Physical ZKP for Connected Spanning Subgraph: Applications to Bridges Puzzle and Other Problems

Suthee Ruangwises[*1] and Toshiya Itoh[†1]

[1]Department of Mathematical and Computing Science, Tokyo Institute of Technology, Tokyo, Japan

Abstract

An undirected graph G is known to both the prover P and the verifier V, but only P knows a subgraph H of G. Without revealing any information about H, P wants to convince V that H is a connected spanning subgraph of G, i.e. H is connected and contains all vertices of G. In this paper, we propose an unconventional zero-knowledge proof protocol using a physical deck of cards, which enables P to physically show that H satisfies the condition without revealing it. We also show applications of this protocol to verify solutions of three well-known NP-complete problems: the Hamiltonian cycle problem, the maximum leaf spanning tree problem, and a popular logic puzzle called Bridges.

Keywords: zero-knowledge proof, card-based cryptography, connected spanning subgraph, Hamiltonian cycle, maximum leaf spanning tree, graph, Bridges, puzzle

1 Introduction

A *zero-knowledge proof (ZKP)* is an interactive protocol introduced by Goldwasser et al. [8], which enables a prover P to convince a verifier V that a statement is correct without revealing any other information. A ZKP with perfect completeness and soundness must satisfy the following three properties.

1. **Perfect Completeness:** If the statement is correct, then V always accepts.

2. **Perfect Soundness:** If the statement is incorrect, then V always rejects.

3. **Zero-knowledge:** During the verification, V gets no extra information other than the correctness of the statement. Formally, there exists a probabilistic polynomial time algorithm S (called a *simulator*), without an access to P but with a black-box access to V, such that the outputs of S follow the same probability distribution as the outputs of the actual protocol.

[*]ruangwises@gmail.com
[†]titoh@c.titech.ac.jp

Goldreich et al. [7] proved that a computational ZKP exists for every NP problem. Several recent results, however, instead considered an unconventional way of constructing ZKPs by using physical objects such as a deck of cards and envelopes. The benefit of these physical protocols is that they allow external observers to check that the prover truthfully executes the protocol (which is often a challenging task for digital protocols). They also have didactic values and can be used to teach the concept of ZKP to non-experts.

Consider a verification of the following condition. An undirected graph G is known to both P and V, but only P knows a subgraph H of G. Without revealing any information about H, P wants to convince V that H is a connected spanning subgraph of G, i.e. H is connected and contains all vertices of G.

A ZKP to verify the connected spanning subgraph condition is important because this condition is a part of many well-known NP-complete problems, such as the Hamiltonian cycle problem, the maximum leaf spanning tree problem, and a famous logic puzzle called *Bridges*. To verify solutions of these problems, P needs to show that his/her solution satisfies the connected spanning subgraph condition as well as some other conditions (which are relatively easier to show).

1.1   Related Work

Most of previous work in physical ZKPs aimed to verify a solution of popular logic puzzles: Sudoku [9, 20], Nonogram [4], Akari [2], Kakuro [2, 14], KenKen [2], Takuzu [2, 13], Makaro [3], Norinori [5], Slitherlink [12], Juosan [13], Numberlink [18], Suguru [17], Ripple Effect [19], Nurikabe [16], and Hitori [16].

The theoretical contribution of these protocols is that they employ novel methods to physically verify specific functions. For example, a subprotocol in [3] verifies that a number in a list is the largest one in that list without revealing any value in the list, and a subprotocol in [9] verifies that a list is a permutation of all given numbers without revealing their order.

Some of these protocols can verify graph theoretic problems. For example, a protocol in [18] verifies a solution of the k vertex-disjoint paths problem, i.e. a set of k vertex-disjoints paths joining each of the k given pairs of endpoints in a graph. In a recent work, a subprotocol in [16] also verifies a condition related to connectivity. However, their protocol only works in a grid graph and also deals with a different condition from the one considered in this paper. (Their protocol only verifies that the selected cells on a board are connected together, not as a spanning subgraph of the whole board.)

1.2   Our Contribution

In this paper, we propose a physical card-based ZKP with perfect completeness and soundness to verify that a subgraph H is a connected spanning subgraph of an undirected graph G without revealing H.

We also show three possible applications of this protocol: verifying a Hamiltonian cycle in an undirected graph, verifying the existence of a spanning tree with at least k leaves in an undirected graph, and verifying a solution of the Bridges puzzle.

$$
\begin{array}{c|cccc}
 & \multicolumn{4}{c}{\text{Column}} \\
 & 1 & 2 & \cdots & k \\
\hline
1 & \boxed{?} & \boxed{?} & \cdots & \boxed{?} & \rightarrow E_k(0) \\
2 & \boxed{?} & \boxed{?} & \cdots & \boxed{?} & \rightarrow B \\
3 & \boxed{?} & \boxed{?} & \cdots & \boxed{?} \\
4 & \boxed{?} & \boxed{?} & \cdots & \boxed{?} \\
\vdots & \vdots & \vdots & \vdots & \vdots \\
m+2 & \boxed{?} & \boxed{?} & \cdots & \boxed{?} \\
 & \downarrow & \downarrow & & \downarrow \\
 & A_0 & A_1 & \cdots & A_{k-1}
\end{array}
$$

Figure 1: An $(m + 2) \times k$ matrix M constructed in Step 1

2 Preliminaries

Each *encoding card* used in our protocol has either ♣ or ♡ on the front side. All cards have indistinguishable back sides.

For $0 \le x < k$, define $E_k(x)$ to be a sequence of consecutive k cards, with all of them being ♣ except the $(x+1)$-th card from the left being ♡, e.g. $E_3(0)$ is ♡♣♣ and $E_4(2)$ is ♣♣♡♣. We use $E_k(x)$ to encode an integer x in $\mathbb{Z}/k\mathbb{Z}$. This encoding rule was introduced by Shinagawa et al. [22].

The cards in $E_k(x)$ are arranged horizontally as defined above unless stated otherwise. In some situations, however, we may arrange the cards vertically, where the leftmost card becomes the topmost card and the rightmost card becomes the bottommost card.

In an $m \times k$ *matrix* of cards, let Row i denote the i-th topmost row and Column j denote the j-th leftmost column.

2.1 Pile-Shifting Shuffle

A *pile-shifting shuffle* on an $m \times k$ matrix shifts the columns of the matrix by a random cyclic shift, i.e. shifts the columns cyclically to the right by r columns for a uniformly random $r \in \mathbb{Z}/k\mathbb{Z}$ unknown to all parties.

This protocol was developed by Shinagawa et al. [22]. It can be implemented in real world by putting the cards in each column into an envelope and applying several *Hindu cuts* to the sequence of envelopes [23].

2.2 Sequence Selection Protocol

Suppose we have k sequences A_0, A_1, ..., A_{k-1}, each encoding an integer in $\mathbb{Z}/m\mathbb{Z}$, and a sequence B encoding an integer b in $\mathbb{Z}/k\mathbb{Z}$. We propose the following *sequence selection protocol*, which allows us to select a sequence A_b (to be used as an input in other protocols) without revealing b.

1. Construct the following $(m + 2) \times k$ matrix M (see Fig. 1).

3

```
                    Column
               | 0  1  2  3  4  5
     ──────────┼─────────────────
            0  |   [1][2][3][4][5]  (actually face-down)
                     
            1  |   [?][?][?][?][?]
                     
Row     2  |[2] [?][?][?][?][?]
                     
            3  |[3] [?][?][?][?][?]
                     
            4  |[4] [?][?][?][?][?]
                 ↑
                 └── (actually face-down)
```

Figure 2: An example of a 4×5 enhanced matrix

(a) In Row 1, place a sequence $E_k(0)$. In Row 2, place the sequence B.

(b) In each Column $j = 1, 2, ..., k$, place the sequence A_{j-1} arranged vertically from Row 3 to Row $m + 2$.

2. Apply the pile-shifting shuffle to M.

3. Turn over all cards in Row 2. Locate the position of a $\boxed{\heartsuit}$. Suppose it is at Column j.

4. Select the sequence in Column j arranged vertically from Row 3 to Row $m + 2$. This is the sequence A_b as desired. Turn over all face-up cards.

After we are done using A_b in other protocols, we can put A_b back into M, apply the pile-shifting shuffle to M, then turn over all cards in Row 1 and shift the columns of M cyclically such that the $\boxed{\heartsuit}$ in Row 1 moves to Column 1. This reverts the matrix back to its original position, so we can reuse the sequences $A_0, A_1, ..., A_{k-1}$, and B.

2.3   Enhanced Matrix

In addition to the encoding cards, we also use *marking cards*, each having a positive integer on the front side. All cards have indistinguishable back sides.

Starting from an $m \times k$ matrix of face-down encoding cards, place face-down marking cards $\boxed{1}, \boxed{2}, ..., \boxed{k}$ from left to right on top of Row 1; this new row is called Row 0. Then, place face-down marking cards $\boxed{2}, \boxed{3}, ..., \boxed{m}$ from top to bottom (starting at Row 2) to the left of Column 1; this new column is called Column 0. We call this structure an $m \times k$ enhanced matrix (see Fig. 2).

2.4   Double-Scramble Shuffle

In a *double-scramble shuffle* on an $m \times k$ enhanced matrix, first rearrange Columns $1, 2, ..., k$ (including the marking cards in Row 0) by a uniformly random permutation unknown to all parties (which can be implemented by putting the cards in each column into an envelope and scrambling all envelopes together). Then, leave Row 1 as it is and rearrange Rows $2, 3, ..., m$ (including the marking cards in Column 0) by a uniformly random permutation unknown to all parties. This protocol was developed by Ruangwises and Itoh [18].

4

2.5 Rearrangement Protocol

A *rearrangement protocol* reverts the rows and columns of an enhanced matrix (after we perform double scramble shuffles) back to their original positions so that we can reuse the cards without revealing them. This protocol was developed by Ruangwises and Itoh [18], although slightly different protocols with the same idea were also used in other previous work [3, 10, 11, 19, 20].

In the rearrangement protocol on an $m \times k$ enhanced matrix, first apply the double-scramble shuffle to the matrix. Then, turn over all marking cards in Row 0 and rearrange the columns such that each marking card with number i will be in Column i. Analogously, turn over all marking cards in Column 0 and rearrange Rows $2, 3, ..., m$ accordingly.

2.6 Neighbor Counting Protocol

Suppose we have an $m \times k$ matrix with each row encoding an integer in $\mathbb{Z}/k\mathbb{Z}$. A *neighbor counting protocol* allows us to count the number of indices $i \geq 2$ such that Row i encodes the same integer as Row 1, without revealing any other information. This protocol was developed by Ruangwises and Itoh [18].

1. Place marking cards to make the matrix become an $m \times k$ enhanced matrix.

2. Apply the double-scramble shuffle.

3. Turn over all encoding cards in Row 1. Locate the position of a $\boxed{\heartsuit}$. Suppose it is at Column j.

4. Turn over all encoding cards in Column j. Count the number of $\boxed{\heartsuit}$s besides the one in Row 1. This is the number of indices that we want to know.

5. Turn over all face-up cards. Apply the rearrangement protocol.

3 Verifying an Undirected Path

In this section, we will explain a *path verification protocol*, which verifies the existence of an undirected path between vertices s and t in an undirected graph G. It is a special case $k = 1$ of the protocol for the k vertex-disjoint paths problem developed by Ruangwises and Itoh [18].[1]

We call s and t *terminal vertices*, and other vertices *non-terminal vertices*. We call a path $(v_1, v_2, ..., v_\ell)$ *minimal* if there are no neighboring vertices v_i and v_j such that $|i-j| > 1$. Observe that given any path between s and t, one can modify it to become a minimal one in linear time, so we can assume that P knows a minimal path between s and t.

Let d be the maximum degree of a vertex in G. In linear time, we can color the vertices of G with at most $d + 1$ colors such that there are no neighboring vertices with the same color. This $(d + 1)$-coloring is known to all parties.

---

[1]Although the k vertex-disjoint paths problem is NP-complete when k is a part of the input, the special case $k = 1$ in solvable in linear time. Hence, this protocol is actually unnecessary since V can easily verifies existence of the path by him/herself given G. However, we explain the details of this protocol in order to show its idea, which will be modified and used in our main protocol in Section 4.

On each terminal vertex v, P publicly places a sequence $E_{d+2}(0)$. On each non-terminal vertex v with the x-th color, P secretly places a sequence $E_{d+2}(0)$ if v is on P's path, or a sequence $E_{d+2}(x)$ if v is not on the path. Let $A(v)$ denote the sequence on each vertex v. Since the path is minimal, every non-terminal cell on the path has exactly two neighbors with a sequence encoding the same number as it (which is 0), while every terminal cell has exactly one such neighbor. On the other hand, every non-terminal cell not on the path has no neighbor with a sequence encoding the same number as it.

The idea is that, for every vertex v with the x-th color, P will add two "artificial neighbors" of v, both having $E_{d+2}(x)$ on it, and show that

1. every non-terminal vertex v (both on and not on the path) has exactly two neighbors with a sequence encoding the same number as $A(v)$, and

2. every terminal vertex v has exactly one neighbor with a sequence encoding the same number as $A(v)$.

Formally, to verify each non-terminal (resp. terminal) vertex v with the x-th color and with degree d_v, P performs the following steps.

1. Construct the following $(d_v + 3) \times (d + 2)$ matrix M.

    (a) In Row 1, place $A(v)$.

    (b) In each of the next d_v rows, place $A(v')$ for each neighbor v' of v.

    (c) In each of the last two rows, place $E_{d+2}(x)$.

2. Apply the neighbor counting protocol to M. V verifies that there are exactly two rows (resp. one row) encoding the same integer as Row 1.

3. Put the sequences back to their corresponding vertices.

    If every vertex in G passes the verification, then V accepts.

4   Verifying a Connected Spanning Subgraph

We get back to our main problem. Let $v_1, v_2, ..., v_n$ be the vertices in G. In order to prove that H is a connected spanning subgraph of G, it is sufficient to show that there is an undirected path between v_i and v_n in H for every $i = 1, 2, ..., n - 1$.

Note that the path verification protocol in Section 3 verifies a path between s and t in a graph G, where G is known to all parties. In this section, we will modify that protocol so that it can verify a path between s and t in a subgraph H of G, where H is known to only P. Then, P will perform the modified protocol for $n - 1$ rounds, with $s = v_i$ and $t = v_n$ in each i-th round.

At the beginning, P secretly places a sequence $B(e)$ on every edge $e \in G$ to indicate whether $e \in H$. ($B(e)$ is $E_2(1)$ if $e \in H$ and is $E_2(0)$ if $e \notin H$.) By doing this, the graph H is committed and cannot be changed later.

Let d be the maximum degree of a vertex in G. Like in the path verification protocol, consider a $(d+1)$-coloring, known to all parties, such that there are no neighboring vertices with the same color.

On every vertex v, P publicly places a sequence $A_0(v)$, which is $E_{d+3}(d+2)$. $A_0(v)$ acts as a "blank sequence" guaranteed to be different from $A_1(v')$ on any vertex v' during any round, which will be defined in the next step.

During each i-th round when P wants to show that there is a path in H between $s = v_i$ and $t = v_n$. First, P selects a minimal path between s and t in H. On each terminal vertex v, P publicly places a sequence $A_1(v)$, which is $E_{d+3}(0)$. On each non-terminal vertex v with the x-th color, P secretly places a sequence $A_1(v)$, which is $E_{d+3}(0)$ if v is on the path and is $E_{d+3}(x)$ if v is not on the path. Note that unlike $A_0(v)$ which remains the same throughout the whole protocol, $A_1(v)$ is changed in every round since it depends on the path selected in each round.

The verification steps are similar to the path verification protocol, except that in Step 1(b), P first applies the sequence selection protocol in Section 2.2 to determine whether to choose $A_0(v')$ or $A_1(v')$ for each neighbor v' of v, depending on whether an edge e between v and v' is in H or not. The idea is that if $e \in H$, then v' is still v's neighbor in H, so P chooses a sequence $A_1(v')$ and the rest works the same way as in the path verification protocol. On the other hand, if $e \notin H$, then v' is not v's neighbor in H, so P chooses a sequence $A_0(v')$ which is guaranteed to be different from $A_1(v)$.

Formally, to verify each non-terminal (resp. terminal) vertex v with the x-th color and with degree d_v, P performs the following steps.

1. Construct the following $(d_v + 3) \times (d + 3)$ matrix M.

   (a) In Row 1, place $A_1(v)$.

   (b) For each neighbor v' of v, let e be an edge between v and v', and let b be a bit encoded by $B(e)$. Apply the sequence selection protocol to choose a sequence $A_b(v')$ and place it in the next row of M. Repeatedly perform this for every neighbor of v to fill the next d_v rows.

   (c) In each of the last two rows of M, place $E_{d+3}(x)$.

2. Apply the neighbor counting protocol to M. V verifies that there are exactly two rows (resp. one row) encoding the same integer as Row 1.

3. Put the sequences back to their corresponding vertices.

   If every vertex in G passes the verification, then V accepts.

This protocol uses $2(d+3)(2n+2) + 2d + 2m$ encoding cards and $2d + 5$ marking cards, where n and m are the numbers of vertices and edges of G, respectively, and d is the maximum degree of a vertex in G. Therefore, the total number of required cards is $\Theta(dn)$.

5 Proof of Correctness and Security

We will prove the perfect completeness, perfect soundness, and zero-knowledge properties of our main protocol in Section 4.

Lemma 1 (Perfect completeness). If H is a connected spanning subgraph of G, then V always accepts.

Proof. Suppose that H is a connected spanning subgraph of G, then there exists a path between v_i and v_n in H for every $i = 1, 2, ..., n-1$.

First, we will prove the correctness of the sequence selection protocol in Section 2.2. Since B encodes the number b, when placing B in Row 2, the ♡ will be at Column $b+1$, the same column as the sequence A_b. After applying the pile-shifting shuffle, they will still be at the same column, so the sequence we get in Step 4 will be A_b.

Now consider the main protocol in each i-th round. In Step 1(b), P always selects a sequence $A_1(v')$ if $e \in H$ and $A_0(v')$ if $e \notin H$. Since $A_0(v')$ is $E_{d+3}(d+2)$ and thus is different from $A_1(v)$, adding $A_0(v')$ to a new row of M does not increase the number of rows encoding the same integer as Row 1. Therefore, the result will remain the same even if in Step 1(b) P adds only the sequences on the vertices such that $e \in H$, which is equivalent to solely applying the path verification protocol in Section 3 to verify a path between v_i and v_n on H.

The perfect completeness property of the path verification protocol has been proved in [18], so we can conclude that V always accepts. □

Lemma 2 (Perfect soundness). *If H is not a connected spanning subgraph of G, then V always rejects.*

Proof. Suppose that H is not a connected spanning subgraph of G, then there exists an index $i \in \{1, 2, ..., n-1\}$ such that there is no path between v_i and v_n in H. In Lemma 1, we have proved that the sequence selection protocol is correct, and the i-th round of the main protocol is equivalent to applying the path verification protocol to verify a path between v_i and v_n on H.

The perfect soundness property of the path verification protocol has been proved in [18], so we can conclude that V always rejects. □

Lemma 3 (Zero-knowledge). *During the verification, V learns nothing about H.*

Proof. To prove the zero-knowledge property, it is sufficient to prove that all distributions of the values that appear when the cards are turned face-up can be simulated by a simulator S without knowing H.

- In the sequence selection protocol:

  - In Step 3, we turn over all cards in Row 2. This occurs right after a pile-shifting shuffle is applied to M. Hence, the ♡ has an equal probability to be at each of the k columns, so this step can be simulated by S.

- In the neighbor counting protocol:

  - In Step 3, we turn over all encoding cards in Row 1. The order of Columns $1, 2, ..., k$ is uniformly distributed among all possible permutations due to the double-scramble shuffle. Hence, the ♡ has an equal probability to be at each of the k columns, so this step can be simulated by S.

  - In Step 4, we turn over all encoding cards in Column j. Suppose there are t ♡ s besides the one in Row 1 (t is now a public information). The order of Rows $2, 3, ..., m$ is uniformly distributed among all possible permutations due to the double-scramble shuffle. Hence, all t ♡s have an equal probability to be at each of the $\binom{m-1}{t}$ combinations of rows, so this step can be simulated by S.

Therefore, we can conclude that V learns nothing about H. □

6   Applications to NP-Complete Problems

6.1   Hamiltonian Cycle Problem

Given an undirected graph G, determining whether G has a Hamiltonian cycle (a cycle that visits each vertex exactly once) is known to be NP-complete [6]. Suppose P knows a Hamiltonian cycle H of G and wants to convince V that G has a Hamiltonian cycle without revealing any information about H.

To prove that H is a Hamiltonian cycle of G, it is sufficient to show that

1. H is a connected spanning subgraph of G, and

2. every vertex in H has degree 2.

At the beginning, P commits H by secretly placing a sequence $B(e)$ on every edge $e \in G$ to indicate whether $e \in H$. ($B(e)$ is $E_2(1)$ if $e \in H$ and is $E_2(0)$ if $e \notin H$.) The first condition can be verified by the protocol in Section 4.

To verify the second condition, P first applies the copy protocol explained in Appendix A.1 to make another copy of a sequence $B(e)$ on every edge e. (Each of the two copies will be used to verify each endpoint of e.) For each vertex $v \in H$, P considers one (unused) copy of a sequence on every edge e incident to v and selects only the leftmost card of it (which is ♣ if $e \in H$ and is ♡ if $e \notin H$). Then, P scrambles all selected cards together and turns over all of them, and V verifies that there are exactly two ♣s among them (which means v has degree 2 in H). V accepts if the verification passes for every vertex in H. This protocol also uses $\Theta(dn)$ cards.[2]

6.2   Maximum Leaf Spanning Tree Problem

Given an undirected graph G and an integer k, the decision version of the maximum leaf spanning tree problem asks whether G has a spanning tree with at least k leaves (vertices with degree 1). This problem is also known to be NP-complete [6]. Suppose P knows a spanning tree H of G with at least k leaves and wants to convince V that the such tree exists without revealing any information about H.

To prove that G has a spanning tree with at least k leaves, it is sufficient to show that

1. H is a connected spanning subgraph of G, and

2. H has at least k leaves.

Note that it is not necessary to show that H itself is a tree. (Even if H itself is not a tree, any spanning tree of H will also be a spanning tree of G, and every leaf of H will still be a leaf of that tree, so G must have a spanning tree with at least k leaves.)

---

[2]There is an alternative way to verify a Hamiltonian cycle: P publicly constructs an $n \times n$ adjacency matrix M of G, then privately selects a permutation σ and rearranges both the rows and columns of M by σ. Finally, P turns over all cards in the form $M(i, i+1)$ and $M(i, i-1)$ to show that they are all 1s. This protocol is simpler and more straightforward, but it requires $\Theta(n^2)$ cards, which is significantly greater than our protocol in sparse graphs.

Figure 3: An example of a Bridges puzzle (left) and its solution (right)

P commits H by the same way as in the Hamiltonian cycle problem, and uses the protocol in Section 4 to verify the first condition.

To verify the second condition, P makes an additional copy of every $B(e)$ like in the Hamiltonian cycle problem. For every vertex v, P selects only the leftmost card of $B(e)$ on every edge e incident to it, scramble these cards, and puts them into an envelope. (If there are less than d cards, P publicly adds more ♡s until there are d cards before scrambling them.) Then, P scrambles all envelopes together. Next, P picks an envelope, opens it and looks at the front side of all cards inside (without V seeing the front side). If there is exactly one ♣ among them, P reveals all cards to let V verify that there is exactly one ♣ (which means the corresponding vertex is a leaf); otherwise, P does not reveal the cards. P repeatedly does this for every envelope. V accepts if there are at least k envelopes with exactly one ♣. This protocol also uses $\Theta(dn)$ cards.

6.3 Bridges Puzzle

Bridges, or the Japanese name Hashiwokakero, is a logic puzzle created by a Japanese company Nikoli, which also developed many other popular logic puzzles including Sudoku, Kakuro, and Numberlink.

A Bridges puzzle consists of a rectangular grid of size $p \times q$, with some cells called *islands* containing an encircled positive number of at most 8. The objective of this puzzle is to connect some pairs of islands by straight lines called *bridges* that can only run horizontally or vertically. There can be at most two bridges between each pair of islands, and the bridges must satisfy the following conditions [15] (see Fig. 3).

1. *Island condition*: The number of bridges connected to each island must equal to the number written on that island.

2. *Noncrossing condition*: Each bridge cannot cross islands or other bridges.

3. *Connecting condition*: The bridges must connect all islands into a single component.

Determining whether a given Bridges puzzle has a solution has been proved to be NP-complete [1]. Suppose P knows a solution of the puzzle and wants to convince V that it has a solution without revealing any information about the solution.

Define a *lip* to be a line segment of a unit length on the Bridges grid that either separates two adjacent cells or lies on the outer boundary of the grid. For each lip ℓ, let

Figure 4: Positions of lips $\ell_1, \ell_2, \ell_3, \ell_4$ surrounding a cell c.

$b(\ell)$ be the number of bridges crossing through ℓ (including bridges coming out of the island from ℓ if ℓ is a lip of an island cell). First, P secretly places on ℓ a sequence encoding $b(\ell)$ in $\mathbb{Z}/3\mathbb{Z}$. Then, P publicly appends six ♣s to the end of the sequence to make it encode $b(\ell)$ in $\mathbb{Z}/9\mathbb{Z}$ (while ensuring V that $b(\ell)$ is at most 2). For each island cell c with a number $n(c)$, P publicly places a sequence encoding $n(c)$ in $\mathbb{Z}/9\mathbb{Z}$ on c.

For each cell c, let $b(\ell_1), b(\ell_2), b(\ell_3), b(\ell_4)$ be the numbers encoded by sequences on the top lip ℓ_1, the right lip ℓ_2, the bottom lip ℓ_3, and the left lip ℓ_4 of c, respectively (see Fig. 4). The steps of verifying P's solution of the puzzle are as follows.

1. For each lip ℓ located on the outer boundary of the Bridges grid, verify that $b(\ell) = 0$ (no bridge goes beyond the grid), which can be shown by simply revealing the sequence on ℓ.

2. For each island cell c with a number $n(c)$, verify that $b(\ell_1) + b(\ell_2) + b(\ell_3) + b(\ell_4) \equiv n(c)$ (mod 9) (the island condition).

3. For each non-island cell c, verify that $b(\ell_1) \equiv b(\ell_3)$ (mod 9) and $b(\ell_2) \equiv b(\ell_4)$ (mod 9) (the number of bridges passing through c is consistent), and also that $b(\ell_1) \cdot b(\ell_2) \equiv 0$ (mod 9) (the noncrossing condition).

Steps 2 and 3 can be performed by applying a combination of copy and arithmetic protocols, which are explained in Appendix A, and the neighbor counting protocol in Section 2.6 (on a 2×9 matrix to verify the congruence).

Finally, construct a public graph G with all islands being vertices of G, and two islands having an edge in G if they are on the same row or column and there is no island between them (i.e. one can construct a valid bridge between them). Let H be a private subgraph of G such that two islands have an edge in H if there is at least one bridge between them in P's solution. P performs the following steps to commit H by placing a sequence $B(e)$, which is either $E_2(0)$ or $E_2(1)$, on every edge $e \in G$ to indicate whether $e \in H$.

1. For each edge $e \in G$ with endpoints u and v, consider any lip ℓ in the Bridges puzzle that lies between the two islands corresponding to u and v.

2. P picks the leftmost card on ℓ and places it as a leftmost card of $B(e)$ without revealing it.

3. P shuffles the second and third leftmost cards on ℓ and looks at the front side of them (without V seeing the front side). Then, P selects a ♣ among them and turns it over to reveal the front side to V. (If both cards are ♣s, P can select any of them; if only one card is a ♣, P must select it.)

4. P places another unselected card in Step 3 as a rightmost card of $B(e)$ without revealing it.

11

Observe that if there are one or two bridges between u and v, then $B(e)$ will be $E_2(1)$; if there is no bridge between them, then $B(e)$ will be $E_2(0)$. Hence, these steps ensure that the subgraph H is compatible with P's solution of the puzzle without revealing any information about it.

Verifying the connecting condition is equivalent to verifying that H is a spanning subgraph of G, which can be done by the protocol in Section 4. In total, this protocol uses $\Theta(pq)$ cards.

7   Future Work

We developed a physical card-based ZKP to verify the connected spanning subgraph condition, and showed applications of this protocol to verify solutions of three well-known NP-complete problems: the Hamiltonian cycle problem, the maximum leaf spanning tree problem, and the Bridges puzzle.

A possible future work is to explore methods to physically verify other NP-complete graph theoretic problems as well as other popular logic puzzles.

References

[1] D. Andersson. Hashiwokakero is NP-complete. *Information Processing Letters*, 109(9): 1145–1146 (2009).

[2] X. Bultel, J. Dreier, J.-G. Dumas and P. Lafourcade. Physical Zero-Knowledge Proofs for Akari, Takuzu, Kakuro and KenKen. In *Proceedings of the 8th International Conference on Fun with Algorithms (FUN)*, pp. 8:1–8:20 (2016).

[3] X. Bultel, J. Dreier, J.-G. Dumas, P. Lafourcade, D. Miyahara, T. Mizuki, A. Nagao, T. Sasaki, K. Shinagawa and H. Sone. Physical Zero-Knowledge Proof for Makaro. In *Proceedings of the 20th International Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS)*, pp. 111–125 (2018).

[4] Y.-F. Chien and W.-K. Hon. Cryptographic and Physical Zero-Knowledge Proof: From Sudoku to Nonogram. In *Proceedings of the 5th International Conference on Fun with Algorithms (FUN)*, pp. 102–112 (2010).

[5] J.-G. Dumas, P. Lafourcade, D. Miyahara, T. Mizuki, T. Sasaki and H. Sone. Interactive Physical Zero-Knowledge Proof for Norinori. In *Proceedings of the 25th International Computing and Combinatorics Conference (COCOON)*, pp. 166–177 (2019).

[6] M.R. Garey and D.S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., 1979.

[7] O. Goldreich, S. Micali and A. Wigderson. Proofs that yield nothing but their validity and a methodology of cryptographic protocol design. *Journal of the ACM*, 38(3): 691–729 (1991).

[8] S. Goldwasser, S. Micali and C. Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, 18(1): 186–208 (1989).

[9] R. Gradwohl, M. Naor, B. Pinkas and G.N. Rothblum. Cryptographic and Physical Zero-Knowledge Proof Systems for Solutions of Sudoku Puzzles. *Theory of Computing Systems*, 44(2): 245–268 (2009).

[10] Y. Hashimoto, K. Shinagawa, K. Nuida, M. Inamura and G. Hanaoka. Secure Grouping Protocol Using a Deck of Cards. *IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences*, 101.A(9): 1512–1524 (2018).

[11] T. Ibaraki and Y. Manabe. A More Efficient Card-Based Protocol for Generating a Random Permutation without Fixed Points. In *Proceedings of the 3rd International Conference on Mathematics and Computers in Sciences and Industry (MCSI)*, pp. 252–257 (2016).

[12] P. Lafourcade, D. Miyahara, T. Mizuki, T. Sasaki and H. Sone. A Physical ZKP for Slitherlink: How to Perform Physical Topology-Preserving Computation. In *Proceedings of the 15th International Conference on Information Security Practice and Experience (ISPEC)*, pp. 135–151 (2019).

[13] D. Miyahara, L. Robert, P. Lafourcade, S. Takeshige, T. Mizuki, K. Shinagawa, A. Nagao and H. Sone. Card-Based ZKP Protocols for Takuzu and Juosan. In *Proceedings of the 10th International Conference on Fun with Algorithms (FUN)*, pp. 20:1–20:21 (2020).

[14] D. Miyahara, T. Sasaki, T. Mizuki and H. Sone. Card-Based Physical Zero-Knowledge Proof for Kakuro. *IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences*, E102.A(9): 1072–1078 (2019).

[15] Nikoli: Hashiwokakero. `https://www.nikoli.co.jp/en/puzzles/hashiwokakero.html`

[16] L. Robert, D. Miyahara, P. Lafourcade and T. Mizuki. Interactive Physical ZKP for Connectivity: Applications to Nurikabe and Hitori. In *Proceedings of the 17th Conference on Computability in Europe (CiE)*, pp. 373–384 (2021).

[17] L. Robert, D. Miyahara, P. Lafourcade and T. Mizuki. Physical Zero-Knowledge Proof for Suguru Puzzle. In *Proceedings of the 22nd International Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS)*, pp. 235–247 (2020).

[18] S. Ruangwises and T. Itoh. Physical Zero-Knowledge Proof for Numberlink Puzzle and k Vertex-Disjoint Paths Problem. *New Generation Computing*, 39(1): 3–17 (2021).

[19] S. Ruangwises and T. Itoh. Physical Zero-Knowledge Proof for Ripple Effect. In *Proceedings of the 15th International Conference and Workshops on Algorithms and Computation (WALCOM)*, pp. 296–307 (2021).

[20] T. Sasaki, D. Miyahara, T. Mizuki and H. Sone. Efficient card-based zero-knowledge proof for Sudoku. *Theoretical Computer Science*, 839: 135–142 (2020).

[21] K. Shinagawa and T. Mizuki. Card-based Protocols Using Triangle Cards. In *Proceedings of the 9th International Conference on Fun with Algorithms (FUN)*, pp. 31:1–31:13 (2018).

[22] K. Shinagawa, T. Mizuki, J.C.N. Schuldt, K. Nuida, N. Kanayama, T. Nishide, G. Hanaoka and E. Okamoto. Card-Based Protocols Using Regular Polygon Cards. *IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences*, E100.A(9): 1900–1909 (2017).

[23] I. Ueda, D. Miyahara, A. Nishimura, Y. Hayashi, T. Mizuki and H. Sone. Secure implementations of a random bisection cut. *International Journal of Information Security*, 19(4): 445–452 (2020).

A   Copy and Arithmetic Protocols

In this appendix, we explain the copy and arithmetic protocols that can be used to verify problems in Section 6.

A.1   Copy Protocol

Given a sequence A encoding an integer a in $\mathbb{Z}/k\mathbb{Z}$, this protocol creates m additional copies of A without revealing a. It was developed by Shinagawa et al. [22].

1. Reverse the $k-1$ rightmost cards of A, i.e. move each $(i+1)$-th leftmost card of A to become the i-th rightmost card for $i = 1, 2, ..., k-1$. This modified sequence, called A', now encodes $-a \pmod{k}$.

2. Construct a $(m+2) \times k$ matrix M by placing the sequence A' in Row 1 and a sequence $E_k(0)$ in each of Rows $2, 3, ..., m+2$.

3. Apply the pile-shifting shuffle to M. Note that Row 1 of M now encodes $-a+r \pmod{k}$, and other rows now encode $r \pmod{k}$ for a uniformly random $r \in \mathbb{Z}/k\mathbb{Z}$.

4. Turn over all cards in Row 1 of M. Locate the position of a $\boxed{\heartsuit}$. Suppose it is at Column j.

5. Shift the columns of M cyclically to the left by $j-1$ columns. Turn over all face-up cards.

6. The sequences in Rows $2, 3, ..., m+2$ of M now encode $r - (-a+r) \equiv a \pmod{k}$, so we now have $m+1$ copies of A as desired.

A.2   Addition Protocol

Given sequences A and B encoding integers a and b in $\mathbb{Z}/k\mathbb{Z}$, respectively. This protocol computes the sum $a+b \pmod{k}$ without revealing a or b. It was developed by Shinagawa et al. [22].

1. Reverse the $k-1$ rightmost cards of A. This modified sequence, called A', now encodes $-a \pmod{k}$.

2. Construct a $2 \times k$ matrix M by placing A' in Row 1 and B in Row 2.

3. Apply the pile-shifting shuffle to M. Note that Row 1 and Row 2 of M now encode $-a + r \pmod{k}$ and $b + r \pmod{k}$, respectively, for a uniformly random $r \in \mathbb{Z}/k\mathbb{Z}$.

4. Turn over all cards in Row 1 of M. Locate the position of a $\heartsuit$. Suppose it is at Column j.

5. Shift the columns of M cyclically to the left by $j - 1$ columns. Turn over all face-up cards.

6. The sequence in Row 2 of M now encodes $(b + r) - (-a + r) \equiv a + b \pmod{k}$ as desired.

A.3 Multiplication Protocol

Given sequences A and B encoding integers a and b in $\mathbb{Z}/k\mathbb{Z}$, respectively, this protocol computes the product $a \cdot b \pmod{k}$ without revealing a or b. It is a generalization of a protocol of Shinagawa and Mizuki [21] to multiply two integers in $\mathbb{Z}/3\mathbb{Z}$.

1. Repeatedly apply the copy protocol and the addition protocol to produce sequences $A_0, A_1, A_2, ..., A_{k-1}$ encoding $0, a, 2a, ..., (k-1)a \pmod{k}$, respectively.

2. Apply the sequence selection protocol to select the sequence A_b encoding $a \cdot b \pmod{k}$.