \dots, \tilde{P}_k) [\Pr([\tilde{P}_1, \dots, \tilde{P}_k, V](x) = \text{accept}) \leq q(|x|)]$.

The parameter $q(|x|)$ is called the soundness error of Π . Soundness can hold against classical provers or against quantum provers sharing entanglements. The former case is called sound against classical provers while the latter is called sound against entangled provers.

Consider a k -prover interactive proof system $\Pi(x)$ (with or without perfect completeness) for L executed with public input $x \notin L$. In this situation, $\Pi(x)$ defines what is called a *quantum game*. The minimum value $q(|x|)$ such that for all P'_1, \dots, P'_k , $\Pr([P'_1, \dots, P'_k, V](x) = \text{accept}) \leq q(|x|)$ is often called the *classical value of game $\Pi[x]$* and is denoted $\omega(\Pi(x))$ when the provers are restricted to be classical and unable to communicate with each other upon public input x . When the provers, still unable to communicate with each other, are allowed to carry their computation quantumly and share entanglements, we denote by $\omega^*(\Pi(x)) \geq \omega(\Pi(x))$ the minimum value $q(|x|)$ such that for all such quantum provers P'_1, \dots, P'_k , $\Pr([P'_1, \dots, P'_k, V](x) = \text{accept}) \leq q(|x|)$. In this case, $\omega^*(\Pi(x))$ is called the *quantum value of game $\Pi(x)$* . A k -prover interactive proof system for L is said to be *symmetric* if V can permute the questions to all provers without changing their distribution. The following result of Kempe, Kobayashi, Matsumoto, Toner, and Vidick[11] shows that the classical value of a symmetric one-round classical game cannot be too far from the quantum value of a *modified* game. Given a symmetric one-round two-prover game Π , one can always add a third prover P_3 and V asks P_3 the same question than P_1 with probability $\frac{1}{2}$ or the same question than P_2 with probability $\frac{1}{2}$. Then, V accepts if P_1 and P_2 would be accepted in $\Pi(x)$ and if P_3 returns the same answer than the one returned by the prover it emulates. We call $\Pi'(x)$ the modified game obtained that way from $\Pi(x)$.

Lemma 1 ([11], Lemma 17). *Let $\Pi(x)$ be a two-prover one-round symmetric game and let $\Pi'(x)$ be its modified version with three provers. If $\omega^*(\Pi'(x)) > 1 - \varepsilon$ then $\omega(\Pi(x)) > 1 - \varepsilon - 12|Q|\sqrt{\varepsilon}$ where Q is the set of V 's possible questions to a prover in Π .*

Lemma 1 remains true for non-symmetric two-prover one-round protocol by first making them symmetric at the cost of increasing the size of Q . This is always possible without changing the classical value of the game and by using twice the number of questions $|Q|$ of the original game (Lemma 4 in [11]).

Let $[P_1, \dots, P_k, V]$ be a k -prover IP. We denote by $\mathbf{view}(P_1, \dots, P_k, V, x)$ the probability distribution of V 's outgoing and incoming messages with all provers according V 's coin tosses.

Definition 3. *Let $[P_1, \dots, P_k, V]$ be a k -prover interactive proof system for L . We say that $[P_1, \dots, P_k, V]$ is perfect zero-knowledge if for all PPT interactive Turing machines \tilde{V} there exists a PPT machine Sim (i.e. the simulator) having blackbox access to \tilde{V} such that for all x ,*

$$\mathbf{view}(P_1, \dots, P_k, \tilde{V}, x) = \text{Sim}(x) \text{ ,}$$

and both random variables are equivalent. In the following, we allow \tilde{V} to be a quantum machine but our simulators will always be classical machines with blackbox access to \tilde{V} . If the zero-knowledge condition holds against quantum \tilde{V} , we say that the proof system is perfect zero-knowledge against quantum verifiers.

2.3 Multi-Prover Commitments with Implicit Unveiling

Our multi-prover proof systems for 3COL use a simple 2-committer commitment scheme with a property allowing to guarantee perfect zero-knowledge. In this section, we give the description of this simple commitment scheme with its important properties for our purposes.

Assume that provers P_1 and P_2 share ℓ values $c_1, c_2, \dots, c_\ell \in \mathbb{F}$ where \mathbb{F} is a finite set. V wants to check that these values satisfy some properties without revealing them all. Assume that \mathbb{F} is a field with operations $+$ and \cdot .

Bit commitment schemes have been used in the multi-prover model ever since it was introduced in [2]. The original scheme was basically $w_i := b_i \cdot r_i + c_i$, a commitment w_i to value $c_i \in \mathbb{F}$ using pre-agreed random mask $b_i \in_R \mathbb{F}$ and randomness $r_i \neq 0$ provided by V. Kilian[14] had a binary version where each bit $c_i := c_i^1 \oplus c_i^2 \oplus c_i^3$ is shared among provers P_1 and P_2 (and therefore \mathbb{F} needs only to be a group). To commit c_i , V samples c_i^h from P_1 and c_i^j from P_2 at random. If $j = h$ but $c_i^j \neq c_i^h$, V immediately rejects the commitment. Otherwise either P_1 or P_2 may unveil by disclosing c_i^1, c_i^2, c_i^3 at a later time. Somehow, bad recollection of [2]’s scheme lead [15] to a similar but different scheme defining $w_i := c_i \cdot r_i + b_i$, a commitment w_i to bit $c_i \in \{0, 1\}$ using pre-agreed bit mask $b_i \in_R \{0, 1\}$ and binary randomness r_i provided by their corresponding verifiers. Although this form of commitment is intimately connected to the CHSH game [16] and the Popescu-Rohrlich box[17], this proximity is not relevant for the soundness and the completeness of our protocols, even against entangled provers. Although the (limited) binding property of these schemes has been established in [3, 18, 5, 19, 4, 6] against entangled provers, we only use this commitment scheme against classical provers, only sharing classical information before the execution of the protocol. The weak binding property of these schemes against entangled provers does not allow us to get sound and complete proof systems against these provers. We shall rather get completeness and soundness against entangled provers using a different technique from [11] that requires a third prover.

For an arbitrary field \mathbb{F} , the commitment scheme produces commitment $w_i := c_i \cdot r_i + b_i$ to field element $c_i \in \mathbb{F}$ using pre-agreed field element mask b_i (specific to value $1 \leq i \leq \ell$) and random field element $r_i \neq 0$ provided by their corresponding verifiers. Many results were proven for this specific form of the commitments. Notice however that the two versions discussed above, $w_i := b_i \cdot r_i + c_i$ in the former case and $w_i := c_i \cdot r_i + b_i$ in the latter have equivalent binding property(left as a simple exercise). Considering, the former as being the degree-one secret sharing [20] of c_i hidden in the degree zero term, while the latter being the degree-one secret sharing of c_i hidden in the degree one term, we decided to use the former (original BGKW form) because all the known results about secret sharing are generally presented in this form. In particular, this form is more adapted to higher degree generalizations such as $w_i := a_i \cdot r_i^2 + b_i \cdot r_i + c_i$ being the degree-two secret sharing of c_i hidden in the degree zero term, and so on.

Moreover, this choice turns out to simplify our (perfect) zero-knowledge simulator. For the rest of this paper, we use $w_i := b_i \cdot r_i + c_i$ where $w_i, b_i, c_i \in \mathbb{F}_3$ and $r_i \in \mathbb{F}_3^*$. Provers therefore commit to trits, one value for each node corresponding to its colour in a 3-colouring of graph $G = (V, E)$. The values shared between P_1 and P_2 are therefore, for each node $i \in V$, the colour c_i of that node.

Suppose that V asks P_1 to commit on the colour c_i of node $i \in V$ using randomness $r \in_R \mathbb{F}_3^*$. Let $w = b_i \cdot r + c_i$ be the commitment returned to V by P_1 . Suppose V asks P_2 to commit on the colour c'_j of node $j \in V$ using randomness $r' \in_R \mathbb{F}_3^*$. Let $w' = b_j \cdot r' + c'_j$ be the commitment issued to V by P_2 . The following 3 cases are possible depending on V’s choices for i, j, r , and r' :

1. (*forever hiding*) if $i \neq j$ then V learns nothing on neither c_i nor c'_j since w and w' hide c_i and c'_j with random and independent masks $b_i \cdot r$ and $b_j \cdot r'$ respectively. Even knowing $r, r' \in \mathbb{F}_3^*$, $b_i \cdot r$ and $b_j \cdot r'$ are uniformly distributed in \mathbb{F}_3 .
2. (*the consistency test*) If $i = j$ and $r = r'$ then V can verify that $w = w'$. This corresponds to the immediate rejection of V in Kilian’s two-prover commitment described above. It allows V to make sure that P_1 and P_2 are consistent when asked to commit on the same value.
3. (*implicit unveiling*) If $i = j$ and $r' \neq r$ then V can learn c_i (assuming $w = b_i \cdot r + c_i$ and $w' = b_i \cdot r' + c_i$) the following way. V simply computes $c_i := 2^{-1} \cdot (w + w')$ (Note that over an arbitrary field $c_i := (wr' - w'r)(r' - r)^{-1}$ whenever $r \neq r'$). Interpreting the meaning of this test can be done when considering a strategy for P_1 and P_2 that always passes the consistency test. In this case, $w = b_i \cdot r + c_i$ and $w' = b_i \cdot r' + c_i$ are satisfied and V learns the committed value c_i .

As long as P_1 and P_2 are *local* (or *quantum non-local*) they cannot distinguish which option V has picked among the three. The consistency test makes sure that if P_1 and P_2 do not commit on identical values for some $1 \leq i \leq \ell$ then V will detect it when V picks the consistency test for commitment w and w' in position i .

3 Classical Two-Prover Protocol

First, consider a small variation over the protocol of Cleve et al. presented in [12]. In their protocol, when P_1 and P_2 both know and act upon the same valid 3-colouring of G , V asks each prover for the colour of a vertex in $G = (V, E)$. Consistency is verified when V asks the same vertex to each prover and compares that the same colour has been provided. The colorability is checked when the provers are asked for the colour of two connected vertices in G . This way of proceeding is however problematic for the zero-knowledge condition. V could be asking two nodes that do not form an edge for which their respective colour will be unveiled. This certainly allows V to learn something about P_1 's and P_2 's colouring. Indeed, repeating this many times will allow V to efficiently reconstruct a complete colouring. To remedy partially this problem, V is instead asking each prover the colouring of an entire edge of G . The colouring is (only) checked when both provers are asked the same edge, while consistency is checked when two intersecting edges are asked to the provers.

3.1 Distribution of questions

Let $G = (V, E)$ be a connected undirected graph. Let us define the probability distribution $\mathcal{D}_G = \{(p(e, e'), (e, e'))\}_{e, e' \in E}$ for the pair $(e, e') \in E \times E$ that V picks with probability $p(e, e')$ before announcing e to P_1 and e' to P_2 . For $e, e' \in E$ such that $e \cap e' = \emptyset$, we set $p(e, e') := 0$ so that V never asks two disconnected edges in G (this would give no useful information).

The first thing to do is to pick $e = (i, j) \in E$ uniformly at random. With probability ϵ (to be selected later), we set $e' = e$, which allows for an edge-verification test. With probability $1 - \epsilon$, we perform a well-definition test as follows. With probability $\frac{1}{2}$, $e' \in \text{Edges}(i)$ uniformly at random and with probability $\frac{1}{2}$, $e' \in \text{Edges}(j)$ uniformly at random. In other words, the well-definition test picks the second edge e' with probability $\frac{1}{2}$ among the edges connecting $i \in V$ and with probability $\frac{1}{2}$ among the edges connecting $j \in V$. It follows that for $e' \in \text{Edges}(i) \cup \text{Edges}(j)$ with $e \neq e'$, we have, for $e = (i, j) \in E$,

$$p(e, e') = \frac{1 - \epsilon}{2|E|} \left(\frac{|\{e'\} \cap \text{Edges}(i)|}{|\text{Edges}(i)|} + \frac{|\{e'\} \cap \text{Edges}(j)|}{|\text{Edges}(j)|} \right) . \quad (1)$$

We also get

$$p(e, e) = \frac{\epsilon}{|E|} + \frac{1 - \epsilon}{2|E|} \left(\frac{1}{|\text{Edges}(i)|} + \frac{1}{|\text{Edges}(j)|} \right) \geq \frac{\epsilon}{|E|} . \quad (2)$$

It is easy to verify that \mathcal{D}_G is a properly defined probability distribution over pairs of edges.

3.2 A Variant Over the Two-Prover Protocol of Cleve et al.

Distribution \mathcal{D}_G produces two edges where the first one is provided to P_1 while the second one is provided to P_2 . Each prover then returns the colour of each node of the edge to V . We denote the resulting protocol $\Pi_{\text{std}}^{(2)}$.

Protocol $\Pi_{\text{std}}^{(2)}[G]$: Two-prover, 3-COL.

Provers P_1, P_2 pre-agree on a random 3-colouring of G : $\{(i, c_i) | c_i \in \mathbb{F}_3\}_{i \in V}$ such that $(i, j) \in E \implies c_j \neq c_i$.

Interrogation phase:

- V picks $((i, j), (i', j')) \in_{\mathcal{D}_G} E \times E$, sends (i, j) to P_1 and (i', j') to P_2 .
- If $(i, j) \in E$ then P_1 replies with c_i, c_j .
- If $(i', j') \in E$ then P_2 replies with $c_{i'}, c_{j'}$.

Check phase:

- **Edge-Verification Test:**
if $(i, j) = (i', j')$ then V accepts iff $c_i = c_{i'} \neq c_j = c_{j'}$.
 - **Well-Definition Test:**
if $(i, j) \cap (i', j') = h \in V$ then V accepts iff $c_h = c'_h$.
-

The perfect soundness of this protocol is not difficult to establish along the same lines of the proof of soundness for the original protocol in [12]. On the other hand, zero-knowledge does not even hold against honest verifiers. V learns the colour of each node contained in any two edges of G . This is certainly information about the colouring that V learns after the interaction. To some extent, the modifications we applied to the 2-prover interactive proof system of [12] leaks even more to V . In the next section, we show that the 2-prover commitment scheme, that we introduced in Sect. 2.3, can be used in protocol $\Pi_{\text{std}}^{(2)}$ to prevent this leakage completely.

4 Perfect Zero-Knowledge Two-Prover Protocol

We modify the protocol of section 3.2 to prevent V from learning the colours of more than two connected nodes in G . The idea is simple, P_1 and P_2 will return commitments for the colours of the nodes asked by V . The implicit unveiling of the commitment scheme described in section 2.3 will allow V to perform both the edge-verification and well-definition tests in a very similar way that in protocol $\Pi_{\text{std}}^{(2)}$. The commitments require V to provide a random nonzero trit for each node of the edge requested to a prover.

4.1 Distribution of questions

We now define the probability distribution \mathcal{D}'_G for V 's questions in protocol $\Pi_{\text{loc}}^{(2)}[G]$ defined in the following section. It consists in one edge and two nonzero trits for each prover:

$$\mathcal{D}'_G = \{(p'(e, r, s, e', r', s'), ((e, r, s), (e', r', s')))\}_{e, e' \in E, r, s, r', s' \in \mathbb{F}_3^*}$$

upon graph $G = (V, E)$ and where (e, r, s) is the question to P_1 and (e', r', s') is the question to P_2 . \mathcal{D}'_G is easily derived from the distribution $\mathcal{D}_G = \{(p(e, e'), (e, e'))\}_{e, e' \in E}$ for the questions in $\Pi_{\text{std}}^{(2)}[G]$, as defined in section 3.1. First, an edge $e \in_R E$ is picked uniformly at random. Together with e , two nonzero trits $r, s \in_R \mathbb{F}_3^*$ are picked at random. Then, as in \mathcal{D}_G , with probability ϵ (to be selected later) the second edge $e' = e$, in which case we always set $r' = -r$ and $s' = -s$. This case allows for an edge-verification test. Finally, with probability $1 - \epsilon$, we pick e' with probability $p(e, e')$ and pick $r', s' \in_R \mathbb{F}_3^*$ so that the couple $((e, r, s), (e', r', s'))$ is produced with probability $\frac{1}{16}p(e, e')$ for all $e, e' \in E$, and $r, s, r', s' \in \mathbb{F}_3^*$. This will allow for a well-definition test. A consequence of (1) is that for $e = (i, j) \in E$, $e' \in \text{Edges}(i) \cup \text{Edges}(j)$ with $e \neq e'$,

$$p'(e, r, s, e', r', s') \geq \frac{1 - \epsilon}{16|E|} \left(\frac{|\{e'\} \cap \text{Edges}(i)|}{|\text{Edges}(i)|} + \frac{|\{e'\} \cap \text{Edges}(j)|}{|\text{Edges}(j)|} \right). \quad (3)$$

According to (2), we also get

$$p'(e, r, s, e, r, s) = \frac{p(e, e)}{4} \geq \frac{\epsilon}{4|E|} . \quad (4)$$

It is easy to verify that \mathcal{D}'_G is a properly defined probability distribution.

4.2 The Protocol

The protocol is similar to $\Pi_{\text{std}}^{(2)}$ except that instead of returning to V the colour for each node of an edge in G , each prover returns commitments with implicit unveilings of these colours. If V asks two disjoint edges then V learns nothing about the values committed by the *forever-hiding* property of the commitment scheme. The resulting 2-prover one-round interactive proof system is denoted $\Pi_{\text{loc}}^{(2)}$.

Protocol $\Pi_{\text{loc}}^{(2)}[G]$: Two-prover, 3-COL

P_1 and P_2 pre-agree on random masks $b_i \in_R \mathbb{F}_3$ for each $i \in V$ and a random 3-colouring of G : $\{(i, c_i) | c_i \in \mathbb{F}_3\}_{i \in V}$ such that $(i, j) \in E \implies c_j \neq c_i$.

Commit phase:

- V picks $((i, j), r, s), ((i', j'), r', s') \in_{\mathcal{D}'_G} (E \times (\mathbb{F}_3^*)^2)$, sends $((i, j), r, s)$ to P_1 and $((i', j'), r', s')$ to P_2 .
- If $(i, j) \in E$ then P_1 replies $w_i = b_i \cdot r + c_i$ and $w_j = b_j \cdot s + c_j$.
- If $(i', j') \in E$ then P_2 replies $w'_{i'} = b_{i'} \cdot r' + c_{i'}$ and $w'_{j'} = b_{j'} \cdot s' + c_{j'}$.

Check phase:

Edge-Verification Test:

- if $(i, j) = (i', j')$ and $(r', s') \neq (r, s)$ then V accept iff $w_i + w'_i \neq w_j + w'_j$.

Well-Definition Test:

- If $(i, j) = (i', j')$ and $\neg((r', s') \neq (r, s))$ then V accepts iff $((w_i = w'_i) \vee (r \neq r')) \wedge ((w_j = w'_j) \vee (s \neq s'))$.
 - if $(i, j) \cap (i', j') = i$ and $r' = r$ then V accepts iff $w_i = w'_i$.
 - If $(i, j) \cap (i', j') = j$ and $s' = s$ then V accepts iff $w_j = w'_j$.
-

Clearly, $\Pi_{\text{loc}}^{(2)}$ satisfies perfect completeness. The following theorem establishes that in addition to perfect completeness, $\Pi_{\text{loc}}^{(2)}$ is sound against classical provers.

Theorem 1. *The two-prover interactive proof system $\Pi_{\text{loc}}^{(2)}$ is perfectly complete with classical value $\omega(\Pi_{\text{loc}}^{(2)}[G]) \leq 1 - \frac{1}{12|E|}$ upon any graph $G = (V, E) \notin \text{3COL}$.*

Proof. Perfect completeness is obvious. Assume $G \notin \text{3COL}$ and let us consider the probability δ that V detects an error in the check phase when interacting with two local dishonest provers $\tilde{\mathsf{P}}_1$ and $\tilde{\mathsf{P}}_2$. $\Pi_{\text{loc}}^{(2)}$ is a one-round protocol where the provers cannot communicate directly with each other nor through V 's questions since they are independent of the provers' answers. It follows that the strategy of $\tilde{\mathsf{P}}_1$ and $\tilde{\mathsf{P}}_2$ can be made deterministic without damaging the soundness error by letting each prover choosing the answer that maximizes her/his probability of success given her/his question. Therefore, consider a deterministic strategy as a pair of arrays $W^\ell[i, r, j, s] \in \mathbb{F}_3^2$ to be used by prover $\tilde{\mathsf{P}}_\ell$ for $\ell \in \{1, 2\}$ (i.e. we only care about the entries where $(i, j) \in E$ upon question $((i, j), r, s)$). For $z \in \{1, 2\}$, $W_z^\ell[\cdot, \cdot, \cdot, \cdot]$ is the z -th component of the output pair $W^\ell[\cdot, \cdot, \cdot, \cdot]$. We let $W_1^\ell(i, r, j, s) = W_2^\ell(j, s, i, r)$, as the order in which the vertices of an edge are

given to a prover is irrelevant (V can always choose the same order). We say that $W[i, r]$ for $[i, r] \in E \times \mathbb{F}_3^*$ is *well defined* if for all j, k such that $(i, j), (i, k) \in \text{Edges}(i) \neq \emptyset$ and $\forall s, t \in \mathbb{F}_3^*$,

$$W_1^1[i, r, j, s] = W_1^2[i, r, k, t] . \quad (5)$$

For $W[i, r]$ well defined, we set $W[i, r] := W_1^1[i, r, j, 1]$ for an arbitrary j such that $(i, j) \in \text{Edges}(i)$.

We now lower bound the probability $\delta_{\text{wdt}} > 0$ that, when $W[i, r]$ is not well-defined for some $i \in V$ and $r \in \mathbb{F}_3^*$, the well-definition test will detect it. When (5) is not satisfied, we have $W_1^1[i, r, j, s] \neq W_1^2[i, r, k, t]$ for some $(i, j), (i, k) \in \text{Edges}(i)$. Let $e = (i, j)$ and $e' = (i, k)$ be these two edges. According to (3) (and (1) when $e = e'$), the well-definition test will then detect an error with probability

$$\Pr(V \text{ picks } e \text{ and } e' \text{ with randomness } r, s, t) = p'(e, r, s, e', r, t) \geq \frac{1 - \epsilon}{16 \cdot |E| |\text{Edges}(i)|} . \quad (6)$$

We can do much better. Consider $W_1^1[i, r, m, u], W_1^2[i, r, m, u]$ for $(i, m) \in \text{Edges}(i)$ and $u \in \mathbb{F}_3^*$. For $i \in V$ and value $r \in \mathbb{F}_3^*$ fixed, three cases can happen:

1. $W_1^1[i, r, m, u] \neq W_1^2[i, r, m, u]$, in which case $e = (i, m)$ and $e' = (i, m)$ are incompatible for values u and u , or
2. $W_1^1[i, r, j, s] \neq W_1^2[i, r, m, u]$, in which case $e = (i, j)$ and $e' = (i, m)$ are incompatible for values s and u , or
3. $W_1^1[i, r, m, u] = W_1^2[i, r, m, u]$ and $W_1^1[i, r, j, s] \neq W_1^2[i, r, k, t]$, in which case $W_1^1[i, r, m, u] \neq W_1^2[i, r, m, u]$ and $e = e' = (i, m)$ are incompatible for value u on both sides.

In other words, if $(i, j), (i, k) \in \text{Edges}(i)$ are such that $W_1^1[i, r, j, s] \neq W_1^2[i, r, k, t]$ then for any $(i, m) \in \text{Edges}(i)$ and for any randomness $u \in \mathbb{F}_3^*$ associated to node m , V catches the provers with probability expressed on the right hand side of (6). It follows that if $W[i, r]$ is not well defined then there are $2 \cdot |\text{Edges}(i)|$ ways for V to catch the provers and each of these has probability at least $\frac{1 - \epsilon}{16 \cdot |E| \cdot |\text{Edges}(i)|}$ to be picked. It follows that,

$$\delta_{\text{wdt}} \geq \frac{2(1 - \epsilon) \cdot |\text{Edges}(i)|}{16 \cdot |E| \cdot |\text{Edges}(i)|} = \frac{1 - \epsilon}{8 \cdot |E|} .$$

Now, assume that for all $i \in E$ and $r \in \mathbb{F}_3^*$, $W[i, r]$ is well-defined, which means that the commitment values produced by the provers satisfy the consistency test. As discussed in section 2.3, when the commitments are consistent, the unique values committed upon are defined by $c_i := 2^{-1} \cdot (W[i, r] + W[i, -r])$. Since $G \notin \text{3COL}$, two of the nodes must be of the same colour at the end-points of at least one edge $(i^*, j^*) \in E$. In this case the edge-verification test will detect it when (i^*, j^*) is the edge announced to both provers and if randomness $(r, s) \in \mathbb{F}_3^* \times \mathbb{F}_3^*$ is announced to P_1 then $(-r, -s)$ is the randomness announced to \tilde{P}_2 . Using (4), the probability δ_{evt} to detect such an edge when $W[i, r]$ is well defined for all $i \in V$ and $r \in \mathbb{F}_3^*$ satisfies

$$\delta_{\text{evt}} \geq \min_{e \in E} (p'(e, r, s, e, r, s)) \geq \frac{\epsilon}{4 \cdot |E|} .$$

Therefore, the detection probability δ of any deterministic strategy for $G \notin \text{3COL}$ satisfies

$$\delta \geq \min(\delta_{\text{wdt}}, \delta_{\text{evt}}) \geq \frac{1}{12 \cdot |E|} \quad (\text{maximized at } \epsilon = 1/3) .$$

The result follows as the classical value of the game $\omega(\Pi_{\text{loc}}^{(2)}[G]) \leq 1 - \delta$. ■

To prove (perfect) zero-knowledge, it suffices to show that if $((i, j), r, s)$ and $((i', j'), r', s')$ are selected arbitrarily, V can determine at most the colours of two nodes (that form an edge). The commitments prevent a dishonest prover \tilde{V} to learn the colours of two nodes that are not connected by an edge in G . Proving this is not very hard and will be done in Section 5.3 for the three-prover case (although with three provers, \tilde{V} may also learn the colour of three nodes that form a triangle). The addition of a third prover will allow, using lemma 1, to get soundness against entangled provers without compromising zero-knowledge. As shown in [12], their protocol is not necessarily sound against two entangled provers. We also do not know whether $\Pi_{\text{std}}^{(2)}$ is sound against two entangled provers.

5 Three-Prover Protocol Sound Against Entangled Provers

The three-prover protocol $\Pi_{\text{qnl}}^{(3)}$, defined below, is identical to $\Pi_{\text{loc}}^{(2)}$ except that P_3 is asked to repeat exactly what P_1 or P_2 has replied. The prover that P_3 is asked to emulate is picked at random by V . An application of lemma 1 allows to conclude the soundness of $\Pi_{\text{qnl}}^{(3)}$ against entangled provers. Zero-knowledge remains since the only way to provide V with the colours of more than two connected nodes is if they form a complete triangle of G . This reveals nothing beyond the fact that $G \in 3\text{COL}$ to V , since all nodes will then show different colours.

5.1 Distribution of questions

The probability distribution \mathcal{D}'_G for V 's questions to the three provers is easily obtained from the distribution \mathcal{D}'_G for the questions in protocol $\Pi_{\text{loc}}^{(2)}[G]$. V picks $((e, r, s), (e', r', s')) \in \mathcal{D}'_G (E \times (\mathbb{F}_3^*)^2)^2$ and sets $e'' = e$, $r'' = r$, and $s'' = s$ with probability $\frac{1}{2}$ or sets $e'' = e'$, $r'' = r'$, and $s'' = s'$ also with probability $\frac{1}{2}$. Defined that way, \mathcal{D}'_G is a properly defined probability distribution for V 's three questions, each one in $E \times (\mathbb{F}_3^*)^2$.

5.2 The Protocol

Protocol $\Pi_{\text{qnl}}^{(3)}[G]$: Three-prover, 3-COL.

Provers P_1, P_2 , and P_3 pre-agree on random values $b_i \in_R \mathbb{F}_3$ for all $i \in V$ and a random 3-colouring of G : $\{(i, c_i) | c_i \in \{0, 1, 2\}\}_{i \in V}$ such that $(i, j) \in E \implies c_j \neq c_i$.

Commit phase:

- V picks $((i, j), r, s), ((i', j'), r', s'), ((i'', j''), r'', s'') \in \mathcal{D}'_G (E \times (\mathbb{F}_3^*)^2)^3$, sends $((i, j), r, s)$ to P_1 , sends $((i', j'), r', s')$ to P_2 , and sends $((i'', j''), r'', s'')$ to P_3 .
- If $(i, j) \in E$ then P_1 replies $w_i = b_i \cdot r + c_i$ and $w_j = b_j \cdot s + c_j$.
- If $(i', j') \in E$ then P_2 replies $w'_{i'} = b_{i'} \cdot r' + c_{i'}$ and $w'_{j'} = b_{j'} \cdot s' + c_{j'}$.
- If $(i'', j'') \in E$ then P_3 replies $w''_{i''} = b_{i''} \cdot r'' + c_{i''}$ and $w''_{j''} = b_{j''} \cdot s'' + c_{j''}$.

Check phase:

Consistency Test:

- If $((i'', j''), r'', s'') = ((i, j), r, s)$ then V rejects if $(w_i, w_j) \neq (w''_{i''}, w''_{j''})$.
- If $((i'', j''), r'', s'') = ((i', j'), r', s')$ then V rejects if $(w'_{i'}, w'_{j'}) \neq (w''_{i''}, w''_{j''})$.

Edge-Verification Test:

- if $(i, j) = (i', j')$ and $(r', s') \neq (r, s)$ then V accept iff $w_i + w'_i \neq w_j + w'_j$.

Well-Definition Test:

- if $(i, j) \cap (i', j') = i$ and $r = r'$ then V accepts iff $w_i = w'_i$.
 - If $(i, j) \cap (i', j') = j$ and $s = s'$ then V accepts iff $w_j = w'_j$.
-

In protocol $\Pi_{\text{qnl}}^{(3)}$, after the three questions picked according \mathcal{D}_G'' by V have been answered by the provers, V accepts if and only if the replies of P_1 and P_2 are accepted in $\Pi_{\text{loc}}^{(2)}$ and in addition, P_3 gave the same reply than the prover it emulates.

The soundness of protocol $\Pi_{\text{qnl}}^{(3)}$ against entangled provers can easily be shown a direct consequence of the soundness of protocol $\Pi_{\text{loc}}^{(2)}$ against classical provers, by an application of Lemma 1. Indeed, the soundness error corresponds to the quantum value of the game when $G \notin 3\text{COL}$ and $\Pi_{\text{loc}}^{(2)}$ is obviously symmetric.

Theorem 2. *The three-prover interactive proof system $\Pi_{\text{qnl}}^{(3)}$ is perfectly complete and has quantum value*

$$\omega^*(\Pi_{\text{qnl}}^{(3)}[G]) \leq 1 - \left(\frac{1}{25|E|} \right)^4 \quad (7)$$

upon any graph $G = (V, E) \notin 3\text{COL}$.

Proof. Assume $G = (V, E) \notin 3\text{COL}$. The contrapositive of Lemma 1 indicates any one-round symmetric game $\Pi_{\text{loc}}^{(2)}[G]$ with classical value $\omega(\Pi_{\text{loc}}^{(2)}[G]) \leq 1 - \delta - 12|Q|\sqrt{\delta}$ is such that the modified game $\Pi_{\text{qnl}}^{(3)}[G]$ has quantum value $\omega^*(\Pi_{\text{qnl}}^{(3)}[G]) \leq 1 - \delta$. The set Q of questions to each player satisfies $|Q| = 4|E|$. Theorem 1 establishes that $\delta + 12|Q|\sqrt{\delta} \geq \frac{1}{12|E|}$, which implies $\sqrt{\delta} \geq \frac{1}{(1+12|Q|) \cdot 12 \cdot |E|} = \frac{1}{12|E|+576|E|^2} \geq \frac{1}{588|E|^2}$, and the result follows. ■

As an immediate consequence of Theorem 2, $\Omega(|E|^4)$ sequential repetitions of $\Pi_{\text{qnl}}^{(3)}$ produces an interactive proof system for 3COL with negligible soundness error. Although the resulting proof system can be implemented on short distances, these many sequential communication rounds need to be performed at high rate for a given proof to be concluded in reasonable time. A few executions of $\Pi_{\text{qnl}}^{(3)}$ could be run in parallel without having to increase (significantly) the distances while reducing the number of sequential rounds. However, we don't know how the soundness error decreases when $\Pi_{\text{qnl}}^{(3)}$ is run only a few times in parallel, even though the results of Kempe and Vidick, a quantum version of Raz's parallel repetition theorem[21], indicate that $\Omega(|E|^4)$ runs in parallel produces a proof system with negligible soundness error[22].

5.3 Proof of Perfect Zero-Knowledge

In this section, we prove that protocol $\Pi_{\text{qnl}}^{(3)}$ is perfect zero-knowledge. As a consequence, $\Pi_{\text{loc}}^{(2)}$ is also zero-knowledge since everything \tilde{V} sees in $\Pi_{\text{loc}}^{(2)}$ can also be observed in $\Pi_{\text{qnl}}^{(3)}$. The proof of zero-knowledge proceeds using the fact that a vertex must appear at least twice to have its colour unveiled. This is the *forever hiding property* of the commitment scheme described in Section 2.3. Notice that this would be enough for \tilde{V} to learn something about the colouring if no extra condition on these three vertices is observed. In fact, we can easily show that only a few cases of colour disclosure are possible and in each of these cases, \tilde{V} learns nothing about the colouring that it could not have computed on its own. \tilde{V} can only learn colour of two connected vertices in G and nothing else or the colours of three vertices forming a triangle in G . In each of these cases, \tilde{V} learns random distinct colours for these vertices, which is to be expected by a valid 3-colouring of G . Let us show why this is enforced by the properties (see Section 2.3) of the commitment scheme. Remember that in order to learn the colour assigned to a vertex $i \in V$, \tilde{V} must ask that vertex to at least 2 distinct provers. Otherwise, \tilde{V} sees only random values returned by the provers. There are 7 cases of figure depending on how \tilde{V} selects the 3 edges asked. Figure 4 shows all cases. The 3 edges indicated for each case are the one picked by \tilde{V} . The colours associated to white vertices remain hidden by the forever hiding property of the commitment scheme. For

these vertices, the committed values received from the provers are just random and independent elements in \mathbb{F}_3 . In each of the 7 cases, the unveiled colours of the vertices are displayed in shade of grey. We see that the only way to unveil the colour of two vertices (cases 2, 3, 4, 5, and 6) is when they are connected by an edge, which means that the colours of both vertices are random but distinct. The only way for \tilde{V} to learn the colour of 3 distinct vertices is when they form a triangle (case 7). In this case, \tilde{V} learns three random and distinct colours. Clearly, this is nothing more than something necessarily true when $G \in 3\text{COL}$.

These properties of the commitment scheme allows, for any quantum polynomial-time dishonest verifier \tilde{V} , an easy simulator for $\mathbf{view}(P_1, P_2, P_3, \tilde{V}, G)$ when $G \in 3\text{COL}$, thus establishing that $\Pi_{\text{qnl}}^{(3)}$ is perfect zero-knowledge.

Theorem 3. *The three-prover interactive proof system $\Pi_{\text{qnl}}^{(3)}$ is perfect zero-knowledge against quantum verifiers.*

Proof. The simulator Sim , see Fig. 3, is classical given blackbox access to \tilde{V} (and \tilde{V} can be quantum). Consider an execution $\text{Sim}(G)$ upon graph $G = (V, E)$. It first picks a random permutation $\text{COL}[\cdot] : \mathbb{F}_3 \mapsto \mathbb{F}_3$ over three colours, each corresponding to a distinct element in \mathbb{F}_3 . Table $\text{MARK}[i, r] \in \{\text{true}, \text{false}\}$, for $i \in V$ and $r \in \mathbb{F}_3^*$, is initialized to false and will indicate if the output of a prover has already been simulated for vertex i with randomness r . Table $\text{COUNT}[i]$, for $i \in V$, counts the number of times vertex i has been asked so far during the simulation. Variable $c \in \mathbb{F}_3$, initialized to 0, indicates the next colour index the simulator should use when a new colour must be unveiled during the simulation.

Simulator $\text{Sim}(G)$: Simulator for \tilde{V} 's view upon graph G in $\Pi_{\text{qnl}}^{(3)}$.
All arithmetic below is performed in \mathbb{F}_3 .

1. Let $\text{COL}[\cdot]$ be a uniform permutation of \mathbb{F}_3 and let $c := 0$.
2. $\forall i \in V, \forall r \in \mathbb{F}_3^*$, let $\text{MARK}[i, r] := \text{false}$ and $\text{COUNT}[i] := 0$.
3. Run \tilde{V} until it returns $((i_1, j_1), r_1, s_1), ((i_2, j_2), r_2, s_2), ((i_3, j_3), r_3, s_3)$.
4. For each $\ell \in \{1, 2, 3\}$ do:
 - Whenever $(i_\ell, j_\ell) \in E$ is provided by \tilde{V} , output $(w_{i_\ell}^\ell, w_{j_\ell}^\ell) \in \mathbb{F}_3 \times \mathbb{F}_3$ to \tilde{V} , both computed as follows:
 - (a) If $\neg \text{MARK}[i_\ell, r_\ell]$ then
 - If $\text{COUNT}[i_\ell] = 0$ then pick $W[i_\ell, r_\ell] \in_R \mathbb{F}_3$.
 - If $\text{COUNT}[i_\ell] = 1$ then
 - * $W[i_\ell, r_\ell] := -\text{COL}[c] - W[i_\ell, -r_\ell]$,
 - * $c := c + 1$.
 - $\text{COUNT}[i_\ell] := \text{COUNT}[i_\ell] + 1$.
 - (b) If $\neg \text{MARK}[j_\ell, s_\ell]$ then
 - If $\text{COUNT}[j_\ell] = 0$ then pick $W[j_\ell, s_\ell] \in_R \mathbb{F}_3$.
 - If $\text{COUNT}[j_\ell] = 1$ then
 - * $W[j_\ell, s_\ell] := -\text{COL}[c] - W[j_\ell, -s_\ell]$,
 - * $c := c + 1$.
 - $\text{COUNT}[j_\ell] := \text{COUNT}[j_\ell] + 1$.
 - (c) $\text{MARK}[i_\ell, r_\ell] := \text{true}$, $\text{MARK}[j_\ell, s_\ell] := \text{true}$.
 - (d) $w_{i_\ell}^\ell := W[i_\ell, r_\ell]$, $w_{j_\ell}^\ell := W[j_\ell, s_\ell]$.

Fig. 3. Simulator for $\Pi_{\text{qnl}}^{(3)}$.

\tilde{V} is then invoked to produce questions $((i_\ell, j_\ell), r_\ell, s_\ell)$ for all provers $P_\ell, \ell \in \{1, 2, 3\}$. Sim now aims at setting the values $(w_{i_\ell}^\ell, w_{j_\ell}^\ell)$ for P_ℓ 's commitments. If $(i_\ell, j_\ell) \notin E$, Sim produces no value for $(w_{i_\ell}^\ell, w_{j_\ell}^\ell)$, exactly as P_ℓ in $\Pi_{\text{qnl}}^{(3)}$.

When $(i_\ell, j_\ell) \in E$, Sim first produces P_ℓ 's commitment $w_{i_\ell}^\ell$ for $i_\ell \in V$ and then produces P_ℓ 's commitment $w_{j_\ell}^\ell$ for $j_\ell \in V$. We show how to compute $w_{i_\ell}^\ell$, $w_{j_\ell}^\ell$ is computed similarly mutatis mutandis:

- if $\text{MARK}[i_\ell, r_\ell]$ then Sim returns the value of $w_{i_\ell}^\ell$ already determined for the simulation of the commitment of an *earlier* prover $P_h, h < \ell$. This ensures that both the commitment's *consistency test* performed and the well-definition test are always successful, as in $\Pi_{\text{qnl}}^{(3)}$ with honest provers.
 - if $\neg \text{MARK}[i_\ell, r_\ell]$ then Sim has never simulated a commitment of the colour for vertex i_ℓ with randomness r_ℓ . The value $\text{COUNT}[i_\ell]$ indicates the number of time prior to this value for ℓ , vertex i_ℓ has been asked:
 - If $\text{COUNT}[i_\ell] = 0$ then $w_{i_\ell}^\ell \in_R \mathbb{F}_3$ is picked uniformly at random, as it should be when the commitment value for the colour of vertex i_ℓ is observed in isolation.
 - If $\text{COUNT}[i_\ell] = 1$ then the colour associated to vertex i_ℓ has been committed to value $w_{i_\ell}^h$ by an *earlier* simulated prover $P_h, h < \ell$ upon randomness $-r_\ell$ (otherwise, $\text{MARK}[i_\ell, r_\ell] = \text{true}$). Sim sets $w_{i_\ell}^\ell = -\text{COL}[c] - w_{i_\ell}^h$, which satisfies the *implicit unveiling* of random colour $\text{COL}[c] = -w_{i_\ell}^\ell - w_{i_\ell}^h$. The current colour c is incremented.
- The value of $\text{COUNT}[i_\ell]$ is increased by one and $\text{MARK}[i_\ell, r_\ell] = \text{true}$, as the colour of vertex i_ℓ with randomness r_ℓ has been committed upon by the simulated prover P_ℓ .

Let $(w_{i_1}^1, w_{j_1}^1)$, $(w_{i_2}^2, w_{j_2}^2)$, and $(w_{i_3}^3, w_{j_3}^3)$ be all commitment values simulated by Sim . As discussed above and shown in Fig. 4, the colours of no more than 3 vertices are unveiled in the process. Sim always unveils as many different colours there are colours unveiled to \tilde{V} . If Sim 's simulated committed values unveils only the colour of one vertex then that colour is random, as it should in this case in $\Pi_{\text{qnl}}^{(3)}$. If Sim 's committed values unveils the colours of exactly 2 vertices then these 2 vertices form an edge in G and the colours are two different random colours, as it should be in $\Pi_{\text{qnl}}^{(3)}$. Finally, when Sim 's committed values unveil the colours of exactly 3 vertices then these vertices form a triangle in G . The 3 colours unveiled by Sim to \tilde{V} are different and assigned randomly to each of the 3 vertices, as it is in $\Pi_{\text{qnl}}^{(3)}$. Otherwise, if w_i^ℓ for $i \in V$ has been generated with only one random value then w_i^ℓ is random and uniform in \mathbb{F}_3 , exactly as it is in $\Pi_{\text{qnl}}^{(3)}$ in the same situation. It is now clear that,

$$\mathbf{view}(P_1, P_2, P_3, \tilde{V}, G) = \text{Sim}(G) ,$$

and $\Pi_{\text{qnl}}^{(3)}$ is perfect zero-knowledge. ■

6 Conclusion and Open Problems

We have provided a three-prover perfect zero-knowledge proof system for **NP** sound against entangled provers that is implementable in some well controlled environment. In order to make it fully practical, it would be better to find a protocol with smaller soundness error and also requiring only two provers. Is it possible? Moreover, we would like to extend our techniques to prove any language in **QCMA** or **QMA**, the natural quantum extensions of **NP**. We would also want to prove whether $\Pi_{\text{std}}^{(2)}$ is sound against entangled provers. Finally, we seek a variant of $\Pi_{\text{std}}^{(2)}$ that would be sound against No-Signalling provers and a variant of $\Pi_{\text{loc}}^{(2)}$ and $\Pi_{\text{qnl}}^{(3)}$ that is both sound against No-Signalling provers and Zero-Knowledge.




Fig. 4. The 7 ways to unveil the colours of at most 3 nodes in $\Pi_{\text{qnl}}^{(3)}$.

Acknowledgements

We would like to thank P. Alikhani, N. Brunner, S. Designolle, A. Chailloux, A. Leverrier, W. Shi, T. Vidick, and H. Zbinden for various discussions about earlier versions of this work. We would also like to thank Jeremy Clark for his insightful comments.

References

1. J. Kilian, “Strong separation models of multi prover interactive proofs,” in *DIMACS Workshop on Cryptography*, 1990.
2. M. Ben-Or, S. Goldwasser, J. Kilian, and A. Wigderson, “Multi-prover interactive proofs: How to remove intractability assumptions,” in *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing*, STOC ’88, (New York, NY, USA), pp. 113–131, ACM, 1988.
3. A. Kent, “Unconditionally secure bit commitment,” *Phys. Rev. Lett.*, vol. 83, pp. 1447–1450, Aug 1999.
4. T. Lunghi, J. Kaniewski, F. Bussi eres, R. Houlmann, M. Tomamichel, S. Wehner, and H. Zbinden, “Practical relativistic bit commitment,” *Phys. Rev. Lett.*, vol. 115, p. 030502, Jul 2015.
5. E. Verbanis, A. Martin, R. Houlmann, G. Boso, F. Bussi eres, and H. Zbinden, “24-hour relativistic bit commitment,” *Phys. Rev. Lett.*, vol. 117, p. 140506, Sep 2016.
6. A. Chailloux and A. Leverrier, “Relativistic (or 2-prover 1-round) zero-knowledge protocol for NP secure against quantum adversaries,” in *Advances in Cryptology – EUROCRYPT 2017: 36th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Paris, France, April 30 – May 4, 2017, Proceedings, Part III*, pp. 369–396, Springer International Publishing, 2017.
7. D. Lapidot and A. Shamir, “A one-round, two-prover, zero-knowledge protocol for NP,” *Combinatorica*, vol. 15, no. 2, pp. 204–214, 1995.
8. U. Feige and J. Kilian, “Two prover protocols: low error at affordable rates,” in *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing, 23-25 May 1994, Montr al, Qu bec, Canada* (F. T. Leighton and M. T. Goodrich, eds.), pp. 172–183, ACM, 1994.
9. A. Chiesa, M. A. Forbes, T. Gur, and N. Spooner, “Spatial isolation implies zero knowledge even in a quantum world,” *Electronic Colloquium on Computational Complexity (ECCC)*, vol. 25, p. 44, 2018.

10. A. B. Grilo, W. Slofstra, and H. Yuen, “Perfect zero knowledge for quantum multiprover interactive proofs,” *Electronic Colloquium on Computational Complexity (ECCC)*, vol. 26, p. 86, 2019.
11. J. Kempe, H. Kobayashi, K. Matsumoto, B. Toner, and T. Vidick, “Entangled games are hard to approximate,” *SIAM Journal on Computing*, vol. 40, no. 3, pp. 848–877, 2011.
12. R. Cleve, P. Høyer, B. Toner, and J. Watrous, “Consequences and limits of nonlocal strategies,” in *Proceedings of the 19th IEEE Annual Conference on Computational Complexity, CCC '04*, (Washington, DC, USA), pp. 236–249, IEEE Computer Society, 2004.
13. D. Chaum, A. Fiat, and M. Naor, “Untraceable electronic cash,” in *Proceedings on Advances in Cryptology, CRYPTO '88*, (Berlin, Heidelberg), pp. 319–327, Springer-Verlag, 1990.
14. J. Kilian, *Uses of randomness in algorithms and protocols*. MIT Press, 1990.
15. G. Brassard, C. Crépeau, D. Mayers, and L. Salvail, “Defeating classical bit commitments with a quantum computer.” arXiv:quant-ph/9806031, June 1998.
16. J. F. Clauser, M. A. Horne, A. Shimony, and R. A. Holt, “Proposed experiment to test local hidden-variable theories,” *Phys. Rev. Lett.*, vol. 23, pp. 880–884, Oct 1969.
17. S. Popescu and D. Rohrlich, “Quantum nonlocality as an axiom,” *Foundations of Physics*, vol. 24, no. 3, pp. 379–385, 1994.
18. C. Crépeau, L. Salvail, J.-R. Simard, and A. Tapp, “Two provers in isolation,” in *Advances in Cryptology – ASIACRYPT 2011: 17th International Conference on the Theory and Application of Cryptology and Information Security, Seoul, South Korea, December 4-8, 2011. Proceedings*, (Berlin, Heidelberg), pp. 407–430, Springer Berlin Heidelberg, 2011.
19. S. Fehr and M. Fillinger, “Multi-prover commitments against non-signaling attacks,” in *Advances in Cryptology – CRYPTO 2015: 35th Annual Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2015, Proceedings, Part II*, (Berlin, Heidelberg), pp. 403–421, Springer Berlin Heidelberg, 2015.
20. A. Shamir, “How to share a secret,” *Commun. ACM*, vol. 22, no. 11, pp. 612–613, 1979.
21. R. Raz, “A parallel repetition theorem,” *SIAM Journal on Computing*, vol. 27, no. 3, pp. 763–803, 1998.
22. J. Kempe and T. Vidick, “Parallel repetition of entangled games,” in *Proceedings of 43rd ACM Symposium on Theory of Computing (STOC)*, pp. 353–362, 2011.